

David Hägele*, Christoph Schulz, Cedric Beschle, Hannah Booth, Miriam Butt, Andrea Barth, Oliver Deussen, and Daniel Weiskopf

Uncertainty visualization: Fundamentals and recent developments

<https://doi.org/10.1515/itit-2022-0033>

Received May 17, 2022; revised August 12, 2022; accepted August 13, 2022

Abstract: This paper provides a brief overview of uncertainty visualization along with some fundamental considerations on uncertainty propagation and modeling. Starting from the visualization pipeline, we discuss how the different stages along this pipeline can be affected by uncertainty and how they can deal with this and propagate uncertainty information to subsequent processing steps. We illustrate recent advances in the field with a number of examples from a wide range of applications: uncertainty visualization of hierarchical data, multivariate time series, stochastic partial differential equations, and data from linguistic annotation.

Keywords: Uncertainty visualization, multivariate data, hierarchical data, partial differential equations, linguistics

ACM CCS: Human-centered computing → Visualization theory, concepts and paradigms

*Corresponding author: David Hägele, University of Stuttgart, Visualization Research Center, D-70569, Stuttgart, Germany, e-mail: david.haegle@visus.uni-stuttgart.de

Christoph Schulz, Daniel Weiskopf, University of Stuttgart, Visualization Research Center, D-70569, Stuttgart, Germany, e-mails: christoph.schulz@visus.uni-stuttgart.de, daniel.weiskopf@visus.uni-stuttgart.de

Cedric Beschle, Andrea Barth, University of Stuttgart, Institute of Applied Analysis and Numerical Simulation, D-70569, Stuttgart, Germany, e-mails: cedric.beschle@ians.uni-stuttgart.de, andrea.barth@ians.uni-stuttgart.de

Hannah Booth, Ghent University, Department of Linguistics, B-9000, Gent, Belgium, e-mail: hannah.booth@ugent.be

Miriam Butt, University of Konstanz, Department of Linguistics, D-78457, Konstanz, Germany, e-mail: miriam.butt@uni-konstanz.de

Oliver Deussen, University of Konstanz, Department of Computer and Information Science, D-78457, Konstanz, Germany, e-mail: oliver.deussen@uni-konstanz.de

1 Introduction

Virtually any information comes with some aspect of uncertainty. The way we visually represent uncertainty can have a strong influence on how we perceive such information. Still, uncertainty is often neglected in general, and thus rarely considered in visual analysis and dissemination processes. There are two reasons for this: First, we tend to interpret visualizations as truthful because they are easier to understand in this case; less straightforward interpretation, as usually accompanied by uncertain data, must be learned. Second, many visualization techniques cannot handle uncertain data; then, the only option is to consider the most likely realization of the data and omit aspects of uncertainty.

Let us illustrate the topic and the aforementioned problems for the simple example of weather forecasts. Ideally, one would like to have a definite forecast concerning temperature and rain, although weather cannot be predicted with pin-point precision. Showing only the most likely temperature and rain, i. e., without the forecast range, it is virtually impossible to assess how much one can trust the weather forecast. Interestingly, there are different ways in which weather forecasts are communicated, for example, in TV news shows in the U.S and Germany. A probabilistic representation, as often used in North America, will tell the viewer the percentage of a certain event (rain in a given region), this way conveying the uncertainty, whereas in Germany, weather charts often just show the most likely outcome of the forecast as a single illustration.

In the example of weather forecasts, we can directly convey uncertainty by numbers or simple visualizations like bar or line charts. However, what can we do if we have complex data like massive tree structures, multivariate data, or simulation results from science and engineering along with uncertainties? Here, it is already challenging to find visual representations that ignore uncertainty, but the problem becomes even more pronounced for uncertainty visualization.

In this paper, we want to provide a brief introduction into the general topic of uncertainty visualization, discussing some background, terminology, and fundamental

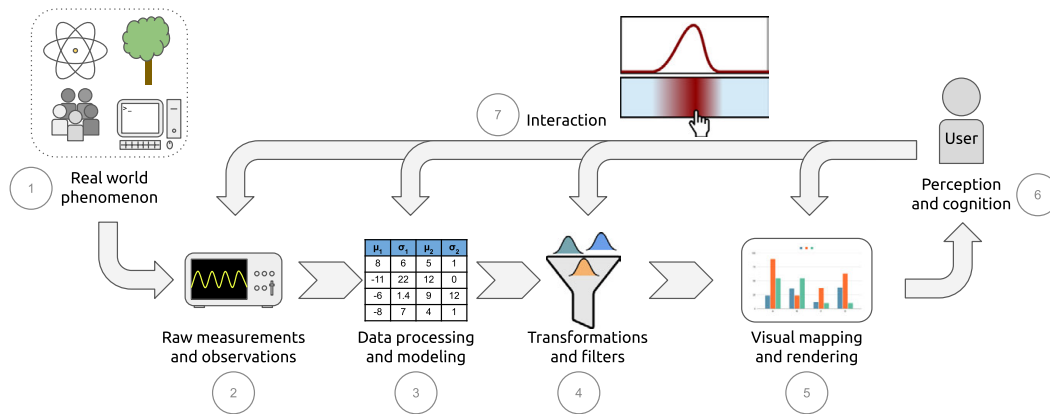


Figure 1: The visualization pipeline where uncertainty can be introduced and propagated in any stage.

concepts. Here, we emphasize the need for making the entire visualization process uncertainty-aware, and for appropriate quantitative modeling and propagation of uncertainty. We also illustrate recent advances in the fields with examples covering a wide range of applications and types of data. Specifically we discuss possibilities to visualize uncertainty in hierarchical data, as well as multivariate time series. Furthermore, we elaborate on modeling uncertain physical phenomena with the example of stochastic partial differential equations, and illustrate the diverse kinds of uncertainty that are present in the application field of linguistic annotation.

2 Overview of uncertainty visualization

In the visualization community, *uncertainty* refers to information from which *disagreement* and *credibility* can be inferred. Orthogonal to these assessment aspects, we distinguish between *measurement* uncertainty (e. g., accuracy and precision), *completeness* uncertainty (e. g., missing values and sampling), and *inference* uncertainty (e. g., prediction and modeling) [41]. Other communities prefer risk, confidence, or trust [38]. The latter term has become increasingly popular in the visualization community since it reduces uncertainty visualization to one question: *What reveals imperfect information and separates it from more trustworthy data?*

These different aspects of uncertainty indicate that there is more to uncertainty visualization than just the direct rendering of images. To this end, we will use the visualization pipeline [22, 14] to illustrate how uncertainty plays a role in the different stages required for visualiza-

tion. For more detailed background information and coverage of work in uncertainty visualization, we refer to survey papers such as [36, 10, 7, 44]. Our discussion adopts the structure of [44], walking through the different stages of the visualization process.

2.1 Visualization pipeline

The process of creating visual representations of data is commonly described by the visualization pipeline; see Fig. 1. Briefly, the process emerges from a real-world phenomenon that we want to study or explore. For this purpose, data is acquired, e. g., by taking measurements or collecting observations about the phenomenon. Usually, the raw data needs further processing or modeling to be represented, stored, and passed to the following stages of the pipeline. Next, filtering and transformation extract meaningful aspects of interest from the data (e. g., subsets, similarities, order, and grouping). To display such distilled data, graphical elements are generated and shown to the user during visual mapping and rendering.

The crucial consideration to be made here is that any stage of the pipeline can introduce uncertainty into the visualization process (cf. Fig. 1). We argue that uncertainties introduced in earlier stages are discarded later in the pipeline when stages are oblivious to the uncertainty, and as a result, the credibility of the visualization is impaired. To provide faithful visualizations each stage needs to be made aware of the uncertainty and explicitly propagate it so that we have a chance to communicate it to the user. In the following subsections, we discuss the different stages and how they may introduce or propagate uncertainty. For uncertainty-aware computing, the uncertainty has to be modeled appropriately. Despite the numerous meanings

of uncertainty, it is often treated as a probability or probability density that quantifies how likely, representative, or credible an observation is.

With this view from the visualization pipeline, we can distinguish between *visualization of uncertainty*, and the *uncertainty in visualization* [10]. While the former is concerned with the display of uncertainty in data and the topic we focus on, the latter considers how the visualization process introduces inaccuracies and distortion of the data, which is also a form of uncertainty. In other words: we need to compute how uncertainty is propagated through the pipeline (starting from uncertain input data) and how the stages of the pipeline add uncertainty.

2.2 Real-world uncertainty

With the different flavors of uncertainty outlined earlier, there are manifold different examples in the real-world phenomena under investigation. We encounter uncertainty in many domains, e. g., in engineering or machining where it is equivalent to *accuracy* or *precision*, in natural language it can mean *ambiguity* of a word or sentence. While uncertainty may already be part of the phenomenon that should be studied ①, it is not required right from the beginning for uncertainty visualization since uncertainty can also be introduced in later stages.

2.3 Uncertainty in data collection

To study a phenomenon we need to observe it in some way ②—for example, by taking measurements of weather conditions using sensors, asking people questions in a survey, gathering historical evidence, or other means of data collection. This data acquisition process can also introduce uncertainty. Measurements may be *inexact* or *noisy*, sensors could temporarily fail, people may be uncertain about their answers, and the sampling may be *unrepresentative*, leading to *missing* and *insufficient* information.

2.4 Data processing and modeling

Usually, the collected data from the previous stage needs to be processed further before use (e. g., data cleansing or unification) ③, which can introduce further uncertainty. For example, incomplete data records are removed and leave a gap, missing data is inferred (interpolated), precision is lost or made up when converting records of different sources to the same format (e. g., nanosecond to a millisecond resolution, high dynamic range imagery to sRGB

color space). The raw data is often used to create models or abstract representations, such as continuous functions through parameter estimation or syntax trees through syntactic parsing of text.

2.5 Uncertainty-aware filtering and transformation

In step ④, the data is prepared for display. Showing all data is often not desired since it is too much information to cognitively process by a human or even too much to display. Instead, subsets of the whole data may be selected that will be subject to display, e. g., using only records of a relational database that match a specific query, or sampling a continuous function with a particular resolution. The data is usually transformed into another representation that can be more easily understood. Such transformations include projecting high-dimensional data to 2D or 3D, or organizing data into groups by clustering. Both examples introduce uncertainty. Projections need to discard or aggregate information to reduce dimensionality and cannot always preserve the relationships between data items. Some clustering algorithms can yield probabilistic assignments to clusters, introducing uncertainty about cluster membership on purpose (e. g., Gaussian mixture model).

Transformations and filters must be made aware of the uncertainty in the pipeline so that they can take it into account, quantify it, and propagate it explicitly. However, due to the nonlinearity of many of these techniques, it is not trivial to do this and uncertainty might be distorted considerably.

2.6 Visual mapping of uncertainty

The transformed data from the previous stage is ready to be brought to screen. This visual mapping step ⑤ translates data into a graphical representation, i. e., it defines how visual variables (such as color, shape, size, and position) are used to convey the data values and their relationships. In a scatter plot, for example, coordinates are mapped to the position of dots, and coloring can be used to encode group membership of the dots.

Well-known examples of uncertainty visualization for simple data records are error bars or box plots that spatially encode data variability. Frequentist approaches to depicting probabilities use a finite set of samples to show possible events or realizations of random objects such as

samples from a bivariate distribution in a scatterplot. Using summary statistics or fuzzy arithmetics, specific measures such as standard deviation, quantiles, or maximal ranges can be shown as for confidence intervals.

Visual mapping is at the core of visualization. Accordingly, there are manifold techniques available for uncertainty visualization; see the survey papers listed earlier. On a generic level, choices of visual variables for encoding uncertainty are discussed and evaluated by MacEachren et al. [33].

2.7 Interaction and integration

We actively include the human in the visualization process, as typical for visual analytics systems. We do this by allowing the user to react to the visualization at any stage of the pipeline ⑦. This could mean loading a different data set, changing data cleansing strategies or assumptions of the data model, refining a data query, or requesting a different chart type.

Such interactions can be quite uncertain, e. g., entering your daily calorie consumption may be based on a rough estimate only. We can support uncertain interaction, e. g., by providing fuzzy selection tools or allowing users to specify their confidence about their input [25, 21].

The integration of uncertainty throughout the pipeline is challenging for many reasons and requires each stage to be made aware of the uncertainties in the system. Explicit propagation to subsequent stages is complicated due to nonlinear transformations and series of transformations. Wu et al. [46] showed an example of uncertainty integration into the visualization process. For uncertainty visualization systems, we recommend assessing the pipeline and targeting the most impactful parts regarding the analysis tasks and goals.

2.8 Perception, cognition, and evaluation

Finding effective mappings for uncertainty that can be understood and read accurately by a human is challenging due to the limited number of visual variables and perceptual restrictions of what a human can process and differentiate (stage ⑥). Even more challenging is the mapping of data with uncertainty, where we have to take special care that the perceived uncertainty matches the data uncertainty. It turned out that even experts struggle with judging the encoded uncertainty of comparably simple visualizations such as error bars correctly [5].

Beyond the perceptual aspects of uncertainty displays, conveying relationships such as covarying probabilities or dependency poses a cognitive challenge [27]. Assessment of the effectiveness, readability, and interpretability of uncertainty visualizations and interaction with the pipeline, is a crucial component. Controlled laboratory and large scale crowd-sourced user studies are means to evaluate approaches for uncertainty visualization [26].

3 Hierarchical data

Let us now illustrate the generic considerations from the visualization pipeline by a few concrete examples of visualization techniques. This section starts by discussing a visual mapping technique for hierarchical data with uncertainty, which is an example of handling uncertainty in the respective stage ⑤ in the visualization pipeline (cf. Fig. 1). There are many specialized methods to visualize hierarchical data; Schulz et al. [40] provide a survey of hierarchy visualizations.

Treemaps have been shown to be effective in conveying implicit hierarchical information [17]. They divide the canvas according to the relative size of sub-hierarchies (aggregated data values). The challenge for such techniques is to find a balanced solution combining readability, compactness, and visual scalability. Uncertainty visualization adds another challenge: certain and uncertain aspects of data values propagate differently, which is related to the modeling stage ③. For example, the mean μ propagates differently from the standard deviation σ . The aggregation of values from children to parent nodes is a summation of the children's random variables: $X_{1,\dots,n} = \sum_i^n X_i$. For propagation of independent probabilities from n child nodes to their parent node, the formulas for mean and standard deviation are:

$$\mu_{1,\dots,n} = \sum_i^n \mu_i, \quad \sigma_{1,\dots,n} = \sqrt{\sum_i^n \sigma_i^2} \quad (1)$$

Thus, the presence of uncertainty violates the visual summation of sub-hierarchies, which is common to many non-uncertainty-aware treemap techniques; therefore, there is a direct implication of the uncertainty model on the requirements for the choice of visual mapping ⑤ later in the pipeline.

Fig. 2 shows one possible solution: an uncertainty-aware Bubble Treemap [19] that deliberately introduces whitespace. Unlike snapshot-oriented financial treemaps,




Figure 2: SP-500 stock data for the course of a single week. Variances are shown by using undulating outlines with an amplitude corresponding to the amount of variance. © 2018 IEEE. Reprinted, with permission, from Görtler et al. [19].

it shows the *Standard & Poor's 500 index* (S&P 500) over the course of one week, grouped by sectors and sub-industries. The size of the circles represent the mean closing price of the stock for the given week. We use the contour to illustrate the standard deviation of each stock. The treemap shows that most stocks were stable during the given period of time, while some had larger variations. By looking at the waviness of the contours, it is relatively easy to identify the stock with the biggest changes, since the variance is reflected in all the contours of the respective sub-systems. In this case, the reason for the big changes were a 5-for-1 stock split, which led to single stock only having a fifth of the original value.

An alternative technique relies on rectangular treemaps that can be made uncertainty-aware, e. g., using well-balanced transparency-based uncertainty masks [42]. This lead to an optimization problem that can be solved to address the readability of rectangular treemaps, which, in particular, depends on the aspect ratio.

Both examples of visualization techniques extend the mapping step ⑤ so that it can explicitly show data uncertainty, thus supporting perception and processing by the human recipient (stage ⑥). They are designed in a way that avoids introducing additional uncertainty, i. e., here, we focus on visualization of data uncertainty, and not on uncertainty in visualization.

4 Multivariate time series

Our next example is the visualization of time series of multivariate data. For background reading on multivariate data, we refer to a survey by Liu [31]. We pick Time Curves [1] as a starting point. They reduce high-dimensional data items to two dimensions while illustrating the temporal succession by connecting consecutive points using a Bézier curve. On point-level, this allows for statements regarding similarity, whereas the geometric characteristics can be interpreted on curve-level (point density, degree of stagnation, oscillation, regularity, etc.).

Here, we want to focus on the uncertainty associated with an ensemble of Time Curves [9]. To assess credibility and disagreement, we can consider the representativeness of a Time Curve. Such an approach requires a quantifiable description to assess if a Time Curve is representative or exceptional. For this purpose, we use a non-parametric statistical approach called *functional band depth* [32] that establishes an order from most-representative to least-representative using convex hull inclusion testing. This refers to the transformation and filtering stage ④ of the visualization pipeline.

Fig. 3 shows a health-related example: 16 patients of the Massachusetts General Hospital/Marquette Foundation (MGH/MF) Waveform data set [45]. When displaying the complete set of Time Curves, cluttering prohibits useful analysis (Fig. 3 left). Since the metric uses convex hulls, we can abstract most curves into two convex hulls, one for 50 % and one for 75 % most-representative patients, retaining the most-representative Time Curve (Fig. 3 middle). This visual mapping (stage ⑤) corresponds to the univariate boxplot (Fig. 3 topmost), hence the name Time Curve Boxplot. Inspection of the outliers (Fig. 3 right), i. e., Time Curves that are not completely contained in the 50 % or 75 % convex hull, shows that not only spatial but also temporal inclusion are of relevance. The scarf plot (Fig. 3 bottom) depicts the true extents of inclusion in the respective high-dimensional hulls, both temporal and spatial. This additional visualization reduces the visualization uncertainty that is inherent to the dimensionality reduction underlying the Time Curves. Works addressing visualization of uncertainty arising from dimensionality reduction [24, 43] explicitly show the error of these methods.

The visual analysis is heavily based on including the human-in-the-loop for exploring the uncertainty from the ensembles of multivariate time series (stage ⑦). More information can be found in the paper by Brich et al. [9].

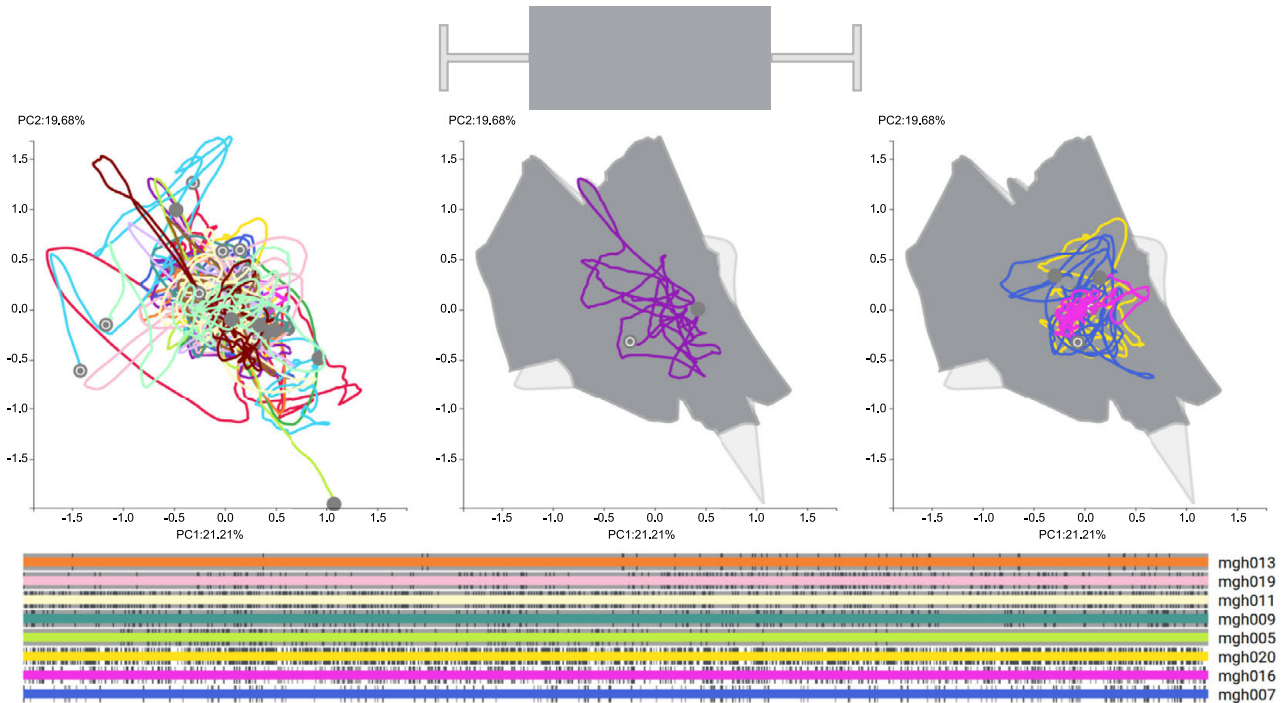


Figure 3: Ensemble of 16 patients from the MGH/MF data set shown by Time Curves (top) and scarf plot (bottom). The Time Curve plots in the center and to the right use the boxplot metaphor to depict the area corresponding to the 50% quantile (dark gray) and 80% quantile (light gray). A regular boxplot glyph with box and whiskers is shown at the top for reference. © 2022 Computer Graphics Forum published by Eurographics – The European Association for Computer Graphics and John Wiley & Sons Ltd. Reprinted, with permission, from Brich et al. [9].

5 Stochastic partial differential equations

Awareness of uncertainty is particularly important when modeling physical phenomena mathematically by partial differential equations (PDEs). In this example, the uncertainty comes from the real-world phenomenon (stage ①) and the corresponding simulation that solves the PDEs (stage ②). Therefore, our discussion focuses on the early stages of the visualization pipeline.

Measuring or looking up data at several distinct spatial points yields information about these points, but may leave uncertain what happens in between these points. These uncertainties are modeled by a stochastic source term or a random coefficient. Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a complete probability space, where Ω denotes the non-empty set of elementary events, \mathcal{A} the σ -algebra of measurable events, and \mathbb{P} the probability measure. On $\mathcal{D} := [0, 1]^2$, the unit square as the spatial domain, a simplified model for subsurface flow through a porous medium with uncertain discontinuous porosity is given by the random elliptic PDE, cf. [3] and the references therein:

$$-\nabla \cdot (a(\omega, x) \nabla u(\omega, x)) = f(x) \text{ for } \omega \in \Omega \text{ and } x \in \mathcal{D}.$$

The solution to this equation is denoted by $u : \Omega \times \mathcal{D} \rightarrow \mathbb{R}$, the random jump-diffusion coefficient by $a : \Omega \times \mathcal{D} \rightarrow \mathbb{R}_+$, and the source term by $f : \mathcal{D} \rightarrow \mathbb{R}$. Together with adequate boundary conditions this is a well-posed mathematical problem admitting a unique weak solution u . The coefficient a models the spatially dependent diffusivity of the porous medium. High values of a correspond to low diffusivity areas and low values of a correspond to high diffusivity areas of the porous medium. The magnitude of these values as well as the position of the diffusivity areas are practically measured for a real-world simulation and are prone to the uncertainties described above. Fig. 4 shows examples for three different coefficient samples $\omega \in \Omega$.

For each fixed $\omega \in \Omega$, when computing a numerical approximation u_h to the solution u , the accuracy of the spatial approximation indexed by h depends strongly on how the uncertain areas are resolved. The areas where the jump from low to high diffusivity occurs lead to steep solution curves, i. e., high solution gradients, and thus need high spatial resolution to be approximated accurately. Fig. 5 illustrates an approach that uses high resolution across the whole domain, however, at the cost of high computation effort. To circumvent this we may start on a coarse initial resolution and refine adaptively only where needed




Figure 4: Different samples of the random coefficient with varying jump magnitudes and positions demonstrating the uncertainty of the problem.




Figure 5: Standard uniform mesh with evenly sized elements (left) and corresponding spatial approximation to the PDE solution (right).

as illustrated in Fig. 6. A-posteriori error estimation techniques, cf. [20], allow us to detect areas in which the approximation error is high and to focus only on them. Even though this approach is more involved it may reduce the overall computational cost to compute a satisfactory numerical approximation u_h to the solution u .

With stochastic PDEs and their numerical solvers, we have a method that allows us to model physical phenomena with uncertainty, such as subsurface flow, and compute an efficient representation. This information can then be processed further along the visualization pipeline. The example of Fig. 6 provides an implicit visualization of the uncertain coefficient via the adaptivity in the mesh: higher resolution in the mesh corresponds to lower numerical error (comparable to the uniform resolution of Fig. 5). However, this visual mapping (stage ⑤) provides only some aspects of uncertainty. A direct visual representation of the full stochastic solutions is infeasible because there is not enough visual space available. Therefore, future work could address appropriate data reduction (stage ④) and corresponding visual mapping (stage ⑤) for stochastic PDEs.




Figure 6: Successively refined meshes generated adaptively w. r. t. to the uncertainty in magnitude and position of the diffusivity areas given by the coefficient (upper left, right, and lower left). Corresponding spatial approximation to the PDE solution (lower, right).

6 Linguistic annotation

Annotated natural language data plays a critical role in theoretical and computational linguistic research by facilitating empirical investigations on specific phenomena and serving as a basis for the development of Natural Language Processing (NLP) models. In recent years, uncertainty has been acknowledged as a problem for linguistic annotation [4, 8, 12], although only relatively few concrete solutions have been proposed [2, 13, 29, 34, 37]. Uncertainty visualization has much to offer in this area, particularly if uncertainties can be resolved as part of an ongoing statistical machine translation [35].

With respect to linguistic data, uncertainty is already present at stage ① of the visualization pipeline, since ambiguity and underspecification is an inherent part of language structure, with ambiguities found at both the token (word) level and higher phrasal levels. Languages also use systems of oppositions, whereby a marked form is contrasted with an unmarked one to signal a semantic effect, as for example in Differential Object Marking, shown below in Example (1) for Marathi. Here, the unmarked (nominative) object is unspecified for definiteness/specificity via the absence of marking. In the presence of overt (accusative) marking, the only available interpretation is that there is a particular elephant that was killed.

- (1) a. *siḡhā=ne hattī mār-lā*
 lion_{M=ERG} elephant._{M.NOM} kill-PERF.3.M.SG
 ‘The lion killed a/the elephant.’
 b. *siḡhā=ne hattī=lā mār-le*
 lion._{M=ERG} elephant._{M=ACC} kill-PERF.3.N.SG
 ‘The lion killed the elephant.’

In the absence of knowledge about this kind of structural markedness in the data, uncertainty about the analysis and interpretational properties can arise. This is particularly true for work with historical corpora, where data are limited and no new data can be adduced.

A common type of linguistic ambiguity is *class ambiguity*, referring to cases where a token allows for more than one classification. One particular example that has been discussed in relation to the annotation of historical linguistic data is the status of clauses introduced by the item *wente* (‘because/that/but’) in Middle Low German (c. 1200–1650 CE) texts, as in the following Example (2) [8].

- (2) *vn̄de ik sach et . vn̄de betugede et . [wente dit is
 and I saw it and attested it WENTE this is
 godes sone]
 god.GEN son
 ‘and I saw it and attested it because/that/but this is
 god’s son’ (Buxteh. Ev., [23])*

Without context, the clause introduced by *wente* can in principle be a main or a subordinate clause, since during this language stage cues such as sentential punctuation, capitalization, and word order do not systematically distinguish between the two. This is just one example of how uncertainty is particularly pertinent when dealing with (typically non-standardized) historical linguistic data, where annotators cannot rely on native-speaker judgments.

Another common source of uncertainty is *boundary ambiguity* at the phrasal level, where linguistic material can be segmented into smaller units in multiple ways, resulting in alternative boundary divisions. Information on how to resolve the ambiguity may not be present in the immediately available data, or may be resolved in a potentially biased manner (e. g., by using learned information on probabilities from a given corpus). For example, (3) illustrates a scope ambiguity that results in two potential readings. Without more context, the ambiguity cannot be resolved, but a language model that has learned likelihood of associations from data/texts that associate African-Americans with criminality in a biased way will always resolve the ambiguity in favor of Reading 2.

- (3) African-American criminals and lawyers presented themselves at court.
 Reading 1: Both the criminals and the lawyers are African-American.
 Reading 2: Only the criminals are African-American.

The occurrence of ambiguity in natural language can thus also lead to a perpetuation of existing biases in data. Making a differentiation and concomitant visualization of ambiguities and the context in which they are (un)resolvable, is an important factor in the on-going discussions on bias in artificial intelligence [11, 30, 6, 39, 18].

During the annotation process, i. e., at the data processing and modeling stage ③, uncertainty arising from linguistic ambiguity is currently usually treated in one of three ways, none of which is fully adequate: (i) stochastic treatment, (ii) assignment of an ‘other’/‘miscellaneous’ label, (iii) left unannotated. Under a stochastic treatment (cf. Most Frequent Class Baseline [28]), the interpretation that is the most likely is chosen, such that the linguistic material in question is assigned a single interpretation and the ambiguity is lost altogether (cf. the scope ambiguity in Example (3)). Under the second approach, the ambiguous material is assigned a special label, as in the case of *wente*-clauses, cf. Example (2), which are annotated with a unique tag in the Corpus of Historical Low German [8]. This goes some way to capturing ambiguity/uncertainty but does not adequately represent the particular type or source of the issue. A third and common approach is to leave ambiguous/uncertain material unannotated, ultimately leading to data loss. Crucially, under all three approaches the nature and source of the ambiguity/uncertainty is not adequately represented.

More sophisticated modeling and in turn visualization of uncertainty in language data via the proposed visualization pipeline has the potential to play an important role in furthering our understanding of natural language and improving language technologies. As mentioned, ambiguity is an inherent property of all natural languages and thus worthy of study in its own right, particularly as, compared to other inherent properties such as variation and change, it has been less exhaustively researched. Moreover, ambiguity is generally acknowledged to play an important role in language change (cf. ‘reanalysis’, [15, 16]), although the precise mechanisms involved remain unclear in the absence of methodologies which explicitly harness ambiguity as a nuanced and multifaceted phenomenon. An interactive system that allows a user to identify and flag instances of unresolved ambiguity and uncertainty as such, as well as to resolve it interactively for purposes of experimentation or as more data becomes available via the vi-

sualization pipeline would serve to further the state of the art considerably.

7 Conclusion

We have given a brief introduction to uncertainty visualization, discussed sources of uncertainty throughout the visualization pipeline, and presented visualization examples and applications. Typical use cases include some kind of probabilistic modeling to quantitatively describe uncertainty, which is critical for propagating and eventually visualizing uncertainty. Our examples come from a wide range of different applications, demonstrating the utility of uncertainty visualization. The examples also show that uncertainty may have different importance at the different stages of the visualization pipeline: some emphasized the earlier stages of data acquisition and modeling, others later stages of filtering, visual mapping, and interaction. Therefore, there is not a single solution that would be able to cover a wide range of uncertainty-related problems. Instead, there is need to take a broader perspective at the visualization problem and select and prioritize an appropriate combination of techniques. Overall, we hope that our paper has triggered awareness for the high relevance of uncertainty in visualization for data analysis and communication.

We see several directions of future research in the field. For example, there is a lack of enough practical software implementations and software ecosystems for widespread application of uncertainty visualization. Similarly, more research is needed to cover further examples of data types and visual analysis with uncertainty propagation. Furthermore, we still need a better understanding of the cognition of uncertainty and improved techniques for evaluating uncertainty visualizations.

Funding: This work was funded by the *Deutsche Forschungsgemeinschaft* (DFG, German Research Foundation)—Project ID 251654672—TRR 161 (Projects A01, B07, and D02). Hannah Booth is funded by a postdoctoral fellowship from the *Research Foundation – Flanders* (FWO), Grant no. 12ZL522N.

References

1. B. Bach, C. Shi, N. Heulot, T. Madhyastha, T. Grabowski, and P. Dragicevic. Time Curves: Folding time to visualize patterns of temporal evolution in data. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):559–568, 2016.
2. F. Barteld, S. Ihden, I. Schröder, and H. Zinsmeister. Annotating descriptively incomplete language phenomena. In *Proceedings of LAW VIII - The 8th Linguistic Annotation Workshop*, pages 99–104, Dublin, Ireland, August 2014. Association for Computational Linguistics and Dublin City University.
3. A. Barth and A. Stein. A study of elliptic partial differential equations with jump diffusion coefficients. *SIAM/ASA Journal on Uncertainty Quantification*, 6(4):1707–1743, 2018.
4. C. Beck, H. Booth, M. El-Assady, and M. Butt. Representation problems in linguistic annotations: Ambiguity, variation, uncertainty, error and bias. In *Proceedings of the 14th Linguistic Annotation Workshop*, pages 60–73, Barcelona, Spain, December 2020. Association for Computational Linguistics.
5. S. Belia, F. Fidler, J. Williams, and G. Cumming. Researchers misunderstand confidence intervals and standard error bars. *Psychological Methods*, 10(4):389–396, 2005.
6. E. M. Bender and B. Friedman. Data statements for natural language processing: Toward mitigating system bias and enabling better science. *Transactions of the Association for Computational Linguistics*, 6:587–604, 2018.
7. G.-P. Bonneau, H.-C. Hege, C. R. Johnson, M. M. Oliveira, K. Potter, P. Rheingans, and T. Schultz. Overview and state-of-the-art of uncertainty visualization. In C. D. Hansen, M. Chen, C. R. Johnson, A. E. Kaufman, and H. Hagen, editors, *Scientific Visualization*, pages 3–27. Springer, 2014.
8. H. Booth, A. Breitbarth, A. Ecay, and M. Farasyn. A Penn-style Treebank of Middle Low German. In *Proceedings of The 12th Language Resources and Evaluation Conference*, pages 766–775, Marseille, France, May 2020. European Language Resources Association.
9. N. Brich, C. Schulz, J. Peter, W. Klingert, M. Schenk, D. Weiskopf, and M. Krone. Visual analytics of multivariate intensive care time series data. *Computer Graphics Forum*, 2022. <https://doi.org/10.1111/cgf.14498>.
10. K. Brodli, R. S. Allendes Osorio, and A. Lopes. A review of uncertainty in data visualization. In J. Dill, R. A. Earnshaw, D. J. Kasik, J. A. Vince, and P. Chung Wong, editors, *Expanding the Frontiers of Visual Analytics and Visualization*, pages 81–109. Springer, 2012.
11. A. Caliskan, J. J. Bryson, and A. Narayanan. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186, April 2017.
12. T. Cassidy, B. McDowell, N. Chambers, and S. Bethard. An annotation framework for dense event ordering. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 501–506, Baltimore, Maryland, June 2014. Association for Computational Linguistics.
13. T. Chen, Z. Jiang, A. Poliak, K. Sakaguchi, and B. Van Durme. Uncertain natural language inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8772–8779, Online, July 2020. Association for Computational Linguistics.
14. E. H. Chi. A taxonomy of visualization techniques using the data state reference model. In *Proceedings of the IEEE Symposium on Information Visualization*, pages 69–75. *IEEE Computer Society*, 2000.
15. H. de Smet. Analysing reanalysis. *Lingua*, 119(11):1728–1755, 2009.

16. H. De Smet and M.-A. Markey. The spark or the fuel? on the role of ambiguity in language change. *Journal of Historical Syntax*, 5(36):1–24, 2021.
17. Carolin Fiedler, Willy Scheibel, Daniel Limberger, Matthias Trapp, and Jürgen Döllner. Survey on user studies on the effectiveness of treemaps. In *Proceedings of the 13th International Symposium on Visual Information Communication and Interaction, VINCI'20*. Association for Computing Machinery, 2020. Article 2.
18. T. Gebru. Race and gender. In M. D. Dubber, F. Pasquale, and S. Das, editors, *The Oxford Handbook of Ethics of AI*. Oxford University Press, Oxford, 2020.
19. J. Görtler, C. Schulz, D. Weiskopf, and O Deussen. Bubble treemaps for uncertainty visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):719–728, 2018.
20. T. Grätsch and K.-J. Bathe. A posteriori error estimation techniques in practical finite element analysis. *Computers & Structures*, 83(4-5):235–265, 2005.
21. M. Greis, H. Schuff, M. Kleiner, N. Henze, and A. Schmidt. Input controls for entering uncertain data: Probability distribution sliders. *Proceedings of the ACM on Human-Computer Interaction*, 1(EICS), June 2017. Article 3.
22. R. B. Haber and D. A. McNabb. Visualization idioms: A conceptual model for visualization systems. In G. M. Nielson, B. D. Shriver, and L. J. Rosenblum, editors, *Visualization in Scientific Computing*, pages 74–93. IEEE Computer Society Press, 1990.
23. J. E. Härd. Syntax des Mittelniederdeutschen. In Werner Besch, Anne Betten, Oskar Reichmann, and Stefan Sonderegger, editors, *Sprachgeschichte: Ein Handbuch zur Geschichte der deutschen Sprache und ihrer Erforschung*, volume 2, pages 1456–1463. de Gruyter, 2000.
24. N. Heulot, M. Aupetit, and J.-D. Fekete. ProxiLens: Interactive Exploration of High-Dimensional Data using Projections. In M. Aupetit and L. van der Maaten, editors, *EuroVis Workshop on Visual Analytics using Multidimensional Projections*. The Eurographics Association, 2013.
25. M. Höferlin, B. Höferlin, D. Weiskopf, and G. Heidemann. Uncertainty-aware video visual analytics of tracked moving objects. *Journal of Spatial Information Science*, 2(1):87–117, 2011.
26. J. Hullman, X. Qiao, M. Correll, A. Kale, and M. Kay. In pursuit of error: A survey of uncertainty visualization evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):903–913, 2019.
27. J. Hullman, P. Resnick, and E. Adar. Hypothetical outcome plots outperform error bars and violin plots for inferences about reliability of variable ordering. *PLOS ONE*, 10(11):1–25, November 2015.
28. D. Jurafsky and J. H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Speech Recognition, and Computational Linguistics*. Prentice-Hall, 2nd edition, 2009.
29. D. Jurgens. Embracing ambiguity: A comparison of annotation methodologies for crowdsourcing word sense labels. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 556–562, Atlanta, Georgia, June 2013. Association for Computational Linguistics.
30. Svetlana Kiritchenko and Saif Mohammad. Examining gender and race bias in two hundred sentiment analysis systems. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 43–53, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
31. S. Liu, D. Maljovec, B. Wang, P.-T. Bremer, and V. Pascucci. Visualizing high-dimensional data: Advances in the past decade. *IEEE Transactions on Visualization and Computer Graphics*, 23(3):1249–1268, 2017.
32. S. López-Pintado and J. Romo. On the concept of depth for functional data. *Journal of the American Statistical Association*, 104(486):718–734, 2009.
33. A. M. MacEachren, R. E. Roth, J. O'Brien, B. Li, D. Swingley, and M. Gahegan. Visual semiotics amp; uncertainty visualization: An empirical study. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2496–2505, 2012.
34. M.-L. Merten and N. Seemann. Analyzing constructional change: Linguistic annotation and sources of uncertainty. In *Proceedings of the Sixth International Conference on Technological Ecosystems for Enhancing Multiculturality, TEEM'18*, pages 819–825. Association for Computing Machinery, 2018.
35. T. Munz, D. Väh, P. Kuznecov, N. T. Vu, and D. Weiskopf. Visual-interactive neural machine translation. In *Proceedings of Graphics Interface 2021*, 2021. Article 30.
36. A. Pang, C. M. Wittenbrink, and S. K. Lodha. Approaches to uncertainty visualization. *The Visual Computer*, 13(8):370–390, 1997.
37. B. Plank, D. Hovy, and A. Søgaard. Learning part-of-speech taggers with inter-annotator agreement loss. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 742–751, Gothenburg, Sweden, April 2014. Association for Computational Linguistics.
38. D. Sacha, H. Senaratne, B. Chul Kwon, G. Ellis, and D. A. Keim. The role of uncertainty, awareness, and trust in visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):240–249, 2016.
39. P. Schramowski, C. Turan, S. Jentzsch, C. Rothkopf, and K. Kersting. The moral choice machine. *Frontiers in Artificial Intelligence*, 3:36, 2020.
40. H. J. Schulz, S. Hadlak, and H. Schumann. The design space of implicit hierarchy visualization: A survey. *IEEE Transactions on Visualization and Computer Graphics*, 17(4):393–411, 2011.
41. M. Skeels, B. Lee, G. Smith, and G. Robertson. Revealing uncertainty for information visualization. *Information Visualization*, 9(1):70–81, 2010.
42. M. Sondag, W. Meulemans, C. Schulz, K. Verbeek, D. Weiskopf, and B. Speckmann. Uncertainty treemaps. In *IEEE Pacific Visualization Symposium, PacificVis 2020*, pages 111–120, 2020.
43. Yong Wang, Qiaomu Shen, Daniel Archambault, Zhiguang Zhou, Min Zhu, Sixiao Yang, and Huamin Qu. AmbiguityVis: Visualization of ambiguity in graph layouts. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):359–368, 2016.
44. D. Weiskopf. Uncertainty visualization: Concepts, methods, and applications in biological data visualization. *Frontiers in Bioinformatics*, 2, 2022.
45. J Welch, P Ford, R Teplick, and R Rubsamen. The Massachusetts General Hospital-Marquette Foundation hemodynamic and electrocardiographic database—Comprehensive collection of

critical care waveforms. *Clinical Monitoring*, 7(1):96–97, 1991.

46. Y. Wu, Guo-Xun Yuan, and Kwan-Liu Ma. Visualizing flow of uncertainty through analytical processes. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2526–2535, 2012.

Bionotes



David Hägele
University of Stuttgart, Visualization
Research Center, D-70569, Stuttgart,
Germany
david.haegle@visus.uni-stuttgart.de

David Hägele is a doctoral researcher at the Visualization Research Center (VISUS) of the University of Stuttgart. He received a master's degree in computer science from the University of Stuttgart. His research topics revolve around quantification in visualization, specifically with respect to uncertainty.



Dr. Christoph Schulz
University of Stuttgart, Visualization
Research Center, D-70569, Stuttgart,
Germany
christoph.schulz@visus.uni-stuttgart.de

Dr. Christoph Schulz is a postdoctoral researcher at the Visualization Research Center (VISUS) of the University of Stuttgart, from where he also graduated. His areas of interests include visualization of uncertainty, computer graphics, and software engineering.



Cedric Beschle
University of Stuttgart, Institute of Applied
Analysis and Numerical Simulation,
D-70569, Stuttgart, Germany
cedric.beschle@ians.uni-stuttgart.de

Cedric Beschle is a doctoral researcher at the University of Stuttgart, where he is part of the working group “Computational Methods for Uncertainty Quantification” under the supervision of Prof. Dr. Andrea Barth at the Institute of Applied Analysis and Numerical Simulation (IANS). He graduated with a master's degree in mathematics at Eberhard Karls University of Tübingen.



Dr. Hannah Booth
Ghent University, Department of
Linguistics, B-9000, Gent, Belgium
hannah.booth@ugent.be

Dr. Hannah Booth is an FWO Postdoctoral Research Fellow at Ghent University. She received her PhD in Linguistics from the University of Manchester, after which she was employed in the group of Prof. Miriam Butt at the University of Konstanz. Her research focusses on language change at the syntax-information structure interface, combining theoretical perspectives with corpus-based approaches.



Prof. Dr. Miriam Butt
University of Konstanz, Department of
Linguistics, D-78457, Konstanz, Germany
miriam.butt@uni-konstanz.de

Prof. Dr. Miriam Butt is Professor for General and Computational Linguistics at the University of Konstanz. She received her PhD from Stanford University in Linguistics. Her research interests include formal linguistics, particularly with respect to morphology and syntax, historical linguistics, grammar engineering and visualization for linguistics (LingVis).



Prof. Dr. Andrea Barth
University of Stuttgart, Institute of Applied
Analysis and Numerical Simulation,
D-70569, Stuttgart, Germany
andrea.barth@ians.uni-stuttgart.de

Prof. Dr. Andrea Barth is a Professor at the University of Stuttgart and graduated at the University of Oslo. She leads the research group “Computational Methods for Uncertainty Quantification” at the Institute of Applied Analysis and Numerical Simulation (IANS). Her research areas include Stochastic Analysis, Numerical Analysis and the quantification of uncertainty in complex systems.



Prof. Dr. Oliver Deussen
University of Konstanz, Department of
Computer and Information Science,
D-78457, Konstanz, Germany
oliver.deussen@uni-konstanz.de

Prof. Dr. Oliver Deussen is a Professor at University of Konstanz and graduated at Karlsruhe Institute of Technology (KIT). He served as visiting professor at the Chinese Academy of Science in Shenzhen (SIAT), was President of the Eurographics Association and is speaker of the Excellence Cluster “Centre for the Advanced Study of Collective Behaviour”. His areas of interest encompass Information Visualization, modeling and rendering of complex systems as well as non-photorealistic rendering.



Photo by Max Kovalenko

Prof. Dr. Daniel Weiskopf
University of Stuttgart, Visualization
Research Center, D-70569, Stuttgart,
Germany
daniel.weiskopf@visus.uni-stuttgart.de

Prof. Dr. Daniel Weiskopf is a Professor at the Visualization Research Center (VISUS) of the University of Stuttgart. He received his Dr. rer. nat. degree in Physics from the University of Tübingen, and the Habilitation degree in Computer Science from the University of Stuttgart. His research interests include visualization, visual analytics, eye tracking, human-computer interaction, computer graphics, and special and general relativity.