TENSE AND ASPECT IN MULTILINGUAL SEMANTIC CONSTRUCTION

A doctoral thesis submitted for the degree of Doctor of Philosophy (Dr.phil.)

presented by
Mark-Matthias Zymla

at the

Universität Konstanz

Faculty of Humanities
Department of Linguistics

Konstanz, 2023

Date of oral examination:
01.03.2023

Referees:
1. Referee: Prof. Dr. Miriam Butt
2. Referee: Prof. Dr. Dag Trygve Truslew Haug
3. Referee: Prof. Dr. Maribel Romero

"Deep in the human unconscious is a pervasive need for a logical universe that makes sense. But the real universe is always one step beyond logic."

– from *"The Sayings of Muad'Dib"* by the Princess Irulan,
in *Dune*, Frank Herbert

Abstract

Tense and aspect categories have played a substantial role in theoretical and computational linguistics for decades. However, in recent years, due to the increased reliance on deep learning models, the role of tense and aspect has somewhat moved into the backseats of linguistic analysis, even though such models famously struggle with certain semantic properties such as modals, negation, and also the categories that lie at the center of the present work. In this thesis, I provide a formally well-defined feature space for the semantic annotation of tense and aspect and a rule-based system for annotating the syntax/semantics interface relevant to the interpretation of tense and aspect.

Cross-linguistic research is thriving and still provides new insights into the commonalities and differences across languages. Specifically, typological work has already shown how rich the variation in systems available across languages is concerning tense and aspect. Furthermore, formal semantics still uncover intricacies in the cross-linguistic variation of the meanings associated with tense and aspect. This thesis aims to investigate recent advances in cross-linguistic research on tense and aspect and compile these insights into a computationally viable annotation scheme for the semantics of these categories. This is done in two steps. First, an appropriate cross-linguistic feature space is developed. Each feature is defined via an exact formal semantic contribution based on cross-linguistic formal semantic research. Secondly, a formal semantic interpretation component is developed for the feature space that translates annotations into formal semantic representations according to their specified formal semantic contribution.

The annotation proposed in this thesis is layered. The first layer provides a general interpretation of the given feature. The morpho-syntactic context primarily affects this interpretation: a direct mapping from syntax to semantics. This thesis pays particular attention to indirect mappings from syntax to semantics, i.e., the second layer of the annotation scheme. This layer of the annotation scheme rewrites semantic features based on the semantic context, mimicking compositional semantic processes. Finally, the third layer of annotation serves as input for contextual resolution – i.e., the annotation

of pragmatic inferences.

The ParTMA annotation scheme is developed as part of the project *Tense and Aspect in Multilingual Semantic Construction*, which was a daughter project of the international ParGram effort, a project to develop cross-linguistically comparable deep linguistic grammars in the Xerox Linguistic Environment(XLE) based on the formalism of Lexical Functional Grammar. The ParTMA annotation scheme is a modular system consisting of a semantics for interpreting annotations and a syntax/semantics interface. While the semantics are entirely independent, the syntax/semantics interface annotation is, to some extent, influenced by the choice of a syntactic parser. In this thesis, two syntactic parsers are considered: primarily the deep grammars from the ParGram project and, secondly, Universal Dependencies (UD). The reasons for comparisons are to show how the system can be made compatible with different syntactic components.

The overall contribution of this thesis is a full implementation of a formal semantic framework for tense and aspect. This implementation uses the flexible Glue semantics formalism developed for dealing with semantics in LFG. The semantic framework is specifically designed to interpret ParTMA style annotations and can, thus, be used to evaluate annotations directly by their formal correctness. Since the ParTMA features provide a cross-linguistically viable feature space, the system for semantic interpretation can provide cross-linguistically viable formal semantic representations with minimal adjustment for language-specific implementations.

Zusammenfassung

Die Kategorien Tempus und Aspekt spielen seit Jahrzehnten eine wichtige Rolle in der theoretischen und computationellen Linguistik. In den letzten Jahren ist ihre Rolle aufgrund des zunehmenden Einsatzes von Deep-Learning-Modellen etwas in den Hintergrund der linguistischen Analyse gerückt, obwohl solche Modelle bekanntermaße mit bestimmten semantischen Elementen wie Modalverben, Negation und auch den Kategorien, die im Mittelpunkt der vorliegenden Arbeit stehen, zu kämpfen haben. Um dem entgegenzuwirken bietet das ParTMA-Annotationsschema klar formalisierte semantische Merkmale, die auf sprachübergreifender Forschung basieren. In dieser Arbeit stelle ich also einen formal wohldefinierten Merkmalsraum für die semantische Annotation von Tempus und Aspekt vor. Außerdem präsentiere ich ein regelbasiertes System für die Annotation des Syntax/Semantik-Interfaces, welches für die Interpretation von Tempus und Aspekt relevant ist.

Die sprachübergreifende Forschung liefert immer noch neue Erkenntnisse über die Gemeinsamkeiten und Unterschiede zwischen Sprachen im Hinblick auf Tempus- und Aspektformen. Insbesondere typologische Arbeiten haben bereits gezeigt, wie vielfältig die in den verschiedenen Sprachen vorhandenen Systeme für Tempus und Aspekt sind. Darüber hinaus deckt die formale Semantik immer noch Feinheiten in der sprachübergreifenden Variation der mit Tempus und Aspekt verbundenen Bedeutungen auf. Ziel dieser Arbeit ist es, die jüngsten Fortschritte in der sprachübergreifenden Forschung zu Tempus und Aspekt zu untersuchen und diese Erkenntnisse in einem computationell realisierbaren Annotationsschema für die Semantik dieser Kategorien zusammenzufassen. Dies geschieht in zwei Schritten. Zunächst wird ein geeigneter sprachenübergreifender Merkmalsraum entwickelt. Jedes Merkmal wird durch einen exakten formal-semantischen Beitrag auf der Grundlage sprachübergreifender semantischer Forschung definiert. Zweitens wird eine semantische Interpretationskomponente für den Merkmalsraum entwickelt, die Annotationen entsprechend ihrer formal spezifizierten Bedeutung in formale semantische Repräsentationen übersetzt.

Eine der primären Anwendungen des vorliegenden Annotationsschemas ist

die Erweiterung der derzeit existierenden ParGram-Grammatiken. Um jedoch die Forschung für ein breiteres Publikum leichter nutzbar zu machen, nimmt diese Arbeit eine generellere Haltung zur morphosyntaktischen Annotation von Tempus und Aspekt ein.

Die in dieser Arbeit vorgeschlagene Annotation ist in sogenannte Tiers oder Ebenen unterteilt. Die erste Ebene liefert eine allgemeine Interpretation eines gegebenen Merkmals. Der morpho-syntaktische Kontext wirkt sich in erster Linie auf diese Interpretation aus. Die erste Ebene ist also eine direkte Abbildung von Syntax auf Semantik. In dieser Arbeit wird ein besonderes Augenmerk auf indirekte Zuordnungen von Syntax zu Semantik gelegt. Dies ist die zweite Ebene des Annotationsschemas. Diese Ebene schreibt semantische Merkmale auf der Grundlage des semantischen Kontexts um und ahmt so semantische Kompositionsprozesse nach. Die dritte Annotationsebene schließlich dient als Input für die Analyse im Kontext, d.h. die Annotation pragmatischer Inferenzen.

Das ParTMA-Annotationsschema wurde im Rahmen des Projekts *Tense and Aspect in Multilingual Semantic Construction* entwickelt, das ein Tochterprojekt des internationalen ParGram-Projekts war. Das ParGram Projekt dient zur Entwicklung sprachübergreifend vergleichbarer detaillierter linguistischer Grammatiken im Xerox Linguistic Environment (XLE) auf der Grundlage des Formalismus Lexical Functional Grammar (LFG). Das ParTMA-Annota-tionsschema ist ein modulares System, das aus einer Semantik zur Interpretation von Annotationen und einem Syntax/Semantik-Interface besteht. Während die Semantik völlig unabhängig ist, wird die Annotation des Syntax/ Semantik-Interfaces in gewissem Maße durch die Wahl einer syntaktischen Analyse beeinflusst. In dieser Arbeit werden zwei syntaktische Analysemöglichkeiten in Betracht gezogen: zum einen die des ParGram-Projekts basierend auf LFG und zum anderen Parser die auf den Universal Dependencies (UD) basieren. Der Grund für diesen Vergleich ist zu zeigen, wie das System mit verschiedenen syntaktischen Komponenten kompatibel gemacht werden kann.

Der Gesamtbeitrag dieser Arbeit ist also eine Implementierung eines formalen semantischen Systems für die Analyse von Tempus und Aspekt. Diese Implementierung verwendet den flexiblen Glue-Semantik-Formalismus, der für den Umgang mit Semantik in LFG entwickelt wurde. Das semantische Framework wurde speziell für die Interpretation von ParTMA-Annotationen entwickelt und kann daher verwendet werden, um Annotationen direkt auf ihrer formale Korrektheit zu überprüfen. Da das ParTMA-Annotationsschema einen sprachübergreifend adäquaten Merkmalsraum bietet, kann das System sprachübergreifend adäquate formale semantische Repräsentationen mit minimaler Anpassung an sprachspezifische Implementierungen liefern.

Acknowledgements

Dear reader,

First of all, thank you for reading this thesis. It was a long journey, and it feels like I am probably not the most skillful navigator when it comes to such huge projects. I have learned a lot throughout the years, and I am grateful to so many people that I am not even sure I remember them all. However, it's always best to start at the beginning.

First, I want to thank Maribel Romero for pulling me into the world of words and meaning. To be honest, I think it was my first class with her (Semantics I) that gave me some sort of confidence that I have chosen a path (Linguistics) that I could walk on for a long time. Maribel made sure that I could enjoy this part of my studies by hiring me as a student assistant and introducing me to many other great people and friends, some of which I will mention later on. *Dear Maribel, thanks for being my supervisor and helping me understand the meaning of words and sentences better.*

Secondly, I want to thank Miriam Butt, who, just like Maribel, was always helpful and supportive despite the various detours I took throughout the years. I think she taught me that sometimes it is okay to be lost as long as you can enjoy the view. She introduced me to the world of computational linguistics, and thanks to her, the circle of great people surrounding me just grew larger. *Dear Miriam, I am very grateful to you, too, for being my supervisor. Thanks for showing me that we can ask computers to do the thinking for us if we are just able to tell them what we want from them properly.*

I am also very grateful to Dag Haug, who, I assume, is not completely sure what he is getting himself into. However, I am more than glad that he accepted to be my external supervisor, and I hope that I can learn a lot from him in the future as someone whose work I admire. *Dear Dag, thanks for being my external supervisor and being patient throughout the last few "moments" of my journey.*

As mentioned above, over the years, I have met many great people to whom I am thankful for one thing or another. I am thankful to the many people from the ParGram project and its behind-the-scenes supporters. I try to name all that come to mind, but please forgive me if I have forgotten anyone. Everyone involved in this grand project: Thank you very much, notably: Tracy Holloway King, Dick Crouch and Valeria de Paiva, Cleo Condoravdi, Mary Dalrymple, Agnieszka Patejuk, Victoria Rosén, Paul Meurer, and, last but not least, Sebastian Sulger who initially took me under his wings at the beginning of my journey. Many of these people I have connected with outside of ParGram as well, and I repeat it again, Thank you, everyone, for all your support and everything else.

Aside from my supervisors, I have met so many great people at the University of Konstanz, for which I am grateful. This includes many of the researchers and other people who are part of the Linguistics department. Thank you, Andreas Walker, Vasiliki Erotokritou, Moritz Meßmer, David Krassnig, Erlinde Meertens, Katerina Kalouli, Christin Beck, Jessica Zipf, Saira Bano, Benazir Mumtaz, and Tina Bögel. *Thanks to everyone at the University of Konstanz for making my time there memorable and enjoyable.*

Finally, I want to thank my family and my friends: first of all, my father, Richard Zymla, who unfortunately died in 2019. It was thanks to him that I had the freedom to try to figure out what I wanted to do with my life. I am sorry that he never got to see the end of my journey, but maybe he is watching from somewhere. I am grateful to my mother, Monika Zymla, for supporting me as well and being there for me over the years. Many thanks to Sabine Zymla, who also supported me in all my endeavors, my siblings, Lisa and Florian Zymla, and the rest of my family. I am grateful to all my friends and to Carina, who may or may not be responsible for finally pushing me over the edge.

Glossary

1 first person
2 second person
3 third person

acc accusative
aux auxiliary
av active voice

cl classifier
comp complementizer

dat dative
dec decessive
del delimitative
det determiner
dm discourse marker

e event
erg ergative

f feminine
foc focus particle
fut future tense

gen genitive

impv imperfective

inf infinitive
intr intransitive

loc locative

m masculine

neg negation
nf non-first person
nom nominative
non-past non-past marker
nonvis nonvisual evidential

obj object
obl oblique

past past tense
perf perfect
pnc punctual
poss possessive marker
pres present tense
prf prefix
prog progressive
prosp prospective aspect
prptcp past participle
prsprt present participle

prt particle
prv perfective
pstptcp past participle

rec recent
reflx reflexive

s state

sg singular
subj subject

topic topic particle/marker
tr transitive

vis visual evidential

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Articulating the flow of time is a skill that is particularly developed within human communication. The human language can express all but the most complex of temporal circumstances. The language system responsible for this ability can be intuitively (and informally) divided into three components: Firstly, a lexical component which conveys more or less explicit references to time such as temporal adverbials like *yesterday*, *a week ago* or *in a little bit*. The second component I have in mind is intertwined with the lexical domain, namely conventionalized elements such as dates and calendar units. The third component is the encoding of temporal information within the grammar. The categories associated with temporal semantics are (predominantly) tense and aspect; these two categories lie at the center of the present thesis.

Concretely, I am interested in a computationally viable annotation scheme for these categories. Various tasks in NLP require or could improve with a proper interpretation of tense and aspect information. This proper interpretation requires a thorough understanding of the underlying semantics. Thus, the primary goal of this thesis is to provide a semantic annotation scheme of tense and aspect that incorporates insights from formal semantics and current typological and computational work.

The project during which the core of this thesis was developed is an extension of an international grammar engineering effort spanning several typologically distinct languages, the ParGram effort (Butt et al., 2002), which will be discussed in detail in chapter 2. This inheritance has a consequence for the main goal of this thesis, namely, particular attention to the cross-linguistic adequacy of the presented annotation scheme. This particular requirement is, of course, not only required within this particular project but tends to a broader need of NLP resources in general.

With the goals listed above in mind, this dissertation is about identifying which different meaning concepts are realized in the surface form of language (i.e., morphosyntax and phonetics/phonology, although the latter does not play a role in this thesis) and labeling the data to encode these concepts. This is not trivial since while grammatical features have a conceptual core (e.g., pastness for past tense), this conceptual core can be employed in several ways to create meanings that differ from the original concept. In this case, it is not always clear how the two interpretations are related. In this thesis, I provide a labeling mechanism that allows us to differentiate between different uses of grammatical features in specific contexts while still preserving the possibility of a common core concept of a given grammatical category. In short, the main goal of this thesis is to provide a computationally viable *annotation scheme* for tense and aspect systems.

An *annotation scheme* is a descriptive tool that aims at adding "interpretive linguistic information" to linguistic data (Wynne, 2005). There, the primary data should always be evident. In this thesis, I propose an annotation scheme for secondary data. I.e., I enrich an existing linguistic annotation, namely, morphosyntactic annotation of linguistic expressions. More concretely, I focus on two types of representations: those derived from Lexical Functional Grammar (LFG) and secondarily those derived via Universal Dependencies (UD). For the sake of this thesis, I call these representations tier-0 annotations: they form the base, the groundwork for the approach followed in this thesis.

The semantics in the ParTMA annotation scheme, presented in this thesis, are encoded in the ParTMA template, a set of strictly semantic features. Within this template, semantic features can be layered based on the annotation of the syntax/semantics interface. These layers are distributed across different tiers. In total, I propose three tiers of annotation in this thesis.

The first tier comprises the meaning contributed by morphosyntactic features in isolation, i.e., what I called their core concept before. More specifically, the first tier comprises rules that map morphosyntactic categories directly to semantic concepts. An intuitively understandable example is the English past tense form (simple past). It usually expresses the temporal anteriority of the described situation. Thus, *Mary kissed Tom* means that at some point in the past, there was a situation in which Mary kissed Tom. Thus, pastness is the core concept of the past tense morphology in English.

However, there are morphosyntactic contexts in which the past tense does not necessarily simply express anteriority. Two examples are i) a past tense clause embedded under a past tense attitude verb. In such contexts, a phenomenon called sequence-of-tense (SOT) occurs in which the embedded clause is temporally limited by the matrix clause (Abusch, 1988, 1997). More concretely, the SOT, shown in (1), introduces an ambiguity in embedded contexts. The embedded clause *Alex was sick* has two possible interpretations: an interpretation where the described state (being sick) holds during the saying eventuality, or an interpretation where Alex was sick before the report. A relative future interpretation is ruled out due to the limitation mentioned above. This phenomenon is not universal but occurs in different disguises across the tense/aspect systems of different languages (see Bochnak et al. (2019) for an overview). The second example, (2), illustrates fake past tense in counter-factual conditionals (see Romero (2014) for a discussion of the underlying semantics). To summarize, the past tense seems to lose its function as a marker for anteriority. Instead, it plays an integral role in expressing counterfactuality, i.e., it indicates that the marked eventuality is not true in the actual world.

(1) Jamie said that Alex was sick.
   a. Jamie said: “Alex is sick”.
   b. Jamie said: “Alex was sick”.

(2) If Susan was sick, she would be at home.

In these cases, a simple mapping from form to meaning is not sufficient. For this purpose, I introduce a second tier, tier-2. Tier-2 annotations may be derived by construction rules which can be interpreted as rules that are part of the grammar and that derive meanings from interactions of morphosyntactic and semantic elements in formal representations. Those rules are, in principle, rewrite rules that elevate default interpretations to higher-tier interpretations when applicable, overwriting the tier-1 default interpretation. This is illustrated in the examples (3) and (4). The first example illustrates the abstract idea that the rule wants to capture. Given two eventualities E_1 and E_2 both marked semantically as referring to a past time point, if they stand in a complementizer `COMP` relation, then the embedded eventuality, E_2 is interpreted as either in the past or simultaneous to the matrix eventuality, E_1, rather than simply as in the past relative to the time at which the sentence was uttered. The rule in (4) is a concrete implementation of this rule in a rewrite system presented as part of this thesis. There, `#a` and `#c` refer to a feature matrix describing the semantics of the eventualities E_2 and E_1 respectively (i.e. the rule is reversed compared to its abstract counterpart). The rule captures the `COMP` relation via the `#a ^(tmp>sem>comp) #b` part, which can be read as: check whether the eventuality that corresponds to `#a` (= E_2) is dominated by some matrix eventuality `#b` via a `comp` relation. The second part of the rule (`#b !(sem>tmp) #c TEMP-REF ‘past’`) checks whether this matrix eventuality also encodes past temporal reference. As explained above, this rule yields a revision of the embedded eventuality's temporal reference to an ambiguous form. Encoding the syntax/semantics interface in this way is one of the key concepts discussed in this thesis. Both the abstract representation and the concrete implementation of these rules are explained in detail. These rules are instances of linguistic rewrite rules since they rewrite the `TEMP-REF` feature.

(3) $\texttt{TEMP-REF}_{E1}$ `‘past’` : $\lambda t.t \prec t_0 \wedge$
$\texttt{COMP}(E_1, E_2) \wedge$
$\texttt{TEMP-REF}_{E2}$ `‘past’` : $\lambda t.\lambda t_0.t \prec t_0 \rightarrow$
$\texttt{TEMP-REF}_{E2}$ $\llparenthesis$ $\texttt{present}_t$, $\texttt{past}_t$ $\rrparenthesis$

(4)
```
#a TEMP-REF ‘past’ &#a ^(tmp>sem>comp) #b !(sem>tmp)
#c TEMP-REF ‘past’
==> #a TEMP-REF {past,present} & #a EVAL #c.
```

Linguistic rewrite rules are very powerful in terms of expressiveness and, in this thesis, are used to model meaning shifts based on compositional processes, such as the shift from past temporal reference, to a simultaneous or past reading. This idea is based on extensive work in deriving semantic representations via rule-based systems in the LFG tradition, such as (Bobrow et al., 2007; Crouch, 2005; Crouch and King, 2006). This thesis aims to present a principled way for defining such rules that makes them more accessible for modeling semantic phenomena.

Finally, a third tier covers pragmatic reasoning, e.g., anaphora resolution and extra-linguistic information. This layer is different from the previous two because it is difficult to put into rules. I will only cover part of what tier-3 is designed for in this thesis, namely, pragmatic default inferences for tense/aspect categories, which can also be modeled in terms of linguistic rewrite rules. Thus, tier-3 is a significant avenue for future work.

The resulting annotation scheme consists of two components: the first component is a semantic annotation for which each feature is typologically attested and formally encoded. This latter condition is crucial. It allows me to provide a concrete semantics for the ParTMA annotation scheme. A major contribution of this thesis is the implementation of a fragment of semantics that can interpret ParTMA style annotations and produce appropriate formal semantic representations. The second part of the annotation scheme consists of a set of principles for defining the syntax/semantics interface for tense/aspect that models the corresponding features in a layered system that captures core contributions of certain categories' core contributions and their variation in the morphosyntactic and semantic context.

This means there are three main contributions of this thesis: i) a cross-linguistically viable feature space for tense and aspect, ii) a concrete semantic framework for interpreting the features modeled in the ParTMA annotation scheme, and iii) a specification for designing a syntax/semantics interface using linguistic rewrite rules. Both the syntax/semantics interface and the semantic framework come with a concrete implementation combining linguistic rewriting and Glue semantics, a computationally viable semantic formalism popularized by research in the LFG community (Dalrymple, 1999).

1.1 Putting things into perspective

This section serves to put this thesis into perspective by explaining its key influences. I start the section by discussing how the present work relates to the topic of semantic annotation. Then, I move to the related topic of grammar development. Finally, I explain the linguistic phenomena that inspired this thesis, namely tense and aspect.

1.1.1 Semantic annotation

Semantic analysis is a crucial component for natural language understanding (Lev, 2007) which in turn finds applications in several NLP and related systems, i.e., information retrieval, machine translation, question-answer systems, artificial intelligence, and so on. The goal of semantic annotation is, simply put, to facilitate the semantic analysis of (primary) data. Consider the following quote:

"The idea behind annotating a text, which dates from long before the digital era, is to add information to a primary text information in order to support its understanding. The semantic annotation of digital source text has a similar purpose, namely to support the understanding of the text by humans as well as by machines."(Bunt, 2015, p. 3)

Investigations like the one by Snow et al. (2008) make clear that simplicity is vital. It may seem that the application of a particular framework to a broad amount of data seems trivial – most sentences behave straightforwardly boring after all – however, formal researchers are often very concerned with the embedding of their idea in their theory of choice. It needs to be made clear that the formal apparatus used, e.g., in formal semantics, makes it highly resource-consuming to annotate large amounts of data. Thus, formal semantic papers usually only focus on a small sample of very complex phenomena with the assumption that the extension of the theory to more simple examples is trivial (or has been proven elsewhere). This is not always the case, especially when considering various linguistic categories, as is desired in interoperable semantic annotations. Thus, careful consideration of the relation between annotation and the underlying assumptions about the semantics is necessary.

Much of the work in linguistic annotation and, in particular, semantic annotation is driven by ISO-standards (Ide and Pustejovsky, 2017). The appeal to the International Organization of Standardization aims to guarantee a cross-linguistically viable and interoperable annotation of linguistic data. Interoperability can be understood as compatibility between different resources. There are several pitfalls when trying to make resources compatible (Bunt, 2015). For example, a temporal annotation scheme and a spatial annotation scheme should not use the same labels to refer to different concepts. Yet another formulation of this constraint is that different annotation schemes should be mutually consistent. Although the annotation scheme presented in this thesis does not (yet) strictly adhere to the ISO standards of linguistic annotation, it takes the core concepts of linguistic annotation as stated in Ide and Pustejovsky (2017) and Bunt (2015) to heart. Thus, cross-linguistic viability and integration with a broader framework of linguistic annotation such as that developed by the SemAF effort (Bunt, 2015) are essential considerations that have affected the development of the present annotation scheme. The present annotation scheme follows the principles mentioned in the literature above by adhering to the CASCADE model for developing semantic annotation schemes.

Developing an annotation requires several steps. The CASCADE model for developing semantic annotation schemes provides a concrete guideline for following these steps (Ide and Pustejovsky, 2017). (5) illustrates the proposed necessary steps for an annotation system. The first point is to provide a *conceptual analysis* of the phenomena to be captured. This thesis can build on a considerable amount of previous work regarding the conceptual analysis of tense and aspect. The contribution of this work is to filter out the innovations that may be useful for temporal annotation tasks and then translate them into an *abstract syntax*, the second step. This, in principle, means that we generate a consistent formal model of annotation objects, e.g., "what are the properties of events and what are the properties of time intervals? What kind of relations do we want to allow between these objects? Temporal relations, causal relations, or some other kind of relation?" and so on. Each of the syntactic properties that are defined also requires a consistent *semantics*. One can apply various semantic systems to the same abstract syntax. This may depend on various factors such as semantic complexity or embedding in downstream applications which require a specific formal semantic representation. Once the system is formally sound, it needs to be translated into a machine-readable and processable representation: the *concrete syntax*. The concrete syntax should be in one-to-one correspondence with the abstract syntax to avoid inconsistencies in the semantics.

(5) a. Conceptual Analysis: Formulate a conceptual view of the information to be captured in annotations. This results in an abstract data model or 'metamodel'.

b. Abstract Syntax: Articulate the conceptual view in the form of an inventory of basic concepts and a formal specification of the possible ways of combining these elements in set-theoretical structures like pairs and triples, called annotation structures.

c. Semantics: Provide a formal semantics for the structures defined by the abstract syntax.

d. Concrete Syntax: Specify a representation format for the annotation structures defined by the abstract syntax;

(Ide and Pustejovsky, 2017)

In this thesis, I not only differentiate between abstract syntax and concrete syntax but also between abstract and concrete semantics. This means I use a compositional semantic framework informed by theoretical formal semantics for dealing with tense/aspect as a foundation to implement a concrete semantics. The concrete semantics differ from the abstract semantics because fully fleshed out computational counterparts do not always exist for the proposals developed in the tense/aspect literature. Ultimately, this thesis presents a concrete semantics that shares labor between a rewrite component and a composition component for dealing with different aspects of formal semantic modeling. Sharing the labor in such a way allows for a formally sufficiently expressive system that can be mapped onto a conceptually simple annotation scheme.

1.1.2 Grammar development

The ParTMA annotation scheme is also heavily influenced by ideas from Lexical Functional Grammar both from a theoretical as well as a computational perspective. The Lexical Functional Grammar movement is one of the driving forces in the development of deep linguistic grammars under the international ParGram project (Butt et al., 2002). This effort is concerned with developing cross-linguistically parallel (i.e., comparable) grammars for morphosyntactic analysis. It is one of the main influences, or more specifically, the raison d'être for this thesis. This influence starts with the fact that the ParTMA annotation scheme does not strictly propose a new annotation of primary data but that it provides a new layer of annotation for the existing morphosyntactic analysis provided by computational grammars written in the Xerox Linguistic Environment (XLE; Crouch et al. (2017)).[1] The philosophy behind LFG, in brief, is that a grammar provides a function that maps a linguistic expression to a meaning representation (Asudeh, 2006). This function may be decomposed into several components that describe particular elements of the grammar of a language in isolation. This is illustrated in Figure 1.1. It provides an LFG analysis for the sentence *John visited Mary*. As can be seen there, the analysis is represented in terms of two different formal representations: a tree structure and an attribute/value matrix. The two structures are called c(onstituent) structure and f(unctional) structure,

[1] The thesis project's primary idea was to improve the treatment of tense and aspect in XLE. In this thesis, in particular, in chapter 5, I show that the annotation scheme is designed to be more general in its purpose. Concretely, I show this by considering Universal Dependencies as an alternative morphosyntactic base for annotating the syntax/semantics interface and the semantics themselves.

respectively. The c-structure encodes syntactic constraints on the surface form of a linguistic expression. On the other hand, the F-structure encodes certain syntactic roles, called grammatical functions and functional features required for further semantic processing.

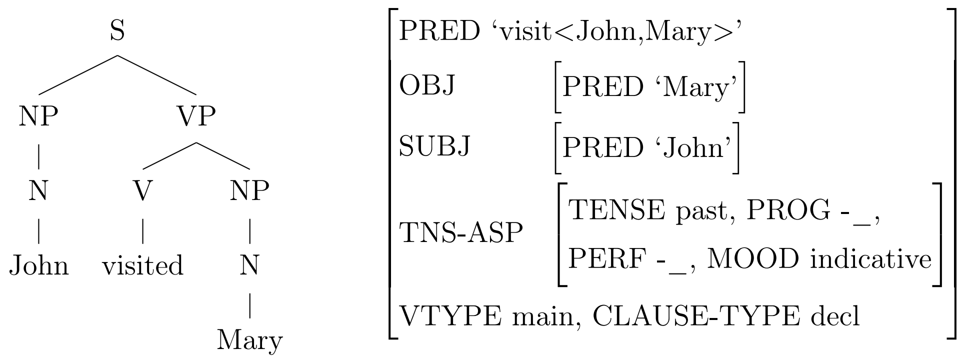

Figure 1.1: C- and F-structure: *John visited Mary*

As the f-structure in Figure 1.1 shows, LFG encodes tense and aspect annotation predominantly in terms of functional features. In this thesis, I analyze the viability of these functional features concerning their capability to allow semantic interpretation and conclude that functional tense/aspect features confound syntactic and semantic information in an unprincipled way. The present thesis argues for a treatment of functional features as exponents of similar morphosyntactic structures rather than as quasi-semantic features to remedy this situation. As a result of this reasoning, I propose a separate semantic representation that captures purely semantic information on tense and aspect: the ParTMA template.

The ParTMA template is an attribute/value structure, similar to LFG's f-structure but also similar to structures used in semantic annotation more generally. It distinguishes itself from the former by attributing a concrete formal semantic contribution for each feature, and it distinguishes itself from the latter in that the underlying semantics preserve compositional semantic processes devised in the formal semantic literature. Concretely, the semantics of the ParTMA annotation scheme are grounded in a description-by-analysis Glue semantics approach with the goal of integrating the resulting semantics in the broader picture of semantic analysis within computational LFG.

When introducing a new level of representation to the LFG formalism, it is necessary to provide an exact formal description of how this intermediate representation relates to the other representations in LFG (Kaplan, 1995b). In general, the co-descriptive approach is popular in LFG. This implies that different structures are built up more or less in parallel. To keep the ParTMA template and the corresponding annotation more modular, I define a mapping from (c- and) f-structures in terms of a set of rules that apply to fully formed syntactic structures. I interpret these rules as a part of the syntax/semantics interface. Thus, the ParTMA template can be understood as part of a semantic projection in LFG.

Separating the ParTMA template from the LFG architecture in this way also allows me to make it more broadly usable. Thus, The ParTMA template could be populated with semantic features based on other syntactic analyses, e.g., Universal Dependencies. Furthermore, it would be possible to devise a mapping directly from primary data to the semantics (in principle, the same possibilities that exist with other annotation schemes: manual annotation or automated annotation). Only the former two approaches with a concrete rule-based mapping from syntax to semantics are considered in this thesis. However, it is important to keep in mind that the ParTMA template is an independent tool for annotating the semantics of tense and aspect.

1.1.3 Tense and aspect: form and meaning

This work is about meaning encoded in linguistic expressions, particularly verb phrase (VP) forms, and their embedding in the linguistic context. When looking at the world's languages, there is a broad spectrum of possibilities. Below are examples from the ParTMA treebank. These sentences have been handpicked to illustrate specific trends in the marking of tense and aspect. They are parallel in their semantics in that they may occur in the same context.[2] In this case, the target semantic feature that these sentences represent is that of *past tense*. More specifically, these sentences are about some **past** occurrence. This can be intuitively tested by combining them

[2]This does not mean that they all are restricted to the same context.

with corresponding temporal expressions as illustrated in table 1.1.[3]

past tense	**present tense**	**future tense**
—	– *now* –	— >
e.g. once,	*e.g. now,*	*e.g. tomorrow,*
a year ago,	*currently,*	*next thursday,*
yesterday	*at present*	*soon*

Table 1.1: Testing for tense meaning

(6) a. The people killed the king.

b. Peter was writing letters. English

(7) a. Peter sedang menulis surat.
Peter PROG write letter
'Peter was writing letters'

b. Orang membunuh raja.
People kill king
'The people killed the king.' Indonesian

(8) a. peter xat likH rahA tHA
Peter letter.SG write PROG be.PAST.M.SG
'Peter was writing letters'

b. logoN=ne baadshaah=ko maar daalaa
people=ERG king.M.SG=ACC hit put.PRV.M.SG
'People killed the king.' Urdu

(9) a. Piotr pisał listy
Peter write-PAST.IMPV letter.3.M.SG
'Peter was writing letters.'

b. Ludzie zabili króla
People.3.M.PL kill.PAST.PRV KING.3.M.SG
'People killed the king.' Polish

[3] While this test serves the present purpose of providing the reader with an intuitive idea about tenses, this test is also tentative. See chapter 3 for a more elaborate understanding of how to identify tenses.

The sentences from different languages above express roughly the same meaning as the sentences in the English example; however, the strategies vary significantly in terms of the word forms used to express the meaning of pastness that is part of each example. I will not go through all examples as the strategies used will be discussed in depth throughout this dissertation. However, I want to clarify some major and some more minor differences across languages. In particular, if the reader investigates the glossing for past tense – intuitively the morphosyntactic exponent of past meaning – they will realize that not all languages seem to obligatorily express the meaning of pastness in their verb form. For example, Indonesian simply omits past morphology altogether. On the other hand, Urdu seems to use its *Past* marker, an auxiliary in some cases but not in other cases – the tense there seems optional.

The sentences also pattern with the English distinction between the *progressive* form and the *simple past* form. Notably, the *Perfective (Prv)* glossing seems to be used when in English the *simple past* is used, while either the *Progressive (PROG)* or the *Imperfective (Impv)* glossing are used in the other languages to represent a meaning similar to the English *progressive*. While Indonesian does not seem to express tense, there is some auxiliary aspectual marking in the form of *sedang* in place. Urdu also has a dedicated *progressive* form. Lastly, the Polish examples directly express tense and aspectual marking in their verb forms. The label *aspect* or *aspectual* covers the forms (*perfective*, *imperfective* and *progressive*). What is their function? It is clear that tense primarily tells us whether a linguistic expression is about the past, the present, or the future, but aspect is more subtle. Some readers may infer from the examples and my remarks that the *imperfective* and the *progressive* share at least some meaning properties. This is indeed the case. The prototypical meaning attributed to the imperfective is that of ongoingness, i.e., it applies in situations where some event, e.g., the *writing of letters* or the *killing a king*, is in the process of happening.

In contrast, the perfective – intuitively – is used in sentences that describe completed events. The examples above do not express whether Peter wrote all the letters he wanted to write or how many he has written. However, in the

	durative	change	endpoint
state	+	-	-
activity	+	+	-
achievement	-	+	+
accomplishment	+	+	+

Table 1.2: Aktionsart

second type of example, the king is dead since the killing event is described as complete by merit of the perfective aspect (see (8b), (9b); in the English example, we infer the completeness of the event, or in other words: the default interpretation of the simple past is a perfective interpretation). Based on these examples, I adopt the view that grammatical aspect encodes whether an event is expressed as a whole (perfective) or only partially (imperfective). This view has been mainly coined by (Comrie, 1976; Smith, 1991) and is widely accepted in the corresponding literature.

The notion of state and action reflect distinctions made at the level of inner or lexical aspect. This can be understood as the third level of semantic interpretation following the level of tense and the level of grammatical (outer) aspect that have been introduced so far. The examples above present two distinctions: momentary actions (*(to) kill*) and prolonged actions (*(to) write a letter*). Together with states (e.g. *(to) love,(to) know*) and actions without inherent endpoint (e.g. *(to) ride a bike, (to) sing*) these distinctions build the groundwork for defining the lexical or inner aspect of verbs and verb phrases. These distinctions are based on Vendler's (1957) seminal work on so-called Aktionsarten, i.e., categorizations for different classes of verbs. As explained above, three properties are crucial: duration (momentary vs. prolonged), change (dynamic vs. static), and the existence of an inherent endpoint. The four classes proposed by Vendler are shown in 1.2.

As the name suggests, lexical aspect is usually grounded more in the lexicon than in the morphosyntax. Thus, differences in the surface structure of linguistic expressions with different verb classes are less apparent. However, as more recent research on the topic has shown, several intricate morphosyntactic reflexes are associated with this deep level of aspectual analysis, which is explored in this thesis.

Overall, the languages presented above vary more or less strongly regarding the richness of their tense/aspect morphosyntax. In addition, to surface structural cues, it is also necessary to consider deeper levels of encoding, particularly encoding at the lexical semantic level, to cover inner/lexical aspect. The richness also affects the meaning that they can express explicitly. Thus, one can assume that the number of contexts where the Polish sentence in (9b) may be used is smaller than the number of contexts of the Indonesian counter-example, (7b). This observation is borne out since the Indonesian example can occur not only in contexts that require a past tense interpretation but also in contexts that allow for present or future tense readings. In other words, the Polish sentence can only be used to describe a past occurrence, but the Indonesian sentence may also be about the present or even the future. How do we know then that the Indonesian example, (7b), is about the past? In fact, we cannot answer this question because it depends on the context in which the sentence is uttered. For example, if the sentence is uttered during a history lesson about Charles I of England, it is most certainly about a past situation. However, uttered by a fortune teller talking to a worried queen, it is probably about the future. Similar reasoning can be used on the level of aspectual marking in Indonesian. As the name suggests, lexical aspect is usually grounded more in the lexicon than in the morphosyntax. Thus, differences in the surface structure of linguistic expressions with different verb classes are less apparent. However, as more recent research on the topic has shown, several intricate morphosyntactic reflexes are associated with this deep level of aspectual analysis, which is explored in this thesis.

I want to repeat that the examples above only illustrate a small part of these languages and their tense and aspect systems. They only serve to express some trends across languages to give the reader a rough idea of the topics of this thesis. I have outlined three essential levels to the analysis of tense/aspect: i) tense, ii) grammatical (outer) aspect, and iii) lexical (inner) aspect. In chapter 3 and 4 the topics outlined here will be discussed in more detail, paying attention to the existing research, in particular, the typological and formal approaches to analyzing tense and aspect.

1.2 The ParTMA annotation scheme

As already mentioned in section 1.1.2, the ParTMA annotation scheme is part of an effort to provide formally sound descriptions for linguistic expressions that are cross-linguistically parallel or as parallel as they can be. This effort has been dubbed the ParGram project (Butt et al., 2002). This thesis is part of the project *Tense and Aspect in Multilingual Semantic Construction*, which I refer to by ParTMA (for parallel tense, aspect, and mood).[4] The project identified a shortcoming concerning the cross-linguistic analysis of tense and aspect in computational linguistics by virtue of a review of the TimeML effort (Pustejovsky et al., 2003, 2010, 2002). The main findings of the review are that TimeML lacks a proper treatment of tense and aspect in particular from a semantic perspective (see chapter 2). At the same time, the treatment of tense and aspect in the ParGram grammars was faced with its own challenges; however, ParGram focused mainly on dealing with morphosyntactic rather than semantic issues due to the lack of a proper semantics component in the XLE, in which all ParGram grammars are developed. One issue that both the TimeML and ParGram have in common is the conflation of morphosyntactic and semantic properties in their features. Both approaches claim only to encode those distinctions that appear overtly in the morphosyntax. However, both approaches also assume that this annotation is sufficient for further semantic processing, as already illustrated in the introduction of this chapter. In this thesis, I propose to separate structural and meaning annotation for these categories strictly. As is common in typological approaches to tense and aspect (Bybee, 1985; Dahl, 1985, 2000), I suggest tense and aspect categories as mappings from form to meaning. I hope this approach clarifies the role of the morphosyntax and the semantics concerning the formation of linguistic expressions. In particular, I think that such an approach makes it easier for a formal description to remain cross-linguistically viable since meaning and structural reflex can be encoded more precisely. As stated above, this is one of the main pitfalls for the existing proposals to semantic annotation

[4] Although mood is not discussed in this thesis; it will be considered for future work.

considered in this thesis.

In this elaboration, I also mentioned a broader goal of the ParTMA project. Namely, providing a semantics for ParGram grammars. This grand goal has been reduced to a more manageable size by focusing only on the categories of tense and aspect. Nonetheless, the results of this thesis and additional work within the ParTMA project serve as an important milestone for introducing a proper semantic analysis to the international ParGram effort and thus to the grammar development community. This is necessary to provide linguistically sophisticated analyses for typologically diverse languages.

1.3 Roadmap

Overall, this work bridges the gap between formal and computational linguistics. Certain issues, in particular tractability or complexity, prevent the direct percolation of formal semantic insights to the domain of natural language processing, one of the main pitfalls being the combinatorial explosion puzzle, which leads to sometimes excessive ambiguity, both intended and spurious (Bunt, 2008; Ebert, 2005; Lev, 2007). Furthermore, semantic frameworks often rely on notions that are difficult to model computationally, such as the notion of possible worlds and modal reasoning in general (however, see Crouch (2005), Crouch and King (2006), and Kalouli and Crouch (2018) for some influential and promising approaches to deal with this issue). Beyond these issues, another obvious problem with formal, in particular, logical models is that of robustness (Martin and Jurafsky, 2009).

In the first half of the next chapter, I discuss the development of standardized semantic annotations as supported by the International Organization for Standardization (ISO). They try to remain independent of formal linguistic research where possible and are mainly concerned with temporal relation typing, i.e., defining temporal objects (events, time intervals) and the temporal relations (before, after, simultaneous) that hold between these temporal objects. There, I will specifically point out that there is a certain state of ignorance towards interpreting tense and aspect features, especially from a cross-linguistic perspective.

The second part of the next chapter (starting at section 2.3) is concerned with efforts beyond the standardized temporal annotation provided by ISO-TimeML. I will discuss the computational treatment of both tense and aspect morphosyntax and semantics, both of which have attracted the interest of researchers for many years. Several people have proposed computational systems to deal with tense and especially its relation to the context, e.g.Asher and Wada (1988), Partee (1984), Passonneau (1988), and Steedman (1995). However, also the purely morphosyntactic treatment has been and still is of interest for computational linguists (Dorr, 1992; Ramm et al., 2017). Most importantly, tense and aspect have proved a challenge for ParGram (Butt et al., 2002) which gave rise to the more specialized ParTMA.

In summary, chapter 2 elaborates on section 1.1 in more detail and introduces the core concepts relevant for this thesis. These concepts are first explored in action in chapter 3, where the linguistic category of tense is discussed in-depth. The chapter is divided into three parts. First, it provides a background on the typology of the category of tense. There, data from several typologically unrelated languages is taken into account. This should give the reader an understanding of the commonalities and differences in tense systems across languages. In the second part of this chapter, this knowledge is translated into a formal system. This involves a cross-linguistically appropriate morphosyntactic treatment but mainly focuses on the semantic formalization of tense and aspect and the syntax/semantics interface. The issue of the syntax/semantics interface is carried over to the third part of chapter 3. There, the annotation scheme that is presented as part of this thesis is explained in detail. This involves several rules and conventions of annotation that are necessary to make the system computationally viable. At this point, the annotation scheme is presented qualitatively with selected examples from the literature on tense.

Chapter 4 follows the same formula as Chapter 3. It discusses the cross-linguistic picture of aspect, including grammatical aspect and lexical aspect. This chapter focuses more on explaining the intricacies of all these categories since annotation principles remain more or less the same. Nonetheless, the understanding aspect is derived from the typological landscape and its adap-

tion in formal linguistics. One novelty in this chapter is altering the semantic framework from a pure interval semantics to a situation semantics. This does not have major repercussions on the previously introduced framework in terms of tense, but it helps us describe aspectual information better.
In Chapter 5, I present a concrete implementation of the annotation scheme described on an abstract level in Chapters 3 and 4. Concretely, I introduce the specification for a linguistic rewrite system that can be applied to abstract syntactic representations. Abstract syntactic representations are graph structures that are derived from various kinds of syntactic input. This is exemplified by illustrating how LFG's syntactic representation and Universal Dependency representation can be translated into abstract syntactic graphs and is based on work by Ide and Bunt (2010). This rewrite system is used to implement the syntax/semantics interface annotation. By combining the linguistic rewrite component with a formal semantic framework, in particular, Glue semantics, a semantic interpretation component for the ParTMA template is developed. This component translates an annotated ParTMA template into a Glue semantics proof (i.e., a semantic object that provides a set of semantic representations and a proof that illustrates how these representations have been compositionally assembled) by extracting partial glue semantics representations based on a description-by-analysis approach Andrews (2008) and Crouch (2005). The implemented system is also illustrated in terms of selected examples discussed in previous chapters. This chapter references the previous sections but should be self-sufficient enough for people who are simply interested in implementing a syntax/semantics interface and semantic interpretation via linguistic rewriting. The final chapter, Chapter 6 concludes this thesis. It summarizes the goals achieved in this thesis, and aveneues for future work.

Chapter 2

Background

In this chapter, I define the extent of linguistic expressions for which I provide an annotation scheme (section 2.1). Then I introduce TimeML, an annotation standard for temporal semantics of which tense and aspect information is one component. I identify challenges for the current TimeML standard with respect to tense and aspect (section 2.2). In the rest of the chapter, I introduce the tools with which I want to remedy this situation. Concretely, I compare TimeML to another annotation standard that is grounded in deep linguistic analysis, namely that provided by the parallel grammar (ParGram) project (starting at section 2.3). Two main points will be discussed: how the two approaches deal with cross-linguistic variation and the level of linguistic analysis. Finally, this chapter introduces some of the fundamental building blocks of this dissertation, specifically, some formal tools necessary to explain the presented annotation scheme and its feature space.

2.1 The extent of tense and aspect marking

Extensive typological work has shown that tense and aspect categories can be expressed in a number of ways (Bybee et al., 1994; Bybee, 1985; Dahl, 1985). Following Dahl (1985), three main distinctions are relevant: inflectional, derivational, and periphrastic marking, and, furthermore, combinations of these marking strategies. For example, in English, past tense may be expressed by the use of the inflectional *simple past* form. The perfect and the progressive aspect can be combined with tense. For this purpose, they have a periphrastic component and an inflectional component. The periphrastic component combines with tense, whereas the inflectional component is inflected for the corresponding aspect form (i.e., present vs. past participle). On the other hand, the future tense is purely periphrastic in the sense that aspectual modification of the main verb is not required. These properties are exemplified in (1). However, languages may lean more towards periphrastic tense and aspect features, like Indonesian, a (mostly) tenseless language that only provides a periphrastic future tense, and several other auxiliaries that express, e.g., perfect and progressive aspect. This is exemplified in (2). There, it is also shown that there is no obligatory tense distinction, as is necessary for English. More concretely, unmarked verb forms can receive either past, present, or future tense interpretations. This means that even the future marker *akan* is not obligatory to express a future interpretation (Arka, 2013; Zymla, 2017).

(1) a. Peter **wrote** a letter. *simple past*

b. Peter **has written** a letter. *present perfect*

c. Peter **is writing** a letter. *present progressive*

d. Peter **will write** a letter. *future tense*

(2) a. Peter **menulis** se.buah surat
Peter **AV.write** one.CL letter
'Peter wrote/is writing/will write a letter.' *no tense/aspect*

b. Peter **akan menulis** se.buah surat
Peter **FUT AV.write** one.CL letter
'Peter will write a letter.' *future tense*

c. Peter **telah menulis** se.buah surat
Peter **PERF AV.write** one.CL letter
'Peter has/had written a letter.' *present/past perfect*

d. Peter **sedang menulis** se.buah surat
Peter **PROG AV.write** one.CL letter
'Peter is/was writing a letter.' *present/past progressive*

(3) a. Peter **pis-ał** list
Peter **write.IMPV-PAST** a letter
'Peter **was writing** a letter' *past imperfective*

b. Peter **na-pis-ał** list
Peter **PRV-write-past** a letter
'Peter **wrote** a letter' *past perfective*

c. Peter **prze-pis-ał** list
Peter **PRV-write-PAST** a letter
'Peter **rewrote/copied** a letter' *past perfective*

d. Peter **pis-ze** list
Peter **write.IMPV-PRES** a letter
'Peter **is writing** a letter' *present imperfective*

e. Peter **na-pis-ze** list
Peter **PRV-write-PRES** a letter
'Peter is going to write a letter (completely)' *future perfective*

f. Peter **będzie pis-ał** list
Peter **be-PRES write.INF** a letter
'Peter will be writing a letter' *future (imperfective)*

The third main variation in tense and aspect marking is derivational in nature. Aspect systems prototypically represent this in Slavic languages. Consider the Polish examples in (3). The aspect system in Polish provides a distinction between *perfective* and *imperfective* aspect. As is typical in Slavic languages, the perfective is (usually) formed derivationally by applying an affix to the imperfective verb stem ((3)a vs. (3)b). For this purpose, usually, multiple affixes exist, some of which contribute additional semantic content. Compare (3)b with (3)c, where the prefix perfectivizing prefix *na-* is replaced by the perfectivizing prefix *prze-*. In this case, *na-* is the default perfectivizing prefix and does not contribute any other semantic meaning. The alternative prefix *prze-* suggests that the action is repeated or copied.[1][2]

[1]Elicited from a native speaker of Polish.

[2]The Polish aspectual system is demonstrated in a very simplified manner here. For a detailed overview, I refer the reader to Kipka (1990), or more recently Młynarczyk (2004). Furthermore, Dickey (1997) provides an overview over Slavic aspect as expressed

In Polish, like in English, tenses are expressed via inflection, except for the future tense. However, for the future tense, two main strategies exist: i) similar to English, Polish uses an *auxiliary* (*(to) be*) and an infinitival form (which is similar to the past imperfective form) to express future tense. ii) By combining the present tense and a perfective marker, a future interpretation is elicited that commits to the full completion of the corresponding action in the future (see (3)e).

To summarize the exploration of the examples above, the main point is that tense and aspect categories can surface in a variety of ways in linguistic expressions, but, for each category, one can find languages that express it and languages that do not (necessarily) express it overtly in their surface structure (Dahl, 1985).

2.1.1 The role of tense and aspect in the grammar

This thesis mainly focuses on tense/aspect marking in matrix clauses. This is only the tip of the iceberg that is tense/aspect marking cross-linguistically. Although the scope of this thesis does not allow us to explore this iceberg fully, it is still crucial to understand what lies beyond the surface of the relative safety of tense/aspect marking on verbs in matrix clauses.

Firstly, tense and aspect marking occurs in different kinds of embedded clauses and modifying, i.e., adjunct clauses. This is particularly relevant from the perspective of tense, as numerous researchers have shown (Abusch (1988), Bochnak et al. (2019), Grønn and von Stechow (2010), and Kusumoto (1999), to name a few). Thus, this thesis touches on the topic of tenses in embedded clauses in chapter 3.

Embedded and modifying clauses can be distinguished between finite and non-finite clauses. Finiteness is a crucial property of languages mainly used to identify what constitutes a full clause. Thus, non-finite forms are seen as incomplete, defective, or dependent (Lowe, 2019). However, the conditions

in different Slavic languages providing a contrastive picture of the Polish aspect. Bittner (2014) provides a more general cross-linguistic comparison between the Polish tense/aspect system and other types of tense/aspect systems across the world. Aspectual prefixes are also picked up again in Chapters 3 and 4 with a focus on their semantic properties.

that need to be met to form a full clause vary across languages (Bittner, 2014; Lowe, 2019). I assume, based on Klein (1998, 2006) that there is a crucial relation between anchoring in time and finiteness, but the exact nature of this relation is not crucial for the goals pursued in this thesis.[3]

As also proposed in Lowe (2019), my hope for the finite/non-finite distinction is that the semantics of non-finite verb forms can generally be derived from their finite counterparts. However, how feasible this view is from a compositional semantic point of view remains to be seen, i.e., is left to be explored in future work.[4]

The role of tense/aspect is further complicated by the introduction of inflectional forms to the nominal domain. This has been already acknowledged by Comrie (1976) and Vendler (1957) and has since been explored in great detail, for example by Koptjevskaja-Tamm (2002, 2008), Nordlinger and Sadler (2004), and Zucchi (2013) just to name a few. As shown in Zucchi (2013) for English, and in Nordlinger and Sadler (2004) from a cross-linguistic perspective, nominal tense/aspect marking provides puzzles for both morphosyntactic analyses and semantic analyses. They are concerned with, for example, whether such nominalized verb forms take arguments or whether they express eventualities or full propositions. Generally, nominal tense/aspect marking provides further challenges and complications that build on top of the previous layers of complication provided by (non-)finiteness and clause type.

All in all, this thesis aims at providing a solid foundation for expeditions that venture deeper into these deeper layers of tense/aspect form and meaning, in particular, from a computational and cross-linguistic perspective. The need for such a solid foundation is made more clear sections 2.2 to 2.4, where I discuss various implementations of tense/aspect annotations.

[3]I do not argue that all languages concretely anchor the content of a linguistic expression in time via finiteness, but rather that the content is anchored in a model in such a way that its temporal orientation (past, present, future) can be inferred.

[4]A particular concern here is the scope of temporal and aspectual operators in combination with other properties of finiteness such as subjecthood and the interaction with different kinds of quantifiers.

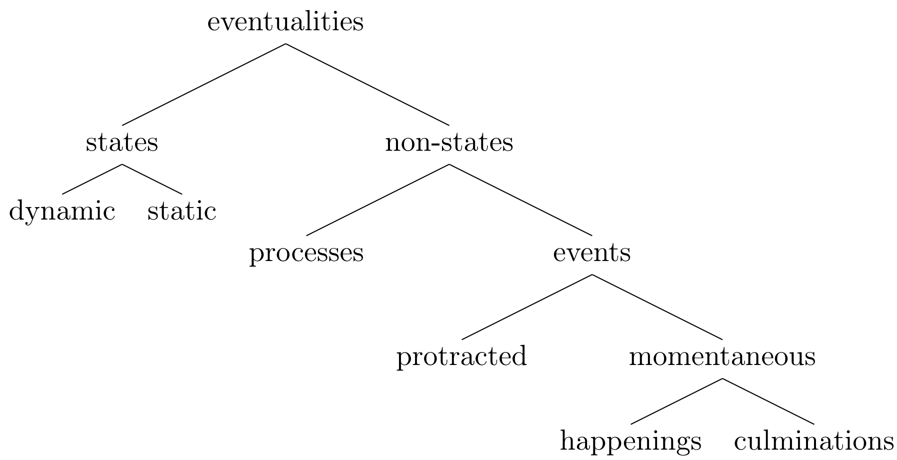

Figure 2.1: Bach's (1986) distinction of eventualities

2.1.2 Tense and aspect – intervals and eventualities

In the previous section, I have mentioned that tense/aspect marking applies to both verbs and nouns, but what is the underlying semantic commonality making this possible? Generally, it is the fact that language is used, among other things, to describe occurrences of actions and states of affairs. This idea has been particularly fleshed out by Davidson (1967), who proposed that we can refer to actions in the same way that we can refer to individuals. This gave rise to what is now typically called event semantics. Event semantics have been closely related to the analysis of tense and particularly aspect since their inception (see for example: Bach (1986), Krifka (1998), Parsons (1990), and Vlach (1981)), however, they are also associated with the question argument structure and other issues (Grimshaw, 1990; Hale and Keyser, 1993; Levin and Hovav, 2005; Ramchand, 2008). As shown in chapter 4, the various challenges surrounding event semantics are not completely independent.

At this point, the notion of eventuality (which I adopt from Bach (1986)) is of crucial importance for the following discussion. As shown in the tree in Figure 2.1, *eventuality* is a generic term that subsumes both states and actions (here: non-states), both of which themselves are subject to more fine-grained distinctions, which are discussed in detail in chapter 4.

Eventualities are mainly derived from the descriptive content of a linguistic expression and the expressed argument structure; thus, the sentence *Addison hugged their friend* describes a hugging eventuality in which Addison is the actor, and their friend participates as the recipient of that hug.
As illustrated in the previous section, tense and aspect is not directly based on the descriptive content but rather encoded morphosyntactically, i.e., via derivation, inflection, or periphrasis (Dahl, 1985). They encode, at least, relations between eventualities and time intervals. Time intervals are semantic objects with a specific start and endpoint, which makes it possible to locate and order them on a timeline (Allen, 1990).
Generally, the role of tense and aspect is to relate eventualities (or rather the time during which they take place (actions) or hold (states)) to specific time intervals expressed and implied in language. Through this, eventualities are anchored in time, situating them in the past, present, or future and relative to other eventualities. The focus of this thesis lies on this first role of tense and aspect. However, tense and aspect also affect the relative ordering of consecutively mentioned eventualities (e.g., in a narrative structure Lascarides and Asher (1993)). That is, eventualities can be temporally ordered in the same fashion as time intervals.
Overall, time intervals and eventualities are crucial components of temporal semantics and the semantics of tense and aspect. This is also reflected in the most prominent effort for annotating temporal semantics, TimeML.

2.2 TimeML: Annotation of Temporal Semantics

At least one goal of a semantic annotation scheme for tense and aspect is to provide a feature set that allows one to generate semantic representations of eventualities, a goal that is shared with the temporal annotation scheme, (ISO-)TimeML (Pustejovsky, 2017; Pustejovsky et al., 2003, 2010, 2002).[5] In fact, the inspection of the TimeML annotation standard and its treatment of tense

[5] Other goals include the anchoring of eventualities in time, space, and the domain of possible worlds.

and aspect has provided the main inspiration for developing the ParTMA annotation scheme presented in this thesis. I lay out the results of this inspection, which are qualitative in nature, as a starting point. Since TimeML is an annotation scheme of temporal properties, the main focus also lies on the temporal properties of tense and aspect features. Before properly starting this section, consider that the result of this inspection confirmed what the TimeML creators (Ingria and Pustejovsky (2002)) essentially already acknowledge: The TimeML annotation standard only provides a basic template for tense and aspect annotation. In this section, I want to address three questions:

(4) What tense and aspect features are available in the TimeML specification for English and across languages?

(5) How well does TimeML cover the temporal properties of semantic tense and aspect features?

(6) How does TimeML deal with cross-linguistic coverage?

To answer these questions, I first discuss the general architecture of TimeML. I focus on the treatment of tense and aspect, concretely, in the English implementation of TimeML. Afterward, I introduce some related efforts developed based on the English TimeML implementation. Finally, I discuss some implementations in other languages and point out some of the weaknesses of TimeML concerning tense and aspect features from a cross-linguistic perspective.

2.2.1 The development of TimeML

TimeML is an important milestone in the development of temporal annotation schemes. In fact, TimeML has pushed the development of semantic annotation in general and is often used as a primary example for new developments in semantic annotation. A myriad of papers reflects this, but for reference, I refer to, e.g., Bunt (2015) and Ide and Pustejovsky (2017), which provide an overview on the topic. Furthermore, TimeML provides the foundation for many NLP tasks related to temporal reasoning, most dominantly, the

TempEval series (Bethard et al., 2015, 2017; UzZaman et al., 2013; Verhagen et al., 2007, 2010).

The development of TimeML can be divided into two stages: the development of TimeML and the development of ISO-TimeML, the latter being an industrial and commercial standard that is not readily open to academic research. Nonetheless, the development of both stages has been documented in the academic literature. I will discuss both stages separately and ask the reader to be wary that the original TimeML standard is still more widespread in the computational linguistics (and NLP) community, mainly due to the commercialization of the latter. Thus, much research is directly based on or inspired by the TimeML standard. Although not central to this thesis, I will mention the main differences between TimeML and ISO-TimeML in the following sections, where appropriate.

2.2.2 TimeML: general structure

This section provides an overview of the linguistic information encoded in TimeML annotations and how this encoding is done. As described briefly in the introduction, semantic annotation schemes consist of a syntax, basically the annotation itself and a semantics, the interpretation of the annotation scheme that is used for semantic tasks, e.g., reasoning (Bunt, 2015). In this section, I discuss the overall syntax of the TimeML annotation scheme and discuss the semantics as far as necessary. As explained in the first chapter, an annotation scheme consists of an abstract syntax and a concrete syntax (i.e., its implementation). Since the concrete syntax is not relevant for this thesis, I will provide an abstract description of TimeML structures, pointing out the relevant features.

The example in Table 2.1 illustrates the main elements of a TimeML-style annotation. Part of the TimeML annotation applies directly to the primary linguistic data, describing spans of text.[6] Our example, illustrates three types of annotations: *events* (indicated by indices starting with *e*), *temporal*

[6] In the original TimeML standard, the text is marked with XML-style tags, e.g. `<TIMEX>2 days<\TIMEX>`.

Text	**Annotation**
John	
left	*e1*, occurrence
2 days	*t1*, duration, $[t2 : t3]$
before	*s1*, signal
yesterday	*t2*, date: 2002-07-10
TimeML objects	
event instance	*ei1* of *e1*, tense: past, aspect: none
temporal expression	*t3*, date: 2002-07-08
temporal link	$t1 \prec t2$
temporal link	$ei1 \subset t3$

Table 2.1: TimeML annotation: *John left two days before yesterday.*

expressions (or TIMEX, indicated by indices starting with t), and *signals* (with indices starting with s). Thus, the TimeML annotations introduce elements relevant to temporal ordering in a given text. These can correspond to semantic objects but also certain functional words. For example, *e1* identifies the verb *left* as a marker of an eventuality. The annotations not only identify linguistic expressions as markers of temporal information but also add various features, e.g., what type of eventuality or temporal expression we are dealing with. This is explained in more detail in the following paragraphs.

Not all information is encoded in the linguistic data directly. Thus, additional objects may be introduced to properly describe temporal information. This includes relational information (e.g., the annotation of *temporal links* or TLINKs) and the annotation of *event instances*. *Temporal links* are, as the name suggests, responsible for temporal ordering. They are discussed in section 2.2.4. *Event instances* are concrete instantiations of the linguistic descriptions mentioned in the text, i.e., of *events*. Generally, Relations between different objects are expressed via reference to their unique id. Thus, *temporal links* refer to two distinct objects that are supposed to be temporally ordered (e.g., in Table 2.1, *t1* precedes *t2*, and the leaving-*event* takes place during *t3*: the run time of *ei1* is included in *t3*). In the same manner, *event instance* annotations refer to the event that they originate from via the event's

Text	**Annotation**
John	
taught	*e1*, occurrence
...	
TimeML objects	
event instance	*ei1 of e1*, tense: past, aspect: none, pos: verb
event instance	*ei2 of e1*, tense: past, aspect: none, pos: verb
...	

Table 2.2: TimeML annotation: *John taught on Mondays and Tuesdays.*

id.

In the next couple of paragraphs, I focus on the annotation of *events* and *event instances* since those are relevant to encoding tense and aspect information. First, it is essential to understand the distinction between an *event* and an *event instance*. All *events* invoke (at least) one *event instance* but can invoke multiple instances, which motivated instances as a separate element of annotation in the first place. An example of this is given in (7). In this example, two different instances of teaching occur: one is described by *taught on Mondays*, and one is described by *taught on Tuesdays*. This is necessary since the two instances have different temporal properties based on their modifier (*Mondays* and *Tuesday* respectively) but are introduced by the same verb phrases headed by the linguistic expression describing a teaching event. Correspondingly, the expected annotation for the verb *taught* is given in Table 2.2.

(7) John taught on Mondays and Tuesdays.

Ingria and Pustejovsky (2002)

The label *event* is applied to linguistic expressions that describe an eventuality. *Events* are annotated with the feature `class` which encodes different types of eventualities expressed in terms of verbs (or nouns), e.g., reporting verbs, perception verbs, states, or simple occurrences. In Figure 2.1 the value `OCCURENCE` is used, which can be interpreted as a simple action. On the

```
class ::= 'REPORTING' |'_PERCEPTION' | 'ASPECTUAL' |
'I_ACTION' | 'I_STATE' | 'STATE' | 'OCCURRENCE'
```

Figure 2.2: Feature space for the attribute `class`

other hand, the `STATE` label describes *states*, i.e., linguistic expressions that describe a state of affairs. Other than for *occurrences*, TimeML only foresees the annotation of states that can be anchored to some other eventuality or temporal expression (including the document time via tense) and states that are identifiably changed (Saurí et al., 2006). The distinction between states and non-states is crucial since non-states have more or less well-defined start and endpoints while states do not. This also motivates the restriction of the annotation of states mentioned before.[7]

The other possible values for the `class` attribute describe linguistic expressions which embed other eventualities in one way or another. The embedding is not encoded within the verbs themselves but in separately annotated relations between them and their embedded elements. Figure 2.2 illustrates all possible values for the `class` feature. Examples are given in (8), including the embedding of an *occurrence*.[8]

(8) **Examples**

a. John *claimed* that they were dancing. `REPORTING`

b. Maria *heard* explosions `PERCEPTION`

c. Peter *began* to climb the tree. `ASPECTUAL`

d. Clara *attempted* to catch a butterfly `I_ACTION`

e. Bart *prepared* for an attack `I_STATE`

[7] The *state/non-state* distinction is usually attributed to lexical or inner aspect. It is discussed as part of Chapter 4.

[8] Syntactic embedding as illustrated in the examples in (8) is crucial for the interpretation of certain tense and aspect features. More precisely, the means of embedding affect the temporal relation between the matrix event and the embedded event. This is discussed in more detail in Chapter 3.

```
tense ::= 'FUTURE' | 'INFINITIVE' | 'PAST' | 'PASTPART' |
          'PRESENT' | 'PRESPART' | 'NONE'
aspect ::= 'PROGRESSIVE' | 'PERFECTIVE' |
           'PERFECTIVE_PROGRESSIVE' | 'NONE'
```

Figure 2.3: `MAKEINSTANCE` features

2.2.3 Tense and Aspect Features in TimeML

The most relevant part of the TimeML annotation scheme for this thesis happens within `MAKEINSTANCE` objects, which are countable extensions of `EVENT`s. The relevant features are shown in Figure 2.3. As pointed out before, this feature space is designed particularly for English. The label `PERFECTIVE` refers to the *perfect* construction in English. The feature `PRESPART` refers to the *-ing* form used independently of the progressive (i.e. in non-finite verb forms) and texttttPASTPART refers to past participle forms respectively.[9]
As the feature space suggests, in TimeML, tense and aspect features are made up from exactly those distinctions that are morphosyntactically motivated. Table 2.3 provides an example for all available configurations in English and is taken from the annotation guidelines (Saurí et al., 2006). The table suggests that only finite forms of verbs receive tense and aspect features intended to be semantically interpreted. Non-finite forms are annotated based on their morphosyntactic status, i.e. as `PRESPART` or `PASTPART`.
A noteworthy property of this feature space is that certain inflections are encoded within the category of tense in TimeML, e.g., the *-ing* form, encoded as `PRESPART`. I assume that this goes back to Comrie's (1985) distinction between absolute tenses and relative tenses, where relative tenses are generally found in non-finite forms.

[9]Currently, there exist conflicting sources with respect to the attribute/value pair `TENSE` in TimeML. The description here adheres to the annotation guidelines (Saurí et al., 2006). The specification of TimeML found at `https://www.cs.brandeis.edu/~cs112/cs112-2004/annPS/TimeML12wp.htm` encodes form (e.g. present participle, past participle) as the feature `nf_morph` suggesting a more strict distinction between form and meaning.

Verb Group	**Tense**	**Aspect**
teaches	PRESENT	NONE
is teaching	PRESENT	PROGRESSIVE
has taught	PRESENT	PERFECTIVE
has been teaching	PRESENT	PERFECTIVE_PROGRESSIVE
taught	PAST	NONE
was teaching	PAST	PROGRESSIVE
had taught	PAST	PERFECTIVE
had been teaching	PAST	PERFECTIVE_PROGRESSIVE
will teach	FUTURE	NONE
will be teaching	FUTURE	PROGRESSIVE
will have taught	FUTURE	PERFECTIVE
will have been teaching	FUTURE	PERFECTIVE_PROGRESSIVE

Table 2.3: Tense and aspect annotation in TimeML

2.2.3.1 Tense and aspect Features in ISO-TimeML

ISO-TimeML still relies on the same basic building blocks for tense and aspect annotation. It behaves similar to LFG in that it treats auxiliary verbs as feature carrying elements (cf. Ide and Pustejovsky (2017) p. 946–948 for TimeML & Butt et al. (1996) for LFG). One of the main innovations of ISO-TimeML is stand-off annotation. This means that rather than having in-line annotations as presented above, for ISO-TimeML, the data gets segmented into word tokens or morpheme tokens (see, e.g., Im et al. 2009. The annotation then can refer to tokens, spans of tokens, and disjoint tokens that contribute to the same semantic element. However, tense/aspect features are accumulated in the head verb as illustrated in (9).[10]

(9) **Text:** did_0 not_1 $teach_2$

Annotation: *e1*, occurrence, offset: token2 type: process, pos: verb, tense: past, aspect: none, polarity: NEG

[10]The change to stand-off annotation has been partially inspired by Korean TimeML, which has proposed a morpheme-based stand-off annotation. There tense/aspect information is also accumulated in an `EVENT` annotation; however, the offset refers to a span of morphemes rather than only the head verb. This suggests that there is a discrepancy between word-level and morpheme-level stand-off annotation.

Verb Group	**Tense**	**Vform**
to teach	NONE	INFINITIVE
taught	PAST	PART
teaching	PRESENT	PARTICIPLE
	NONE	GERUNDIVE
to be taught	NONE	INFINITIVE
being taught	NONE	GERUNDIVE
having been taught	PAST	PART

Table 2.4: ISO-TimeML non-finite verb forms (Pustejovsky, 2017)

```
relType ::= 'BEFORE' | 'AFTER' | 'INCLUDES' | 'IS_INCLUDED' |
            'DURING' | 'SIMULTANEOUS' | 'IAFTER' | 'IBEFORE' |
            'IDENTITY' | 'BEGINS' | 'ENDS' | 'BEGUN_BY' | 'ENDED_BY'
```

Figure 2.4: Values of the `relType` attibute for `TLINK`s

ISO-TimeML incorporates further morphosyntactic information for non-tensed verb forms. This is illustrated in Table 2.4. Since this thesis focuses on finite constructions, a discussion of non-finite forms is left for future work.

2.2.4 Semantic Links in TimeML

The most prominent link in TimeML is the `TLINK` which indicates the temporal ordering of two semantic objects. This includes `TIMEX-EVENT`, `TIMEX-TIMEX`, and `EVENT-EVENT` pairs. Semantic links, such as the `TLINK`, link two objects based on relation types, encoded in the `relType` attribute. For `TLINK`s these are temporal relations, such as *before* and *after*, inspired by the classification by Allen (1990). The full feature space is given in Figure 2.4. As can be seen there, TimeML uses several symmetric relations (e.g. *before* and *after*). This is done to make the labels conceptually more easy to understand and, thus, annotate. Formally, they express the same relation. Thus, TimeML distinguishes eight different temporal relations.
As shown above, tense and aspect are annotated as features, but their effect

```
relType ::= 'INITIATES' | 'CULMINATES' | 'TERMINATES' |
            'CONTINUES' | 'REINITIATES'
```

Figure 2.5: Values of the `relType` attibute for `ALINK`s

on temporal ordering is not annotated in terms of `TLINK`s.[11] Only explicit temporal links are annotated by TimeML. This includes temporal links introduced by temporal order relators such as *before* (see Figure 2.1 above). However, section 2.2.7 presents some proposals that make more apparent the relation between temporal ordering via `TLINK`s and tense/aspect marking. TimeML also introduces `SLINK`s and `ALINK`s, subordination and aspectual links respectively. These links are also relevant for encoding tense and aspect information. They are discussed in the next two sections.

2.2.4.1 Aspectual links

`ALINK`s or *aspectual links* are elements that express relations between auxiliary aspectual verbs and the event expressed by the main verb. Aspectual verbs are those verbs which refer to certain phases of a given eventuality, as illustrated in (10). In this example, the aspectual verb *(to) begin* refers to the initiation of a *sinking* eventuality. In terms of aspectual verbs, the TimeML specification distinguishes between *initiation*, *culmination*, *termination*, *continuation*, and *reinitiation* verbs (see Figure 2.5; Saurí et al. (2006)). The important aspect here is that aspectual verbs – and other embedding verbs, i.e., those that introduce an `SLINK` (see below) – affect the interpretation of the embedded event. For example, in (10), the initial phase of the *sinking* event, or, more formally, the temporal starting point, is restricted to the past (the sinking could still go on at the *now*, the time of evaluation). The according annotation is shown in Table 2.5. As shown there, an `ALINK` generally holds between an aspectual verb, annotated in the `EVENT` tag, that specifies the relation type of the `ALINK` and an embedded verb. The embedding relation may be indicated by a specific signal – *to* in the present example.

[11] see also section 2.2.6.

Text	**Annotation**
The boat	
began	*e1*, aspectual
to	*s1*, signal
sink	*e2*, occurrence
TimeML objects	
event instance	*ei1 of e1*, tense: past, aspect: none
event instance	*ei2 of e2*, tense: none, aspect: none, nf_morph: INFINITIVE
aspectual link	*ei1 initiates ei2*

Table 2.5: TimeML annotation: *The boat began to sink.*

```
relType ::= 'MODAL' | 'EVIDENTIAL' | 'NEG_EVIDENTIAL' |
            'FACTIVE' | 'COUNTER_FACTIVE' | 'CONDITIONAL'
```

Figure 2.6: Values of the `relType` attibute for `SLINK`s

According to Bunt and Overbeeke (2008a), the `ALINK` encodes temporal information but also indicates that the embedded eventuality (*sink* in Table 2.5) is the THEME of the main argument. Thus, the `ALINK` combines semantic argument and temporal information.

(10) The boat began to sink. Ingria and Pustejovsky (2002)

2.2.4.2 Subordination links

`SLINK`s or subordination links express semantic subordination, e.g. in a factive or counter-factive contexts. The full feature space of relation types is given in Figure 2.6. Examples (11) and (12) illustrate a different subordination link by virtue of a conditional construction. As shown in the concrete example in Figure 2.6, `SLINK`s are used independently of `TLINK`s. As with aspectual and temporal links they can be signaled by specific tokens, here *if*.
These links do not themselves encode temporal information but rather are occupied with the modal dimension, as indicated by the corresponding feature space. A concrete proposal for `SLINK`s using a possible world semantics is

Text	**Annotation**
If	*s1*, signal
Graham	
leaves	*e1*, occurrence
today	*t1*, date
he will not	
hear	*e2*, occurrence
Sabine	
TimeML objects	
event instance	*ei1* of *e1*, tense: present, aspect: none
event instance	*ei2* of *e2*, tense: future, aspect: none, polarity: NEG, modality: WILL
subordination link	*ei1* $\rightarrow$ *ei2*, signaled: *s1*
temporal link	*ei1* $\prec$ *ei2*

Table 2.6: TimeML annotation: *If Graham leaves today, he will not hear Sabine.*

made in Katz (2007). In section 2.2.6, I briefly discuss some challenges that `SLINK`s pose for a semantics of tense and aspect. The details are worked out in chapter 3. There tense in different embedded contexts is discussed.

(11) If Graham leaves today, he will not hear Sabine.

Ingria and Pustejovsky (2002)

a. e_1 = *leaves*
   e_2 = *not hear*
   `SLINK(e2,e1) = 'CONDITIONAL'`

(12) If Susan was sick, she wouldn't be at work.

a. e_1 = *was sick*
   e_2 = *would not be at work*
   `SLINK(e2,e1) = 'COUNTER_FACTIVE'`

2.2.5 Cross-linguistic Variation in TimeML

The original TimeML standard lays much of the groundwork for ISO-TimeML, which is arguably more oriented towards cross-linguistic applicability. As shown in the previous section, the original TimeML annotation scheme was mostly tailored towards English, particularly in its treatment of tense and aspect. However, some adaptions to other languages have taken place, e.g. Bittar et al. (2011), Caselli et al. (2011), and Spreyer and Frank (2008). Some of these directly adapt the treatment of tense and aspect in English (e.g. Spreyer and Frank (2008)), while others accept the difference in tense and aspect systems and modify the annotation guidelines accordingly, e.g., Bittar (2010), Caselli et al. (2011), and Im et al. (2009)). TimeML acknowledges cross-linguistic variation in tense and aspect in the literature; however, they only provide a basic paradigm (see Ingria and Pustejovsky (2002)). The creators of TimeML ask efforts that develop concrete implementations of TimeML in different languages to model tense and aspect features according to the specific needs of the respective language. This approach is sensible, but the lack of specification potentially hurts comparability between languages. To expand on this point, I provide a small survey on the existing TimeML corpora in languages other than English.

The existing literature on cross-linguistic TimeML primarily consists of insights from six languages that participated in the TempEval-2 task (Verhagen et al., 2010): English, Chinese, French, Italian, Korean and Spanish. However, further resources have been generated, e.g., for German, Japanese, Portuguese, Persian, Romanian, and Arabic. Most of these languages have specified features for tense and aspect annotation that suit the corresponding tense and aspect systems. In fact, Korean has triggered a complete shift in TimeML from in-line annotation to standoff annotation to account for the Korean morphological system (Im et al., 2009).[12] On the other hand,

[12] This methodology has been later adapted as a new standard by ISO-TimeML. This is highly relevant since "[t]his is in accordance with the general methodology to create interoperable annotation languages that do not modify the text being annotated."" (Pustejovsky et al., 2010). This concretely means that ISO-TimeML allows for various parallel annotations added to the same text without interference with the source material. It also

the Japanese (Asahara et al., 2014) and Chinese (Cheng et al., 2007; Xue and Zhou, 2010) documentations of their treatment of tense and aspect acknowledge the specific tense/aspect systems, but concrete specifications are not readily accessible. In fact, in the case of the Japanese language, it was determined that aspect especially provides a challenge that does not seem to have been addressed since (Asahara et al., 2014).
Overall, there is no openly available cross-linguistically valid specification of tense and aspect features for TimeML. Conversely, the existing TimeML resources do not conform to an underlying cross-linguistically uniform frame. This makes the comparison between languages with respect to these features challenging. As Im et al. (2009) have shown, it is not only important to address the space of available tense and aspect features but also the mode of annotation. Ideally, this mode of annotation is uniform across languages. Some of these issues have been addressed in ISO-TimeML; however, this standard is not openly available. Thus, although some of the insights presented in relation to ISO-TimeML benefit the computational linguistics community, there exists no publicly available standard for tense and aspect annotation cross-linguistically.

2.2.6 Semantic interpretation of TimeML structures

There exist several proposals for a semantics for (ISO-) TimeML some of which have already been mentioned before, e.g., Bunt (2010), Bunt and Overbeeke (2008b), Katz (2007), and Lee (2008). As for example discussed in Katz (2007), the main challenge for the TimeML annotation scheme is that flat representations provide a challenge for a scope-based semantics (see also Bunt 2019). However, with respect to tense and aspect, only rudimentary semantics are generally proposed. This is illustrated in Table 2.7 from Bunt and Overbeeke (2008a). There, the values for the feature `TENSE` are shown. Semantic interpretation of tense/aspect features in TimeML runs into more

means that the ISO-TimeML annotation is independent of syntactic well-formedness. This makes it inherently more robust since the pre-processing only requires proper segmentation of a given text.

tense **value**	**Translation**
`tense="PAST"`	$\lambda e.Before(ET(e), T_0)$
`tense="PRESENT"`	$\lambda e.Inside(T_0, ET(e))$
`tense="FUTURE"`	$\lambda e.Before(T_0, ET(e))$

Table 2.7: Translation table for the `EVENT` tag attribute `tense`

issues when combined with other linguistic phenomena, e.g., when used in conditionals that are analyzed as `SLINK`s. As shown in example (12), the English past tense morphology may induce a counterfactual reading. In the process of inducing a counterfactual reading, the past tense loses its past interpretation (see Iatridou (2000) for an overview; chapter 3 for a more detailed discussion on the role of tense in counterfactual conditionals). As a result, the sentence in (11) is about the present rather than the past.[13] It is not clear how this is accounted for in TimeML and can lead to an incoherent semantics since the default assumption is that tense and aspect are annotated based on the morphosyntactic input. If the features are semantically interpreted as illustrated in Table 2.7, the annotation then produces an incorrect semantics, bar an intricate semantic analysis for conditionals, for example, along the lines of Romero (2014). Generally, this means that tense and aspect features potentially behave semantically different if embedded in such subordination structures, which is not straightforwardly accounted for in TimeML.

Overall, the semantics of tense and aspect are incorporated inconsistently into TimeML in the sense that sometimes tense and aspect are taken to be morphosyntactic features without effect on the resulting semantic representations, whereas in other instances, as the one illustrated above, tense and aspect only receive basic interpretations which run into trouble when embedded in more complex contexts. The problem that this poses becomes

[13] Counterfactuals are not available in future contexts. Intuitively, this is because it is impossible to tell what is factual and what is counterfactual in situations that have not yet happened (Iatridou, 2000).

even more apparent when considering the relatively free stance towards the cross-linguistic annotation of tense and aspect in TimeML described in the previous section.

In the introduction, I described the CASCADE model for semantic annotation schemes. As established there, a semantic annotation scheme requires an appropriate semantic interpretation, i.e., a formal semantics that works in tandem with the syntax of the annotation scheme. This means it would be *technically* necessary to establish both a feature space and a corresponding semantics for TimeML, when developing a resource for a new language. However, the annotation guidelines generally only describe the syntax of the corresponding TimeML resources. Thus, one of the main desiderata of semantic annotation schemes, namely, a fleshed-out concrete formal semantics, is not reinforced within the TimeML community for tense and aspect, particularly from a cross-linguistic perspective. One of the main reasons for this is the lack of resources for validating the semantic representations produced by TimeML annotations.

2.2.7 Alternative approaches to tense and aspect in TimeML

In this section, I focus on two approaches to expanding TimeML with respect to tense/aspect annotation. On the one hand, tense/aspect annotation can be integrated into the annotation of temporal semantics by adding additional objects (i.e. `TIMEX`s and `TLINK`s) to the annotation to increase granularity of the temporal ordering. On the other hand, tense can be treated as a separate layer for constraining temporal semantics.

The former approach has been advocated by Gast et al. (2016, 2015). The main idea they pursued was the introduction of the Kleinian model of tense and aspect (Klein, 1994) into the TimeML standard. This leads to a more fine-grained granularity of temporal expressions in any given text. By making topic/reference times of verbs explicit in the annotation, semantic theories that distinguish between event run time and reference time (Reichenbach, 1947) are supported within the TimeML architecture. However, one needs to consider whether this goes against the philosophy of only annotating those

temporal expressions for which a human annotator can find a concrete value (Ferro et al., 2001).[14]
Another issue of this system is that tenses are often ambiguous, and allow for various different temporal relations (Derczynski, 2016). For example, in German, the present tense morphology is compatible with both a future-oriented relation between the speech and eventuality time and a simultaneous relation between the speech and eventuality time. Derczynski (2017) describe the same issue for the simple future in English. Thus, another proposal to incorporate the semantic contribution of tense and aspect into the architecture for temporal relations provided by the `TLINK`s is to interpret tense as constraints on `TLINKs`. In such a framework, the annotation of `TLINK`s is restricted by the tense value of a given event instance. This approach allows the system to handle the sometimes ambiguous or vague constraints imposed by tense and aspect on temporal ordering, a crucial factor in annotating the semantics of tense and aspect. Derczynski (2017) points out that an annotation scheme encoding the Reichenbachian tense framework is cross-linguistically applicable and can serve as a useful tool to constrain TimeML annotations.

2.2.8 Summary on TimeML

In summary, TimeML provides an annotation scheme for the temporal properties of linguistic expressions. The handling of tense and aspect is treated as a language-specific phenomenon and more of a pre-processing step. TimeML does not enforce a cross-linguistically viable feature space for the treatment of tense and aspect, leading to various approaches to dealing with these categories in different languages. Finally, alternative treatments of tense and aspect have been discussed briefly. The approach in this thesis is closest in nature to that pursued in Derczynski (2017), which proposes to treat tense and aspect as constraints on temporal relations.

[14] It is noteworthy that this philosophy is somewhat relaxed in TimeML. Thus, TimeML actually allows the annotation of implicit time intervals (Saurí et al., 2006).

2.3 Tense and aspect annotation in deep linguistic parsing

In this section, I describe an effort that strives towards linguistic analyses that are parallel across languages based on Lexical Functional Grammar (LFG), (Bresnan et al., 2015; Bresnan and Kaplan, 1987; Kaplan, 1995a), namely, the Parallel Grammar (ParGram) project Butt et al. (2002). In this project, LFG grammars for several typologically diverse languages have been and are developed.

The difference between TimeML and the related efforts discussed in the previous section and the ParGram project is that the latter explicitly aims towards cross-linguistically comparable feature spaces at different levels of linguistic interpretation (in particular, morphology and syntax). Thus, the ParGram project stands in stark contrast to the TimeML effort since the latter, as mentioned in the last section, argues that tense and aspect should be treated as a language-specific issue.

This section is structured as follows: first, the linguistic framework employed in the ParGram grammars, i.e., LFG, requires an introduction. Thus, in the next section, I will focus on more common and basic linguistic phenomena to explain the architecture of LFG. Readers familiar with this topic may skip the next section and go directly to section 2.3.4. The treatment of tense and aspect within the ParGram grammars is discussed in detail in section 2.3.3 with reference to both the theoretical as well as the practical literature on LFG. Finally, I compare LFG to the universal dependencies (UD) effort, which also provides a cross-linguistically comparable syntactic analysis, however, on a more shallow level. I present both UD and LFG in this thesis to show that the present annotation scheme can be applied to different syntactic formalisms. More generally, I want to carve out the formal underpinnings necessary to provide an annotation scheme that covers both morphosyntactic and semantic tense/aspect information.

2.3.1 The ParGram Project

The late 1990s and early 2000s were prime years for deep linguistic parsing, giving rise to workshops directly focused on computational syntax, e.g., *Treebanks and linguistic theory* and benefiting from international projects such as the ParGram effort Butt et al. (2002), which first started in 1996. This project has (and is) pushing theoretical and applied research in the LFG community and provides grammars for various languages:

(13) a. Larger grammars for English, German, French, Norwegian, Chinese, Japanese, Polish

b. Smaller grammars for Greek, Indonesian, Malagasy, Turkish, Welsh, Wolof, Urdu, Georgian, Hungarian

As the name suggests, the goal of the project is to provide cross-linguistically parallel (or comparable) linguistic analyses for semantically comparable linguistic expressions. It makes use of the unique architecture of LFG to distinguish between language-specific and universal properties. Although my focus lies in presenting the insights on tense/aspect gained from the ParGram project, it is first necessary to present LFG.

2.3.2 Lexical Functional Grammar

LFG has a long-standing tradition in both theoretical and computational linguistics. Although it is most well-known for its contribution to syntactic research, it is, in fact, a full-fledged formalism of language covering phonology, morphology, syntax, and semantics. In this thesis, I will refer to the morphosyntactic component of Lexical Functional Grammar (LFG) with the label LFS (Lexical Functional Syntax). More specifically, I focus in particular on those elements of LFG that have been implemented computationally. Concretely, I rely on the implementation of LFG in terms of the Xerox Linguistic Environment (XLE), which has been developed for the implementation of large-scale LFG grammars (Crouch et al., 2017). It is the most widely used platform for computational LFG research and is also the platform for the biggest cross-linguistic grammar development project in the LFG community, the ParGram project (Butt et al., 2002). The implementation of LFG in terms of XLE underpins the formal soundness of the formalism.

2.3.2.1 Projection structure

LFS is a constraint-based syntactic grammar formalism based on two different levels of formal representations: the constituent, or c-structure, and the functional core structure, usually referred to as f-structure. These two structures can be understood as parts of a function that maps linguistic form to meaning, i.e., a semantic model. Thus, each of these structures is itself derived via a function usually called projections in LFG. Traditionally, the c-structure is a constituent parse of the input string, and the projection pointing from form to c-structure is called the π-projection. The f-structure is a more abstract flat syntactic representation that encodes relational information as well as certain functional linguistic features. It is derived from the c-structure by placing additional constraints on constituent rules. The corresponding projection is called the ϕ-projection.

The projection or correspondence functions are intermediate representations that serve to explore certain types of constraints (Kaplan, 1995b). Thus, the structures provided by LFG are more of a tool for linguistic research rather than any sort of mental representation of language. Due to this insight, LFG allows for the introduction of novel intermediate representations between form and meaning. However, these intermediate levels have to follow strict mathematical rules. Most importantly, in the series of successive functions from form to meaning, for each function, its domain is the range of the previous one (Asudeh, 2006). Figure 2.7 illustrates a potential decomposition of the function from form to meaning that LFG aims to describe. It shows that the range of the π function is any input string. Furthermore, this illustration shows that the ϕ-projection has been decomposed into the μ,α, and λ projection providing an intermediate morphological structure(Butt et al., 1996) and an argument structure (Butt et al., 1997).

This kind of decomposition of functions is particularly valuable for computational approaches since it allows us to distribute particular steps of linguistic

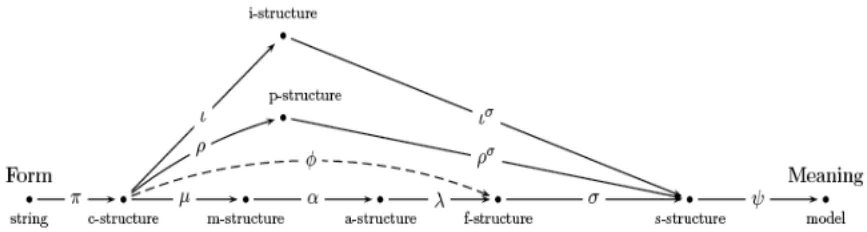

Figure 2.7: Parallel projection structure as shown in Asudeh (2006)

analysis among separate modules and has been implemented to great effect in the XLE grammars.

2.3.2.2 Syntactic representations

In the previous section, I have mentioned two different representations that are central to the research done in LFG. The first of these two structures is the c-structure, a constituent parse of a given input string. More concretely, the c-structure is provided by a context-free grammar encoded in terms of phrase structure rules. These rules describe the set of admissible trees for a given language. In other words, the output of the π-projection is the set of trees that is available for a specific input string. This means the c-structure allows for syntactic ambiguity as, for example, the well known prepositional phrase (PP) ambiguities.[15] The most important role of the c-structure is to encode cross-linguistic variation by capturing linear order. Furthermore, it encodes constituency and hierarchical relations between constituents. In an LFG grammar, where terminal nodes of the c-structure are full form lexical entries, c-structure also encodes morphological variation.

In traditional LFG, lexical entries are stored based on their surface string, as shown in the sample grammar in 2.9. However, nothing prevents LFG from introducing another projection that deals with morphological analysis, as has been discussed above. In fact, arguments for a M(orphological)-structure

[15] Usually, the c-structure grammar loosely follows the rules of X'-theory (Jackendoff, 1977), but can, in principle, be substituted for by other kinds of tree structures expressing categorical information (Bresnan et al., 2015).

have been made in (Butt et al., 1996; Frank and Zaenen, 2002) and various approaches to morphology in LFG have been undertaken since. The currently best-articulated architecture for the morphology-syntax interface is found in Dalrymple (2015), and this thesis largely adopts that architecture, which also describes that of the morphology-syntax interface used by XLE.

As described in Butt et al. (1999), the morphology uses, in principle, exactly the same mechanisms as the c- and f-structure. Morphological tags receive their own lexical entries together with relevant f-structure annotations. These tags are combined into full words via sub-lexical tree rules, specified via finite-state machines (Kaplan et al., 2004). Those rules produce flat tree structures that impose a certain order on morphological tags. The translation of morphological features into functional features is discussed in section 2.3.3 with a focus on the treatment of tense/aspect morphosyntax in LFG.

The c-structure and morphological analysis encode properties of the surface structure of a given input string, such as morphological features, constituency, word order, and syntactic ambiguity. On the other hand, the functional structure encodes functional relations and features which mediate between surface structure and semantic interpretation. Thus, the f-structure provides a level of abstraction from the surface structure. Its main purpose is to describe syntactic relations and functional features that are cross-linguistically available. The main idea is that different languages use different surface structure cues to express certain grammatical functions (e.g., subject, object, etc.); however, the role of these grammatical functions is (more or less) invariant across languages and related to the semantic argument structure. Similarly, functional features (e.g., tense, mood, case, etc.) result from similar structural reflexes cross-linguistically. Functional features also mediate between surface structure and semantic interpretation of an input string. This dual role leads to a tension that is very similar to what we had seen when the treatment of tense and aspect in TimeML was discussed.[16]

In abstracting away from the surface structure, the f-structure also discards

[16] I assume that similar arguments can be made for the role of grammatical functions, but this is of limited concern of this thesis.

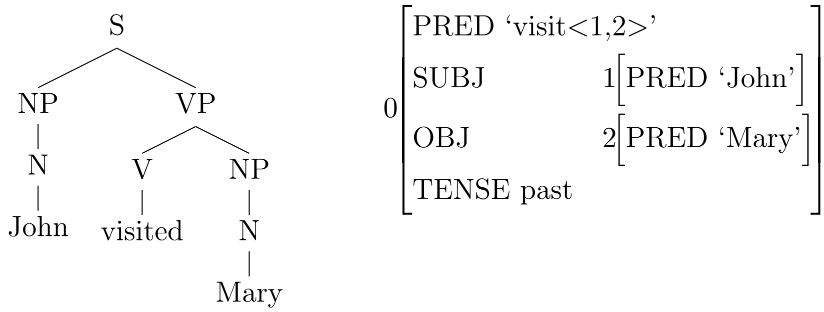

Figure 2.8: c- and f-structure: *John visited Mary*

hierarchical ordering to an extent. It is represented as an attribute/value matrix that still encodes certain dominance constraints, but only those that reflect the role of grammatical functions with respect to predicate/argument structure. The importance of the argument structure is also reflected in the `PRED` value, which plays a special role in the f-structure: for predicates and arguments, it provides the string by which the element is encoded in the lexicon, and its arguments, creating a *semantic form*. Every `PRED` value introduces a new instance of a given object. Thus, each value is unique even if they share the same semantic form (Börjars, 2020). The semantic form also indicates precisely which grammatical functions need to be represented in the f-structure (in other words, the semantic form provides a frame for subcategorization). An example f-structure is given in Figure 2.8. As shown there, the `PRED` indicates that the verb *visit* takes two arguments, both of which have to be instantiated as sub-structures in the f-structure governed by their grammatical functions (in this case: `SUBJ` and `OBJ`).
The functional structure is provided by local co-description (Bresnan et al., 2015) on c-structure rules. This means that the ϕ projection is built in parallel to the c-structure by following certain annotations of c-structure rules that constrain the f-structure. An example of this is given in Figure 2.9, which provides a grammar consisting of a lexicon and annotated c-structure rules for the sentence *John visited Mary*. The annotations on the c-structure rules pick up the partial f-structures described in the lexicon and assemble the full f-structure via unification. In the current example, the VP rules

S	→	NP (↑SUBJ) = ↓	VP ↑=↓
NP	→	N ↑ = ↓	
VP	→	V ↑ = ↓	NP (↑OBJ) =↓

John	N	(↑PRED) =‘John’
Mary	N	(↑PRED) =‘Mary’
visited	V	(↑PRED) = ‘visit<(↑SUBJ),(↑OBJ)>’ (↑TENSE) = past

Figure 2.9: Grammar for *John visited Mary*

generate an f-structure that contains the `PRED` value, a feature tense, and a sub-structure embedded via the grammatical function `OBJ`. This intermediate structure is combined with the `SUBJ` NP via the top-level S-rule.

The process described above is managed via the meta-variables ↑ and ↓, which point at dominating f-structures and dominated f-structures nodes that are aligned with the c-structure via the correspondence phi, respectively. As mentioned above, these variables are instantiated in XLE via the unification of functions and equations. In the present example, the two grammatical functions (`SUBJ` and `OBJ`) are encoded in terms of functions of the form $f(x) = y$, where f is the grammatical function, and x and y correspond to the indices in the syntactic structure. For example, the function (↑OBJ) =↓ equates the f-structure of the dominated NP with the structure corresponding to the `SUBJ` of the whole f-structure. The equation ↑=↓ simply passes the properties of the dominated f-structure to the dominating f-structure. Furthermore, the `TENSE` feature of V is passed up to S from its instantiation directly in the lexicon. Thus, it occurs at the top level of the corresponding f-structure.[17]

This section has shown that LFG is a constraint-based grammar formalism that encodes grammatical properties as constraints on linguistic input expressions. These constraints are distributed across various layers (projections), such as the c- and f-structure explained above. In the next section, I concretely describe the annotation of tense and aspect in LFS in particular. Special attention is paid to the encoding of tense/aspect as functional features, i.e., features in the f-structure, which are assumed to be primary input for further semantic processing.

2.3.3 Tense and aspect in LFG

LFG uses the f-structure to abstract away from the surface form to encode cross-linguistically robust syntactic information. The main concern in the LFG community was the distinction between analytic and synthetic forms of tense and aspect categories (Butt et al., 1996; Frank and Zaenen, 2002). In Dahl's (1985) work, this roughly refers to the distinction between morphologically constructed categories and periphrastically constructed categories. The current treatment in XLE is to abstract away from such distinction at the level of f-structure. I first present this in terms of an example from the ParGram corpus shown in (14). In this example, the English past tense represents a morphologically constructed category, whereas the Urdu past tense is realized as an auxiliary construction.

(14) a. Peter had built this house.

b. | piʈər=ne | je | gʰər | bənɑjɑ |
|---|---|---|---|
| Peter.M.SG.=ERG | this | house.M.SG.NOM | build.PRV.M.SG |

t̪ʰɑ.
be.PST.M.SG

'Peter had built this house.' *Urdu*

[17] At this point, I only explain the basics of encoding C-structure and F-structure in the grammar. A full explanation of the possibilities provided by XLE is given in Crouch et al. (2017) (more specifically: `https://ling.sprachwiss.uni-konstanz.de/pages/xle/doc/notations.html`). Some more advanced operations are discussed in chapter 6 when they become relevant.

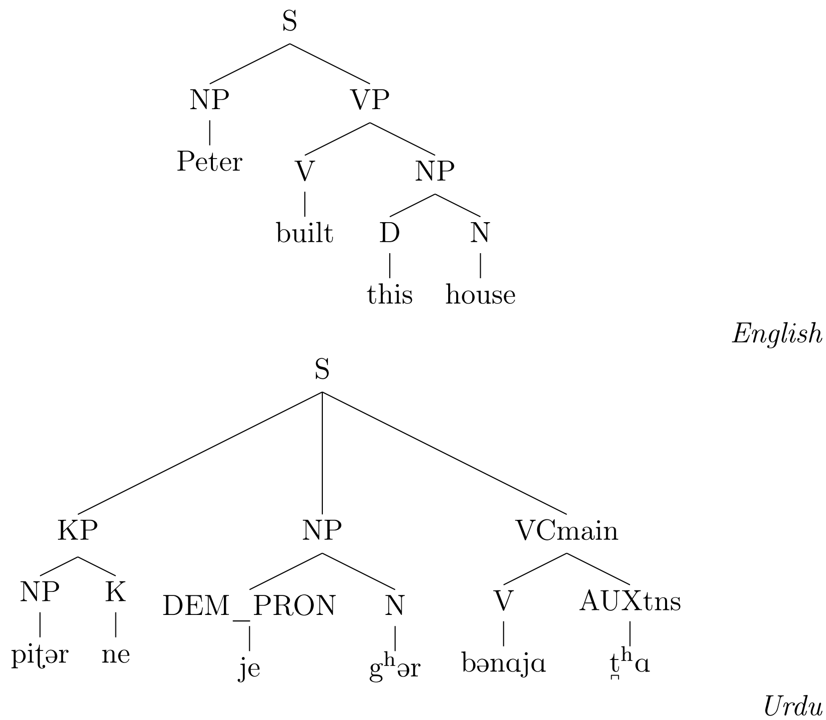

Figure 2.10: Urdu and English C-structure: *Peter built this house*

In this comparison, the c-structure encodes the distinction between the realization of the tense feature in the two languages (see Figure 2.10). While the English c-structure expresses the verb form in terms of a simple *V*-node, the Urdu c-structure composes the tense auxiliary and the main verb into a verb complex (`VC`). This complex receives the tense feature from the auxiliary and behaves, in principle, like the English `V`-node in the way it percolates the `TENSE` feature to the top-level node `S`, thus, making it available as a functional feature for the corresponding top-level f-structure. In the f-structure, both of these constructions are mapped onto the same functional `TENSE` feature as shown in (15).

(15) $\begin{bmatrix}\text{TNS-ASP} & \begin{bmatrix}\text{TENSE past, ...}\end{bmatrix}\end{bmatrix}$

Example 16 provides some further illustrations of this practice for future tense. All of these sentences receive the f-structure attribute/value pair `TENSE` `fut`, (17), despite significant differences in the surface structure (as indicated by the glossing).

The Indonesian example and the French example illustrate a contrast similar to that seen between English and Urdu for past tense. The Polish example provides a special case; however, as indicated by the glossing, the verb form is not marked for future explicitly but rather implicitly by the combination *PRV + PRES*. This supports the idea that the f-structure provides a semantically interpretable feature that mediates between surface structure and an actual semantic representation. In fact, the process witnessed here by virtue of the Polish example provides a precedent for the accumulation of semantic features in the ParTMA annotation scheme, as will be discussed in detail in Chapter 3.

(16) a. Peter akan menulis surat
Peter FUT write letter
'Peter (is/was) going to write a letter.' *Indonesian*

b. Peter écri-ra une lettre
Peter write-FUT a letter
Peter will write a letter. *French*

c. Peter **na-pis-ze** list
Peter **PRV-write-PRES** a letter
'Peter is going to write a letter (completely)' *Polish*

(17) $\left[\text{TNS-ASP} \quad \left[\text{TENSE fut, ...}\right]\right]$

The approach covers not only tense features but also aspectual features. Particular attention has been paid to perfect and progressive constructions within ParGram (Butt et al., 1999, 1996). In fact, they served as the driving force behind the simplification and cross-linguistic unification of f-structures illustrated above. This is because auxiliaries have been traditionally treated as complement-taking verbs (see, e.g., Falk (1984)). However, this treatment of auxiliaries has gone out of fashion since at least King (1995) as summarized in Butt et al. (1996). As explained there, the old analysis that treats aspectual auxiliaries similar to raising verbs provides several challenges both from a syntactic and a semantic point of view. The main argument is that such an analysis obscures the parallelism between languages at the level of f-structure. Consider again the french example in (16). In french, the future tense is not encoded by means of an auxiliary but as an inflectional element. This also affects its behavior when combined with aspectual features. As the examples in (18) show, English and German distinguish between tense (*will/wird*) and aspect (*have/haben*) auxiliaries which leads to two levels of embedding at f-structure (see 19a). On the other hand, the french auxiliary *aura* encodes both tense and aspect information and thus only leads to one level of embedding (see 19b).

(18) The driver will have turned the lever.
Der Fahrer wird den Hebel gedreht haben — German
Le conducteur aura tourné le levier — French
(Butt et al., 1996, p. 3)

(19) a. $[_{fut}$ XCOMP ... $[_{perf}$ XCOMP $[_{turn}$ ... $]]]$
b. $[_{fut}$ XCOMP ... $[_{turn}$ ... $]]$

To reiterate: the different levels of embedding in the f-structure enforce a different predicate (i.e., semantic) structure across languages. This is an undesired consequence of this kind of analysis since it undermines the parallelism strived towards at the level of f-structure. This is another instance of tension between the encoding of structural information and semantic information that troubles tense/aspect annotation efforts more generally.
The approach presented above proposes some additional machinery, the morphological structure (m-structure, μ-structure). This additional projection uses the same kind of equations as the f-structure on a morphological level. The result is an attribute-value matrix encoding morphological dependencies, i.e., agreement. See also Frank and Zaenen (2002) for some further discussion on this topic. The use of the m-structure cleans up the f-structure considerably by enforcing the idea that the f-structure abstracts away from surface-level syntactic cues. Furthermore, it helps to resolve long-distance dependencies, as, for example, in (20), where the VP is topicalized, i.e., moved to a frontal position (compare with (21)).

(20)

[Den	Hebel	gedreht]	wird	der	Fahrer	haben.
the.ACC	lever	turn.PSTPTCP	will.PRES.SG	the	driver	have.INF

'The lever turned, will have the driver.'

German

Butt et al. (1996, p. 9)

(21)

Der	Fahrer	wird	den	Hebel	gedreht	haben.
the	driver	will.PRES.SG	the	lever	turn.PSTPTCP	have.INF

'The driver will have turned the lever.' *German*

Within the ParGram grammars, morphology is not described in a separate projection. It uses the `CHECK` feature, which stores attribute value pairs relevant for well-formedness for different components of an f-structure. Thus, the idea is similar, but encoding the morphological information in the `CHECK` feature, which appears internally at the level of f-structure, leaves it more accessible within a given grammar. Furthermore, this approach simplifies debugging since it is not necessary to keep track of information across multiple projections. However, the information in `CHECK` is not a part of the f-structure but only an artifact used to ascertain the well-formedness of the morphological surface structure (among other things). Thus, the information encoded at this level is preserved in computational representations of LFG parses, although they are supposed to be only used XLE internally (Crouch et al., 2017). Additionally, sub-lexical information is stored as part of the c-structure making morphological features overtly available for further processing.
At this point, it is important to note that the treatment of morphology is a point of tension between theoretical and computational approaches to LFG. As becomes evident in the triangular structure of Figure 2.7, in LFG, there exists a tension between parallelism and direction of analysis. The approach used in XLE decomposes word forms via sub-lexical rules so the system can derive functional features from morphological information, i.e., XLE is, to an extent, directional. On the other hand, Dalrymple (2015) (among others) proposes a fully parallel treatment of morphology. This approach is prevalent in theoretical LFG but is of little concern for the present thesis since the goal is to extend the semantic and not the morphological capabilities of XLE.

This kind of parallelization is what lies at the heart of the ParGram project and has been applied to various tense and aspect categories, such as the perfect, the progressive, and tense. An example of this are long-distance dependencies between auxiliaries and their main verbs (Butt et al., 1996; Frank and Zaenen, 2002). However, this is a syntactic issue that is not relevant to the semantic analysis.

2.3.4 Insights from the ParGram project

As explained in the previous section, tense/aspect information is encoded across various levels of representation in LFG. Structural properties like linear word order and constituency are primarily encoded in the c-structure, which is decomposed into a full-form constituent parse and a morphological analysis providing sub-lexical features for the terminal nodes of said parse. Thus, the lexicon in large-scale XLE grammars contains pseudo full-form entries composed based on sub-lexical flat tree rules. In LFG theory, further possibilities of decomposing structural information into morphological and syntactic information have been put forward, but their implementation within XLE remains, at best, implicit. More interestingly, LFG's f-structure is supposed to provide semantically interpretable features based on the morpho-syntactic input. In other words, f-structure features must be directly discernible from a given input string's surface structure and must be semantically interpretable. Within the ParGram project, these thoughts have resulted in the set of semantic features for tense, aspect, and mood given in Figure 2.11.
This feature structure illustrates how features may be embedded in an f-structure and their respective values. In this thesis, I will ignore the `VOICE` matrix and focus solely on the `TAM` matrix. This structure alone contains a wide variety of features, and it is impossible to cover all of them in detail. Thus, I am mainly concerned with those features that are part of the three major levels of tense/aspect analysis: i) tense, ii) grammatical (outer) aspect, and iii) lexical (inner) aspect. This means I will not provide a detailed description of `MOOD`, `EV-VALUATION`, and `MODALITY`.[18]

[18] Nonetheless, these categories will play a role in some analysis presented in this thesis.

The semantic foundation of these features is a Reichenbachian system for tense/aspect (Reichenbach, 1947). Similar approaches have been made for semantic annotation within TimeML and related efforts, as we have seen before. Reichenbach's system (and subsequent efforts building on similar ideas, most notably, Klein (1994)) deal with the temporal properties contributed by tense and aspect. This reflects in the values for `TENSE` and `ASPECT` which correspond to various potential temporal relations between points in time. In addition to `ASPECT`, `VIEWPOINT` also contributes to the level of grammatical, or outer aspect, however, the concrete semantic remain undefined. Each of these categories plays an important role in temporal semantics. This role is explored in detail in chapter 3. In addition to `VIEWPOINT`, `EV-STR` and `EV-EXEC` are discussed in chapter 4, focusing on aspectual properties that indirectly affect temporal semantics.

There is a reason I do not go into detail in explaining these features at this point: although they provide a fine-grained system of semantically informed features, most of them have never been actually implemented within the ParGram project. This is indicated by the `STANDARD_COMMON_FEATURES` template provided by the ParGram grammars. Similar to the template in Figure 2.11, the features provided there decompose tense and aspect into distinct features as proposed by Butt et al. (1999). However, the feature space is much more limited, namely, to the morphosyntactically distinguishable features available in English. The exact feature space is shown in Figure 2.12. With this they generally follow this insight: "The f-structure thus now encodes exactly those distinctions which are made overtly in each of the languages without attempting to second guess a semantic analysis [...]. Butt et al., 1999, p. 69"

This has consequences as to how the f-structure plays into the semantic analysis of tense and aspect in LFG: According to the quote above, the tense/aspect parameters are reduced to a bare minimum, namely to exactly those that can be grammatically differentiated and are not dependent on

However, I will not propose substantial changes in handling these features.

TAM	TENSE	pres/past/fut/non-past/non-fut
	TENSE-DIST	neutral/remote/immediate
	ASPECT	follow/overlap/precede/non-follow/non-precede
	VIEWPOINT	imperfective/perfective
	MOOD	indicative/subjunctive/conjunctive(?) imperative/hortative/optative
	EV-EXEC	EXTEND +/- ITER-DUR +/-/iter/dur HASTY +/- DEIXIS away/towards HABITUAL +/-
	EV-STR	STATIVE +/- INIT +/ PROCESS +/- RESULT +/-
	EV-VALUATION	evidential/mirative/certain/ quotative/inferential
	MODALITY	unspec/deontic/epistemic/potential/ desiderative/permissive/...
VOICE	PASSIVE	+/-
	PASSIVE-TYPE	personal/impersonal/adversative
	MIDDLE	+/dir-refl/indir-refl/caus-refl/anti-caus

Figure 2.11: Tense-Aspect-Mood-Voice feature space ParGram (2010)

TNS-ASP	TENSE fut, null, past, pres PROG +,-,+_,-_ PERF +,-,+_,-_ MOOD imperative, indicative, subjunctive, successive

Figure 2.12: ParGram common features (2006)

contextual disambiguation (for example, temporal reference in Indonesian).[19] I interpret this in the following way: Morphosyntactic features receive their respective abstract representation in the f-structure based on their form rather than their meaning. The meaning of a form contributes more to the choice of the arbitrary symbolic label given to it rather than any semantic distinction. For example, suppose the primary use of a given morphosyntactic form is that of anteriority. In that case, the label `TENSE` with the value `past` is employed. However, non-past forms, i.e., forms that are used to express both present and future tense, get the value `pres` (for present tense), which is more in line with the older naming conventions for this grammatical form when compared with other languages. An example of this is found in the German present tense.

(22) Peter besuch-t (*gestern/gerade/morgen) seine Mutter.
Peter visit-PRES (yesterday/currently)/tomorrow) his mother
'Peter is visiting/will visit his mother'

Additional challenges are provided by languages that have a more fine-grained tense/aspect system. They deviate from the proposed feature space (e.g., Polish, Urdu, Greek, etc.), usually referring back to the more detailed feature space in Figure 2.11. However, the uniformity across grammars is less well achieved than for the features concretely specified by the common features template. For example, both Urdu and Polish encode viewpoint aspect (i.e. *perfective* vs. *imperfective*) in the feature `ASPECT`, rather than `VIEWPOINT`. Although those are minor inconsistencies, they seriously affect the possibility of developing a cross-linguistically applicable semantic component for handling tense/aspect in XLE grammars when they pile up. I summarize some further inconsistencies in the ParGram tense/aspect annotations in the next few paragraphs.

As explained above, the ParGram features are designed to mark morphosyntactically overt tense/aspect features. However, there are some instances within the ParGram grammars where this view is violated. This sometimes

[19] The ParGram grammars thus, ultimately, take a similar stance as TimeML.

leads to more or less severe confusion when analyzing structures. Consider the following example from Polish, which exemplifies one possibility of expressing future tense by combining a verb marked for perfective aspect with the present tense marker *-ze*. In comparison, (24) shows that the present tense and the imperfective together form a sentence about the present. This suggests that the contribution of the marker *-ze* is associated with present tense and future tense, respectively, in the different examples. This is since the f-structure here encodes the semantic interpretation rather than the contribution of a single morpheme. The future reading in (23) is derived from the semantic properties of the perfective and semantic properties of the present tense. Thus, although this is a fair analysis of the construction, it can lead to potential confusion concerning the analysis of the marker *-ze* as either present or future tense if one assumes a strictly morpho-syntactic analysis.

(23) a. Peter **na-pis-ze** list
Peter **PRV-write-PRES** a letter
'Peter is going to write a letter (completely)' Polish

b. [TNS-ASP [TENSE fut, ASPECT perf]]

(24) Peter **pis-ze** list
Peter **write-PRES** a letter
'Peter is writing a letter.

[TNS-ASP [TENSE pres, ASPECT imperf]]

Another area of conflict is the cross-linguistic comparison of categories. In Urdu, past and present tense are only ever encoded by virtue of an auxiliary, whereas these tenses in Germanic languages are encoded inflectionally. This means the feature `TENSE` encodes a different morphosyntactic form in these languages. Consider the two examples given in (26) and (27). They are semantically similar, at least to the extent that they can occur in the same context. The English tense morphology is used to express the meaning of pastness. Urdu, on the other hand, does not necessarily carry a tense feature, as illustrated by the difference between (27) and (28a). However, if the feature `past` is used in Urdu, it refers to an auxiliary construction with a past perfect interpretation (Butt and Rizvi, 2010). Due to this last property, a problem arises: the two instances of `TENSE` are not cross-linguistically comparable.[20]

(25) a. **CONTEXT:** [Assuming that B was going to meet Peter, A asks:]
A: Did you meet Peter?

(26) a. B: I met Peter.

b. [TNS-ASP [TENSE past, MOOD indicative]]

(27) a. mæ̃. piʈər=se mɪla.
I.SG.MOM Peter.M.SG=with meet.PRV.M.SG
'I met Peter.' *Urdu*

b. [TNS-ASP [MOOD indicative, ASPECT perf]]

(28) a. mæ̃. piʈər=se mɪla t̪ʰɑ.
I.SG.MOM Peter.M.SG=with meet.PRV.M.SG be.PAST.M.SG
'I had met Peter.' *Urdu*

b. [TNS-ASP [MOOD indicative, TENSE past, ASPECT perf]]

This means that while the ParTMA tense/aspect labels stem from the desire to analyze these languages in parallel, they do not, in fact, encode either syntactic or semantic parallelism. It is impossible to make the same kind of generalizations about these two seemingly congruent features across languages. All in all, it is not straightforward which properties of a morphosyntactic form should be generalizable across language for them to retrieve the same label at f-structure and, if so, whether the label is cross-linguistically consistent or language-specific on a semantic level.

In summary, although agreed-upon approaches exist to encode tense/aspect morphosyntax in LFG grammars, the implementation of tense and aspect features in the ParGram grammars has not been ideally regulated. This is

[20] This problem is further reinforced if we consider the distinction between optional and obligatory tenses, possibly invoking further semantic differences between the two instances of past tense (Cable, 2015; Plungian and van der Auwera, 2006).

potentially due to the focus on other more general issues of parallelism that took center stage in the ParGram research. Furthermore, the rich variation in tense/aspect systems is difficult for parallel approaches. Nonetheless, the general insights on encoding tense/aspect that have been gained so far will form a valuable foundation for developing a new treatment of tense/aspect in XLE (and in general) in the following chapters.[21]

In the next section, I explore a second grammar formalism with respect to its encoding of tense/aspect, namely Universal Dependencies (UD). UD serves as a similar but more shallow approach to syntactic parsing and can thus serve as an alternative input for a semantic analysis.

2.4 Tense and aspect in Universal Dependencies

Like LFG and the ParGram project, universal dependencies strive to create cross-linguistically comparable syntactic analyses focusing on further semantic processing. In this section, I investigate dependency parsers in a fashion similar to the previous section on LFG. More concretely, I point out some crucial commonalities and differences between the two formalisms. This means that I explore their potential as input for a semantic annotation scheme of tense and aspect.

2.4.1 The development of UD

UD is the result of a combination of efforts, namely, the (universal) Stanford dependencies (De Marneffe et al., 2014; De Marneffe and Manning, 2008), the Google universal part-of-speech tagset (Petrov et al., 2011), and the Interset interlingua for morphosyntactic tagsets (Zeman, 2008). Two milestones in the development of UD are represented by version one (Nivre et al., 2016) and version two (Nivre et al., 2020) of the UD tagset that are represented by large cross-linguistic treebanks.[22] Similar to the ParGram project, the UD project

[21]Note also that this is not an exhaustive list of issues with tense/aspect in ParGram, but rather a series of examples that illustrate the general trends underlying the conflation of syntactic and semantic features within (some of) the grammars.

[22]Generally, the facts listed below refer to UD v2, i.e., the most recent version of UD.

strives toward cross-linguistic coverage. In fact, the UD project provides a wider coverage of languages. According to the UD homepage, "UD is an open community effort with over 300 contributors producing nearly 200 treebanks in over 100 languages."[23] The size of these treebanks varies greatly across languages, from under a thousand words to millions of words. Overall, UD provides a much broader spectrum compared to ParGram, but the coverage across languages varies significantly in both projects. This can be attributed to the most glaring difference between ParGram and UD. ParGram provides carefully handcrafted grammars with a deep level of analysis, whereas UD grammars often rely on machine learning methods to generate large-scale grammars.[24] Such grammars are often more robust than their handcrafted counterparts and are the de facto standard for modern NLP applications. As such, it makes sense to tailor a semantic resource that relies on syntactic input toward compatibility with UD. This will become especially relevant in chapter 6 when I present a concrete implementation of the annotation scheme developed in this thesis.

2.4.2 UD architecture and philosophy

The goal behind UD is to provide syntactic analyses that are ready for further semantic processing, similar to LFG's f-structure. However, UD stores all relevant information in a single dependency structure, which preserves the surface structure of an input string by applying the analysis to it in its linear order. This is illustrated in Figure 2.13, where the linear order is indicated by the enumeration of the tokens that make up the input string. A UD parse (minimally) consists of a directed graph created from typed, i.e., labeled dependencies, such as `obj`, `nsubj`, or `det` and a set of leaf nodes corresponding to the words in a given sentence and their respective POS-tags. Version two of universal dependencies also encourages further morphological

[23] https://universaldependencies.org/#language-v2

[24] Note, however, that the underlying data and annotations still need to be derived somehow. UD proposes two main possibilities:
i) Convert annotations on an existing treebank to UD annotations
ii) Generate data from scratch via annotation

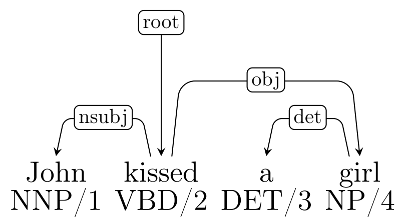

Figure 2.13: UD structure: *John kissed a girl*

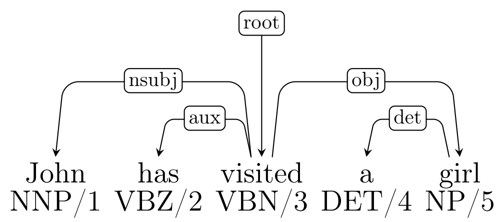

Figure 2.14: UD structure: *John has visited a girl*

annotation of leaf nodes. While they provide a set of basic features (of which tense and aspect are a part; see section 2.4.2.1), the general idea is that additional annotations are added based on the needs of a given language and/or the needs of downstream applications.

The semantic orientation of UD representations is perpetuated by the assumption that UD prefers content words as heads. Thus, function words like the quantifier in Figure 2.13 are subordinate to the noun that they modify. This also holds for aspectual auxiliaries, as illustrated by means of the perfect construction in Figure 2.14. This general principle of UD that content words form primary heads leads to some linguistically unsatisfying analyses. A simple example is given in Figure 2.15. On the left side, the UD structure suggests that *laughed* can subcategorize for an NP object, which is modified by a preposition. On the right side, a typical phrase-structure analysis shows that *laughed* may be combined with a PP, which is the subcategorization expected from the intransitive *(to) laugh*. Similar examples and additional inconsistencies of UD are discussed in Osborne and Gerdes (2019).

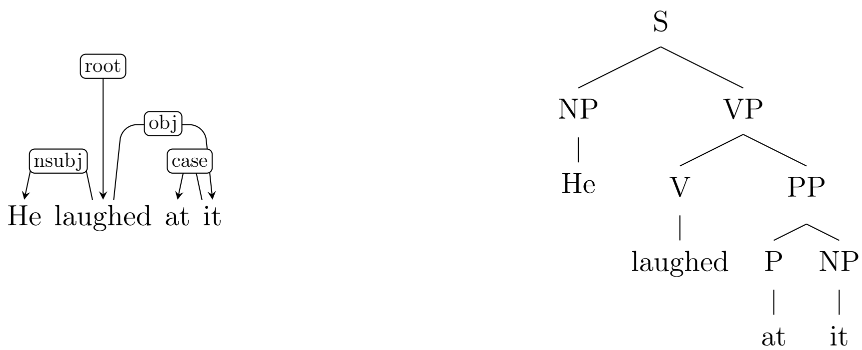

Figure 2.15: UD vs. phrase structure analysis: *He laughed at it.*

The UD effort acknowledges these issues. As the following quote suggests, syntactic correctness is sacrificed for other, presumably semantic, relations.

> "The notion of dependency has limits: Not all grammatical relations can be reduced to binary asymmetric relations between a syntactic head and a subordinate element, and some of our typed "dependency" relations therefore must be understood as convenient encodings of other relations without implications about syntactic headedness."
> https://universaldependencies.org/u/overview/syntax.html

In general, UD parses are aimed more toward usability for further semantic processing. Nonetheless, dependencies are encoded in terms of syntactic rather than semantic relations. This makes UD representations fairly similar to LFG's f-structure representations. Both approaches strive towards a level of parallelism across analyses for different languages encoded by elevating the role of content words in linguistic analyses. Functional words are encoded as sub-structures of those content words in UD and as sub-structures or attribute/value pairs in LFG.[25,26] Overall, given the fact that UD parses omit

[25]However, note that LFG preserves headedness information in some cases at f-structure. Taking the example in Figure 2.15, the corresponding f-structure specifies a substructure for *at it* that embeds *it* as an `OBJ` of *at*. Thus, the grammatical function satisfied by the PP (e.g., `OBL` or potentially `ADJUNCT`) is correctly headed by the preposition rather than the content word. In comparison, quantifiers are usually embedded under their corresponding content words in XLE analyses.

[26]Some other differences between f-structures and UD structures are based on the fact

certain syntactic information such as headedness but provide an output similar to f-structures, UD parses are generally simpler in terms of the encoding of linguistic information than LFG analyses which distribute different types of information across different projections. Thus, although UD preserves information about linear order and grammatical functions (i.e., it is encoding information from both c- and f-structure), its core assumption to favor content words as heads requires it to discard certain linguistic insights leading to sometimes striking differences in terms of analysis.

2.4.2.1 Universal depenencies: tense and aspect

This section focuses on the encoding of tense/aspect in UD in more detail and compares it to the encoding of tense/aspect within LFG presented in the previous section.

As, for example, shown by Meurer (2017), dependencies share certain commonalities with LFG's f-structure; namely, they encode grammatical relations in a flat structure. Furthermore, in both formalisms, structures are usually headed by content words (PREDs in the case of LFG, which are derived from content words in the c-structure via the lexicon). However, in contrast to LFG, UD preserves all functional words in its representation. This has been shown by virtue of the perfect construction in Figure 2.14. As I have explained in section 2.3.3, in LFG, some function words are reduced to functional features at the level of f-structure. This is possible due to the sharing of labor between f-structure and c-structure. However, UD also differs from the alternative analysis of auxiliaries in LFG that treats them

that LFG is a lexicon-driven formalism. As has been shown in the previous section, the PRED value encodes the semantic form of a content word, i.e., its arguments. This serves to preserve the felicity conditions of f-structures (uniqueness, completeness, and coherence), which do not hold in UD structures. Note also that the role of PRED structures at the syntactic level could be transferred to a proper semantic component, as argued in Andrews (2008). Another difference that UD handles differently is the distinction between argument and adjuncts, which I have not discussed before. This is because it is i) not relevant for the analyses presented in this thesis, and ii) there are ongoing arguments on the necessity of the argument/adjunct distinction in LFG and UD (Przepiórkowski, 2017; Przepiórkowski and Patejuk, 2018). Note that both of these points of tension are the results of a larger tension between syntax and semantics.

VB	Verb, base form
VBD	Verb, past tense
VBG	Verb, gerund or present participle
VBN	Verb, past participle
VBP	Verb, non-3rd person singular present
VBZ	Verb, 3rd person singular present

Table 2.8: Penn treebank POS tags

as elements that select for complements. In fact, the auxiliaries in UD are more easily reduced to attribute/value pairs on the main verb rather than converted to complement-taking verbs. This underlines the point that UD parses are relatively close to LFG f-structure in terms of the representation of tense/aspect. This is explored in more detail in this section.

UD grammar minimally encode tense/aspect information via different POS-tags in addition to the `aux` constructions illustrated above. The example mentioned above relies on the Penn treebank POS-tagset, which is illustrated in Table 2.8.

The Penn Treebank tag-set is specifically designed for verb forms in the English grammar and, as pointed out before, was adapted and modified as necessary for the needs of the UD project. However, as shown in the next paragraph, the original tag sets are streamlined via Interset (Zeman, 2008). In UD v2, POS-tags are replaced or expanded upon by the addition of morphological features at the word level, such as `tense` and `aspect`. These features are encoded on a language-specific basis, i.e., no cross-linguistically viable tag-set has been proposed to my knowledge. Thus, UD v2 follows the same philosophy as TimeML, treating tense/aspect as a language-specific issue. As noted in the annotation guidelines for UD v2, similarly to the ParGram effort, UD does not encode compound features (e.g., pluperfect) and only explicitly expressed features. As I have said before, in UD, this information is not accumulated at the head level of a UD representation.[27] However, Ramm et al. (2017) have provided a tense/aspect/mood annotation

[27] https://universaldependencies.org/u/feat/Tense.html

finite	**mood**	**tense**	**voice**
yes	ind	present presProg presPerf presPerfProg past pastProg pastPerf pastPerfProg future1 future1Prog futureII futureIIProg	act pass
	subj	condI condIProg condII condIIIProg	
no	-	-	-

Table 2.9: TMV combinations for English (Ramm et al., 2017)

for UD that shows that UD features can, in principle, be packed together in a fashion similar to LFG to describe verb complexes.[28] The resulting analyses for German, English, and French are comparable in terms of their level of detail to the tense and aspect LFG analyses employed in the ParGram project. In their system morphosyntactic tense/aspect/mood/voice systems are categorized in terms of a Tense/Mood/Voice(TMV) table. Table 2.9 shows the TMV feature space for English.

Although the approach by Ramm et al. (2017) provides a promising level of abstraction over UD analyses, it deviates from the idea of encoding tense and aspect separately. As work in ParGram has shown, this approach is challenging for cross-linguistically comparable annotation of tense/aspect.

[28]A similar approach is followed in Myers and Palmer (2019), providing a more precise tool for deriving the same annotations. Ramm et al.'s (2017) is rule-based, whereas Myers and Palmer (2019) train a neural network to deduce the relevant information.

Furthermore, it is not completely clear how a UD+TMV representation compares to an LFG (i.e., c- + f-structure representation) since a TMV table strips a verb complex of information irrelevant to determining tense/aspect. A concrete correspondence between UD parse and a TMV table would have to be established to compare the two approaches in more detail.
Overall, universal dependencies favor an explicit annotation of word-level features for tense/aspect. However, this approach can be enriched by using additional uses to abstract from word-level representations such as, for example, presented in Myers and Palmer (2019) and Ramm et al. (2017).

2.4.2.2 Conclusion on tense/aspect in linguistic annotations

I have presented two additional approaches to encoding tense/aspect information. Although both LFG and UD are technically syntactic formalisms, the information they encode is highly relevant for further semantic processing. On the other hand, I presented TimeML in section 2.2, which provides a semantic annotation for tense and aspect that is directly based on the morphosyntax. For each of the three systems presented so far, TimeML, LFG, and UD, there has been one central issue: the tension between expressing syntactic and semantic tense/aspect information. This has been clearly shown for TimeML, but similar points can be made for LFG and UD. The tension is particularly strong in LFG. Although the f-structure is considered part of the syntactic module of LFG, it sometimes abstracts away from morphosyntactic features to provide a semantic feature – a semantic feature with a concrete morphosyntactic exponent. However, this abstraction puts the features encoded within the f-structure on a continuum between more structurally motivated features and semantically motivated features. One of the main goals of this thesis is to explore this tension and relax it by providing a concrete proposal for semantically interpreting tense/aspect information. For this purpose, the next section lays some groundwork for the semantic framework pursued in this thesis. Rather than developing an entirely new framework, the semantics I propose for tense and aspect is grounded mainly in the formal semantic research on tense and aspect.

2.5 Motivating the ParTMA annotation scheme

This thesis aims to develop a cross-linguistically viable annotation scheme for the semantics of tense and aspect. Although the insights that drive this annotation scheme are heavily rooted in theoretical and practical approaches to LFG, the need for proper treatment of tense and aspect is more general. In this section, I motivate the ParTMA annotation scheme inside and outside of the LFG framework. In the process, I elaborate on some of the issues mentioned in the discussion of the three primary influences of this thesis, namely, the tension between form and meaning when encoding linguistic information as features. In my opinion, this point is explained best based on LFG's form-to-meaning mapping. Thus, I discuss the need for different layers of linguistic annotation in general and the need for a separate semantic annotation of linguistic features specifically. More concretely, the goal is to motivate a separate semantic feature structure that is independent of LFG's f-structure, but that can be connected to it via semantic mapping rules. Importantly, however, it can also be integrated into different pipelines. I.e., it is not assembled relative to some other structure (e.g., f-structure is assembled relative to c-structure), but rather is a template filled out by given syntax/semantics mapping.

The next part of this section explains the exact makeup of the semantic structure I propose for tense and aspect semantics. It is inspired by Reichenbach's (1947) approach to handling tense and aspect, briefly introduced in section 2.6.1. As I show in section 2.6.5, Reichenbach is often implemented following a hierarchical structure. By combining this structure with insights from LFG's treatment of semantics, I can propose a modular encoding of tense and aspect represented via a feature structure while preserving the hierarchical structure in the underlying semantics. The arguments for this are summarized in section 2.6.6.

Finally, section 2.5.2 explains how the syntax/semantics interface is envisioned based on insights that have been gained from the previous discussion. As proposed in the introduction to this thesis, I propose a layered system that uses various annotation tiers that roughly correspond to structurally, semantically, and pragmatically entailed meanings.

2.5.1 The tension between, structure, function, and meaning

The first main point is to clarify the role of functional features with respect to the surface structure and the semantic interpretation, which has been left open in the discussion of lexical functional syntax relying only on the basic assumption that functional features inform semantic analysis. The second point is to establish the semantics developed in this thesis such that they are usable within the projection structure of LFG but are also usable independently for researchers interested in the research of tense and aspect more generally.

In accordance with the philosophy of LFG, we can assume (at least) three distinct levels of features: structural features, functional features, and meaning features. Example (29) illustrates the intuitive distinction between a structural feature and a meaning feature. A structural feature can be understood as a function providing a surface form given a particular `PRED` value.[29] A semantic feature is a feature that corresponds to a (partial) semantic representation here encoded in terms of a lambda expression(i.e., also a function). In the simple example in (29b), the semantics of *leave* correspond to a function that describes a time interval during which a leaving event takes place. The *past* feature corresponds to a function that locates this time interval in the past with respect to the speech time encoded as t_0 via the precedence relation.[30]

(29) a. *leave* + *past* → *left*

b. *leave* + *past*
$\lambda t.leave(t) \quad + \quad \lambda P.\exists t[t \prec t_0 \wedge P(t)] \quad \rightarrow \quad \exists t[t \prec t_0 \wedge leave(t)]$

We have seen that LFG's functional structure is situated between surface

[29] Of course, the example is simplified and designed so that it ignores other features that affect the surface form, such as person and number, as well as features that ensure proper agreement between verb and subject.

[30] The semantics is also responsible for ensuring that the correct semantic expressions combine in the right order. However, this is not relevant to the present argument and remains implied at this point.

form and meaning, but what is the exact role of a functional feature? In the literature, the functional feature in LFG is expected to be able to do both form and meaning. On the one hand, generating a surface form string for any f-structure should be possible given the grammar that produced it. On the other hand, functional features should inform the semantic analysis, i.e., give rise to semantic representations (Butt et al., 1999).

In the following paragraphs, I discuss the idea of semantic features. Functional features are semantic features in that they inform semantic interpretation, but they are not purely semantic. Consider the present tense in German. The German language forms an opposition between inflectional past tense and inflectional present tense. However, in the semantic domain, the present tense seems to express non-pastness rather than present temporal reference. The examples in (30) illustrate this by combining past and present tense with compatible temporal modifiers.

″

(30) a. Peter besuch-t (*gestern/jetzt/morgen) seine Mutter.
Peter visit-PRES (yesterday/now/tomorrow) his mother
'Peter is visiting/will visit his mother'

b. Peter besuch-te (gestern/*jetzt/*morgen) seine Mutter.
Peter visit-PAST (yesterday/now/tomorrow) his mother
'Peter is visiting/will visit his mother'

German

On the morpho-syntactic side it makes sense to distinguish between (30a) and (30b) as values of the same feature since the surface structure for the verb varies accordingly. However, semantically, the *present tense* feature is less precise. Of course, we could rename the feature to accommodate these insights, but this solution is not always available. Again, German serves as an example. Löbner (2002) argues that the German present perfect is ambiguous between a semantic past tense (see (31) and an aspectual interpretation, allowing for a non-past interpretation of the present tense auxiliary in the perfect construction. This is illustrated in examples (32a) and (32b).

(31) Peter hat mich gestern angerufen.
Peter have.PRES me yesterday call.PSTPTCP
'Peter called me yesterday.'

(32) a. Ruf mich, sobald der Computer **abgestürzt ist**.
call me as soon as the computer **crash.PSTPTCP is**
'Call me as soon as the computer will have crashed.'
(Löbner, 2002)

b. Ruf mich, sobald der Computer die Datei heruntergeladen hat.
Call me, as soon as the computer the file download.PSTPTCP have.PRES
'Call me as soon as the computer has downloaded the file.'

Consequently, the functional feature of present tense can occur in the context of past, present, and future temporal reference given appropriate syntactic and semantic contexts. Thus, the problem of functional vs. semantic features is not only an issue of appropriate naming conventions but rather one of providing the exact interpretations for different configurations of functional features. This is precisely the goal of this thesis: to provide a system for exploring tense/aspect features not only as certain forms of verb (phrases) but rather as features sensitive to their syntactic and semantic context.

2.5.2 The syntax/semantic interface

The purpose of the syntax/semantics interface is to relax the tension between form and meaning, resulting in a more adequate representation of tense/aspect categories that is applicable cross-linguistically. For this modularization to be successful, a good mapping from one module to the other has to be specified, as has been established by the highly-modular LFG framework.

Inspired by LFG's overall architecture, the ParTMA annotation scheme proposes a syntax/semantics interface that can be understood as a correspondence function between a structural representation and a meaning representation. I.e., for every semantic annotation, there should be a morphosyntactic annotation that accompanies it. Furthermore, there should be a module that creates associations between a given morphosyntactic structure and a given semantic annotation. In this thesis, this module is covered by a rule-based system to illustrate the initial generation of resources. However, given enough resources, such a module could be implemented in terms of a machine learning pipeline that takes a morphosyntactic annotation as input and outputs a semantic annotation structure. This is left for future work.

The rule-based system mentioned above is inspired by works in computational semantics conducted with the XLE, e.g., Bobrow et al. (2007), Crouch (2005), and Crouch and King (2006). These systems make use of XLE's transfer component to derive semantic representations. The transfer component provides a platform for so-called description-by-analysi approaches to generating semantic representations. This approach was first pursued in Halvorsen (1983) and has since been explored in various works within the computational and theoretical LFG community (Andrews, 2008, 2009, 2010; Bobrow et al., 2007; Crouch, 2005; Crouch and King, 2006). The idea behind description-by-analysis is simple: an input structure is analyzed and translated into a semantic structure based on this analysis. For this, (generally) two kinds of methods exist: destructive methods (Bobrow et al., 2007; Crouch, 2005; Crouch and King, 2006) and non-destructive methods (Andrews, 2008, 2009, 2010). In the former, the input structure is destroyed, whereas, in the latter, the input structure is preserved.

The system presented in this thesis preserves the input structure but crucially uses recursive rules to develop multiple layers of analysis for a given input. This allows the system to implement a syntax/semantics interface with different levels of granularity in the rule-based approach. These levels are called tiers throughout this thesis. Tier 0 is provided by the syntactic input, Tier 1 comprises features that are generated directly from the morphosyntactic input, and Tier 2 contains features that are generated on the basis of the syntactic AND semantic context, i.e., rules that combine Tier 0 and Tier 1 (and potentially Tier 2 due to the recursive nature of the rules) features to generate a new feature. As example (33) shows, the syntax/semantics interface refines the foundational idea of LFG to decompose the form-to-meaning mapping. I have already suggested in this section that f-structure features are not strictly morphosyntactic in nature. Thus, when considering LFG as input, the analyses have to be evaluated concerning the question of whether they correspond to a tier 0 feature or a tier 1 feature. This is not necessary from a practical perspective, but it clarifies the generalizations concerning the syntax/semantics interface made by the corresponding rules.

(33)	Tier 0	Tier 1	Tier 2
	Form	$\Longleftrightarrow$	Meaning

It should be noted that there is no requirement for any feature to be evaluated at a higher level than Tier 1. Tier 1 corresponds to the default semantic interpretation of a feature, and it is thus by no means a necessity to further process it. This means that the ParTMA annotation template ultimately consists of the accumulation of the most specific feature from the various tiers for each attribute. In general, the further right in the diagram in (33), the more precise a feature is. In other words, the correspondence function between syntactic input and semantic output is decomposed into multiple tiers that ultimately all contribute to the same semantic structure.

One could consider a third tier, Tier 3, that enriches a ParTMA template with pragmatic information. I will entertain this idea at different stages of this thesis, but this thesis will not provide a fully fleshed-out formalism for this potential third, pragmatic tier. In this thesis, I focus mainly on purely semantic issues and issues at the interface of syntax and semantics.

2.6 Semantics in the ParTMA annotation scheme

The gist of the proposal in this thesis is simple so far: define semantic features whose values rely on their syntactic and semantic context. As the previous section explains, these features are determined via rewrite rules that populate a semantic feature template. Finally, formal semantics are used to validate the semantic annotations by generating semantic representations based on them. The remainder of this section serves to introduce the basic formal concepts of tense/aspect semantics and their role in the ParTMA annotation scheme. The following section is concerned with explaining tense and aspect from the perspective of temporal semantics.

2.6.1 The semantics of tense/aspect forms

In this thesis, I use (and have been using) the term *temporal reference* to describe the temporal semantic contribution of tense/aspect forms. Temporal reference has been explored formally in the domain of temporal logic at least since (Prior, 1962). However, an older, more descriptive work serves as the main influence for developing modern formal semantic accounts of temporal reference (Tonhauser, 2015). As already pointed out before, Reichenbach's (1947) approach to encoding the semantics of tense and aspect is integral to the work presented in this thesis. In general, tense/aspect annotation follows a common trend, namely, following the well-known system proposed by Reichenbach (1947) and its successors (as mentioned before, most famously Klein (1994)). This trend exists for a good reason. As Steedman (1997) bluntly puts it: "there [often] is one early modern piece of insightful descriptive work which most theories build upon, and which those who ignore seem doomed to reconstruct." Thus, this thesis focuses on relating current formal approaches to tense/aspect forms to Reichenbach's insights.

Reichenbach classifies simple and complex tense/aspect forms as tenses. The distinction between simple and complex here refers to the distinction between simple tenses (e.g., simple past) and compound forms such as the past perfect.[31] He proposed to use three distinct temporal variables to account for different tense/aspect forms: the evaluation or speech time S, which generally denotes the point of utterance, the reference time R, which provides a temporal perspective from which the event is perceived, and the event time E, i.e., the time during which the event takes place, or a state holds. In his preliminary work, Reichenbach (1947), he maps relations between these three time points onto specific tense-aspect forms. Consider the following constellations as representations of well-known tenses:

(34)	simple past:		E-R	S	
	simple present:			S-E-R	
	simple future:			S	E-R
	perfect:		E	S,R	
	past perfect:	E	R	S	

Simple past tense holds if the reference time and the event time coincide both preceding the speech time. In the case of a present tense interpretation, all three variables, S-E-R, coincide temporally. Future tense meaning is achieved by allocating the compound of E-R after the speech time. These three tenses do not necessarily warrant the existence of R. However, Reichenbach (1947) treats the English perfect constructions as tenses as well. The reference time is crucial to model them in his framework. This is in particular shown by the past perfect. The past perfect expresses – in this formulation – that a certain event has happened before a particular point in the more recent past. Similarly, the relative past future, *was going to* + infinitive, construction may be represented as $R \prec E \prec S$.
The system has been considerably refined over the years feeding into the Neo-Reichenbachian tradition leading to a unified analysis of tense and aspect in the sense that R and E are independently motivated by both the category of tense and the category of aspect (Klein, 1994). Here, the role of tense and aspect is clearly defined: tense relates the reference point R (Klein calls it topic time) to the evaluation time S, whereas aspect relates the evaluation

[31] This distinction is fleshed out in more detail in Chapter 3 considering also a cross-linguistic perspective.

time to the reference time. This means that Reichenbach's system encodes two levels of temporal reference, capturing both the contributions by tense as well as the contributions made by aspect.
In chapter 3, I introduce a relatively uncontroversial formal semantics for tense and aspect that combines Reichenbach's insight with a compositional system for assembling meaning structures. However, tense and particularly aspect do not only contribute temporal information. Thus, we need to extend the semantics further in 4.

2.6.2 The semantics of grammatical aspect

Studies in grammatical and inner aspect have revealed that tense/aspect semantics incorporate more than just temporal information. This means that it is not enough to simply map tense/aspect features on a temporal semantics as proposed by Reichenbach (1947) and (to some extent) by his successors. The perfective and imperfective aspects, or, more generally, grammatical aspect does not typically shift the evaluation time relative to the reference time, other than prospective aspect and perfect aspect mentioned in the previous section (exemplified in (35)).[32] Thus, for grammatical aspect, we need to specify how to interpret the relation described by E-R. While this relation implies temporal overlap between reference time and evaluation time, it also contributes to another axis of meaning.

(35)	past progressive:		E-R		S
	past + *going to* future:	R		E	S

This has been shown in particular for imperfective aspect and, more concretely, progressive aspect that can be seen as a sub-class of imperfective aspect (Comrie, 1976).

[32] What precisely this means for classifying such constructions in comparison to the perfective/imperfective opposition is discussed in detail in chapter 4.

(36) Aspectual oppositions by Comrie (1976)

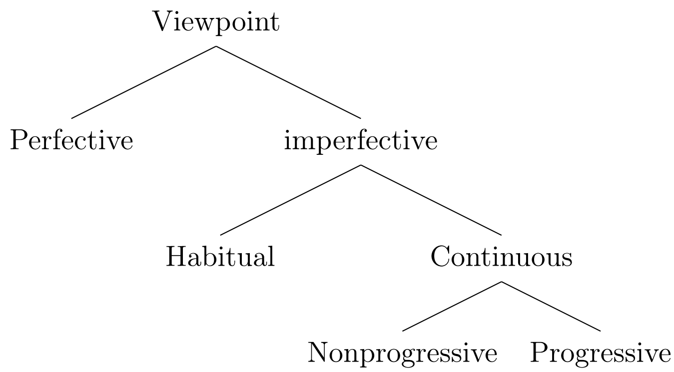

In broad strokes, approaches to analyzing the progressive can be divided into two categories: those that derive the progressive from the non-progressive form and those that do it the other way round. Arguably, the more popular approach is to derive the semantics of the progressive from the non-progressive or perfective semantics (Portner, 2019). For this reason, the progressive non-progressive distinction serves as a simple way of introducing the differences between imperfective and perfective aspect more generally.

A second division in formal analyses of the progressive is that between modal approaches and event structure approaches. The first approach – deriving the progressive meaning from the non-progressive meaning – is usually taken by researchers deriving the semantics of aspect from event structure, whereas the latter approach is supported by researchers who analyze the progressive as an instance of modality (see Portner 2019 for an overview over both of the two approaches and possible overlaps; see also Hallman 2009).

This thesis is based on the former approach, which is also in line with much of the descriptive literature on the progressive. Thus, the most basic observation about the progressive is still that it provides a partial view of an event based on the perspective provided by the reference time (Smith, 1991). In the (Neo-)Reichenbachian tradition, this relation boils to a subset relation between an eventuality's reference time and run time (derived via the function τ).

(37) $t \subseteq \tau(e)$

2.6.2.1 A second axis of meaning: existence

Semanticists aim to shed more light on the distinction between progressive and perfective by explaining the inference pattern called the imperfective paradox (Dowty, 1977) illustrated in (38). This phenomenon also serves to explain why the simple analysis in (37) runs into various challenges.

(38) a. John was drawing a circle. $\not\rightarrow$ John drew a circle.

b. John was pushing a cart. $\rightarrow$ John pushed a cart.

(39) The chicken was crossing the road (when it got hit by a truck).

The imperfective paradox is concerned with a completeness inference when applying the progressive to different kinds of eventuality descriptions (Dowty, 1977). When applied to a description that describes an action with a concrete endpoint, a telic description, then the completion of that eventuality is not implied. This intuition is most easily recognizable when combining the progressive with a creation verb such as in (38a), or a motion verb with a specific goal as in (39). However, when the progressive is applied to an open-ended process like *pushing a cart* in (38b without a specific goal and, thus, a specific endpoint, then the event description is entailed in the progressive form. Example (39) makes the intuition that no endpoint is necessarily reached concrete. The progressive can be used even if we actively deny the existence of the expected result state(Landman, 1992).

Thus, for eventualities with an inherent endpoint, the question is: how do we know that what we describe as a partial event is, in fact, part of a larger eventuality that didn't fully unfold yet? The imperfective paradox above indicates that we use an expression that leaves open the veridicality of parts of the event description in the actual world.

In section 2.1.2, I have introduced time intervals and eventualities as primary building blocks for tense and aspect information. Example (37) reflects this as well. The intricacies of the progressive are concerned with the question of whether an eventuality based on an event description is (partially) real (the modal approach) or whether we map a state onto an event description that fits our cognitive/perceptual input under it (Ramchand, 2018).

Analyses of the progressive show that modality, i.e., hypothetical reasoning, is the potential second axis of meaning covered by grammatical aspect. An early proposal for a modal analysis of the progressive has been made by Dowty (1977). Subsequently, modal analyses of the progressive rose to prominence within the tense and aspect literature, e.g. Arregui et al. (2011), Cipria and Roberts (2000), Deo (2009, to appear), Landman (1992), and Portner (1998) competing with an event-based semantics (e.g. Krifka 1998; Parsons 1990; Szabó 2004).

Ramchand (2018) proposes yet another approach to the semantics of the progressive. She argues that the progressive does not rely on a modal system but rather on a system based on identifying state specified via cognitive or perceptual information of an event property. In other words, this means that the progressive maps an event description onto a state that is identifiable with the event. This automatically accounts for the fact that the progressive coerces an event description into a stative description (De Swart, 1998). However, this insight is captured in any of the proposals mentioned above in one way or another. More interestingly, theories of imperfective and progressive aspect deal with the question of to what extent an eventuality exists in the real world.

In Chapter 4, I introduce a semantics that is mechanically similar to the modal approach but which crucially uses event structure to derive the entailments imposed by the imperfective paradox. As this suggests, the approach presented here is similar to Hallman (2009) approach, which also combines an event structure and modal approach. However, the reasoning for this choice is somewhat different. As it turns out, this choice makes the semantics more expressive from a computational perspective for which hypothetical reasoning by virtue of a modal semantics is untenable.[33]

Similar observations can be made for other flavors of imperfective aspect as well. They are discussed in detail in Chapter 5. The important basic observations here are that to account for the imperfective paradox, we require a semantics that allows us to account for the (partial) existence of eventualities

[33] See also section 2.6.4.

and entities (the semantics of viewpoint aspect) and a semantics that helps us distinguish different kinds of eventualities, such as actions, processes, and states (the semantics of lexical or inner aspect). The basics of inner aspect are discussed in the next section.

2.6.3 The semantics of inner aspect

As the imperfective paradox shows, analyses of the progressive aspect rely crucially on the semantics below viewpoint aspect (Dowty, 1977; Parsons, 1990; Ramchand, 2018), i.e., the semantics of the verb phrase. The corresponding aspectual features are subsumed under the term lexical or inner aspect, often also called Aktionsarten, often attributed to Vendler (1957).[34] Section 2.1.2 has already introduced a classification for eventualities based on Bach (1986). The linguistic tests for the temporal properties of verb phrases have been coined by Vendler (1957). One of these has already been discussed in the previous section: compatibility with continuous tenses, i.e., the progressive in English. The progressive cannot generally be applied to states bar some that can dynamically change, as illustrated in the following example.

(40) a. *John is knowing mathematics.
b. The socks are lying on the bed.

This gives us the first main distinction, namely, that between states and non-states. However, as the imperfective paradox has illustrated, there must be more fine-grained distinctions at play. In particular, I have distinguished between verb phrases that describe an eventuality with an inherent endpoint like *draw a circle* and verb phrases that describe an eventuality without an inherent endpoint, e.g., *push a cart*. This property of having a natural endpoint is generally called telicity as proposed by Garey (1957).The most prominent test for telicity in English has been introduced by Vendler (1957) and concerns the availability of different temporal modifier phrases for telic

[34]According to Filip et al. (2009), Agrell (1908) first coined the term; however, Vendler (1957) is generally accepted as the basis of the term in modern linguistics.

and atelic predications. This is illustrated in the following examples.

(41) a. Jamie drew a circle in ten seconds/?for ten seconds.

b. Jordan pushed the cart for an hour/*in an hour.

In (41a), the modifier *for ten seconds* is infelicitous if the sentence describes a single complete occurrence of the drawing eventuality.[35] On the other hand, in (41b), the continuation *in an hour* is completely infelicitous.
As the examples above illustrate, the telicity of a predication does not necessarily rely on the verb itself – the verb might not specify telicity at all – but on its arguments, in particular, its theme argument (Dowty, 1991). Consider the intransitive verb *(to) walk*. It interacts differently with the previously introduced test depending on its optional arguments. In example (42), the existence of a natural endpoint depends on the path described by the optional argument and the preposition that introduces it (Verkuyl, 1993). In (42b), the exact path remains unspecified leading to a predication that does not entail a natural endpoint, whereas the path described in (41a) has a clear endpoint.[36]

(42) a. Sam walked to the store in thirty minutes.

b. Sam walked in the park for thirty minutes.

As Krifka (1998), among others, has shown, the path need not be a literal path. Different kinds of arguments introduce different kinds of paths. The important distinction is between non-quantized and quantized arguments (Krifka, 1998). Dowty (1991) dubs the latter kind of arguments *incremental themes*. The idea is that, for a predicate to be telic, the event contributed by the verb must incrementally affect its theme argument. Thus, each part of the event affects a different part of the theme. An intuitive test for this is

[35]One could imagine the example to be acceptable in a situation where Jamie is drawing a circle multiple times (either the same or different circles).

[36]When observing verbs in their intransitive form, it is still possible to distinguish between telic and atelic forms, e.g., *walk* is atelic, whereas *collapse* is telic. Generally, bare verbs seem to have a disposition towards either telicity or atelicity, but their meaning varies in the syntactic, semantic, pragmatic, and extra-linguistic context.

that for the subinterval property (Bennett and Partee, 1978). In (43a), the allotted time roughly corresponds to the distance travelled. I.e., we would expect that Sam has only traveled roughly two-thirds of the way after twenty minutes. Thus, each arbitrarily small interval of the thirty minutes maps onto a different part of the path to the store, which in turn maps onto a different sub-event.[37] In comparison, in (43b), the path is not incremental. In fact, Sam could walk the same fifty-meter stretch back and forth for an hour. As long as she stays in the park, the predication is true. In other words, in (43a), *Sam walked to the store* is not true of all its sub-parts, but in (43b), *Sam walked in the park* is true of all sufficiently large sub-parts which becomes apparent by mapping the predication on the given time scale.[38] Aside from this path example, a similar observation can be made for singular indefinites vs. bare plurals (see (44)).

(43) a. Sam walked to the store in thirty minutes. $\not\rightarrow$ Sam walked to the store in twenty minutes.

b. Sam walked in the park for thirty minutes. $\rightarrow$ Sam walked in the park for twenty minutes.

(44) a. Jordan ate an apple in five minutes. $\not\rightarrow$ Jordan ate an apple in two minutes.

b. Jordan ate apples for five minutes $\rightarrow$ Jordan ate apples for two minutes.

For Vendler (1957) the *in/for* test distinguishes two groups of verbs classes. The first group consists of states and activities. Other than states, activities are compatible with the progressive, e.g.:

(45) John was learning for an hour.

[37]It is generally assumed that the progression happens monotonically, but this is not strictly the case for all eventualities (Croft, 2012).

[38]As, e.g., Champollion (2017) points out that properties such as the subinterval property adhere to certain restrictions on the granularity of the underlying event. If we look for subevents of the predication *(to) walk*, at some point, we will reach a level where the property does not hold anymore. For example, it is not clear if *taking one step forward* is already a *walking* eventuality. However, for the purposes of this thesis, we do not have to pay particular attention to these fine-grained distinctions.

However, as supported by the *in/for* test, activities do not have an inherent endpoint and can be continued, in principle, indefinitely. thus, *eat apples* and *walk in the park* are also activities. The intuition here is that activities describe an undirected process.
The second group consists of achievements and accomplishments. Both of these verb classes have a process component, making them compatible with the progressive. They are also both telic. I.e., they are compatible with *in*-adverbials, but not *for*-adverbials. The difference between these two verb classes lies in their duration. Achievements are momentary eventualities like *reach the summit* or *arrive*. Accomplishments, on the other hand, are protracted events like *write a letter* or *read a book*. Both classes are compatible with *in*-adverbials. Thus, even momentary events are available for temporal modification. In other words, they can be used to describe eventualities that take place over an extended period of time. The difference here can be diagnosed with a test similar to the subinterval property test. However, this time the grammatical aspect in the entailed sentence is changed to the progressive. Thus, if someone wrote a letter, he has been writing that letter at each point during the eventuality. In comparison, if someone has reached the summit, then the process of getting to the summit is part of the predication (the part that can be temporally modified), but only the final part of that predication describes that the summit has been reached.[39]

(46) a. Alex wrote a letter in ten minutes. $\rightarrow$ Alex was writing a letter for ten minutes.

b. Alex reached the summit in ten minutes $\not\rightarrow$?Alex was reaching the summit for five minutes.

In summary, inner aspect is determined by the properties of the verb and its arguments, particularly the theme argument in the sense of (Dowty, 1991). Table 2.10 summarizes the main distinctions of inner aspect in the tradition of Vendler (1957). States are distinguished from non-states, a heterogenous

[39] This last distinction also sheds further light on the progressive, which seems to make a predication atelic. However, a detailed examination of the interaction between grammatical and inner aspect is postponed until chapter 4.

	durative	change	endpoint
state	+	-	-
activity	+	+	-
achievement	-	+	+
accomplishment	+	+	+

Table 2.10: Aktionsart

group consisting of activities, achievements, and accomplishments. These three verb classes have a process component (*change* in table 2.10), but only achievements and accomplishments have an inherent endpoint. Thus, they are telic. They are distinguished from each other in terms of whether they are durative. Achievements are momentary, whereas accomplishments are protracted.

Inner aspect is explored in detail in chapter 4. This includes more fine-grained distinctions within the English language but also the state of research of inner aspect cross-linguistically. In fact, we will see that prevalent tests like the *in/for* test do not behave consistently across languages in the sense that they do not allow for the systematic distinction of verb classes according to grammatical reflexes. Thus, while tense and grammatical aspect already have a relatively promising foundation for cross-linguistic research, inner aspect has to be investigated more closely.

2.6.4 Modal semantics

In the previous sections, the basics of grammatical and inner aspect have been discussed. There, we have seen that inner aspect is defined based on the verb and its argument(s) and that grammatical aspect interacts differently with different verb classes. In section 2.6.2.1, I have introduced the idea of treating imperfective aspect, particularly the progressive, as an instance of modality. This allows us to cover both the properties of the eventuality and its arguments that might or might not come into existence under the progressive. In this section, I describe the semantics of modality more generally to finalize the ingredients for the semantics of tense and aspect proposed in this thesis. Modality has played a role in linguistic analyses as early as Lewis (1918) when considering the semantics of certain connectives in traditional logic.[40] These early proposals relied on the notion of possibility and necessity. The traditional building blocks of modal logic.[41]

Possibility and necessity are prototypically exemplified by virtue of modal auxiliaries in (47). Example (47a) expresses the possibility that Jordan visits her aunt and (47a), conversely, describes the necessity of Jordan visiting her aunt. Intuitively, possibility and necessity make assertions about actual and hypothetical situations. Possibility asserts that there is at least one hypothetical situation (that is presently relevant) in which Jordan would visit her aunt. Necessity expresses that Jordan visits her aunt in every hypothetical situation that is presently relevant.[42]

(47)
  a. Jordan may visit her aunt.
  b. It is possible that May visits her aunt.

  a. Jordan must visit her aunt.
  b. It is necessary for Jordan to visit her aunt.

The distinction between possibility and necessity is called modal force and only reflects the grammaticalization of modality in English. Other languages may allow for additional distinctions requiring a more fine-grained system for determining modal force.[43]

The second ingredient of modality is the modal flavor. This property of modals is not grammaticalized in English. Thus, *may* is of weak modal force but can receive various modal flavors. This is exemplified in (48. In (48a),

[40] Actually, even before that (see Ballarin 2021), but I take this as the point of entry for the modern semantics of modality.

[41] The basis for modern modal logics has been provided by Kripke (1963, 2014). However, exploring the formalities in detail goes beyond the scope of this conceptual introduction.

[42] In this section, I use the term situation colloquially. In modal logic, necessity and possibility refer to properties of the accessibility relation that holds between elements, usually possible worlds, in a model (Priest, 2008). Our actual world is one such possible world, and modal logic presupposes that other such worlds exist. We, then, reason about possibility and necessity by referring to these possible worlds and their relation to the actual world.

[43] For example, more recent advances in modal semantics propose a degree-based analysis that allows researchers to capture gradual changes in modal force and comparative differences between modal forces (Kratzer, 2012).

every passenger is allowed to only bring one piece of hand luggage (though they are free not to bring any luggage at all). On the other hand, in (48b), Jordan is not allowed to miss the flight. Rather, the utterance suggests that the speaker thinks it is possible that Jordan has missed the flight according to what they know or what evidence they have.

(48) a. Every passenger may bring at most one piece of hand luggage.

b. Jordan may have missed the flight.

Generally, modal flavor is dictated by the context.[44] We can make this point more apparent by looking at more fine-grained distinctions between various modals. For example, they can be used to express wishes (bouletic modality), as in (49), belief (doxastic modality), or knowledge (epistemic modality; in (50)), requirements imposed by some kind of law (deontic modality; in (51)), or abilities.

(49) John wishes to own a fast car.

(50) Jordan believes that Jamie is sick.

(51) Every citizen has to pay taxes.

(52) Sam is able to open the door.

Regarding example (47a), I claimed that *must* expresses necessity. However, this necessity is only one part of the meaning of the modal auxiliary. In fact, depending on the context, *must* can cover most of the modal interpretations mentioned above and more. As Kratzer (2012, Ch. 1) describes, we can make this explicit by outsourcing the context into an *in view of*-phrase, for example:

(53) a. In view of her family duties, Jordan must visit her aunt. *deontic*

[44] Example (48) incidentally also shows that modal flavor is also sensitive to the more immediate syntactic and semantic context. In combination with the perfect in (48b), the interpretation based on knowledge/evidence seems to be the most sensible choice. In example (48b), the context favors the *permission* interpretation but a elaborate (albeit maybe contrived) context may allow for other interpretations as well.

b. In view of what I know about her plans, Jordan must visit her aunt. *epistemic*

c. In view of Jordan's love for her aunt, she must visit her aunt. *bouletic*

This suggests that at least part of the meaning of *must* is highly context-sensitive and leads to ambiguity in many cases since concrete specifications such as the one in (53) are not the norm in natural language. This apparent flexibility of modal auxiliaries is further expanded when we look at the different flavors more closely. To exemplify this, let us look at deontic modals. They further vary depending on the underlying law. Is it a law of nature, a law imposed by the government, or just the conventions a family has settled for? All in all, modal auxiliaries are extremely sensitive to the context in which they occur.

Furthermore, modality is associated with counterfactual reasoning, which is often associated with certain conditionals such as the one in (54a), but counterfactuals, i.e., expressions that describe states-of-affairs that oppose our view on reality also occur in certain embedding verbs, e.g., *wish*, see (54b).

(54) a. If Jordan were sick, she would not be at work.

b. Kim wishes she had a fast car.

Consider (54a) in the context that Jordan is at work. A co-worker claims that she is sick, but (54a), as a response, denies the truth of this claim. The antecedent of the conditional is counterfactual: Jordan is not actually sick, but if she were sick, she would not be at work. The counterfactuality in (54b) is even more clear. If the embedded proposition that Kim has a fast car were true, then there would be no need for Kim to have this wish.

Incidentally, in English, counterfactuality is closely related to the past tense. Although both (54a) and 54b) do not make assertions about past situations but about hypothetical situations that describe a situation that is untrue in the actual world. Fleischman (1989) shows that it is not uncommon, particularly for tenses, to expand their basic meaning to other dimensions.

Recent research on modality suggests that there exist countless types of modality that can be fine-tuned in one way or another (Kratzer et al., 1991). Thus, the kind of modality explained in the previous section is just a small branch of a big tree. The examples given in this section only cover some significant categories. However, the main ingredients have been provided: modal force and modal flavor. They provide a solid basis for discussing modal semantics in relation to tense and aspect in this thesis.

2.6.5 Tense and aspect in formal semantics

I have provided a descriptive introduction to the semantic categories explored in this thesis in the previous sections. This section finalizes the introduction to semantics by touching on some key points in formalizing these categories. First of all, this thesis follows a typical formal semantic approach assuming compositionality (Frege et al., 1892) and a truth-conditional semantics in the Montague tradition (Montague, 1970a,b). However, for the sake of this thesis, I presuppose an understanding of the underlying ideas and their implications. Thus, without further ado, I will delve into some concrete assumptions about the formalization of tense/aspect semantics.

The foundation for a (Neo-)Reichenbachian tense/aspect approach in formal semantics has been laid out by Kratzer (1998) based on Klein (1994). In her work, tense and aspect are assumed to apply in a particular order to descriptions of eventualities. The semantic interpreter I present in this thesis makes crucial use of this order given in (55).

(55)

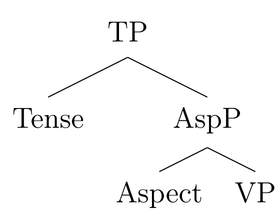

In Kratzer (1998) this structure is motivated through functions that select for different kinds of arguments (e.g., aspect selects for a verb phrase and tense selects for an aspect phrase). Traditionally, this hierarchy is stipulated by the syntax as proposed in the Minimalist Program (see Chomsky (1957)).

Ramchand and Svenonius (2014) argue for the existence of different sources of ordering in such a functional hierarchy, ranging from syntactic to semantic, as well as extra-linguistic factors. This approach is pushed further in Ramchand (2018). No matter which position one subscribes to, the structure given in (55) serves as the foundation for many formal semantic approaches to tense/aspect.

The hierarchical approach in (55) stands in contrast to various approaches that treat tense/aspect information as constraints in a more flat semantic representation (generally based on DRS-representations; Kamp (1981)), most prominently the approach presented in Partee (1984) and subsequent work.[45] Similar approaches have also been used for the semantics employed in TimeML (Bunt, 2010; Bunt and Overbeeke, 2008b), but also within the LFG community, e.g., Bary and Haug (2011a), Bary and Haug (2011b), Fry (2005), Haug (2008), and Lowe (2019), among many others.

These approaches present a salient alternative to the semantics pursued in this thesis. However, and this is important, they are not strictly incompatible with the approach presented in this thesis either.[46] Furthermore, although the theoretical adequacy is still subject to discussion, establishing and following the tree simplifies the underlying semantics. For example, in a flat structure, the order of application of semantic features is arbitrary, which leads to much spurious ambiguity in the semantic composition process (if not accounted for elsewhere). On the other hand, interpreting the semantics following the functional hierarchy (to some extent) filters out such ambiguities before they even arise.

From a purely mechanical standpoint, this thesis does not deviate from the norm, but Chapter 4 will show how to implement viewpoint aspect and inner aspect in a self-contained way, expanding on Zymla (2019). The idea roughly follows Ramchand (2018) in that the progressive is explained as taking an

[45]Note, that an DRS approach does not necessarily imply a flat tense/aspect structure. For example, De Swart (1998, 2000) implement Progressive aspect as a scope-taking operator.

[46]By this, I mean that the approach presented here can produce the same semantic representations as those proposed by the cited work, although the process (i.e., the semantic composition) by which the present system gets there may be different.

atemporal event description dictated by cause and effect (Ramchand, 2008) as argument and realizing only the part of if that matches the truth conditions an event description marked with the progressive. As shown in Chapter 4, this approach allows us to capture the insights provided by modal analyses of the imperfective using a *modal part-of* relation (Arregui, 2005; Arregui et al., 2011) while also preserving ideas from event-based approaches.

2.6.6 A flat representation of the tense/aspect hierarchy

For the semantics, similar to the TimeML annotation scheme, I distinguish between an annotation in terms of semantic features (the syntax) and a corresponding interpretation (the semantics of the annotation scheme). The syntax is encoded in terms of an attribute/value matrix similar to LFG's f-structure. This structure is supplied with a concrete (and complete) semantic interpretation. For the sake of this thesis, I focus on simple first-order logic representations that are derived compositionally from the building blocks of the ParTMA annotation scheme. This means concretely that each annotated attribute/value pair in the ParTMA annotation scheme "fills a hole" in the compositional framework that underlies the ParTMA template. The feature space of this template is presented in Figure 2.16.

The feature space is divided into three main categories: `TEMP-REF`, `VIEWPOINT`, and `LEX-ASP`. Each of these categories has an independent internal structure and feature space that is motivated in chapters 3, 4, and 5 of this thesis (theoretically in chapters 3 and 4, and practically in chapter 5). Although the structures can be annotated independently, ultimately, the subtleties encoded in tense/aspect semantics arise from the interaction of these categories. This part is covered by the "backend" of the annotation scheme, i.e., a semantic interpretation component that, as the name suggests, serves as the interpreter of a given annotated ParTMA structure. This interpreter follows the hierarchy specified in the previous section based on the Reichenbachian tense system and its implementation in formal semantics.

```
[ID] [ LEX-ASP    [ class ::= 'event' | 'process' | 'state' | 'unspec'
                    extent ::= 'mom' | 'dur'
                    FPS [ initP  [ init ::= 'fid'
                                   i-subj ::= 'fid' ]
                          procP  [ proc ::= 'fid'
                                   p-subj ::= 'fid' ]
                          resP   [ res ::= 'fid'
                                   r-subj ::= 'fid' ]
                          rheme  [ head ::= 'fid'
                                   type ::= 'location' | 'path' |
                                   'i-theme' | 'unspec' ] ] ]
       TEMP-REF   [ t-ref ::= 'past' | 'present' | 'future' |
                    'non-past' | 'non-present' | 'unspec'
                    t-restr ::= 'immediate' | 'recent' | 'non-recent'
                    eval ::= 'fid' ]
       VIEWPOINT  [ aspect ::= 'perfective' | 'imperfective'
                    restr ::= 'progressive' | 'habitual' | 'ongoing' |
                    'terminated' | 'culminated' ] ]
```

Figure 2.16: ParTMA eventuality template

In sum, this thesis follows the steps down the functional hierarchy presented in (55). This means that the temporal node is discussed in the next chapter, whereas both grammatical and inner aspect are discussed in 4. By the end of this journey, we will arrive at a formally and theoretically sound, cross-linguistically viable representation of tense and aspect. The computational needs of these semantics are catered to in chapter 5. There, a computationally viable compositional framework for the tense/aspect semantics developed in this thesis is described. This system uses two layers of interpretation i) linguistic rewriting and ii) Glue semantics. Attaching the formal semantics to a Glue semantics framework allows me to preserve the insights gained from the formal literature on tense/aspect while providing sufficient flexibility to work with a flat feature structure as input that has been motivated in this section.

2.7 Summary

This section has set up the stage for the following chapters. This was achieved by presenting several computational linguistic approaches to encoding tense/aspect information at a syntactic, semantic, and intermediate level, namely, TimeML, UD, and LFG. Through this, I have set up one of the main topics of this paper, namely, the tension between the syntactic and semantic encoding of tense and aspect in its current computational treatments. TimeML has established that semantic features should be backed by a concrete semantic framework within an annotation scheme. Such a semantic framework has neither been developed for UD nor computational LFG, as represented by the ParGram project. On the other hand, TimeML ignores the intricacies that arise from the interaction of tense/aspect with the surrounding morphosyntactic, semantic, and pragmatic context. This thesis remedies this by combining the two syntactic frameworks described in this section with a concrete semantics. The semantics is represented in terms of an attribute-value matrix that encodes a well-established formal framework for tense/aspect semantics. Although the resulting structure can be, in principle, annotated as a stand-alone semantic object, putting it on top of a syntactic analysis allows me to implement the syntax, semantics, and the interface between the two clearly and concisely that is grounded in formal linguistic research.

In the following two chapters, I establish a detailed semantics for tense (chapter 3) and aspect, including both grammatical and lexical aspect (chapter 4) and a syntax/semantics interface that explains the relation between syntax and semantics. Finally, in chapter 5, I present a concrete implementation of this system based on the groundwork that I laid out in this chapter and the more detailed exploration of the categories of tense and aspect in the following chapters.

Chapter 3

The annotation of temporal reference

> [...]when the Rabbit actually took a watch out of its waistcoat-pocket, and looked at it, and then hurried on, Alice started to her feet, for it flashed across her mind that she had never before seen a rabbit with either a waistcoat-pocket, or a watch to take out of it, and burning with curiosity, she ran across the field after it.
>
> Lewis Caroll, *Alice in Wonderland*

This chapter and the following chapter describe the development of the ParTMA annotation scheme for tense and aspect. The main goal of this chapter is to explore the semantic property of temporal reference. As already shown in the previous chapter, this semantic property relates two distinct time points or intervals provided in the context of evaluating an utterance, i.e., evaluation time(S), reference time(R), and eventuality time(E) (Reichenbach, 1947). There generally exist two types of categories that are relevant for temporal reference: those that relate the evaluation time and reference time (the classic tenses primarily examined in this chapter) and those that determine the relation between reference time and eventuality time.[1]

[1] The relation between the time of evaluation and the eventuality time naturally emerges from these two relations.

Tenses are the prototypical category associated with temporal reference. In the absence of aspectual modifiers, tenses are usually sufficient to temporally locate R and E since they overlap in this case. This is illustrated in (1), where the content described by the verb is situated in the past.

(1) The people kill-**ed** the king.
kill-PAST

The second type of category relevant for determining temporal reference is usually seen as a subset of aspectual categories. As I will explain in section 3.1.6, aspectual categories form a less homogenous group than tenses. We can distinguish between categories that specify overlap between R and E, generally instances of grammatical aspect, and categories where there is no overlap between R and E. Examples of this are the perfect construction and prospective aspect, as illustrated in (2). The dashed lines describe the contribution of the past tense in both instances and the unbroken lines encode the contribution of the aspectual category.

(2) a. Jamie **had** **finish-ed** his paper.
have.PAST finish-PSTPTCP
Temporal progression diagram:

— *finish his paper* — *R* — *S* —>

b. Sam was **going to visit** her aunt.
be.PAST PROSP visit.INF
Temporal progression diagram:

— *R* — *visit her aunt* — *S* —>

In this chapter, I primarily discuss this second type of aspectual categories. Grammatical aspect and inner or lexical aspect are discussed in chapter 4. This is due to the fact that grammatical and inner aspect are not explicit temporal ordering relators but rather invite certain inferences concerning temporal reference (this is elaborated upon in section 3.1.6). For the reminder of this thesis, I call aspectual categories that introduce explicit temporal ordering aspectual tenses.

Temporal reference encompasses at least three different meanings: anteriority, simultaneity, and posteriority relative to some temporal anchor, usually the speech/evaluation time, often also called deictic center (Comrie, 1985). However, temporal reference also emerges in terms of meanings such as non-future, non-past, or non-present, which can be understood as combinations of the three basic meanings attributed to tense. For example, non-future then means that the respective time interval is either simultaneous to the now or precedes it (or both).[2] Temporal reference encompasses the often-used term semantic tense (semantic tense is temporal reference contributed by a morphosyntactic tense feature; Von Stechow 2009). However, as we will see in this chapter, temporal reference is not always handled by tense (or by tense alone). From a cross-linguistic perspective, the strategies vary heavily both on a syntactic and semantic level. Thus, the category of tense only serves as a stepping stone to understanding the broader semantic category of temporal reference, which also includes aspectual tenses, grammatical aspect, mood markers, and the context.

This chapter is structured as follows: First, I provide a descriptive overview of the morphosyntactic and semantic properties of tense and other categories that affect temporal reference in section 3.1 on a descriptive level. Then, I explain how these properties have affected the formation of formal theories describing the structural and meaning features of tense/aspect categories in section 3.2. This, ultimately, leads to a formal semantics for temporal reference. Exploring the formal tools used to describe tense and aspect in form and meaning provides the backbone of the annotation developed for the category. The annotation scheme consists of two components: a syntax/semantics mapping and a template for semantic annotation grounded in formal semantic research. This is explained in detail in section 3.4.

[2] In general, the treatment of temporal reference in formal semantics can be traced back to (at least) Prior's (1962) temporal logic. However, as mentioned in the previous chapter, the most influential framework is that of Reichenbach (1947) and its successors. As discussed in the previous chapter, computational approaches have, furthermore, been influenced by (Allen, 1990) and refer to more general temporal relations going beyond those expressed in tense and aspect.

3.1 Tense/aspect descriptively

The default grammatical device for encoding temporal reference in many languages is the category of tense. It is, thus, crucial to understand this category before extending the understanding of temporal reference proposed in this thesis to other categories. In this section, I provide a systematic introduction to the structural variation of linguistic elements called tenses. First, it is important to declare what I understand to be an instance of *tense*: Early works such as Reichenbach (1947) understand tenses as all possible structural variations of the form of the verb affecting temporal interpretation, thus, including both temporal and aspectual markers. This means *present perfect* and *past perfect* are tenses just as the *simple present* or the *simple past*. However, more recently, tense categories and aspect categories are analyzed more modularly (Comrie, 1976). Thus, following Comrie (1985), tense categories are those categories that temporally relate the semantic content of a linguistic expression with some point of evaluation (usually, but not necessarily, the time of speaking).

Consequently, I use the terms tense and aspect roughly as proposed by Comrie in the previously cited works. This means that tenses in the sense of Reichenbach (1947), let's call them compound tenses, can be decomposed into building blocks contributing tense information and building blocks that provide aspectual information. To avoid confusion, I call compound tenses *tense/aspect forms*, distinguishing between simple and complex tense/aspect forms. The *simple past* (*V+ed)* in English is a simple tense/aspect form, since it cannot be (morphosyntactically) decomposed into a tense and an aspect component. On the other hand, the *past progressive* (*be+Past V+ing*) is a complex tense/aspect form, since it can be decomposed into *past* and *be V+ing*.[3] Thus, combinability also plays an essential role in understanding the structural properties of tense (and aspect) categories.

Before diving in to concrete examples, one last point is wort mentioning:

[3]Note that this is also the approach that has generally been pursued in the annotation schemes presented in the previous chapter.

simple tense/aspect forms may still contribute complex meanings, i.e., both temporal and aspectual meanings. In such a case, I point out the existence of additional semantic properties beyond temporal reference but a more exact description of these additional properties is postponed to the next chapter, or left for future work. Generally, it is crucial to distinguish between structural and meaning features associated with tense and aspect as already highlighted in the background chapter.

3.1.1 The core structural properties of tense/aspect forms

Extensive typological studies on tense and aspect systems such as the one conducted by Dahl (1985) show that tense varies heavily concerning its structural properties. In fact, tense can be encoded in various forms on the continuum of structural representation that reaches from lexical expressions to syntactic and morphological expressions (Bybee and Dahl, 1989).
Following Dahl (1985) and Bybee (1985), I consider three main distinctions with respect to tense/aspect morphosyntax: inflectional categories, periphrastic categories, and derivational categories. As proposed in Bybee and Dahl (1989), only the former two are considered as structural exponents of the category of tense (encompassing past, present, and future-oriented tenses). In addition, to this distinction, zero-marked tenses are also considered ((Bybee, 1994); see Jóhannsdóttir and Matthewson (2008) and Reis Silva and Matthewson (2007) for concrete examples).[4] Although the future tense is often periphrastic in Germanic languages, more generally, the distribution between periphrastic future tenses and inflectional future tenses seems to even out cross-linguistically (Bybee and Dahl, 1989). Derivational morphology plays a role in chapter 4 since it is associated with aspectual properties (as,

[4]Bybee (1985), Bybee and Dahl (1989), and Dahl (1985) argue that certain meanings gravitate towards certain modes of expression. This is further supported by early work in the semantics of tense and aspect, for example, in Comrie (1976, 1985). An example of this would be that the past tense is prototypically encoded in terms of an inflectional category. On the other hand, temporal aspectual categories (aspectual tenses) like the perfect or the progressive are more likely to be expressed periphrastically. (I tentatively use the the term temporal aspectual category for the perfect and progressive since they bundle temporal reference with additional aspectual features (Portner, 2019).)

for example, observable in many Slavic languages (Dahl, 1985)).
In tense/aspect systems, several different structural devices are often used to cover the full spectrum of different meanings associated with both tense and aspect categories. Thus, it is rarely possible to say that a category is an inflectional category in some language. For example, in the case of tense, it is usually necessary to describe past, present, and future tense independently. This is the approach that I take for the reminder of this chapter.

3.1.2 Inflectional and periphrastic tenses

The most commonly researched instances of the category of tense are probably found in the Germanic tense systems represented, for example, in terms of the languages German and English. Tense marking is obligatory and part of the inflectional system of these languages; furthermore, as pointed out in the previous chapter, it plays a crucial role in determining finiteness. Example (3) shows the dedicated past tense forms in English and German.[5]

(3) a. The people kill**-ed** the king.

b. Das Volk töte-**te** den König
the people kill-PAST the king
'The people killed the king' *German*

Inflectional tenses may carry specific references to past, present, or future but, in some cases, also express less specific meanings such as non-past or non-future. For example, the German past tense and present tense form a paradigm that is complementarily distributed among past and non-past contexts (exemplified in (4).

(4) a. Peter besuch-t (*gestern/jetzt/morgen) seine Mutter.
Peter visit-PRES (yesterday/now/tomorrow) his mother
'Peter is visiting/will visit his mother'

[5]Inflectional paradigms usually have irregular forms. In this thesis, I try to rely on regular forms whenever possible to avoid confusion in the structural domain. See Bybee (1995) for an overview of the apparent non-difference between regular and irregular forms.

b. Peter besuch-te (gestern/*jetzt/*morgen) seine Mutter.
Peter visit-PAST (yesterday/now/tomorrow) his mother
'Peter is visiting/will visit his mother'

German

In addition, there exists a periphrastic form in German that can be used for future temporal reference. However, Comrie (1985) describes these auxiliary future forms in German and other European languages as secondary at best (p. 49). He elaborates on this by explaining that these forms are merely used to counteract the vagueness of present tense, which belongs to the primarily grammaticalized tenses. The German future tense is illustrated in example (5). Thus, in Comrie's (1985) view, German can be seen as a language that exemplifies a past/non-past distinction at the level of tense but possesses additional means of encoding future temporal reference.
According to Ehrich (1987), the *wird + V.INF* future tense form provides an alternative rather than a secondary form. She claims that the future tense may always be used, but it competes with the present tense and temporal adverbs. The approach in this thesis supports this insight: temporal reference may be expressed in several ways and is influenced by grammatical and inner aspect (see Chapter 4), as well as temporal modification (i.e., temporal adverbs) and the context.[6]

(5) a. Das Volk wird den König töten
The people be.FUT the king kill.INF
'The people will kill the king' *German*

In contrast, English is more strict in this regard. Although it is not entirely impossible to express future time reference in English using the present tense, it is usually restricted to present time reference except in particular contexts (Comrie, 1985). Despite the variation in tense systems, there is little controversy in the literature about how these obligatory inflectional tenses

[6]Note, however, that the present annotation scheme makes no claims about any psycholinguistic realities. The distinction between structural and meaning features advocated for in this thesis is focused on descriptive accuracy and computational viability.

are to be interpreted, namely, as overt instances of temporal reference, as specified in the introduction to this chapter.

Inflectional tense forms in Germanic languages are outsourced to an auxiliary in complex tense/aspect forms, as illustrated for English in (6). As shown there, inflectional tense applies to the outermost aspectual auxiliary. In comparison, the future in English (and the periphrastic future in German) is applied externally to the aspectual modification as is presented in (7). Thus, we can further distinguish inflectional tenses from periphrastic tense forms.[7]

(6) a. John visit-**ed** his mother. **past**

b. John *ha*-**s** *visited* his mother. **present** + *past participle*

c. John *ha*-**d** *visited* his mother. **past** + *past participle*

d. John ***is*** *visiting* his mother. **present** + *present participle*

e. John ***was*** *visiting* his mother. **past** + *present participle*

f. John ***has/d*** *been* *visiting* his mother.
pres/past + *past participle* + present participle

(7) a. John **will** visit his mother. **future**

b. John **will** *have visited* his mother. **future** + *perfect*

c. John **will** *be visiting* his mother. **future** + *progressive*

d. John **will** *have* *been visiting* his mother.
future + *perfect* + progressive

However, the distinction between aspect and tense auxiliary is not always as clear when considering a cross-linguistic perspective. Recall example (8). In German, especially in southern German dialects, most uses of the past tense have been replaced by present perfect constructions (Musan, 2002). Thus, an alternative (and more contemporary) way of expressing the statement in (3b) would be:

[7] Abusch (1998) argues that future temporal reference is still marked inflectionally in English and German. However, rather than applying to an aspectual auxiliary, future tense applies to a (presumably) modal auxiliary. I do not discuss this distinction any further in the scope of this thesis.

(8) Das Volk hat den König getötet.
the people have-PRES the king kill.PSTPTCP
'The people killed the king' *German*

Although the present perfect can be used to express simple past time reference in German, it nonetheless retains its aspectual interpretation in specific contexts. Thus, it can be argued that the German perfect is ambiguous between a tense-like interpretation and an aspectual interpretation (Löbner, 2002). This becomes apparent in examples where the present tense contributes a future-oriented interpretation. In (9), the present tense and the perfect construction make independent contributions to the interpretation of the clause. This stands in contrast to example (8), where tense and aspect together elicit a past temporal reference interpretation. The temporal contribution of perfect aspect is discussed in section 3.2. Further aspectual properties are discussed in chapter 4.

(9) Future perfect

a. Ruf mich, sobald der Computer **abgestürzt** **ist**.
call me as soon as the computer **crash.PSTPTCP** **is**
'Call me as soon as the computer will have crashed.'
(Löbner, 2002)
German

Similar effects can be observed more transparently in the Romance language family (see, e.g., Abeillé and Godard 2002; Giorgi and Pianesi 1997). In these languages, there often exist multiple past tense forms, some of which are periphrastic tense forms as shown in (10) for French.

(10) Les gens ont tué le roi
The people have.3.PL kill.PSTPTCP the king
'The people killed the king.' *French*

There, the respective construction, called *passé composé* is also ambiguous between an aspectual and a temporal interpretation.[8] As for German, this

[8]The emergence of periphrastic (past) tense forms can, at least in some cases, be

can be shown by combining the *passé composé* with future temporal reference. In the case of French, this is provided by an inflectional future tense form which applies to the auxiliary. Thus, temporal reference is provided by the inflectional marking on the auxiliary and the aspectual meaning is contributed by the combination of an auxiliary and a past participle. The French inflectional future tense and the corresponding future perfect are shown in (11a) and (11a) respectively.

(11) a. Le conducteur tourne-ra le levier
The conductor turn-FUT the lever
'The conductor will turn the lever'

b. Le conducteur au-ra tourn-é le levier
The conductor have-FUT turn-PAST the lever
'The driver will have turned the lever' French
(Butt et al., 1996)

Generally, these examples show that inflectional tense forms apply to the main verb. However, in the case of aspectual modification, the tense marking moves to the auxiliary, or even out of the scope of the aspectual auxiliaries (as shown for the English future tense form *will + V*).

3.1.3 Non-obligatory tenses

Aside from tense systems that form a paradigm consisting of both periphrastic and inflectional tenses, non-obligatory tense languages exist. There, tense is not required to form a full clause. One such example is Urdu, where the past tense can either be expressed explicitly in terms of a periphrastic construction with the past form of the auxiliary verb *be* (similarly for present and future tense) or by the meaning of anteriority that is (non-obligatorily) carried by the perfective marker (see (12a); temporal reference of the perfective

attributed to a difference in their diachronic development (see Squartini and Bertinetto (2000) for an overview). There, the distinction between perfect, perfective, and past tense plays a crucial role (Condoravdi and Deo, 2008). The German language also may be in the process of developing from perfect to past tense, considering it is frequently used interchangeably with the simple past in many contexts (Musan, 2002).

is further discussed in section 3.1.6.1). Another example is given in (12b). There the perfective marker occurs in a simple main verb rather than a complex predicate, showing that the perfective marker is a proper inflectional category.

(12) a. logõ=ne bɑd̪ʃɑh=ko mɑr ɖɑlɑ.
people.PL=ERG king.M.SG=ACC kill put.PAST.M.SG
'People killed the king.'

b. əndʒʊm ɑdʒ. sʊbɑh skul ʧəl-i.
Anjum.F.SG.NOM today morning school.SG.LOC walk-PRV.F.SG
'Anjum walked/went to school this morning.'

(Butt and Rizvi, 2010)
Urdu

(13) a. əndʒʊm=ne əd̪nɑn=ko d̪ekʰɑ
Anjum.F.SG=ERG Adnan.M.SG=ACC see-PRV.M.SG
t̪ʰɑ.
be.PAST.M.SG
'Anjum had seen Adnan'

b. bɪʧʰʧʰu pɑr kərne=kɑ t̪əriqɑ ɖʰunɖ.
scorpion.SG.NOM other-side do.INF.OBL=GEN method search
rəhɑ t̪ʰɑ dʒəb ʊs=ne əpne pɪʧʰe ek
PROG be.PAST.M.SG when 3.SG=ERG Gen.REFLX behind one
mẽɖək=ko d̪ekʰɑ.
frog.SG.ACC see.PRV
'The scorpion was looking for a way to cross, when he noticed a frog behind him.'

Urdu

Similarly to German, the auxiliary construction can be used to express perfect readings, as shown in (13a) from Butt and Rizvi (2010). Furthermore, as shown in the first clause of (13b), the past auxiliary is used to modify aspectual constructions such as the progressive temporally. In general, there is a much richer variation concerning the morphosyntax of the modified (complex) verb. For example, the auxiliary can combine with a verbal noun to express a recent event. This is illustrated in Butt and Rizvi's (2010) example, taken from Glassman (1983).

(14) dʒəb ḍɔkṭər sɑhəb bol-ne=ko t̪ʰe
when doctor sahab.M.NOM speak-INF.OBL=ACC be.PAST-M.PL
t̪o səb lɔg tʃʊp ho gə-je.
though all people.NOM quiet become go-PRV.M.PL
'When the doctor was about to speak, everybody fell quiet.'
(Glassman, 1983, p. 233)
Urdu

The combinatory possibilities of tense auxiliaries and different types of verbs, e.g., verbal nouns, complex predicates, adjectival predications, etc., are, to my knowledge, not fully explored yet; thus, close examination of the Urdu tense auxiliary provides an avenue for further research.
As mentioned above, clauses can be formed without a proper tense in Urdu. As Butt and Rizvi (2010) describe, the tense paradigm is only applied to auxiliaries, however, in terms of tense/aspect marking, Urdu allows for the formation of clauses that are only marked for grammatical aspect as shown in the example in (15); see section 3.1.6.1 for more details.

(15) mæ̃=ne kʰɑnɑ xərid̪ɑ.
I.SG=ERG food.M.SG.NOM buy.PRV.M.SG
'I bought food.' *Urdu*

According to Butt and Rizvi (2010), the perfective morphology is the default past tense marker in Urdu.[9] If tensed auxiliaries in Urdu are optional, this begs the question of whether tenses can be optional more generally.
In fact, non-obligatory tense marking represents an important hallmark in the morphosyntactic variation of tense since this type of variation extends to the semantic domain as well, as optional tenses are argued to carry

[9] Butt and Rizvi (2010) do not distinguish between perfect and perfective aspect. I use the label perfective morphology here since it forms an opposition with imperfect(ive) aspect morphology (on main verbs). Furthermore, as mentioned above, they also associate perfect semantics with constructions involving a tensed auxiliary. This topic is further discussed in section 3.1.6.1.

implications that obligatory tense markers lack (Plungian and van der Auwera, 2006). A case study on these implications is presented in Cable (2015). I will explain the insights gained from this study briefly through examples (16) to (18b) taken from Cable (2015) (and originally based on Leer 1991).

(16) K̲uwak'ei.
IMPV.good.weather
'The weather is/was nice.'
(Leer, 1991, p. 464)

(17) K̲uwak'eiy**een**.
IMPV.good.weater.**DEC**
'The weather was nice (but turned bad)'

(18) a. I tláa áwé x̲washáa**yin**.
your mother FOC 3SG.OBJ.**PFV**.1.SG.SUBJ.marry.**DEC**
'I married your mother (but we're not married any more)'
(Leer, 1991, p. 468)

b. Du x̲'eis áwé weit'át x̲walawaas**ín**.
his mouth.for FOC that.thing 3.OBJ.**PFV**.1.SG.SUBJ.roast.**DEC**
'I roasted that for him (but he didn't want to eat it).'
(Leer, 1991, p. 468)
Tlingit

As in Urdu, the aspectual opposition between perfective and imperfective aspectual morphology is sufficient to form a full clause in Tlingit. By adding the decessive (DEC) marker, the content of the respective linguistic expression is located in the past (compare (16) and (17)) and furthermore receives either a *cessation implication* when combined with the imperfective or a *canceled result* (18a) implication or an *unexpected result* implication (18b), when combined with perfective aspect. Thus, generally, the type of implication that an expression receives is based on its grammatical aspect marking. This is in line with observations made by Plungian and van der Auwera (2006). Ultimately, Cable (2015) shows that these implications are cancellable, and the grammaticalized contribution of the decessive marker is that of past time reference. Nonetheless, it is important to keep the semantic differences between obligatory and optional tenses in mind.

3.1.4 Relative tenses

This section presents the phenomenon of *relative tense*. Relative tenses are recognized by the trait that they may pick up a time other than the speech time as a temporal anchor. This definition invites the use of the term relative tenses for different kinds of tense/aspect forms.

For example, the typical examples discussed in the previous sections receive relative interpretations in some contexts, namely, when embedded under complements. This is exemplified in the Japanese example in (19). This utterance expresses that Mary was sick at some point before Peter reported that situation.

(19)

pītā	wa	mearī	ga	byōki-da	to	it-ta.
Peter	TOPIC	Mary	S	sick-PAST	COMP	say-PAST

‘Peter said that Mary was sick’ *Japanese*

However, at this point, I want to focus on relative tenses that apply on top of regular tense markers.[10] A straightforward example of this is provided by two constructions that express future temporal reference in English: The *Aux + going to + infinitive* construction and the auxiliary *would*.

The former combines with an additional tense similar to the perfect construction as shown in (20). On the other hand, *would* may pick up some sort of antecedent as is shown in (21). This is indicated in terms of the co-indexation. Abusch (1998) proposes to decompose *would* into a modal and a temporal component, where the temporal component is that of past temporal reference. This means the general semantics of the two examples presented below is that of a future temporal reference relative to some point in the past provided by a past-tense-like component.

(20) a. John **was** going to visit his aunt. [PAST [FUT]]

(21) John said$_i$ that he would visit his aunt. [FUT$_i$]

Indonesian provides a more extreme example of a relative tense. Arka (2013) establishes only one proper temporal marker in Indonesian, namely, *akan*,

[10] Relative tenses in the sense demonstrated in (19) are discussed in more detail in 3.2.

a future tense marker. This future tense marker can be evaluated both relative to the speech time but also relative to some past time, i.e., it can be understood as a combination of *non-future* + *future* temporal reference.[11,12]

(22) Peter **akan menulis** se.buah surat
Peter FUT **AV.write** one.CL letter
'Peter will write a letter.'
'Peter was going to write a letter.' *Indonesian*

According to Comrie (1985), English *would* and Indonesian *akan* apply to some contextually provided reference point, whereas examples such as the English *AUX + going to + V* and the relative tense in Japanese are not cases of relative temporal reference in the same sense. For this purpose, I distinguish between compositional relative temporal reference and contextual relative temporal reference. This means we have seen two types of relative tenses. Those that combine with an openly expressed tense to add a secondary temporal orientation (e.g., the *Aux + going to + Inf*), or those that combine with an apparently invisible element to provide a relative temporal orientation (e.g., *would* in English).[13]

[11] Indonesian aspectual markers like the perfect marker *telah* behave similarly. Thus, example (1) can express both present perfect and past perfect based on the context.

(1) Saya telah membeli makanan
1.SG PERF AV.buy food
'I have bought food'/'I had bought food' *Indonesian*

[12] Only the future temporal reference is openly expressed. The origin of the non-future temporal reference is explored in more detail in the next section (3.1.5).

[13] This section leaves a few open questions. Firstly, it seems like aspectual tenses are generally instances of compositional temporal reference. However, as we have just seen, this is not the case for Indonesian aspectual markers (see (1)). Cross-linguistically, the picture is obscured further when we also consider non-finite tense/aspect forms (Comrie, 1985). Secondly, if we accept *akan* as a proper tense, then we would have to explain how the contextual shift of the evaluation time works. It certainly is an instance of temporal reference since it situates the evaluation time, but so far we have only discussed tenses that either shift the reference time or the eventuality time. An in depth analysis of the Indonesian tense/aspect markers goes beyond the scope of this thesis but they provide an important avenue for further research. For the remainder of this thesis, I treat *akan* as an aspectual tense operator. Thus, I assume tense to be a purely contextual in Indonesian (see section 3.1.5). The backshifted reading then results from a contextually provided past

3.1.5 Zero-marked tenses

In the previous section, I have already illustrated the existence of anchors for temporal reference that are not expressed in the linguistic surface structure. Consider these additional examples from Indonesian, where the same verbal form occurs in combination with different temporal adverbials, each exclusive to one of the three basic tenses we considered at the beginning of this section.

(23) a. Dia datang (besok).
3.SG come tomorrow
'She will come (tomorrow).'

b. Dia datang (kemarin).
3.SG come yesterday
'He came in (yesterday).'

c. Dia datang (sekarang).
3.SG come now
'She is coming (now).' *Indonesian*

These examples show that temporal reference is generally provided by the context rather than by some grammaticalized device. Nonetheless, the discussion of the future marker has shown that Indonesian conceptualizes different times and relations between them from a meaning perspective but does not, or rather, only optionally conceptualize them in the grammar.
It is also possible that only part of a tense paradigm is zero-marked. This is, for example, claimed for Gitksan, a Tsimshianic language spoken in northwestern British Columbia, Canada (Moseley, 2008, 65f) by Matthewson (2012). Example (24a) illustrates that sentences with no overt tense marking are compatible with either past time or present time reference (due to the lack of morphological marking, a so-called null tense). On the other hand, there is an obligatory future marker *dim* that is required in future contexts, as illustrated in (24b) – dropping the future marker results in infelicity when referring to a future point in time (Matthewson, 2012).

temporal reference combined with aspectual future temporal reference.

(24) a. siipxw=t James (k'yoots)
sick=DM James (yesterday) /
'James was sick (yesterday)'
'James is sick.'

b. *(dim) siipxw=t James t'aahlakw
*(FUT) sick=DM James tomorrow
'James will be sick tomorrow.' *Gitksan*

This means that the zero-marked non-future is identified in terms of contrast, other than in Indonesian, where the future marker is an optional alternative to the unmarked variant. This further shows that temporal reference is encoded on a spectrum from morphosyntactic to pragmatic mechanisms.

3.1.5.1 A quick note on temporal remoteness

Something I have not mentioned before is that various languages allow for even more fine-grained distinctions of temporal reference, usually subsumed under the label of temporal remoteness (Bochnak and Klecha, 2015; Klecha and Bochnak, 2016), or graded tenses (Mucha, 2015). In these languages, the temporal distance between the time of evaluation and the reference time plays a crucial role and can affect grammatical marking, usually allowing for a handful of distinctions (e.g., immediate, near, and remote). Although temporal remoteness may also be associated with categories such as present and past perfect (Fleischman, 1989), a detailed examination is left for future work. This is mainly because temporal remoteness is often highly context-dependent (Bochnak and Klecha, 2015).

3.1.6 Temporal reference beyond tense

Research on categories not traditionally analyzed as tenses further supports the linguistic spectrum from form to meaning on which temporal reference can be encoded cross-linguistically.

3.1.6.1 Temporal reference via grammatical aspect

For example, Smith (2006) proposes that grammatical and lexical aspect encode different default temporal interpretations of untensed clauses. A popular example of this type of language is Chinese, where a perfective interpretation, i.e., an interpretation where an event is described as completed, results in a (contextually or explicitly) defeasible past time reference. On the other hand, imperfective and stative descriptions receive default present time reference but can also be modified to refer to past or future situations.

(25) Ta dapuo yi-ge huaping.
he break one-CL flower vase
'He broke a flower vase.' (Lin 2003, p. 262)

(26) Ta hen congming.
he very clever
'He is very clever.' (Lin 2003, p. 263)
Chinese

I already mentioned the tendency of perfective aspect to express pastness or, more rarely, future temporal reference. To elaborate on this point, let us look at Urdu again. In this language, perfective morphology can occur both in past and future tense constructions. The perfective marker by itself is generally used to invoke a past tense interpretation. However, when combined with a temporal modifier indicating present temporal reference, as shown in (27), then the marked eventuality is shifted to the (immediate) future. Combining perfective morphology with grammaticalized present temporal reference can achieve the same effect. Example (28) illustrates this for Polish.

(27) mæ̃. əb^{h}i ɑ-i.
I.NOM now came-PRV.F.SG
'I'll be right there (come right away).' (Butt and Rizvi, 2010)
Urdu

(28) Peter **na-pis-ze** list
Peter **PRV-write-PRES** a letter
'Peter is going to write a letter (completely)' *Polish*

While perfective aspect usually associates with non-present temporal reference, imperfective aspect is predominantly associated with a present tense interpretation. This is observable in language with and without overt tense marking. For example, many romance languages have a dedicated past imperfective form but not a dedicated past perfective form.

(29) John construis-ait une maison.
John construct-PAST.IMPV a house
'John was building a house.'

In the same vein, the English progressive is compatible with present, past, and future tense markers. However, the simple tense forms in English illustrate a clear congruence with either perfective or imperfective aspect: the simple past is assumed to have a default perfective interpretation as predicted by the observations we made so far (Aloni and Dekker, 2016). On the other hand, the simple present in English and French (among other languages) is used primarily to describe habituality, i.e., regularly occurring habits, which is assumed to be related to imperfectivity (discussed in more detail in Chapter 4). This is in accordance with the general picture depicting the relation tense and grammatical aspect.

As the last example, I draw from Błaszczak and Klimek-Jankowska's (2016) examples of Polish. They discuss the combination of perfective vs. imperfective with non-past tense. In line with the pattern worked out above, they point out that non-past imperfective verbs receive a present tense interpretation (see (30)), while non-past perfective verbs receive a future tense interpretation as shown in (28). This contrasts their respective combinations with past tense, where aspect does not influence the temporal interpretation.

(30) Jan pisze artykuł
John write.IMPV.NON-PST article
'John is writing an article.'

The examples above allow us to draw some conclusions about the interaction between grammatical aspect and temporal reference. These are summarized in Table 3.1.

	perfective	imperfective
present	✗	✓
past	✓	✓
future	✓	✓

Table 3.1: Temporal reference and grammatical aspect

As a result of the generalization in table 3.1, we will pay close attention to grammatical aspect as a potential restrictor on temporal reference. Perfective aspect, in particular, seems to imply some strong constraints on temporal reference. In comparison, imperfective aspect seems to combine more freely with temporal reference. However, its affinity with present tense, as opposed to the perfective aspect, makes it potentially just as relevant to interpreting temporal information in languages without overt tense marking.

3.1.6.2 Temporal reference and evidentials

Another example of inferential meaning may be provided by evidentials. Evidential expressions are expressions that encode different relations between the speaker and the knowledge the speaker wants to convey. For example, evidential markers might indicate that the speaker has direct (e.g., visual or auditory) evidence for his utterance.

(31) a. Juse irida di-manika-ka
José football 3SG.NF-play-REC.PAST.VIS
'José has played football (we saw it)'

b. Juse irida di-manika-mahka
José football 3SG.NF-play-REC.PAST.NONVIS
'José has played football (we heard it)' (Aikhenvald, 2004, p. 2)
Tariana

According to Aikhenvald (2004) evidentials pattern with present or past interpretations but rarely if never with future temporal reference. In other words, evidential semantics may interact compositionally with the future tense, e.g., to produce hypothetical modality in the case of a non-first-hand evidential. I.e., this indicates that the speaker only has indirect evidence. as illustrated in the example below. Interestingly, the meaning associated with the future tense marking seems to be canceled.

(32) a. a:che chuge-ge jo:dude-t ejrie-l'el-te-j
deer track-LOC turn-IMPV walk-NF-FUT-INTR:3.SG
'He is probably walking along deer tracks'

b. locil-ŋin lebie-d emej-ŋin tadi:-nu-**l'el**-te-m
fire-DAT soil-AT mother-DAT give-IMPF-NF-FUT-TR:3.SG
'Probably, he used to give it to the Fire, to Mother of the Earth'

(Aikhenvald, 2004)
Yukaghir

Generally speaking, the association of evidentials with non-future contexts makes sense since we usually do not have evidence for future occurrences. The exact interaction of tense and evidentiality is highly language-specific. Still, it should be noted that even categories beyond tense and aspect – which traditionally affect temporal interpretation – affect the interpretation of verbal predicates in terms of their temporal semantics.
Overall, the variation of the category of tense is rich. Particular instantiations of temporal reference exist across and sometimes within languages. Linguistic material might carry temporal reference as a by-product of its conventionalized meaning leading to the emergence of new morphosyntactic constructions that express temporal meaning. To better understand the category of temporal reference and its variation, we require several parameters that restrict the use of a given category on a morphosyntactic level, a semantic level (and a pragmatic level). This is the topic of section 3.2.

3.1.7 Interim summary: encoding temporal reference

Up until this point, I have illustrated various ways of encoding temporal reference cross-linguistically. I do not claim that I have provided a fully comprehensive picture. For this reason, I use this section to summarize the main points that I assume to be important for annotating temporal reference and mention some loose ends.

Temporal reference can be expressed explicitly; the corresponding category is usually called tense. The main distinctions are either between past and non-past, future and non-future, or between past, present, and future tense. Temporal reference has at least two layers. The first layer relates the time of evaluation with some point reference. The second layer connects the point of reference with the time at which an eventuality occurs (the eventuality time). Certain aspectual categories, e.g., perfect and prospective aspect, are instantiations of the latter but can also be interpreted as constraints on the relation between the reference time and evaluation time. For example, the German present perfect is ambiguous between the two readings.
Tenses are not obligatory cross-linguistically. Thus, temporal reference can be expressed by various kinds of linguistic encoding, e.g., grammatical aspect, evidentiality, or simply by contrast (i.e., tenses can be zero-marked). These different forms may entail or imply temporal reference. In fact, languages may rely on a pragmatic system for inferring temporal reference.
Bohnemeyer (1998) argues that temporal ordering is not a semantic primitive. Rather, temporal reference can be deduced from phasal markers referring to different stages of an eventuality, such as pre- and post-state. However, while I agree that temporal ordering must not be grammaticalized, I believe that temporal ordering is an essential component of human reasoning and, thus, temporal reference is, at the very least, a relatively basic semantic feature that needs to be captured to model natural language reasoning.

3.2 Tense/aspect syntax and semantics formally

In the previous section, I have introduced varying ways of expressing temporal reference. Based on Dahl (1985), we can distinguish between inflectional, periphrastic, derivational, and zero-marked markers for temporal reference. Furthermore, we have established that these markers can be either optional or obligatory.

form	inflectional	periphrastic	derivational	zero-marked
optional(+/-)				
meaning	anteriority	simultaenity	posteriority	unspecified
relative(+/-)				

Table 3.2: Temporal reference: basic feature distinctions

3.2.1 Distinguishing syntactic and semantic features

The main point of the ParTMA annotation scheme is to annotate form and meaning more modularly. Thus, in this section, I describe a simple classification for markers of temporal reference that makes this distinction explicit. I will focus on exploring the basic meanings associated with temporal reference: past (anteriority), present (simultaneity) and future (posteriority). Overall, the interpretation of tense markers cannot be distributed among the three basic tenses, but it seems that we require at least the existence of non-future and non-past in the paradigm. The last logical alternative is non-present. However, non-present readings seem to be rarely, if ever, available. This is not necessarily surprising. As I have pointed out, there is an apparent asymmetry between past and future temporal reference. Thus, any given form compatible with either but not with the present might lean more towards one of the two alternatives depending on its precise meaning. This is also reflected in the fact that future temporal reference without morphosyntactic marking requires, at the very least, a solid contextual restriction and often either the use of future auxiliaries or lexical elements.

Table 3.2 describes the basic descriptive feature distinctions worked out in this section for instances of tense/aspect categories. An instance of a tense aspect category refers to the occurrence of that category in a particular morphosyntactic context (e.g., matrix clauses) where it behaves predictably. At the level of morphosyntax (more generally surface form), we have seen distinctions between inflectional, periphrastic, and zero-marked tenses. We have further proposed that certain other categories contribute temporal reference, which may occur in terms of derivational morphology, thus, covering the basic spectrum of variation proposed by Dahl (1985). An instance can be either optional or obligatory. In Germanic languages, aspectual markers are usually optional. However, I also have presented some examples of optional tenses from Plungian and van der Auwera (2006) and Cable (2015).

On the level of semantics or meaning, we have distinguished between *anteriority*, *simultaenity*, and *posteriority*. As I have shown, in some cases, adjacent feature combinations can occur, e.g., non-past would be encoded as + *simultaenity* and + *posteriority*. The value *unspecified* corresponds to a combination of all three features but boils down to there being no restriction at the level of semantics on the respective instance. Based on the discussion in section 3.1.4, I have also highlighted a distinction between compositionally relative tenses and contextually relative tenses. This distinction is not highlighted here but will play a role in section 3.2.

Consider the example of the German complex tense/aspect form *present perfect* again. In a (non-subordinate) main clause, I have proposed an ambiguity between a past tense meaning and a perfect meaning, following Ehrich (1987), Löbner (2002), and Musan (2002), among others. As explained in section 3.1.2, the present perfect proper can be decomposed into two components, namely a temporal and an aspectual component, both of which contribute temporal reference (although the perfect provides some additional semantic features that are discussed in chapter 4). Here, the only relevant feature of the perfect is that it is an instance of relative anterior temporal reference. It is relative to the present tense component. This fact is not encoded in the basic table. Thus, the syntactic and semantic features of these two components are given in table 3.3 and 3.4 respectively.

form	inflectional	periphrastic	derivational	zero-marked
optional(+/-)		+		
meaning	anteriority	simultaenity	posteriority	unspecified
relative(+/-)	+			

Table 3.3: Basic features: the *perfect*

The past tense interpretation of the German present perfect combines these two categories on the syntactic level but is semantically different in that it generally conveys a non-relative anterior temporal reference making it similar

form	inflectional	periphrastic	derivational	zero-marked
optional(+/-)	+			
meaning	anteriority	simultaenity	posteriority	unspecified
relative(+/-)		+	+	

Table 3.4: Basic features: the *present tense*

to the plain past tense in German. Note that every cell that receives a value is obligatory to form the instance of the tense/aspect form. Thus, the *optional* feature refers to the properties of the morphosyntactic components rather than the properties of the complex tense/aspect form.

form	inflectional	periphrastic	derivational	zero-marked
optional(+/-)	-	+		
meaning	anteriority	simultaenity	posteriority	unspecified
relative(+/-)	-/+			

Table 3.5: Basic features: *present perfect as past*

We have seen that past tense markers, although usually evaluated relative to the speech time, can sometimes receive relative interpretations, for example, in complement clauses. This has been concretely shown for Japanese. In the present basic table, when we want to unify two contexts (in this case, matrix clauses and complement clauses), we can use the -/+ feature to express that a given category instance is morphosyntactically relative, i.e., relative in well-defined morphosyntactic contexts.

form	inflectional	periphrastic	derivational	zero-marked
optional(+/-)	+			
meaning	anteriority	simultaenity	posteriority	unspecified
relative(+/-)	-/+			

Table 3.6: Basic features: *Past tense in Japanese*

In comparison, the relative future marker *akan* in Indonesian can be annotated as shown in Table 3.7 since it is both optional and relative to some contextually provided point in time.

form	inflectional	periphrastic	derivational	zero-marked
optional(+/-)		+		
meaning	anteriority	simultaenity	posteriority	unspecified
relative(+/-)			+	

Table 3.7: Basic features: *Future tense in Indonesian*

In summary, the table template given in table 3.2 serves as a guideline for categorizing tense/aspect categories on a descriptive level. This is the first step to providing a fully fleshed-out annotation scheme incorporating syntactic and semantic information. In the next section, I work out the formal underpinnings of the semantics underlying this annotation scheme.

3.2.2 Towards a formalization of temporal reference

In this section, I provide a compositional semantics for tense and aspect that is used as the backbone of the ParTMA annotation scheme. In the previous chapter, I have established three main criteria of this compositional semantics: i) it follows roughly the analysis of tense/aspect forms as proposed by Reichenbach (1947), ii) It follows a particular order of composition given in (33), and iii) it uses lambda calculus to assemble meanings.

(33)

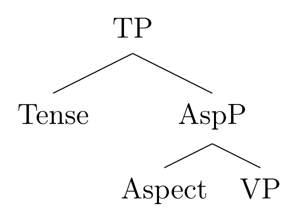

None of these criteria makes the present semantics special compared to the numerous existing formal approaches to tense and aspect. However, by providing a concrete way of composing tense/aspect meanings, we can infer which features are necessary to derive the same semantics via annotation.
Before we can use this structure to assemble meaning, we need to decide how temporal reference is encoded formally in the semantics. As before, we start with the prototypical way of encoding temporal reference: tenses.

3.2.3 Tenses: pronouns or quantifiers?

To begin with, a small overview of the development of tense semantics is in order. However, I will not go back all the way to Prior (1967) in this review (For a comprehensive summary, I refer the reader to Kusumoto (1999)). Rather, I constrain myself to what I refer to as *modern temporal semantics*. This stage of semantic research on temporal reference begins with the observation of Barbara Partee that tenses share several properties with pronouns (Partee, 1973), hence referred to as temporal pronoun hypothesis TPH. A claim that quickly gained traction in the formal semantic literature and has been entertained up to this day (Abusch, 1997; Heim, 1994; Hinrichs, 1986; Kamp, 2017; Kratzer, 1998).[14]

Partee's idea questioned the traditional treatment of tenses in terms of simple (existential) quantifiers. A popular approach pioneered by Prior's (1967) tense logic. The gist of her idea is that linguistic expressions are not just about some arbitrary time as introduced by quantifiers but rather about particular time intervals. The example in (34) elegantly displays the problem in the interaction between a tense operator and negation. Scopal ambiguity results in two readings, neither of which is very appealing. Consider the following paraphrases: there is no time at which I turned off the stove ($\neg\exists$), and there is a time at which I did not turn off the stove ($\exists\neg$). The first paraphrase shows that wide scope of the negation makes far too strong of a prediction for the meaning of the sentence in (34). The second paraphrase, however, makes a prediction that is almost trivially true. In fact, nearly all past times will be times at which I did not turn off the stove. In conclusion, neither of the two readings is what we would intuitively associate with the sentence in (34). Partee's solution is to assume that tenses do not behave like sentential operators as illustrated below in (35a) but instead appear as variables in the object language.

[14] I cite Kamp (2017) in place of several related papers often cited in the literature. The present paper clarifies and summarizes many of Kamp's observations and provides a modern perspective on the impact of his work on the development of DRT and with respect to temporal semantics. Thus, it provides a suitable starting point for tense semantics in DRT and a good overview for the present purposes.

(34) I didn't turn off the stove. (Partee, 1973, p. 602)

The original goal of Partee (1973) was to argue for a representation of temporal variables in the object language. Thus, time variables would behave in parallel to variables over entities, allowing Partee to explain the similarities between nominal and temporal anaphoric reference (Partee, 1984). I will not argue in particular whether tense should be represented in the extension of a lexical entry or whether it should be treated as an operator, where the meaning of tense is expressed in the meta language since the claim that tenses behave similarly to pronouns does not necessarily warrant one approach over the other (Kusumoto, 1999; Partee, 1984). In (35), I illustrate both approaches of representing past tense as an existential quantifier over times. (35a) represents the traditional Priorian tense operator, while (35b) illustrates an explicit quantification over tenses in the object language.
The consequence of Partee's original approach (Partee, 1973) is to put temporal variables in the denotation of verbal predicates as is illustrated in (36). Correspondingly, we can treat temporal reference as an existential quantifier, the only difference being that it is explicitly stated in the object language.

(35) a. $[\![PAST\phi]\!]^{g,t,w} = 1$, iff there is a time t' such that $t' \prec t$ and $[\![\phi]\!]^{g,t',w} = 1$.

b. $[\![PAST]\!]^{g,w} = \lambda P \in D_{<i,t>}[\lambda t \in D_i[$ there is a time t' such that $t' \prec t$ and that $P(t') = 1$ $]$ (Kusumoto, 2005)

(36) **Example:** John slept.

a. $[\![$ *John sleep* $]\!]^{w,g} = \lambda t \in D_i[$ John sleeps at t $]$

b. $[\![$ *PAST(John sleep)* $]\!]^{w,g} = \lambda t \in D_i[$ there is a time t' such that $t' \prec t$ and that John sleeps at $t' = 1$ $]$

Even if we accept the existence of temporal variables in the object language, there is still the question as to how exactly tense should be formalized. Example (36) obviously doesn't incorporate Partee's TPH. An implementation of her approach needs to capture that a time interval is an element of the domain that is referred to using a temporal pronoun introduced by an instance of temporal reference (e.g., a tense marker). This can be formalized as in

(37), where the assignment function assigns the temporal variable. However, we then require a way to differentiate between the different orientations of temporal reference (past, present, future, etc.). The traditional way to solve this is to introduce a presupposition restricting a given pronominal tense to a past time (Heim, 1994; Kratzer, 1998).

(37) $[\![PAST]\!]^{g,w} = g(i)$, $g(i)$ is only defined iff $g(i) \prec t_0$

This approach allows us to produce a valid semantic representation for the stove example with the paraphrase: it is not the case that I forgot to turn off the stove at t, where t is a specific past time provided by the assignment function. If the assignment function, i.e., the context, does not provide an appropriate value for t, i.e., some time in the past, the meaning of the utterance is undefined.

3.2.4 Composing temporal reference

Either of the two approaches entertained above is compatible with Reichenbach's system. However, how do they work in the compositional framework proposed in this thesis? Generally, I assume that the sister node of the tense operator provides a property of times (type $< i, t >$) that is compatible with the temporal operator given in (35b) but would also be compatible with a framework in which tense is treated as a simple temporal variable that is defined by the assignment function.

3.2.4.1 Temporal reference via tense

Kusumoto (1999, 2005) proposes such a semantics that combines the two approaches presented in the previous section. Tenses are constructed of two elements: a temporal variable as specified in (37), and a temporal quantifier that locally accommodates the presuppositional content (Romero, 2014).[15] In this framework, Kusumoto (2005) proposes the following denotation for

[15] Kusumoto's (2005) original approach argues that the temporal variable does not carry a temporal reference presupposition by itself and that the temporal quantifier properly asserts the temporal restriction.

verbs and nouns Importantly, both assume a temporal variable.

(38) $[\![\textit{ sleep }]\!]^{w,g} = [\lambda x \in D_e.[\lambda t \in D_i.\text{ x sleeps at t }]]$

(39) $[\![\textit{ man }]\!]^{w,g} = [\lambda x \in D_e[\lambda t \in D_i.\text{ x is a man at t }]]$

Based on the explanation above, the corresponding order of composition is presented in (40). The tense variable saturates the time slot of the verb. In Kusumoto (2005) proposal, the temporal variable needs to be licensed by a temporal quantifier, making the quantifier obligatory in her semantics.[16]

(40)

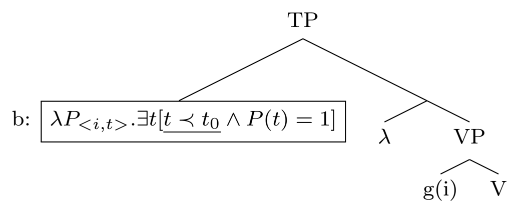

This is the general semantic composition that I will take as a basis for temporal reference (in the sense that there is no anaphoric component). However, there is still no contextual restriction on the reference time t since the anaphoric element is abstracted over to be compatible with the quantificational tense. To remedy this issue, I follow Kusumoto (2005) and subsequent work in introducing a function C that restricts the available reference times to those that are contextually salient (based on Von Fintel (1994)).

(41) $[\![PAST]\!]^{g,w} = \lambda C_{<i,t>}.\lambda P_{<i,t>}.\exists t'[C(t') \wedge t' \prec t_0 \wedge P(t) = 1]$

The lexical entry in (41) could be properly used for absolute tenses. However, as I have shown, there are several arguments for the existence of relative tenses as they, for example, appear in complement sentences. This issue is quickly solved by introducing a lambda binder for t_0 and replacing it with an arbitrary variable:

[16]The structure, thus, bares resemblance to a structure where an NP quantifier is moved to a higher position, indicating a wide scope reading and leaving behind a trace to locally saturate the verbal predicate (Kratzer and Heim, 1998). In other words, the lower variable can be intuitively understood as a trace of the temporal quantifier. Thus, it is more of a technical trick to get the scope of the quantifier right rather than a necessarily real contribution of tense. This idea becomes relevant when computationally implementing the present framework in chapter 5.

(42) $[\![PAST]\!]^{g,w} = \lambda C_{<i,t>}.\lambda t_i.\lambda P_{<i,t>}.\exists t'[C(t') \wedge t' \prec t \wedge P(t') = 1]$

In Kusumoto's (2005) work, absolute tenses introduce a temporal variable referring to the evaluation time that she calls $s*$. Thus, simple past tense sentences will be evaluated in the following way.[17]

(43) John visited his mother.

a. $\exists t'[t' \prec t^* \wedge visit(j, m, t') = 1]]$

Relative tenses, particularly compositionally relative tenses, do not introduce such a separate evaluation time but rather have their evaluation time specified compositionally. This is most easily understood by looking at how complex tense/aspect forms are analyzed in this framework. This is illustrated in (44a). The evaluation time of the lower temporal quantifier is saturated by the higher temporal quantifier, thus establishing a relative evaluation.

(44) John had visited his mother.

a. $\exists t'[t' \prec t^* \wedge \exists t''[t'' \prec t' \wedge visit(j, m, t'') = 1]]$

There exist different contexts for compositionally relative tenses, as we will see in section 3.3. However, it is also necessary to distinguish these from absolute tenses in embedded contexts, as is, for example, the case with tensed relative clauses as in:

(45) Hillary married a man who <u>became</u> the president of t he U.S.
(Kusumoto, 2005)

Kusumoto (1999, 2005) argues at length that the underlined tensed predicate in the relative clause needs to be interpreted temporally independent from the matrix tense. This means the temporal relation between the predicate *married* in the main clause and *became* remains unspecified (the man may have become president before marrying her, or only after he became Hillary's

[17] Although I have introduced the contextual restriction, I will generally omit it in the following examples to keep the semantic representations simpler.

husband).[18] This independence can be achieved by introducing two different absolute tenses, i.e., each tense introduces its own evaluation time. Conversely, a variable for the evaluation time is not provided at a sentential level but by absolute tenses.

Furthermore, to get the semantics of the relative clause right, the NP quantifier needs to scope outside the tensed relative clause but inside the matrix tense. This has two reasons: i) to ascertain co-binding between the entity in the relative clause and the entity in the main clause, ii) to not violate movement constraints of quantifiers (she illustrates this by embedding the relative clause under an NPI; for more details see Kusumoto (2005)). Concretely, the following semantics are desired for the sentence in (45):

(46) $[\exists t[t \prec t^* \wedge \exists x[man(x,t) \wedge$
$\exists t'[t' \prec t^* \wedge \textit{become-president(x,t')}] \wedge marry(h,x,t)]]]$

How is this meaning achieved? In the present proposal, nominal predicates like verbs, are of type $< e, it >$, and, thus, quantifiers are of type $<< e, it >, < e, it >, it >$. It is, thus, simple to derive the semantics of the main clause. However, the relative clause provides a small challenge. Typically, the semantics of relative clauses are derived via predicate modification (Kratzer and Heim, 1998). Due to the condition that the relative clause needs to be fully tensed, the corresponding semantic element ($[\![$ *who*$_i$ *became president* $]\!]^{g,w}$) is of type $< e, t >$, which needs to be combined with *man* of type $< e, it >$ via predicate modification, which is impossible due to the type clash. For this purpose, Kusumoto (2005) introduces a specific rule for combining a tensed relative clause with its NP head. Since this is not of immediate relevance to the topics discussed in this chapter, I will leave the discussion at that here. However, I will pick it up in chapter 5 when I propose a concrete implementation for the semantics presented here.

In summary, the present system consists of, at most, two levels of temporal reference. One that specifies the temporal relation between speech/evaluation

[18]There is also a potential reading, where the man was in the process of becoming president during their wedding. However, this reading is based on the lexical and aspectual semantics of *(to) become* and is thus ignored for reasons of simplicity.

time and reference time, and one level that specifies the relation between the reference time and eventuality time. The latter is always compositionally relative to the former. Relative temporal reference is distinguished from absolute temporal reference in that they introduce the speech time as an additional element. Finally, in main clauses, NP quantifiers scope between TP and VP, but due to the specific semantics of relative clauses, there, the scope order is reversed. Example (47) reiterates the proposed semantics for temporal reference, exemplified by past temporal reference.

(47) a. **Absolute temporal reference:**
$[\![PAST]\!]^{g,w} = \lambda C_{<i,t>}.\lambda t_i.\lambda P_{<i,t>}.\exists t'[C(t') \wedge t' \prec t \wedge P(t') = 1]$
$[\![now]\!]^{g,w} = t^*$

b. **Relative temporal reference:**
$[\![PAST]\!]^{g,w} = \lambda C_{<i,t>}.\lambda t_i.\lambda P_{<i,t>}.\exists t'[C(t') \wedge t' \prec t \wedge P(t') = 1]$

3.2.4.2 Present and Future tense

In this section, I want to touch briefly on present and future temporal reference. First, I will touch on present temporal reference. A quantificational approach does not seem to be necessary. The most straightforward analysis is to introduce a temporal variable that denotes the *now* and insert it into the semantic computation at the same place as the temporal quantifier proposed above. Alternatively to the actual now, there is also the possibility that some tenses express a *extended now* (Dowty, 1982), which technically can be described as a case of `non-future`. This notion has mainly been used to explain the semantics of the perfect, which do not allow for temporal adverbials that do not include the *now* (the so-called present perfect puzzle (Klein, 1992)). Within the formal semantic literature, this is a fairly standard assumption (Aloni and Dekker, 2016). It is, in principle, used disregarding whether the perfect should be an aspectual or a temporal operator. I follow the approach by (Arregi and Klecha, 2015; Klecha, 2016), which treats the perfect as a temporal operator. I bring this up since the same question can be put forward for future tenses and prospective aspect, as I have discussed in Zymla (2017). It is difficult to make a decisive judgment at the level of sentential semantics, to say the least. For the remainder of this chapter, I treat the perfect and prospective aspects as tenses, as illustrated by virtue of the past perfect in the previous section. Furthermore, I assume that there exist not only past, present, and future tenses but also non-past, non-future (and non-present) tenses. This allows me to cover the tenses in terms of the Reichenbachian tradition, which we have established as inspiration for most of the current semantic analyses. Thus, I can circumvent (to a degree) such issues as the present perfect puzzle mentioned before.

I have mentioned before that the future tense frequently comes with an additional modal component or some other type of operator, e.g., a plan operator (Abusch, 1998; Copley, 2009). I won't provide a suitable semantics for all of these different meanings. Rather, I constrain myself to purely temporal semantics for the future.

Consequently, present and future temporal reference deserve closer investigation in future work. For the present purposes, I find it useful to express present time reference as a proper relation between two time intervals.Just as past temporal reference, simple present time reference is then derived by combining the present tense with a specified evaluation time. This is done so that past, present, and future temporal reference behave similarly from a type-logical perspective. The benefit of this will be worked out in more detail in chapter 5.

(48) a. $[\![PRESENT]\!]^{g,w} = \lambda P_{<i,t>}.\lambda t.\exists t'[t' \circ t \wedge P(t) = 1]$

b. $[\![FUTURE]\!]^{g,w} = \lambda C_{<i,t>}.\lambda P_{<i,t>}.\lambda t_i.\exists t'[C(t') \wedge t \prec t' \wedge P(t) = 1]$

3.2.4.3 Zero-marked tenses

In principle, Zero-marked tenses behave the same as the tenses introduced in the previous section. In fact, there are no apparent differences between the two variations since I have not specified a syntax/semantics interface yet. However, I will generally analyze contextually relative tenses such as the Indonesian *akan* and *telah* as instances of temporal reference between the reference time and eventuality time. The higher temporal reference node provides the invisible anchor that specifies the reference time.

An alternative way of dealing with Indonesian would be to let *akan* be an instance of temporal reference that introduces a time $t^* \neq now$, similar to an absolute tense. However, then it would need to be explained what shifts the temporal orientation of that evaluation time. This means, ultimately, we need two different instances of temporal reference in either case. To my knowledge, and as reflected in the Reichenbachian tense system, there exists no language that introduces three separate instances of temporal reference in its grammar.[19]

In any case, the resulting semantics should be sufficiently similar. This is also reflected in the following quote: "temporal reference in tenseless languages under tenseless analyses is just as specific as in tensed analyses of tensed(and tenseless)" languages (Tonhauser, 2015, p. 147). This might change in the future when more fine-grained distinctions of temporal reference need to be considered, but for the sake of this thesis, I subscribe to this statement.

Now that I have explained the basic formal semantic assumptions about tense, I turn to some more complex phenomena discussed in the formal semantic literature. In other words, I illustrate some challenges that any account explaining temporal semantics has to face.

[19] It is not impossible to make *TEMP-REF* simply a recursive category. (The distinction between S,R and R,E is then void. It remains implicit due to the specific articulation of the evaluation time by some instances of temporal reference and the combination with VPs in other cases).

3.3 Alternative interpretations of tenses

One of the main arguments for the annotation scheme and, in particular, the syntax/semantics interface presented in section 3.4 is that the semantic interpretation of tense/aspect forms can systematically shift in specific syntactic and semantic contexts. In other words, tense/aspect forms do not only vary cross-linguistically in how they encode temporal reference, but their semantic interpretation also shifts according to the immediate linguistic context.
This section exemplifies this by virtue of some well-known phenomena in the domain of temporal reference. Furthermore, it exemplifies how the semantics of these meaning shifts can be captured formally.

3.3.1 Sequence-of-tense

The first phenomenon I want to present is the so-called Sequence-of-tense(SOT) phenomenon. The SOT-phenomenon is a language-specific behavior of the past tense (or past temporal reference) in complement clauses, i.e., it does not occur in all languages (Grønn and von Stechow, 2010). Nonetheless, it is still a fairly robust occurrence cross-linguistically (Bochnak et al., 2019). For the sake of this thesis, I focus on its effects on temporal reference regardless of the mode of encoding. However, Bochnak et al. (2019) provide some interesting insights on the cross-linguistic variation of the SOT-phenomenon, showing that different ways of encoding temporal reference, as described in section 3.1 lead to the same general behavior with respect to the SOT in certain languages.
Consequently, I see the sequence-of-tense phenomenon as a parameter on temporal reference rather than on tense. Nonetheless, the phenomenon is illustrated most easily by inspecting a prototypical way of encoding temporal reference. Consider example (49). As indicated there, embedded tenses are ambiguous between two distinct ways of constraining temporal reference. The first reading indicates that the *saying* eventuality is simultaneous to the embedded eventuality of *being sick*. The second reading is the typical reading we expect from an embedded relative tense, namely, a back-shifted reading

which I have illustrated for Japanese in section 3.1.4.

(49) John said that Karen was sick.

a. John said: Karen is sick.

b. John said: Karen was sick.

With the semantics we have introduced so far, we seem to be able to at least derive one of these readings. The back-shifted reading where Karen was sick before the saying event.Assume the order of composition given in (50) for this.

(50) $[t^* \; [_{TP}$ PAST $[_{VP}$ ... $[_{COMP}$ say $[_{TP}$ PAST $[_{VP}$ ... $]]]]]]$

As shown there, the matrix tense is related to the evaluation time. On the other hand, the embedded clause is relative to the matrix clause or, more specifically, to the eventuality time of the propositional attitude verb. To clarify this proposal, first, let us compute the two main clauses of the sentence. This concretely illustrates that the embedded tense does not introduce its own evaluation time since it is provided compositionally by the matrix clause. The resulting semantics in (51c), thus, capture the desired back-shift between the time of *saying* (*t'*) and the time of *being sick* (*t"'*).

(51) a. $[\![\text{Tom said that Q}]\!]^{w,g} = \lambda t.\exists t'[t' \prec t \wedge say(t', tom, Q(t') = 1)]$

b. $[\![\text{Karen was sick}]\!]^{w,g} = [\![\text{Q}]\!]^{w,g} = \lambda t''.\exists t'''[t''' \prec t'' \wedge \textit{be-sick}(t''',karen)]$

c. $\lambda t.\exists t'[t' \prec t \wedge say(t', tom, \exists t'''[t''' \prec t' \wedge \textit{be-sick}(t''',karen)])]$

The second reading can simply be computed by omitting the embedded tense operator. This means the embedded clause is treated simply as a function of type $< it >$ after having its NP arguments saturated. The embedded clause is then evaluated with respect to the tense of the matrix predicate with which it shares the temporal argument, just as before.

(52) a. $[\![\text{Karen} \; \sout{\text{was}} \; \text{sick}]\!]^{w,g} = [\![\text{Q}]\!]^{w,g} = \lambda t''.\textit{be-sick}(t'',karen)$

b. $\lambda t.\exists t'[t' \prec t \wedge say(t', tom, \textit{be-sick}(t',karen))]$

This approach to derive the two readings of the SOT-phenomenon is based on Kusumoto (1999, 2005). However, the underlying core-principles have been discussed in the literature for a while now (Abusch (1988, 1997) and Heim (1994), Ogihara (1994) and subsequent work, Grønn and von Stechow (2010)). There is an alternative line of analyzing SOT-sentences based on the competition of morphological features and pragmatic inferences based on them. Furthermore, the notion of stativity plays a role in these approaches (Gennari (1999) and subsequent work, Altshuler and Schwarzschild (2012)). While the debate does not seem to be settled, there are certain benefits in following a structural approach, as entertained above. This will become apparent in section 3.4 when I encode the observations made in this section in a computationally viable annotation scheme.

3.3.1.1 The modal component of propositional attitudes

In the literature on the SOT-phenomenon, a second important factor is the modal nature of propositional attitude verbs (Hintikka, 1969). In the proposal made in this chapter, I ignore modality under the assumption that appropriate modeling of the semantics of propositional attitude verbs accounts for potential mismatches in the temporal interpretation of embedded clauses. In this paragraph, I briefly illustrate cases where a temporal reference configuration is compatible with different seemingly conflicting temporal coordinates, i.e., examples where a modal coordinate may be used to account for differing interpretations concerning the temporal allocation of the eventuality time, reference time, and evaluation time.

Let us assume that Zuzana woke up at 6.00 am and saw that there was still somebody in Vashek's bed. However, in Zuzana's belief worlds, it is 7.00 am because she forgot to consider daylight saving time. This means that it is 6.00 am in the actual world and 7.00 am in Zuzana's belief worlds. Given this scenario, the following two sentences are true:

(53) a. Zuzana thought that Vashek was asleep at 7.00 am.

b. Zuzana thought that Vashek was asleep at 6.00 am.

The first reading, the *de dicto* interpretation, is true in case the complement

is about Zuzana's belief worlds at the time when she woke up. The second, *de re* interpretation is true if the embedded clause is about the actual world. In this thesis, I generally assume that E,R, and S refer to the same entity across different worlds. However, by employing certain modal elements (as shown in 53), we can reach a situation where that entity can receive different temporal interpretations across different possible worlds.
The analysis presented here is completely compatible with this insight (see Kusumoto 1999, 2005). However, for this thesis, I do not fully implement a modal semantics for propositional attitude verbs. Generally, formal approaches to modal semantics need to be handled with care from the perspective of NLP. I will illustrate this by virtue of an analysis of grammatical aspect in the next chapter. There, the semantic analysis of the category is more intrinsically intertwined with the notion of modality, other than here, where the more obscure cases fall out from the interaction with certain modal elements. At the same time, the core semantics provided by temporal reference do not have to be changed in the presence of these elements.

3.3.2 Double-access readings

A different kind of semantic shift, predominantly reported for (some) Germanic languages, is encoded in terms of a phenomenon called *double-access* interpretation. *Double access* here refers to the property that a reference time is related to two different evaluation times. Consider the following example:

(54) Peter said that Mary <u>is</u> pregnant.

a. $\exists t'[t' \prec t^* \wedge say(t', p, \exists t''[t'' \circ t' \wedge t''' \circ t^* \wedge pregnant(t'', m)])]$

As shown in (54a), the embedded present tense is simultaneously (and obligatorily) evaluated with respect to both the speech time and the relative evaluation time provided by the propositional attitude verb. In other words, the pregnancy's duration must overlap with the reporting event and the time of speech. Ogihara (1994), for this purpose, suggests duplicating the present tense. In this thesis, I adopt this approach, disregarding the modal components of this analysis.[20] As the formula in (54a) suggests, the two different constraints are bound by the same temporal quantifier that provides the reference time for both of them but only provides the evaluation time for one of them. To account for this, I propose to discern the temporal quantifier and the temporal constraint:[21]

(55) $\lambda P_{<i,it>}.\lambda Q_{<i,t>}.\lambda t^0.\exists t[P(t^0)(t) \wedge Q(t)]$

a. $\lambda t.\lambda t'.t' \prec t$ *past*

b. $\lambda t.\lambda t'.t' \circ t$ *pres*

c. $\lambda t.\lambda t'.t \prec t'$ *future*

In cases where it is possible to combine two temporal reference restrictions along the lines of the double access semantics presented above, the second constraint is treated as a modifier on the restrictor of the temporal reference operator. Intuitively, the present tense in the present example is simultaneously interpreted as absolute and a relative present temporal reference in the present framework.[22]
A complex restrictor can be created by using the device in (56). It works similarly to predicate modification. The important feature here is that it does not affect the reference time (i.e., temporal reference introduces at most one reference time), but only affects its relation to other points in time.

(56) $\lambda P_{<i.t>}.\lambda Q_{<i.t>}.\lambda t.P(t) \wedge Q(t)$

To get the semantics right, the two arguments must be partially saturated with distinct evaluation times. As pointed out above, this is accounted for if we assume one instance of relative present temporal reference and one instance

[20]See Romero (2014) for a similar proposal that can be raised somewhat easily to form an account that includes the modal elements of the double-access analysis.

[21]As, for example, also proposed in Haug (2008).

[22]This begs the question of whether existential quantification is contributed by temporal reference, or by some other mechanism (e.g., a finiteness effect or some other kind of clause boundary). This question arises since the modifying instance of present temporal reference does not introduce its own temporal quantifier (it could, but that seems uneconomical). For the sake of this thesis, I aim to propose the most economical solution given the underlying framework.

of absolute present temporal reference. To be more precise, the mechanism I propose for this, in principle, uses a trace for the relative temporal reference constraint and t^* for the absolute temporal reference constraint. The semantic composition is illustrated in (57). It results in the quantifier in (58), which can be combined with the VP it scopes over and then with the propositional attitude verb to derive the desired reading of (54).

(57)

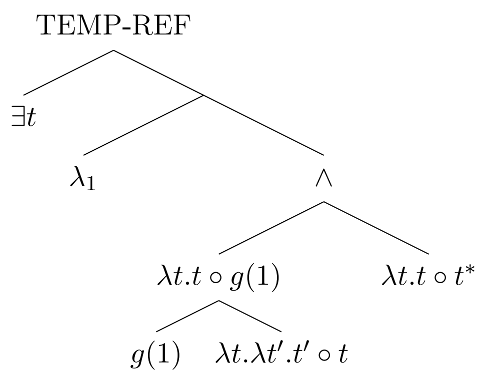

(58) $\lambda Q.\lambda t^0.\exists t[t \circ t^0 \wedge t \circ t^* \wedge Q(t)]$

Decomposing temporal reference in such a way is beneficial for analyzing the double-access reading and also allows for dealing with other potential modifiers of temporal reference. This idea is revisited again in chapter 5.[23]

3.3.3 Tense as marker for counterfactuality

A final instance where tenses, specifically past tense, behave oddly in English is in counterfactual contexts, as, e.g., in counterfactual conditionals but also when embedded under modals with a counterfactual flavor, as in the following example:

(59) I wish I <u>had</u> a car.

[23]Note also that this picture is somewhat complicated when an analysis of the modal component of propositional attitude verbs comes into play: as Bary et al. (2018) show, double access readings are also possible in a context where the embedded belief is not true at present in the actual world, i.e., the temporal anchor can be shifted on a modal axis. Figuring out the exact semantics of such an account is left for future work.

In this example, the underlined past-tense marker is not interpreted temporally but counterfactually. In some sense, it seems to "agree" with the modal that dominates it.[24] In such cases, I assume that temporal reference is simply empty. Thus, the semantics of (59) would be as follows:[25]

(60) $\forall s[s \in BUL(x, t^*) \rightarrow have(x, c)]$

This executive decision is made mainly based on the fact that counterfactuality is expressed differently cross-linguistically. Furthermore, it is not clear if counterfactuality works in the same way in the example above as it works, for example, in counterfactual conditionals such as:[26]

(61) a. If John was/(were) there now, it would be fun. (Romero, 2014)

Overall, from a temporal perspective, in both cases, the tense is not evaluated in terms of temporal reference. Thus, this is another argument that the syntax/semantics interface must allow for changing the semantic contribution of a morphosyntactic marker. This is a topic for section 3.4.

3.3.4 Interim summary

The discussion in this section has focused on English since the center stage was occupied by the formal semantics rather than the languages that they are applied to. However, each of these phenomena has received considerable cross-linguistic coverage. Sequence-of-tense has been discussed cross-linguistically, e.g., by Bochnak et al. (2019), Grønn and von Stechow (2010), Kusumoto (1999), Mucha (2015), and Ogihara (1994). In addition to the previously cited Ogihara (1999), many of these sources have also discussed double-access readings in some capacity, since sequence-of-tense and double-access can be subsumed under the general topic of tense in complement clauses. Romero

[24]Note that the requirement of the embedded tense is also somewhat affected by the context. Thus, in a scenario where you utter a wish to a djinn, you could also say: *I wish I have a car* or, less colloquially: *I wish to have a car*.

[25]*(to) wish* is treated as a quantifier over situations with a bouletic modal base appealing to the desires of the subject following a simplified Hintikka (1969)-style modal semantics.

[26]See Fleischman (1989) for a basic overview of different instances of non-temporal uses of tenses.

(2014) made use of the insights gained from the analysis of tenses in embedded contexts to provide a formal semantics for counterfactual conditionals. Thus, the formal devices presented here are not only supported by cross-linguistic research but have also been explored across different linguistic contexts.

In summary, I have introduced a temporal semantics that is based on the formal semantic advances made in recent years. I will translate this logic into an approach compatible with computational approaches to syntactic and semantic analysis. There, my goal is to strip the annotation of theory-driven assumptions as much as possible and provide a general, powerful tool to generate semantic representations from syntactically parsed inputs. This is done in chapter 5.

The next chapter aims to establish a syntax/semantics interface that allows us to capture the facts outlined in the previous sections of this chapter. The interface should allow us to encode simple instances of temporal reference and the more complex examples where temporal reference undergoes a semantic shift, as illustrated in this section.

3.4 Annotation of temporal reference

As Comrie (1985) says, a theory of tense must be "sufficiently flexible to accommodate tense distinctions that recur across the languages of the world while being sufficiently constrained to exclude logically possible distinctions that are not, in fact, possible in human language" (p. 50). The annotation scheme presented in this work does not serve as a full-fledged theory of tense in that sense. Rather I present a system that is deliberately capable of describing all logical possibilities. Thus, the annotation can serve the exploration of distinctions of linguistic categories without being constrained by theoretical assumptions.[27]

In the previous sections, I presented cross-linguistic variation of tempo-

[27] The annotation examples presented in this work will depend on certain theoretical assumptions made in the literature. However, this should be understood as a possible representation of the behavior of specific linguistic categories, which provides a basis for empirical testing of exactly these assumptions.

ral reference at both the syntactic and semantic level. I have shown that tense/aspect forms are combinations of different structural reflexes with distinct but sometimes overlapping semantics. Importantly, there is rarely a one-to-one correspondence between form and meaning. I have shown that semantics can be derived from the morphosyntactic context of a tense/aspect marker or even the semantic context (e.g., sequence-of-tense, embedding under a counterfactual modal). Thus a system for generating intermediate meanings is indispensable. These intermediate meanings can then be used to construct the final meaning of an utterance if direct interpretation from the morphosyntactic input is insufficient. In theoretical linguistics, this step is interwoven with the principles of compositionality; thus, modeling semantic phenomena is a delicate matter.

I intend to provide a more independent approach to deriving semantic representations for language. To be more precise, I propose a model more along the lines of Bybee (2006) (although Bybee's work aligns meaning with form in terms of phonetic representations rather than morphosyntactic representations). Where I go further is that I assume that the existence of certain meanings licenses additional inferences which can be interpreted as being constructed from the original meanings. Thus the principles of compositionality are not entirely lost. The resemblance (if rather superficial) to (Bybee, 2006) model stems from the idea that there should be an easy-to-understand relation between form and meaning that may be easily quantified. Furthermore, I reduce the load on the compositional apparatus by representing intermediate steps simply in terms of rules rather than assuming a full semantic analysis of the underlying meaning construction (sometimes called description-by-analysis, in particular, in the LFG literature (Andrews, 2008; Crouch, 2005; Crouch and King, 2006; Wedekind and Kaplan, 1993)). Another factor that differentiates the semantic construction that I assume from the more traditional approaches in generative grammar is that my system is not designed for a pairwise combination of meaning features but allows for inferences of meaning features based on an unrestricted number of premises either syntactically or semantically.

In short, compositional processes are simplified where possible. Furthermore, meaning contributions of different levels (i.e., semantic vs. pragmatic) are treated similarly in their composition. To illustrate the simplified composition, compare the two structures in (62). Assuming that there is no external motivation for the existence of δ', the computation in (62b)contains superfluous steps in reaching the desired meaning. A binary tree is necessary for a compositional framework based on functional application (all semantics based on functional application ultimately result in a (set of) tree(s)). However, by operating on a constraint or feature-based level, we can combine multiple features non-compositionally, i.e., we collapse certain parts of the analysis into construction rules that generate either semantic features (for further semantic construction) or semantic representations if no other semantic construction rules apply.

(62) Composing features vs. composing functions

a. $\alpha \wedge \beta \wedge \gamma \rightarrow \delta$

b.
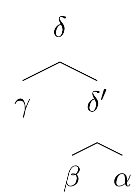

This is the foundation of the syntax/semantics interface designed to translate morphosyntactic features into semantic features. It consists of a set of rules that translates syntactic features into semantic features and a set of rules that translates semantic features into semantic representations. Before I discuss the details of annotating the syntax/semantics interface, I will discuss the feature space that is relevant in this chapter in the next section. I will continue to discuss how these values are derived from the morphosyntactic, semantic and pragmatic context, starting in section 3.4.2.

3.4.1 Feature space for the category *temporal reference*

Based on the observations made in the previous sections, I employ the feature space I have organized in figure 3.1 as an attribute/value matrix. As shown there, the annotation of tense is deconstructed into three elements: Temporal reference, modifiers on temporal reference, and the evaluation time. In

$$\begin{bmatrix} \text{TEMP-REF} & \begin{bmatrix} \textbf{t-ref} ::= \text{`past'} \mid \text{`present'} \mid \text{`future'} \mid \\ \text{`non-past'} \mid \text{`non-present'} \mid \text{`unspec'} \\ \textbf{t-restr} ::= \text{`immediate'} \mid \text{`recent'} \mid \text{`non-recent'} \\ \textbf{eval} ::= \left\{ \text{`sid',...} \right\} \end{bmatrix} \\ \text{ASP-TREF} & \begin{bmatrix} \textbf{a-ref} ::= \text{`past'} \mid \text{...} \\ \textbf{a-restr} ::= \text{`immediate'} \mid \text{...} \end{bmatrix} \end{bmatrix}$$

Figure 3.1: Temporal reference feature template

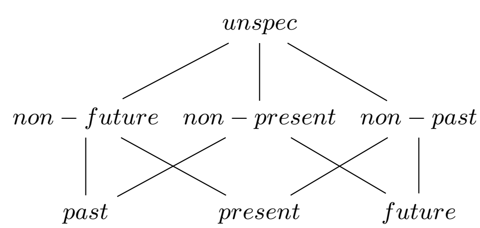

Figure 3.2: Tense features

addition, I propose, for ease of annotation, `ASP-TREF`, which is an instance of temporal reference that is always relative to the reference time provided by `TEMP-REF`. Thus, it provides a more restricted version of `TEMP-REF` that applies recursively to tense in the fashion discussed in section 3.2.4.1. The evaluation time of a tense is specified in `EVAL` and refers either to the speech time (*now*) or a compositionally specified reference time, as discussed in the previous section.

As I have shown, there are several special properties of temporal reference: i) temporal reference can be vague, e.g., non-past or non-future, ii) temporal reference can be ambiguous, as, e.g., in the sequence-of-tense phenomenon, and iii) temporal reference can be semantically shifted or modified(as exemplified by double-access readings).

Figure (3.2) summarizes the features associated with temporal reference in a lattice. In this framework, the intermediate level is used to express vague temporal reference, and the top level is used to express a temporal anchor that is not temporally restricted. For example, a vague temporal reference is provided by tenses, where the reference time can follow one of three configurations (depending on the value). For example, *non-past* can indicate that the reference time is the speech time, some time in the future, or an interval that is started by the speech time and ends at some future time. This is illustrated in the following example:

(63) Temporal progression chart:

	—	S	—>
[	NON-FUT	]	
		[	NON-PAST]

The requirement for annotating vague semantic values is tricky. It is difficult to specify whether a reference time simply overlaps with the now or extends beyond it. A possible candidate for testing vague tenses are temporal expressions associated with the extended now (Bennett and Partee (1978) and Dowty (1982) among others). These expressions specify a particular start point or end point, but the respective other end of the specified time interval is the utterance time, i.e., the present (although they might be used relatively). Examples of these kinds of expressions are *since yesterday* for non-future and and *until 4:00pm* for non-past. Taking these kinds of temporal expressions as a test for vague tenses is an enterprise that needs to be handled with care, but it provides a reasonable enough starting point.
Ambiguity, in principle, occurs when a set of mutually exclusive meanings apply to the same target. For example, if a given feature can have either a present or future interpretation, then it can be annotated in the following way (the feature, there, is called *non-past*$_b$, which makes a stronger semantic claim than the previously introduced vague *non-past*).

(64) a. $[\![non - past_b]\!] = (\!|present, future|\!)$

Example (64a) introduces a notion for sets of meanings that I call meaning family, i.e., packed representations of meaning (Lev, 2007). This distinction is necessary to differentiate it from the typical use of sets of meanings in the semantics of questions. The benefit of this notion is that we can pack meanings together that map onto different semantic representations.

Thus, the intuitive idea that ambiguous expressions differ at some level of representational structure is preserved. On the other hand, vague expressions do not invoke multiple representational structures.
The precise implications of this variation depend on the semantic formalism that is stacked on top of the annotation. However, it is particularly well suited in combination with XLE grammars since it extends the idea of packed representations to the domain of semantics. In the semantics I have proposed, meaning families can be combined with other elements in the calculation via pointwise functional application (PWA) to yield a family of meanings with different semantic representations.

(65) $[\![\text{PWA}(\phi \subseteq D_{<\rho,\tau>}, \psi \subseteq D_{\tau})]\!] = (\!| \, f(x) \in D_{\tau} : f \in \phi \wedge x \in \psi \, |\!)$
(Shimoyama, 2006)

Encoding the resulting semantic representations in a packed manner is a separate issue. In this work, the resulting semantics are ultimately unpacked (see chapter 5). In the next section, we will encounter another type of ambiguity, i.e., ambiguities at the level of semantic structure.
While ambiguous and vague tenses constrain the possibilities of locating the reference time, double-access introduces multiple evaluation times. Thus, the evaluation time is encoded as a set of different evaluation times as specified compositionally or by the context. Evaluation times are referred to via a unique identifier. This is illustrated in more detail in chapter 5.
With the feature space presented in this section, we can deal with the discussed semantic properties of temporal reference, such as vague, ambiguous, or semantically unspecified temporal relations. The next section explains how these features can be derived via a rule-based syntax/semantics interface.

3.4.2 The syntax/semantics interface in ParTMA

This section addresses the second part of the ParTMA annotation scheme: the semantic construction rules. These rules are designed to simplify the compositional progress on the one hand and, on the other hand, anchor contextual information in the interpretation of constructions that are underspecified in terms of their overt morphosyntactic realization of meanings. To illustrate the concept, I begin with constructing meanings for simple tense/aspect forms. I will then turn to more complex expressions, where the meaning of a given tense/aspect form is shifted, such as the SOT-phenomenon and counterfactuals.

3.4.3 Basic mappings from syntax to semantics

To begin, I provide straightforward rules that map a specific syntactic form directly on their meaning feature, i.e., temporal reference. For this purpose, I will again rely on the well-discussed tense system of the English language. As per usual, I introduce examples of cross-linguistic variation where appropriate. Mapping structure directly to meaning seems to be a redundant enterprise that just serves to duplicate a given feature, but, as I have elaborated in the previous section: the semantics of a tense category are more complex than a simple feature: We need not only specify a temporal ordering with respect to the evaluation time but also that very evaluation time.
Thus, the English past tense morphology can be simply mapped onto past tense, as illustrated in (67), but we also need to specify the evaluation time. This example visualizes a rule describing an implication relation between morphology and syntax. The semantic features also receive a quasi-formal interpretation to illustrate their formal meaning.[28]

(66) People killed the king.

(67) $\begin{bmatrix}\text{TNS-ASP} & \begin{bmatrix}\text{TENSE past}\end{bmatrix}\end{bmatrix}$ $\qquad$ $\begin{bmatrix}\text{TEMP-REF} & \text{`past'} : t \prec t_0 \\ \text{EVAL}\left\{\text{t}^*\right\} & \end{bmatrix}$

(68) `TENSE past` $\rightarrow$ `TEMP-REF 'past'` : $t \prec t^* \wedge$ `EVAL {`t^*`}`

The same type of rule can be used to compose multiple morphosyntactic features into a semantic interpretation, for example, the past temporal reference interpretation of the German present perfect:

[28]However, this is only to show that the annotations result in the correct temporal constraints. As shown in Chapter 5, the actual relation between semantic features and semantic interpretation is more modular.

(69) Das Volk ha-t den König getötet
The people have-PRES the king kill.PSTPTCP
'The people killed the king.' *German*

(70) $\left[\text{TNS-ASP} \quad \left[\text{TENSE pres, PERF} +\right]\right] \begin{bmatrix} \text{TEMP-REF} & \text{'past'} : t \prec t^* \\ \text{EVAL} \left\{t^*\right\} & \end{bmatrix}$

(71) `TENSE pres` $\wedge$ `PERF +` $\rightarrow$ `TEMP-REF 'past'` : $t \prec t^* \wedge$ `EVAL {t*}`

So far, I have discussed temporal reference as a relation between the reference time and evaluation time. Let us now take a look at instances of tenses with two layers. Firstly, the basic example of past perfect constructions. As explained in the previous section, relative tenses that apply directly to a reference time can be simply annotated with the `ASP-TREF` feature, as shown below. In this case, the `ASP-TREF` behaves as if it was embedded under `TEMP-REF`.

(72) Rory had left. $\rightsquigarrow$ [PAST[PAST VP]]
Klecha and Bochnak (2016)

(73) $\left[\text{TNS-ASP} \quad \left[\text{TENSE past, PERF} +\right]\right] \quad \begin{bmatrix} \text{TEMP-REF} & \begin{bmatrix} \textbf{t-ref} ::= \text{'past'} \\ \textbf{eval} ::= \left\{t^*\right\} \end{bmatrix} \\ \text{ASP-TREF} & \begin{bmatrix} \textbf{a-ref} ::= \text{'past'} \end{bmatrix} \end{bmatrix}$

(74) `TENSE past` $\wedge$ `PERF +` $\rightarrow$ `TEMP-REF 'past'` : $t \prec t^* \wedge$ `EVAL {t*}` $\wedge$ `ASP-TREF 'past'` : $t' \prec t$

More generally, this can also be expressed via multiple instances of `TEMP-REF`, which would be more faithful to the initial inspiration provided by Klecha (2016) and Klecha and Bochnak (2016). Consequently, although we work with the two distinct relations provided by Reichenbach, they behave formally similarly in the temporal domain.

(75) $\left[\text{TNS-ASP} \quad \left[\text{TENSE past, PERF} +\right]\right] \quad \begin{bmatrix} \text{TEMP-REF} & \text{r}\begin{bmatrix} \textbf{t-ref} ::= \text{'past'} \\ \textbf{eval} ::= \left\{t^*\right\} \end{bmatrix} \\ \text{TEMP-REF} & \begin{bmatrix} \textbf{t-ref} ::= \text{'past'} \\ \textbf{eval} ::= \left\{r\right\} \end{bmatrix} \end{bmatrix}$

In this example, the order of application is indicated by the evaluation time. The symbol r stands for the identifier of the respective structure. This means the `TEMP-REF` structure corresponds to a reference time or, more generally, describes the temporal constraints on a time interval.

3.4.4 Structural ambiguity in the semantic structure

The German present perfect is ambiguous between a present perfect ($t' \prec \circ t \prec t^*$) and a past ($t \prec t^*$) reading.[29] In this case, simply applying ordered rewrite rules is not sufficient to capture the ambiguity. Rather, we need to introduce a structural ambiguity in the annotation structure that can be resolved later. Let us first discuss how to arrive at this structure:

(76) a. `TENSE pres` $\rightarrow$ `TEMP-REF 'pres'` : $t \circ t^*$

b. `TENSE past` $\rightarrow$ `TEMP-REF 'past'` : $t \prec t^*$

(77) a. `TENSE pres` $\wedge$ `PERF +` $\rightarrow$ `TEMP-REF 'past'` : $t \prec t^*$

b. `TENSE pres` $\wedge$ `PERF +` $\rightarrow$ `TEMP-REF 'non-future'` : $t' \prec \circ\, t$

The rules in (76) are our typical tense rules. However, the perfect introduces two different possible semantic interpretations: the perfect-as-past interpretation that replaces the temporal information provided by (76a) with a past interpretation, and a perfect interpretation that interacts compositionally with present temporal reference.

Written as rules: the first rule in (77) rewrites the relation between the reference time t and the evaluation time $t*$ resulting in the perfect-as-past interpretation. The second rule describes the relation between a second temporal variable t' and the reference time, i.e., the temporal semantics that the perfect contributes when interpreted as an aspectual marker. Both of these rules are possible interpretations. This leads to a structural ambiguity in the semantic feature structure, as illustrated in Figure 3.3, which displays a set of possible (sub)structures of a given semantic structure. There, **A1** and **A2** refer to two sets of mutually exclusive constraints on the semantic

[29] $t' \prec \circ\, t$ here stands for non-future. A simple conception of the perfect's temporal semantics.

$$\left\{ \begin{array}{l} \textbf{A1:} \begin{bmatrix} \text{TEMP-REF} & \begin{bmatrix} \textbf{t-ref} ::= \text{'past'} \\ \textbf{eval} ::= \left\{ t^{*} \right\} \end{bmatrix} \end{bmatrix}, \\ \textbf{A2:} \begin{bmatrix} \text{TEMP-REF} & \text{r}\begin{bmatrix} \textbf{t-ref} ::= \text{'pres'} \\ \textbf{eval} ::= \left\{ t^{*} \right\} \end{bmatrix} \\ \text{TEMP-REF} & \begin{bmatrix} \textbf{t-ref} ::= \text{'non-future'} \\ \textbf{eval} ::= \left\{ r \right\} \end{bmatrix} \end{bmatrix} \end{array} \right\}$$

Figure 3.3: Ambiguous semantic structure

structure. **A1** refers to the past interpretation, whereas **A2** refers to the proper present perfect interpretation.

This kind of ambiguity has been studied extensively in the computational LFG literature, and the XLE provides tools to handle this (Crouch et al., 2017; King et al., 2000). Crouch (2005) concretely describes ambiguity management in the domain of rewriting. Further information is given in chapter 5 when discussing the concrete implementation of these rules.

Although the examples above explain the basic considerations that go into developing semantic mapping rules, the predictions made there are not fully congruent with the observations made in the first part of this chapter. To refine the analysis above, the German present tense should receive the vague value `non-present`. This clarifies that although three basic tense forms are available in German, the present tense is not mutually exclusive with the future form. However, there are certainly instances where the present tense in German receives either a present or future tense interpretation. This can be made explicit by combination with temporal adverbs.

(78) a. Paul komm-t gerade.
Paul come-PRES currently.
'Paul is currently arriving.'

b. Peter komm-t morgen.
Paul come-PRES tomorrow.
'Paul kommt morgen.' *German*

In such cases, the tense is restricted by the temporal orientation of the temporal modifier. Thus, a future-oriented interpretation is ruled out for (78a) and a present-oriented interpretation is ruled out for (78b). This is a refinement of the original (and still correct) value. This, in principle, is not an issue for rule-based approaches to mapping syntax to semantics but may lead to contradicting annotations when manually annotating the semantics based on a syntactic input (In this thesis, I will focus on the rule-based syntax/semantics interface; see also section 3.4.8 for some additional thoughts on this issue).
As I have shown in section 3.3, there are several additional semantic properties of tenses, particularly in embedded contexts. In the next section, I will discuss the annotation of these instances and how they can be derived via a rule-based system of semantic construction rules, as presented in this section. While the rules in this chapter have exemplified semantic construction via graph expansion (the rules illustrated above have been used to introduce new elements of the semantic structure), the next section will turn to semantic annotation via rewriting, producing a layered annotation of temporal reference features.

3.4.5 Secondary meanings of temporal reference in ParTMA

The main goal of this section is to firmly establish the rules I have introduced above by applying them to more elaborate semantic phenomena. I want to show that these rules allow me to cover a broad range of examples of morpho-syntactically expressed meanings but also more complex cases which require an in-depth semantic analysis.

3.4.6 Semantic construction: Sequence of tense

I have introduced the sequence-of tense phenomenon in section 3.3. There, I briefly argued for a structural account over a pragmatic account, as is recently debated in the literature. In this section, I show how to implement a structural account of the SOT-phenomenon that allows us to capture its semantic properties successfully. First, I will deal with the most basic examples of sequence-of-tense as reiterated below in (79). The main property that needs to be captured is that the embedded eventuality can be interpreted either simultaneously or back-shifted with respect to the matrix eventuality.

(79) Tom said that Karen was sick.

a. Tom said: Karen is sick.

b. Tom said: Karen was sick.

The basic requirements our annotation of the SOT-phenomenon should fulfil are the following:

(80) i. Annotate the SOT phenomenon as a instance of ambiguity

ii. Root the inference of the semantic labels in the morphosyntax and semantics

iii. Applicable cross-linguistically

The first point in (80) was discussed in section 3.3 . Point ii.) guarantees that we, in fact, create a structurally motivated annotation rather than one that is rooted in the pragmatics. The first two points are, of course, intertwined. The third point is important since it makes it easier to compare linguistic categories based on their annotations in different languages. The more properties can be streamlined across languages, the better the cross-linguistic validity of the proposed annotation is verifiable, and the better different languages are comparable.

The prevalent treatment of the SOT-phenomenon in formal semantics is based on coindexing the tense in the complement clause with the tense in the matrix clause in one reading and introducing an embedded tense operator for the other reading. In other words, the embedded tense is not realized as such but is simply a vacuous element in one of the readings. This ascertains, without any additional assumptions, that the upper-limit constraint entertained by Abusch (1988) is met (i.e., this constraint guarantees that there can be no future-shifted interpretation of the embedded eventuality). I want to stipulate a similar approach that only differs in the technicalities. Firstly, the matrix tense introduces a tense relative to the speech time, i.e., an absolute tense. This tense stems from the basic English tense rule:

(81) `TENSE past` $\rightarrow$ `TEMP-REF 'past'` : $t \prec t_0$

We cannot apply the same rule to the embedded tense since the two tenses would behave completely independently of each other (although, in this case, we would explain the SOT patterns in terms of a type of pragmatic approach). This independence would result in more readings than are typically licensed by the SOT phenomenon, in particular, a reading where the embedded tense situates the event it restrains between the time of the matrix event and the speech time. The allowed readings are summarized in the table below, where t^m stands for the reference time of the matrix event, t^c stands for the reference time of the event in the complement clause, and t^0 corresponds to the speech time:

Table 3.8: SOT vs NON-SOT language

Readings	**SOT**	**NON-SOT**
$(t^c \prec t^m) \prec t^0$	+	+
$(t^c \otimes t^m) \prec t^0$	+	-
$t^m \prec t^c) \prec t^0$	-	-

The structural co-indexing approach to the SOT-phenomenon hinges on the assumption that the embedded tense can be interpreted relative to the reference time of the matrix tense. Accordingly, I propose a relative tense treatment for the embedded tense as well. However, I do not argue for a simple deletion of the embedded tense but rather for interpreting the embedded tense as ambiguous between a relative present tense and a relative past tense. The present tense then fulfills the task of co-indexing the embedded tense with the reference time of the matrix tense. This assumption has its own implications. Mainly, the matrix tense is interpreted semantically as such (or at least has some temporal interpretation). This is necessary since it is impossible to define the embedded tense as relative without an interpreted matrix tense. Based on these assumptions, the sequence-of-tense rule must interact with the basic tense rule in such a way that the semantic properties of the two tenses are met, as illustrated in (82).

(82) a. $[\![PAST_{matrix}]\!] = \lambda t.t \prec t_0$

b. $[\![PAST_{comp}]\!] = (\!|\lambda t'.t' \prec t, \lambda t'.t' \otimes t|\!)$

To achieve this goal, the SOT-rule (in (84)) rewrites the basic temporal reference value provided by the simple tense rule. The SOT-rule, as motivated by cross-linguistic evidence suggests, takes into account that the SOT-phenomenon results from two instances of temporal reference rather than the underlying morphosyntactic representation. Consider the German example in (83). At least in areas where the perfect construction robustly carries a past tense meaning, the two readings attributed to the SOT-rule emerge.

(83) Paul hat gesagt, dass Julia krank gewesen ist.

To reiterate, this means that the SOT-phenomenon is not restricted to tenses but is invoked when two instances of temporal reference interact with each other. Further evidence for this line of reasoning is provided by cross-linguistic research on understudied languages (Bochnak (2016), Bochnak et al. (2019), and Mucha (2015) and others). Thus, I propose to analyze the SOT-phenomenon in the following way:

(84) $\texttt{TEMP-REF}_{matrix}$ `'past'` : $\lambda t.t \prec t_0 \wedge$
$\texttt{COMP}(E_1,(E_2) \wedge$
$\texttt{TEMP-REF}_{comp}$ `'past'` : $\lambda t'.t' \prec t \rightarrow$
$\texttt{TEMP-REF}_{comp}$ $(\!|$ $\texttt{present}_t$, $\texttt{past}_t$ $|\!)$

The rule in (84) states that given a complement construction with a matrix eventuality $E1$ and a complement eventuality $E2$ where both eventualities have a semantic past restriction (`TEMP-REF`), then the embedded tense, $\texttt{TEMP-REF}_{comp}$, is rewritten into a ambiguous feature, $(\!|$ $\texttt{present}_t$, $\texttt{past}_t$ $|\!)$. This rule weakens an existing meaning (past temporal reference), transforming it into a more general meaning (past or present relative temporal reference). The rule above is, of course, not the end of the story, but it provides the basic template for the SOT-rule that is cross-linguistically applicable. The important fact that this section illustrates is that features can be layered, which is modeled via rewriting in the present system. This is further described in the next section, where I model a possible analysis of counterfactuals in a similar fashion.

3.4.7 Semantic construction: counterfactuals

Certain instances of past tense in English sometimes receive an interpretation that is not associated with temporal distance but rather with counterfactuality (Fleischman, 1989). In this section, I want to discuss two instances of this. The first method is invoked via embedding under certain modals. A more complex case is provided by counterfactual conditionals, where multiple layers of tense can occur in the form of the past perfect.

Consider example (85). This statement is clearly about the present and potential future state of affairs. Nonetheless, past tense morphology is used in the embedded clause *I had a car*. Thus, in this case, we would have to produce a rule that contradicts our initial idea that morphological past tense in English implies a semantic past. However, this does not seem to be a case of compatibility since the non-past interpretation is the only available one.[30] There is no past interpretation in this case.

(85) I wish I had a car.

It is true that there appears to be a shift in temporal meaning here. This shift in meaning appears in very clearly defined morphosyntactic contexts. In this case, the embedding is under a bouletic modal, i.e., a modal that expresses the desires/wishes of the speaker. Generally, non-past uses of the English past tense are all firmly restricted to certain morphosyntactic contexts. Thus, it makes sense to write a rule that considers these specific contexts in determining the value of the semantic tense. Consider the schematic rule in (86), which tells us that embedded under a predicate, an eventuality marked for past tense may express the meaning `non-past`.

(86) `PRED 'wish'` $\wedge$ `TENSE past` $\rightarrow$ `TEMP-REF 'non-past'`

[30] Of course, if one assumes the temporal restriction in 85 to be a non-past restriction, then one could devise a rule to choose one of the two in accordance with the context.

Compare example (86) to our initial rule for the English past tense morphology. The difference between the two rules is apparent: The rule for the counterfactual use of the English past tense is more complex than the rule for the prototypical use of the English past tense morphology, as was also the case for the SOT-rule. However, since this rule does not rewrite the default value of past temporal reference, this rule again introduces a structural ambiguity. This ambiguity seems to be language specific. For example, American English tends to allow utterances like the one in (87) colloquially. There, the semantics of past tense and counterfactuality are combined in the simple past tense marking.

(87) I wish I had a car back then.

This suggests that tense forms should retain their original semantic interpretation in cases where rewriting would delete the original value. The previously discussed SOT-example preserved past temporal reference in the rewritten value. However, in the example presented here, the suggested meaning shift is from past to non-past. This seems to be defeasible by concretely specifying a past temporal context. Thus, here, it might make more sense to introduce an ambiguity at a structural level that can be disambiguated by the context later on.[31]

Furthermore, considering an annotation for modals, it might make more sense to design the rule such that it is sensitive to certain modal features rather than to a specific `PRED` value, making it more robust. Generally, this is one of the main assumptions of this annotation scheme. Grammatically distinct meanings also differ in the complexity of their rules. Frequent, prototypical uses are described in terms of simpler rules, while more complex rules describe more complex, often less frequent, meanings. Thus, there is a nice interaction between complexity and frequency that is well worth exploring.

The non-past interpretation of the past tense in a counterfactual context can

[31] Generally speaking, the effect discussed here may be highly language and context-specific. Thus, this simply illustrates a possible pitfall of rewriting semantic values rather than providing a comprehensive semantic analysis of the issue or similar issues that may be found in other tense/aspect forms.

also be moved to the past explicitly by adding perfect morphology. This is shown in (88b).

(88) a. I wish I owned a fast car.

b. I wish I had owned a fast car.

The examples in (88) illustrate that the counterfactual requires an additional layer of tense to shift the temporal reference of the embedded predicate to the past, leaving the second instance of temporal reference as the origin for counterfactuality. A similar effect can be observed for counterfactual conditionals.

(89) a. If John had come yesterday, it would have been fun.

b. If John was/(were) there now, it would be fun. (Romero, 2014)

Other than SOT, the fake tense in counterfactuals is not a parameter that results in a two-way distinction of languages across the world. Rather, even among languages that carry use a fake tense, the expression of counterfactuality varies. I will briefly sketch an annotation for this phenomenon and compare it to the approach I have chosen in the previous section with SOT. Often, counterfactuals seem to pattern with imperfective morphology primarily although nonetheless the role of the past tense morphology that often occurs in these types of sentences has been thoroughly researched (Ferreira (2016), Romero (2014), Iatridou (2000) and subsequent work, Arregui (2005) and subsequent work, Ippolito (2003) and subsequent work). I will not be able to do all of this work justice at this point, but I will sketch an idea for implementation in the ParTMA annotation scheme. What I want to show is how to annotate the temporal properties of counterfactual conditionals while leaving a blank space for whatever is responsible for the counterfactual interpretation of such event descriptions.

There exist different analyses of counterfactual conditionals in the literature: a temporal remoteness line (Romero (2014) and subsequent work) and a modal remoteness line (Iatridou (2000) and subsequent work). The former of these two approaches assumes that the counterfactual past is still interpreted as a semantic past tense. The modal remoteness line assumes an ambiguity of the past tense between a modal operator and a past tense operator. Pursuing the temporal remoteness line could result in a treatment of counterfactual conditionals in the same fashion as the SOT phenomenon. Consider the following rule as a possible annotation. Again we have two relevant event descriptions. One for the antecedent event E_{ant} and one for the consequent event E_{cons} each with specific syntactic restrictions.

(90) a. If Susan was sick$_{\text{ant}}$, she would be at home$_{\text{cons}}$.

b. If Susan had been sick$_{\text{ant}}$, she would have been at home$_{\text{cons}}$.

(91) $\texttt{TENSE}_{\texttt{ant}}$ `past` $\rightarrow$ $\texttt{TEMP-REF}_{\texttt{ant}}$ `'past'` : $t \prec t_0$

(92) $\texttt{TENSE}_{\texttt{cons}}$ `pres` $\wedge$ $\texttt{VTYPE}_{\texttt{cons}}$ `modal` $\rightarrow$ $\texttt{TEMP-REF}_{\texttt{cons}}$ `'pres'` : $t \prec t'$

(93) `ADJUNCT(ant)` $\wedge$ $\texttt{PRED}_{\texttt{ant}}$ `'if'` $\wedge$ $\texttt{TEMP-REF}_{\texttt{ant}}$ `'past'` $\wedge$ $\texttt{TEMP-REF}_{\texttt{cons}}$ `'pres'` $\rightarrow$ $\texttt{TEMP-REF}_{\texttt{cons}}$ `'non-past'` : $\neg(t \prec t_0)$ $\wedge$ $\texttt{TEMP-REF}_{\texttt{ant}}$ `'non-past'` : $\neg(t \prec t_0)$

These rules model the counterfactual in such a way, that a semantic past tense is attributed to E_{ant} which in turns feeds into the generation of the temporal interpretation of tense in the counterfactual conditional ($\texttt{TEMP-REF}_{\texttt{cons}}$ and $\texttt{TEMP-REF}_{\texttt{ant}}$). The rewrite rule in (93) crucially relies on the morpho-syntactic context of the conditional based on the `ADJUNCT` structure with the `PRED` 'if'. A modal remoteness line-type of approach could be modeled by omitting the first step, i.e., the rule in (91). Thus, the counterfactual interpretation would simply rely on the morphosyntactic input. In this case, the temporal interpretation of the counterfactual tense would simply be part of an alternative mapping from form to meaning without any implications on the existence of a semantic past tense. In fact, the fake tense analysis in counterfactuals only describes a subset of languages. Other languages rely on different ways of subjunctive marking. Consider counterfactuals in Indonesian, which are expressed using a specific conditional marker (Sneddon et al., 2012). In this case, there is no conceivable reason why the counterfactual interpretation should stem from a past temporal interpretation. There are three distinct markers that seem to be associated with counterfactuality as shown in (94).[32]

(94) Seandainya/andaikata/sekiranya musyafir itu kembali ia tetap
If traveller DET return he still
tidak akan memperoleh air minum.
not FUT get water drinkable
'Supposing/if that traveller returned he still wouldn't be given water to drink.' (Sneddon et al., 2012)

Thus, even with both of these possibilities given, there are still languages that will require a different architecture for temporal interpretation in counterfactuals.

From this short discussion, making an executive decision is difficult. What is important at this point is that the approach I advocate here is compatible with either interpretation, although counterfactual rules do not align as neatly cross-linguistically as SOT rules.[33] Nonetheless, I have shown how different types of meanings shifts are dealt with within the ParTMA annotation scheme. Thus, I have shown, in general, how ParTMA can be used to annotate the specific syntactic and semantic context that licenses certain variations in meaning. However, cross-linguistic variation still requires careful examination and modeling of the corresponding rules mapping syntax to semantics.

3.4.8 Pragmatic reasoning with ParTMA

As of now, I have only discussed morphosyntactically realized instances of tenses. However, as elaborated in 3.1 and 3.2, sometimes languages make use of covert tenses or unmarked instances of temporal reference. Such tenses are, too a large degree, context-dependent. Although the main work of this thesis is to provide an annotation scheme that reflects the formal semantics of temporal reference, it is necessary to consider an extension to pragmatic reasoning.

[32] The glossing is my own.

[33] I discussed only some very distinct cases and only very briefly. The cross-linguistic variation in counterfactuals would provide me with many less clear-cut cases, and even in the cases that I presented here, I have only introduced some basic examples. All in all, this requires further investigation.

In this section, I analyze a distinct language concerning its encoding of temporal meaning compared to the mainly Germanic languages I have used so far to illustrate the annotation scheme, namely Indonesian. Indonesian is discussed as an example of a tenseless language. Through this, I will explain the annotation of pragmatic inferences.

Indonesian does not contain any obligatory tense marking; however, some auxiliary items may be used to construct periphrastic tense/aspect forms. I will focus on particular on the aspectual markers *sudah/telah* (perfect) and *sedang* (progressive) as well as the future marker *akan* which complement a mostly pragmatic tense and aspect paradigm in Indonesian (Arka, 2013).

The first thing that needs to be clarified is whether the Indonesian, basically zero-marked verb forms do interact in some way with temporal meanings. I and many others argue that they, in fact, do.[34] Evidence comes from the fact that they combine with the various lexical items – a test that I have mentioned before. Especially the future marker *akan* provides interesting evidence as shown in section 3.1. Consider the example in (95a), where the future marker *akan* picks up a past time to express that something happened after that past time but is still in the past (a relative future temporal reference).

(95) [Q: When you visited Peter yesterday, what did he do after you had dinner? ANSWER:]

a. Peter akan menulis se.buah surat
Peter FUT AV.write one.CL letter
'Peter wrote a letter.' (Zymla, 2017)

(96) Temporal progression chart:

—- (*dinner*) —- (*writing a letter*) —- (*conversation*) – ST

The Indonesian *akan* behaves similarly to the English *going to* construction (Eckardt, 2006; Traugott and Dasher, 2001).[35] It provides an instance of temporal reference directed to the future relative to some salient time point.

[34] For an overview, see Tonhauser (2015) and citations therein.

[35] The analysis by Eckardt (2006) suggests a greater similarity to the progressive form than I assume for this comparison. More concretely, Eckardt (2006) assumes an analysis

Thus, it could be interpreted either as a relative tense or a prospective aspect (which are formally similar in the present system; (Zymla, 2017)). I do not want to extend this discussion here. However, what remains is the fact that there is some element in the expression in (95a) that picks up a contextually salient time point. A property that we attribute to grammatical tense. The same observation can be made for the perfect aspect, which can pick up either a present or a past time point. This is explicitly shown in the examples below.

(97) a. Dia sudah pergi (sekarang). (E $\prec$ S,R)
3.SG PERF go now
'S/he has left (now).'

b. Dia sudah pergi kemarin (E $\prec$ R $\prec$ S)
3.SG PERF go yesterday
'S/he (had) already left yesterday.' (Arka, 2013, p. 26)

I conclude from these examples that clauses in Indonesian carry a temporal trace that does not necessarily coincide with the utterance time. In other words, Indonesian instances of temporal reference provide a reference time that may be contextually or explicitly restricted. This is confirmed in Arka (2013). The challenge here is that the temporal anchor is naturally compatible with any of the three basic tenses; thus, Indonesian verbal forms can only ever entail unspecified temporal reference. The question is then how to describe the relationship between form and meaning in terms of the rule system I have proposed above. There is seemingly no morphosyntactic exponent which we could faithfully map onto the semantics. In such cases, the mapping relies somewhat on the morphosyntactic input. I propose a technical solution based on the architecture of the Indonesian XLE grammar Arka (2012): The relevant verbal forms that receive a tensed interpretation are exactly those that are marked for `MOOD` in the XLE grammar. Ascribing the temporal

that is more in line with the preparatory progressive (see, e.g., Rivero and Arregui (2010)) rather than a temporal relation. Concretely, this means that the *going to* construction relates a preparatory event stage with the reference time rather than relating the proper eventuality time with the reference time (Aloni and Dekker, 2016). I will show in the next chapter that such analyses are also possible in the present annotation scheme.

interpretation to the PRED value would imply that the temporal interpretation belongs to the lexical semantics of the verb. However, verbal forms can occur both in tensed and non-tensed environments. Consequently, the following rule is the base rule for tense and aspect interpretation in Indonesian:

(98) MOOD indicative → TEMP-REF 'unspec'

However, this rule does not tell us anything except that there is a tense node that modifies the underlying eventuality. Explicit specification through lexical items or pragmatic reasoning is necessary to resolve the underspecified tense value as illustrated in (97) and (95a), respectively. Aspect is not marked obligatorily in Indonesian either. Thus it is not possible to draw inferences from the aspectual markers available in Indonesian in many cases either.
Consider again example (95a) repeated in (99a). Without the specified context, this sentence has the two readings paraphrased in (99b) and (99c). This suggests, at least, that contextual temporal reference is constrained to non-future contexts (i.e., it is not possible to produce a future-shifted based on a reference time in the future). Of course, such syntactic interactions have to be investigated for each auxiliary separately.

(99) a. Peter akan menulis se.buah surat
Peter FUT AV.write one.CL letter
b. Tom will write a letter.
c. Tom was going to write a letter.

In the present case, the interpretation seems to rely on the existence of a salient non-future time. In Zymla (2017), I entertain the idea that the future marker *akan* is a reflex of a existing temporal relation in the context. If a clause in Indonesian would simply pick up some reference time from the context, then the use of the future tense appears odd since no future temporal variable is salient. What is salient is a past time variable – introduced by *did do* – that is in the future relative to some other past variable, namely the *dinner*. Thus, *akan* does not refer to an anaphoric element alone but also to the salient temporal relations that are available with respect to that anaphoric element. While the absolute temporal allocation of a temporal variable is

always available, this is not always the case for relative temporal allocation. One could argue along the lines of Grice (1975) that the future marker is picked up as a direct consequence of the maxim of relevance and quantity since it identifies better with the anaphoric temporal element introduced by the question in (95) which is also allocated in time via a posteriority relation. This claim puts a heavy load on the theoretic assumptions that underly the resolution of temporal anaphora and anaphora in general. In Zymla (2017), I argued for an approach based on situation semantics (although it might be easier to give up variables altogether and entertain a conceptual approach as in Bobrow et al. (2007), Bobrow et al. (2009), and Crouch and King (2006)). The basic idea was that anaphora resolution would work via conceptual similarity. In this framework, verbal expressions refer to antecedent situations with similar semantic properties. Thus, a situation that is described as being temporally preceded by another situation is referred to as a situation that is temporally preceded by another situation as illustrated in (95a).
Although this goes beyond the scope of this thesis, I would suggest analyzing inter-sentential phenomena such as the one discussed above in the context of discourse representation theory, following a rich literature (Kamp, 2017; Lascarides and Asher, 1993; Partee, 1984). Consequently, pragmatic reasoning still requires additional attention and is left for future work. However, an important takeaway of this section is that zero-marked tenses are affected by the temporal and aspectual auxiliaries in the context, as has been shown for *akan*, which rules out contextual future temporal reference concerning the reference time.

3.5 Summary

All in all, I have described the three steps necessary to represent the morphosyntactic and semantic properties of language adequately. i) I have investigated the empirical landscape to understand the types of variation that may occur on various levels of grammar. ii) I have explored the formal literature and have identified crucial notions in distinguishing different kinds of linguistic phenomena. iii) Ultimately, I put forward annotations for various phenomena based on empirical and formal insights I explored up to this point. I have introduced a temporal annotation scheme that mimics formal semantic thought processes while simplifying or dropping many theory-driven assumptions. In fact, I have proposed semantic labels that may be mapped onto different semantic formalisms depending on the requirements of the application. The annotation scheme is not fully independent of theory-driven approaches and occasionally requires executive decisions in annotating semantic phenomena and the syntax/semantics interface. However, it is important to note that the annotation system is very powerful and can implement various formal analyses while simplifying each of them more or less considerably.

One of the main benefits of the present annotation scheme is the strict modularization of syntactic and semantic features. The idea of mapping forms to meanings, even in trivial cases, allows us to explore how far semantic meaning may take precedence over structural representations in terms of identifying cross-linguistic generalizations. To this end, we first have to identify the relations between form and meaning and then test if the same semantics – based on different form-to-meaning mappings – allow for the same generalizations within and across languages. Ultimately, we also want to be able to quantify these generalizations, which is part of the role of linguistic annotation. The role of the present annotation scheme in this machinery is to ascertain the soundness of the semantic side by embedding it in a proper formal semantics (see also Chapter 5).

Chapter 4

Beyond tense: grammatical and lexical aspect

> The rabbit-hole went straight on like a tunnel for some way, and then dipped suddenly down, so suddenly that Alice had not a moment to think about stopping herself before she found herself falling down what seemed to be a very deep well.
>
> Lewis Caroll, *Alice in Wonderland*

In the previous chapter, I provided a cross-linguistic examination of tense and, more generally, temporal reference to provide a general motivation for the annotation scheme presented in this thesis. This foundation has led to the development of a cross-linguistically viable semantic feature space for annotating temporal configurations as proposed by (Reichenbach, 1947). Based on this feature space, a formal semantics has been worked out, grounded in cross-linguistically motivated formal semantic research, that accommodates the possible features and their interactions. The formal semantics assure the proper interaction of the features within the annotation scheme. Furthermore, I have defined a procedure for developing rules describing a syntax/semantics mapping and semantic construction for specific semantic interactions.

In this chapter, I extend the ParTMA annotation framework to incorporate the categories of viewpoint and inner aspect as well. In the previous chapter, I have already shown that grammatical and inner aspect carry certain entailments and implications with respect to temporal reference. For example, a bare perfective marker usually provides past temporal reference. In contrast, perfective + present temporal reference encodes future temporal reference due to the incompatibility between the instantaneous time interval corresponding to the speech time (or the *now*) and protracted time intervals bounded by perfective aspect. This is exemplified again for Polish in (1).

(1) Peter **na-pis-ze** list
Peter **PRV-write-PRES** a letter
'Peter is going to write a letter (completely)' Polish

In this chapter, the main focus lies on two other layers of meanings. One of these layers is associated with the modal location of eventualities. This is shown in (2): the expression *drawing a circle* alludes to a possible view of the world (*w'*) in which the described eventuality has happened in its entirety (as indicated by the use of the simple past)[1]. One way to establish this intuition is by comparing the resulting state of the respective worlds. In *w'*, *res'* entails the existence of a circle in that world. On the other hand, *res* in w^0 does not commit to the existence of a circle (although it does not deny it either). The contribution of the possible world, *w'*, allows us to explain why we describe an arbitrary *drawing* eventuality as a *drawing a circle* eventuality. It shows a world that describes our expected continuation of the eventuality in w^0.

(2) a. Jordan was drawing a circle.

b. Temporal and modal progression diagram:

—	***drew a circle***	*res'*	—	—$>_{w'}$
	$\supseteq$			
—	<u>*drawing (a circle)*</u>	*res*	*S*	—$>_{w^0}$

As indicated by the progression chart, this alternative world does not exist at an arbitrary point in time but rather at the same point in time in a belief

[1]Usually, modal approaches to the imperfective universally quantify over all suitable possible worlds. *w'* can be understood as a prototypical world representing the relevant set.

	durative	change	endpoint
state	+	-	-
activity	+	+	-
achievement	-	+	+
accomplishment	+	+	+

Table 4.1: Aktionsart

world.[2] This means that we generally expect two dimensions to be relevant for modeling temporal reference and grammatical aspect: the temporal dimension and the modal dimension.[3] The former has already been discussed in detail in the previous chapter.

The second layer of meaning associated with aspectual semantics is concerned with the internal makeup of an eventuality. In chapter 2, I used table 4.1 based on Vendler (1957) to explain inner aspect (see section 2.6.3). Thus, its representation is more multidimensional.[4] These dimensions can vary across languages (Bittner, 2014; Croft, 2012), as will be shown in this chapter.

Table 4.1 classifies verb phrases according to various *verb classes*. In this chapter, I propose a more fine-grained distinction that includes other aspects of the lexical semantics of verbs, particularly argument structure (Dowty, 1991; Hale and Keyser, 1993; Levin and Hovav, 2005; Ramchand, 2008). The

[2]One might argue that belief worlds need not overlap temporally with the actual world (This becomes more apparent when explicitly involving belief worlds by use of propositional attitude verbs; Von Stechow 2009). As explained in the previous chapter, I assume that temporal reference does not happen based on absolute temporal values (such as dates and times) but is based on relative anchoring. Here, I extend this view. Temporal anchors – such as E, R, and S – of belief worlds are mapped onto their counterparts in the actual world.

This treatment does not rule out a more fine-grained modeling of the temporal properties of E, R, and S, but it allows us to omit temporal coordinates in belief worlds when not mentioned explicitly. Since the modeling of belief worlds is limited in scope in the first place, as I will show in this chapter, I believe this to be a reasonable treatment of the issue.

[3]This is also reflected in many analyses of the progressive in formal semantics. They often use world/time pairs rather than only worlds or only times (Dowty, 1986) (There also exists a variety of approaches employing situations that can have both modal and temporal coordinates, e.g., Arregui et al. (2011) and Cipria and Roberts (2000)).

[4]Croft (2012) proposes a two-dimensional model of inner aspect. He uses the dimensions of time and quality. As indicated by Vendler's (1957) linguistic tests for inner aspect, time seems to be an obvious choice. However, the dimension of quality is quite opaque.

resulting semantics also considers a pre- and post-state.

More generally, the treatment of verb classes I propose for this chapter falls under the label of decompositional event semantics (Pustejovsky, 1991; Ramchand, 2008). The goal of this enterprise is to formally compose eventuality descriptions based on the available sub-events. For example, activities, achievements, and accomplishments must have a *process* subevent. as already mentioned back in section 2.6.3, this can be tested by applying *progressive* marking to the VP in question, since it requires the *process property* (Portner, 2019). This insight is obviously covered by Vendler's (1957) distinctions. However, what we can deduce from his classification is the relation between the aspectual property of having a *process* sub-eventuality and argument structure. Following Ramchand (2008), I treat the two notions as semantically connected. Thus, an eventuality can (generally) only have a *process* subevent, if we can also identify an *undergoer*, i.e., someone who is affected by the eventuality (see (3a); the index i maps the argument onto the relevant subevent). Similarly, to have an *initiation* subevent, an *initiator* is required, as illustrated in (3b).

(3) a. The ball$_i$ rolled. $process_i$

b. Jamie$_i$ rolled the ball$_j$. $initiation_i, process_j$

Verb classes are then defined based on different configurations of the participants of an eventuality. This is described in detail in section 4.3.2, where I also include a discussion of *result* states already hinted at in the discussion of (2). While the focus of this thesis does not lie on the semantics of argument structure, this approach allows for the description of subtle interactions between syntax and semantics.

In this chapter, I explore these categories in a fashion similar to the previous chapter. Temporal ordering will become less relevant, and the focus will shift toward the question of how aspectually marked clauses relate to the actual world in terms of internal properties of the event. As Vendler's (1957) tests show, these properties affect temporal interpretation more or less directly but can not always be described by virtue of interval semantics alone. As explained above, a modal dimension or something similar is also required.

For this reason, there will also be a shift in the semantics in this chapter. Rather than working with time intervals, this chapter introduces situation semantics to describe the semantics of viewpoint aspect, and event semantics to describe the semantic contributions of inner aspect.[5]
Starting in section 4.1, this chapter follows the same structure as chapter 3: By looking at the cross-linguistic variation, I provide a general picture of the spectrum of marking grammatical and inner aspect and the semantics of these categories. The corresponding formal framework is developed in sections 4.3 and 4.4. This includes both morphosyntactic considerations, as well as the semantic properties of these categories. The implications of the formal framework for implementing aspectual semantics within the ParTMA annotation scheme are discussed in section 4.5.

4.1 Cross-linguistic variation of grammatical aspect

This section provides an intuitive description of grammatical aspect, focusing on its realization across languages as documented within the ParGram project but also in the linguistic literature more generally. As in the previous chapter, this part is structured based on different morphosyntactic forms, starting with grammatical aspect via derivational morphology.

4.1.1 Grammatical aspect via derivational morphology

The main distinction pertaining to grammatical aspect is that between perfective and imperfective grammatical aspect. The terms *perfective* and *imperfective* have been coined by research on aspectual systems in Slavic languages. In Slavic languages, verbs form aspectual pairs, with a perfective and an imperfective variant. Consider the Russian examples in (4) and (5) from Romanova (2007). They display infinitival forms of some Russian verbs. As these examples suggest, perfective aspect is often derived from

[5]This shift does not affect any of the insights gained in the previous chapter. Time intervals can be thought of as situations that only contain information about the time interval that they correspond to. Conversely, situations CAN have temporal properties that make them susceptible to temporal operators. More details are given in section 4.4.

imperfective aspect via prefixation.

(4) za-pisatj, pro-citatj, po-sidetj, vy-ucitj
in-write^P.INF PRF-read^P.INF DEL-sit^P.INF out-learn^P.INF
'write down, read completely, sit for a while, learn (by heart)'
perfective

(5) pisatj, citatj, sidetj, ucitj
write^I.INF read^I.INF sit^I.INF learn^I.INF
'write, read, sit, learn' *imperfective*

The distinction between perfective and imperfective verbs usually encodes a distinction between a bounded (e.g., *po-sidetj: sit for a while*) eventuality or an eventuality with an inherent endpoint (e.g., *pro-citatj: citaj*) (Comrie, 1976; Dahl, 1985). Although this provides a fairly reasonable classification of the aspectual opposition at the lexical level, the system is far more complex than this classification suggests. The infinitival forms in (4) are subject to various kinds of (further) aspectual marking. Concretely, the Slavic aspectual system allows for the stacking of various affixes, as illustrated in (6).[6] As shown there, a given verb can alternate between perfective and imperfective interpretation based on the affixes employed. As the glossing in (6) suggests, the stacking of aspectual morphemes does not simply alternate between perfective and imperfective aspect but introduces more fine-grained distinctions such as, for example, iterativity.

(6) a. pf. *dat'* 'to give'
b. pf. *dat'* > ipf. *da-va-t'* 'to give'
c. ipf. *vy-da-va-t'* > pf. *po-vy-da-va-t* 'to hand out, distribute'
(Gehrke, 2008)

Grammatical aspect in Russian is tied to a complex system of affixation. This system is complicated further by the observation that Slavic aspectual systems seem to contain two layers of aspectual modification: lexical and

[6] As explained in the last chapter, this system interacts with an inflectional tense system in various Slavic languages (e.g., Polish, and Russian).

superlexical modification, both of which are associated with the morphosyntax of the perfective/imperfective opposition (Gehrke, 2008; Ramchand, 2004; Romanova, 2007).[7] This is shown in example (7) from Gehrke (2008, p. 162). There, the internal prefix changes the semantics of the verb. Generally, it is argued that internal affixes affect the lexical semantics and "participate in structuring the event" (Gehrke, 2008, p. 162). External affixes serve as modifiers (for the temporal properties) of the event description they are applied to. For example, the external affixes in (7b) mark a temporally bounded event and the inception of an event respectively.

(7) *Perfective counterparts of (Russian) ipf. pisat' 'to write'*

a. pf. *pod-pisat'* 'to sign' (lit.: under-write) — INTERNAL

b. pf. *po-pisat'* 'to write (for a while)' — EXTERNALS
   pf. *za-pisat'* 'to (begin to) write'

(Gehrke, 2008)

To better understand the argument structure modifying effect of lexical affixes, consider the following example from Ramchand (2004, p. 337). It illustrates how lexical or internal modification may affect the argument structure of a given verb while also perfectivizing it (as indicated by the superscript P). In each of these cases, the lexically modifying prefix adds an object to the argument structure, which is unavailable with the corresponding imperfective form (shown on the right of each example and marked with a superscript I).

(8) a. v-rezat' zamok v dverj — (*rezat' zamok)
into-cut^P *lock*.ACC *in door*.ACC — *cut*^{I} *lock*.ACC
'insert a lock into a door'

b. vy-bit' glaz — (*bit' glaz)
out-beat^P *eye*.ACC — *beat*^{I} *eye*.ACC
'hit an eye out'

c. pro-gryzt' dyru — (*gryzt' dyru)
through-gnaw^P *hole*.ACC — *gnaw*^{I} *hole*.ACC
'gnaw a hole in something'

[7] Called internal and external modification in Gehrke (2008)

(Ramchand, 2004)

The main caveat on the Russian aspectual system suggested by Romanova (2007) is, that there is "no uniform morphological indication of perfective or imperfective aspect" (Romanova, 2007, p.3). This means that in Russian grammatical aspect is not a morphological category but is derived via additional morphosyntactic and semantic aspects. This is, for example, reflected in the fact that affixes in Slavic languages can be stacked and is, furthermore, underlined by the distinction between internal and external modifiers. This indicates that aspectual modifiers interact not only with different levels of aspectual interpretation (i.e., grammatical and lexical aspect) but also with the syntactic and semantic context, as has been exemplified in terms of the potential modification of argument structure.

As the examples above illustrate, the distinction between perfective and imperfective grammatical aspect is far from clear-cut. The fundamental distinction is between bounded or completed forms and unbounded forms, but different configurations of aspectual affixes can affect various semantic properties of the described eventuality, including the available participants. We have also seen that there is a mapping from properties of grammatical aspect (i.e., boundedness) to external modification and a mapping from internal modification to properties of lexical semantics (e.g., telicity). However, so far, we have no concrete proposal for isolating the relevant semantic features beyond the basic distinctions made in the introduction to this chapter and the more general introduction in 3. In other words, the contributions of markers of grammatical aspect, at least in Russian, are not always as transparent as the contribution of tenses (even in simple sentences). This is due to the many subtle interactions between bare verb forms and internal and external modification.

In the following few sections, I enrich the cross-linguistic picture on grammatical aspect. This allows us to get a better understanding of the relevant semantic features. However, as already hinted, we will see an even stronger interplay between grammatical aspect and inner aspect features compared to their respective interactions with temporal reference.

4.1.2 Grammatical aspect via auxiliaries

In the previous chapter, we have seen that Germanic tense/aspect systems heavily rely on the use of auxiliaries to encode temporal and aspectual distinctions. This section examines several examples of aspectual modification, mainly in Germanic languages, to exemplify aspectual auxiliaries.

In chapter 3, I have already introduced the auxiliary systems of German and English. There, I have shown that auxiliaries follow a fixed order that roughly moves from temporal to aspectual modification. The relevant auxiliary constructions in English were primarily the perfect construction ($Aux + V_{past\_participle}$) and the progressive construction ($Aux + V_{present\_participle}$). In previous mentions, I have labeled the progressive an instance of imperfective aspect. However, the perfect is not an indicator of perfective aspect (Condoravdi and Deo, 2014). In fact, perfective aspect is not encoded periphrastically in English and usually also not in other languages that employ auxiliaries to encode (certain) aspectual distinctions. Particularly the progressive aspect regularly emerges in terms of an aspectual auxiliary cross-linguistically (Dahl, 1985). The following examples illustrate this.

(9) John was building a boat.

(10) vo nɑõ bənɑ rəhɑ hæ.
he boat.NOM make PROG.M.SG be.PRS.SG
'He is building a boat.' *Urdu*

(11) kare wa fune o tsukut-te imasu
He SUBJ boat OBJ build-PRSPRT be.NON-PAST
'He is building a boat.' *Japanese*

(12) Dia sedang membuat perahu
3.SG PROG build boat
'He is building a boat.' *Indonesian*

In the examples above, English and Japanese form the progressive with a present participle form and an auxiliary carrying temporal information. Similarly, Urdu combines the infinitival main verb with *rahaa* (derived from *(to) stay*) and a temporally marked auxiliary. Due to the lack of overt tense marking in Indonesian, the progressive marker *sedang* that serves as an aspectual auxiliary does not carry any temporal information. Furthermore, the main verb does not have a dedicated progressive form (such as, e.g., the present participle in English). Thus, the Indonesian progressive marker is similar to a lexicalized aspectual marker such as *gerade* in German.[8]

(13) Er bau-t gerade ein Boot.
3.SG build-PRES PROG a boat.
'He is building a boat.' *German*

4.1.3 Grammatical aspect via inflectional categories

As already shown before (e.g., in chapter 3), certain languages tend to combine temporal and aspectual information in the same morphosyntactic form. In this case, grammatical aspect seems more likely to occur as part of an inflectional category.

This tendency is more developed (but not exclusively) in Romance languages such as Spanish, Italian, and French. Tense and grammatical aspect are also more intertwined in Modern Greek, for example. These languages provide a dedicated *past imperfective* marker such as the *imperfetto* in Italian. Italian furthermore distinguishes between a progressive marker and the *imperfetto*, which is used for general imperfective uses. This is illustrated in example (14) from Giorgi and Pianesi (1997). Following the trend we have seen before, the Progressive is formed using an auxiliary verb combined with the gerund of the main verb (however, the present participle is distinct from the gerund form in Italian).

(14) a. Quanto Artù entrò, Merlino creava un unicorno.

[8]The claim that *gerade* is a progressive marker, might be to strong (Ebert, 2000). As Schaden (2010) notes, *gerade* has several temporal uses. Most prominently, an immediate past reading, when combined with the perfect:

(1) Peter ha-t gerade seinen Helm aufgesetzt.
Peter have-PRES just his helmet put on.PSTPTCP
'Peter has just put on his helmet.' *German*

When Arthuer entered, Merlin created(IMPV) a unicorn.

b. Quanto Artù entrò, Merlino stava creando un unicorno.
When Arthuer entered, Merlin was creating(PROG IMPV) a unicorn.
(Giorgi and Pianesi, 1997)

General imperfective uses are illustrated in the following example from Ferreira (2016). As shown there, pure imperfective marking is ambiguous between a habitual interpretation and a continuous interpretation, as indicated by the glossing.

(15) a. Eperne to farmako
take-PAST-IMPV the medicine
'He was taking the medicine/He used to take the medicine.' *Greek*

b. Gianni fumava
Gianni smoked-IMPV
'Gianni was smoking/Gianni used to smoke.' *Italian*
Ferreira (2016)

The continuous interpretation of the imperfective aspect is relatively similar to the progressive in that both describe an unbounded eventuality. However, in Italian, the imperfective typically patterns with activities (such as smoking), whereas the progressive patterns with achievements and accomplishments, i.e., telic predicates to express the ongoingness of a given eventuality.
Generally, habituality is expressed when one describes the tendency of someone (or something) to do something regularly (Carlson, 2012). However, the habitual does not describe a state of affairs at a specific reference time but instead in a broader time frame. The next section expands on this explanation.

4.1.4 Contrastive inference of aspectual features

Inflectional forms are often used to form a contrast with other aspectual forms (a similar process has been identified for some zero-marked tenses in the previous chapter). As we have seen above, only the progressive aspect is marked explicitly in English. It stands in contrast with the simple present and the simple past. However, both of these have default aspectual interpretations. The realization of habitual aspect in English is indicated by the aspectually unmarked present tense and (optionally) plural marking on the object. Thus, (16c) represents a prototypical habitual interpretation in English. Ferreira (2016) points out that the plural marking is important for the habitual interpretation. However, the plurality does not necessarily have to come from the *object*. Thus, (16d) can also encode habitual aspect if the underlying event is separated in multiple non-continuous instances. Importantly, none of these instances has to overlap with the evaluation time.[9] Thus, both (16c) and (16d) are true even if Jamie is not engaged in the activity described by the verb phrase. Habituality can also be expressed by combining the past tense with plural marking on the object. However, as explained in the next few paragraphs, the simple past has a default perfective interpretation. Thus, in the past tense, there exist explicit lexical markers of habituality, e.g., the phrase *used to*. As these examples suggest, aspectual marking is much more context-dependent compared to temporal reference, at least in English.

(16) a. Jamie wrote a letter. *perfective*

b. Jamie is writing a letter. *imperfective/progressive*

c. Jamie writes letters. *imperfective/habitual*

d. Jamie builds a house. *imperfective/habitual*

e. Jamie used to write letters. *imperfective/habitual*

Compared to imperfective aspect, perfective aspect is generally incompatible with a present tense interpretation (De Wit, 2016). Thus, the English simple past is the main contender for perfective marking in English. However, as mentioned above, the simple past can receive a habitual interpretation in the right context. Thus, it would be more precise to say that the simple past can express perfective aspect in a certain morphosyntactic and semantic context. In other words, there is no concrete morphosyntactic marker for perfective in English, but the default interpretation of the simple past is perfective.

[9] This is a tentative description of habitual semantics.

4.1.5 Zero-marked grammatical aspect

Zero-marked grammatical aspect is arguably more prevalent than zero-marked tenses. I have already pointed out a superficial similarity between German in Indonesian with respect to progressive marking (see section 4.1.2). Following this thread, we can say that both German and Indonesian allow for inferring aspectual properties contextually or from the interaction with temporal modifiers, as illustrated in (17).

(17) Dia makan sambil menonton TV.
3.SG eat while AV.watch TV
'He was/is eating while watching TV.' *Indonesian*

(18) Er aß während er fernsah.
3.SG eat.PAST while 3sg watch.PAST TV
'He was eating while watching TV' *German*

Thus, these languages provide more extreme examples of contextual aspectual reference compared to English, where the aspectual semantics are inferred from the contrast with other, overtly marked aspectual categories. In the previous chapter, I mentioned that the optional marking of temporal reference suggests additional implications. However, I am not aware of any such implications associated with the markers presented here. However, this is not conclusive evidence. Thus it provides another avenue for future research.

4.1.6 Complex interactions of aspectual features

Already in the discussion of Slavic-style aspectual marking, I have pointed out that aspectual marking can be stacked. When stacked, the semantics of such expressions do not simply switch from bounded to unbounded (or vice versa). Rather complex stacking of aspectual features affects different semantic properties between grammatical aspect and inner aspect. One example of this was *iterativity*, which also emerges when combining grammatical aspect with plural or pluractional markers. Pluractionality marks the repeated occurrence of an action, for example, via reduplication as shown in (20); Lasersohn (1995). As the example indicates, in this case, the progressive marker *sedang* is optional to express progressive aspect (Arka, 2013).

(19) John is building boats.

(20) Ali (sedang) me-mukul-i kepala=nya sendiri.
Ali PROG AV.hit-I head=3SG.POSS self
'Ali is beating his own head.' *Indonesian*

Various languages in the ParGram project provide avenues for further research on complex interactions between aspectual features. However, while the interaction with plural and pluractional markers is relatively straightforward cross-linguistically, interactions between aspectual markers can induce more fine-grained semantic distinctions. One noteworthy example is that of Urdu. We have seen that aside from the progressive form that is derived from the verb stay (Butt and Rizvi, 2010), there is a second similar form that allows for aspectual stacking (see (21)).

(21) a. vo nɑõ bənɑt̪ɑ rəhɑ.
he boat.NOM build-IMPV stay-PERF
'There is a boat building event which continued past a point at which it was supposed to end.'

b. vo nɑõ bənɑt̪ɑ rɛht̪ɑ hæ.
he boat.NOM make-IMPV stay-IMPV be.PRS.SG
'He is "obsessively" building boats' *Urdu*

As these examples show, aspectual stacking can lead to highly specific readings. For the sake of this thesis, I do not inspect such examples more closely. However, I believe that the syntax/semantics interface introduced in the previous chapter is especially well-suited to capture such interactions. A more concrete investigation of this is left for future work. In this thesis, I mainly focus on the more primary semantics associated with grammatical and inner aspect since they are also more well-represented cross-linguistically. Furthermore, it is unclear how well the meanings in (21) contribute to the resolution of basic inference patterns such as the imperfective paradox.

All in all, we see a similar general pattern as with temporal reference in the previous chapter: Aspectual categories can be encoded on a spectrum from structurally to contextually encoded meanings. Similar to complex tense/aspect forms such as the past perfect, aspectual features can be stacked, but rather than encoding some sort of relativity, the combination of aspectual features allows for the expression of more fine-grained and varied meanings. This trend becomes even more apparent in the next section, where I discuss the encoding of inner aspect cross-linguistically.

4.2 Cross-linguistic variation of inner aspect

Although aspectual marking allows for encoding various complex meanings, it is still possible to attribute certain features to specific forms. The semantic contribution of inner aspect is less specific. In other words, the semantic properties that affect the linguistic surface representation of eventualities concerning their inner aspect are less clear-cut. In chapter 2, I introduced some linguistic tests for inner aspect, such as the (non-)compatibility with the progressive aspect or the combination with specific temporal adverbials (i.e., the *in/for* alternation). Such tests indirectly elicit information about the inner aspect of a given verb phrase. By this, I mean that the surface form of the eventuality does not change corresponding to the inner aspectual properties but only to the constraints on modification as imposed by these tests (e.g., progressive aspect is formed in the same way for all verb classes compatible with it). However, specific properties related to inner aspect can be illustrated in terms of cross-linguistic variation, particularly when arguments of the verb are taken into consideration as well. Consider the verb *(to) walk*, which seems to be underspecified for telicity, as it is combinable with different arguments to produce either a telic or atelic interpretation. This is illustrated in the following example.

(22) a. Jordan walked to the store (?for thirty minutes/in thirty minutes)
  b. Jordan walked around in the park (for thirty minutes/*in thirty minutes)

However, the intransitive version of the verb, as, e.g., in *Jordan walked*, seems to be strictly atelic. This suggests that argument selection is tightly related to inner aspect and that we cannot simply apply labels of inner aspect to only the verb itself but also need to factor in its immediate syntactic and semantic context (e.g., Dowty 1982; Krifka 1998; Verkuyl 1972; Verkuyl 1993). We also need to consider the internal makeup of these arguments, which is illustrated particularly well by the distinct properties of the PPs attached in (22). The first PP, *to the store*, clearly expresses a specific goal which effectively suggests the existence of a clearly specified endpoint. On the other hand, the PP *around in the park* describes an undirected path with no linguistically specified goal, thus, implying atelicity.

In other languages, for example, Hungarian, telic and atelic readings are sometimes grammatically distinguished by using particles. This is illustrated in (23b), where the particle *le* indicates that the result stage is reached (the door is completely painted). However, Hungarian also may introduce other means by which telicity is distinguished, as is shown in (23c). There, the adjectival modifier forces a telic interpretation where the natural endpoint is a completely white door (Kardos, 2016).

(23) a. Péter 10 perc-ig /*10 perc alatt festett egy ajitó-t
P.NOM 10 minute-for /*10 minute under painted a door-ACC

'Peter was painting/painted a door for 10 minutes' *atelic*

b. Péter 10 perc alatt /*10 perc-ig le-festett egy ajitó-t
P.NOM 10 minute under /*10 minute-for PRT-painted a door-ACC

'Peter painted a door in 10 minutes.' *telic*

c. Péter 10 perc alatt /*10 perc-ig fehér-re festett egy ajitó-t
P.NOM 10 minute under /*10 minute-for white-into painted a door-ACC

'Peter painted a door white in ten minutes' *telic*

(Kardos, 2016)

Hungarian

4.2.1 The question of cross-linguistic uniformity

However, the classification of Aktionsart or aspectual class runs into further problems from a cross-linguistic perspective. Not because they do not occur in other languages but because different classes may be specified based on how they interact with the grammar. This point is, for example, elaborated on in Bittner (2014), who proposes a new set of aspectual classes for Chinese.[10] The classes and their diagnostics are shown in (24) and (25) respectively.

(24) **Chinese verb classes**

a. predicates of *events*

b. predicates of *gradable states*

c. predicates of *point events*

d. predicates of *point-scale states*

(25) **Diagnostics**

a. Compatibility with measure phrases

b. Compatibility with phase verbs

c. Compatibility with degree modifiers

Similarly to Vendler's (1957) classification, verb phrases can be distributed across four different classes according to tests based on the compatibility with different types of modifiers. However, the resulting distinctions are quite different. The following examples from Bittner (2014) illustrate the application of the tests to examples from Chinese. The first test involves the compatibility with event-measure phrases such as, for example, *three times*. This is only available for events (marked by E/) in the glossing. Consequently, states can not be pluralized in Chinese.

(26) [±*event*]: (in)compatibility with event-measure *cì* 'M$_{evt}$'

a. Mingzi, Lisi xie san-ci le
name, Lisi E/write three-M$_{evt}$ PNC

[10] As Bittner (2014) herself points out, there are several proposals that provide an alternative classification of verb classes. This only serves as an illustrative example.

‘His name, Lisi wrote three times.’ event

b. Mingzi, Lisi xie.wan san-ci le
name, Lisi E/write.finish three-M$_{evt}$ PNC
‘His name, Lisi finished writing three times.’ point event

c. Lisi *hen lei san-cì (le)
Lisi POS S/tired three-M$_{evt}$ (PNC)
gradable state

d. Lisi lèi.si san-cì (le)
Lisi S/tired.die three-M$_{evt}$ (PNC)
point-scale states

Chinese

The second test concerns the compatibility with phase verbs. According to Bittner (2014), the item *zài* can be interpreted as a progressive marker. As predicted, this marker is compatible with events (note that the progressive marker makes all predications it applies to stative. Thus, the examples in (27) do not provide any E/ annotations). However, it is not compatible with point events.[11] It also is not compatible with the two kinds of states identified by Bittner (2014).

(27) [$\pm$*point* $-$ *event*]: (in)compatibility with phase-verbs, e.g. *zài* ‘S/be.in.prg’

a. Lisi zài xie mingzi
Lisi S/be.in.prg write name
‘His name, Lisi wrote three times.’ event

b. Lisi zài *xie.wán mingzi
Lisi S/be.in.prg *write.finish name
INTENDED: ‘Lisi is finishing writing his name’ point event

c. Lisi *zài hen lèi
Lisi S/be.in.prg POS S/tired
gradable state

[11]Similarly, not all achievements in English are readily compatible with the progressive,e.g., *?Jamie was finding her keys.*

	measure phrase	phase verb	degree modifier
event	+	+	-
point event	+	-	-
gradable state	-	-	+
point-scale state	-	-	-

Table 4.2: Alternative classification of Chinese inner aspect

d. Lisi *zài lèi.si
Lisi S/be.in.prg S/tired.die

point-scale states

Chinese

Ultimately, Bittner's (2014) tests lead to the classification in Table 4.2. Although her approach and Vendler's classification do not seem to be too far apart, it challenges the idea of a cross-linguistically valid treatment of inner aspect. Not necessarily because eventualities are perceived fundamentally different in Chinese, but rather because the perception of eventualities focuses on different properties, which leads to different interactions with other elements in the grammar.

Bittner (2014) raises a critical question: Do we simply try to fit Chinese into a framework that is mainly based on research in Germanic languages, or should we specify a separate aspectual system based on rigorous linguistic testing? In the next section, I propose a uniform system for annotating inner aspect based on Ramchand (2018) and Ramchand (2008). The gist of this system is that it provides more basic building blocks for inner aspect via a decompositional event semantics. Thus, the insights gained here provide an overview of the topic of inner aspect but can be broken down more for the sake of formalization.

4.2.2 Interim summary

The main point of this section is to show that semantic properties of inner aspect usually emerge based on the possible interactions with other elements in the grammar rather than some sort of grammatical marking. Concretely, aspectual properties arise from the interaction of properties of the verb itself and the interaction with its internal arguments. However, some languages furthermore allow for different strategies for indicating aspectual properties. For example, it was shown for Hungarian that telicity can be encoded in terms of particles. However, the exact aspectual distinctions might vary cross-linguistically. An examination of this is potentially left for future work. Compared to the other categories, this section is fairly short. This is in part due to the language-specific nature of inner aspect. However, rather than losing myself in minute details of aspectual interpretations across languages, this thesis focuses on the bigger cross-linguistic picture. Thus, in the next two sections, I work out a formal semantics that encodes distinctions of inner aspect and explain how this framework can potentially be extended to more difficult cases, such as the Chinese aspectual system.

4.3 Towards a formalization

The formalization of grammatical and inner aspect provides a more complex challenge than temporal reference. Particularly, the many interactions between grammatical aspect and inner aspect enhance the difficulty of finding a cross-linguistically common core. In this section, my goal is to line out concrete distinctions between semantic properties contributed by grammatical aspect and semantic properties contributed by inner aspect. More concretely, this section is concerned with the question of whether it is possible to identify some basic building blocks for annotating grammatical and inner aspect.

4.3.1 Grammatical aspect

As illustrated in the previous section, grammatical aspect, at the very least, forms a dual distinction. This distinction has been traditionally based on the Slavic aspectual system, where verb forms are categorized into two major categories (Romanova, 2007): perfective and imperfective aspect. This view is also perpetuated in the distinction by Comrie (1976), a more cross-linguistic study of aspect, shown in (28):

(28) Aspectual oppositions by Comrie (1976)

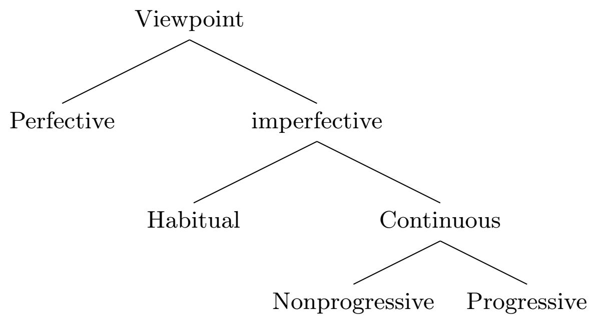

In the previous section, I have shown that each of these distinctions has been attested cross-linguistically. As the hierarchical structure of (28) indicates, the progressive and the habitual are more specific readings of the imperfective aspect. Accordingly, we would want our semantics to capture this intuition. More concretely, our goal is to formalize imperfective aspect so that a pure imperfective can receive either a habitual or a progressive interpretation, or rather a habitual or continuous interpretation. However, the continuous interpretation of an imperfective marker can also be a progressive interpretation. According to Comrie (1976), the reverse should not be possible.[12] I.e., it should not be possible to interpret a habitual form as a progressive form. In English, this seems to be, in fact, the case.

---

[12]A possible challenge to this view comes from other uses of the present participle form, e.g.:

(1) Once there was a scorpion, *standing by a river*. adjunct

(2) *Hunting deer* is a difficult endeavor. gerund

The adjunct clause indicates a continuous interpretation. This is acceptable if we assume that the English progressive encodes *continuous aspect*. However, the gerund does not describe an actual situation but makes a generic assertion about hunting deer. This suggests a semantics closer to that of habituals, which also do not commit to the fact that the described situation takes place at the time of reference. For example, *John smokes* does not mean that he smokes right now, whereas *John is smoking* indicates that he is doing it at the time of speech. Of course, whether these examples should be compared to progressive marking (i.e., *Aux V+ing*) in the first place is an open question (Ramchand (2018) seems to answer it positively, though). This requires a more thorough analysis of these forms.

Overall, this thesis does not commit to the fact that the formal system proposed here is also transferable to non-finite forms without any modifications. As mentioned in chapter 2, the analysis of non-finite forms is left for future work.

While imperfective aspect is often researched more thoroughly due to its heterogeneous nature, I believe that perfective aspect deserves more attention. In the previous chapter, I have shown that the perfective aspect has strong implications concerning temporal reference. While these have been addressed in the literature, perfective aspect is usually simply described as perceiving an eventuality as an unanalyzed whole (Comrie, 1976; Smith, 1991). However, Singh (1991) presents evidence that the perfective aspect limits an eventuality rather than entailing the achievement of its endpoint.[13]

(29) a. mæ̃=ne ɑdʒ. əpnɑ kæk kʰɑjɑ or bɑki kəl
I=ERG today mine cake eat-PAST and remaining tomorrow
kʰɑũ-gɑ
eat-FUT
'I ate my cake today and I will eat the remaining tomorrow.'
Hindi

b. wakash-ita keredo wakanakatta
boil-PRV though boil.NEG
'I boiled the water, but it did not boil' Singh (1991)
Japanese

Similar observations can be made for other languages, as well. For example, in the discussion of the Slavic aspectual system, I have presented examples that indicate that the perfective entails a boundary rather than an endpoint when combined with a stative verb (see (4)). Example (30) shows the same for Urdu.

(30) kəmrɑ gənd̪ɑ rəh-ɑ.
room.SG.NOM dirty stay-PRV
'The room stayed dirty'/ 'There was a period in which the room stayed dirty.' *Urdu*

The same is also the case for the simple past which can be applied to activities and states to induce a bounded interpretation in English. Consequently, I

[13]Another possibility is that the relaxed semantics of the perfective are due to a language change from past perfective to a pure past tense marker (Condoravdi and Deo, 2014).

propose to annotate the perfective more distinctively, in particular, from the semantic perspective. Thus, I suggest that perfective markers can receive the following semantic features:

(31) perfective
complete bounded

To my knowledge, this distinction is not made overtly in the Grammar of any language. However, in many languages, when combined with telic predicates, i.e., predicates with an inherent endpoint, the perfective aspect usually entails that this endpoint is achieved. Thus, the examples in (29) represent very specific cases that need to be considered when annotating grammatical aspect. For the same reason, it is important to properly distinguish the semantic properties of grammatical aspect and inner aspect.

4.3.2 Inner aspect

The main distinction between inner aspect and grammatical aspect is that the former is a property of the verb phrase before temporal modification, while the latter plays a part in anchoring an eventuality with respect to its temporal and, primarily, modal dimension. I.e., grammatical aspect affects the relation between eventuality time and reference time, as already explained in the previous two chapters. Thus, the semantics of grammatical aspect are, in a sense, also relational like those of temporal reference, as we will see in section 4.4.
According to the Cambridge handbook of formal semantics (Rothstein, 2016), telicity is also affected by aspectual markers such as the progressive. Linguistic tests involving the specification of a concrete temporal endpoint support this view (Jackendoff, 1996; Tenny, 1987) Tenny (1987) shows this by virtue of the following example:[14][15]

[14]In essence, these tests correspond to the same tests proposed by Vendler (1957) and many researchers following his work.

[15]As Jackendoff (1996) points out, some predicates can be coerced into a telic or an atelic interpretation in the sense of Rothstein (see also Moens and Steedman 1988).

(32) a. Patricia climbed the tree.
b. It took Patricia two hours to climb the tree.

(33) a. Patricia was climbing the tree.
b. ??It took Patricia one hour to climb the tree.

Accordingly, there must be a property that is sensitive to these temporal tests but is it telicity? The answer to this question, of course, depends on the definition. In my view, there exist two layers of telicity: i) temporal telicity (there exists an endpoint, and it is reached within the reference time), and ii) aspectual telicity (there exists an endpoint, it must not have been reached in the actual world; i.e., during the reference time – see section 4.4). Temporal telicity affects the interaction with temporal modifiers as, e.g., in (32) or with the *for/in*-adverbial test. In comparison, aspectual telicity (or aspectual entailments of the telicity feature) affects the underlying aspectual inferences. Consider the examples in (34). Tenny (1995) explains that motion verbs such as *(to) push* are not specified for telicity. However, by concretely specifying the path as in (34b), the description becomes telic. If the progressive would cancel this (i.e. making the predicate atelic), then we must explain the discrepancy between (34a) and (34b) in a different way.

(34) a. Jordan was pushing a cart. $\not\rightarrow$ Jordan pushed a cart.
b. Jordan was pushing a cart to the store $\not\rightarrow$ Jordan pushed a cart to the store.

I see these two layers of telicity as two sides of the same coin. The interaction with grammatical aspect dictates which side of the coin is relevant for a given situation. Thus, when combined with the simple past, a telic predicate *realizes* its endpoint because of the default perfective interpretation, and when combined with the progressive (imperfective aspect), a telic predicate does not necessarily realize its endpoint. Generally, an atelic predicate combined with perfective aspect induces a bounded reading, whereas a combination with imperfective aspect induces an unbounded reading.[16]

[16] As, for example, Krifka (1998), shows, this view is not necessarily controversial, but I

As shown by the various examples in this section, telicity is not a property of verbs themselves but rather arises from the interaction between the verb and its arguments (Krifka, 1998; Verkuyl, 1972; Verkuyl, 1993). More specifically, a relatively simple notion like quantizedness of the complement of a verb (Krifka, 1998) is not sufficient to infer the telicity of an eventuality description. In example (34), both complements of *(to) push* are quantized. Intuitively, this means that if we dismantle the objects described by the complements, then the resulting sub-parts do not describe the same object, e.g., parts of a cart are not a cart, and parts of the path to the store are not paths that lead to the store.[17] Thus, as Krifka (1998) describes, the telicity of an eventuality description relies on the mapping of the sub-events to the sub-parts of the complement. In example (34a), each sub-part of the pushing event is mapped onto the same cart (not its sub-parts). In comparison, example (35) contains a *creation* verb. There, each sub-part of the drawing corresponds to another part of the circle being drawn.

(35) John was drawing a circle. $\not\rightarrow$ John drew a circle.

Similarly, as shown in section 4.2, telicity can be expressed overtly by using certain particles. This is also possible in English, e.g., *climb up the tree*. In fact, telicity can be expressed in numerous ways. Important examples, based on Hay et al. (1999), indicate telicity properties of various adjectives:

(36) They are straightening the rope. $\not\rightarrow$ They have straightened the rope.

(37) They are cooling the soup. $\rightarrow$ They have cooled the soup.

These examples are based on the degree to which the underlying property is achieved, i.e., the property of being straightened or the property of being cooled. This property can be measured according to a scale. *(To) Cool* induces an open scale. This means that scale of cooling has no inherent maximum. In comparison, *(to) straighten (a rope)* has a maximum: a completely straight rope (Kennedy and McNally, 1999).

find it worth spelling out for the sake of this thesis, due to the conflicting accounts found in the literature.

[17] Except maybe the final sub-part.

Furthermore, certain adjectival verbs can be combined with an adjective to induce a maximum of the scale attributed to the verb (Wechsler, 2005). This is shown in (38).

(38) a. The metal was hammered flat.
b. The puddle froze solid.

(Wechsler, 2005)

The common notion across all the examples so far is that they involve some kind of incrementality (Dowty, 1991), i.e., they follow some kind of path. This path can be provided by a scale (Kennedy and McNally, 1999), a measure phrase (Krifka, 1998), an actual spatial path (Jackendoff, 1996), among other means. Objects that provide that scale (either the verb itself or its arguments or modifiers) are generally called incremental themes (Dowty, 1991). Tenny (1995) calls these objects elements that *measure out*.
Let us call the line of work that defines telicity based on a path, actual or abstract, path-based telicity. Path-based telicity requires the following properties (taken from Wechsler (2005)):

(39) a. Some **property** of the **affected theme** argument changes by degrees along a **scale** due to the action described by the verb, until it reaches a **bound**.
b. The telic event and the path must be (a) homomorphic (parts of the event must correspond to parts of the path and vice versa) and (b) coextensive (the event must begin when the affected theme is at the start of the path and end when the affected theme reaches the end of the path).
c. The affected theme must be an argument of the event-denoting predicate.

4.3.2.1 Result states

An alternative view on telicity (or maybe a view on different core properties of telicity) is associated with the attainment of a result state Dowty 1977;

Pustejovsky 1991; Rappaport Hovav and Levin 1998). This is, in principle, compatible with the path-based telicity view. However, as Ramchand (2008) points out, verbs might differ concerning whether they lexically entail a result state, i.e., whether it is part of the verb itself. According to her, verbs like, for example, *(to) break sth.*, *(to) arrive*, or *(to) enter* specify their result state as part of their lexical semantics, in comparison to, e.g., motion verbs where the result is calculated from the interaction between the verb and its arguments (Ramchand, 2008). Thus, in addition to the properties above, a telic eventuality description must entail or allow for calculating a result state. This is simple in the case of verbs like *(to) arrive* or *(to) enter*. There, the result state can be described as a transition from *negp* to *p* for a given predicate (Pustejovsky, 1991). I.e., the result state of the predicate is described by the predicate itself. Compare this with some kind of motion verb, e.g., *(to) push x to store*. There, the result state is something like *at-store(x)*. *(To) hammer x flat* results in *flat(x)*. Generally, the result state is specified either by the verb itself or by the type of path/scale provided by the verb's argument.
In summary, telic predicates describe the congruent movement of an eventuality and its arguments towards the maximal or boundary point of a path. Some verbs entail their result state, whereas, for other verbs, the result state is calculated based on the verb and its arguments (i.e., telic pair formation, a compositional process according to Higginbotham 2009). As explained at the beginning of this section, telicity is determined before aspectual modification. There is no commitment to a result state being reached. This is determined by grammatical aspect.

4.3.2.2 The process property

One of the most fundamental distinctions of inner aspect is that between a state and a non-state (Bach, 1986). In chapter 2, I have explained that the progressive in English only combines with verbs that describe a process, i.e., have the process property according to Portner (2019). According to the Vendlerian classification of verb phrases, the process property can be described as some kind of dynamicity or change. In accordance with the observations in the previous section, change can be described as the movement across a path. Of course, this path does not need to have a boundary. Thus, both of the following examples express a process:

(40) a. Kim walked to the store.

b. Kim walked around in the park.

Consequently, processes themselves are unbounded but can become bounded when combined with an argument providing an appropriate path. Verbs like *(to) enter* or *(to) arrive*, which also entail a result state are still compatible with the progressive, which suggests that they also entail a process.[18]

Generally, I assume that *processes* are an important building block for distinguishing between states and non-states. Progressive aspect can be used to test for a process state. As mentioned above, a process is the part of an eventuality that affects its argument relative to the scale it provides. Thus, the process property is also essential for encoding telicity.

4.3.2.3 Eventuality initiation

As of now, I described a process stage and a result stage as possible building blocks for eventualities. Ramchand (2008) proposes a third building block for verb descriptions, namely the *initiation* stage. The initiation stage can be seen as the pre-state of an eventuality. In this pre-state, there is not yet any qualitative change that would affect telicity. It could also be understood as a preparatory stage.

While not affecting telicity per se, the `INITIATION` stage potentially affects the interpretation of the English progressive. As Dowty (1977) points out, the progressive has a futurate reading, where the described process is shifted into the future. This is illustrated in examples (41) and (41).

[18] In the previous section, I used the term transition based on Pustejovsky (1991). Arguably, achievements like *(to) arrive* instantiate such momentary transitions. For the sake of this thesis, I assume that a transition simply refers to the relation between the process and the result of a given eventuality. Typically, the process part is momentary but can be protracted, for example, by applying the progressive.

(41) John is visiting his family tomorrow.

(42) I am eating an apple now.

If we assume that the progressive refers either to the initiation stage or the unfinished process stage, such examples will fall out naturally without requiring the assumption that the present tense encodes non-past in these cases. Rivero and Arregui (2010) similarly describes the potential of the imperfective aspect, particularly the progressive, to refer to preparatory stages.[19]

Next to the *process* property and the *result state*, the *initiation* stage is the third building block, I assume for the semantics of inner aspect. As shown above, these properties are important in determining crucial notions of inner aspect and its interaction with grammatical aspect. In the next section, I will recapitulate the main inference patterns that I want to capture with the semantics worked out in this chapter.

4.3.3 Aspectual inference patterns

Giving a complete overview of all semantic properties associated with grammatical and inner aspect goes far beyond the scope of this thesis. For this reason, I limit the scope of this chapter mainly to some of the previously introduced inference patterns, mainly the *imperfective paradox.*

4.3.3.1 Inference patterns of progressive aspect

The imperfective paradox is concerned with the question of why some kinds of eventualities but not others allow for an inference from the progressive (PROG) form to a form where the completion of the given eventuality is implied. Incidentally, this is the meaning that is generally (but not always) associated with the English simple past. I call this a perfective (PRV) interpretation of the simple past. The inference pattern is illustrated in (43).

(43) a. John was drawing a circle. $\not\rightarrow$ John drew a circle.

[19] The English *going to expression* might be a result of this property of the progressive (see Eckardt 2006).

i. *PROG(draw a circle)* $\not\rightarrow$ a circle came into existence.

ii. *PRV(draw a circle)* $\rightarrow$ a circle came into existence.

b. John was pushing a cart. $\rightarrow$ John pushed a cart.

(44) John was knocking on the door $\rightarrow$ John knocked on the door (at least once).

Following the discussion in section 4.3.2, it should become clear that the imperfective paradox distinguishes between (path-based) telic and atelic sentences. When combined with an eventuality that describes a path with a maximal boundary point, then the inference from the progressive form to the perfective form fails. In comparison, the inference from progressive to perfective form holds when the path is undirected/unbound or not specified (the default interpretation of unspecified paths seems to be unbounded).
To leverage this information, it is necessary to be able to calculate the paths as they are constructed for different verb classes. For example, for *creation* and *consumption* verbs, the results are *creation* and *consumption* respectively. However, this is almost tautological. A more general path would be from *non-existence* to *existence* and vice versa. As explained before, the result of an actual path entails the existence of the underlying argument at the specified goal.
For adjectival verbs, the result state is lexically described, and the same is true for most aspectual particles. On one hand, consider example (45a). Obviously, the result state of the banana is not *up*. The particle's contribution solely indicates that the described eventuality has a specific boundary. On the other hand, *off* in example (45b) might be sufficient to describe the result state of the underlying result.[20]

(45) a. Jordan ate up the banana.

b. The ball rolled off.

Consequently, to properly model the imperfective paradox, it is not only necessary to annotate grammatical and inner aspect but also provide annota-

[20] Rather than specifying an endpoint of a process, *off* seems to describe a transition similar to the previously discussed achievements.

tions specific to various verb classes. In many cases, this calculation is bound to be context-sensitive. In this thesis, I rely on the fairly well-understood examples introduced so far, being well aware that this is a considerably large avenue for future research. [21]

4.3.3.2 Inference patterns of habitual aspect

Similar to the progressive, the habitual also provides indicative inference patterns when compared to the other grammatical aspects (Katz, 2011). Some of the basic inferences that I aim to capture with the ParTMA annotation scheme are presented in (46).

(46) a. Jordan eats apples. $\not\rightarrow$ Jordan is eating an apple.
b. Jordan eats apples. $\stackrel{?}{\rightarrow}$ Jordan ate apples.
c. Jamie ate apples. $\rightarrow$ Jamie was eating apples.
d. Jamie ate apples. $\stackrel{?}{\rightarrow}$Jamie was eating apples yesterday.

The notable difference here is between habituals in the scope of past temporal reference and present temporal reference. Under past temporal reference, the habitual entails an actual occurrence of the described eventuality. However, this is only true if the past temporal is indefinite, i.e., the reference time is not specified in more detail (compare with (46d)).
Similarly to the temporally constrained past temporal reference, present temporal reference also does not assert a commitment to the described eventuality in the actual world at the time of evaluation. Consider the following examples from Fara (2005).

(47) a. This orange-crushing machine crushes oranges (when it's switched on.)
b. This electric toaster crushes oranges (when it's switched on.)

If there is an expected causing relation (i.e., a causes b) between the subject of a sentence and the verbal predicate, then the expected activity attributed

[21] Not to speak of the integrating quantifiers into the mix, as discussed, e.g., in Verkuyl (1993).

to the subject does not need to occur. Furthermore, suppose the causing relation is not expected, such as in (47b). In that case, two readings are acceptable: one where the unexpected event happened, let's call it precedence reading, and second, a challenging reading that is based on an unlikely, unverified claim that may or may not be proved in a creative way (In the current example, one could expect that the person putting up the challenging claim turns on the toaster picks it up the toaster and crushes the orange by hammering on it. All in all, there seems to be no claim about an actual instance of the described event to have taken place in the case of present tense habituals.[22]

Consequently, the habitual does not entail the truth of the described eventuality during the time of reference. An exception to this occurs if the past temporal reference is completely indefinite. Compared to the imperfective paradox, the habitual inference patterns are simpler since they are not affected in the same way by telicity. In the next section, I develop a semantics that captures the basic facts outlined in this section and the previous sections.

4.4 Semantics for aspectual properties

In this section, I describe the formalization of grammatical and inner aspect as proposed for implementation in the ParTMA annotation scheme. Since I have pointed out several flaws of the ParGram system in chapter 2, I will not say much more in this chapter. My main reason for this is that many discrepancies in the ParGram grammars are a matter of naming conventions of features.[23] In some cases, e.g., the Polish perfective non-past, a semantic analysis has found its way into f-structure. However, whether this step is implemented in the grammar or the semantic rewrite component is not of immediate concern for the overall architecture of the ParTMA annotation scheme. Such smaller inconsistencies may even be desired on an application-to-application

[22] While this example supports the main point I am trying to make, I do not expect that a computational system could discern the two interpretations proposed by Fara (2005).

[23] Accounting for these belongs into a separate document, e.g., a specification for tense/aspect annotation in ParGram.

basis. I.e., maybe some applications benefit from a partial semantic analysis directly at f-structure. As long as there is an entailment relation between morpho-syntax and semantics, this should not be a problem.[24] [25]

As in the previous chapter, I propose a compositional semantics for the aspectual properties discussed in this chapter so far. As announced in the introduction, I use formal semantics that describes aspectual properties in terms of situations and events. More concretely, grammatical aspect is modeled in terms of situation semantics following (Arregui, 2005; Arregui et al., 2014; Cipria and Roberts, 2000), and lexical aspect is modelled in terms of a decompositional event semantics (Ramchand, 2004; Ramchand and Svenonius, 2014; Ramchand, 2008).

(48)

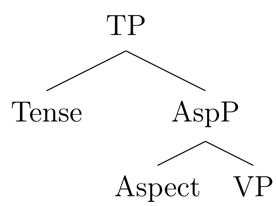

4.4.1 Situations and events

I continue the work on the ParTMA semantics under the assumption that the *TP* now denotes a set of situations, i.e., is of type *s,t*. This does not affect the points made in the previous section due to the assumption that situations are (potentially) spatio-temporal entities that can be reduced to only one property, e.g., their temporal anchoring (Kratzer, 2017), making them, essentially, time intervals. This stands in contrast to a semantics that uses world-time pairs for tense and aspect (see, for example, Deo (to

[24] As noted in chapter 2, it is still an issue from the perspective of cross-linguistic comparability. Thus, the fact that I focus on the semantics rather than the morpho-syntactic analysis is merely a concession and should be taken more seriously in future work.

[25] One important property of the aspectual systems presented in this chapter is the ability to recursively apply them (de Swart, 2016; De Swart, 1998). Slavic tense/aspect systems in particular have shown that recursive stacking does not result in a simple *perfective* or *imperfective* value, but rather that each morpheme contributes a specific meaning to the overall interpretation (e.g., iterativity, distributivity; Gehrke 2008; Romanova 2007).

appear), Dowty (1982), and Kusumoto (2005)). In the present work, I assume sufficient overlap between the two approaches to capture relevant insights from both strains of research but encode them in a situation semantics to keep the ontology (arguably) more straightforward. [26] Thus, both tense and grammatical aspect apply to and return sets of situations.

With respect to the level of the VP, I assume an event semantics that is *post-davidsonian* as claimed by Ramchand (2005). However, this does not affect the most basic claim that verbs, in their lexical semantics, denote an event variable (Davidson, 1967). What is different in the event semantics assumed for this thesis is the way in which arguments are related to this event. Generally, the approach I assume follows a more Neo-Davidsonian formalization, where (internal) arguments are not interpreted as inherent properties of the verb but rather as predicates identifying relations between the event and its arguments (e.g., Parsons 1990). Furthermore, as I have done throughout this thesis, events are eventualities, i.e., they describe both stative and non-stative situations (Higginbotham, 1985, 2000).[27]

An important property of event semantics is that the event closure is generally required to apply at the lowest possible scope (Champollion, 2015). This means that NP quantifiers generally outscope eventualities. However, since they affect telicity (Verkuyl, 1972), the event closure cannot be associated with telicity directly. I will not discuss this issue in this chapter, but chapter 5 presents an analysis that circumvents some of the pitfalls from a formal perspective mentioned in Champollion (2015). In section 4.4.3, I present a

[26] This consideration gets more weight if we take into account that we want to describe lexical aspect in terms of a decompositional event semantics (see below), since then events and situations together form a sufficient ontology. A potential counter-argument to situation semantics is the fact that the order of composition does not fall out naturally from the typing. This widely accepted practice has been coined by Kratzer (1998). There, grammatical aspect takes as argument a property of eventualities and returns a property of times, which in turn is saturated by a time variable (or a temporal existential closure operator, as has been shown in the previous chapter). However, this argument does not affect the present framework since we assume the rigid structure in (48) independently of compositional utility. For a recent discussion on the order of composition, see Ramchand (2018).

[27] Events are still distinguished from states, however, not in terms of their ontological status but instead based on the properties of a given eventuality (see section 4.4.3).

decompositional event semantics that strictly applies to the event variable and, thus, bears no repercussions for the treatment of grammatical aspect and temporal reference.

The last component we briefly need to discuss is the mediator between eventualities and situations. I follow the traditional way of mediating between eventualities and situations by means of exemplification (Kratzer, 2017). With this, I follow approaches to defining aspectual semantics in a situation semantics framework such as Cipria and Roberts (2000) and Arregui et al. (2011, 2014). I rely on the definition of exemplification used in the former approach shown in (49a). Crucially, the corresponding component in the compositional framework is shown in (49b). Exemplification applies at the top of VP to an existentially closed event predicate, as illustrated in (49c).[28] Note that the proposition embedded under eventuality exemplification is not yet marked for tense and aspect. The result is a situation that is just big enough to capture all constraints imposed on the situation by the event semantics representation (see Cipria and Roberts (2000, p. 318) for a more elaborate discussion of the notion of exemplification concerning events).

(49) a. *Eventualities that Exemplify Propositions* If s is any possible situation and p any proposition, then s is an eventuality that exemplifies p iff for all s' such that $s' \leq s$ and p is not true in s', there is an s" such that $s' \leq s'' \leq s$, and s" is a minimal situation in which p is true.

(Cipria and Roberts, 2000, p. 318)

b. $[\![\textit{ exemplification }]\!]^{g,w} = \lambda p_t.\lambda s.exemplify(s,p)$

c. $[\![\textit{ Brutus stab Caesar }]\!]^{g,w} =$
$\lambda s.exemplify(s, \exists e[stab(e) \wedge ag(e) = Brutus \wedge pt(e) = Caesar])$

The plan for this section is as follows: Firstly, I motivate and describe the formal semantics of grammatical aspect proposed for use in the ParTMA

[28] The description of the event participants in this example does not reflect the semantics in this chapter but rather a standard Neo-Davidsonian treatment along the lines of Parsons (1990). Furthermore, this example does not make any assumptions about the telicity of the described eventuality.

annotation scheme in section 4.4.2. Then, I propose a semantics for inner aspect in section 4.4.3 that is capable of interacting with cross-linguistically attested morphosyntactic reflexes of this category. As mentioned before, this goal is achieved by using a first-phase syntax (decompositional event semantics) as a mediator between the syntax and the semantics of inner aspect.

4.4.2 Compositional semantics for grammatical aspect

Following the TAM structure down the tree in (48), the first step in developing a compositional semantics for aspectual properties is to define a semantics for grammatical aspect. In this section, I propose a semantics for perfective aspect and then discuss a semantics for imperfective aspect and the distinctions therein, such as the habitual and the progressive.

4.4.2.1 Formal semantics for perfective aspect

In section 4.3.1, I have adapted the view that the main contribution of grammatical aspect is to distinguish whether a situation is described as bounded or unbounded. There, perfective aspect was associated with a bounded interpretation. Furthermore, we have seen that the boundary contributed by the perfective coincides with the natural endpoint of a given action in most cases. Describing the perfective as a boundary marker was supported on the one hand on work by (Singh, 1991), but we have also seen more general examples of perfective markers indicating boundaries. This property of the perfective aspect is also reflected in the most basic, traditional representations of perfective aspect, here shown as a constraint in (50a).

(50) a. *Perfective:* $\tau(e) \subseteq t$

b. *Imperfective:* $\tau(e) \supset t$

This formulation of the constraints imposed by grammatical aspect is based on Neo-Reichenbachian tense semantics and, in particular, the work done by Wolfang Klein (Klein, 1994, 2009). The most peculiar property of the perfective, namely, that it is incompatible with the present tense, has been discussed in the previous chapter. In the present tense, reference or topic time, t in (50), coincides with the speech or evaluation time and is, thus, assumed to be a moment in time. This means basically that the start and the endpoint of this time interval are indiscernible. As such, it is not possible to find a time interval that fits within the duration of the speech time and of which a description of some eventuality holds true. Thus, I model the perfective aspect as a part-of relation that holds between situations akin to the traditional treatment of the perfective outlined above.

(51) $[\![\textit{perfective}]\!]^{g,w} = \lambda s.\lambda s'.s \leq s'$

Rather than treating (51) as a constraint, as for example, in DRT (Kamp, 1979; Partee, 1984) [29], I treat perfective aspect as an existential quantifier. Thus, the full denotation of the perfective in the present annotation scheme is as in (52). [30]

(52) $[\![\textit{perfective'}]\!]^{g,w} = \lambda p_t.\lambda s_0.\exists s[s \leq s_0 \wedge exemplify(s,p)]$

As shown above, encoding the perfective aspect in terms of the part-of relation is a good first approximation that captures the traditional insights in the literature. To account for cases such as those presented in Singh (1991), where the perfective actually denotes a cessation rather than a culmination point, the semantics can be refined. In the previous chapter, I presented a framework in which the semantic differences between operators rely on defining the restrictor of a given temporal quantifier. Similarly, I argue that a more fine-grained semantics for the perfective (and later the imperfective and progressive) can be achieved by providing different denotations for the restrictor of the existential quantifier shown in (52).

Concretely, I propose that there are two distinct interpretations of the perfective. One where the event(s) exemplified by a given situation s culminate (in

[29] Partee (1984) does not specifically describe aspectual properties but refers to a Reichenbachian tense system, where aspect plays a role in relating the reference time (constrained by tense) and the event run time.

[30] Under this view, situations constrained by tense are not pure time intervals, else, the sub-part relation would always turn out false, but situations with some spatio-temporal content, thus, making the approach comparable to frameworks that use world/time pairs.

the sense of Parsons (1990); i.e., they reach their natural endpoint), and one where the exemplified event(s) terminate. Termination simply means that an eventuality is bounded for states and processes. For accomplishments, this means that they are potentially terminated prematurely. In other words, the perfective slices out a relevant junk of an eventuality and presents that as a complete whole. Given this explanation, I propose the definitions in (53) as potential variations of the simple part-of restrictor used before.[31]

(53) a. $[\![$ *cul* $]\!]^{g,w} = \lambda s.\lambda s'$: *s'* exemplifies a telic eventuality. *s* is a situation, such that all eventualities that begin in *s'* reach their natural endpoint in *s*.

b. $[\![$ *bounded* $]\!]^{g,w} = \lambda s.\lambda s'$. *s* is a situation such that all eventualities that begin in *s'* terminate in *s*.

The restrictors defined in (53) carry some universal force to deal with pluralities exemplifying multiple eventualities. This is necessary to account for Singh's (1991) partitive telic readings. From a temporal perspective, the *cul* restrictor corresponds to the part-of relation. However, the *bounded* restrictor also allows for the reverse part-of relation, i.e., similar to imperfective aspect, the reference time may be included in the run time of an eventuality (the eventuality time), or in terms of situations, the topic situation may be contained in the situation exemplifying the event. To avoid clashes with the modal analysis of the imperfective in the next section, I further refine this analysis in section 4.5. The gist of this idea is that bounded readings are incompatible with lexically telic verbs (e.g., *(to) arrive*). Alternatively, they are only compatible with eventuality descriptions that describe a protracted process stage, i.e., an eventuality that follows an incremental path).

4.4.2.2 Formal semantics for imperfective aspect

The idea of associating grammatical aspect with quantificational force is commonplace in the formal semantics literature. In fact, the semantics

[31] In principle, the part-of operator corresponds to an unspecified restrictor, whereas the more specific readings are only available if inferable via the syntax/semantics interface. This is explained in more detail in section 4.5.

described above have been developed as a counter-part to the much more discussed treatment of imperfective aspect as a universal quantifier over situations, or world time pairs (Arregui et al., 2011, 2014; Cipria and Roberts, 2000; Deo, to appear; Dowty, 1977). The semantics I propose here are closest in spirit to work by Arregui et al. (2011, 2014) and encode the cross-linguistic variation of imperfective aspect in terms of the restrictor of a universal quantifier over situations. The restrictor of this quantifier, called modal base in Arregui et al. (2011, 2014); based on Kratzer (1981), can be understood as a modal accessibility relation, relating two situations based on certain conditions. This is, in fact, the same procedure that was used in the previous section. However, the modal bases available for imperfective aspect are generally distinct from those available for perfective aspect.[32]

(54) $[\![\textit{ imperfective' }]\!]^{g,w} = \lambda p_{st}.\lambda s_0.\forall s[MB(s)(s_0) \rightarrow p(s)]$

In accordance with Comrie's (1976) typological picture, I start with the distinction between continuous and habitual aspect. Incidentally, Arregui et al. (2014) defines the ongoing aspect – continuous aspect without an event in progress reading – in terms of the same part-of relation as the default imperfective aspect. This means, the denotation for an ongoing interpretation is as in (56).

(55) $\lambda s.\lambda s'.s \leq s'$

(56) $[\![\textit{ ongoing }]\!]^{g,w} = \lambda p_{st}.\lambda s_0.\forall s[s \leq s_0 \rightarrow p(s)]$

This analysis of the *ongoing* aspect ascertains that the topic situation is homogenous for states and processes since the predicate in question must be true of all sufficiently large sub-situations. For telic eventualities, the sub-situations are also restricted with respect to their granularity. The sub-situations have to be at least large enough to allow the truth of the telic eventuality in question. Thus, the given telic eventuality is iterated

[32]Making any claims about unifying modal bases across perfective and imperfective aspect lies beyond the scope of this dissertation, and I mention them, at best, as potential ideas for future work.

back-to-back within the topic situation (Arregui et al., 2014), resulting in an iterative interpretation.

Generic and habitual readings function similarly. The modal base is shown in 57. The important notion here is that of a characteristic (sub-)situation. Characteristic sub-situations are determined by virtue of the content of the utterance, as well as the context (Arregui et al., 2014; Cipria and Roberts, 2000). Obviously, the sub-situations must also be large enough to accommodate for the eventuality described by the predication that the imperfective scopes over. However, other than the ongoing modal base, this modal base does not scope over actual but possible situations that do not (necessarily) happen back-to-back. According to Arregui et al. (2014) this explains, why the habitual modal base allows for the utterance of non-accidental generalizations (see Arregui et al. (2014, 316ff)).[33]

(57) $[\![\ MB_{generic} \ ]\!]^{g,w} = \lambda s.\lambda s'.s'$ is a characteristic part of s.

The event-in-progress or progressive interpretation of the imperfective and dedicated progressive markers differs crucially from the two previously discussed modal bases in one regard: the topic situation is a part of the situation exemplifying an eventuality. In comparison, the habituaL/generic interpretation and the ongoing interpretation treated the topic time as a super-situation of the eventuality-exemplifying situations. This is in line with work by (Deo, 2009), who makes similar claims.[34] [35] Thus, the progressive clearly distinguishes itself from the other imperfective interpretations. The relevant modal base is defined in (58).

(58) Event Inertia
$\text{MB}_{\text{E-inertia}} = \lambda s.\lambda s'.s'$ is an **Event-Inertia** situation for s,

[33]Deo (2009) provides a more formally more fleshed out account for defining the domain of quantification of the imperfective, but since this is not relevant to the main point of this section: Formalizing imperfective in terms of a universal quantifier and its different flavors as different modal bases. I leave this note regarding the topic for the interested reader.

[34]An important insight from Deo's (2009) is that these part-of relations are constrained to the temporal domain, as she models the imperfective as a quantifier over world/time pairs.

[35]A more relaxed view would be to state that the generic modal base identifies characteristic situations that are accessible from the reference situation (see also section 4.5).

> where for any two situations *s* and *s'*, *s'* is an *Event-inertia* situation for *s* iff all the events that have actually started in *s* continue in *s'* as they would if there were no interruptions.
>
> (Arregui et al., 2014, p. 327)

With this, I have covered the basic cases of imperfective aspect. However, I have also presented some additional variations in the previous section. As mentioned before, I support the claim in the literature that further variation can be encoded by inserting different modal bases (Arregui et al., 2014). I exemplify this in terms of a final example, the preparatory imperfective, which is used when an eventuality is anticipated to occur in the (relative) future. For this, Arregui et al. (2014) assumes that preparatory stages of eventualities and the eventuality itself can be sufficiently distinguished (the decompositional event semantics in the next section provide one way of doing this). Given this distinction, the modal base for future-oriented imperfectives can be defined as in (59). For further aspectual distinctions, I refer the reader to Arregui et al. (2011, 2014) and Cipria and Roberts (2000).

(59) Preparatory Inertia
$[\![\ MB_{P\text{-}inertia} \ ]\!]^{g,w} = \lambda s.\lambda s'.s'$ is a **Preparatory-inertia** situation for s,
where for any two situations *s* and *s'*, *s'* is a *Preparatory-inertia* situation for *s* if all the events that are in preparatory stages in *s* continue in *s'* as they would if there were no interruptions.

(Arregui et al., 2014, p. 327)

In summary, I have presented a compositionally simple approach to encoding grammatical aspect that ascribes the perfective, the core meaning of an existential quantifier over situations, and the imperfective aspect, the meaning of a universal quantifier over situations. Fine-grained distinctions in aspectual systems are encoded in terms of variations in the restrictor of those quantifiers. These restrictors serve as accessibility relations between situations. However, not all of the accessibility relations are necessarily modal in the sense that they refer to non-actual or possible situations rather than situations that are part of the actual timeline, including past and present eventualities that have actually taken place. The proposal is built on existing literature that is still discussed in current research and thus serves as a modern foundation for the semantic annotation of tense and aspect.

4.4.3 Formalizing the semantics of inner aspect

As exemplified in section 4.3.2, inner aspect is only partially encoded in the lexicon. In fact, verbs are often underspecified concerning typical notions of inner aspect. Rather, it is encoded at the level of VP. This means that inner aspect is (minimally) derived from features on the verb and its NP arguments (in particular, its internal arguments, in the sense of Kratzer (1996)). This insight has been confirmed syntactically and semantically by many researchers, e.g. Hinrichs (1985), Kiparsky (1998), Kratzer (2004), Krifka (1998), Ramchand (1997), Verkuyl (1972), and Verkuyl (1993), just to name a few.

As Ramchand (1997), among others, argues, the morphosyntax of both the verb and its internal arguments needs to be taken seriously to systematically describe the interaction between form and meaning with respect to inner aspect, which is, in fact, regular and predictable in many cases. To this end Ramchand (2008) develops, what she calls, the first-phase syntax, which is, in fact, a functional structure designed for a decompositional event semantics (see also Ramchand 2004, 2018; Ramchand and Svenonius 2014).

The main idea behind this approach is to cover both syntactic and semantic reflexes encoded in the description of an event with a small set of primitives. As discussed in the previous section, the first-phase syntax, for example, reveals that only a small set of verbs are truly telic. More often than not, telicity is a property of the whole verb phrase rather than only the verb itself. In the view adopted in this dissertation, eventualities are abstract objects without any spatio-temporal anchor. They are only ready for anchoring via tense and grammatical aspect, after a situation exemplifies them. Thus, telicity is also an abstract concept that acknowledges the existence of a natural endpoint but not whether this endpoint is actually reached. As seen in the previous section, this is the role of grammatical aspect. This distinction is highlighted by encoding lexical aspect purely in terms of a decompositional event semantics.

"Change" is seen as one of the fundamental properties encoded in lexical aspect (Filip, 2012). In the present framework, eventualities are divided into two categories: those that describe a state and those that describe an eventuality that contains some internal change. They are identified via the respective predications as shown in (60).

(60) a. State(e) : e is a state
b. Process(e) : e is an eventuality that contains internal change

The semantics presented here decompose eventualities into up to three separate eventualities. Those eventualities are related via causal implication (Hale and Keyser, 1993). For the present approach, two types of causal implications are considered. Firstly, a state may lead up to a process, and secondly, a process leads to a result.

(61) $IF, \exists e_1, e_2[State(e_1) \wedge Process(e_2) \wedge e_1 \rightarrow e_2]$, then by definition Initiation(e_1)

(62) $IF, \exists e_1, e_2[State(e_1) \wedge Process(e_2) \wedge e_2 \rightarrow e_1]$, then by definition Result(e_1)

This distinction corresponds to the distinction between causing an event and transitioning to the result state of an event, i.e., causative semantics vs. resultative semantics (Ramchand, 2008). This assumption plays a core role in defining the first-phase syntax and semantics. Generally, such transitions are captured via the event composition rule shown in (63).

(63) Event Composition Rule:
$e = e_1 \rightarrow e_2$: e consists of two subevents, e_1, e_2 such that e_1 casually implicates e_2
(cf. Hale and Keyser (1993)) from Ramchand (2008)

The event decomposition rule makes available sub-eventualities of a given eventuality for modification. More concretely, Ramchand (2008) proposes

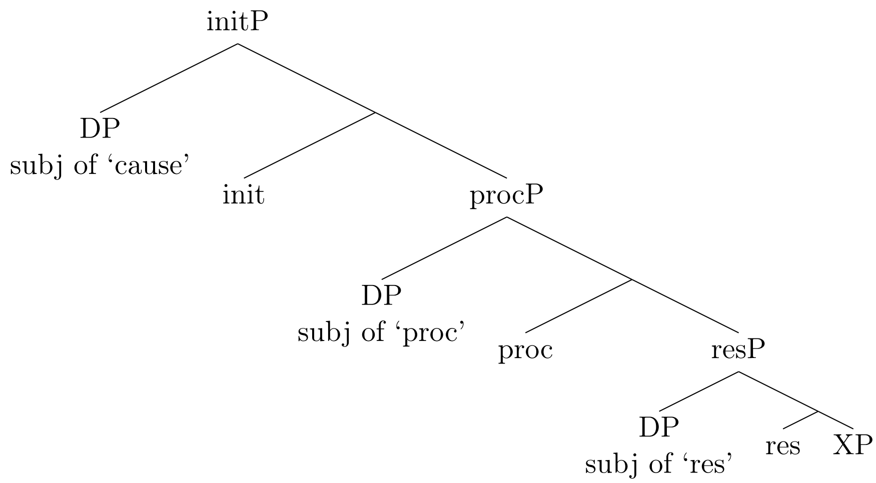

Figure 4.1: First-phase-syntax template (Ramchand, 2008)

(at most) three sub-eventualities for a given predicate. This is illustrated in Ramchand's (2008) example in (64). There, the predicate *(to) defuse the bomb* is decomposed into an *initiation* state, a *process*, and a *result state*. This is the underlying structure of Ramchand's (2008) first-phase syntax.

(64) 'defuse-the-bomb'(e) *where* e = $e_1 \rightarrow (\ e_2 \rightarrow e_3)$:
$[initiate - defuse(e_1 \wedge process - defuse(e_2 \wedge result - of - defusing(e_3)]$

(Ramchand, 2008, p. 36)

The tree in Figure 4.1 illustrates the syntactic structure underlying this event composition. Notably, the structure does not only define sub-eventualities in terms of *init*, *proc*, and *res* but also defines subject positions for each of them. Each of these sub-eventualities correspond to a different projection in the first-phase syntax, called *initP*, *procP*, and *resP* respectively that stand in a hierarchical structure from initiation to result.

4.4.4 The first-phase syntax

The subject positions of each projection correspond to a small set of primitives roughly comparable to thematic roles (Dowty, 1991). However, compared to other approaches for defining thematic roles (in computational semantics), they provide a more general set of semantic roles that are tied to concrete syntactic diagnostics (Hautli et al., 2015). The concrete specification for each of the subject positions in relation to the first-phase syntax is given in (65).

(65) a. Subject(x,e) and Initiation(e) entails that x is the `INITIATOR` of e.

b. Subject(x,e) and Process(e) entails that x is the `UNDERGOER` of the process.

c. Subject(x,e) and Result(e) entails that x is the `RESULTEE`.

Finally, the rhematic position is represented in terms of *XP*. Rhematic material or `RHEMES` are components of a predication that are not part of the sub-eventuality structure. Ramchand (2008) argues that they are part of the description of the predicate, and can appear in different forms, as we will see throughout this section[36] If any rhematic material is specified, it takes the complement position of the lowest available projection. This means that rhematic material can apply to each of the three projections depending on the configuration of the underlying verb. This is explained in more detail below.

4.4.4.1 The initiation

Telicity in the first-phase syntax is affected by the configuration of the sub-eventualities as well as the configuration of the different subject positions and a path. Predicates that concretely specify a result state are telic at some level, as is the case for *(to) defuse the bomb.* On the other end of the spectrum are stative, i.e. atelic, predications, which Ramchand (2008) argues, only encode a *initP*. This is exemplified in (66). The *init* sub-eventuality describes a state and determines John as the cause that this state holds.

(66) a. John loves Mary.

[36] Recall also section 4.3.2). The rhematic material affects the specification of the path variable as in, e.g., the optional PP phrase in *Jordan pushed the cart to the store.*

b.

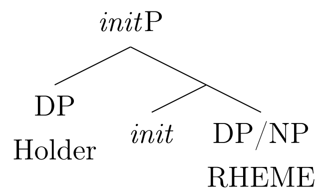

In this case, *Mary* takes the place of the rheme. Ramchand (2008) does not explain how the *init* head combines with its rhematic material in examples such as (66). I tentatively suggest encoding such objects as themes in terms of a Neo-Davidsonian event semantics. Thus, a sentence like *John loves Mary* would have the semantics in (67).

(67) $\exists e[love(e) \wedge init(e) \wedge state(e) \wedge subject(j, e) \wedge theme(e) = m]$

As initially mentioned, the first-phase syntax allows various kinds of configurations of different heads. This structure also affects the semantic denotation. For example, in (66), the *init* head combines with rhematic material, an entity of type *e,vt*, e.g. via predicate modification. However, in the presence of a *proc* head, the *init* head is required to combine with an object of type $< v, t >$, according to Ramchand. Thus, the semantic denotation of first-phase syntax heads is directly affected by the immediate syntactic context, a property that is particularly interesting from the perspective of the semantic construction rules, as we will see in section 4.5. For the *init* head this means, we need to presuppose (at least) two different denotations shown in (68a) and (68b).

(68) a. $[\![\ init\ ]\!]^{g,w} = \lambda P_{vt}.\lambda x_e.\lambda e_v.P(e) \wedge init'(e) \wedge State(e) \wedge Subject(x, e)$

b. $[\![\ init\ ]\!]^{g,w} = \lambda P_{vt}.\lambda x_e.\lambda e_v.\exists e_1, e_2[P(e_2) \wedge init'(e_1) \wedge State(e_1) \wedge e = (e_1 \rightarrow e_2) \wedge Subject(x, e_1)]$

Let us look at the second possible denotation of the *init* head, (68b), more closely. As it illustrates, the *init* head decomposes an eventuality into an initiation state, e_1, and a second eventuality, e_2, that is specified by the first argument of the function. Accordingly, this object is of type $< v, t >$ and is provided by the next lower head: the *proc* head.

4.4.4.2 The process

With the process head, it is possible to produce *process* and *initiation, process* verbs. Consider the example in (70). As shown there, the subject *Jordan* takes the role of `INITIATOR` and the object *cart* appears as `UNDERGOER`. Generally, the `UNDERGOER` follows the path dictated by the `RHEME` in the XP.[37]

(69) Jordan pushed the cart to the store.

(70)

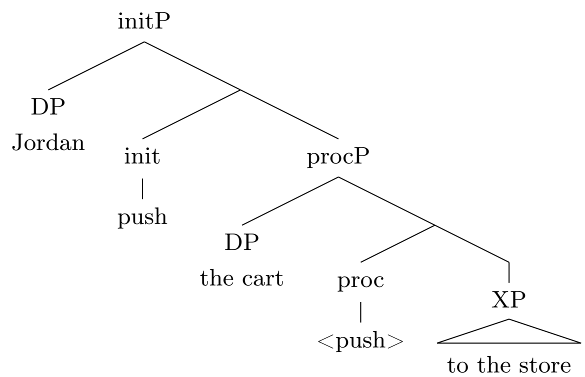

From the perspective of aspectual properties, the important role contributed by the *proc* head is the component of "change". The kind of change that occurs depends on how the subject position, i.e., the `UNDERGOER` position, is filled. Consider the case of consumption/creation verbs. Their objects are often referred to as *incremental themes* (Dowty, 1977). The relation between an eventuality and an incremental theme object can be defined via a homomorphism relation as proposed by Krifka (1998). Thus, the change in the eventuality is described by virtue of the change of the incremental theme. However, in the first-phase syntax, the `UNDERGOER` of the change is not the incremental theme but the entity responsible for causing the change. This is due to the fact, that the `UNDERGOER` is supposed to serve as a *continuous experiencer* of the eventuality (see Ramchand (2008, p. 60)).

[37]If the path describes a final place, i.e., is bounded, then it should be possible to also specify a result state. The following example encodes a trajectory rather than a final place:

(1) Jordan pushed the cart towards the store.

For a concrete treatment of path PPs, I refer to (Ramchand, 2008, ch. 5). However, I also make a tentative proposal for analyzing them in 5.

However, consumption and creation prevent and cause existence, respectively, disallowing continuous change. [38]

(71) Jamie ate an apple.

(72)

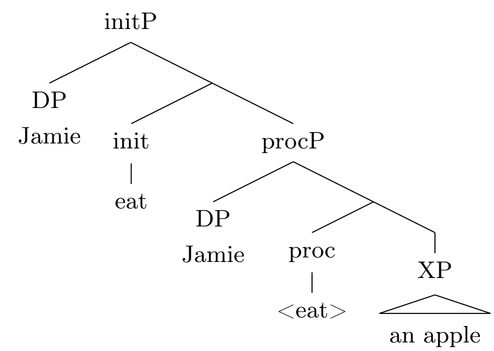

Due to this fact, Ramchand (2008) argues that the telicity in creation/consumption verbs is an implication derived from the semantic relation between the *proc* sub-eventuality and the rhematic material. Building on Krifka (1998), Ramchand (2008) encodes objects that are gradually affected (for example, incremental themes) as paths, and relates eventualities and paths via the `Path/2` relation, that takes an entity and an eventuality as argument. According to the definition by Ramchand (2008) in (73), for an entity x, there must exist some monotonic property of that entity that is also monotonic with respect to the part-whole structure of an event. This monotonic property of x, is a set of measures D_x, which is defined such that for any two measures d,d_1, there exists a linear order for them in D_x, based on the part of the entity, x, that they measure: Thus, iff x_1,$x_2 \subset x$, and $x_1 \subset x_2$ then $d \leq d_1$.[39]

(73) PATH(x,e) =
$_{def}$ $\exists R \exists D_x [\forall e, d, d' [R(e,d) \& d' \leq d \rightarrow \exists e' \subseteq e \& R(e', d')]$
(mapping to measures) &
$\forall e, e', d' [R(e,d) \& e' \subseteq e \rightarrow \exists d' [d' \leq d \& R(e', d')]]$ (mapping to events)

[38] Ramchand (2008, 162f) also provides further syntactic evidence for such a treatment of creation/consumption verbs from a cross-linguistic perspective.

[39] Ramchand (2008) treats the path variable (X in (73) as an entity of type e. More concretely, she proposes to pragmatically infer a measure based of the denotation of a given NP or PP (e.g., *eat an apple* or *push a cart to the store*. She does not explain how she reduces both NPs and PPs to entities of type e.

As explained in section 4.3.2, telicity entailments are generally based on the properties of a given path. However, the presented implementation does not allow us to distinguish these properly since we cannot infer from Ramchand's denotation whether the described path has a maximal endpoint (it depends on the set of measures). Furthermore, I do not believe such a system can be feasibly implemented computationally, or, at least, is much too detailed to be testable in any real-world scenario. Thus, rather than formalizing the relationship between an eventuality and the path it takes in terms of some kind of mapping function, I believe it to be sufficient to describe the path in terms of a set of features (see section 4.5). As a result, for the remainder of this section, I treat $Path(x, e)$ as defined by Ramchand (2008) as a black box that gives us information about the a path specified according to some variable x with respect to an eventuality e.
Throughout the past few sections, I have mentioned that the path is mainly tied to the process stage of an eventuality. Thus, Ramchand proposes two different denotations for processes. One that picks up a *path* and one that combines with a *res* head, i.e., a result state. The semantics of the two variants are shown in (74) and (75) respectively.[40]

(74) $[\![\textit{ proc }]\!]^{g,w} = \lambda P_{vt}.\lambda x_e.\lambda e_v.\exists e_1, e_2[P(e_2) \wedge proc'(e_1) \wedge Process(e_1) \wedge e = (e_1 \rightarrow e_2) \wedge Subject(x, e_1)]$

(75) $[\![\textit{ proc }]\!]^{g,w} = \lambda y_e.\lambda x_e.\lambda e_v.Path(y, e) \wedge proc'(e) \wedge Process(e) \wedge Subject(x, e)$

In the first phase syntax, subjects can be merged into different positions, consequently allowing for duplicate entries of an NP/DP. This process allows for marking compound thematic roles such as *initiator-undergoers* (as in (72)). For the sake of introducing the relevant concepts, I accept this treatment. However, in chapter 5, I introduce a process by which compound roles are formed via co-binding the respective variables.

[40] A potential avenue for future work is to merge these to denotations since, even if an eventuality expresses a process and a result, it would still need to follow the path.

4.4.4.3 The result

This behavior also extends to the result head. In Ramchand (2008), she mainly focuses on result states that are part of the lexical semantics of a given verb. She identifies these by attaching rhematic material that describes a final state. One example she gives is:

(76) Sam broke the stick <u>into pieces</u>.

Accordingly, the first-phase syntax of *broke the stick* is given in (78). This example also shows another instance of a compound thematic role, the *undergoer-resultee*.

(77) Sam broke a stick.

(78)

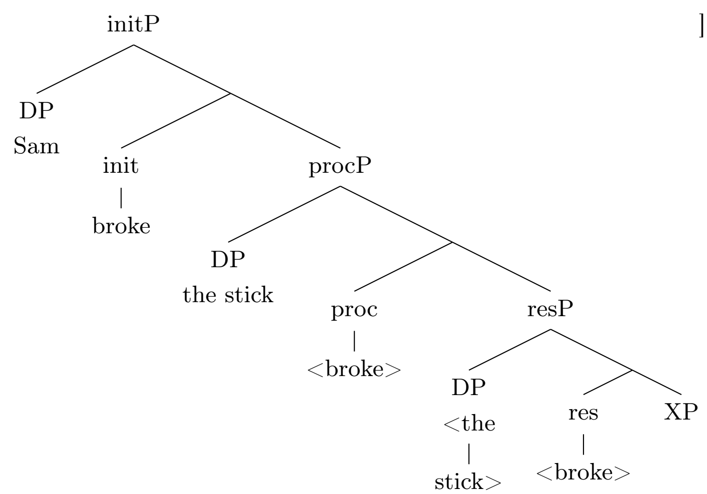

The semantics of the *result* head are fairly simple, as shown in (79). The event variable there is generally saturated by the *proc* head (see (74)). Furthermore, these semantics can combine with a state description to model the result state more specifically.

(79) $\lambda P_{v,t}.\lambda x_e.\lambda e_v.\exists res'(e) \wedge subject(e,x) \wedge P(e)$

4.4.4.4 The rhematic material

The rhematic material is essential for describing a given eventuality's telicity properties. Rhematic material on the process allows for the inference of

telicity, i.e., by providing a path with a concrete goal. When combined with the *result head*, the rhematic material is expected to simply describe a final state. What does not become apparent in the first-phase syntax is the relation between *path* and *final state* since no path is specified if a result head is available. As already mentioned in section 4.4.4.2, this treatment is somewhat confusing and is a potential avenue for further research. An added benefit of this would be to reduce the number of differing *proc* semantics, thus, simplifying the system.
As examples (70) and (72) have shown, a telicity inference is not congruent with a *result* head in the syntax. However, as described in section 4.3.2, a result state can be calculated based on the boundary of the provided path. This is not worked out in detail in Ramchand's (2008) work. However, in the next section and 5, I illustrate some possibilities for treating telic pair formation.

4.4.5 Interim summary

In this section, I have presented the tools provided by Ramchand (2008) to describe the lexical semantics of verbs. More concretely, I have shown how thematic roles are related to particular eventuality sub-events. Thus, the present semantics tie together argument selection and aspectual properties. For this fact alone, I believe the present system is worth pursuing. In the next section, I describe how the present system can be implemented in a mainly feature-based annotation scheme. Furthermore, I show how these annotations can be linked with the annotation of grammatical aspect to provide a detailed semantics of both of these categories.

4.5 Semantic annotation of aspectual properties

In this section, I explain how the semantic insights presented in the last section can be integrated into the ParTMA annotation scheme. As in the previous section, I focus on deriving a semantic feature structure.[41] In the previous section, I presented a formal semantics for grammatical aspect. However, I have also indicated that it requires some adjustments to fit together with the decompositional event semantics presented in 4.4.3. Thus, I first comment on the annotation of inner aspect. Then, I explain how to leverage this information to provide a formalization of grammatical aspect that captures the basic facts of a modal theory of imperfective aspect as presented in section 4.4.2 and translates into a computationally viable formal framework.

4.5.1 Inner aspect

The annotation of inner aspect is based on the first-phase syntax developed by Ramchand (2008). As she argues, some aspectual information is encoded directly in the lexical semantics, whereas other information arises from the interaction between the verb and its arguments in a more compositional manner. More specifically, internal arguments are mapped onto a *path* thematic role, whose properties describe the mode of affectedness of the argument and, potentially, a boundary, indicating telicity.

In addition to the aspectual information, the first-phase syntax presents a unique way to capture thematic roles. They are specified as *subjects* of the sub-events identified by the first-phase syntax: *initiation, process* and *result*. The configuration of the sub-events and their results can be understood as subcategorization templates. Some of these are illustrated in (80).

(80) *(to) roll* : $[init_i]$
(to) push sth.: $[init_i,\ proc_j]$
(to) eat sth. : $[init_i,\ proc_i]$
(to) arrive : $[init_i, proc_i, res_i]$

[41] Concrete details on the interpretation of the semantic feature structure are given in the next chapter.

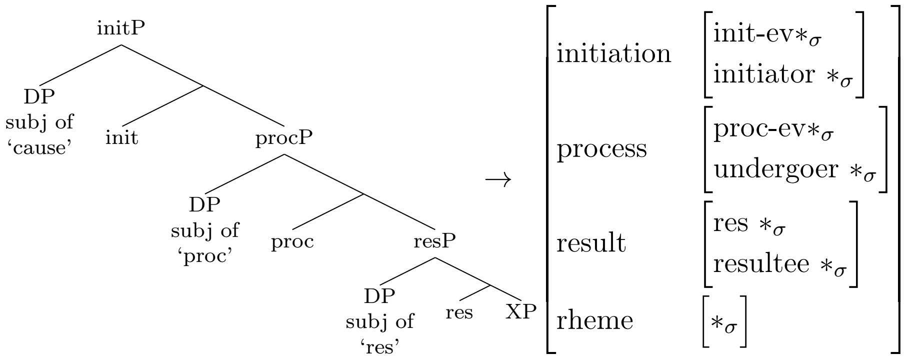

Figure 4.2: Flattening the first-phase syntax

These templates describe the available sub-events and the roles of their arguments. For example, *(to) push* has a separate `INITIATOR` and `UNDERGOER`. In comparison, *(to) eat* licenses a compound argument, an `INITIATOR-UNDERGOER`. This kind of encoding only captures the very basic facts of the first-phase syntax. Thus, while it provides an intuitive way to refer to different verb classes, it is not sufficient for a full-fledged semantics.

For a detailed annotation of the first-phase syntax, I propose the flat feature structure in Figure 4.2 on the right. It does away with the framework-specific assumptions of Ramchand's (2008) work that are of no benefit when combined with the computational grammars used in this thesis.[42]

A concrete annotation example is given in (81). As shown there, the values of the annotation are indices in the f-structure. Thus, the general principle for this annotation is that the first-phase annotation embeds partial f-structures (or their semantic counterparts) to enrich them with semantic information.

(81) a. Sam pushed the cart. $[init_i, proc_j]$

[42]The hierarchy is preserved internally for semantic interpretation. This is discussed in more detail in the next chapter.

b.
$$
{}_{0}\begin{bmatrix}
\text{PRED} & \text{‘push<Sam,cart>’} \\
\text{SUBJ} & {}_{1}\begin{bmatrix}\text{PRED ‘Sam’}\end{bmatrix} \\
\text{OBJ} & {}_{2}\begin{bmatrix}\text{PRED ‘cart’}\end{bmatrix} \\
\text{TENSE} & \text{past}
\end{bmatrix}
\begin{bmatrix}
\text{initiation} & \begin{bmatrix}\text{init-ev}[0] \\ \text{initiator } [1]\end{bmatrix} \\
\text{process} & \begin{bmatrix}\text{proc-ev}[0] \\ \text{undergoer } [2]\end{bmatrix}
\end{bmatrix}
$$

As explained in the previous section, such an annotation would also receive a default unbounded path variable, resulting in an atelic eventuality description. Let us compare this to an example based on telic pair formation and a lexically telic predicate.

Example (82) describes a lexically telic predicate. As pointed out in the previous section, this can be tested by attaching PPs that describe a final location. For example, *(to) arrive at the store* sounds fine, while *(to) arrive to the store* sounds weird.[43] In short, the example describes a $init_i, proc_i, res_i$ verb. As the mapping from syntax to semantics indicates, this can be encoded by mapping grammatical functions to nodes in the first-phase syntax.

(82) a. Sam arrived. $[init_i, proc_i, res_i]$

b.
$$
{}_{0}\begin{bmatrix}
\text{PRED} & \text{‘arrive<Sam>’} \\
\text{SUBJ} & {}_{1}\begin{bmatrix}\text{PRED ‘Sam’}\end{bmatrix} \\
\text{TENSE} & \text{past}
\end{bmatrix}
\begin{bmatrix}
\text{initiation} & \begin{bmatrix}\text{init-ev}[0] \\ \text{initiator } [1]\end{bmatrix} \\
\text{process} & \begin{bmatrix}\text{proc-ev}[0] \\ \text{undergoer } [1]\end{bmatrix} \\
\text{result} & \begin{bmatrix}\text{res-ev}[0] \\ \text{resultee } [1]\end{bmatrix}
\end{bmatrix}
$$

However, in example (83), there is no result state specified. I assume this is a case where semantic construction rules come into play. As the annotation indicates, I propose to decompose the PP *to the store* into a path component and a goal component. From the perspective of telicity, the `GOAL` the final bound of the path, whereas `PATH` describes the trajectory (based on Kracht

[43] Although, *to the store* can be understood as a PP that describes a path and a final state.

(2002) and Zwarts (2005), as proposed by Ramchand 2008).

(83) a. Sam pushed the cart. $[init_i, proc_j]$

b.
$$
{}_0\left[\begin{array}{ll}
\text{PRED 'push<Sam,cart>'} & \\
\text{SUBJ} & {}_1\left[\text{PRED 'Sam'}\right] \\
\text{OBJ} & {}_2\left[\text{PRED 'cart'}\right] \\
\text{ADJUNCT} & \left\{ {}_3\left[\begin{array}{ll} \text{PRED 'to<store>'} & \\ \text{OBJ} & {}_4\left[\text{PRED 'store'}\right] \end{array}\right] \right\} \\
\text{TENSE past} &
\end{array}\right]
\left[\begin{array}{ll}
\text{initiation} & \left[\begin{array}{l} \text{init-ev[0]} \\ \text{initiator [1]} \end{array}\right] \\
\text{process} & \left[\begin{array}{l} \text{proc-ev[0]} \\ \text{undergoer [2]} \end{array}\right] \\
\text{RHEME} & \left[\begin{array}{l} \text{PATH [3]} \\ \text{GOAL [4]} \\ \text{TYPE 'spatial'} \end{array}\right]
\end{array}\right]
$$

Similarly, other verbs also can be decomposed into a final `GOAL` state, as well as a path, as argued by Ramchand (2008). Thus, in examples like *(to) eat up the banana*, *up* indicates that there is a final goal. However, since *up* is not an entity (as, e.g., *store*), it does not describe the goal itself but rather allows for the calculation of the result state relative to the type of verb. For example, with respect to a consumption verb, the final state is that of being completely eaten or not existing in its original form.[44] Accordingly, the presence of a `GOAL` in the annotation leads to a telic interpretation. In contrast, the existence of a `PATH` does not commit to the existence of an inherent endpoint.

4.5.2 Grammatical aspect

Properly understanding the annotation in the previous section requires its combination with the annotation of grammatical aspect. The tree in Figure 4.3 summarizes the distinctions that the annotation scheme is supposed to capture. As explained in section 4.3, perfective and imperfective aspect are affected by telicity. Thus, in this section, I explain how the two types of telicity presented in the previous section interact with grammatical aspect.

[44] Different consumption verbs might need to be decomposed into more detailed types since, for example, *(to) swallow* does not imply that the underlying object is destroyed.

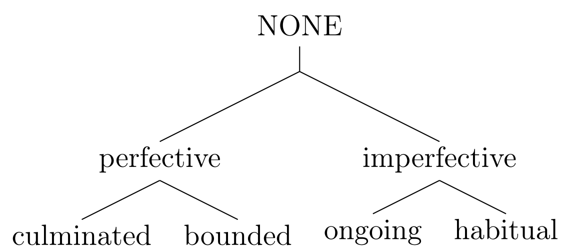

Figure 4.3: Feature space for grammatical aspect

4.5.3 Semantic annotation: perfective aspect

In this section, I want to plead for a semantics that distinguishes between a perfective that realizes a unique natural endpoint and a perfective that simply realizes an endpoint, i.e., a boundary. Generally speaking, this distinction is only available for verbs that do not encode their result state lexically. This means that the existence of a result state combined with perfective aspect automatically results in a culminated interpretation.
However, verbs that do not specify their endpoint directly but are based on following a path can be affected by contextual factors. An example of this is given in (84) taken from Jackendoff (1996).

(84) Sam walked the Appalachian Trail.

The NP *the Appalachian Trail* can receive to distinct interpretations. One that refers to a part of the trail and one that refers to the complete trail. A possible semantics is given in the next example. As shown there, when the NP simply describes part of the path, it combines with the *process* head, thus, indicating a process. When a *goal* is specified, it introduces a result state description which can only be accommodated by semantically constructing a *result* state.

(85) a. $[\![\ Goal\ ]\!]^{g,w} = \lambda e_v.\textit{at-end(e)}\wedge Goal = \iota x[trail(x)]$

b. $[\![\ Path\ ]\!]^{g,w} = = \lambda x_e.\lambda e_v.Path(\iota x[trail(x)], e) \wedge proc'(e) \wedge Process(e) \wedge Subject(x, e)$

Accordingly, the annotation *culminated* perfective forces the semantic construction of a result state, whereas the *bounded* perfective neither presupposes a lexically specified nor semantically constructed result state. Conversely, if perfective aspect remains underspecified (i.e., simply $s < s'$), then telicity annotation automatically derives the appropriate semantics.

Since both of these possibilities are available for the example in (84), the annotation is entirely reliant on the annotation of the *path*: the Appalachian Trail. Consequently, the example does not allow the inference from the progressive to its result state based on grammatical aspect alone. The inference is based solely on the annotation of the underlying path.

Generally, this section has shown that the perfective aspect interacts tightly with telicity, as observed in the literature. Depending on the exact annotation, various entailments need to be resolved (e.g., that a culminating perfective requires the presence of a result state). Consequently, the imperfective paradox is also based on the existence of such a result state (this has already been postulated in section 4.3.

4.5.4 Semantic annotation: imperfective aspect

In the previous section, I have shown that the perfective aspect, in certain instances, confirms the presence of a result state. Thus, the perfective interpretation monotonically adds information to the eventuality description. Consequently, the perfective and the imperfective are, in principle, derived from the same underlying eventuality description.

4.5.4.1 The continuous aspect

In accordance with the formalization discussed in section 4.4.2, the continuous aspect simply fills up the reference situation with instances of the described eventuality. For this, the semantics do not have to be modified. However, it is worth noting that iterativity can be introduced in a number of ways in the morphosyntax, e.g., plural subject, plural object, and possibly verbal modifiers (although in these cases, the iterative interpretation might stem from the semantics of the modifier rather than the aspectual marking. Thus, the interpretation might appear iterative because of the context. The context is not making the interpretation iterative). Thus, it introduces a variety of variables such as pluractionality (e.g., Lasersohn 1995) and plurality (e.g., Link 1998). Since this interpretation is not related to the main inference patterns described in 4.3.3, I leave it for future work.

4.5.4.2 The progressive aspect

In terms of the annotation, the progressive is straightforward since it uses the same basic mechanisms as temporal reference. This is due to the fact that it also introduces a quantifier where differences in annotation correspond to differences in the restrictor of the underlying quantifier. Thus, I use this section to focus on explaining the interaction between the progressive aspect and the first-phase syntax in the present annotation scheme.

The general idea is that the progressive, similarly to the perfective, relies on the existence of a syntactically or semantically constructed result state. However, rather than describing the complete eventuality as part of the reference interpretation, the progressive only entails the existence of either the *initation* stage, or the *process* stage at the reference situation. Accordingly, the modal bases need to be modified to account for these semantics. Since both the progressive and the futurate progressive rely on *inertia* semantics in the original proposal, they receive similar interpretations based on the first-phase syntax rather than proper inertia. In (86a), the situation describes both the *initiation* and the *progress* in cases where an *initiation* sub-event exists. Since most processes are the result of some kind of initiation state (except maybe in expressions like *the ball rolled*), this state must have occurred for the process to be observable at the reference situation.[45]

(86) a. $[\![\ MB_{progressive}\ ]\!]^{g,w} = \lambda s_s.\lambda s'_s.$ *s describes the initiation and process of the eventuality described by s'*

b. $[\![\ MB_{preparatory}\ ]\!]^{g,w} = \lambda s_s.\lambda s'_s.$ *s describes the initiation of the eventuality described s'*

[45] To my knowledge, there are no tests that are sensitive to whether an eventuality has actually come to be as predicted by the linguistic description.

Concretely, the first phase syntax allows to explain in more detail which part of a given eventuality takes place during the reference situation and which part of an eventuality takes place in the inertia situations. Since we also know which entities take place in which sub-events, we can also calculate the existence or non-existence of the corresponding entities in the different situations. This idea is difficult to formulate precisely in the semantics used in this thesis. However, a possible implementation has been described in Zymla (2019).

4.5.4.3 The habitual aspect

The original implementation of the habitual describes it as sub-parts of the reference situation. However, based on Zymla (2019), I propose to describe the habitual in terms of an accessibility relation, as in (87a). The accessible situations are such that they exemplify the underlying eventuality, but the underlying action must not be part of the reference situation.

(87) a. $[\![\; MB_{habitual} \;]\!]^{g,w} = \lambda s_s.\lambda s'_s.$ *s' is a situation accessible from s*

This straightforwardly allows capturing the intuition that the habitual does not entail its progressive or perfective counterpart at the reference situation (since it simply describes an empty situation with respect to the exemplification of the described eventuality). In this case, the semantics of inner aspect play a lesser role in relation to grammatical aspect than with the perfective and the progressive. However, as with the continuous aspect, the habitual aspect interacts with various elements in the syntax and semantics that have not been discussed in this thesis. Thus, there is room for future work.

4.6 Summary

This section had three goals: i) describe the cross-linguistic variation of grammatical and inner aspect, ii) provide an overview of the basic demands of a semantic theory of tense and aspect, and iii) show how these can be encoded formally and in the ParTMA annotation scheme.

To achieve goals ii) and iii), I have introduced a situation semantics for grammatical aspect and proposed to encode inner aspect in terms of the first-phase syntax (Ramchand, 2008). As shown in section 4.5, the two layers of meaning are closely connected. The result is, at its core, a semantics for aspect that explains aspectual inferences by virtue of lexically encoded or semantically constructed result states. Such result states are often not computed based on a given verb but rather on its complementary material, including the internal argument and further rhematic material.

This chapter has addressed only some fundamental aspects of aspectual semantics, i.e., inference patterns between expressions marked for different grammatical aspects. This thesis argues that the presented semantics allow for capturing these inferences, but a concrete description of the underlying algorithm/system is left for future work.

The next section further illustrates this by implementing the semantics presented in this and the previous chapter. However, as noted towards the end of this chapter, future work might benefit from a paradigm shift to another semantic framework. A potential candidate is a semantics that has been used to model the effects of modal operators in Bobrow et al. (2007), Condoravdi et al. (2001), Crouch (2005), Crouch and King (2006), and De Paiva et al. (2007). Zymla (2019) builds on this to propose a more suitable semantics for grammatical aspect and inner aspect. However, discourse representation theory, or at least one of its many incarnations, might also provide a more suitable semantics for future work since it would make the incorporation of discourse effects much easier. All in all, aspectual features still provide a huge avenue for future research.

Chapter 5

Implementation and Applications

In the last two chapters, I have established the ParTMA annotation scheme for the semantics of tense and aspect. There, the main point was to illustrate the cross-linguistic variation in the realization of tense and aspect features and how this variation can be captured. The annotation scheme presented in this thesis has two components: a feature set and a ruleset. The feature set, of course, provides descriptive information about tense/aspect features, whereas the ruleset has shown to be necessary (or at least useful) to implement *semantic construction*. *Semantic construction* covers a wide array of interactions between tense and aspect categories and tense/aspect with other categories that affect the interpretation of semantic features. In the tense domain, this was illustrated primarily in terms of the Sequence-of-Tense phenomenon and the treatment of tense in (counterfactual) conditionals. In the domain of grammatical aspect, default interpretations have been constructed from tense features and features of lexical aspect. Concerning lexical, or inner, aspect, semantic construction has been illustrated in terms of telicity, which was derived from the interaction of the verb and its rhematic material in a first-phase syntax.

In this chapter, I describe an implementation of both of these components, starting with the generation of semantic construction rules. These have been implemented in terms of a system of graph-rewriting, which is inspired by the packed rewrite system used in XLE (Crouch et al., 2017) and graph-matching algorithms used for querying treebanks, in particular, Rosén et al.'s (2012) *Infrastructure for the Exploration of Syntax and Semantics* (INESS). This system is described in section 5.1. The second component, the feature space of the annotation scheme, has been motivated in the previous two chapters. In section 5.2, I describe the compositional framework through which the semantic features are composed into semantic representations based on first-order logic. Concretely, I present a compositional system based on LFG's Glue semantics (Dalrymple, 1999), grounded in the description-by-analysis tradition (Andrews, 2008; Crouch, 2005; Crouch and King, 2006; Wedekind and Kaplan, 1993). I show that by employing Glue semantics as an underlying semantic framework for the annotation scheme, the annotation can be kept relatively simple while still being able to systematically encode certain hierarchies involved in the composition of tense and aspect, for example, the $[T...[Asp...[VP]]]$ hierarchy, or the first-phase syntax. This means the abstract syntax of the annotation scheme (in the sense of Bunt 2015) may remain straightforward, while the semantics of the annotation scheme can become more flexible in dealing with various complex semantic phenomena.

5.1 Semantic construction via graph rewriting

In this section, I first discuss the implementation of semantic construction rules. For this purpose, I provide a formal system that is compatible with both XLE's output and the output of dependency parsers (in particular, the universal dependencies parser that is part of the Stanford CoreNLP library (Chen and Manning, 2014)). In other words, the system presented in this chapter is compatible with the two grammar formalisms described in the second chapter, LFG in 2.3 and UD in 2.4. It has two properties:

(1) **Property 1:** Identify partial syntactic (and semantic) structures

(2) **Property 2:** Add additional structural and semantic information to an input structure

To kick off this section, I explain how the present system implements **Property 1**. Then, I will illustrate how the system can be extended to produce semantic construction rules or, in other words, how to deal with **Property 2**. The properties in 1 are fairly general. However, I will decompose them in the next few sections to provide a detailed understanding of the implementation presented in this chapter.

5.1.1 Identifying partial syntactic structures

The identification of partial syntactic structures not only plays a role in syntactic parsing itself (especially in lexical approaches such as, for example, LFG, HPSG (Pollard and Sag, 1994), or LTAG (Joshi and Schabes, 1997)). It is also used to analyze syntactically annotated corpora (see Lai and Bird (2004) for a comprehensive overview) or to add additional syntactic structures to existing representations (e.g., the enhancement of universal dependencies (Schuster and Manning (2016); based on Chambers et al. 2007). Techniques of the latter kind have also been used to derive full-fledged semantic representations from the syntactic output of the XLE system (Crouch, 2005; Crouch and King, 2006).

What all of the approaches mentioned above have in common is, that they introduce a *query* language for identifying partial syntactic structures. For this thesis, three query languages serve as the main inspiration: the query language used in the infrastructure for the exploration of syntax and semantics (INESS; Rosén et al. 2012), and two components from the XLE system (Crouch et al., 2017). First, the query language used to identify partial f-structures in the unification-based implementation of local co-description, and second, the query language used in the packed rewrite system (PRS). This system was initially envisioned for translating syntactic representations between languages but was also used extensively for deriving semantic representations from XLE's syntactic representations (Crouch and King, 2006). More generally, the identification of partial syntactic structures can be understood as identifying partial linguistic annotations. As illustrated in Ide and Bunt (2010), one can abstract away from certain types of linguistic annotations to achieve an abstract annotation language, which, in turn, can be queried by a corresponding query language. I compile components from the three inspirations mentioned above into a query language for abstract syntactic representations and show how this query language is applicable to syntactic annotations provided by the XLE and the Stanford CoreNLP dependency parsers. Concretely, I adapt Ide and Bunt's (2010) approach in that I translate the respective syntactic representations into abstract syntactic graphs. This is explained in the next section. Understanding these structures then allows us to define a query language to identify partial linguistic structures.

5.1.2 Abstract syntactic graphs

Ide and Bunt (2010) have shown that linguistic annotation schemes are suited to be represented as graph structures. In the system implemented for this thesis, I use this insight to generate graph structures that I have dubbed abstract syntactic graphs. In this section, I explain the formal underpinnings of these graphs, representations for the unified annotation of XLE's syntactic structures, and universal dependency structures, based on ideas from (Zymla, 2018).[1,2] Both UD and XLE's syntactic structures use enumeration to identify core elements of the annotation (see Figure 5.1). However, in the case of UD, the enumeration also encodes word order, as shown in Figure 5.1 on the left. On the other hand, in XLE, integers simply serve as unique indices for partial f-structures. For both of these structures, indexation plays an integral role.[3] In abstract syntactic graphs, the indices used in the enumeration of XLE and UD structures serve to encode nodes in the graph, thus preserving

[1] More concretely, abstract syntactic graphs serve to unify UD structures and XLE's f-structure representations, since the f-structure is assumed to be the input to further semantic processing (Butt et al., 2002; King et al., 2005).

[2] Generally, syntactic parses consist of directed, possibly cyclic, graphs. For this reason, the ParTMA query language is also tailored toward use with directed graphs. The benefits of also covering graphs that are not directed still need to be explored and are, thus, left for future work.

[3] This indexation is, in fact, relevant for relating c-structure and f-structure to each other, but this is not relevant for the present purposes.

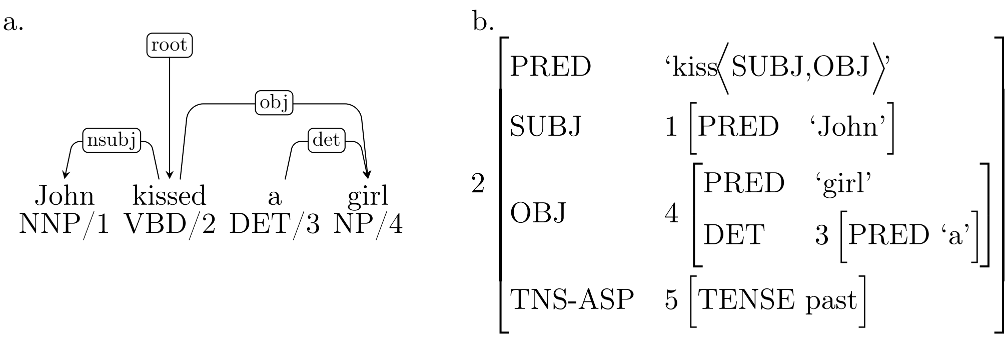

Figure 5.1: Enumerated syntax for: *John kissed a girl*

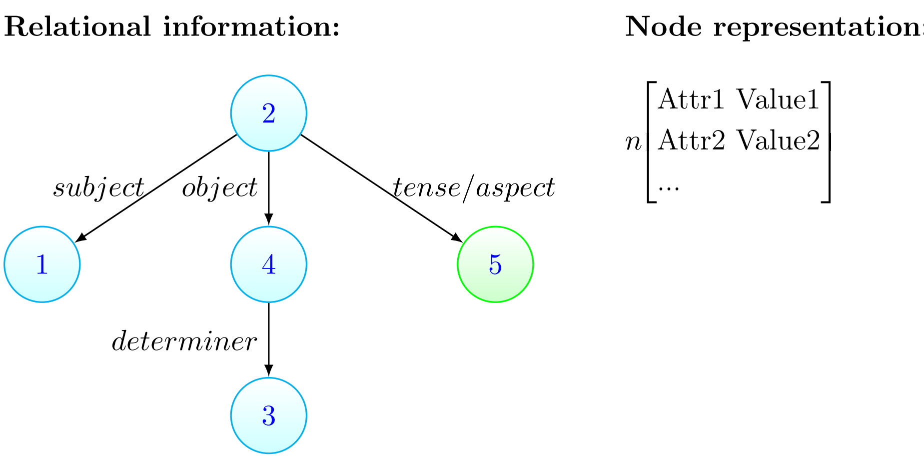

Figure 5.2: Abstract syntactic graph for *John kissed a girl.*

the additional information encoded in UD structures.

Abstract syntactic graphs differentiate between *relational* information and attribute/value pairs, the two cornerstones of linguistic annotation (Bunt, 2015). An example is given in Figure 5.2 for the syntactic representation in Figure 5.1b). The translation is very straightforward. Dependencies and grammatical functions are translated into labeled edges, whereas features associated with specific nodes are encoded in corresponding attribute/value matrices.

Before more complex cases are discussed, let us first define an abstract syntactic structure S. A tuple S consists of a set of graph node indices N, a set of edges E, a set of attributes A, and a set of values V.

(3) a. $S = (N, E, A, V)$

b. N is a set of integers referring to indices in a given syntactic representation.

c. an *attribute* in A is a symbol (i.e. a string) representing a valid relation, or attribute in a given syntactic representation.[4]

d. an edge is defined, for i,$j \in N$ and $a \in A$, as: $i \rightarrow_a j$

e. For each $n \in N$, there exists a (potentially empty) set of attribute value pairs, P^N. For any $p \in P^N$, $p = (i, j)$, where $i \in A$, $j \in V$.

A graph adhering to the formal outlines presented above can also be represented as a set of facts F, as specified in (4). This is the computational representation of f-structure graphs used in XLE. More specifically, it is the f-structure representation that XLE's transfer system uses – a simplification of the actual f-structure representation.

(4) F is a list of facts, such that every $f, f \in F$, is a triple of the form: $(x \in N, y \in A, z :$ either $z \in N$,or $z \in V)$

Figure 5.3 illustrates the fact-based f-structure representation used in XLE, where a given fact is encoded in terms of the `eq/2` predicate. Each fact is wrapped in a choice fact, `cf/2`, which is used to handle ambiguities. In an ambiguous structure, invariant facts receive the choice 1, whereas facts that pertain to specific readings of an ambiguity receive separate choices. This allows for the expression of ambiguous structures without the need to duplicate unnecessary information (Maxwell III and Kaplan, 1989).

In the examples presented here, I do not deal with ambiguous XLE-structures. The main reason for this is to preserve parallelism between f-structures and

[4]Attributes could be divided into *relations* and *features* formally, but this step is not required.

```
cf(1,eq(attr(var(2),'PRED'),semform('kiss',2,[var(1),var(4)],[]))),
cf(1,eq(attr(var(2),'SUBJ'),var(1))),
cf(1,eq(attr(var(2),'OBJ'),var(4))),
cf(1,eq(attr(var(2),'TNS-ASP'),var(5))),
cf(1,eq(attr(var(1),'PRED'),semform('John',1,[],[]))),
cf(1,eq(attr(var(4),'PRED'),semform('girl',7,[],[]))),
cf(1,eq(attr(var(4),'DET'),var(3))),
cf(1,eq(attr(var(3),'PRED'),semform('a',5,[],[]))),
cf(1,eq(attr(var(5),'MOOD'),'indicative')),
cf(1,eq(attr(var(5),'PERF'),'-_')),
cf(1,eq(attr(var(5),'PROG'),'-_')),
cf(1,eq(attr(var(5),'TENSE'),'past')),
```

Figure 5.3: Fact-based Encoding of an f-structure (Prolog style)

UD structures. Ambiguous f-structures can be represented efficiently using choice packing (Kaplan and Maxwell, 1995), but this is not readily available for UD parses. Thus, being able to deal with ambiguities is a desirable goal; however, it is left for future work.

In the next section, I work out the operations I use to identify sub-graphs in the kind of graphs discussed in this section. The first building block of the rule system is presented in this chapter. For the sake of this thesis, the corresponding query language is called ParTMA query language.

5.1.3 Basic query operations

The ParTMA query language is designed by taking inspiration from treebanking applications (in particular INESS; Rosén et al., 2012), as well as LFG and general graph matching techniques.

As described in section 5.1.2, the syntactic input is translated into a list of facts describing a graph. The most simple elements of the ParTMA query language are such facts, where any graph node indices are replaced by corresponding variables, indicated by a # followed by a lower-case letter, e.g., `#f`. This is exemplified in (5), which also indicates the instantiation of the given facts with respect to the syntactic structure in Figure 5.4.[5]

[5] In the following examples, I make use of universal dependency representations for ease

(5) a. `#a POS 'NNP'`
$\{\#a \rightarrow 1\}$

b. `#a POS 'VBD'`
$\{\#a \rightarrow 2\}$

c. `#a POS 'NP'`
$\{\#a \rightarrow 4\}$

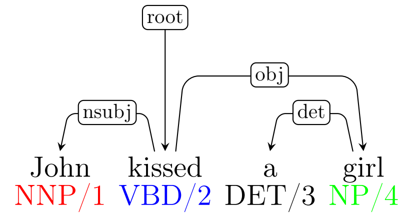

Figure 5.4: Dependency analysis of *John kissed a girl.*

The system does not only encode variables over f-structures but also allows variables for values that follow the same rules as f-structure variables. A % sign precedes such variables for values. Strings can be compared via the operators `==` (equals) and `!=`(not equal). Thus, a convoluted way of writing (5a) is shown in (6).

(6) `#a POS %a & %a =='NNP'`

The query language presented in this chapter also allows for the definition of paths through graphs as illustrated in (7), as applied to Figure 5.4. This example, furthermore, illustrates that multiple queries can be combined, sharing the same variable bindings via the symbol `&`. With this, more complex sub-graphs can be queried.

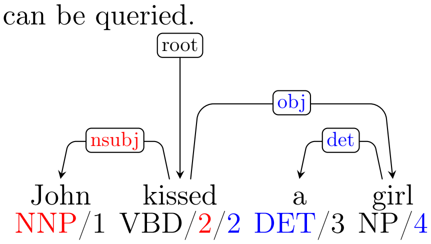

Figure 5.5: Example of path matching

(7) `#a NSUBJ #b POS 'NNP'` & `#a OBJ #c DET #d POS 'DET'`
$\{\#a \rightarrow 2, \#b \rightarrow 1, \#c \rightarrow 4, \#d \rightarrow 3\}$

In addition to the simple graph matching outlined above, the system also allows for queries inspired by LFG. In particular, the system allows users

of exposition.

to apply techniques based on LFG's functional uncertainty and inside-out functional uncertainty. This is presented in the next section.

5.1.4 LFG-inspired query operations

Functional uncertainty expands on the basic *functional application* equation which is used in LFG to identify whether some f-structure node either governs some grammatical function or contains some attribute/value pair. To be more precise, functional application in LFG checks whether an f-structure node satisfies a path constraint defined by some string. The formal definition for this is equation is given (8). Other than the paths outlined above, in LFG's FA, only the equation's start and endpoint are defined. Intermediate nodes that are part of the path are not instantiated. The ParTMA query language allows for a similar mechanism, as illustrated in the example in (8a). The exclamation mark expresses a direction (namely downwards; toward sub-structures), and the following string surrounded by parenthesis defines a path through the queried syntactic structure. In this case, the query succeeds if a syntactic node, `#h`, is found that is reached via the relations `COMP` and `SUBJ` in that order. A corresponding example is visualized below in Figure 5.6 and (12).

(8) **Basic equation** (corresponds to functional application; FA):
$(f\ a) = v$ holds if and only if f is an f-structure, α is an attribute, and the pair $< av > \in f$. (Bresnan et al., 2015)

a. **Example query:** `#f !(COMP>SUBJ) #h`

Functional uncertainty uses the same mechanism as functional application. However, rather than applying it to a single string, functional uncertainty operates on a set of strings as the function's argument. A set of strings can be invoked by the use of the *-operator. This operator looks for one or more instances of the preceding relation label (`COMP` in the example).[6] For

[6] The *-operator is handled differently in this implementation than in XLE, where the *-operator corresponds to the Kleene star, meaning zero or more instances of the preceding element (Crouch et al., 2017).

each possible match, e.g. `COMP>SUBJ`,`COMP>COMP>SUBJ`, etc, a new solution is generated by the query as long as corresponding `COMP` elements exist in the underlying syntactic structure. A concrete example is shown in Figure 5.6. Compare the structure with the available solutions for the functional uncertainty query in (12a).

(9) **Functional uncertainty (FU):**
For α any set of strings of attributes, $(f\ \alpha) = v$ holds if and only if for some x in the set of strings α, $(f\ x) = v$. (Bresnan et al., 2015)

a. **Example query: `#f !(COMP*>SUBJ) #h`**

The previous two techniques are convenient tools for checking if certain sub-structures (partial f-structures in LFG, sub-graphs in the present work) exist in a given syntactic representation. LFG, furthermore, provides tools for checking whether certain *super* structures exist, i.e., whether a given syntactic node is governed by specific grammatical functions (relations). The formal definition is given in (10).
In the present query language inside-out functional application follows the same syntax as functional application, however, the direction is specified as upwards by means of the `^`. Thus, the query in (10a), checks whether there exists any super structure `#h` from which `#f` can be accessed via the FA-path `(!COMP>SUBJ)`. This is illustrated in (12b) in relation to Figure 5.6.

(10) **Inside-out functional application (IOFA):**
For any f-structure f' and attribute α, $(\alpha\ f')$ designates some f-structure f such that $(f\ \alpha) = f'$.

a. **Example query: `#f ^(SUBJ>COMP) #h`**

Finally, an inside-out version of functional uncertainty also exists, which indicates multiple inside-out paths, just as its previously introduced counterpart, functional uncertainty. The concrete formal definition from LFG is given in example (11). As with FU, IOFU queries return all paths that satisfy the query. The example in (11a) is illustrated concretely in (13b), as the previous examples, with respect to Figure 5.6.

(11) **Inside-out functional uncertainty (IOFU):**
For any f-structure f', attribute α, and string of attributes x, $(\epsilon f') = f'$ (where ϵ is the empty string) and $(x\ \alpha\ f') = (x(\alpha f'))$. For any set of strings of attributes, $(\alpha\ f) = v$ holds if and only if for some x in the set of strings α, $(x\ f) = v$.

a. **Example query:** `#f ^(SUBJ>COMP*) #h`

(Bresnan et al., 2015)

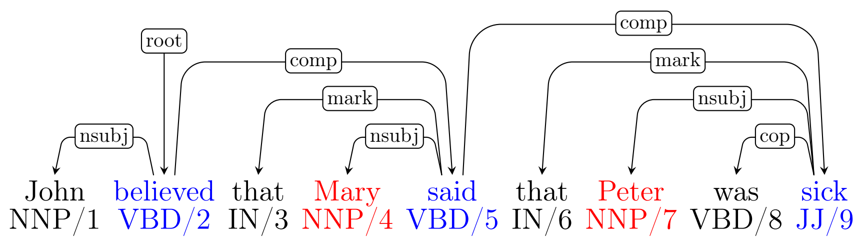

Figure 5.6: Example for LFG inspired queries

(12) **a. Example query (FA):**
`#f !(COMP>SUBJ) #h`
$\{\#f \to 2, \#h \to 4\}$,
$\{\#f \to 5, \#h \to 7\}$

b. Example query(IOFA):
`#f ^(SUBJ>COMP) #h`
$\{\#f \to 4, \#h \to 2\}$,
$\{\#f \to 7, \#h \to 5\}$

(13) **a. Example query (FU):**
`#f !(COMP*>SUBJ) #h`
$\{\#f \to 2, \#h \to 4\}$,
$\{\#f \to 2, \#h \to 7\}$,
$\{\#f \to 5, \#h \to 7\}$

b. Example query(IOFU):
`#f ^(SUBJ>COMP*) #h`
$\{\#f \to 4, \#h \to 2\}$,
$\{\#f \to 7, \#h \to 2\}$,
$\{\#f \to 7, \#h \to 5\}$

The functionalities borrowed from XLE, from functional application to inside-out functional uncertainty, follow a much simpler implementation than in the computational LFG framework since the input to the query language is always a finite f-structure as checked by the corresponding parsers (Kaplan and Maxwell III, 1988; Kaplan and Zaenen, 1989). This minimizes the computational challenges for the query language presented here.

5.1.5 Dictionary look-up

The final functionality of the present system, with respect to querying syntactic structures, is the use of dictionaries. In the sense of the present query system, dictionaries are external files that categorize relations and values according to specific tags. By calling a specific dictionary, relations and values can be replaced with a specific category if the corresponding element belongs to that category. Consider the toy example in (14a). It presents a dictionary that specifies the temporal orientation of temporal modifiers into either *past* or *future* orientation. This can be used to check the agreement between the tense in a sentence and any available temporal modifier. An example of how the lexicon is being accessed is shown in (14c) based on the node given in (14b). What happens there is that for the lexicon with the name `TNS-AGR`, the system returns the category that contains the value given to the `lex/2` function. Thus, in the current example $\%a$ is instantiated to *past* since it is in the corresponding set shown in (14a). As of now, this functionality is designed for dictionaries with mutually exclusive categories. Ambiguous entries are left for future work.

(14) a. `TNS-AGR` `past` $=$ {`yesterday,recently,...`}
`future` $=$ {`tomorrow, soon,...`}

b. $_0\begin{bmatrix}\text{TOKEN ‘yesterday’}\\ \text{TNS-AGR ‘past’}\end{bmatrix}$

c. `#a TOKEN %a & #a TNS-AGR lex(%a,TNS-AGR)` $\rightarrow$
`#a TNS-AGR past`
$\{\#a \rightarrow 0, \%a \rightarrow$`past`$\}$

We now have a query language for identifying partial syntactic structures. The system uses simple graph matching techniques and additionally implements some query operations inspired by LFGs equations on f-structure constraints. Finally, the system also allows for the definition of dictionaries that can be used to categorize groups of features. This last functionality is less practical for simple querying tasks but proves helpful in section 5.1.6. In that section, the ParTMA query language is integrated into a rule-writing system that enables the user to expand on nodes that match a given query with additional edges and attribute/value pairs. This is used to model the syntax/semantics interface. More concretely, the rules are used to generate semantic structures related to the given syntactic input.

5.1.6 Formal underpinnings of rewriting and expanding ASGs

Now that I have established how to identify partial syntactic structures via the ParTMA querying language, it is also possible to expand those structures by introducing new structures with the same query language, following the idea behind XLE's transfer system. In such a system, the variables identified in a query may be used to add additional relation between nodes, (15a), add attribute/value pairs to a specific node, (15b), or rewrite values of specific attributes, (15c).[7] This means new features are constructed from queries by instantiating the variables used in a query to particular graph representations and using the nodes of these graphs to introduce new relations and attribute-value pairs.

(15) a. $a\ A_1\ b_n \wedge b\ A_2\ c_n \rightarrow a\ A_3\ c_n$

b. $a\ A_1\ b_n \rightarrow b\ A_2\ c_v$

c. $a\ A_1\ b_v \rightarrow a\ A_1\ c_v$

In addition to the operations on existing nodes and attribute/value pairs exemplified in (15), the system also is capable of introducing new structures. This is done by using unbound variables, which automatically instantiate a new, not previously used, index. This is exemplified in (16). The rule shown there introduces a new fact. Example (16b) illustrates this in terms of a set of facts before and after the application of the rule.

(16) a. $a\ A_1\ b_n \rightarrow b\ A_2\ c_n$

b. $\{1\ A_1\ 2\} \rightarrow \{1\ A_1\ 2,\ 2\ A_2\ 3\}$

[7] In this example, $a...z$ are variables over nodes or values (indicated by a subscript n, or a v respectively, where necessary, else they are variables over nodes). $A_1...A_n$ are specific attributes differentiated by the integer in the subscript.

In principle, this system also allows the introduction of separate graphs, but in the implementation presented in this chapter, this capability is not explored (see section 5.2.4 for more details).
The right-hand side of a semantic construction rule is a list of facts. This means that, compared to the left-hand side, which comprises a set of graph paths, the right-hand side encodes partial graphs as a list of facts, e.g.:

(17) $a\ A_1\ b \wedge c\ A_2\ d \wedge e\ A_3\ f \wedge \ldots$

In summary, the syntax of a semantic construction rule consists of a list of query expressions describing a partial graph, as discussed in the previous section, and a list of new facts that are introduced by that rule's right-hand side. In the next section, I explain how the semantic rules designed throughout this thesis can be implemented within the rule system outlined above.

5.1.7 From theory to implementation

In the theoretical chapters of this thesis, rules were encoded without references to nodes. These references have remained implicit to make the rules more easily readable. However, in some cases, indexation was used to differentiate distinct annotations of two (or more) eventualities within the same sentence. Consider the following example from Chapter 3, (19). $E1$ and $E2$ refer to the matrix eventuality and the embedded eventuality respectively, as indicated in (18). The rule says that if a semantic past tense is embedded under another semantic past tense via a `COMP` relation, then this embedded semantic past is interpreted as ambiguous between relative present tense or relative past tense.

(18) John *said*$_E$1 that Mary *was sick*$_E$2

(19) $\texttt{TEMP-REF}_{E1}$ `'past'` : $\lambda t.t \prec t_0 \wedge$
$\texttt{COMP}(E_1, E_2) \wedge$
$\texttt{TENSE}_{E2}$ `'past'` : $\lambda t.\lambda t_0.t \prec t_0 \rightarrow$
$\texttt{TEMP-REF}_{E2}$ $\llparenthesis\ \texttt{present}_t, \texttt{past}_t\ \rrparenthesis$

In the previous section, I presented a concrete syntax (i.e., machine-readable) syntax for semantic construction rules. This concrete syntax of the ParTMA annotation scheme uses graph matching techniques to identify partial syntactic and semantic graphs and rewrites or adds new nodes and features to existing graphs. A concrete example of this is given in Figure 5.7. The blue graph represents the syntactic analysis of a sequence-of-tense sentence such as the one given in (18). The two red graphs are semantic annotations of the two eventualities described in the sentence. The graph headed by the node with the index 6 corresponds to the matrix eventuality, while the graph headed by the node with the index 10 corresponds to the embedded eventuality of Mary being sick. Both of these graphs are anchored in the syntactic graph via the projection relation, a special kind of edge that relates syntactic to semantic graphs. The two red arrows are instantiations of the following default past tense rule in (20a). Example (20b) illustrates the concrete implementation of this rule in the present system. The left-hand side of the rule checks whether there exists an annotation template for the ParTMA annotation scheme, and if so, it assigns the `REF` feature of the `TEMP-REF` sub-graph the value `past`. Checking for existing annotations is relevant for rewriting existing values, a mechanism whose theoretical role in the present annotation scheme has been discussed in Chapter 3.

(20) a. [`TNS-ASP` [`TENSE past`]]→ `TEMP-REF` : ‘`past`’ : $\lambda t.\lambda t_0.t \prec t_0$

b.
```
#a TNS-ASP #b TENSE 'past' &
#a proj #c tmp #d TMP-REF %a
==> #d TMP-REF 'past'
```

As can be seen in the graph, the rule in (20) applies twice to the complex graph given in Figure 5.7. Once for $E1$ and once for $E2$. The proper application of this rule relies on the exact relations between the nodes in the graph (e.g., it would be implausible to link the `TNS-ASP` node of the embedded eventuality to the corresponding semantic node of the matrix eventuality and vice versa). These relations are implicitly marked in the rule format used in Chapters 3 and 4 but are made concrete in this section.

To make this point more clear, I present a second example building on the first one: The sequence-of-tense rule (seen in (84)), a tier-2 rule that applies after the default rule for past tense has already been applied. The rule

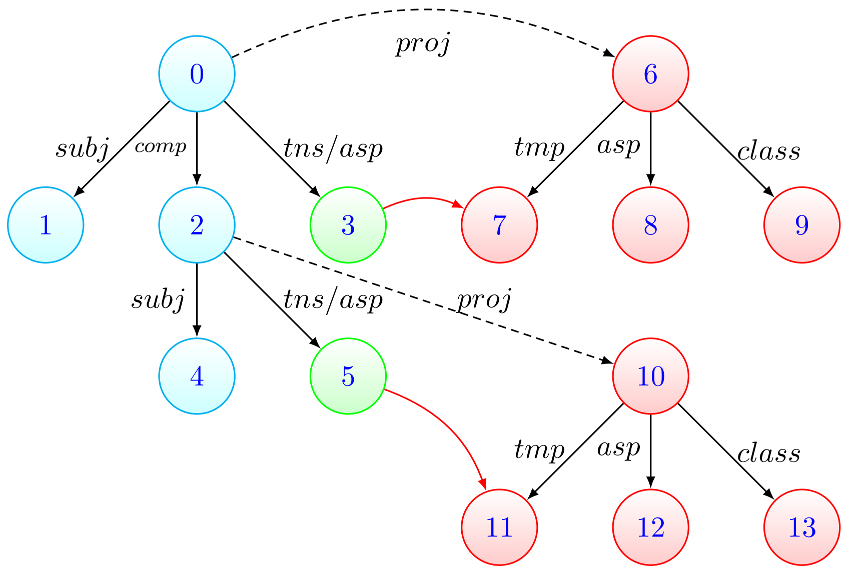

Figure 5.7: Example graph for sequence-of-tense sentence (tier-1)

uses functional application and inside-out functional application to check whether a matrix clause, marked for past tense, governs a past tense in an embedded clause. If so, the rule rewrites the value of the embedded tense to the ambiguous $\{past, present\}$, which corresponds to the meaning family ⦇ `present`, `past` ⦈. Furthermore, the rule specifies that the evaluation time of the embedded tense is the topic time of the matrix tense. The concrete implementation of this rule using the tools presented in this chapter is shown in (21).

(21)
```
#a TEMP-REF ‘past’ &#a ^(tmp>proj>comp) #b !(proj>tmp) #c
TEMP-REF ‘past’
==> #a TEMP-REF {past,present} & #a EVAL #c.
```

The application of this rule is visualized in Figure 5.8. The loop on node 11 illustrates the actual goal of this rule. It indicates that the rule rewrites some information within itself, in this case, the `TEMP-REF` feature. The path highlighted in red from node 11 to node 7 corresponds to the antecedent of the rule in (21). As with the default tense rule, the tier-2 rule instantiates the

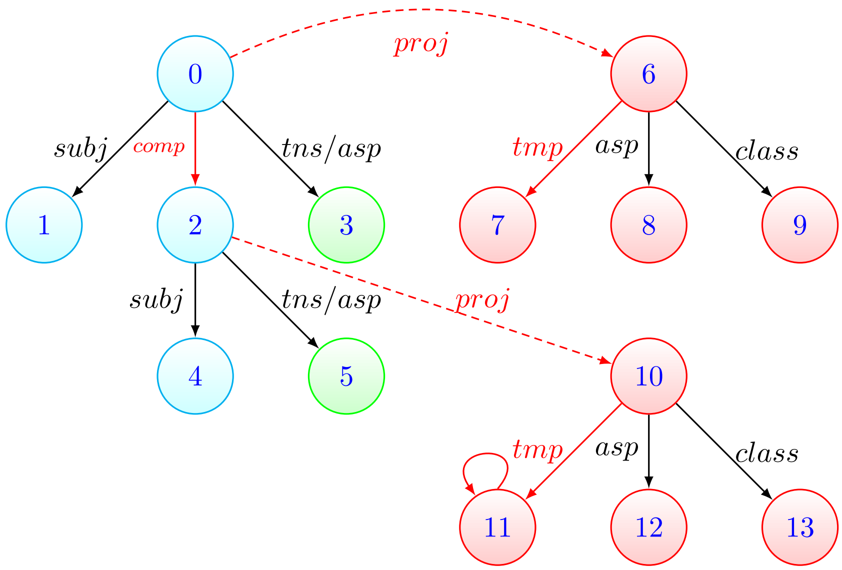

Figure 5.8: Example graph for sequence-of-tense sentence (tier-2)

feature that is annotated on the left-hand side and is defined on the right-hand side. Generally, the procedure follows the following rewrite path via the tier-1 rule and then the tier-2 rule: $unspec \rightarrow tier-1 \rightarrow tier-2$. For the embedded tense of the SOT-example this means: $unspec \rightarrow past \rightarrow \{past, present\}$.
Generally, the goal of the annotation rules, exemplified via the tense example in this section, is to populate the annotation template that has been discussed throughout this thesis and is presented in Figure 5.9. In this section, I have shown how the template can be populated by implementing the rules presented in this thesis within a rewrite system for linguistic information encoded in terms of graphs. I have appealed to research on LFG, universal dependencies, the relation between the two, and research on linguistic annotation more generally.
For the remainder of this chapter, I present the second use of the rewrite system presented in this paper – namely that of semantic interpretation. As of now, I have only presented semantic attribute/value pairs whose interpretation I have explained in chapters 3 and 4. However, to complete the semantic annotation scheme presented in this paper, I also show how these semantic

```
ID [ LEX-ASP    [ class ::= 'event' | 'process' | 'state' | 'unspec'
                  extent ::= 'mom' | 'dur'
                  FPS  [ initP   [ init ::= 'fid'
                                   i-subj ::= 'fid' ]
                         procP   [ proc ::= 'fid'
                                   p-subj ::= 'fid' ]
                         resP    [ res ::= 'fid'
                                   r-subj ::= 'fid' ]
                         rheme   [ head ::= 'fid'
                                   type ::= 'location' | 'path' |
                                   'i-theme' | 'unspec' ] ] ]
     TEMP-REF   [ t-ref ::= 'past' | 'present' | 'future' |
                  'non-past' | 'non-present' | 'unspec'
                  t-restr ::= 'immediate' | 'recent' | 'non-recent'
                  eval ::= 'fid' ]
     ASP-TENSE  [ a-ref ::= 'past' | ...
                  a-restr ::= 'immediate' | ... ]
     VIEWPOINT  [ aspect ::= 'perfective' | 'imperfective'
                  restr ::= 'progressive' | 'habitual' | 'ongoing' |
                  'terminated' | 'culminated' ] ]
```

Figure 5.9: ParTMA eventuality template

features can be translated into formal semantic representations that can be integrated into a broader semantic framework. For this purpose, I use the framework of Glue semantics, which has the additional desirable effect of being compatible with several different semantic formalisms. This is explained in detail in the next section.

5.2 A semantics for the ParTMA annotation scheme

In chapters 3 and 4, I presented an abstract syntax for the ParTMA annotation scheme. In the previous two sections, I have added a concrete syntax that implements the semantic annotation in terms of a query language inspired by tools for syntactic and semantic analysis, such as, for example, treebank analysis tools. The result of the annotation rules is a set of semantic attribute-value pairs, corresponding to Figure 5.9.

The template in Figure 5.9 is, like the f-structure in LFG, a flat structure. This is at odds with the semantic framework postulated for chapters 3 and 4, which is based on the hierarchical tense/aspect structure shown in (22).

(22)

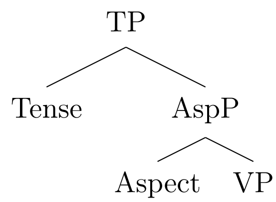

In this section, I provide a semantics for the template given above that follows the general idea presented in the previous chapters and is fully implemented. For this implementation, a semantic formalism called Glue semantics is used. Like much work of this thesis, Glue semantics is also inspired by LFG. In fact, Glue semantics was specifically designed for LFG to account for the flat f-structure, which was assumed to be the input to further semantic processing (Asudeh, 2006; Dalrymple, 1999; Dalrymple et al., 1993).
Glue semantics provides a great deal of flexibility in more than one respect. On the one hand, it has been shown to be an interesting semantic representation for various syntactic formalisms, as, for example, dependencies (Garrette and Klein, 2009; Gotham and Haug, 2019), HPSG (Asudeh and Crouch, 2002), and minimalist syntax (Gotham, 2015). On the other hand, Glue semantics allows for the use of different semantic formalisms, such as DRT (Garrette and Klein, 2009), and its variants (Kokkonidis, 2006; van Genabith and Crouch, 1997), and simple first-order lambda calculus (Dalrymple, 1999). In fact, the only real constraint imposed by Glue semantics onto a given semantic formalism is that it requires a mechanism that parallels functional application and abstraction over variables (or some equivalent mechanism). In the next section, I provide a high-level overview of Glue semantics that describes how these properties are achieved within the framework and will then continue to explain how these properties can be used in the context of this thesis. For a in-depth discussion of the properties of Glue semantics, I refer the reader to Andrews (2009), Asudeh (2006), Crouch and van Genabith (2000), Dalrymple (1999), and Lev (2007), with Andrews (2009) being most relevant to the present discussion of a description-by-analysis Glue approach.

5.2.1 Glue semantics: Compositionality via linear logic

The flexibility of Glue semantics stems from the distinction between a glue language that guides semantic composition and a meaning language whose assembly is guided by this glue language. In LFG, the object that combines these two components is called a meaning constructor and consists of a meaning representation $\mathcal{M}$ and glue representation $\mathcal{G}$. For ease of exposition, I adopt the term meaning constructor for such a pair for the rest of this thesis.

(23) $\mathcal{M} : \mathcal{G}$

As explained above, the meaning language I use in this thesis is that of first-order logic. I follow the traditional formal semantic approach of using lambda expressions to describe partial formulas that must be composed into complete first-order formulas.
The glue language guides this compositional process and dictates the order of combination steps in the meaning language. This role, in Glue semantics, is attributed to linear logic, a so-called resource-conscious logic (Girard, 1987, 1995). To be more precise, various fragments of linear logic have been proposed for use in Glue semantics, although they often converge with respect to certain properties (Crouch and van Genabith, 2000; Dalrymple et al., 1999, p. 98f). The gist of these properties is that they ensure a level of parallelism between the linear logic side and the semantic side of a meaning constructor. How this is achieved is explained in the next section. However, in this thesis, I do not provide a full overview of glue semantics and linear logic but rather provide a sufficient introduction for its use in formal semantics, and particularly in this thesis.[8] To this end, section 5.2.2 introduces the *implicational* fragment of linear logic.

[8] Linear logic has been conceptualized by Jean-YvesGirard, who envisioned it as a tool mainly for (theoretical) computational scientists (Girard, 1987), but also suggested a wide array of uses in other science, e.g., chemistry (Girard, 1995).

5.2.2 The implicational fragment of linear logic

The traditional (initial) approach to Glue semantics relies on deduction rather than lambda calculus to combine meanings (Dalrymple et al., 1993). In this thesis, I follow the more recent trend to guide computations within the domain of lambda calculus via linear logic formulas (Dalrymple et al., 1999). The crucial notion underlying this approach to Glue semantics is that of the Curry-Howard isomorphism (Curry et al., 1958). This isomorphism describes the relationship between inference rules in linear logic and semantic operations in the meaning language. More specifically, linear logic inference rules may be paired with corresponding semantic operations (Crouch and van Genabith, 2000).

Up until now, I have described Glue semantics as a resource-conscious approach. Let us first clarify this term. Resource-consciousness in linear logic is achieved by removing the rules of contraction and weakening from traditional logic. The former rule is used to duplicate premises, and the latter one is used to discard premises. The rules are specified in 24. As a result, in a linear logic proof, all available resources also have to be used, and each resource can only be used once.

(24)

Contraction:	**Weakening:**
$\dfrac{\Gamma, A, A \vdash B}{\Gamma, A \vdash B}$	$\dfrac{\Gamma \vdash B}{\Gamma, A \vdash B}$

In this thesis, I make use of the implicational fragment of linear logic, the only relevant operator of which is that of linear implication: $\multimap$. Atomic and complex terms, *A,B,...* may be combined via this operator to produce complex formulas similar to the regular implication in traditional logic. The resource-consciousness is also reflected in the definition of the rule for linear implication, which generally corresponds to modus ponens (without weakening and contraction as described above). Thus, the formula $A \multimap B$ describes a complex term that *consumes* a resource A to produce a resource, B. Consequently, the deduction in the following example is valid.

(25) $A, A \multimap B \vdash B$

However, the inference $A, A \multimap B \vdash A, B$, which is valid in traditional logic, is not valid in linear logic. This is due to the fact that the resource A is consumed in the inference process, i.e., it disappears from the proof. A common analogy is that of a snack machine. Consider a situation in which you have only one coin, A, left in your purse. The snack machine consumes this coin to return a specifically labeled snack of your choosing, B, that matches the price of the coin. As a result of this process, the coin is gone, but you gained a new resource, namely, the snack of your choosing.[9]

As explained in the previous section, Glue semantics pairs linear logic with a meaning language. In accordance with the Curry-Howard isomorphism (Howard, 1980), the deduction in (25) is paired with functional application, as is illustrated in figure 5.10, on the right. As shown there, a formula containing a linear implication generally corresponds to a function f. The antecedent of the linear implication identifies the required argument (which itself can be any well-formed linear logic term). Given such a resource, the deduction in (25) can be applied, and, correspondingly, the function is applied to the argument a.

The left side of the same figure illustrates the process that corresponds to lambda abstraction: implication introduction. The figure is to be interpreted as follows: One may introduce an assumption A with corresponding semantics x. This assumption receives an index i and can be used freely in any other computations leading to the resource B, where x is an argument in B. In a resource B generated on the basis of the assumption A, this assumption must be reintroduced at some point in the derivation to preserve resource sensitivity. We achieve this by executing a lambda abstraction step on the semantic side and a linear implication on the logic side of a given Glue semantics term in which the assumption forms the antecedent.

There is an apparent relationship between linear logic and typed lambda calculus. In fact, the glue language can be understood as providing types for the meaning language. The following example makes this correspondence

[9] Girard (1995) uses a similar analogy but also explains more complex cases, e.g., such cases in which multiple, or possibly even infinitely many coins exist. However, this is not relevant for the present purposes.

$$\dfrac{\begin{matrix}[x : A]^i \\ \vdots \\ f(x) : B\end{matrix}}{\lambda x.f(x) : A \multimap B} \multimap_{I,i} \qquad \dfrac{f : A \multimap B \qquad a : A}{f(a) : B} \multimap_E$$

Figure 5.10: Implication introduction and elimination

concrete: a transitive verb consumes two resources of type e to form a statement of type t. However, by simply using types as atomic elements on the glue side, certain constraints imposed by the syntax are lost. Thus, the premises shown in (26) yield two different solutions, depending on which argument first combines with the verb.

(26) **John** $e : john$
Mary $e : mary$
kiss $(e \multimap (e \multimap t)) : \lambda x_e.\lambda y_e.kiss(x, y)$

(27) a. $kiss(john, mary)$
b. $kiss(mary, john)$

5.2.3 Linear logic and semantic types

In Glue semantics proper, the linear logic side of meaning constructors does not only encode the type but also makes reference to the role of the underlying resource in the structure of the sentence (Dalrymple et al., 1999). To put it simply, atomic elements in the glue languages can be understood as pairs of types and references to the syntax, here exemplified in terms of the f-structure. Consider (28), there g can be understood as an entity of type e that fulfils the role of the subject of the sentence, whereas h is also of type e but refers to the object of the sentence. The verb is a complex formula that needs to consume g and h to produce some resource f that corresponds to the meaning of the full sentence.

(28) John kisses Mary.

$$f\begin{bmatrix}\text{PRED} & \text{`kiss}\left\langle\text{SUBJ,OBJ}\right\rangle\text{'} \\ \text{SUBJ} & g\begin{bmatrix}\text{PRED} & \text{`John'}\end{bmatrix} \\ \text{OBJ} & h\begin{bmatrix}\text{PRED} & \text{`Mary'}\end{bmatrix}\end{bmatrix}$$

John	$g : john$
Mary	$h : mary$
kiss	$(g \multimap (h \multimap f)) : \lambda x_e.\lambda y_e.kiss(x, y)$

(29) a. $kiss(john, mary)$

In short, this means that the linear logic side of a meaning constructor contains both type information and a reference to the syntax. Within the literature on LFG in Glue semantics, there are two distinct ways of linking meaning constructors to the syntax. In one, the semantic representation is part of the lexicon. This is called the co-descriptive approach and is more prevalent in theoretical Glue semantics (Dalrymple (1999) and Halvorsen and Kaplan (1988) provide some examples as well as central arguments for co-description). Just recently, Dalrymple et al. (2020) have provided a system for integrating co-descriptive Glue semantics within XLE, opening up possibilities for exploring it within the computational linguistics community. On the other hand, there is the description-by-analysis approach, which has been more widely used in computational Glue semantics (most famously Crouch, 2005; Crouch and King, 2006). However, the description-by-analysis approach has also been thoroughly investigated from a theoretical perspective (e.g., Andrews, 2008, 2010; Gotham and Haug, 2019; Wedekind and Kaplan, 1993). In this thesis, I use this latter approach, which is in line with the idea of implementing the syntax/semantics interface as a rule-based mapping from form to meaning, rather than the idea that form and meaning exist as parallel projections. The exact process by which the syntax/semantics interface is established is explained in the next section.

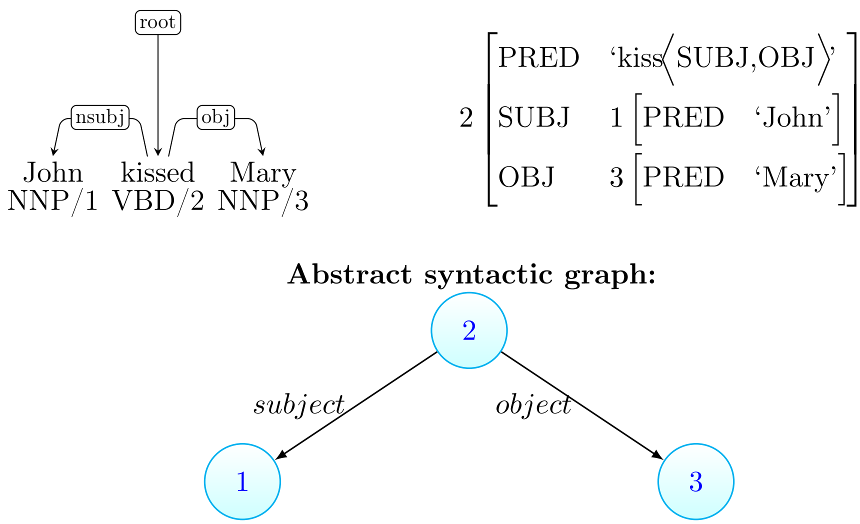

Figure 5.11: Parallelized syntax for: *John kissed Mary*

5.2.4 The syntax/semantics interface in Glue semantics

The syntax/semantics interface relies on the indexation of syntactic structures. The previous section shows that this indexation is also the foundation for generating Glue semantics terms. This process is highly relevant for the rest of this thesis. For this reason, I focus mainly on this feature of Glue semantics. Note that the process of establishing the syntax/semantics interface is parallel in LFG and UD. In fact, the syntax/semantics interface presented in this section is basically designed for abstract syntactic graphs extracted from either LFG or UD structures. Consider Figure ?? in which nodes in the dependency tree have been aligned with the labels of partial f-fstructures via their indices (1,2,3). We can use these indices to form the meaning constructors in (30).

(30) **John** $1 : john$
Mary $3 : mary$
kiss $(1 \multimap (3 \multimap 2)) : \lambda x_e.\lambda y_e.kiss(x, y)$

The proof that can be derived with these meaning constructors is given in 5.12, where each line corresponds to a deduction step, with the result written below the line (i.e., a proof tree). The result may be recursively picked up for a new deduction step resulting in the given proof structure. It clarifies the relation between implication elimination on the linear logic side and its lambda calculus counterpart functional application on the meaning side.

$$\dfrac{\dfrac{\lambda x.\lambda y.loves(x,y) : 1 \multimap (3 \multimap 2) \qquad john : 1}{\lambda y.kiss(john,y) : 3 \multimap 2} \qquad mary : 3}{loves(john,mary) : 2}$$

Figure 5.12: Derivation of *John loves Mary.*

In the spirit of LFG, we would expect there to be some function or projection that maps the syntactic representations to their semantic counterpart. Although the end result is intuitively clear, at least in the simple example above, the mapping (function) from syntax to semantics still needs to be explained. As stated in section 5.2.3, the syntax/semantics interface proposed in this thesis follows the tradition of description-by-analysis. The main hallmarks of the approach that this thesis implements are:

1. A two-stage process for processing the semantics
   - i) A set of rules mapping syntactic representation onto semantic features
   - ii) A set of rules interpreting semantic feature structures
2. Use of propositional linear logic

The first hallmark describes the conscious decision of treating the syntactic input as such: syntactic information. This information is translated into semantic information that is stored in the s-structure, i.e., the template for semantic features introduced in Figure 5.9. In Chapter 2, I noted that there exists a tension between form and meaning in current approaches to annotating tense and aspect. I have explored the corresponding issues in chapters 3 and 4. There, I focused on establishing semantic feature structures based on formal semantic insights. In section 5.1.7, I showed how the semantic feature structure can be populated via syntactic rewriting, i.e., the first component of the semantic processing component.

This section explains the foundations of the second hallmark, which is crucial for establishing the second stage of the first hallmark. The next few paragraphs are concerned with explaining the semantic interpretation rules in more detail. These rules use the same general mechanism as the syntax/semantics interface rules. The difference is that the rules explained now do not map onto semantic features but on Glue semantics meaning constructors for calculating meaning representations. Thus, Glue language terms are instantiated via querying the feature-based semantic representation that is the output of the machinery described in section 5.1, i.e., the template in Figure 5.9. Rather than semantic construction rules, these rules can be dubbed semantic interpretation rules and are cross-linguistically invariant. They are simply a set of rules that translate a semantic feature structure into meaning constructors. In principle, this boils down to adding a new node to each sub-structure of the annotation template that has the attribute `GLUE` and whose value is the corresponding meaning constructor.

This is demonstrated in (32a) in terms of the interpretation rule for a past tense annotation.[10] The right-hand side of this rule describes the resulting meaning constructor. It illustrates that the linear logic side of the meaning constructor is instantiated based on the reference to the variables used in the left-hand side query. Thus, given the indexation in (31), this rule yields the meaning constructor in (32b).[11] This meaning constructor picks up a resource, $2 \multimap 1$, and produces the resource $3 \multimap 1$. The resource 3 corresponds to the evaluation time and 1 to the semantics of the whole eventuality. This means the meaning constructor takes a proposition of type *st* and returns a proposition of *st*. Due to the Curry-Howard isomorphism, this can be confirmed by looking at the meaning side of the meaning constructor. The resulting proposition of type $3 \multimap 1$ expects as input the evaluation time

[10] This rule serves illustrative purposes and is revised in section 5.3.

[11] The string representation used in the meaning constructor is designed for ease of use, i.e., not relying on any special symbols. In this chapter, I generally present a more conventional representation of the corresponding expression, here in (32b).

of the given annotation, here, simply the *now*. This gives us a semantics that relates the reference time of the annotated eventuality to the time of evaluation, i.e., the prototypical semantics of past tense.

(31)
$$0\left[\begin{array}{ll} \text{PRED} & \ldots \\ \ldots & \\ \text{SEM} & 1\left[\begin{array}{ll}\text{TEMP-REF} & 2\left[\begin{array}{ll}\textbf{t-ref} & ::= \text{'past'} \\ \textbf{t-restr} & ::= \text{'unspec'} \\ \textbf{eval} & 3\left[\text{TIMEX 'now'}\right]\end{array}\right]\end{array}\right]\end{array}\right]$$

(32) a.
```
#a SEM #b TEMP-REF #c T-REF 'past' & #c EVAL #d ==>
#c GLUE
[/P_<s,t>.[/s_s.Er_s[before(r,s) & P(r)]]] :
((#c -o #b) -o (#d -o #b)).
```

b. $\lambda P_{st}.\lambda s_s.\exists r_s[r \prec s \wedge P(r)] : (2 \multimap 1) \multimap (3 \multimap 1)$

As mentioned above, this approach is heavily inspired by Andrews's (2008) description-by-analysis approach (see also Andrews, 2010). One of the main points of Andrew's approach to description-by-analysis is that the rules are bi-directional. This means that in addition to producing a semantic representation from a syntactic input, these rules also allow for creating a syntactic representation off of a Glue semantics derivation. This is not integrated into the present system – implementing and testing this is left for future work.[12]

A second property that the present system shares with Andrews's (2008) is that each rule has a dedicated central node. This means that the left-hand side is required to form a connected graph structure for which one (or more) paths exist between any pair of nodes. More concretely, Andrews says that each left-hand side of a given semantic interpretation rule (he calls them Semantic Lexical Entries: SLEs) requires a central node from which all other

[12]Although not relevant for the present purposes, this has interesting implications for generation from a glue semantics proof, making the underlying proofs first-class semantic objects (Asudeh et al., 2008).

nodes are accessible. This principle is and has been adhered to in the rules presented in this thesis. Andrews, 2008 proposes additional constraints on the relation between the center and the other nodes in a rule, but these are not of central relevance here.

Using the simple description-by-analysis glue semantics framework introduced above, we can define semantic interpretation rules (corresponding to Andrew's SLEs). As initially explained, these rules apply to the feature representation resulting from the tense/aspect annotation. In the next section, I show how Glue semantics can be used to derive the appropriate semantics from the attribute-value structure of the ParTMA template in Figure 5.9.

5.2.5 Tense/aspect annotation: from structure to meaning

Recall that the goal of the present annotation scheme is to provide a simple attribute/value matrix as frontend for annotation while relying on a more articulated semantic structure for interpretation in the back. Adhering to this conceptualization, the tense/aspect annotation template can be translated into a tree structure, as illustrated (in a simplified manner) in Figure 5.13. As shown there, each node in the hierarchical tree structure is related to the top node in terms of its linear logic representation. The tense and aspect operators lift the calculation to the next higher node. Type-wise, both tense and aspect take an element of type $< s, t >$ and return an element of type $< s, t >$. Note that the left node of each tree corresponds to some annotation, i.e., a meaning constructor (see 32a). The branching nodes, on the other hand, refer to intermediate results of the computation, eventually resulting in the element 0 of type t. The linear logic types given in the tree are derived from the nodes of the tense/aspect template shown in (33), a simplified version of the template in Figure 5.9. Note that the final step in the interpretation is the combination with the time of evaluation $t*$ that is part of the annotation of temporal reference.

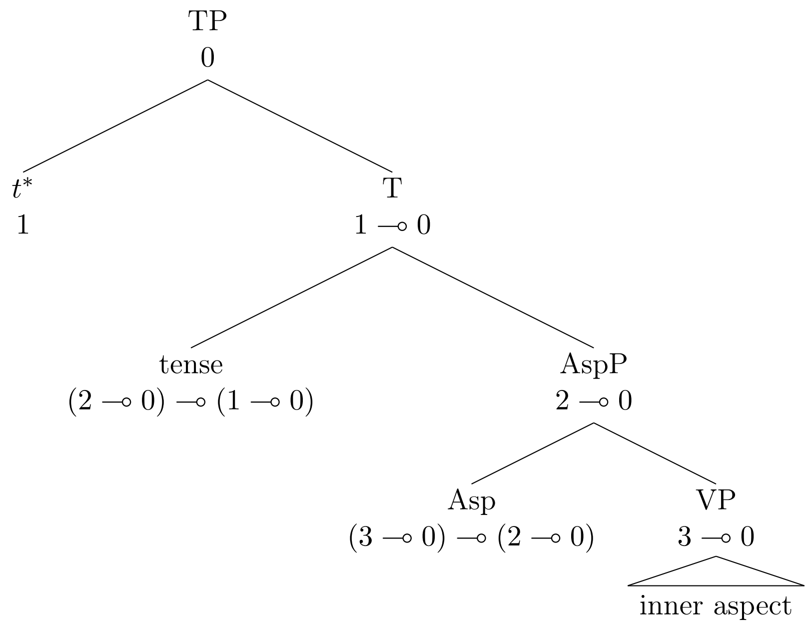

Figure 5.13: Order of composition in the ParTMA annotation template

(33)
$$
0\begin{bmatrix}
\text{TEMP-REF} & 2\begin{bmatrix} \text{...} & \\ \text{EVAL} & 1\begin{bmatrix}\text{...}\end{bmatrix} \end{bmatrix} \\
\text{VIEWPOINT} & 3\begin{bmatrix}\text{...}\end{bmatrix} \\
\text{LEX-ASP} & 4\begin{bmatrix}\text{...}\end{bmatrix}
\end{bmatrix}
$$

Although this simple structure provides a nice approximation of the system I want to present, it omits several crucial details for the treatment of some of the phenomena discussed in this thesis. In the next few sections, I elaborate on these details. As illustrated in chapter 3 and 4, tense and aspect both can be represented relatively straightforwardly as quantifiers and, thus, can be treated similarly from a formal perspective. The specifics of these two categories are discussed in sections 5.3 and 5.4. Lexical aspect, on the other hand, provides a more significant challenge and is discussed more extensively in section 5.4.2.

5.3 The semantics of temporal reference

Once again, I begin from the top of the hierarchical TAM tree to explain the exact treatment of tense within the whole tense/aspect structure, i.e., first, I explain how the temporal reference feature structure is translated into a quantifier over tenses. In the process, I address how temporal restrictors, e.g., temporal remoteness morphemes and contextual restriction on time variables, are implemented in the present semantic framework. Secondly, I explain the treatment of absolute and relative tenses. Finally, I address how to deal with aspectual tenses, e.g., the perfect or prospective aspect.

5.3.1 Temporal reference and temporal restriction

To begin with, let us recall the general architecture of our temporal semantic framework, which is based on the quantificational tense approach introduced in Kusumoto (1999, 2005). For this purpose, recall example (32a), here repeated in (34), and the corresponding partial attribute/value matrix that the rule applies to, repeated in (35).

(34) a.
```
#a SEM #b TEMP-REF #c T-REF 'past' & #c EVAL #d ==>
#c GLUE [/P_<s,t>.[/s_s.Er_s[before(r,s) & P(r)]]] :
((#c -o #b) -o (#d -o #b)).
```
b. $\lambda P_{st}.\lambda s_s.\exists r_s[r \prec s \wedge P(r)] : (2 \multimap 1) \multimap (3 \multimap 1)$

(35)
$$_{0}\left[\begin{array}{ll}\text{PRED} \ldots & \\ \ldots & \\ \text{SEM} & _{1}\left[\text{TEMP-REF} \;\; _{2}\left[\begin{array}{ll}\textbf{t-ref} ::= \text{'past'} & \\ \textbf{t-restr} ::= \text{'unspec'} & \\ \textbf{eval} & _{3}\left[\text{TIMEX 'now'}\right]\end{array}\right]\right]\end{array}\right]$$

The structure given in (35) provides two indices, 2 and 3. These correspond to the topic (or reference) time and the evaluation time, respectively. As the meaning constructor in (34) indicates, the tense quantifier saturates the topic time of its argument and returns an element that requires an evaluation time to form the meaning corresponding to the full eventuality. This is in line with Kusumoto's (2005) quantificational tense account presented in chapter 3. However, this iteration of the tense rule does not allow for the treatment of temporal restrictors. Let us examine the case of temporal remoteness morphemes. As mentioned in Chapter 3, Temporal remoteness markers indicate a certain distance between the evaluation time and the reference time leading to more a more fine-grained encoding of temporal reference (Bochnak and Klecha, 2015; Cable, 2013). For the sake of simplicity, I model these in terms of the two-place predicates `close/2`, `recent/2`, and `far/2`, partially borrowed from Bochnak and Klecha (2015).[13]

(36) a. $\lambda t.\lambda t'.close(t,t')$, *iff* the temporal distance between t and t' is close.

b. $\lambda t.\lambda t'.recent(t,t')$, *iff* the temporal distance between t and t' is moderate.

c. $\lambda t.\lambda t'.far(t,t')$, *iff* the temporal distance between t and t' is far.

The main idea of temporal restrictors is that they, as the name suggests, modify the restrictor of the temporal quantifier. To achieve this goal, we have to disentangle the restrictor from the quantifier to make it modifiable. On the meaning side, this process is straightforward, as shown in (37).

(37) $\lambda P_{st}.\lambda s_s.\exists r_s[r \prec s \wedge P(r)]$

a. $\lambda P_{sst}.\lambda Q_{st}.\lambda s_s.\exists r_s[P(r)(s) \wedge Q(r)]$

b. $\lambda t_s.\lambda t'.t \prec t'$

Since the temporal quantifier now takes two arguments instead of one, this needs to be reflected in the glue side due to the Curry-Howard isomorphism, as explained in the previous section. Since the restrictor of the temporal

[13] Other than in the cited work, those predicates do not return a degree but rather yield a truth value. Through this, some benefits of Bochnak and Klecha's (2015) work, pertaining to the context sensitiveness and vagueness, are lost. Integrating a more detailed semantic analysis is left for future work. At this point, I focus on the general idea of integrating such modifications into the present work.

quantifier is relating the topic time to the evaluation time, we can straightforwardly supply it with an appropriate linear logic formula. In accordance with the indexation in (35), the linear logic side of the restrictor is derived from the index associated with the topic time and the index associated with the evaluation time. Since this element is now missing from the quantifier, we have to modify the quantifier so that it takes the restrictor as its first argument (see (39)).

(38) $\lambda t_s.\lambda t'.t \prec t' : (2 \multimap (3 \multimap 4))$

(39) $\lambda P_{sst}.\lambda Q_{st}.\lambda s_s.\exists r_s[P(r)(s) \wedge Q(r)]$:
$(2 \multimap (3 \multimap 4)) \multimap ((2 \multimap 1) \multimap (3 \multimap 1))$

As shown in section 5.2.4 (example (34)), `GLUE` nodes are anchored in the indices of the tense/aspect attribute/value structure. The decomposition of the temporal quantifier leads to a surplus of meaning constructors that need to be anchored in the tense/aspect structure. To remedy this situation, the semantic interpretation rules add covert nodes to the structure that serve as references for the new resources that are introduced to the computation. This is visualized in the tree in (40). Two new indices, 4 and 5, are introduced to the structure using covert nodes. This is indicated by marking them with a hyphen, as in `T-REF'`, which points to the additional covert node for the `T-REF` feature. Similarly, `T-RESTR'` provides an anchor for the modification encoded in the restrictor. These attributes denote values rather than nodes and thus cannot serve as anchors for their respective meanings by themselves. The missing subscript indicates this in the tree.

The structure in (40) corresponds to the *tense* node in Figure 5.13. This can be verified by combining the resources provided in the tree structure in (40). The result is the same as the linear logic type given to the tense node in Figure 5.13, except for different indices: $(2 \multimap 1) \multimap (3 \multimap 1))$.

(40)

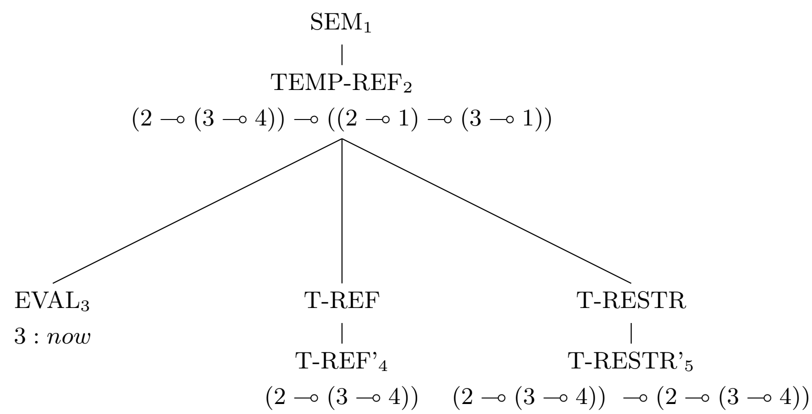

This means the general framework given in Figure 5.13 holds up under the assumption that we expand its subtrees such that we have enough resources to provide all meaning constructors necessary for the computation. This expansion is local within the sub-tree and thus does not require modification of the overall tense/aspect annotation. In fact, such treatment is commonplace in Glue semantics, as the following quote confirms: "One use of the semantic projection is to provide a place to put some non-syntactic attributes, so that each atomic sub-formula of the constructors [...] involve a different structural location" (Andrews, 2010). In this quote, the semantic projection can be understood as an expansion of certain nodes in the f-structure, the book-keeping of which is attributed to a separate projection (albeit, in practice, the semantic projection and the ϕ-projection use the same formal tools in computational LFG, in particular, in XLE).[14]

One last remark needs to be made to close off the topic of modifying the restrictor of a temporal quantifier. In principle, an arbitrary number of modifications may be imposed on the restrictor. For example, a contextual restriction may be added (on top of a temporal remoteness restriction). These modifications always have the same linear logic type. This means contextual

[14] The introduction of additional nodes could be removed or at least reduced by storing meaning constructors in sets rather than as single values of the `GLUE` feature. In the present example, the node with the index 5 is not required to specify any meaning constructors but only exists to attach the meaning constructor for the temporal restrictor. This node could be removed, and the temporal restrictor could be added to the set of meaning constructors in the `T-REF` node. Such optimizations are left for future work.

restriction can be simply added like so:

(41) $\lambda P_{sst}.\lambda t_s.\lambda t'_s.C(t) \wedge P(t)(t') : (2 \multimap (3 \multimap 4)) \multimap (2 \multimap (3 \multimap 4))$

The full interpretation for the annotation of absolute tenses (tenses with *now* as evaluation time) is given in (42) in terms of three rules. The first rule introduces the resources for the temporal quantifier, the quantifier itself, and its restrictor. The second rule applies if the feature `T-RESTR` has been annotated, indicating that the restrictor needs to be modified. It then introduces a corresponding resource. The final rule saturates the second argument of the tense quantifier with the evaluation time, in this case, simply the constant *now*.[15]

(42) a.
```
#a SEM #b TEMP-REF #c T-REF 'past' & #c EVAL #d ==>
#c T-REF' #e &
#c GLUE
[/P_<s,<s,t».[/Q_<s,t>.[/s_s.Er_s[P(r,s) & Q(r)]]]] :
((#c -o (#d -o #e)) -o ((#c -o #b) -o (#d -o #b))) &
#e GLUE
[/t_s.[/r_s.before(t,r)]] :  (#c -o (#d -o #e)).
```
b.
```
#a SEM #b TEMP-REF #c T-RESTR 'remote' & #c EVAL #d
==> #c T-RESTR' #e &
#e GLUE
[/P_<s,<s,t».[/t_s.[/r_s.far(t,r) & P(t)(r)]]] :
((#c -o (#d -o #e)) -o (#c -o (#d -o #e))).
```
c.
```
#a EVAL #b TIMEX 'now' ==> #b GLUE now :  #b.
```

5.3.2 The evaluation time in matrix and embedded contexts

The annotation of the evaluation time is another noteworthy feature of the implementation presented in this chapter. This annotation allows us to

[15] Of course, these three rules do not cover the whole spectrum of tense annotation but rather are examples of the three main steps: i) Introduce a temporal quantifier and a restrictor in accordance with the tense annotation, ii) modify the restrictor where appropriate, and, iii) closing of the tense by saturating the evaluation time.

differentiate between absolute and relative tenses, as well as (some) temporal anaphora. Consider the following examples (taken from Chapter 3). The first two examples, (43) and (44), based on Kusumoto (2005), shed more light on the treatment on absolute and relative tenses, including the ambiguity resulting from the SOT phenomenon (see (44)). Examples (45) to (46) exemplify the treatment of aspectual tenses. Specifically, prospective aspect is illustrated in (45) and treatment of (past) perfect aspect is demonstrated in (46).

(43) Hillary married a man who became president of the U.S.

a. $\exists t[t \prec t_0 \wedge \exists x[man(x,t) \wedge marry(h,x,t) \wedge$
$\exists t'[t' \prec t_0 \wedge become - president(x,t')]]]$ (Kusumoto, 2005)

(44) Tom said that Karen was dancing.

a. $\exists t[t \prec t_0 \wedge say(tom, \exists t'[t' \prec t \wedge dance(k,t')], t)]]$

b. $\exists t[t \prec t_0 \wedge say(tom, dance(k,t), t)]]$ (Kusumoto, 2005)

(45) Peter was going to write a letter.

a. $\exists t[t \prec t_0 \wedge \forall t'[t \prec t' \wedge plan(p,t') \rightarrow write(p, letter)]]$

(46) We arrived at Mwangi's house, but he had (already) left.

a. $\exists t[t \prec t_0 \wedge arrive(we, house, t) \wedge \exists t'[t' \prec t \wedge leave(m,t')]]$
(Cable, 2013)

The gist of example (43) is that the two tenses are independently interpreted absolute tenses. The machinery presented above naturally accounts for this. An important feature is that the `EVAL` feature provides an independent copy of the externally provided temporal expression referring to the speech time. As shown in more detail in Chapter 3, this captures the specifics of the analysis presented in (Kusumoto, 2005).

Let me now show how the semantics of the annotation can be straightforwardly integrated into a basic Glue treatment of relative clauses to deal with cases such as (43). The treatment of relative clauses is based on Dalrymple (2001), which treats the relative pronoun *who* as an instance of predicate modification in the sense of Kratzer and Heim, 1998). This allows us to deal with examples such as The approach is shown in (47). The instantiation is based on Figure 5.14.

(47) $\lambda P_{et}.\lambda Q_{et}.\lambda x_e.P(x) \wedge Q(x) : (6 \multimap 7) \multimap ((3 \multimap 8) \multimap (6 \multimap 7))$

Using this meaning constructor in our established system leads to the proof in Figure 5.15. The resulting meaning is given in (48), which corresponds to the predicted meaning given in (44).

(48) $\exists t[t \prec \mathit{now} \wedge \exists x[\mathit{man}(x,t) \wedge$
$\exists t'[t' \prec \mathit{now} \wedge$ *become-president(x,t')* $] \wedge \mathit{marry}(h,x,t)]]$

As this calculation shows, the semantic interpretation can be built up so that it can do both, preserve the specifics of Kusumoto's (2005) approach to quantificational tenses, as well as maintain compatibility with general strategies developed within the Glue semantics literature. I will maintain this approach for the rest of this section to illustrate the compatibility of the system presented in this thesis with general ideas from formal semantics.

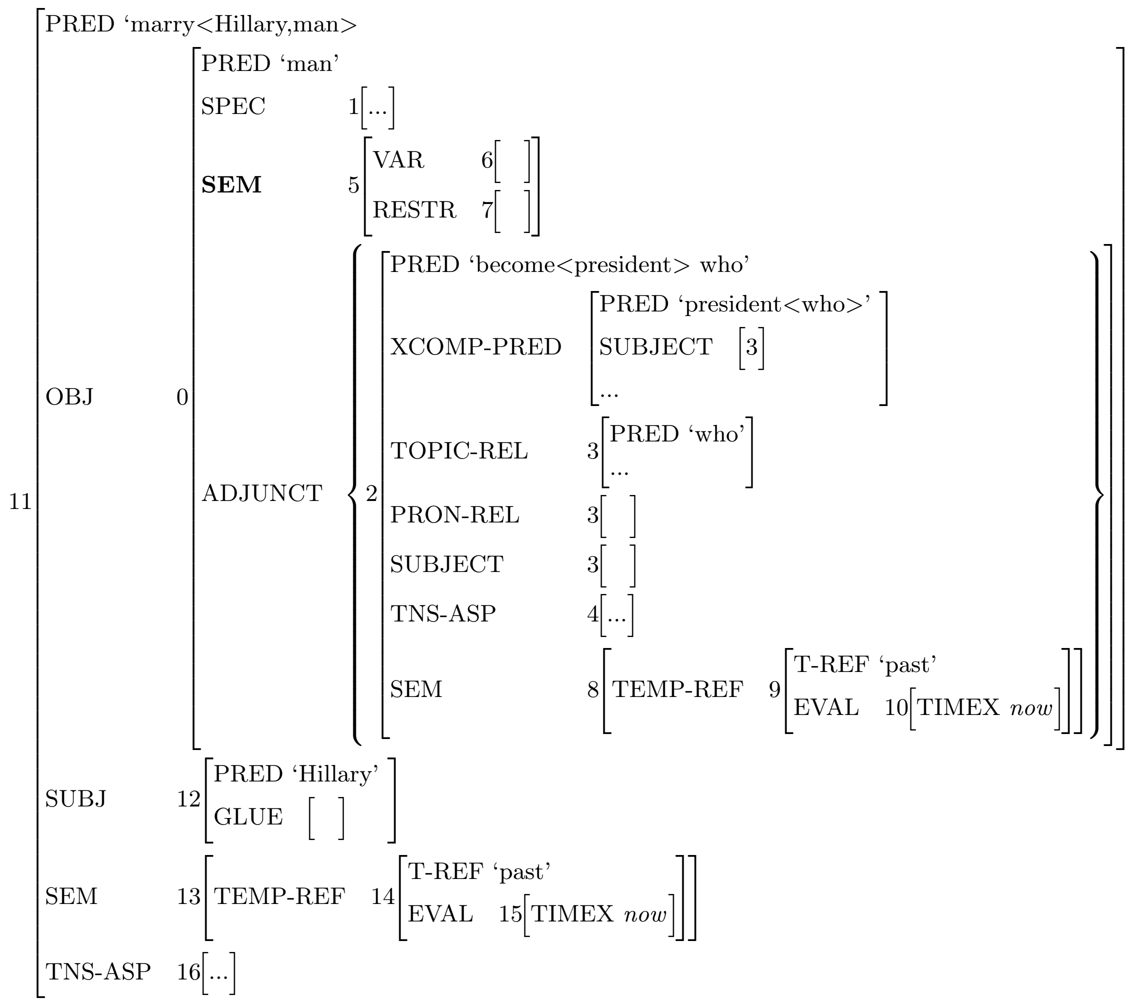

Figure 5.14: Annotation: *Hillary married a man who became president.*

5.3.3 Interpreting SOT-sentences

The second example through which I want to illustrate the implementation of formal semantic insights in the present annotation scheme is the annotation of the SOT phenomenon, where the tense of the matrix clause constrains the embedded tense. I have proposed to model this constraint as an ambiguity between a relative present tense and a relative past tense. Thus, it deviated slightly from its role model, which encoded the ambiguity in terms of a difference between a relative past tense and a deleted tense, essentially leading to a situation where the matrix tense also binds the temporal variable of the embedded predicate. The implementation presented in this chapter will make the choice of the former system more transparent.

The general structure for the SOT phenomenon is given in Figure 5.16 in terms of an enumerated example. A crucial role in this kind of example is played by the propositional attitude verb, which serves as an eventuality in its own right but also introduces the important `COMP` relation between itself and the embedded eventuality. In chapter 3 I provided a very simple semantics for the propositional attitude verb *(to) say*:

(49) $[\![\ say\ ]\!]^{g,w} = \lambda P_{st}.\lambda x_e.\lambda t_s.say(x, P(t), t)$

The main role of this lexical entry is to saturate the evaluation time of the embedded eventuality.[16] Provided this information, the corresponding linear logic side is given in (50), relative to Figure 5.16. The attitude verb takes as an argument the embedded proposition with an unsaturated evaluation time slot and returns a resource that requires a subject to yield a proposition that can be modified by a tense quantifier (type according to the template: $(7 \multimap 6) \multimap (8 \multimap 6)$).

(50) a. $\lambda P_{st}.\lambda x_e.\lambda t_s.say(x, P(t), t) : (7 \multimap 9 \multimap (1 \multimap (7 \multimap 6))$

The `T-REF` feature ⦇past,present⦈ provides a so-called meaning family that leads to an ambiguous reading. In the current implementation, meaning

[16] Recall from chapter 3 that the lexical entry is stripped of its modal ingredients as far as possible, but nothing hinges on this simplification.

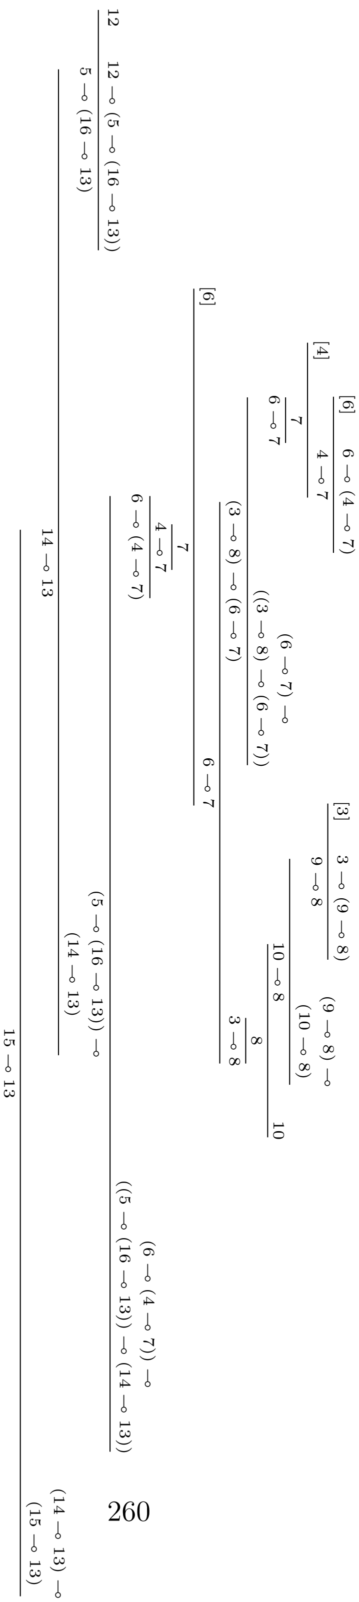

Figure 5.15: Derivation for *Hillary married a man who became president*

$$
0\left[\begin{array}{ll}
\text{PRED} & \text{`say<Tom,dance>} \\
\text{SUBJ} & 1\left[\text{PRED `Tom'}\right] \\
\text{COMP} & 2\left[\begin{array}{ll} \text{PRED} & \text{`dance<Karen>} \\ \text{SUBJ} & 4\left[\text{PRED `Karen'}\right] \\ \text{TNS-ASP} & 5\left[...\right] \\ \text{SEM} & 6\left[\text{TEMP-REF}\quad 7\left[\begin{array}{l}\text{T-REF `past'} \\ \text{EVAL } 8\left[\text{TIMEX } \textit{now}\right]\end{array}\right]\right]\end{array}\right] \\
\text{TNS-ASP} & 3\left[...\right] \\
\text{SEM} & 9\left[\text{TEMP-REF}\quad 10\left[\begin{array}{l}\text{T-REF } (\!|\text{past,present}|\!) \\ \text{EVAL } \left[7\right]\end{array}\right]\right]
\end{array}\right]
$$

Figure 5.16: Annotation: *Tom said that Karen was dancing*

families are implemented as sets of semantic representations of the meaning side of a meaning constructor. Thus, the meaning constructors satisfy the condition of pointwise functional application that all elements of the set are of the same type. Given the hidden node `T-REF'` which introduces a new index, 11, for Figure 5.16, we get the meaning constructor in (51a) for the restrictor of the temporal quantifier of the embedded tense. When combined with the temporal quantifier, the ambiguity in the annotation is percolated up the computation, as shown in (51c).

(51) a. $10 \multimap (7 \multimap 11) : (\!|\lambda t.\lambda t'.t \prec t', \lambda t.\lambda t'.t \circ t'|\!)$

b. $(10 \multimap (7 \multimap 11) \multimap ((10 \multimap 9) \multimap (7 \multimap 9)) :$
$\lambda P_{sst}.\lambda Q_{st}.\lambda s_s.\exists r_s[P(r)(s) \wedge P(r)]$

c. $(10 \multimap 9) \multimap (7 \multimap 9) :$
$(\!|\lambda Q_{st}.\lambda s_s.\exists r_s[s \prec r \wedge Q(r)], \lambda Q_{st}.\lambda s_s.\exists r_s[s \circ r \wedge Q(r)]|\!)$

In Figure 5.16 the annotation is congruent with the semantics that is intended for the sentence. The evaluation time of the embedded eventuality is annotated as the topic time of the matrix eventuality. Since there is no free copy of the topic time of the matrix clause to saturate the embedded clause, the computation requires that the topic time of the embedded predicate is saturated elsewhere, namely, via the semantics of the propositional-attitude predicate, as shown in (50).

There is one caveat to this solution: Rather than sharing the topic time between matrix and embedded eventuality in the simultaneous reading, the topic time of the matrix clause plays the role of the evaluation time of the embedded tense. This is a stipulation that arguably is not desired from a theoretical perspective since it introduces an additional temporal variable that is not needed in Kusumoto's (2005) approach. Furthermore, this can be seen as introducing a covert present tense into the derivation. To circumvent this problem, let me briefly sketch potential solutions more in line with the formal semantic literature.

(52) $[\![\textit{ undefined }]\!]^{g,w} = \lambda t_s.\lambda t'_s.t = t'$

This value simply equates whatever is given by the matrix `COMP` verb to the topic time of the embedded eventuality. Aside from this minor difference, this approach remains the same as the one previously illustrated and maintains exactly the same compositional process, resulting in the solution given in (53). In this approach, the only difference is the choice of restrictor for the embedded tense.

(53) $(10 \multimap 9) \multimap (7 \multimap 9)$:
$(\!|\lambda Q_{st}.\lambda s_s.\exists r_s[s \prec r \wedge Q(r)], \lambda Q_{st}.\lambda s_s.\exists r_s[s = r \wedge Q(r)]|\!)$

The second way to stay closer to the actual formal semantic insights cited as inspiration (Kusumoto, 1999, 2005) is to provide a modified interpretation for the feature *undefined*, namely, the identity function. This would make it so that the `T-REF` node in the annotation would simply percolate the embedded semantic object up the tree without modifying it. The corresponding lexical entry is given in (54).

(54) $[\![\textit{ undefined }]\!]^{g,w} = \lambda P_{st}.\lambda t_s.P(t)' : (10 \multimap 9) \multimap (7 \multimap 9)$

Two distinct semantic structures would need to be generated in such a framework. One that introduces the meaning constructors and invisible nodes associated with the *past* feature(in particular, `T-REF'` and the corresponding `GLUE` node) and one structure associated with the *undefined* feature, which simply attaches an identity function to the `TEMP-REF` node.

The benefit of meaning families is that ambiguity in the ParTMA annotation template is restricted to the introduced meaning constructors. As shown in (53), meaning families are, in fact, intended to be encoded on the meaning side of a meaning constructor. Moving the ambiguity to the temporal quantifier rather than the restrictor thus poses a challenge since there is a type mismatch between a temporal quantifier and the identity function resulting in the aforementioned structural ambiguity within the semantic feature structure.

We now have two possibilities for dealing with the *undefined* feature. The first one preserves the general architecture of the ParTMA annotation scheme and its interpretation. The second one results in a marginally more simple semantic representation by omitting an additional quantifier and equation constraint.

For the semantics in this thesis, I propose to use the second variant of the *undefined* feature value but make it mutually exclusive to any other feature or feature set.[17] Thus, the benefits of the feature can be reaped while the pitfalls are avoided. Concretely, in an annotation where temporal reference is unspecified, the computation does not break down due to a missing link in the hierarchical structure proposed for the semantic composition process.

This is also in line with the feature table presented in chapter 3. The only shortcoming of this approach is that it makes the analysis of the SOT illustrated in (53) impossible.[18] Overall, the annotation and interpretation of the SOT phenomenon can be accounted for in different ways. In this thesis, I treat it as an ambiguity of the restrictor of a temporal quantifier. This

[17]Of course, there is also the possibility to introduce a temporal reference feature , for example,*equal* which provides the restrictor in (52), in addition to the special *undefined* feature.

[18]It might make sense, to implement `TEMP-REF` in such a way, that its structure can be shared between clauses. This approach would naturally account for the SOT phenomenon but introduces structural ambiguity to the ParTMA template, resulting in all kinds of different complexities. For this reason, I leave this line of thought for future work.

captures the explored features of the SOT phenomenon while preserving the key features proposed for the ParTMA annotation scheme, i.e., encoding ambiguity at the level of semantic representation or only as a meaning family of features. Further refinements are left for future work.

5.3.4 Aspectual tense

The last example of temporal annotation that I want to discuss in this chapter is that of aspectual tenses, such as the prospective aspect and the perfect aspect. In this thesis, I assume a framework where aspectual tenses are interpreted between actual tenses and aspect in the semantic composition and are treated generally in the same fashion as tenses, i.e., existential quantifiers with a temporal restrictor (Klecha, 2016). In other words, they are also simply instances of temporal reference. They follow the principles presented in the previous section. Thus, for a sentence like (55), we would expect two instances of temporal reference. There, the order would be determined by annotating the evaluation time accordingly. As the annotation in Figure 5.17 illustrates, the reference time provided by the past tense encoded in 2 plays the role of the evaluation time of the future temporal reference encoded in 4. This is further spelled out in the meaning constructors in (56a) and (56a).

(55) John **was** going to visit his aunt. [PAST [FUT]]

(56) a. $(2 \multimap 1) \multimap (3 \multimap 1)$:
$\lambda Q_{st}.\lambda s_s.\exists r_s[s \prec r \wedge Q(r)]$

(57) a. $(4 \multimap 1) \multimap (2 \multimap 1)$:
$\lambda Q_{st}.\lambda s_s.\exists r_s[s \prec r \wedge Q(r)]$

Together with the SOT example from the last section, this shows that temporal reference is treated uniformly across corresponding markers. This also illustrates the difference between the different possible triggers of temporal reference. Traditional semantic tenses check whether they are syntactically governed by a temporal anchor. Some languages then are interpreted as *non-future* tenses (SOT vs. non-SOT languages). Compositional relative tenses

```
0 [ PRED      'go<John,visit>'
    TNS-ASP   [ TENSE past PROG +_, PERF -_ MOOD indicative ]
    XCOMP     [ PRED 'visit<John,mother>'
                SEM 1 [ TEMP-REF 2 [ t-ref ::= 'past'
                                     t-restr ::= 'unspec'
                                     eval 3 [ TIMEX 'now' ] ]
                        TEMP-REF 4 [ t-ref ::= 'future'
                                     t-restr ::= 'unspec'
                                     eval [ 2 ] ]
                        ... ]
                ... ] ]
```

Figure 5.17: Partial annotation: *John was going to visit his mother*

such as the *going to* future require a directly governing instance of temporal reference that provides their evaluation time. Contextual relative tenses work similarly, but their evaluation time may be specified via a temporal modifier or externally similar to the evaluation time. Since the overall mechanism remains the same, this is not illustrated explicitly.

Overall, this section has shown how to deal with different instances of temporal reference focusing on issues concerning the interpretation of corresponding annotations. In other words, this section has shown how the interpretation component of the ParTMA annotation scheme assembles meanings compositionally based on the syntactic and semantic input. In the following sections, this procedure is extended to grammatical and lexical aspect.

5.4 The semantics of aspect

In the introduction to this section, I have pointed out that the category of temporal reference and the category of aspect use fairly similar formal tools in the semantics I pursue in this thesis. For this reason, the next section, dealing with grammatical aspect, will be considerably shorter than the section on temporal reference. More focus is put on section 5.4.2, which focuses on implementing the first-phase syntax as a tool to encode inner aspect.

5.4.1 Grammatical aspect

Aspect employs the same basic framework as tense, except that no specific situation is provided for evaluation. This is because the corresponding situation is always provided by the `TEMP-REF` node. Furthermore, the input situation is attached to the `LEX-ASP` node. This makes the task of grammatical aspect to provide the link between propositions exemplifying events and temporal anchoring explicit. This is visualized in Figure 5.18. (58) shows the corresponding attribute-value matrix representation.

(58)
$$ {}_{0}\left[\begin{array}{ll} \text{PRED} \ldots & \\ \ldots & \\ \text{SEM} & {}_{1}\left[\begin{array}{ll} \text{TEMP-REF} & {}_{2}[\ldots] \\ \text{VIEWPOINT} & {}_{3}\left[\begin{array}{l} \textbf{a-ref} ::= \text{`perfective',`imperfective'} \\ \textbf{a-restr} ::= \text{`bound',`culminating',`ongoing',} \\ \text{`inertia',`habitual'} \end{array}\right] \end{array}\right] \end{array}\right] $$

Since the `VIEWPOINT` node provides the link between the `TEMP-REF` node and the `LEX-ASP` node, the concrete instantiation of certain indices required for the linear logic side of the meaning constructors is different from the instantiation that we have seen for tense. However, all the components have the same semantic types, as shown in (59).

(59) a. $\lambda P_{sst}.\lambda Q_{st}.\lambda s_s.\exists r_s[P(r)(s) \wedge Q(r)]$ *perfective*

b. $\lambda P_{sst}.\lambda Q_{st}.\lambda s_s.\forall r_s[P(r)(s) \wedge Q(r)]$ *perfective*

c. $\lambda s_s.\lambda s'_s.MB_x(s,s') = 1$

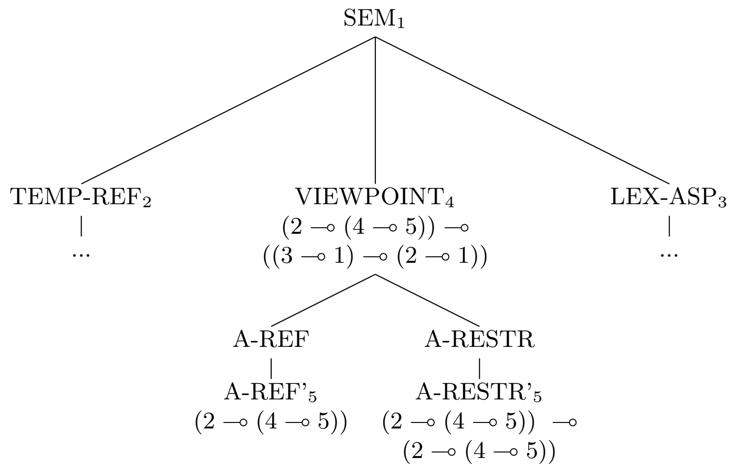

Figure 5.18: Composition of grammatical aspect

There is one other crucial difference between the semantics outlined for tense and the semantics for aspect discussed in this section. More specifically, there is a difference in terms of permitted annotations. In principle, it is possible to annotate perfective aspect with restrictors intended to combine with imperfective aspect and vice versa. This is primarily the responsibility of the annotator, but it is also possible to design a set of rules that run a consistency check over certain parts of the annotation and which, if necessary, rewrite values into an appropriate default value. Consider the example in (60), which combines a generally perfective viewpoint combined with an inertia modal base, which is associated with imperfective aspect to yield a progressive semantics.[19]

(60)
$$0\begin{bmatrix} \text{...} & \\ \text{VIEWPOINT} & 3\begin{bmatrix} \textbf{a-ref} ::= \text{`perfective'} \\ \textbf{a-restr} ::= \text{`inertia'} \end{bmatrix} \end{bmatrix}$$

[19]Recall that an inertia modal base provides situations that are continuations of the origin situation if that situation would continue without any interruptions. This would mean that, in the present configuration, we would say that there exists a situation that is a continuation of the reference/topic situation where the underlying eventuality is true. This goes against the intuitions of perfective aspect, which generally encodes culmination or termination in the actual situation rather than in an arbitrary possible situation.

A possible rewrite rule to fix this annotation is given in (61). This rule simply replaces the `a-restr` annotation with a value that suits the more general `A-REF` annotation. However, such rules must be carefully considered since it is not always clear which of the two values is wrongly annotated. A more safe strategy would be to use the rule in (62), the result of which makes it so that the viewpoint aspect annotation is ignored completely during the interpretation. This is a general technique that may be used to reduce the risk of nonsense annotations in favor of simpler semantic representations. This is possible since, as explained in the previous section, a *unspec* annotation results in an identity function, which means the semantic representation that is provided as input from lower in the tree, i.e., by the `LEX-ASP` node is simply percolated to the next stage of interpretation, namely, the interpretation of the `TEMP-REF` node as discussed in the previous section. To put it more concisely, rewrite rules can be used to check for incompatible annotations and rewrite them into identity functions.[20]

(61) `#a A-REF ‘perfective’ & #a A-RESTR ‘inertia’`
`==> #a A-RESTR ‘bounded’.`

(62) `#a A-REF ‘perfective’ & #a A-RESTR ‘inertia’`
`==> #a A-REF ‘unspec’ & #a A-RESTR ‘unspec’.`

The major difference between tense and aspect in the formal system presented here is that aspect always receives a reference situation from within a clause, for example, by tense (63a), but possibly also by particular embedding verbs, e.g., raising verbs (63b). However, since an undefined tense feature raises the aspectually modified clause to the appropriate level, we do not need to implement a separate anchoring mechanism for such clauses.

(63) a. Peter was writing a letter.

b. John seemed to be writing a letter.

[20] Of course, these rules should be a last resort to correct annotations. However, annotation errors are rarely completely avoidable, neither in rule-based systems nor systems based on machine learning.

This section has shown that the formal treatment of tense and grammatical aspect is fairly similar. By linking the glue semantics representations to elements in a tree structure, the order of application is preserved in a glue semantics proof. The last task is to explain the semantic interpretation of the annotation of lexical aspect to finalize the set of semantic interpretation rules. This will be discussed in the next section.

5.4.2 Inner aspect and first-phase syntax

In this section, I illustrate the implementation of Ramchand's (2008) first-phase syntax in Glue semantics and how it can be integrated with the broader semantic framework used in this thesis. This means concretely, a neo-Davidsonian event semantics embedded in a situation semantics framework for tense and aspect.

The first-phase syntax is a syntax for decomposing verb complexes, including its arguments and other rhematic material, into sub-eventualities that specify specific roles for event participants in the syntax. As shown in chapter 4 and Zymla, 2019, this kind of decomposition may be used to model different inference patterns associated with lexical and grammatical aspect, in particular, the imperfective paradox (Dowty, 1977). Furthermore, the first-phase syntax, in combination with LFG's f-structure, allows for sharing labor between syntax and subevent structure with respect to defining thematic roles, as I will show in this section.

The system is designed in such a way that the first-phase syntax decomposes the Davidsonian event argument into the initiation phase, the process phase, and the result phase. Verb classes and interaction with rhematic material determine which of these phases are realized. Each level of the first phase syntax has a subject position that, in the present implementation, is aligned with event arguments, i.e., thematic roles.

Ultimately, the goal of the annotation is to enable the generation of semantic subcategorization templates in the sense of Asudeh et al. (2008) and Haug (2008). The gist of the implementation here is that these templates are sensitive to the content of the first-phase syntax and saturate the subject

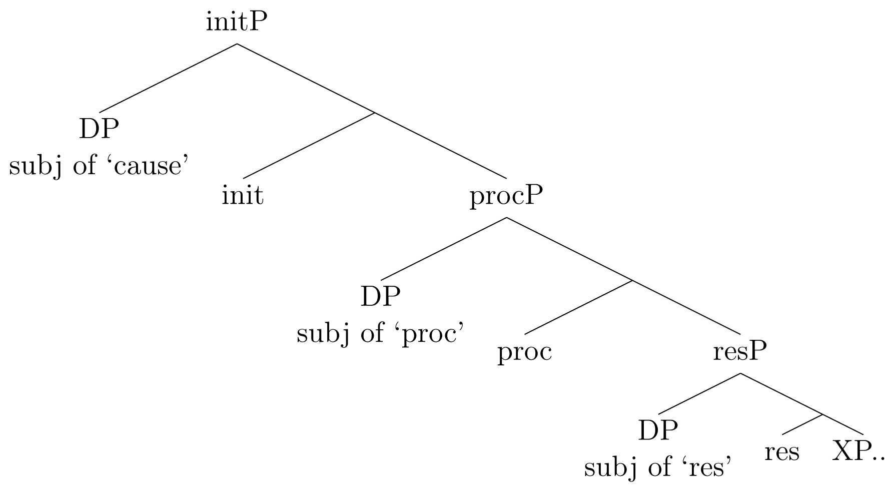

Figure 5.19: First-phase-syntax template (Ramchand, 2008)

positions of the first-phase syntax with the same variables that saturate grammatical functions in a syntactic representation (i.e., subject, object, etc.). The main idea outlined above relies equally on the Glue semantics machinery, as well as on the appropriate use of semantic construction rules.[21] In the next few paragraphs, I explain the details of the formal system and the sharing of labor between those two components.

To begin with, I discuss the event decomposition in Ramchand (2008) from an aspectual perspective, the main point for this thesis, then I show how this structure interacts with the syntactic f-structure and what this tells us about thematic roles in the traditional sense.

Starting in section 5.2.5, I have explained that, although designed for flat structures, Glue semantics inherently models hierarchical structures in terms of its proofs and, thus, also lends itself to model compositional processes guided by tree structures. I apply the same technique to Ramchand's (2008) first-phase syntax. This, conversely, means that the structure given in Figure 5.19 can be reduced to a flat structure from the perspective of annotation

[21] This is due to the fact that the first-phase syntax is not developed with resource sensitivity in mind and makes regular use of a copying mechanism, for example, when different projections share the same subject (Ramchand, 2008).

```
[ LEX-ASP [ class ::= 'event' | 'process' | 'state' | 'unspec'
            extent ::= 'mom' | 'dur'
            FPS [ initP  [ init ::= 'fid'
                           i-subj ::= 'fid' ]
                  procP  [ proc ::= 'fid'
                           p-subj ::= 'fid' ]
                  resP   [ res ::= 'fid'
                           r-subj ::= 'fid' ]
                  rheme  [ head ::= 'fid'
                           type ::= 'location' | 'path' |
                           'i-theme' | 'unspec' ] ] ] ]
```

Figure 5.20: Annotation of the first-phase syntax

(see Figure 5.20).

As the template in Figure 5.20 shows, the annotation values for each feature of the fps-structure are `fids`, i.e., references to the syntactic f-structure. Through this, the annotation references the `PRED` value. I assume that the `PRED` value is the key to the (lexical) encyclopedic content. Furthermore, this allows us to determine the semantics of the rhematic material based on the syntactic properties of different kinds of rhemes specified at the level of f-structure(e.g., NPs, PPs, etc.). In other words, we can use known Glue semantics analysis for the f-structure and embed the FPS in it.

In the present analysis (recall chapter 4), the goal is to provide a small set of primitives that allows for describing grammatical distinctions, similar to the approach postulated, e.g., by Dowty (1991) or Jackendoff (1987). However, by aligning the present approach with a thematic role based semantics, the present system can also be integrated with further resources, e.g. PropBank (Palmer et al., 2005), or VerbNet (Schuler (2005); see also Hautli et al. (2015) for comparable work).[22]

[22] Such alignments between different theories of interpreting thematic roles, furthermore,

This means aligning the first-phase syntax with neo-Davidsonian event semantics at a formal level is an exercise in interoperability of different semantic levels that is necessary for any semantic annotation scheme.

Consider the simple sentence *Mary ate the mango* from Ramchand (2008). According to her, the event described in this sentence can be decomposed into the following structure. The important feature here is that Mary simultaneously takes the role of the initiator of the event as well as the undergoer (by virtue of being the subject of the *proc* sub-event). However, in a neo-Davidsonian analysis, Mary would simply describe the *agent* role, while the mango would take the role of the theme, which is also often associated with the notion of *undergoing* or *receiving* the action of the verb.

Compare this to the example in (66). It receives the same Neo-Davidsonian analysis. However, as the event decomposition annotation indicates, the cart is the undergoer of the event. The path is unspecified which allows for an atelic interpretation, whereas the rheme in (64) provides a quantized path (annotated as `incr_theme` for incremental theme; cf. Krifka, 1998). This means concretely that the rhematic material is a crucial factor in determining the aspectual properties of a given expression and specifies the compatibility with aspect features higher in the composition, such as the available modal bases for imperfective aspect. As illustrated in the previous section, such *presuppositions* can be encoded in terms of axioms that ensure the soundness of the annotation. An example of this is the axiom that the progressive aspect is only available if the embedded element has a *proc* head and either a *res* head or some rhematic material that specifies a quantized path, i.e., the progressive modal base is only available if the element it applies to is telic.

(64) Mary ate a mango.

(65) $\exists e[eat(e) \land ag(e) = Mary \land the(e) = \iota x[mango(x)]]$

lends itself to the study of exactly the question, how thematic roles are to be understood at the level of syntax as well as the level of semantics (Dowty, 1991).

$$
{}_{0}\begin{bmatrix}
\text{PRED} & \text{`eat<Mary,mango>'} \\
\text{SUBJ} & {}_{1}\begin{bmatrix}\text{PRED `Mary'}\end{bmatrix} \\
\text{OBJ} & {}_{2}\begin{bmatrix}\text{PRED} & \text{`mango'} \\ \text{SPEC} & \begin{bmatrix}\text{PRED `the'}\end{bmatrix}\end{bmatrix} \\
\text{SEM} & {}_{3}\begin{bmatrix}\text{FPS} & {}_{4}\begin{bmatrix}
\text{initP} & {}_{5}\begin{bmatrix}\textbf{init} ::= {}_{0}[\;] \\ \textbf{i-subj} ::= {}_{1}[\;]\end{bmatrix} \\
\text{procP} & {}_{6}\begin{bmatrix}\textbf{proc} ::= {}_{0}[\;] \\ \textbf{p-subj} ::= {}_{1}[\;]\end{bmatrix} \\
\text{rheme} & {}_{7}\begin{bmatrix}\textbf{head} ::= {}_{2}[\;] \\ \textbf{type} ::= \text{'incr\_theme'}\end{bmatrix}
\end{bmatrix}\end{bmatrix}
\end{bmatrix}
$$

Figure 5.21: FPS annotation: *Mary ate a mango*

(66) a. John pushed a cart to every store.

b. John pushed a cart.

(67) a. $\forall y[store(y) \rightarrow \exists x[cart(x) \wedge \exists e[push(e) \wedge ag(e) = John \wedge th(e) = x \wedge goal(e) = y]]]$

$\exists x[cart(x) \wedge \forall y[store(y) \wedge \exists e[push(e) \wedge ag(e) = John \wedge th(e) = x \wedge goal(e) = y]]]$

b. $\exists x[cart(x) \wedge \exists e[push(e) \wedge ag(e) = John \wedge th(e) = x]]$

As illustrated above, `subjects` can be shared between different heads. This leads to variability with respect to the available resources from a Glue semantics perspective.

In this thesis, I provide a system that uses optional co-binding. More concretely, I use semantic construction rules to check for structure sharing in the annotation (as in (64)) and introduce an appropriate meaning constructor accordingly. This means the different heads have multiple possible meaning constructors depending on whether subjects are shared or distinct. The

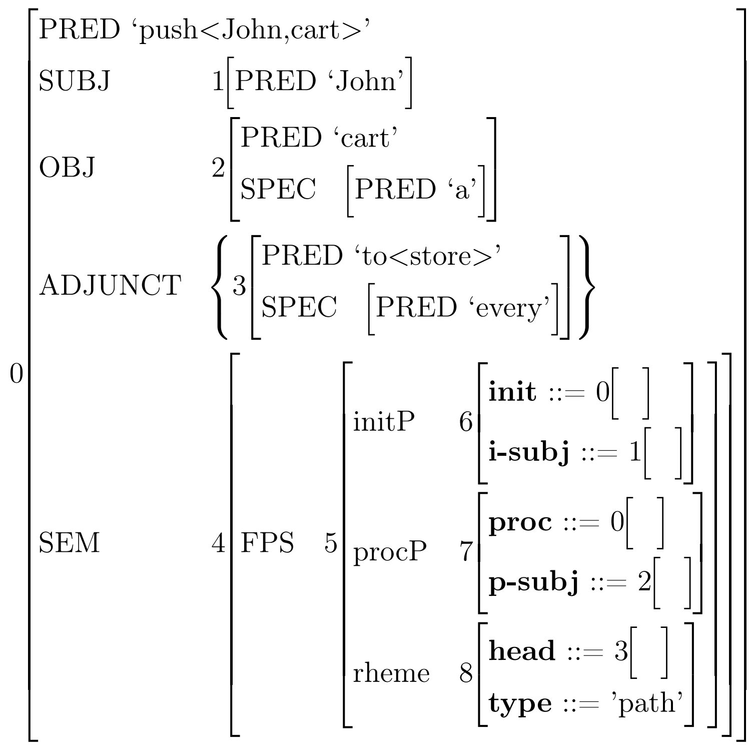

Figure 5.22: FPS annotation: *John pushed a cart to every store.*

semantics of the *init* head are given in (68).[23]

(68) a. $\lambda P_{vt}.\lambda x_e.\lambda e_v.\exists e', e''[eat'(e') \wedge e = e' \rightarrow e'' \wedge init(e', x) \wedge P(e'')]$: $(0 \multimap 7) \multimap (1 \multimap (0 \multimap 6))$

b. $\lambda P_{e,vt}.\lambda x_e.\lambda e_v.\exists e', e''[eat'(e') \wedge e = e' \rightarrow e'' \wedge init(e', x) \wedge P(x)(e'')]$ $(1 \multimap (0 \multimap 6)) \multimap (1 \multimap (0 \multimap 5))$

As the formulas in (68) show, the difference lies in the type of the next lower head (the *proc* head). When the same subject appears in *init* and *proc*, indicated by reference to the same resource at f-structure, then the formula in (68b) is used, else the formula in (68a) is used. To complete the semantics of the previous two examples, we need the following semantics of the *proc* head.[24]

(69) $\lambda y.\lambda x.\lambda e.Path(y, e) \wedge eat'(e) \wedge proc(e, x)$

a. Instantiation for (64):
$(2 \multimap (1 \multimap (0 \multimap 6))$

b. Instantiation for (66):
$(3 \multimap (2 \multimap (0 \multimap 7))$

For the sake of simplicity let us assume the semantics in (70) and (71) for the NPS in the two running examples. What is important is, that the semantics of *proc* map the entity, that is given to it as its first argument, to the event via the `path/2` predicate in accordance with Krifka, 1987(see chapter 4; Ramchand, 2008).

[23] As specified in chapter 4, I make use of the following abbreviations of Ramchand's formula (Ramchand, 2008).

(1) a. Subject(x,e) and Initiation(e) ↔ init(e,x)

b. Subject(x,e) and Process(e) ↔ proc(e,x)

c. Subject(x,e) and Result(e) ↔ res(e,x)

[24] Recall that, in accordance with Ramchand (2008), I have postulated two different *proc* semantics, depending on the existence of a *res* head.

(70) a. *Mary*:

MC: Mary : $s1$

$$\begin{bmatrix} \text{SUBJ} & 1\begin{bmatrix}...\end{bmatrix} \end{bmatrix} \qquad \text{s1}\begin{bmatrix} \ \end{bmatrix}$$

b. *a mango*:

MC: $\lambda Q_{et}.\exists x[mango(x) \wedge Q(x)] : (s2 \multimap 0) \multimap 0$

$$\begin{bmatrix} \text{OBJ} & 2\begin{bmatrix}...\end{bmatrix} \end{bmatrix} \qquad \text{s2}\begin{bmatrix} \text{VAR} & \text{s3}\begin{bmatrix} \ \end{bmatrix} \\ \text{RESTR} & \text{s4}\begin{bmatrix} \ \end{bmatrix} \end{bmatrix}$$

(71) a. *John*:

MC: John : $s1$

$$\begin{bmatrix} \text{SUBJ} & 1\begin{bmatrix}...\end{bmatrix} \end{bmatrix} \qquad \text{s1}\begin{bmatrix} \ \end{bmatrix}$$

b. *a cart*:

MC:$\lambda Q_{et}.\exists x[cart(x) \wedge Q(x)] : (s2 \multimap 0) \multimap 0$

$$\begin{bmatrix} \text{OBJ} & 2\begin{bmatrix}...\end{bmatrix} \end{bmatrix} \qquad \text{s2}\begin{bmatrix} \text{VAR} & \text{s3}\begin{bmatrix} \ \end{bmatrix} \\ \text{RESTR} & \text{s4}\begin{bmatrix} \ \end{bmatrix} \end{bmatrix}$$

c. *every store*:

MC: $\lambda Q_{et}.\exists x[store(x) \rightarrow Q(x)] : (s3 \multimap 0) \multimap 0$

$$\begin{bmatrix} \text{OBJ-TH} & 3\begin{bmatrix}...\end{bmatrix} \end{bmatrix} \qquad \text{s5}\begin{bmatrix} \text{VAR} & \text{s6}\begin{bmatrix} \ \end{bmatrix} \\ \text{RESTR} & \text{s7}\begin{bmatrix} \ \end{bmatrix} \end{bmatrix}$$

As a mediator between the first-phase syntax and the neo-Davidsonian event semantics, I make use of templates mimicking the subcategorization of a given verb based on Asudeh et al. (2008). However, rather than being only dependent on the subcategorization at the f-structure, the template is also sensitive to the first-phase syntax. More specifically, the system works such that a given template takes as a first argument a first-phase syntax representation and provides it with a subcategorization template compatible with a neo-Davidsonian event semantics. In other words, the first argument of a template varies in terms of its semantic type, depending on how the arguments are encoded in the first-phase syntax. The result of applying a template to an fps representation is a subcategorization template in the sense of Asudeh et al. (2008) and Haug (2008). In the process of applying the template, it shifts the indices relevant for the semantic computation, similar to the function of a manager resource (Asudeh, 2004).[25] In (72), the two templates for the two examples are given.

As these examples show, the templates project the syntactic indices of the f-structure to the semantic indices introduced in (70) and (71). This technique allows the system to keep the semantics modular while still accounting for the possibility of co-binding via the use of *assumptions*, i.e., the use of linear abstraction and implication introduction. This is shown explicitly in the proofs, (5.23) and (5.24).

(72) a. $\lambda R_{eevt}.\lambda x_e.\lambda y_e.\exists e_v[R(x)(y)(e) \wedge ag(e) = x \wedge th(e) = y]$:
$1 \multimap (2 \multimap (0 \multimap 5)) \multimap (s1 \multimap (s2 \multimap 0))$

b. $\lambda R_{eeevt}.\lambda x_e.\lambda y_e.\lambda z_e.\exists e_v[R(x)(y)(z)(e) \wedge ag(e) = x \wedge th(e) = y \wedge goal(e) = z]$:
$1 \multimap (2 \multimap (3 \multimap (0 \multimap 6))) \multimap (s1 \multimap (s2 \multimap (s3 \multimap 0)))$

The resulting semantic representations are shown in (73a) and (73b).

(73) a. Mary ate a mango.
$\exists x[mango(x) \wedge \exists e[\exists e', e''[eat'(e') \wedge e = e' \rightarrow e'' \wedge init(e', Mary) \wedge Path(x, e'') \wedge eat'(e'') \wedge proc(e'', Mary)] \wedge ag(e) = Mary \wedge th(e) = x]]$

b. John pushed a cart to every store.

i. $\forall y[store(y) \rightarrow \exists x[cart(x) \wedge \exists e[\exists e', e''[push'(e') \wedge e = e' \rightarrow e'' \wedge init(e', John) \wedge Path(y, e'') \wedge push'(e'') \wedge proc(e'', x)] \wedge ag(e) = John \wedge th(e) = x \wedge goal(e) = y]]]$

ii. $\exists x[cart(x) \rightarrow \forall y[store(y) \wedge \exists e[\exists e', e''[push'(e') \wedge e = e' \rightarrow e'' \wedge init(e', John) \wedge Path(y, e'') \wedge push'(e'') \wedge proc(e'', x)] \wedge ag(e) = John \wedge th(e) = x \wedge goal(e) = y]]]$

[25] A manager resource can be understood as a (partial) identity function, which does not affect the meaning side of a meaning constructor but changes the linear logic side of a meaning constructor to behave differently within a glue derivation. For example, to change the order in which certain elements apply.

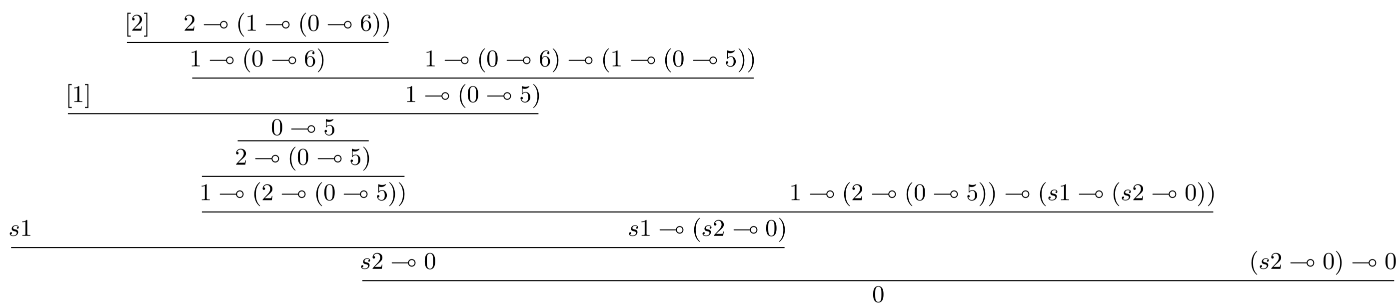

Figure 5.23: Derivation (untensed) for *Mary ate a mango*

As the previous two examples show, the semantics of the first-phase syntax can be assembled via Glue semantics meaning constructors for each head, as well as a template that takes the role of aligning the *subjects* of the first-phase syntax with the event semantics used at the level of f-structure. This means the semantics presented here provide an interface between first-phase syntax, thematic roles, and f-structure more generally and show how the present approach can be reconciled with existing insights from the Glue semantics literature.

As discussed in chapter 4, the *proc* head has two different denotations based on the existence of a *res* head. Furthermore, I have shown that to achieve the desired variable bindings, the *init* and *proc* heads need to have two different denotations based on shared subjects between the different heads. At the beginning of this section, I pointed out that this choice is reflected in the annotation. Thus, it is possible to use semantic construction rules to check for the appropriate meaning constructors for a given annotation. I will explain this in more detail next.

To begin with, I examine the use of the two *proc* heads postulated by Ramchand (2008) and shown in (74).

(74) a. $\lambda y.\lambda x.\lambda e.Path(y, e) \wedge \%p'(e) \wedge proc(e, x)$

b. i. $\lambda P_{evt}.\lambda x_e.\lambda e_v.\exists e', e''[P(x)(e'') \wedge \%p'(e') \wedge proc(e', x) \wedge e = (e' \rightarrow e'')]$

ii. $\lambda P_{vt}.\lambda x_e.\lambda e_v.\exists e', e''[P(e'') \wedge \%p'(e') \wedge proc(e', x) \wedge e = (e' \rightarrow e'')]$

The condition for applying the first vs. the second semantics is simple. (74b) is used, when the *res* head is unspecified, whereas (74a) is used, when no *res* head is specified. Thus, the default semantics are given in (74a). They are replaced by a more specific rule checking for a *res* head. The argument of the *res* semantics in (74a) is always an entity of type *e*. If there is no overtly specified rhematic material, then a contextual value is provided. This means that a single rule is sufficient to cover first-phase syntax representations without a *res* head. On the other hand, when there is a *res* head, the semantics in (74b) are used. In this case, the exact semantics vary

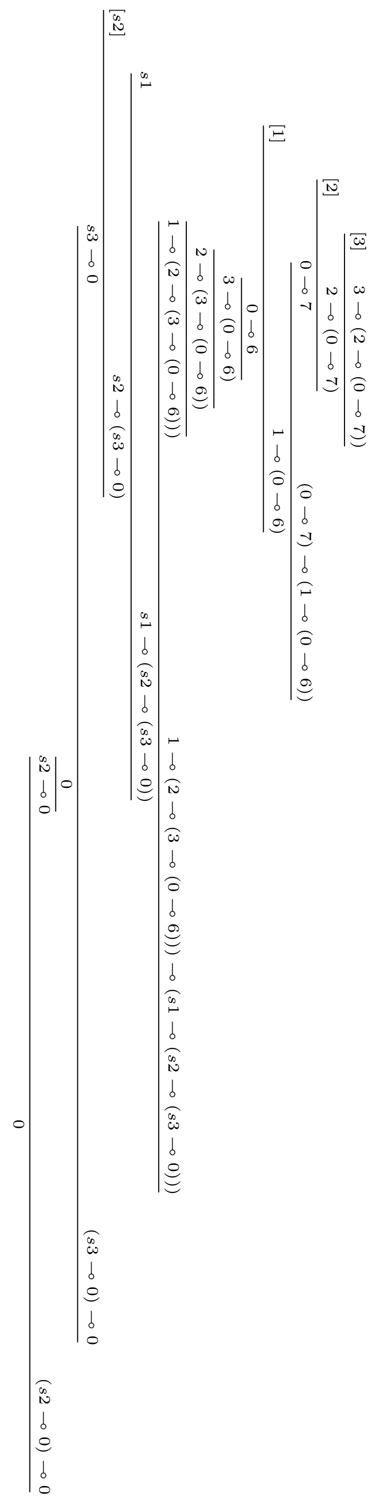

Figure 5.24: Derivation (untensed) for *John pushed a cart to every store*

depending on whether the *proc* head and the *res* share their subject, (74b-i), or whether they have different subjects, (74b-ii).

The *init* head also has this division. By providing these two different semantic denotations for *init* and *proc* respectively, all available co-binding configurations in the first-phase syntax are covered. Note also that each of the heads only needs to co-bind the head directly beneath it. This means the head encodes whether it wants to co-bind the next lower projection in the first-phase syntax, independent of the denotation of that lower head. The same process can be extended to the res-head to fully cover the semantics of the first-phase syntax. This architecture reduces context sensitivity for each head in the first-phase syntax to a minimum.

5.5 Summary

In this chapter, I have presented how the annotation scheme and the corresponding rules developed throughout this thesis can be implemented. For this purpose, I have presented a novel system for querying linguistic structures and adding annotations relative to those structures. The added annotation can be represented as a graph or f-structure. The system also allows for creating specific semantic representations, namely, glue semantics representations. This is grounded in theoretical and practical research on Glue semantics. By combining the possibilities provided by the rule system worked out in this chapter and Glue semantics, I have provided a concrete syntax and semantics for the annotation scheme presented in this thesis spanning a cross-linguistically valid feature set for tense, viewpoint aspect, and lexical aspect. Particular attention was paid to avoiding structural ambiguities in the annotation template in virtue of ambiguities enabled by the concepts worked out for the semantics of tense and aspect (and generally) in the previous chapters (3 and 4. The present system is compatible with manual and rule-based annotation efforts but also allows parTMA features to be learned via machine learning or deep learning approaches. The validity of such approaches can then be verified by applying semantic construction rules on the resulting feature set and inspecting the resulting formula. A concrete implementation is left for future work.

The next step is to generate more expansive resources based on the current system. In particular, the goal is to provide additional treebanks for the ParGram grammars and to provide a venue for inspecting both the syntactic and semantic annotation contained in those treebanks, thus making the work more accessible to a broader audience.

Chapter 6

Conclusions and future work

This thesis aimed to present a computationally viable annotation scheme of tense and aspect. The main focus was on cross-linguistic applicability and a grounding in formal semantics. Thus, the ParTMA annotation scheme, developed as part of this thesis, is informed by various sources in the linguistic research landscape. It combines insights from cross-linguistic computational and typological research and the growing literature on cross-linguistically informed formal semantics for tense and aspect.

The project that gave birth to this thesis was a follow-up project for the international ParGram project with the goal of providing cross-linguistically comparable syntactic analyses for typologically diverse languages. Although tense and aspect have been heavily discussed within the project, only a few key insights have been implemented in the ParGram grammars. So, for example, the treatment of auxiliaries across languages (Butt et al., 1996), or the separation of the categories tense and aspect into independent features (Butt et al., 1999). All in all, the existing research on tense and aspect from the perspective of the ParGram project has focussed on structural, i.e., morphosyntactic, issues rather than semantic issues. The goal of this thesis was to pick up the efforts of the ParGram project at exactly this point.

6.1 Summary

In the introductory chapters, 1 and 2, I examined the state of affairs with respect to the treatment of tense and aspect within ParGram and in semantic annotation. As the main idea of this thesis was to explore the semantic properties of tense and aspect, chapter 2 has focused on providing a broad overview of tense/aspect annotation schemes used in computational linguistics and compared them to work done in the ParGram effort. The main shortcoming of the ParGram grammars was that they lacked an explicit formal semantics, a crucial component for semantic annotation schemes (Bunt, 2015). Thus, the background chapter solidified the main goal of this thesis to provide exactly that: an annotation scheme with a concrete formal semantics supporting it.

The resulting annotation scheme, the ParTMA annotation scheme, consists of a syntax/semantics interface encoded in terms of multilayered rewrite rules and a semantic interpretation component for annotated structures that produces formal semantic representations based on the input annotation. Thus, the ParTMA annotation scheme can be understood as a separate projection for LFG in the sense of Kaplan (1995b): it serves as a useful intermediate representation that does not necessarily capture a psycho-linguistic reality but can be seen as a tool for linguists to describe specific properties of language. In this case, the semantic properties are provided by tense and aspect. The foundations for this annotation scheme have been extensively discussed in chapters 3 and 4.

The third chapter of this thesis provides an overview of both the structure as well as the possible meanings encoded via the category of tense. The category of tense has been chosen to exemplify the main features of the ParTMA annotation scheme to the reader for the first time in 3. Tense is arguably the more intuitive of the two major categories discussed in this thesis and is mainly associated with locating situations described by linguistic expressions marked with tense relative to some evaluation point, i.e., primarily, temporal semantics. I proposed a system that isolates the temporal ordering properties of tense and connects them to their respective morphosyntactic reflexes, be it temporal ordering induced by aspectual marking or temporal ordering by tense proper. Both have shown to be available cross-linguistically. Furthermore, I have discussed the annotation of languages without overt tense marking. I dubbed these cases of inferential tense. An extreme case of this has been presented in terms of Indonesian, which has neither obligatory tense nor aspect marking, making temporal ordering a purely contextual endeavor.

Based on the insights I gained from examining the tense and aspect systems of the languages covered by the ParGram grammars and beyond, I have proposed to separate structural and meaning properties of the category of tense and modularize its semantic contribution. This yielded the semantic category of *temporal reference*, which is an ordering function for temporal intervals (inspired by the term semantic tense). I have illustrated this function by embedding it in a compositional interval semantics based on work by Kusumoto (1999, 2005), which has recently gained traction in cross-linguistic research by, for example, Bochnak et al. (2019). The resulting semantics model tense as a contextually restricted quantifier over times and temporal reference as the restrictor of this quantifier. I have related the semantics to the morphosyntax via a set of rules called semantic construction rules. Semantic construction rules generate semantic features from the morphosyntactic and semantic context. These rules can be applied iteratively, resulting in a tiered system that calculates structural and semantic context dependencies as they become available. The result is a set of semantic features from which a semantic representation can be derived.

I have illustrated this style of annotation by applying it to two different phenomena that have been discussed in the literature. Tense in subordinate clauses and tense in counterfactual conditionals. Furthermore, I have discussed the treatment of cases of inferential temporal reference. I have proposed a set of axiomatic rules that reduce underspecified values for temporal reference to default values, where this is possible based on the morphosyntactic and lexical context, e.g., via concrete specification (yesterday, tomorrow, etc.), or via default inferences induced by, for example, grammatical aspect. Thus, chapter 3 illustrated the general methodology pursued in this thesis of determining relevant morphosyntactic and semantic features, providing a formal semantics for these features, and, finally, a procedure for mapping form to meaning.

This procedure has been applied to the categories of grammatical and lexical aspect in chapter 4. The semantics for grammatical aspect have been inspired mainly by Arregui et al. (2011, 2014), denoting grammatical aspect as quantifiers over situations. This has been motivated in the literature, mainly by exploring the various meanings of the imperfective aspect, which has proven to be a more significant challenge than the perfective aspect. Ultimately, the result was a semantics that mainly distinguished perfective and imperfective aspects in terms of quantificational force. Perfective aspect is encoded in terms of an existential quantifier over situations, and imperfective aspect is encoded in terms of a universal quantifier over situations. Cross-linguistic variation of these categories has been encoded as accessibility relations encoded in the restrictors of those quantifiers. Thus, the treatment of tense and grammatical aspect yields some parallels.

The treatment of lexical aspect has been based on a decompositional event semantics, more concretely, on a framework for capturing the morphosyntactic reflexes of event structure of which lexical aspect is an integral part. The first-phase syntax that has served as the underlying framework for the modeling of lexical aspect was devised in Ramchand (2008). The goal of this approach was to reduce the properties of event structure to a small set of primitives that can cover a large number of configurations by virtue of different combinations of these primitives.

Chapter 4 points out some interesting synergies between the treatment of grammatical aspect and the treatment of lexical aspect. Concretely, I have proposed redefining the accessibility relations used in the quantifiers in terms of the sub-event structure of the predicates they select. This, for example, allowed me to provide a coherent distinction between the event-in-progress and preparatory interpretation of imperfective aspect.

In the last part of Chapter 4, the insights made about grammatical and lexical aspect have been incorporated into the ParTMA annotation scheme. Grammatical aspect has been embedded in the ParTMA annotation scheme in a fashion similar to tense. The first-phase syntax has been integrated in terms of an attribute-value structure, thus, reducing the hierarchical first-phase syntax with the goal of encoding hierarchical effects in the semantic composition. This is elaborated upon in Chapter 5 where a concrete implementation of the ParTMA annotation scheme is discussed.

Chapter 5 describes a concrete syntax and semantics for the ParTMA annotation scheme. More specifically, I presented a linguistic rewrite system that exemplifies a concrete implementation of the tiered semantic construction rules described in chapters 3 and 4. By combining this rewrite system with the capabilities of the Glue semantics workbench, a tool for evaluating proofs in the semantic formalism of Glue semantics, I have furthermore presented a tool that allows for the semantic interpretation of the ParTMA annotation scheme resulting in full-fledged semantic representations. I have shown that the semantics of the ParTMA annotation scheme can be carefully engineered to capture the possible combinations of its feature space and thus capture the linguistic insights coded into the semantics in the previous chapters. To illustrate this point, I have revisited various examples discussed in the previous chapters to highlight both the role of the syntax/semantics interface and the role of the semantic interpretation component.

6.2 Future work

In this section, I discuss some potential and necessary venues for future work with respect to the ParTMA annotation scheme. This includes the generation of more resources, refinement of the underlying tense/aspect semantics, expanding the annotation scheme to discourse-level phenomena, and embedding the tense/aspect semantics in a general theory of formal semantics, particularly Glue semantics.

6.2.1 The syntax/semantics interface

This thesis presents a fully fleshed-out formal semantics for tense and aspect couched in a simple annotation scheme based on attribute/value structures. I made a proposal to the effect that the annotation scheme is populated via a rule-based syntax/semantics interface. However, a ParTMA template may be filled out by any other means of annotation, such as manual annotation or automated annotation. The decision to provide a rule-based system stemmed from the fact that tense/aspect categories heavily interact with one another even within the confines of a single clause. The rule-based system was designed to explain these interactions. Although the semantics are fully fleshed out, I have only exemplified the syntax/semantics interface by virtue of carefully crafted examples, such as the Sequence-of-tense phenomenon. Thus one of my future goals is to further investigate and scale up the applicability of rewrite rules to encode the syntax/semantics interface for tense and aspect and in general.

6.2.2 Refining tense/aspect semantics

There are several points of interest where the present annotation scheme could be refined even further. For example, future tense has only been discussed briefly and has been reduced to the semantic property of future orientation. Similarly, the perfect has been reduced to its temporal properties. The latter, in particular, has received this treatment based on the ongoing debate in the literature with respect to its status as tense or aspect. On the other hand, doing the future tense justice would require expanding the present annotation scheme to modals, opening a whole other can of worms.

The annotation for lexical aspect has been based on a decompositional event semantics. Although I have shown that these semantics interact with the syntax in an interesting way and, thus, provide a basis for annotating lexical aspect in relation to its morphosyntactic reflexes, I have only discussed the annotation of rhematic material by some chosen examples. More concretely, there is future work to do in working out the semantics of rhematic material, particularly the rhematic material that describes the result state of some eventuality. Another issue that I have mentioned is the question of how to deal with pluralities of either eventualities or the corresponding actors. In this thesis, I have only modeled pluralities external to the eventuality argument; however, plurality and pluractionality may affect the VP at different levels, creating a complex picture that needs to be examined in more detail.

The ParTMA annotation scheme mainly focuses on sentence-level phenomena. This is an obvious shortcoming since many discourse effects are induced by tense and aspect and, thus, call for a formal pragmatics that captures these insights. This is another obvious avenue for future work that includes extensive research in Glue semantics and the discourse effects of tense and aspect in general.

Finally, some smaller refinements can be made, for example, concerning the definition of the accessibility relations for the aspectual operators. These and other smaller improvements of the annotation scheme, I believe, will happen naturally with time, as more resources are created based on the system.

6.2.3 Resources

As of now, the ParTMA annotation scheme has produced a small number of treebanks on the INESS webpage. In addition, to these resources, a resource for encoding the syntax/semantics interface on linguistic rewriting has been developed. Based on this resource, a second tool has been developed for deriving semantic representations from ParTMA style annotations. My next goal is to provide more corpora annotated with the annotation tools that I have made available. Producing more corpora will also serve as an exercise in scalability of the present annotation scheme and the underlying semantic framework.

An important factor that I want to pay particular attention to is working on a packed system of semantic representations that spans from XLE's output to the final output in terms of glue semantics proofs. In other words, I want to pay particular focus on working on ambiguity management for the ParTMA annotation scheme and the related resources. This also entails the continuous work on the related Glue semantics workbench that I have used to produce semantic representations.

Overall, the results of this thesis provide lots of avenues for further research that synergizes with different areas of linguistics, including corpus linguistics, formal linguistics, and computational linguistics. My main goal is to work on computational tools. My conviction is that improving these tools with a concise goal, such as the annotation of tense and aspect, will not only improve the computational viability but also its benefits for the related formal linguistic research fields.

6.3 Concluding remarks

The goal of this thesis is to breach the gap between formal semantic insights and computational models of dealing with tense and aspect. As Steedman (1997) said, research in tense and aspect runs the danger of repeating the same points again and again. The annotation scheme in this thesis does not reject this claim. Quite the contrary, many of the insights compiled in this thesis rely on systems that are prevalent both in formal and computational frameworks of tense and aspect. Thus, for example, temporal reference is a rendition of the widely used Neo-Reichenbachian tense system, following the tradition of Reichenbach (1947) and some refining proposals by Klein (1994). The main contribution of this thesis is to embed these insights in a compositional semantic framework that can be derived from a simple annotation structure. Thus, the present thesis combines the appeal of devising different kinds of annotation procedures for different tasks and the support of a fleshed-out formal semantic framework that allows users to evaluate the resulting annotation. Thus, the thesis follows general principles for semantic annotation but provides a more intimate relationship between the underlying formal semantic framework and the corresponding semantic annotation scheme. For this reason, I believe that this thesis has achieved the goal it was set out to do: bringing computational and formal linguistics closer together, particularly in the domain of semantics.

Bibliography

Abeillé, Anne and Danièle Godard (2002). "The Syntactic Structure of French Auxiliaries". In: *Language*, pp. 404–452.

Abusch, Dorit (1988). *Sequence of Tense, Intensionality and Scope*. Stanford University.

— (1997). "Sequence of Tense and Temporal De Re". In: *Linguistics and Philosophy* 20.1, pp. 1–50.

— (1998). "Generalizing Tense Semantics for Future Contexts". In: *Events and Grammar*. Springer, pp. 13–33.

Agrell, Sigurd (1908). *Aspektänderung und Aktionsartbildung beim polnischen Zeitworte: ein Beitrag zum Studium der indogermanischen Präverbia und ihrer Bedeutungsfunktionen*. Lund, Sweden.

Aikhenvald, Aleksandra (2004). *Evidentiality*. Oxford University Press.

Allen, James F (1990). "Maintaining Knowledge about Temporal Intervals". In: *Readings in Qualitative Reasoning about Physical Systems*. Elsevier, pp. 361–372.

Aloni, Maria and Paul Dekker, eds. (2016). *The Cambridge Handbook of Formal Semantics*. Cambridge Handbooks in Language and Linguistics. Cambridge University Press.

Altshuler, Daniel and Roger Schwarzschild (2012). "Moment of Change, Cessation Implicatures and Simultaneous Readings". In: *Proceedings of Sinn und Bedeutung*. Vol. 17, pp. 45–62.

Andrews, Avery D (2008). "The Role of PRED in LFG + Glue". In: *Proceedings of the LFG08 Conference*, pp. 47–67.

— (2009). *Yet Another Attempt to Explain Glue*.

Andrews, Avery D (2010). "Propositional Glue and the Correspondence Architecture of LFG". In: *Linguistics and Philosophy* 33.3, pp. 141–170.

Arka, I Wayan (2012). "Developing a Deep Grammar of Indonesian within the Pargram Framework: Theoretical and Implementational Challenges". In: *Proceedings of the 26th Pacific Asia Conference on Language, Information, and Computation*, pp. 19–38.

— (2013). "On the Typology and Syntax of TAM in Indonesian". In: *Tense, Aspect, Mood and Evidentiality in Languages of Indonesia*. Ed. by John Bowden, pp. 23–40.

Arregi, Karlos and Peter Klecha (2015). "The Morphosemantics of Past Tense". In: *North East Linguistic Society (NELS)*. Vol. 45, pp. 53–66.

Arregui, Ana (2005). *On the Accessibility of Possible Worlds: The Role of Tense and Aspect*. University of Massachusetts.

Arregui, Ana, María Luisa Rivero, and Andrés Salanova (2011). "Imperfectivity: Capturing Variation across Languages". In: *MIT Working Papers in Linguistics* 75, pp. 53–68.

— (2014). "Cross-Linguistic Variation in Imperfectivity". In: *Natural Language & Linguistic Theory* 32.2, pp. 307–362.

Asahara, Masayuki, Sachi Kato, Hikari Konishi, Mizuho Imada, and Kikuo Maekawa (2014). "BCCWJ-TimeBank: Temporal and Event Information Annotation on Japanese Text". In: *International Journal of Computational Linguistics & Chinese Language Processing, Volume 19, Number 3, September 2014*.

Asher, Nicholas and Hajime Wada (1988). "A Computational Account of Syntactic, Semantic and Discourse Principles for Anaphora Resolution". In: *Journal of Semantics* 6.1, pp. 309–344.

Asudeh, Ash (2004). "Resumption as Resource Management". PhD thesis. Stanford University.

— (2006). "Direct Compositionality and the Architecture of LFG". In: *Intelligent linguistic architectures: Variations on themes by Ronald M. Kaplan*, pp. 363–387.

Asudeh, Ash and Richard Crouch (2002). “Glue Semantics for HPSG”. In: *Proceedings of the 8th international HPSG conference, Stanford, CA. CSLI Publications.*

Asudeh, Ash, Mary Dalrymple, and Ida Toivonen (2008). “Constructions with Lexical Integrity: Templates as the Lexicon-Syntax Interface”. In: *Proceedings of the LFG08 Conference*, pp. 68–88.

Bach, Emmon (1986). “The Algebra of Events”. In: *Linguistics and Philosophy* 9.1, pp. 5–16.

Ballarin, Roberta (2021). “Modern Origins of Modal Logic”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Fall 2021. Metaphysics Research Lab, Stanford University.

Bary, Corien, Daniel Altshuler, Kristen Syrett, and Peter de Swart (2018). “Factors licensing embedded present tense in speech reports”. In: *Proceedings of Sinn und Bedeutung 22, vol. 1*, pp. 127–142.

Bary, Corien and Dag Haug (2011a). “Inter-and Intrasentential Anaphora: the Case of the Ancient Greek Participle”. In: *Semantics and Linguistic Theory*. Vol. 21, pp. 373–392.

Bary, Corien and Dag Trygve Truslew Haug (2011b). “Temporal Anaphora across and Inside Sentences: The Function of Participles”. In: *Semantics and Pragmatics* 4, pp. 8–1.

Bennett, Michael and Barbara H Partee (1978). *Toward the Logic of Tense and Aspect in English*. Blackwell Publishing Ltd.

Bethard, Steven, Leon Derczynski, Guergana Savova, James Pustejovsky, and Marc Verhagen (2015). “Semeval-2015 task 6: Clinical tempeval”. In: *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval 2015)*, pp. 806–814.

Bethard, Steven, Guergana Savova, Martha Palmer, and James Pustejovsky (Aug. 2017). “SemEval-2017 Task 12: Clinical TempEval”. In: *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*. Vancouver, Canada: Association for Computational Linguistics, pp. 565–572. DOI: 10.18653/v1/S17-2093. URL: https://www.aclweb.org/anthology/S17-2093.

Bittar, André (2010). “ISO-TimeML Annotation Guidelines for French Version 1.0”. In: *ALPAGE & Université Paris* 7.

Bittar, André, Pascal Amsili, Pascal Denis, and Laurence Danlos (2011). “French TimeBank: an ISO-TimeML Annotated Reference Corpus”. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies: short papers*. Association for Computational Linguistics, pp. 130–134.

Bittner, Maria (2014). *Temporality: Universals and variation*. John Wiley & Sons.

Błaszczak, Joanna and Dorota Klimek-Jankowska (2016). “Aspectual Coercion versus Blocking: Experimental Evidence from ERP Study on Polish Converbs”. In: *How Categorical are Categories*.

Bobrow, Daniel G., Bob Cheslow, Cleo Condoravdi, Lauri Karttunen, Tracy Holloway King, Rowan Nairn, Valeria de Paiva, Charlotte Price, and Annie Zaenen (2007). “PARC’s Bridge and Question Answering System”. In: *Proceedings of the GEAF 2007 Workshop*, pp. 1–22.

Bobrow, Daniel G, Cleo Condoravdi, Lauri Karttunen, and Annie Zaenen (2009). “Learning by Reading: Normalizing Complex Linguistic Structures onto a Knowledge Representation”. In: *AAAI Spring Symposium: Learning by Reading and Learning to Read*, pp. 3–6.

Bochnak, M Ryan (2016). “Past Time Reference in a Language with Optional Tense”. In: *Linguistics and Philosophy* 39.4, pp. 247–294.

Bochnak, M Ryan, Vera Hohaus, and Anne Mucha (2019). “Variation in Tense and Aspect, and the Temporal Interpretation of Complement Clauses”. In: *Journal of Semantics* 36.3, pp. 407–452.

Bochnak, M Ryan and Peter Klecha (2015). “Temporal remoteness and vagueness in past time reference in Luganda”. In: *To appear in Selected Proceedings of ACAL* 45, p. 1.

Bohnemeyer, Jürgen (1998). *Temporal reference from a radical pragmatics perspective: Why Yucatec does not need to express’ after’and’before’*.

Börjars, Kersti (2020). “Lexical-Functional Grammar: An Overview”. In: *Annual Review of Linguistics* 6, pp. 155–172.

Bresnan, Joan, Ash Asudeh, Ida Toivonen, and Stephen Wechsler (2015). *Lexical-functional syntax*. Vol. 16. John Wiley & Sons.

Bresnan, Joan and Ronald Kaplan (1987). *Lexical-Functional Grammar*. Stanford University, Linguistic Institute.

Bunt, Harry (2008). "Semantic Underspecification: Which Technique for what Purpose?" In: *Computing Meaning*. Springer, pp. 55–85.

— (2010). "A Methodology for Designing Semantic Annotation Languages Exploring Semantic-syntactic ISO-morphisms". In: *Proceedings of the Second International Conference on Global Interoperability for Language Resources (ICGL 2010), Hong Kong*, pp. 29–46.

— (2015). "On the principles of interoperable semantic annotation". In: *Proceedings of the 11th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, pp. 1–13.

— (2019). "A Semantic Annotation Scheme for Quantification". In: *Proceedings of the 13th International Conference on Computational Semantics-Long Papers*, pp. 31–42.

Bunt, Harry and Chwhynny Overbeeke (2008a). "An Extensible Compositional Semantics for Temporal Annotation". In: *The Workshop Programme*.

— (2008b). "Towards Formal Interpretation of Semantic Annotation". In: *LREC*. Citeseer.

Butt, Miriam, Mary Dalrymple, and Anette Frank (1997). "An Architecture for Linking Theory in LFG". In: *Proceedings of the LFG97 Conference*. University of California San Diego, pp. 1–16.

Butt, Miriam, Helge Dyvik, Tracy Holloway King, Hiroshi Masuichi, and Christian Rohrer (2002). "The Parallel Grammar Project". In: *Proceedings of the 2002 Workshop on Grammar Engineering and Evaluation*. Vol. 15. Association for Computational Linguistics, pp. 1–7.

Butt, Miriam, Tracy Holloway King, María-Eugenia Niño, and Frédérique Segond (1999). *A Grammar Writer's Cookbook*. CSLI Publications.

Butt, Miriam, María-Eugenia Niño, and Frédérique Segond (1996). "Multilingual Processing of Auxiliaries within LFG". In: *KONVENS*. Vol. 96, pp. 111–122.

Butt, Miriam and Jafar Rizvi (2010). "Tense and Aspect in Urdu". In: *Layers of Aspect*, pp. 43–66.

Bybee, Joan (1994). "The Grammaticization of Zero". In: *Perspectives on grammaticalization* 109, p. 235.

— (1995). "Regular Morphology and the Lexicon". In: *Language and Cognitive Processes* 10.5, pp. 425–455.

Bybee, Joan, Revere Perkins, and William Pagliuca (1994). "The Evolution of Grammar: Tense, Aspect, and Modality in the Languages of the World". In: *Aspect, and Modality in the Languages of the World, Chicago*.

Bybee, Joan L (1985). *Morphology: A Study of the Relation Between Meaning and Form*. Vol. 9. John Benjamins Publishing.

— (2006). "From Usage to Grammar: The Mind's Response to Repetition". In: *Language* 82.4, pp. 711–733.

Bybee, Joan L and Östen Dahl (1989). *The Creation of Tense and Aspect Systems in the Languages of the World*. John Benjamins Amsterdam.

Cable, Seth (2013). "Beyond the Past, Present, and Future: Towards the Semantics of 'Graded Tense'in Gĩkũyũ". In: *Natural Language Semantics* 21.3, pp. 219–276.

— (2015). "The Tlingit Decessive and 'Discontinuous Past': The Curious Implicatures of Optional Past Tense". In: *Under revision for Natural Language and Linguistic Theory*.

Carlson, Greg (2012). "Habitual and Generic Aspect". In: *The Oxford handbook of Tense and Aspect*. Ed. by Robert I. Binnick. Oxford University Press.

Caselli, Tommaso, Valentina Bartalesi Lenzi, Rachele Sprugnoli, Emanuele Pianta, and Irina Prodanof (2011). "Annotating Events, Temporal Expressions and Relations in Italian: The It-TimeML Experience for the Ita-TimeBank". In: *Proceedings of the 5th Linguistic Annotation Workshop*. LAW V '11. Association for Computational Linguistics, pp. 143–151. ISBN: 978-1-932432-93-0. URL: http://dl.acm.org/citation.cfm?id=2018966.2018984.

Chambers, Nathanael, Daniel Cer, Trond Grenager, David Hall, Chloe Kiddon, Bill MacCartney, Marie-Catherine de Marneffe, Daniel Ramage, Eric Yeh, and Christopher D. Manning (2007). "Learning Alignments and Leveraging Natural Logic". In: *Proceedings of the ACL-PASCAL Workshop on Textual Entailment and Paraphrasing*. USA: Association for Computational Linguistics, 165–170.

Champollion, Lucas (2015). "The Interaction of Compositional Semantics and Event Semantics". In: *Linguistics and Philosophy* 38.1, pp. 31–66.

— (2017). *Parts of a Whole: Distributivity as a Bridge between Aspect and Measurement*. Vol. 66. Oxford University Press.

Chen, Danqi and Christopher Manning (2014). "A Fast and Accurate Dependency Parser Using Neural Networks". In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 740–750.

Cheng, Yuchang, Masayuki Asahara, and Yuji Matsumoto (2007). "Constructing a Temporal Relation Tagged Corpus of chinese Based on Dependency Structure Analysis". In: *14th International Symposium on Temporal Representation and Reasoning (TIME'07)*. IEEE, pp. 59–69.

Chomsky, Noam (1957). *Syntactic Structures*. De Gruyter Mouton.

Cipria, Alicia and Craige Roberts (2000). "Spanish imperfecto and Pretérito: Truth Conditions and Aktionsart Effects in a Situation Semantics". In: *Natural Language Semantics* 8.4, pp. 297–347.

Comrie, Bernard (1976). *Aspect: An Introduction to the Study of Verbal Aspect and Related Problems*. Vol. 2. Cambridge university press.

— (1985). *Tense*. Vol. 17. Cambridge University Press.

Condoravdi, Cleo, Dick Crouch, John Everett, Valeria Paiva, Reinhard Stolle, Danny Bobrow, and Martin van den Berg (2001). "Preventing Existence". In: *Proceedings of the International Conference on Formal Ontology in Information Systems - Volume 2001*. FOIS '01. Ogunquit, Maine, USA: ACM, pp. 162–173. ISBN: 1-58113-377-4. DOI: 10.1145/505168.505184. URL: http://doi.acm.org/10.1145/505168.505184.

Condoravdi, Cleo and Ashwini Deo (2008). "Aspect Shifts in Indo-Aryan". In: *Proceedings of the 18th International Congress of Linguistics*. Vol. 151, pp. 1–20.

— (2014). "Aspect shifts in Indo-Aryan and Trajectories of Semantic Change". In: *Language Change at the Syntax-semantics Interface*, pp. 261–292.

Copley, Bridget (2009). *The Semantics of the Future*. Routledge.

Croft, William (2012). *Verbs: Aspect and causal structure*. OUP Oxford.

Crouch, Dick, Mary Dalrymple, Ronald M. Kaplan, Tracy Holloway King, John T. Maxwell III, and Paula Newman (2017). *XLE Documentation*. Palo Alto Research Center.

Crouch, Richard (2005). "Packed Rewriting for Mapping Semantics to KR". In: *Proceedings of the Sixth International Workshop on Computational Semantics (IWCS-6)*. Tilburg, pp. 103–114.

Crouch, Richard and Tracy Holloway King (2006). "Semantics via F-Structure Rewriting". In: *Proceedings of the LFG06 Conference*. Ed. by Miriam Butt and Tracy Holloway King. Stanford, CA: CSLI Publications, pp. 145–165.

Crouch, Richard and Josef van Genabith (2000). *Linear Logic for Linguists. Introductory Course, ESSLLI-00*.

Curry, Haskell Brooks, Robert Feys, William Craig, J Roger Hindley, and Jonathan P Seldin (1958). *Combinatory Logic*. Vol. 1. North-Holland Amsterdam.

Dahl, Östen (1985). *Tense and Aspect Systems*. Oxford: Blackwell.

— (2000). *Tense and Aspect in the Languages of Europe*. Walter de Gruyter.

Dalrymple, Mary (1999). *Semantics and Syntax in Lexical Functional Grammar: The Resource Logic Approach*. MIT Press.

— (2001). *Lexical Functional Grammar*. Vol. 34. Syntax and Semantics. New York: Academic Press.

— (2015). "Morphology in the LFG Architecture". In: *Proceedings of the LFG'15 Conference*. CSLI Publications, pp. 64–83.

Dalrymple, Mary, Vaneet Gupta, John Lamping, and Vijay Saraswat (1999). "Relating Resource-based Semantics to Categorial Semantics". In: *Semantics and Syntax in Lexical Functional Grammar: the Resource Logic Approach*, pp. 261–280.

Dalrymple, Mary, John Lamping, and Vijay Saraswat (1993). "LFG Semantics via Constraints". In: *Proceedings of the Sixth Conference on European Chapter of the Association for Computational Linguistics*. EACL '93. Utrecht, The Netherlands: Association for Computational Linguistics, 97–105. ISBN: 9054340142. DOI: 10.3115/976744.976757. URL: https://doi.org/10.3115/976744.976757.

Dalrymple, Mary, Agnieszka Patejuk, and Mark-Matthias Zymla (2020). "XLE+Glue – A New Tool for Integrating Semantic Analysis in XLE". In: *Proceedings of the LFG'20 Conference, Australian National University*. Ed. by Miriam Butt, Tracy Holloway King, and Ida Toivonen. Stanford, CA: CSLI Publications.

Davidson, Donald (1967). "The Logical Form of Action Sentences". In.

De Swart, Henriëtte (2016). "Aspectual Composition and Recursion". In: *Mood, Aspect, Modality Revisited: New Answers to Old Questions*, p. 314.

De Marneffe, Marie-Catherine, Timothy Dozat, Natalia Silveira, Katri Haverinen, Filip Ginter, Joakim Nivre, and Christopher D Manning (2014). "Universal Stanford Dependencies: A Cross-linguistic Typology." In: *LREC*. Vol. 14, pp. 4585–4592.

De Marneffe, Marie-Catherine and Christopher D Manning (2008). *Stanford Typed Dependencies Manual*. Tech. rep. Technical report, Stanford University.

De Paiva, Valeria, Daniel G. Bobrow, Cleo Condoravdi, Richard Crouch, Lauri Karttunen, Tracy Holloway King, Rowan Nairn, and Annie Zaenen (2007). "Textual Inference Logic: Take Two". In: *Contexts and Ontologies Representation and Reasoning*, p. 32.

De Swart, Henriette (1998). "Aspect Shift and Coercion". In: *Natural Language & Linguistic Theory* 16.2, pp. 347–385.

— (2000). "Tense, Aspect and Coercion in a Cross-Linguistic Perspective". In: *Proceedings of the Berkeley Formal Grammar Conference*. Vol. 20. University of California, Berkeley.

De Wit, Astrid (2016). *The Present Perfective Paradox Across Languages*. Vol. 4. Oxford University Press.

Deo, Ashwini (2009). "Unifying the imperfective and the progressive: partitions as quantificational domains". In: *Linguistics and philosophy* 32.5, pp. 475–521.

— (to appear). "Imperfectivity". In: *Wiley's Companion to Semantics*.

Derczynski, Leon RA (2016). *Automatically Ordering Events and Times in Text*. Vol. 677. Springer.

— (2017). "Using a Framework of Tense and Aspect". In: *Automatically Ordering Events and Times in Text*. Springer, pp. 139–177.

Dickey, Stephen McCartney (1997). *Parameters of Slavic Aspect*. Indiana University.

Dorr, Bonnie J. (1992). "A Parameterized Approach to Integrating Aspect with Lexical-semantics for Machine Translation". In: *Proceedings of the 30th Annual Meeting on Association for Computational Linguistics*. ACL '92. Newark, Delaware: Association for Computational Linguistics, pp. 257–264. DOI: 10.3115/981967.982000. URL: https://doi.org/10.3115/981967.982000.

Dowty, David (1991). "Thematic Proto-roles and Argument Selection". In: *Language* 67.3, pp. 547–619.

Dowty, David R (1977). "Toward a Semantic Analysis of Verb Aspect and the English 'Imperfective' Progressive". In: *Linguistics and philosophy* 1.1, pp. 45–77.

— (1982). "Tenses, Time Adverbs, and Compositional Semantic Theory". In: *Linguistics and Philosophy* 5.1, pp. 23–55.

— (1986). "The Effects of Aspectual Class on the Temporal Structure of Discourse: Semantics or Pragmatics?" In: *Linguistics and philosophy* 9.1, pp. 37–61.

Ebert, Christian (2005). "Formal Investigations of Underspecified Representations". PhD thesis. University of London.

Ebert, Karen H (2000). "Progressive Markers in Germanic Languages". In: *Empirical approaches to language typology* 6, pp. 605–654.

Eckardt, Regine (2006). *Meaning Change in Grammaticalization: an Enquiry into Semantic Reanalysis*. Oxford University Press.

Ehrich, Veronika (1987). "The Generation of Tense". In: *Natural Language Generation*. Springer, pp. 423–440.

Falk, Yehuda N (1984). "The English Auxiliary System: A Lexical-Functional Analysis". In: *Language*, pp. 483–509.

Fara, Michael (2005). “Dispositions and Habituals”. In: *Noûs* 39.1, pp. 43–82. ISSN: 00294624, 14680068. URL: http://www.jstor.org/stable/3506173.

Ferreira, Marcelo (2016). “The Semantic Ingredients of Imperfectivity in Progressives, Habituals, and Counterfactuals”. In: *Natural Language Semantics* 24.4, pp. 353–397.

Ferro, Lisa, Inderjeet Mani, Beth Sundheim, and George Wilson (2001). *TIDES Temporal Annotation Guidelines-Version 1.0.2*. Tech. rep.

Filip, Hana (2012). “Lexical aspect”. In: *The Oxford handbook of tense and aspect*, pp. 721–751.

Filip, Hana, C Maienborn, K von Heusinger, and P Portner (2009). “Aspectual Class and Aktionsart”. In: *Noun Phrases and Verb Phrases*, p. 274.

Fleischman, Suzanne (1989). “Temporal distance: A basic linguistic metaphor”. In: *Studies in language* 13.1, pp. 1–50.

Frank, Anette and Annie Zaenen (2002). “Tense in LFG: Syntax and Morphology”. In: *How we say when it happens: Contributions to the Theory of Temporal Reference in Natural Language*, pp. 17–52.

Frege, Gottlob et al. (1892). “Über Sinn und Bedeutung”. In: *Zeitschrift für Philosophie und philosophische Kritik* 100.1, pp. 25–50.

Fry, John (2005). “Resource-logical Event Semantics for LFG”. Manuscript.

Garey, Howard B. (1957). “Verbal Aspect in French”. In: *Language* 33.2, pp. 91–110. ISSN: 00978507, 15350665. URL: http://www.jstor.org/stable/410722.

Garrette, Dan and Ewan Klein (2009). “An Extensible Toolkit for Computational Semantics”. In: *Proceedings of the Eighth International Conference on Computational Semantics*. Association for Computational Linguistics, pp. 116–127.

Gast, Volker, Lennart Bierkandt, Stephan Druskat, and Christoph Rzymski (2016). “Enriching TimeBank: Towards a More Precise Annotation of Temporal Relations in a Text”. In: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*. Ed. by Nicoletta Calzolari (Conference Chair), Khalid Choukri, Thierry Declerck, Sara Goggi, Marko Grobelnik, Bente Maegaard, Joseph Mariani, Helene

Mazo, Asuncion Moreno, Jan Odijk, and Stelios Piperidis. Portorož, Slovenia: European Language Resources Association (ELRA). ISBN: 978-2-9517408-9-1.

Gast, Volker, Lennart Bierkandt, and Christoph Rzymski (2015). "Creating and retrieving tense and aspect annotations with GraphAnno, a lightweight tool for multi-level annotation". In: *Proceedings 11th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, p. 23.

Gehrke, Berit (2008). "Ps in motion". PhD thesis. University of Utrecht.

Gennari, Silvia (1999). "Tense, Aktionsart and Sequence of Tense". In: *Empirical Issues in Formal Syntax and Semantics* 2.

Giorgi, Alessandra and Fabio Pianesi (1997). *Tense and Aspect: From Semantics to Morphosyntax*. Oxford University Press.

Girard, Jean-Yves (1987). "Linear Logic". In: *Theoretical Computer Science* 50.1, pp. 1 –101. ISSN: 0304-3975. DOI: https://doi.org/10.1016/0304-3975(87)90045-4. URL: http://www.sciencedirect.com/science/article/pii/0304397587900454.

— (1995). "Linear Logic: Its Syntax and Semantics". In: *London Mathematical Society Lecture Note Series*, pp. 1–42.

Glassman, Eugene H (1983). *Spoken Urdu*. Nirali Kitaben Publishing House.

Gotham, Matthew (2015). "Towards Glue Semantics for Minimalist Syntax". In: *Cambridge Occasional Papers in Linguistics* 8, p. 5683.

Gotham, Matthew and Dag Trygve Truslew Haug (2019). "Glue Semantics for Universal Dependencies". In.

Grice, H Paul (1975). "Logic and Conversation". In: *Syntax and Semantics Volume 3: Speech Acts*.

Grimshaw, Jane (1990). *Argument Structure*. the MIT Press.

Grønn, Atle and Arnim von Stechow (2010). "Complement Tense in Contrast: The SOT Parameter in Russian and English". In: *Oslo Studies in Language* 2.1.

Hale, Kenneth and Samuel Jay Keyser (1993). "On Argument Structure and the Lexical Expression of Syntactic Relations". In: ed. by Kenneth Hale and Samuel Jay Keyser. Cambridge, MA.: MIT Press, pp. 53–109.

Hallman, Peter (2009). “Proportions in Time: Interactions of Quantification and Aspect”. In: *Natural Language Semantics* 17.1, pp. 29–61.

Halvorsen, Per-Kristian (1983). “Semantics for Lexical-Functional Grammar”. In: *Linguistic Inquiry* 14.4, pp. 567–615.

Halvorsen, Per-Kristian and Ronald M. Kaplan (1988). “Projections and Semantic Description in Lexical-Functional Grammar”. In: *Proceedings of the International Conference on Fifth Generation Computer Technology*. Reprinted 1995 in Dalrymple et al. (eds) *Formal Issues in Lexical-Functional Grammar*, pp. 1116–1122.

Haug, Dag (2008). “Tense and Aspect for Glue Semantics: The Case of Participial XADJs”. In: *Proceedings of the LFG08 Conference*. Ed. by Miriam Butt and Tracy Holloway King. CSLI Publications, pp. 291–311.

Hautli, Annette, Tracy Holloway King, and Gilian Ramchand (2015). “Encoding Event Structure in Urdu/Hindi VerbNet”. In: *Proceedings of the The 3rd Workshop on EVENTS: Definition, Detection, Coreference, and Representation*, pp. 25–33.

Hay, Jennifer, Christopher Kennedy, and Beth Levin (1999). “Scalar structure underlies telicity in" Degree achievements"”. In: *Semantics and linguistic theory*. Vol. 9, pp. 127–144.

Heim, Irene (1994). “Comments on Abusch’s Theory of Tense”. In: *Ellipsis, Tense and Questions*, pp. 143–170.

Higginbotham, James (1985). “On Semantics”. In: *Linguistic Inquiry* 16.4, pp. 547–593.

— (2000). “On Events in Linguistic Semantics”. In: *Speaking of Events* 49, p. 80.

— (2009). *Tense, Aspect, and Indexicality*. OUP Oxford.

Hinrichs, Erhard (1986). “Temporal Anaphora in Discourses of English”. In: *Linguistics and Philosophy* 9.1, pp. 63–82.

Hinrichs, Erhard W (1985). “A Compositional Semantics for Aktionsarten and NP Teference in English”. PhD thesis. The Ohio State University.

Hintikka, Jaakko (1969). “Semantics for Propositional Attitudes”. In: *Models for Modalities*. Ed. by Donald Davidson, Jaako Hintikka, G. Nuchelmas, and Wesley C. Salmon. Dordrecht: D. Reidel Publishing Company, pp. 87–111.

Howard, William A (1980). “The Formulae-as-types Notion of Construction”. In: *To HB Curry: Essays on Combinatory Logic, Lambda calculus, and Formalism* 44, pp. 479–490.

Iatridou, Sabine (2000). “The Grammatical Ingredients of Counterfactuality”. In: *Linguistic Inquiry* 31.2, pp. 231–270.

Ide, Nancy and Harry Bunt (2010). “Anatomy of Annotation Schemes: Mapping to GrAF”. In: *Proceedings of the Fourth Linguistic Annotation Workshop*, pp. 247–255.

Ide, Nancy and James Pustejovsky (2017). *Handbook of Linguistic Annotation.* Springer Netherlands. ISBN: 9789402408799.

Im, Seohyun, Hyunjo You, Hayun Jang, Seungho Nam, and Hyopil Shin (2009). “KTimeML: Specification of Temporal and Event Expressions in Korean Text”. In: *Proceedings of the 7th Workshop on Asian Language Resources (ALR7)*, pp. 115–122.

Ingria, Robert and James Pustejovsky (2002). *TimeML Specification 1.2.* URL: \url{http://www.timeml.org/publications/timeMLdocs/timeml_1.2.1.html}.

Ippolito, Michela (2003). “Presuppositions and Implicatures in Counterfactuals”. In: *Natural Language Semantics* 11.2, pp. 145–186.

Jackendoff, Ray (1977). *X-Bar Theory: A Study of Phrase Structure.*

— (1987). “The Status of Thematic Relations in Linguistic Theory”. In: *Linguistic inquiry* 18.3, pp. 369–411.

— (1996). “The Proper Treatment of Measuring out, Telicity, and perhaps even Quantification in English”. In: *Natural Language & Linguistic Theory* 14.2, pp. 305–354.

Jóhannsdóttir, Kristın and Lisa Matthewson (2008). “Zero-marked Tense: The Case of Gitxsan”. In: *Proceedings-nels.* Vol. 37. 1, p. 299.

Joshi, Aravind K and Yves Schabes (1997). “Tree-Adjoining Grammars”. In: *Handbook of formal languages.* Springer, pp. 69–123.

Kalouli, Aikaterini-Lida and Richard Crouch (2018). “GKR: the Graphical Knowledge Representation for semantic parsing”. In: *Proceedings of the*

Workshop on Computational Semantics beyond Events and Roles, pp. 27–37.

Kamp, Hans (1979). “Events, Instants and Temporal Reference”. In: *Semantics from Different Points of View*. Springer, pp. 376–418.

— (1981). “A Theory of Truth and Semantic Representation”. In: *Formal Semantics – The Essential Readings*, pp. 189–222.

— (2017). “Events, Discourse Representations and Temporal Reference”. In: *Semantics and Pragmatics* 10.

Kaplan, Ronald and John T Maxwell (1995). “A Method for Disjunctive Constraint Satisfaction”. In: *Formal Issues in Lexical-Functional Grammar*, pp. 381–401.

Kaplan, Ronald M (1995a). “The formal Architecture of Lexical-Functional Grammar”. In: *Formal Issues in Lexical-Functional Grammar* 47, pp. 7–27.

— (1995b). “Three Seductions of Computational Psycholinguistics”. In: *Formal Issues in Lexical-Functional Grammar* 47.

Kaplan, Ronald M and John T Maxwell III (1988). “An Algorithm for Functional Uncertainty”. In: *Coling Budapest 1988 Volume 1: International Conference on Computational Linguistics*.

Kaplan, Ronald M, John T Maxwell III, Tracy Holloway King, and Richard Crouch (2004). “Integrating Finite-state Technology with Deep LFG Grammars”. In: *Proceedings of the workshop on combining shallow and deep processing for NLP (ESSLLI)*, pp. 11–20.

Kaplan, Ronald M and Annie Zaenen (1989). “Long-Distance Dependencies, Constituent Structure, and Functional Uncertainty”. In: *Alternative conceptions of Phrase Structure* 17, p. 42.

Kardos, Éva (2016). “Telicity Marking in Hungarian”. In: *Glossa* 1.1, p. 1.

Katz, Graham (2007). “Towards a Denotational Semantics for TimeML”. In: *Annotating, Extracting and Reasoning about Time and Events*. Springer, pp. 88–106.

— (2011). *Generics and Habituals and Textual Inference*. URL: \url{http://web.stanford.edu/~cleoc/Sem-Text-Inf/katz.pdf}.

Kennedy, Christopher and Louise McNally (1999). "From Event Structure to Scale Structure: Degree Modification in Deverbal Adjectives". In: *Semantics and linguistic theory*. Vol. 9, pp. 163–180.

King, Tracy Holloway (1995). "Configuring Topic and Focus in Russian". PhD thesis. Stanford University.

King, Tracy Holloway, Stefanie Dipper, Anette Frank, Jonas Kuhn, and John Maxwell (2000). "Ambiguity Management in Grammar Writing". In: *Workshop on linguistic theory and grammar implementation*, pp. 5–19.

King, Tracy Holloway, Martin Forst, Jonas Kuhn, and Miriam Butt (2005). "The Feature Space in Parallel Grammar Writing". In: *Research on Language and Computation* 3.2-3, pp. 139–163.

Kiparsky, Paul (1998). "Partitive Case and Aspect". In: *The projection of arguments: Lexical and compositional factors* 265, p. 307.

Kipka, Peter Francis (1990). "Slavic Aspect and its Implications". PhD thesis. Massachusetts Institute of Technology.

Klecha, Peter (2016). "Past is Perfect, Give or Take Adverbials". In: *Proceedings of Sinn und Bedeutung 20*.

Klecha, Peter and M Ryan Bochnak (2016). "Temporal Remoteness and Relativity". In: *Proceedings of NELS46*.

Klein, Wolfgang (1992). "The Present Perfect Puzzle". In: *Language*, pp. 525–552.

— (1994). *Time in Language*. Psychology Press.

— (1998). "Assertion and Finiteness". In: *Issues in the Theory of Language Acquisition: Essays in Honor of Jürgen Weissenborn*. Peter Lang, pp. 225–245.

— (2006). "On Finiteness". In: *Semantics in Acquisition*. Springer, pp. 245–272.

— (2009). "How time is encoded". In: *The expression of time* 3, pp. 1–43.

Kokkonidis, Miltiadis (2006). "A Glue/λ-DRT Treatment of Resumptive Pronouns". In: *Proceedings of the Eleventh ESSLLI Student Session*, pp. 51–63.

Koptjevskaja-Tamm, Maria (2002). *Nominalizations*. Routledge.

— (2008). *Action Nominal Constructions in the Languages of Europe*. De Gruyter Mouton.

Kracht, Marcus (2002). “On the Semantics of Locatives”. In: *Linguistics and Philosophy* 25.2, pp. 157–232.

Kratzer, Angelika (1981). “The notional category of modality”. In: *Formal semantics: The essential readings*, pp. 289–323.

— (1996). “Severing the External Argument from its Verb”. In: *Phrase Structure and the Lexicon*. Ed. by Johan Rooryck and Laurie Zaring. Dordrecht: Springer Netherlands, pp. 109–137. ISBN: 978-94-015-8617-7. DOI: 10.1007/978-94-015-8617-7_5. URL: https://doi.org/10.1007/978-94-015-8617-7_5.

— (1998). “More Structural Analogies between Pronouns and Tenses”. In: *Semantics and linguistic theory*. Vol. 8, pp. 92–110.

— (2004). *Telicity and the Meaning of Objective Case*. Citeseer.

— (2012). *Modals and Conditionals: New and Revised Perspectives*. Vol. 36. Oxford University Press.

— (2017). “Situations in Natural Language Semantics”. In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta. Winter 2017. Metaphysics Research Lab, Stanford University.

Kratzer, Angelika and Irene Heim (1998). *Semantics in Generative Grammar*. Blackwell Oxford.

Kratzer, Angelika et al. (1991). “Modality”. In: *Semantics: An International Handbook of Contemporary Research* 7, pp. 639–650.

Krifka, Manfred (1987). “Nominal Reference and Temporal Constitution: Towards a Semantics of Quantity”. In: *Proceedings of the 6th Amsterdam Colloquium, University of Amsterdam*, pp. 153–173.

— (1998). “The Origins of Telicity”. In: *Events and Grammar*. Springer, pp. 197–235.

Kripke, Saul A (1963). “Semantical Analysis of Modal Logic I Normal Modal Propositional Calculi”. In: *Mathematical Logic Quarterly* 9.5-6, pp. 67–96.

— (2014). *Semantical Analysis of Modal Logic II. Non-normal Modal Propositional Calculi*.

Kusumoto, Kiyomi (1999). "Tense in Embedded Contexts". PhD thesis. University of Massachusetts, Amherst.

— (2005). "On the Quantification over Times in Natural Language". In: *Natural Language Semantics* 13.4, pp. 317–357.

Lai, Catherine and Steven Bird (2004). "Querying and Updating Treebanks: A Critical Survey and Requirements Analysis". In.

Landman, Fred (1992). "The Progressive". In: *Natural Language Semantics* 1.1, pp. 1–32.

Lascarides, Alex and Nicholas Asher (1993). "Temporal Interpretation, Discourse Relations and Commonsense Entailment". In: *Linguistics and Philosophy* 16.5, pp. 437–493.

Lasersohn, Peter (1995). *Plurality, Conjunction and Events*. Dordrecht: Kluwer.

Lee, Kiyong (2008). "Formal semantics for interpreting temporal annotation". In: *Unity and Diversity of Languages*, pp. 97–108.

Leer, Jeffry A (1991). "The Schetic Categories of the Tlingit Verb". In.

Lev, Iddo (2007). "Packed Computation of Exact Meaning Representations". PhD thesis. Stanford University.

Levin, Beth and Malka Rappaport Hovav (2005). *Argument Realization*. Cambridge University Press.

Lewis, Clarence Irving (1918). *A Survey of Symbolic Logic*. University of California press.

Link, Godehard (1998). *Algebraic Semantics in Language and Philosophy*. CSLI Publications.

Löbner, Sebastian (2002). "Is the German Perfekt a Perfect Perfect". In: *More than Words: A Festschrift for Dieter Wunderlich*. Ed. by Ingrid Kaufmann and Barbara Stiebels. Akademie Verlag.

Lowe, John J (2019). "The Syntax and Semantics of Nonfinite Forms". In: *Annual Review of Linguistics* 5, pp. 309–328.

Martin, James H and Daniel Jurafsky (2009). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Pearson/Prentice Hall.

Matthewson, Lisa (2012). "On the (Non)-Future Orientation of Modals". In: *Proceedings of Sinn und Bedeutung*. Vol. 16, pp. 431–446.

Maxwell III, John T and Ronald M Kaplan (1989). "An Overview of Disjunctive Constraint Satisfaction". In: *Proceedings of the First International Workshop on Parsing Technologies*, pp. 18–27.

Meurer, Paul (2017). "From LFG Structures to Dependency Relations". In: *Bergen Language and Linguistics Studies* 8.1.

Młynarczyk, Anna Katarzyna (2004). "Aspectual Pairing in Polish". PhD thesis.

Moens, Marc and Mark Steedman (1988). "Temporal Ontology and Temporal Reference". In: *Computational Linguistics* 14.2, pp. 15–28.

Montague, Richard (1970a). "English as a Formal Language". In.

— (1970b). "Universal Grammar". In: *Theoria* 36.3, pp. 373–398. DOI: 10.1111/j.1755-2567.1970.tb00434.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1755-2567.1970.tb00434.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1755-2567.1970.tb00434.x.

Moseley, Christopher (2008). *Encyclopedia of the World's Endangered Languages*. Routledge.

Mucha, Anne (2015). "Temporal Interpretation and Cross-Linguistic Variation". PhD thesis. University of Potsdam.

Musan, Renate (2002). *The German Perfect: Its Semantic Composition and its Interactions with Temporal Adverbials*. Vol. 78. Springer Science & Business Media.

Myers, Skatje and Martha Palmer (2019). "ClearTAC: Verb Tense, Aspect, and Form Classification Using Neural Nets". In: *Proceedings of the First International Workshop on Designing Meaning Representations*, pp. 136–140.

Nivre, Joakim, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman (2020). "Universal Dependencies v2: An Evergrowing Multilingual Treebank Collection". In: *arXiv preprint arXiv:2004.10643*.

Nivre, Joakim, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D Manning, Ryan McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. (2016). "Universal Dependencies v1: A Multilingual Treebank Collection". In: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pp. 1659–1666.

Nordlinger, Rachel and Louisa Sadler (2004). "Nominal Tense in Crosslinguistic Perspective". In: *Language*, pp. 776–806.

Ogihara, Toshiyuki (1994). "Adverbs of Quantification and Sequence-of-Tense Phenomena". In: *Semantics and Linguistic Theory*. Vol. 4, pp. 251–267.

— (1999). "Double-access Sentences Generalized". In: *Semantics and Linguistic Theory*. Vol. 9, pp. 224–236.

Osborne, Timothy and Kim Gerdes (2019). "The Status of Function Words in Dependency Grammar: A Critique of Universal Dependencies (UD)". In: *Glossa (Online)*.

Palmer, Martha, Daniel Gildea, and Paul Kingsbury (2005). "The Proposition Bank: An Annotated Corpus of Semantic Roles". In: *Computational Linguistics* 31.1, pp. 71–106.

Parsons, Terence (1990). *Events in the Semantics of English*. Vol. 5. MIT Press.

Partee, Barbara H. (1984). "Nominal and Temporal Anaphora". In: *Linguistics and Philosophy* 7.3, pp. 243–286. ISSN: 01650157, 15730549. URL: http://www.jstor.org/stable/25001168.

Partee, Barbara Hall (1973). "Some Structural Analogies between Tenses and Pronouns in English". In: *The Journal of Philosophy* 70.18, pp. 601–609.

Passonneau, Rebecca J (1988). "A Computational Model of the Semantics of Tense and Aspect". In: *Computational Linguistics* 14.2, pp. 44–60.

Petrov, Slav, Dipanjan Das, and Ryan McDonald (2011). "A Universal Part-of-Speech Tagset". In: *arXiv preprint arXiv:1104.2086*.

Plungian, Vladimir A and Johan van der Auwera (2006). "Towards a Typology of Discontinuous Past Marking". In: *STUF–Sprachtypologie und Universalienforschung* 59.4, pp. 317–349.

Pollard, Carl and Ivan A Sag (1994). *Head-driven Phrase Structure Grammar*. University of Chicago Press.

Portner, Paul (1998). "The Progressive in Modal Semantics". In: *Language*, pp. 760–787.

— (2019). "Perfect and Progressive". In: *Semantics-Noun Phrases and Verb Phrases*, p. 313.

Priest, Graham (2008). *An introduction to non-classical logic: From if to is*. Cambridge University Press.

Prior, Arthur N (1962). "Tense-logic and the Continuity of Time". In: *Studia Logica: An International Journal for Symbolic Logic* 13, pp. 133–151.

— (1967). *Past, Present and Future*. Vol. 154. Clarendon Press Oxford.

Przepiórkowski, Adam (2017). "Hierarchical Lexicon and the Argument/Adjunct Distinction". In: *The Proceedings of the LFG'17 Conference*, pp. 348–367.

Przepiórkowski, Adam and Agnieszka Patejuk (2018). "Arguments and Adjuncts in Universal Dependencies". In: *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 3837–3852.

Pustejovsky, James (1991). "The Syntax of Event Structure". In: *Cognition* 41.1-3, pp. 47–81.

— (2017). "ISO-TimeML and the Annotation of Temporal Information". In: *Handbook of Linguistic Annotation*. Ed. by Nancy Ide and James Pustejovsky. Dordrecht: Springer Netherlands, pp. 941–968. ISBN: 978-94-024-0881-2. DOI: 10.1007/978-94-024-0881-2_35. URL: https://doi.org/10.1007/978-94-024-0881-2_35.

Pustejovsky, James, José M Castano, Robert Ingria, Roser Sauri, Robert J Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir R Radev (2003). "TimeML: Robust specification of Event and Temporal Expressions in Text". In: *New directions in question answering* 3, pp. 28–34.

Pustejovsky, James, Kiyong Lee, Harry Bunt, and Laurent Romary (2010). "ISO-TimeML: An International Standard for Semantic Annotation". In: *LREC*. Vol. 10, pp. 394–397.

Pustejovsky, James, Roser Saurí, Andrea Setzer, Rob Gaizauskas, and Bob Ingria (2002). "TimeML Annotation Guidelines". In: *TERQAS Annotation Working Group* 23.

Ramchand, Gillian (1997). *Aspect and Predication: The Semantics of Argument Structure*. Oxford University Press.

— (2004). "Time and the Event: The Semantics of Russian Prefixes". In: *Nordlyd* 32.2.

— (2005). "Post-Davidsonianism". In: *Theoretical linguistics* 31.3, pp. 359–373. DOI: doi:10.1515/thli.2005.31.3.359. URL: https://doi.org/10.1515/thli.2005.31.3.359.

— (2018). *Situations and Syntactic Structures: Rethinking Auxiliaries and Order in English*. MIT Press.

Ramchand, Gillian and Peter Svenonius (2014). "Deriving the Functional Hierarchy". In: *Language Sciences* 46, pp. 152–174.

Ramchand, Gillian Catriona (2008). *Verb Meaning and the Lexicon: A First Phase Syntax*. Vol. 116. Cambridge University Press.

Ramm, Anita, Sharid Loáiciga, Annemarie Friedrich, and Alexander Fraser (2017). "Annotating Tense, Mood and Voice for English, French and German". In: *Proceedings of ACL 2017, System Demonstrations*, pp. 1–6.

Rappaport Hovav, Malka and Beth Levin (1998). "Building Verb Meanings". In: *The projection of arguments: Lexical and compositional factors*, pp. 97–134.

Reichenbach, Hans (1947). "The Tenses of Verbs". In: *Elements of Symbolic Logic*, pp. 287–298.

Reis Silva, Amelia and Lisa Matthewson (2007). "An Instantaneous Present Tense in Blackfoot". In: *Proceedings of SULA* 4, pp. 197–213.

Rivero, María Luisa and Ana Arregui (2010). "Aspectual Microvariation: the Case of Slavic Imperfectives". In: *Proceedings of the 2010 Annual Conference of the Canadian Linguistic Association*.

Romanova, Eugenia (2007). "Constructing Perfectivity in Russian". PhD thesis. Universitetet i Tromsø.

Romero, Maribel (2014). "'Fake Tense' in Counterfactuals: A Temporal Remoteness Approach". In: *The Art and Craft of Semantics: A Festschrift for Irene* 2. Ed. by HeimLuka Crnic and Uli Sauerland, pp. 47–63.

Rosén, Victoria, Koenraad De Smedt, Paul Meurer, and Helge Dyvik (2012). "An Open Infrastructure for Advanced Treebanking". In: *META-RESEARCH Workshop on Advanced Treebanking at LREC2012*. Ed. by Jan Hajič, Koenraad De Smedt, Marko Tadić, and António Branco, pp. 22–29.

Rothstein, Susan (2016). "Aspect". In: *The Cambridge Handbook of Formal Semantics*. Ed. by Maria Aloni and Paul Dekker. Cambridge University Press, pp. 342–368.

Saurí, Roser, Jessica Littman, Bob Knippen, Robert Gaizauskas, Andrea Setzer, and James Pustejovsky (2006). "TimeML Annotation Guidelines". In: *Version* 1.1, p. 31.

Schaden, Gerhard (2010). "German Gerade: A Unified Analysis of its Temporal Uses". In: *Layers of Aspect*, pp. 167–188.

Schuler, Karin Kipper (2005). "VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon". In.

Schuster, Sebastian and Christopher D Manning (2016). "Enhanced English Universal Dependencies: An Improved Representation for Natural Language Understanding Tasks". In: *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pp. 2371–2378.

Shimoyama, Junko (2006). "Indeterminate Phrase Quantification in Japanese". In: *Natural Language Semantics* 14.2, pp. 139–173. ISSN: 1572-865X. DOI: 10.1007/s11050-006-0001-5. URL: https://doi.org/10.1007/s11050-006-0001-5.

Singh, Mona (1991). "The Perfective Paradox: Or How to Eat your Cake and Have it Too". In: *Annual Meeting of the Berkeley Linguistics Society*. Vol. 17. 1, pp. 469–479.

Smith, Carlota (2006). "The Pragmatics and Semantics of Temporal Meaning". In: *Proceedings of the 2004 Texas Linguistics Society Conference*, pp. 92–106.

Smith, Carlotta (1991). *The Parameter of Aspect*. Dordrecht: Kluwer Academic Publishers.

Sneddon, James Neil, K Alexander Adelaar, Dwi N Djenar, and Michael Ewing (2012). *Indonesian: A Comprehensive Grammar*. Routledge.

Snow, Rion, Brendan O'Connor, Daniel Jurafsky, and Andrew Y. Ng (2008). "Cheap and Fast—but is It Good?: Evaluating Non-expert Annotations for Natural Language Tasks". In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. EMNLP '08. Honolulu, Hawaii: Association for Computational Linguistics, pp. 254–263. URL: http://dl.acm.org/citation.cfm?id=1613715.1613751.

Spreyer, Kathrin and Anette Frank (2008). "Projection-based Acquisition of a Temporal Labeller". In: *Proceedings of the Third International Joint Conference on Natural Language Processing: Volume-I*.

Squartini, Mario and Pier Marco Bertinetto (2000). "The Simple and Compound Past in Romance Languages". In: *Tense and Aspect in the Languages of Europe*, p. 403.

Steedman, Mark (1995). "Dynamic Semantics for Tense and Aspect". In: *International Joint Conference on Artificial Intelligence*, pp. 1292–1298.

— (1997). "Chapter 16 - Temporality". In: *Handbook of Logic and Language*. Ed. by Johan van Benthem and Alice ter Meulen. North-Holland, pp. 895 –938.

Szabó, Zoltán Gendler (2004). "On the Progressive and the Perfective". In: *Noûs* 38.1, pp. 29–59.

Tenny, Carol L (1995). "How Motion Verbs are Special: The Interaction of Semantic and Pragmatic Information in Aspectual verb meanings". In: *Pragmatics & Cognition* 3.1, pp. 31–73.

Tenny, Carol Lee (1987). "Grammaticalizing aspect and affectedness". PhD thesis. Massachusetts Institute of Technology.

Tonhauser, Judith (2015). "Cross-Linguistic Temporal Reference". In: *Linguistics* 1, pp. 129–154.

Traugott, Elizabeth Closs and Richard B Dasher (2001). *Regularity in Semantic Change*. Vol. 97. Cambridge University Press.

UzZaman, Naushad, Hector Llorens, Leon Derczynski, James Allen, Marc Verhagen, and James Pustejovsky (2013). "SemEval-2013 Task 1: Tempeval-3: Evaluating Time Expressions, Events, and Temporal Relations". In: *Second Joint Conference on Lexical and Computational Semantics (* SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pp. 1–9.

Van Genabith, Josef and Richard Crouch (1997). "How to Glue a Donkey to an G-structure or Porting a Dynamic Meaning Representation Language into LFG's Linear Logic based Glue Language Semantics". In: *International Workshop for Computational Semantics, Tilburg, Proceedings*. Citeseer, pp. 52–65.

Vendler, Zeno (1957). "Verbs and times". In: *The philosophical review* 66.2, pp. 143–160.

Verhagen, Marc, Robert Gaizauskas, Frank Schilder, Mark Hepple, Graham Katz, and James Pustejovsky (2007). "SemEval-2007 task 15: Tempeval Temporal Relation Identification". In: *Proceedings of the 4th International Workshop on Semantic Evaluations*. Association for Computational Linguistics, pp. 75–80.

Verhagen, Marc, Roser Sauri, Tommaso Caselli, and James Pustejovsky (2010). "SemEval-2010 task 13: TempEval-2". In: *Proceedings of the 5th International Workshop on Semantic Evaluation*. Association for Computational Linguistics, pp. 57–62.

Verkuyl, Hendrik Jacob (1972). *On the Compositional Nature of the Aspects*. Vol. 15. Springer Science & Business Media.

Verkuyl, Henk J (1993). *A Theory of Aspectuality: The Interaction between Temporal and Atemporal Structure*. 64. Cambridge University Press.

Vlach, Frank (1981). "The Semantics of the Progressive in Tense and Aspect". In: *Syntax and Semantics* 14. Ed. by Philip Tedeschi and Annie Zaenen, pp. 271–292.

Von Fintel, Kai (1994). "Restrictions on Quantifier Domains". PhD thesis. University of Massachusetts at Amherst.

Von Stechow, Arnim (2009). "Tenses in Compositional Semantics". In: *The Expression of Time* 2009, pp. 129–166.

Wechsler, Stephen (2005). "Resultatives under the 'Event-Argument Homomorphism' Model of Telicity". In: *The Syntax of Aspect: Deriving Thematic and Aspectual Interpretation* 10, p. 255.

Wedekind, Jürgen and Ronald Kaplan (Jan. 1993). "Type-Driven Semantic Interpretation of f-Structures." In: pp. 404–411. DOI: 10.3115/976744.976791.

Wynne, Martin (2005). *Developing Linguistic Corpora: a Guide to Good Practice*. Vol. 92. Oxbow Books Oxford.

Xue, Nianwen and Yuping Zhou (2010). "Applying Syntactic, Semantic and Discourse Constraints in Chinese Temporal Annotation". In: *Coling 2010: Posters*, pp. 1363–1372.

Zeman, Daniel (2008). "Reusable Tagset Conversion Using Tagset Drivers". In: *LREC*. Vol. 2008, pp. 28–30.

Zucchi, Alessandro (2013). *The Language of Propositions and Events: Issues in the Syntax and the Semantics of Nominalization*. Vol. 51. Springer Science & Business Media.

Zwarts, Joost (2005). "Prepositional Aspect and the Algebra of Paths". In: *Linguistics and Philosophy* 28.6, pp. 739–779.

Zymla, Mark-Matthias (2017). "Cross-Linguistically Viable Treatment of Tense and Aspect in Parallel Grammar Development". In: *Proceedings of the LFG17 Conference*. CSLI Publications.

— (2018). "Annotation of the Syntax/Semantics Interface as a Bridge between Deep Linguistic Parsing and TimeML". In: *Proceedings 14th Joint ACL-ISO Workshop on Interoperable Semantic Annotation*, pp. 53–59.

— (2019). "Aspectual Reasoning in LFG – A Computational Approach to Grammatical and Lexical Aspect". In: *Proceedings of the LFG'19 Conference, Australian National University*. Ed. by Miriam Butt, Tracy Holloway King, and Ida Toivonen. Stanford, CA: CSLI Publications, pp. 353–373. URL: http://cslipublications.stanford.edu/LFG/2019/lfg2019-zymla.pdf.