

MultiSegVA: Using Visual Analytics to Segment Biologging Time Series on Multiple Scales

Philipp Meschenmoser, Juri F. Buchmüller, Daniel Seebacher, Martin Wikelski, and Daniel A. Keim.

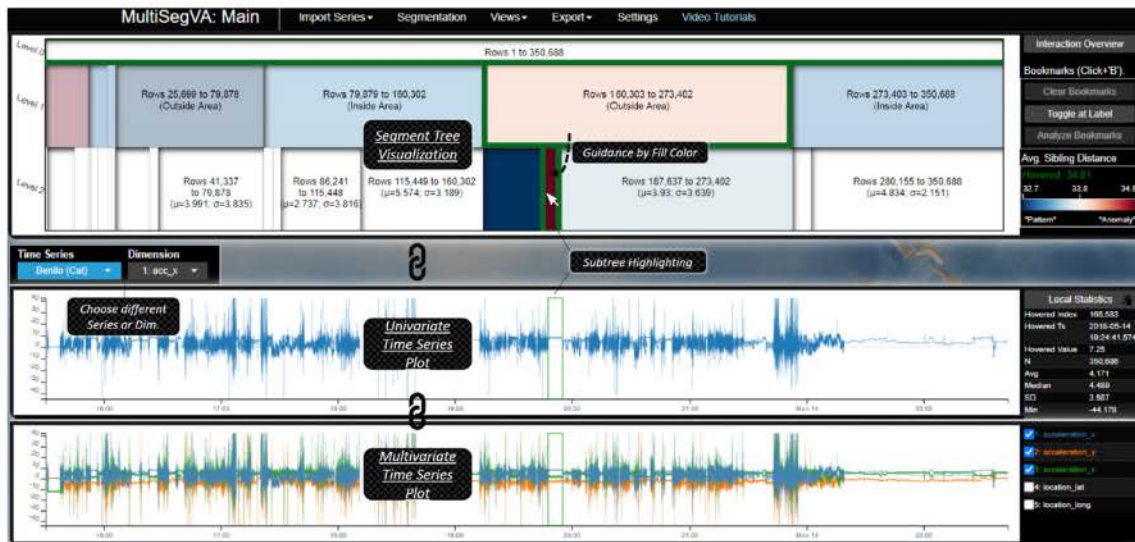


Fig. 1. The main window of our MultiSegVA platform for the visual-interactive multi-scale segmentation of biologging time series. After arranging scale-wise segmentation techniques by a new visual query language, the analyst can explore results in tailored visualizations and refine parameters. Here the tree relies on segmentation by geographical area, then recursively by acceleration change points.

Abstract—Segmenting biologging time series of animals on multiple temporal scales is an essential step that requires complex techniques with careful parameterization and possibly cross-domain expertise. Yet, there is a lack of visual-interactive tools that strongly support such multi-scale segmentation. To close this gap, we present our MultiSegVA platform for interactively defining segmentation techniques and parameters on multiple temporal scales. MultiSegVA primarily contributes tailored, visual-interactive means and visual analytics paradigms for segmenting unlabeled time series on multiple scales. Further, to flexibly compose the multi-scale segmentation, the platform contributes a new visual query language that links a variety of segmentation techniques. To illustrate our approach, we present a domain-oriented set of segmentation techniques derived in collaboration with movement ecologists. We demonstrate the applicability and usefulness of MultiSegVA in two real-world use cases from movement ecology, related to behavior analysis after environment-aware segmentation, and after progressive clustering. Expert feedback from movement ecologists shows the effectiveness of tailored visual-interactive means and visual analytics paradigms at segmenting multi-scale data, enabling them to perform semantically meaningful analyses. A third use case demonstrates that MultiSegVA is generalizable to other domains.

Index Terms—Visual analytics, time series segmentation, multi-scale analyses, movement ecology

1 INTRODUCTION

Time series often include patterns and semantics on very different temporal scales, but a majority of segmentation techniques is applied on a single scale with global parameters. More complex multi-scale techniques commonly imply careful parameterization and possibly the need for cross-domain expertise, yet there are no visual-interactive tools with strong support for segmenting time series on multiple scales.

- Philipp Meschenmoser, Juri F. Buchmüller, Daniel Seebacher, and Daniel A. Keim are with the University of Konstanz, Germany. E-mail: {philipp.meschenmoser, juri.buchmueller, daniel.seebacher, keim}@uni-konstanz.de
- Martin Wikelski is with the Max Planck Institute of Animal Behavior, Germany. E-mail: wikelski@ab.mpg.de

We present the MultiSegVA platform that facilitates multi-scale time series segmentation by tailored visual-interactive features, established VA principles, and a new visual query language. In this publication, we focus on biologging time series of moving animals: these time series have prototypical multi-scale character and include widely unexplored behaviors, which are hidden in high resolutions and cardinalities. Additionally, biologging-driven movement ecology is an emerging field [10, 59, 76, 77], triggered by technical advances that enable academia to address open questions in innovative ways. The biologging time series stem from miniaturized tags and give high-resolution information about, e.g., an animal's location, tri-axial acceleration, and heart rate. Here, semantics are typically distributed on diverse temporal scales, including life stages, seasons, days, day times, and (micro)movement frames. These temporal scales are complemented by spatial scales concerning, e.g., the overall migration range, migration stops, and foraging ranges. There are complex scale- and context-specific conditions [6, 44, 65], implying different energy expenditures, driving factors, and decisions for behavior. Hence, segmenting such

time series on a single scale with global parameters does not sufficiently address their multi-scale character.

The relevance of multi-scale segmentation can be further motivated by three reasons. First, analysts can deepen their understanding of how scales relate to each other: e.g., in terms of nesting relations, next to relative scale sizes and types. A multi-scale perspective can even enable “to gain an insight on an entire knowledge domain or a relevant sub-part” [51]. Second, even without labeled data or thoroughly parameterized single-scale techniques, it is possible to identify fine-grained patterns that are wrapped by lower-scale, context-yielding patterns. Such fine-grained and context-aware patterns are crucial to enrich existing classification and prediction models. Third, demands for more multi-scale analyses originate from domain literature. Such demands can be found in movement ecology and analysis [4, 20, 36, 44], but also in, e.g., medical sciences [1] and social sciences [15].

However, in practice, segmenting time series on multiple scales is often impeded by several factors. First, multi-scale techniques rely on more in-depth, theoretical foundations and inherent parameters that need to be carefully adapted. Therefore analysts (e.g., movement ecologists) might require cross-domain expertise in statistical multi-scale time series analysis. Second, even with such expertise, it is difficult to decide on scale properties (e.g., size, dimension, number of scales) and further parameters. Third, we observe a lack of suitable visual-interactive approaches in related works (Section 3.2) that could strongly support and promote segmenting time series on multiple scales.

To close this gap, we present our web-based *MultiSegVA* platform that allows analysts to visually explore and refine a multi-scale segmentation, which results from a simple way of setting segmentation techniques on multiple scales. In the context of multi-scale segmentation, *MultiSegVA* primarily contributes the use of tailored visual-interactive features and established VA paradigms (contribution C1). To flexibly configure segmentation techniques and parameters, *MultiSegVA* includes a new visual query language (VQL, C2) that links a variety of segmentation techniques across multiple scales. These techniques stem in the present case from a set (C3) that was derived together with movement ecologists and covers typical domain use cases.

Section 2 outlines typical tasks and requirements from movement ecology that are contextualized and backed up by the related work Section 3. Then Section 4 shows how we address the first requirements, namely by describing and justifying the platform design. Further requirements are fulfilled by the VQL and its specification in Section 5. Three real-world use cases and domain expert feedback from movement ecologists compose Section 6, before discussing several aspects.

In the following, we assume single, discrete, and unlabeled time series, denoted as pair (t, v) and with n records. Let the array $t = [t[1], \dots, t[n]]$ specify timestamps $\subset \mathbb{R}$ with $\forall i \in \{2, \dots, n\} : t[i-1] < t[i]$. Then $v[i]$ is the (multi-dimensional) value record for timestamp $t[i]$. In general, segmentation shall be the problem of finding a set of index intervals $[\text{from}, \text{to}] \subset \{1, \dots, n\}$, such that each index $i \in \{1, \dots, n\}$ is included in exactly one of these index intervals.

2 DOMAIN BACKGROUND, TASKS, AND REQUIREMENTS

The *MultiSegVA* platform was developed in close collaboration with domain experts from movement ecology who are working with biologging time series on a day-to-day basis.

Our domain experts are researchers from the Max Planck Institute of Animal Behavior that is one of the major driving forces for the biologging-based, in-depth understanding of animal movements. Hence, the institute coordinates Movebank [77], an online repository for biologging time series, and the ICARUS initiative [76], where an ISS antenna gathers signals of global animal movements from space. Likewise, recent advances in biologging technology are opening up a new era in understanding the dynamics and specifics of animal movements. Biologging tags “have and will become smaller, cheaper and more accessible, new satellite tracking technologies are introduced, data download methodologies become more efficient, battery life increases” [20]. The domain experts work with resulting time series with GPS and acceleration dimensions, for contributing to subdomains such as animal migration, animal societies, or computational ecology.

Such contributions help in better understanding some of today’s most urgent issues: e.g., climate change [53, 58], species decline (cf. [61, 64]), disease transmission [32]. Thus the movement ecologists are highly experienced in data acquisition, ecological modeling, and corresponding hypothesis testing, while multi-scale time series statistics and advanced machine learning are often less focal points. By interviewing our experts, we could extract the following tasks and related challenges.

Tasks. In a typical workflow, our experts mostly follow the KDD pipeline [24]: they begin with sampling, cleaning, standardizing sensor data, before deriving new dimensions (e.g., pitch, ODBA [30]) and plotting GPS as well as acceleration dimensions. The workflow continues by the main tasks of visual exploration (T1), segmentation (T2), and annotation with behavior labels (T3). For these tasks, the experts often sequentially apply R and Python scripts, whose adaptation and chaining are tedious. Also, their tools in use do not consider the multi-scale nature of a biologging time series, and they often lack in visual-interactive features, automated techniques, and scalability. Consequently, the existing analysis means cannot sufficiently exploit the true potential and information hidden in the enormous amounts of accruing data [76, 77]. For instance, the experts use a tool that shows tri-axial acceleration data in a zoomable and multivariate time series plot; see the schematic in Figure 2. The experts interact with the plot, consider contextual data sources, and often try to project themselves into the animal with its motion and movements. To identify and annotate segments of the time series, the experts have to follow a manual process: segments have to be defined by a rectangle selection, before predefined behavior labels are individually assigned. At millions of records, this process can take, as reported, several months. To handle T1–T3 while strongly reflecting the multi-scale nature of the time series, the following requirements evolved.

Requirements. Through regular meetings with the movement ecologists, we derived requirements for the *MultiSegVA* platform by means of semi-structured interviews. Initially, we proposed a set of requirements, based on literature research and supported by interface sketches. These requirements were discussed and refined together with the domain experts. In the following, we describe a high-level set of requirements as result of our discussions. This set complements the requirements in [7] with the aim of satisfying multi-scale aspects.

R1 Integrating Approach. Insights in time series analysis are typically hidden on different temporal scales: e.g., daily foraging versus bi-annual migration. At the same time, each scale comes with scale-specific segmentation objectives, data distributions, patterns, and semantics. To handle these aspects, the experts now tediously combine techniques from different scripts and tools. For improving the situation, an integrating approach that incorporates various techniques into one scheme is needed.

R2 Multi-Scale Structure & Specifics. A temporal scale is typically surrounded by larger temporal scales that have context-yielding patterns, events, anomalies, and semantics. For segmentation, the analyst requires a mean to cover the specifics of surrounding temporal scales and to extract the multi-scale structures of a time series. Current tools in use support only one temporal scale at a time.

R3 Flexible Parameterization. Time series can be based on several possible multi-scale structures and often the analyst has to explore which structure is most adequate to answer a specific question. Using scripts for this exploration tends to be inflexible and error-prone as code has to be re-arranged. Hence, the analyst requires means for the flexible and effortless parameterization of a segmentation task.

R4 Visual-Interactive Features. To facilitate configuration and reasoning, the analyst needs extensive visual-interactive features that clearly go beyond colored line plots. The methods shall strongly support result exploration, parameter refinement, and segment annotation.




Fig. 2. Manually defining and labeling segments in tri-axial acceleration data.

3 RELATED WORK

This section backs up the elicited requirements, contextualizes, and further motivates our work. We initially overview techniques for time series segmentation, before depicting VA approaches for a visual-interactive segmentation process and outlining related VQLs.

3.1 Techniques for Time Series Segmentation

Many segmentation techniques are applied on a single scale with global parameters and return a simple segment alignment. Keogh et al. present a survey [40] on geometric segmentation: with bottom-up, top-down, sliding-window algorithms, next to the efficient hybrid *SWAB*, that try to meet segmentation criteria concerning piecewise linear approximations.

Also change point detection is often applied to split a time series into homogeneous units, see the surveys in [3, 66]. The detection is usually steered by a cost function (e.g., maximum-likelihood-based), a search function (e.g., optimal PELT), and additional constraints. With usually more complex foundations in probability theory, Hidden Markov Models (HMMs, [14, 55, 74]), Hidden Semi-Markov Models, and State Space Models (SSMs) are options that tend to capture inherent semantics better than previous techniques.

Next to these generic techniques, there are more domain- or data-specific variants and developments. For example, techniques specifically for (time-annotated) trajectories are depicted in [20] and are part of the framework [11, 12] by Buchin et al. Additionally, techniques with relatively novel ways of functioning have recently emerged, e.g., semantic segmentation based on matrix profiles [29].

Though, a common aspect of the above techniques is that their default variants do not emphasize the multi-scale character of time series, as they are often applied on a single temporal scale with global parameters. Such an emphasis can be obtained by more complex variants, e.g., by hierarchical HMMs, multi-level SSMs, or multi-scale change point detection [21, 27]. So, [18] depicts a non-parametric Wavelet model to identify change points in periodograms of multiple scales. In general, alternative approaches can emerge from the field of multi-scale time series modeling [25, 26]. Hence to identify multi-scale structures, Vámos proposes a technique [67] based on a time series' monotony spectrum: i.e., "the variation of the mean amplitude of the monotonic segments with respect to the mean local time scale during successive averagings of the time series" [67]. Such techniques have strong theoretical foundations and their overall value is certainly unquestionable. Yet, their increased complexity demands stronger efforts at parameterization and possibly cross-domain expertise. With these impeding factors, the next subsection regards to which extent visual-interactive approaches could support the analyst at multi-scale segmentation.

3.2 VA Approaches for Time Series Segmentation

	R1	R2	R3	R4	Heterogeneous Data	Generalizability	Geo-Context
MultiSegVA	•	•	•	•	•	•	•
Alsallakh et al. 2014, [2]	•	•	•	•	•	•	•
Röhlig et al. 2015, [57]	•	•	•	•	•	•	•
Spretke et al. 2011, [62]	•	•	•	•	•	•	•
Gao et al. 2013, [28]	•	•	•	•	•	•	•
Zhao et al. 2011, [72]	•	•	•	•	•	•	•
Walker et al. 2015, [68]	•	•	•	•	•	•	•
Bernard et al. 2016, [7]	•	•	•	•	•	•	•

Table 1. Comparison of related approaches. No related approach fulfills all of the necessary requirements **R1–R4**.

The following VA approaches shall enable interactive time series segmentation, where "humans and machines cooperate using their respective distinct capabilities for the most effective results" [38]. Several of these approaches are intended for movement and motion analyses.

Spretke et al. present their VA tool *AnimalEcologyExplorer* [62], while emphasizing that VA "can help to empower the animal tracking community and to foster new insight into the ecology and movement of tracked animals" [62]. Their tool is intended for the iterative enrichment and segmentation of biologging time series. The outcomes are communicated in trajectory views, horizon graphs, and line charts, and can be refined by a feedback loop. However, *AnimalEcologyExplorer* enables only value range segmentation that is applied on a single scale.

Alsallakh et al. describe an approach [2] where the user decides between various segmentation techniques and parameters. The user specifies intervals for their local application, observes outcomes in interactive visualizations, and can follow a feedback loop. In principle, the local application allows segmentation on multiple nested temporal scales, but the proposed one-dimensional color stripe visualization does not promote multi-scale analyses. In contrast, the *KronoMiner* [72] system of Zhao et al. has many visual-interactive features for the multi-scale exploration of segments, which are arranged in a hierarchical circular layout. However, defining these segments is a fully manual task without an automated segmentation technique involved.

Röhlig et al. [57] enable the analysis of the impact of different parameter configurations in an interactive prototype. Here, a time series is segmented multiple times with differently parameterized SSMs. The segmentation results are encoded by aligned color stripes, having several possible color encodings, aggregations, and interaction techniques. A related approach [7] by Bernard et al. aligns color stripes from various segmentation techniques. Though these approaches lack algorithmic and visual multi-scale support.

Further approaches do not consider typical segmentation techniques inherently, but still help to structure the time series. The *TimeClassifier* tool [68] of Walker et al. structures time series via template matching and visually encodes confidence of matched segments. A tool [28] of Gao et al. identifies and annotates segments based on a hierarchical support vector machine (SVM). This tool distinguishes between high- and low-level motion patterns but has few visual-interactive features.

The above VA approaches can be compared to MultiSegVA directly, see Table 1. Besides the requirements from Section 2, we also compare support for heterogeneous data and geographical context, as well as generalizability to other domains and diverse kinds of time series. There are several gaps inside the table, and in particular, no related approach fulfills each of the necessary **R1–R4**. This aspect can be traced back to different overall research objectives.

Further distinction would be possible by comparing the scalability in terms of temporal resolution, duration, and amount of different time series to be analyzed in context. Yet, these properties are not depicted in all compared works. The range for the amount of displayable time series goes from 1 with high temporal resolution and duration of a few hours [68] to many, but with lower temporal resolution [62]. We have found no tool that scales to multiple time series with high temporal resolutions (> 1 Hz) over durations longer than a few hours.

To conclude, to the best of the authors' knowledge, there is no VA approach that simultaneously puts a strong emphasis on intelligent multi-scale time series segmentation and the additional support by tailored visualizations and interactions. To fill this gap and address further properties, this publication describes a comprehensive VA system with tailored visual-interactive features and a new visual query language for multi-scale time series segmentation.

3.3 Related Visual Query Languages

Visual query languages (VQLs) allow a wide range of users to search databases by replacing textual queries with visual abstractions, thus offloading the complexity of query building from the user to the system.

Despite many existing VQLs, there is a need for further research at VQLs for temporal domains [16]. The main focus of current research about VQLs for temporal domains lies on event sequences. Decisionflow [31] is a VA solution for high-dimensional temporal event sequence data that provides a VQL to find event sequences by entering preconditions, episodes, and outcomes. A related approach by Kraus et al. [41] includes the visual-interactive segmentation and filtering of event sequences. Another approach, Segmentifier [22], iteratively applies filter operations to subdivide click-stream sequence data into smaller segments by imposing constraints for further detail inspection. Though, these aforementioned approaches focus on event sequences and not unlabeled time series data. In contrast, Stein et al. [63] consider time-annotated trajectories and present a VQL for the hierarchical definition of events. However, their approach only focuses on movement data and does not entail multi-scale segmentation.

4 THE MULTISEGVA PLATFORM

With MultiSegVA, the analyst uses tailored visual-interactive features to explore and refine the result of a time series' multi-scale segmentation (R2, R4). This result is a segment tree and stems from a new visual query language (Section 5) that enables to specify the hierarchical application order of diverse segmentation techniques. Section 4.1 gives an overview on MultiSegVA and its workflow. Section 4.2 describes and justifies visualization and interaction choices in the platform's main window, while detail windows are focused in Section 4.3.

4.1 System Overview and Feedback-based Workflow

MultiSegVA has one main window where the analyst builds visual queries and analyzes segmentation results inside a hierarchy visualization, closely linked to uni- and multivariate time series plots. The platform includes two additional windows for detail analysis of interesting segments. See Figures 1 and 3 for an overview on this scheme.



Fig. 3. The main window (with its segment tree visualization and time series plots) that enables to send segments to $n+m$ detail windows.

One of these additional windows focuses on temporal details and local anomalies. In presence of geographical dimensions, the analyst can open the other additional window that has an interactive trajectory map and gives the spatial context. MultiSegVA ensures a close linking between these windows, even when the analyst opens them on different computers. The intended workflow with MultiSegVA contains the following six steps, also shown in Figure 4. The workflow involves a feedback loop for iterative analysis and refinement.

1. Series Selection. In the main window, the analyst initially decides on a time series: there are four options. (A) The analyst imports a biologging time series from Movebank [77], which is enabled by a Movebank-MultiSegVA linkage. In principle, the analyst can select a time series out of a set of more than 7,320 studies, with > 2.2 billion locations in total. (B) The analyst imports an arbitrary time series from a curated set that MultiSegVA offers. (C) Or else, the analyst uploads a .csv file to analyze an own time series. (D) With multiple series imported, it is possible to switch between the series by a select menu.

2. Browse Dimensions. Now the selected series' first numerical dimension is shown in a univariate time series plot. The plot's y-domain matches to the dimension's minimum and maximum value. To contextualize, all dimensions are shown in a multivariate plot below. At both plots, the analyst can flexibly select the dimension (subset) to show, followed by interactive zooming and panning. Thereby the analyst can develop an initial understanding of dimension-specific properties and can already identify (ir)relevant dimensions and intervals.

3. Query Building. The analyst arranges segmentation techniques on multiple scales, see the VQL in Section 5. Step 2 is helpful again.

4. Query Processing. This step is not steerable for the analyst.

5. Overview + Detail Visualizations. By a hierarchy visualization, the main window provides an *overview* on the segmentation result, see Figure 1 and Section 4.2. This visualization encodes, e.g., a segment's start index, stop index, and segment tree level. By color encoding a similarity measure, the analyst gets guided to interesting segments, indicating possible patterns and anomalies. With a right-click, the analyst sends an interesting segment to the additional windows, tailored for *detail* analysis (Section 4.3). Here the analyst inspects already

zoomed-in data values, point-based anomalies, and the spatial context of even very short segments. Simultaneously, the main window shows the *overview*, and steadily enables to select segments for *detail* analysis.

6. Reasoning. The analyst gets an intuition for scale properties, specific segment lengths and distributions, temporal subtleties, anomalies, spatial context, and corresponding influences. A reasoning step follows, especially in case of domain expertise. Thus in the main window, the analyst can label segments or refine the query. The platform enables such actions by a strong feedback loop: the analyst can dynamically trace back to step 1, 2, 3, or 5. Alternatively, the analyst can export the segment tree to a .csv file and continue externally by own analysis methods or extrapolations.

Other Workflow Mappings. This workflow also has smaller loops, e.g., between steps 2 and 3. Furthermore, like in [2], it is possible to map Keim's VA mantra onto the workflow: "analyze first - show the important - zoom, filter, analyze further - details on demand" [39]. Finally, the entire workflow can be also integrated into higher-order workflows, e.g., where the analyst starts with a straightforward segmentation query and iteratively increases the complexity after step 6.

4.2 Main Window: Visualization and Interaction

Besides the query building dialog, the main window's centerpiece is an icicle visualization [43] that encodes the segment tree for a time series with n records. While there are additional time series plots, the focus of this section is on describing and justifying the icicle visualization design and then the interaction with it.

Visualization. The icicle visualization recursively partitions display space along the x -axis and aligns child stripes below their parent stripe. In MultiSegVA, a stripe's vertical position encodes the respective segment tree level. The stripe height is constant, except for the root stripe that is visually less prominent as it covers all n records. A stripe's x -position is given by the linear $xPos(\text{from})$ with $xPos: [1, n] \rightarrow [x_{\min}, x_{\max}]$ in the simplest case. It is $xPos(\text{from}) < xPos(\text{from} + 1)$, and $xPos(\min(\text{to} + 1, n)) - xPos(\text{from})$ yields stripe width. By using record indices for position and size, a compact view is ensured even for time series with large temporal gaps.

The icicle technique is an adequate choice for two reasons. First, the technique provides an intuitive mapping of segment start index and length onto x -position and width. Segment start index and length are attributes that asked ecologists emphasized. Second, the icicle technique tends to be space-filling. Here, the icicle technique outperforms competing node-link diagrams that imply often unusable white-space between nodes. While node-link diagrams usually yield less intuitive mappings for a segment's start index and length, these diagrams could be more usable at very large-scale time series.

To guide the analyst to pattern and anomaly segments, a stripe's fill color encodes a similarity value: the segment's mean sibling distance \bar{d} in a given dimension (set). Siblings of a segment a with $[\text{from}, \text{to}]$ on level i are those other segments on level i that share with a their parent segment $[\text{from}', \text{to}']$ on level $i-1$. It is $\text{from}' = \min(\text{from})$ and $\text{to}' = \max(\text{to})$. Then MultiSegVA calculates \bar{d} between each segment and its siblings. For this, segments are normalized and then compared by dynamic time warping [8, 56]. A relatively low, respectively high, \bar{d} can indicate a segment-based repeating pattern, respectively anomaly.

A linear color scale $c: D \rightarrow R$ maps \bar{d} to an corresponding fill color. The color scale's range R is given by ColorBrewer's [33] diverging red-white-blue scheme. A blue stripe can point to a repeating pattern, a red stripe to an anomaly. There are four reasons for choosing this color scheme. (1) The diverging scheme reflects the natural break point $(d_{\min} + d_{\max})/2$: a whitish stripe indicates a relatively neutral segment. (2) The color scheme can intuitively encode the semantics of \bar{d} , especially for anomalies and neutral segments. (3) The colors are well distinguishable, and (4) the scheme is colorblind friendly.

MultiSegVA applies the color-encoding for a hovered stripe and its siblings, while colors for other stripes are faded out. The color scale's domain D is given by the sibling-based minimum and maximum of \bar{d} . This choice optimally ensures valid segment comparisons due to relatively similar segment lengths and contextual semantics. Deciding on D was an enduring process where variants were tested and discussed




Fig. 4. Feedback-based MultiSegVA workflow, as described in Section 4.1.

with domain experts: e.g., global color scale domain and mapping, as well as multiple level-wise color scales and legends (see Section 7).

Two kinds of labels support understanding and analyzing the icicle visualization. First, labels at the visualization’s left denote segment tree levels. Second, labels inside the stripes refer to the start index, stop index, and further annotations. While MultiSegVA derives annotations automatically (e.g., “inside value range $[r_{\min}, r_{\max}]$ ”, cf. **T3**), the analyst can also assign own labels. In both cases an uncluttered view is ensured, as a label only appears if it horizontally fits into its stripe.

Interaction. Next to manual labeling, manifold features for interacting with the icicle visualization are described in a thorough documentation, tutorial videos, and the following lines.

The used zooming and panning behavior ensures quick interaction with the icicle visualization. A left-click moves the clicked stripe into the focus and excludes potentially irrelevant stripes. Then the clicked stripe consumes much space, but the stripe’s parent, its direct neighbors, and children are still available for zooming and panning. This behavior involves animated transitions that aid at constructing a mental map and “maintaining object constancy” [5]. Note that panning is also doable by arrow keys, thus not needing pointing precision at thin stripes.

Next, hovering over stripes triggers two functions. First, there is subtree highlighting where the hovered stripe and each of its wrapping ancestor stripes obtain a green border. Second, hovering links the colored stripes to actual data values: namely, it triggers visual cues inside the aligned time series plots. The visual cues encode start and stop indices of hovered stripes and possible child stripes. Vice versa, hovering the time series plots triggers subtree highlighting. Besides these features, double-clicking a stripe synchronizes the below time series plots to the respective time domain.

Another interaction feature relates to placing bookmark icons on stripes by a mouse and keyboard interaction. These bookmarks are intended for three uses. (1) Bookmarks enable the icicle’s partial growth (Section 5.2). (2) With bookmarks it is trivial to remember stripes of interest, even after zooming, panning, or segmenting other time series. (3) By an additional button click, the analyst triggers local and ad-hoc anomaly detection for the bookmarked stripes. Section 4.3 describes the anomalies’ visual mappings in the temporal detail window.

Lastly, right-clicking a stripe triggers that the segment is sent to, and directly focused within, the temporal and geographical detail windows. There the analyst can effortlessly inspect, e.g., zoomed-in values and local anomalies, while retaining the overview in the main window.

4.3 Detail Windows: Visualization and Interaction

By visiting session-based URLs, the analyst opens additional windows for inspecting (a) temporal details and anomalies, as well as (b) the trajectory and spatial context of a forwarded segment. The scheme enables to retain the overview, extend display space, and do group-based analyses at multiple computers. This subsection describes and justifies the design of first the temporal then geographical detail window.

Temporal Detail Window. This window has by default one enriched time series plot and a linked stacked bar chart, both referring to one dimension of a forwarded segment, cf. Figure 5. The analyst can dynamically select previously forwarded segments, respectively dimensions. When the analyst triggers local anomaly detection in the main window, the enriched time series plot shows visual mappings of, e.g., point-based anomalies. The stacked bar chart aggregates these anomalies and shows their distribution over time.

The time series plot is an interactive line chart that offers several interpolation options and can hide line segments at temporal gaps. These properties hold also for the main window’s time series plots. In addition, the *enriched* time series plot attaches squares to individual

data points, for showing point-based anomalies. A square’s categorical fill color encodes the anomaly type: e.g., density-based local outlier [9], global outlier by *MAD* [45] or *S-H-ESD* [34], innovative outlier or temporary change by [17]. The square size is density-dependent to avoid visual overlap. To still guide the viewer to suspicious intervals and allow their cognitive processing [49], 7 equal-size rectangles are placed above the line chart. There, ColorBrewer’s *reds* scheme encodes the overlaid number of anomalies that are in a $\frac{1}{7}$ interval of the current time domain. Lastly, bars in the top are used for encoding motifs [71] and events [37].



Fig. 5. The temporal detail window in its expert mode with two enriched time series plots (top, middle) for comparing segments, and a stacked bar chart for showing anomaly distributions in forwarded segments.

For zooming and panning the plots, the analyst can decide whether the y-domain should adjust to local extrema. To improve perception, the lines become thicker if data density is below a threshold. Details (i.e., concrete data values, timestamps) are part of line-snapping tooltips that are density-dependent to avoid flickering effects. Further, the analyst can open an expert mode where two enriched time series plots are aligned. Here it is possible to automatically synchronize time domains and compare two different dimensions of a segment. The zoomed-in time domain is then indicated in the following stacked bar chart.

The stacked bar chart partitions the global time domain of the selected segment into equal-width bins. The bar heights encode normalized anomaly counts per bin b and anomaly type a . Again, fill color encodes anomaly type a . Beside absolute anomaly counts, various normalizations can address the following questions. Given bin b , how many % of the contained anomalies are of type a ? Given anomaly type a , how many % of type a anomalies occur in bin b ?

For inspecting small bars, the chart allows zooming in the y-direction. In x-direction, we prevent this interaction to retain the global time domain and yield the context for the (zoomed-in) enriched time series plot(s). By sorting items in an attached legend, the analyst changes the bars’ vertical order to inspect another anomaly type’s distribution at the baseline. Clicking a bar synchronizes the enriched time series plot(s).

Geographical Detail Window. In presence of geographical dimensions, the analyst can also forward a segment to the geographical detail window, inspect the directly zoomed-in segment on an interactive map, and observe the spatial context. Thereby a segment can be enriched by information about, e.g., land use, driving and disturbing factors. Again, it is at any time possible to select and thereby directly zoom-in two different, previously forwarded, segments.

A dark blue, respectively red, polyline encodes (lat, long) tuples of a focused segment. For understanding how a focused segment relates to

the overall movement, another polyline encodes all (lat, long) tuples of the time series, wrapping the focused segment. This polyline has a lower z-index and paler hue to avoid overlap and make it less prominent. To get a better feel for movement direction, polylines are enriched by equidistant arrows that rely on geographical headings.

Another of these arrows snaps to the $(\text{lat}[i], \text{long}[i])$ that is closest to the mouse cursor's location. Details for the i -th record are to the left of the map. To ensure a strong linkage to the main window, a map marker is placed at the location of a record that is hovered in previous time series plots. Also, on the interactive map the analyst can use Google Maps utilities to draw markers, circles, rectangles, and custom polygons. Their geometry coordinates can be copied to the clipboard and inserted at the following VQL to segment a time series by a custom geographical area. Finally, it is possible to vary polyline width, toggle map labels, select tile lightness and layout.

5 VQL FOR MULTI-SCALE SEGMENTATION

This section answers how to obtain the main window's segment tree. To model different multi-scale structures and allow flexible parameterization (R3), MultiSegVA includes a new visual query language that links various segmentation techniques (R1) across multiple scales. Here the analyst defines the hierarchical application order (R2) of techniques by a simple play with building blocks. So, technique B shall locally operate on the segments resulted from technique A.

As a major strength, the VQL avoids drawbacks that come with textual queries and code snippets. While the exact multi-scale structure is often unknown a priori, changing the nesting of a text query or code can be error-prone, relatively tedious, and commonly does not entail a compact visual representation. Further strengths (e.g., technique variability, query expressiveness) are depicted within this section.

While in principle the VQL can consume any indices-serving technique, we exemplify a domain-oriented set of segmentation techniques that was derived in close collaboration with movement ecologists: see the overview with domain-specific use cases in Table 2. Besides, the VQL also has operators to consider important relations between segments on the same scale. A multi-scale segmentation query, including techniques and possibly operators, becomes recursively parsed and evaluated in MultiSegVA's Java backend. The resulting segment tree references a segment's start and stop record index, next to child nodes.

Section 5.1 specifies preconditions and the VQL's elementary concept. To make concepts more tangible, the section also depicts query building. Section 5.2 specifies one VQL selector and several operators.

5.1 Hierarchical Application and Query Building

Preconditions. An array element $m[i]$ shall describe technique type and parameters (e.g., relevant dimension, k) for applying a segmentation technique, such as from Table 2. Then the method call $\text{getSplitIndices}(m[i], (t, v), \text{from}, \text{to})$ shall locally apply $m[i]$ at the index interval $[\text{from}, \text{to}]$. The outcome shall be a list l with unique, local, and sorted split indices $\in \{0, \dots, \text{to} - \text{from} + 1\} \subset \mathbb{N}$. The list head shall be padded by 0 and the list tail by $\text{to} - \text{from} + 1$.

Concept: Hierarchical Application. The VQL's underlying concept is to iterate over segmentation techniques, and use split indices resulting from $m[i-1]$ to get the application ranges for the next $m[i]$. So, assume that list l contains a split indices of the globally applied technique $m[1]$; it is $[\text{from}, \text{to}] = [1, l.\text{length}]$. To obtain split indices of a subsequent technique $m[2]$, one calls $\text{getSplitIndices}(m[2], \dots)$ exactly $(a-1)$ times. Here one iterates over new application ranges, i.e., $[\text{from}, \text{to}] = [l[j] + o, l[j+1] + o - 1]$. The offset o is relevant for projecting from local to global indices. Obviously one can apply this concept recursively and develop hierarchies with many levels.

This way of specifying hierarchical application has two benefits. First, technique and dimension variability is ensured. Various segmentation techniques can operate in different dimensions, but the techniques' results are still assimilable into one tree. Second, when segmenting with $m[i]$ there is only one *external* dependence: the application range $[\text{from}, \text{to}]$ given by $m[i-1]$. Apart from this, $m[i]$ is fully independent of other segmentation steps from any tree level. This independence can be exploited for subquery caching and parallelization.

Interaction: Query Building. Before providing further formalization, the remaining subsection shall make the VQL's fundamental concept even more tangible. To this end, Figure 6 shows the query building dialog. The query in the figure's right segments a time series initially by temporal gaps, then recursively by changes in categorical values, and finally calls a fuzzy operator on tree level 3. Building such a query relies on three simple and effortless interactions.



Fig. 6. The query building dialog is shaped by three steps: choose a segmentation technique, set parameters, arrange the technique/hierarchy.

(1) Initially the analyst decides on a segmentation technique, see Figure 6, left. The techniques are grouped into meaningful categories, and their documentation is linked by question mark icons. (2) Now the analyst parameterizes the chosen technique. The number of required parameters is kept low, partly given with optional settings and suggested values (e.g., number of motifs). Also, MultiSegVA automatically suggests dimensions with suitable data type for the chosen technique. Longitude and latitude dimensions are internally identified by column headers. Invalid input is caught before backend query processing. (3) Lastly the analyst places the parameterized technique in the query hierarchy (Figure 6, right). Here the technique is represented by a building block, whose position is steered by a blinking rectangle. Clicking on an existing building block triggers its blinking and enables to change parameters or replace the technique. Besides, building blocks can be removed via right-click and hierarchy levels are interactively sortable along the y-axis.

With this base, the specification of further query language features, e.g., of the red AFTER operator in Figure 6, follows.

5.2 Selector and Operators

This section describes the intuition and relevance of one query language selector and several operators.

Bookmark Selector. This selector enables a tree's partial growth. Given k indices from $\text{getSplitIndices}(m[i-1], (t, v), \text{from}, \text{to})$, this selector ensures that the next level's technique $m[i]$ does not perform on each of the $k-1$ new application ranges $[\text{from}, \text{to}]$. With this selector it is possible to exclude application ranges that are irrelevant for further analysis, and thereby, to reduce processing times.

Operators. The VQL has operators (OR, AND, AFTER, NEAR, NOT) that enable to link techniques also on a *single* tree level (i.e., "horizontal" linking). The operators are implemented by the template in Algorithm 1 that consumes the array m^* of techniques to link, as well as the time series and application range. The template iterates over m^* , extracts split indices in $[\text{from}, \text{to}]$, and aggregates them by the operator-dependent function f . In the following we focus on the intuition and relevance of OR, AND, AFTER, before referring to further (conceivable) operators.

OR Operator. This operator merges split indices from multiple techniques and parameterizations. So one could, e.g., segment a series by manually specifying multiple geographical areas, apply pattern matching with multiple query patterns (in various dimensions), segment by anomalously high or low value ranges (in various dimensions). At OR, the template's f is responsible for list concatenation.

After this first example, it is evident that an operator description can serve as the first parameter at $\text{getSplitIndices}(\dots)$, just like the description for a common technique from Table 2. This property enables nested queries on a single tree level: e.g., a $m^*[i]$ could relate to a second OR operator. Thus highly complex and expressive queries are possible, also with the following template-based operators.

Algorithm 1 Operator Template

```

1: procedure OPERATOR( $m^+, (t, v)$ , from, to)
2:   outIndices  $\leftarrow$  initEmptyList()
3:   for each  $m^+[i]$  in  $m^+$  do
4:     newIndices  $\leftarrow$  getSplitIndices( $m^+[i]$ ,  $(t, v)$ , from, to)
5:     outIndices  $\leftarrow$   $f$ (newIndices, outIndices)
6:   end for
7:   return UniqueAndSort(outIndices)
8: end procedure

```

AND Operator. This operator retains only these split indices that are returned by all specified techniques in m^+ . Thus this operator is useful for obtaining exact co-occurrences of split indices, given different techniques or parameterizations. At AND, the template’s f corresponds to list intersection.

AFTER Operator. This operator also focuses on co-occurrences but it is order-dependent and not as restrictive as AND. Namely, AFTER allows co-occurrences with a temporal tolerance $\theta \geq 0$. Given the techniques $m^+[1]$ and $m^+[2]$, this operator returns a split index a from $m^+[1]$ if there exists a split index b from $m^+[2]$ with $b \geq a$ and $b - a \leq \theta$. Here, f is used to build and extend θ -chains, starting at indices from $m^+[1]$ and ending at the indices from the last $m^+[i]$.

Further Operators. NEAR is similar to AFTER but allows also $b \leq a$. NOT is to be combined with other operators: here, f removes indices from $U = \{1, \dots, \text{to} - \text{from}\}$. We focused on the integration of above operators to avoid overwhelming interaction choices; still, further (e.g., density-based) operators can be well realized by the template.

6 USE CASES AND DOMAIN EXPERT FEEDBACK

MultiSegVA has been developed in close collaboration with movement ecologists and fulfills our audience’s requirements (Section 2). We show MultiSegVA’s usefulness and applicability by two extensive real-world use cases that were discussed with our domain experts. The use cases originate from live sessions where the domain experts and we were together using MultiSegVA. A third use case from power systems shows the general applicability. Finally, we depict expert feedback that was also used to iteratively improve MultiSegVA.

6.1 Use Case: Environment-Aware Behavior Analysis

This use case shows how MultiSegVA helps movement ecologists at identifying environment-aware behaviors in multivariate time series. While behavior is the response to complex environments with many temporal and spatial scales, our experts currently do not have a tool to identify highly specific, environment-aware behaviors. Instead, they apply script chaining that is inflexible at complex environments. Meanwhile, MultiSegVA enables the seamless multi-scale exploration and identification of such behaviors. For this use case, a biologging time series of one migrating Himalayan vulture is focused, including dimensions for location, acceleration, and environmental conditions. See the schematic workflow and scale properties in Figure 7.

Himalayan vultures belong to the heaviest flying birds [60] and have their main habitat on “the Roof of the World”, the Tibetan Plateau. Especially juvenile Himalayan vultures perform outstanding altitudinal migrations as part of post-breeding dispersal. Here, the juveniles migrate to warmer regions of Northern India and back to the heights of the Himalaya and Tibetan Plateau. Partly the vultures reach altitudes up to 9,000 m AMSL. Such altitudinal migrations come with a bandwidth of different environmental conditions: including air densities, wind speeds, temperatures, and resource availability. Movement ecologists are interested in how vultures migrate through these conditions by minimum energy expenditure. Are beneficial environmental conditions always exploited? Are there context-specific differences in acceleration?

The used time series [78] stems from Movebank, covers one year, and has 798,120 records. Each record has two GPS dimensions, one altitude dimension, three acceleration dimensions, and several environmental dimensions. The acceleration data are given on average every 16.16 minutes ($\sigma = 52.19$) by a burst of usually 40 records at a frequency of $f \text{ Hz} \in \{10.54, 18.74\}$. “By recording at high resolution [...]

but short duration [...] this strategy aims to sample just one behavior type and avoid behavioral transitions that could complicate automated classification statistics” [10]. Movebank’s Env-DATA service [23] was used to derive environmental dimensions (e.g., orographic uplift).

The analyst imports such a time series by uploading a .csv file. Then in the geographical detail window, the analyst visually explores the trajectory that shows the altitudinal migration from the Tibetan Plateau to India, and vice versa. To distinguish migration stages in 798,120 records, the analyst starts query building in the main window: change point detection at the altitude dimension. Provided time series plots help in determining the number of change points $k = 2$ (i.e., three segments). Visual exploration in the third segment shows that the vulture ascended approximately 5,000 meters within less than four days.

The analyst wants to see how the ascent and high-altitude flight correlate with beneficial environmental conditions. Eastward wind speed, thermal uplift, and in particular orographic uplift have a strong rise during the ascent, shown in the plots and indicating possible migration aids (Figure 7, left). Still, there are intervals with adverse uplift conditions and behaviors on high altitudes. To identify these intervals, the analyst places a bookmark on the high-altitude segment and applies segmentation by the orographic uplift’s value range on tree level 2.




Fig. 7. From 798,120 to 40 records: the query and scale-wise specifics at our first use case.

As the vulture could also rest on high altitudes, the analyst needs to distinguish between different behaviors, such as fly, rest, or restless. Therefore, segmenting by the y -acceleration’s value range on level 3 is sufficient, as earlier quantitative experiments could indicate. Adequate value ranges can be determined by visual exploration, some available ground-truth annotations, and domain expertise.

Visual exploration also points out temporal gaps due to burst-wise sampling; thus fully-automated segmentation by temporal gaps is applied to facilitate upcoming analyses. For continuing with statistical tests, the analyst exports the resulted 4-level segment tree to a .csv file. Finally, MultiSegVA returns bursts of flight patterns on high altitudes during adverse weather conditions, and besides, the platform can yield the time series’ full multi-scale structures, cf. Figure 8, a).

MultiSegVA facilitates answering questions as before for several reasons. (1) Without chains of scripts or complex models, MultiSegVA enables high flexibility and multivariate, environment-aware analyses by scale- and dimension-specific segmentation techniques and parameters. (2) Visualizations and interactions are supportive at query building and their scalability is ensured at 798,120 records. (3) MultiSegVA has strong I/O abilities and can be integrated into other analysis workflows.

6.2 Use Case: Clustering and Behavior Analysis

For our experts, analyzing spatial density clusters is essential as these regions can point to eating and hence hunting behaviors. Yet, the experts’ current analysis means do not enable (1) to flexibly define intervals relevant for the clustering, and (2) to directly apply further intra-cluster analyses. MultiSegVA handles (1) and (2) by progressive clustering: specifically, the analyst recursively chooses relevant intervals on multiple scales, applies one or more spatial clustering steps, and analyzes the intra-cluster acceleration. For this use case, a time series of a free-ranging domestic cat is focused, with data about GPS location and acceleration in more than 1.2 million records. An intermediate outcome of a 3-level segmentation is given by Figure 8, b).

Free-ranging cats have considerable ecological effects on wildlife. For instance, cats interfere in the habitats of wild predators [42], and in

Segmentation by	Category	Description/Reference	Movement Ecology Application Cases
Temporal Gaps	Value-Independent	Split at anomalous time difference between records.	Bursts, different recording sessions, changes in sampling rate.
Bins	Value-Independent	Equal or varying depth bins.	Comparisons, splitting by calendar units, discarding irrelevant segments, adding offsets.
Change Points	Generic Value-Based	See Section 3.1 and [3,66].	Posture changes, unintended collar shifts, simple distinction activity vs. non-activity.
Numerical Value Range	Generic Value-Based	Split when entering or leaving the value range $[r_{min}, r_{max}]$.	Flight altitudes or weather conditions, patterns regular in overall shape but elongated.
Categorical Values	Generic Value-Based	Split when the categorical value changes.	Validating automated or manual annotations, see the Movebank [77] attribute dictionary.
Seasonal Patterns	Generic Value-Based	Based on power-maximizing frequency in periodogram.	Regular short-time patterns (e.g., flapping), daily fluctuations, latent seasonalities.
Motif Representatives	Generic Value-Based	Split when a motif representative (see [70]) begins or ends.	Potentially latent recurrences without high assumptions on temporal structure, starting point for further analyses [73]. Further application cases in [46,48,50].
Pattern Matches	Generic Value-Based	Split when a matched pattern instance (see [56]) begins or ends.	Follow-up after finding interesting patterns or anomalies, analyzing temporal distribution of matching instances as well as their influence on temporal context.
Geographical Area	GPS-Based	Split when entering or leaving the specified polygon.	A priori estimated home ranges, analyzing relations to specific objects in space, exclude areas with high GPS error.
Density Clusters	GPS-Based	Split when entering or leaving a HDBSCAN [13] cluster.	Automated identification of migration habitats, home ranges, regions with high resource availability. Excluding noise regions.
Local Minima of First-Passage Time	GPS-Based	FPT is "the time taken by an animal to cross a circle with a given radius" [54].	Strongly incorporating time component, analyses from "explorative movements to area-restricted search" ([52], p. 59).

Table 2. Domain-oriented set of techniques for structuring biologging time series. Value-independent techniques consider only timestamps or indices. Some techniques are inherently no typical segmentation techniques, still their output can be mapped to the definition in Section 1.

particular, the cats' hunting behavior leads to significant reductions of wildlife populations [35], e.g., by up to 20.7 billion mammals per year in the US [47]. Yet, so far there is little known about the cats' exact hunting behavior, based on high-resolution biologging data and without disturbing the cats. In fact, identifying detailed hunting and eating behaviors of cats is a very recent topic. Our experts are currently working on this topic and emphasize the value of spatial density clusters.

While the GPS and acceleration dimensions rely on high sampling rates (continuous 1 Hz and 16 Hz), the time series has recordings from three days that are separated by temporal gaps. Clearly, the three days can reflect entirely different environmental factors and the clusters should not include records from distinctly separated times. As visual exploration suggests, the analyst applies segmentation by temporal gaps on tree level 1. This step extracts the two (out of > 1.2 million) record indices where the time difference to the next record is anomalous.

Focusing one or more days separately, the analyst now excludes the time series' segments that correspond to irrelevant movements in the trajectory. So on tree level 2, the analyst applies segmentation by a drawn geographical area (see Table 2) to exclude segments inside buildings, which is well feasible at free-ranging cats. This step is useful as there is relatively less prey and the cat is disturbed, but especially, the GPS signal is erroneous inside buildings. To additionally reflect segments outside home ranges, it suffices to link these areas via OR.

Now bookmarking only relevant segments, the analyst locally applies a DBSCAN variant at GPS dimensions on level 3. The local clustering avoids irrelevant data, is faster, and can be again applied on several further levels to better capture hierarchical cluster components.

Finally, for relevant cluster segments, value range segmentation or change point detection at acceleration dimensions helps in distinguishing coarse behaviors. Exported results can then serve as input at a recent project, i.e., for behavior classifiers with more than 50 classes. To conclude, this use case shares benefits from Section 6.1, but it also shows how MultiSegVA works within the clustering paradigm and how it unifies more diverse techniques into one compact workflow.

6.3 Use Case: Power Systems

To outline MultiSegVA's value even beyond movement ecology, we briefly present a use case from the power systems domain. For this, we focus on time series that stem from ENTSO-E [75] and give country-wide total power loads. The time series are sampled by every 15 minutes and have a time range of several years, up to near realtime. At visual exploration with MultiSegVA, seasonal patterns become evident on various scales: annual (e.g., summer valleys), weekly (e.g., weekend valleys), and daily (e.g., lunch break). To cope with such patterns, MultiSegVA and its VQL already include techniques such as seasonality detection or binning (Table 2). MultiSegVA shows the nestedness of seasonal patterns and ensures the structured and guided multi-scale exploration. Also, MultiSegVA provides features for in-depth segment comparisons and local anomaly detection (e.g., Christmas trend) in the temporal detail window. Referring to only two segmentation techniques, this use case already indicates that MultiSegVA variants for other domains are well conceivable. Hereto, specifically tailoring domain-oriented sets of techniques in collaboration with domain experts (e.g., grid operators) is a promising way for most insightful analyses.

6.4 Domain Expert Feedback

Iterative feedback by reputable experts in movement ecology allowed us to tailor MultiSegVA for several months and fulfill the elicited requirements. Now, the use cases become complemented by excerpts of this iterative feedback and summarizing opinions.

The project began by eliciting requirements, their implementation was reviewed and refined by weekly meetings with the movement ecologists. Here one could discuss relevant segmentation techniques, their parameters, categorization, and application cases in movement ecology. For the main visualization, the ecologists emphasized attributes that appear most important to encode: segment duration, position, and similarity. After prototyping, it was even possible to decide together on the subtleties of visualization and interaction.

After the main development phase, structured live demos of Multi-

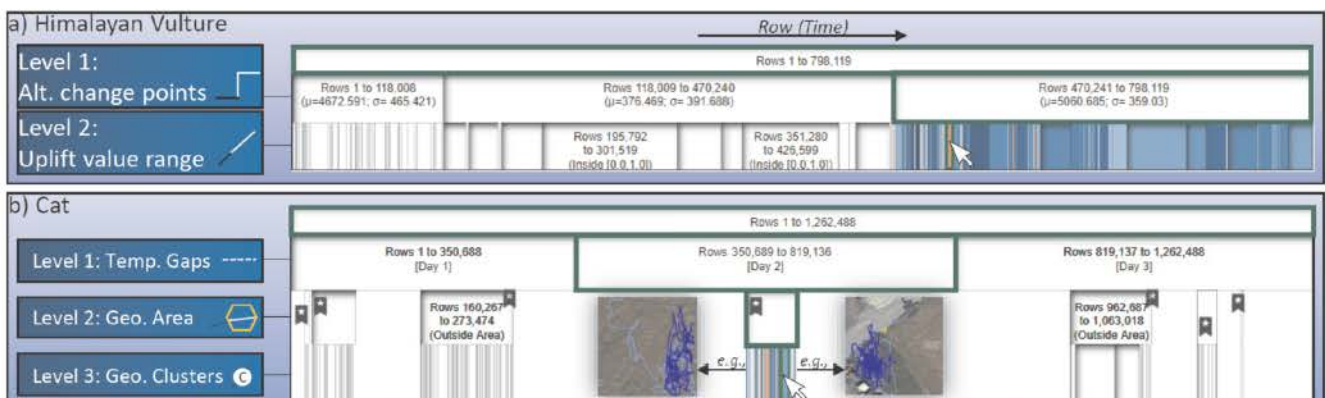


Fig. 8. a) Full multi-scale segmentation at splitting the vulture's time series by altitude change points and uplift value range. b) Segmenting the cat's time series by temporal gaps, then locally by area and finally by clusters. Bookmarks allow partial growth and can be automatically attached to labels.

SegVA were given to more than ten movement ecologists who had not been involved until this point. Here we could gather informal feedback that was distinctly positive but still included requests for improvements.

The movement ecologists showed much interest, highlighted the usefulness of given multi-scale segmentation, next to the ease of query building. The query building dialog was denoted as comprehensive and well ordered. The experts liked the close linkage between the segmentation task and visual-interactive features, as well as the linkage between visualizations across computers. Several experts were surprised about MultiSegVA's scalability and short processing times. According to them, MultiSegVA could have saved much time at a recent project.

However, the movement ecologists depicted the need for more multivariate functions, especially for tri-axial acceleration data. Hence the multivariate plot was added to the main window. While VQL operators already supported multivariate analyses, several segmentation techniques were extended to inherently handle multivariate cases. Anyhow, time series with reduced or projected dimensions can be uploaded. Also, the movement ecologists asked how to know, which interaction features are given and how to use all of them. Thus video tutorials, an overview sheet, and extensive documentation were added. Moreover, the experts asked for export functions that were then added for segment trees and queries. Lastly, the icicle's initially global color scheme appeared difficult to understand, thus leading to the sibling-based scheme.

After these improvements, three movement ecologists tested MultiSegVA at their own computers without further guidance. Their feedback can give a first intuition on how simple or difficult the independent use is. The domain experts share the opinion that it is straightforward to build queries and segment time series. Still, the experts agree that at least basic prior knowledge about time series segmentation is needed.

We conclude this section by informal feedback that was gathered throughout the entire project: about (a) possible application cases and (b) perceived novelty. Regarding (a), the movement ecologists see MultiSegVA with its current techniques as a very useful tool for multi-scale exploration and segmentation. One domain expert sees the platform's use for modeling an individual's life history. In general, analyzing populations or with further techniques shall follow after exporting results. One expert would be interested in synchronizing video recordings with the platform, another expert would use MultiSegVA for validation. The experts were also asked (b) to assess the platform's novelty, based on their perspective and domain experiences. Certainly, such opinions depend on the exact background and sub-domain specializations, but still, they give another intuition on the platform's value. The easy segmentation with the emphasis on the multi-scale setting and the strong linkage to visual-interactive methods were seen as most novel. Several movement ecologists agree that MultiSegVA is a progress to what is currently available to them and that it allows them to do more semantically meaningful multi-scale analyses.

7 DISCUSSION

MultiSegVA enables comprehensive exploration and refining of multi-scale segmentations by tailored visual-interactive features and VA paradigms. MultiSegVA includes segment tree encoding, subtree highlighting, guidance, density-dependent features, adapted navigation, multi-window support, and a feedback-based workflow. The VQL facilitates exploring and parameterizing different multi-scale structures. Still, few aspects remain for further reflection.

The icicle visualization meets expert requests and has several benefits. Yet, guiding the user by color to interesting parts of the segment tree is a challenging task. We tested global, level-based, and sibling-based guidance variants and according color fills. We chose sibling-based guidance (i.e., all siblings of one hovered segment are colored) that optimally captures local similarities, while requiring more navigation effort across levels and nodes. Upcoming works will include an even more effective variant, i.e., guidance to local similarities with little interaction and one fixed color scale.

Our VQL makes it trivial to *build* a multi-scale segmentation. Query building is a play with building blocks that benefits from strong abstraction and simple interactions. Rather it is difficult to *decide* which multi-scale structure and building blocks are most appropriate: a deci-

sion that depends on data, analyst, and tasks. MultiSegVA facilitates this decision by extensive documentation, technique categorization, few technique parameters, and short processing times in a compact workflow. For further support, we plan predefined queries, instant responses at query building, next to parameter and technique suggestions.

For suggesting parameters, we will apply estimators [19, 69] for the number of change points as well as the elbow method for *knn*-searches. While motif length and HDBSCAN's *minPts* [13] optimally benefit from domain expertise, suggesting other parameters will simplify the interaction and can address another limitation. Now, a technique processes each segment of one scale with the same parameters; thus slight data-dependent parameter modifications will be examined. For technique suggestions, we envision for each technique a scale-wise relevance score that reflects data properties and is part of a rule-based prioritization, shaped by domain expertise and meaningful hierarchies.

It is essential to depict the semantics into which MultiSegVA can give insights. First of all, MultiSegVA illuminates diverse multi-scale structures and gives insights on how scales relate to each other. Coarse behaviors can be distinguished by relatively simple techniques, motifs show repetitive behaviors, and *knn*-searches allow the matching with already explored segments. Segment lengths and similarities can be explored, next to local anomalies and spatial contexts. However, with the current techniques it is difficult to broadly capture deeper, behavioral semantics (e.g., chew, scratch). Hereto more complex or learning techniques (e.g., HMMs, SVMs) will be needed that neither overfill the interface nor delimit generalizability due to the lack of learned patterns.

The latter point goes hand in hand with our major limitation and the corresponding implication for upcoming work: integrating even more intelligent methods and automatisms. These plans all relate to aspects from above, i.e., better guidance, technique and more parameter suggestions, as well as techniques for deeper behavioral semantics.

MultiSegVA relies on requirements by movement ecology experts and stands for an iterative, extensively collaborative and interdisciplinary process. We could gather domain feedback on several stages, derive a domain-oriented set of techniques, and even link MultiSegVA to Movebank with > 2.2 billion animal locations. With this application domain focused, MultiSegVA underpins the value of multi-scale analyses and is certainly another step forward "to empower the animal tracking community and to foster new insight into the ecology and movement of tracked animals" [62]. Meanwhile, our third use case shows that MultiSegVA variants for other domains are conceivable, especially with tailored domain-oriented technique sets. This generalizability is promoted by the platform's I/O features and its ability to handle heterogeneous time series with > 1.2 million records.

8 CONCLUSION

We presented the web-based MultiSegVA platform that facilitates multi-scale segmentation of biologging time series and enables various semantic analyses. To explore results and refine parameters, MultiSegVA primarily contributes the use of visual-interactive features and VA paradigms that are specifically tailored for multi-scale segmentation. To flexibly model multi-scale segmentations, our VQL is a simple play with building blocks and the second contribution. As an input, the VQL takes techniques out of a domain-oriented set (third contribution) that can fulfill various segmentation objectives, specifically for scale and dimension. MultiSegVA is shaped by a fruitful collaboration and movement ecology experts see MultiSegVA as a very useful approach for semantically meaningful multi-scale analyses. We are looking forward to integrating even more intelligent methods where interdisciplinary collaboration builds again the basis for effectiveness and quality.

ACKNOWLEDGMENTS

This work has been supported by the German Research Foundation DFG within Priority Research Program 1894. We thank the many inspiring researchers from the MPIAB and related institutions who have contributed valuable feedback. They have given us fascinating views into their domain.

REFERENCES

- [1] M. Alber, A. B. Tepole, W. R. Cannon, S. De, S. Dura-Bernal, K. Garikipati, G. Karniadakis, W. W. Lytton, P. Perdikaris, L. Petzold, et al. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *npj Digital Medicine*, 2(1):1–11, 2019.
- [2] B. Alsallakh, M. Bögl, T. Gschwandtner, S. Miksch, B. Esmael, A. Arnaout, G. Thonhauser, and P. Zöllner. A visual analytics approach to segmenting and labeling multivariate time series data. *Proc. of EuroVA*, 14:31–35, 2014.
- [3] S. Aminikhanghahi and D. J. Cook. A survey of methods for time series change point detection. *Knowledge and information systems*, 51(2):339–367, 2017.
- [4] G. Andrienko, N. Andrienko, P. Bak, D. A. Keim, and S. Wrobel. *Visual Analytics of Movement*. Springer Berlin Heidelberg, 2013.
- [5] B. B. Bederson and A. Boltman. Does animation help users build mental maps of spatial information? In *Proceedings 1999 IEEE Symposium on Information Visualization (InfoVis '99)*, pp. 28–35. IEEE, 1999.
- [6] S. Benhamou. Of scales and stationarity in animal movements. *Ecology letters*, 17(3):261–272, 2014.
- [7] J. Bernard, E. Dobermann, M. Bögl, M. Röhlig, A. Vögele, and J. Kohlhammer. Visual-interactive segmentation of multivariate time series. In *EuroVis Workshop on Visual Analytics (EuroVA)*. Eurographics, 2016.
- [8] D. J. Berndt and J. Clifford. Using dynamic time warping to find patterns in time series. In *KDD workshop*, vol. 10, pp. 359–370. Seattle, WA, 1994.
- [9] M. M. Breunig, H.-P. Kriegel, R. T. Ng, and J. Sander. Lof: identifying density-based local outliers. In *ACM sigmod record*, vol. 29, pp. 93–104. ACM, 2000.
- [10] D. D. Brown, R. Kays, M. Wikelski, R. Wilson, and A. P. Klimley. Observing the unwatchable through acceleration logging of animal behavior. *Animal Biotelemetry*, 1(1):20, 2013.
- [11] M. Buchin, A. Driemel, M. Van Kreveld, and V. Sacristán. Segmenting trajectories: A framework and algorithms using spatiotemporal criteria. *Journal of Spatial Information Science*, 2011(3):33–63, 2011.
- [12] M. Buchin, H. Kruckenberg, and A. Kölsch. Segmenting trajectories by movement states. In *Advances in spatial data handling*, pp. 15–25. Springer, 2013.
- [13] R. J. G. B. Campello, D. Moulavi, and J. Sander. Density-based clustering based on hierarchical density estimates. In *Pacific-Asia conference on knowledge discovery and data mining*, pp. 160–172. Springer, 2013.
- [14] O. Cappé, E. Moulines, and T. Ryden. *Inference in Hidden Markov Models*. Springer Series in Statistics. Springer New York, 2006.
- [15] D. W. Cash, W. N. Adger, F. Berkes, P. Garden, L. Lebel, P. Olsson, L. Pritchard, and O. Young. Scale and cross-scale dynamics: governance and information in a multilevel world. *Ecology and society*, 11(2), 2006.
- [16] T. Catarci. *Visual Query Language*, pp. 3399–3405. Springer US, Boston, MA, 2009. doi: 10.1007/978-0-387-39940-9_448
- [17] C. Chen and L.-M. Liu. Joint estimation of model parameters and outlier effects in time series. *Journal of the American Statistical Association*, 88(421):284–297, 1993.
- [18] H. Cho and P. Fryzlewicz. Multiscale and multilevel technique for consistent segmentation of nonstationary time series. *Statistica Sinica*, pp. 207–229, 2012.
- [19] T. T.-L. Chong. Estimating the locations and number of change points by the sample-splitting method. *Statistical papers*, 42(1):53–79, 2001.
- [20] U. Demšar, K. Buchin, F. Cagnacci, K. Safi, B. Speckmann, N. Van de Weghe, D. Weiskopf, and R. Weibel. Analysis and visualisation of movement: an interdisciplinary review. *Movement Ecology*, 3(1):5, 2015.
- [21] H. Dette, T. Schüller, and M. Vetter. Multiscale change point detection for dependent data. *arXiv preprint arXiv:1811.05956*, 2018.
- [22] K. Dextras-Romagnino and T. Munzner. Segmentifier: Interactive refinement of clickstream data. In *Computer Graphics Forum*, vol. 38, pp. 623–634. Wiley Online Library, 2019.
- [23] S. Dodge, G. Bohrer, R. Weinzierl, S. C. Davidson, R. Kays, D. Douglas, S. Cruz, J. Han, D. Brandes, and M. Wikelski. The environmental-data automated track annotation (env-data) system: linking animal tracks with environmental data. *Movement Ecology*, 1(1):3, 2013.
- [24] U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, et al. Knowledge discovery and data mining: Towards a unifying framework. In *KDD*, vol. 96, pp. 82–88, 1996.
- [25] M. A. Ferreira, A. Marco, and H. K. Lee. *Multiscale modeling: a Bayesian perspective*. Springer Science & Business Media, 2007.
- [26] M. A. Ferreira, M. West, H. K. Lee, D. M. Higdon, et al. Multi-scale and hidden resolution time series models. *Bayesian Analysis*, 1(4):947–967, 2006.
- [27] K. Frick, A. Munk, and H. Sieling. Multiscale change point inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(3):495–580, 2014.
- [28] L. Gao, H. A. Campbell, O. R. Bidder, and J. Hunter. A web-based semantic tagging and activity recognition system for species’ accelerometry data. *Ecological Informatics*, 13:47–56, 2013.
- [29] S. Gharghabi, Y. Ding, C.-C. M. Yeh, K. Kamgar, L. Ulanova, and E. Keogh. Matrix profile viii: domain agnostic online semantic segmentation at superhuman performance levels. In *2017 IEEE International Conference on Data Mining (ICDM)*, pp. 117–126. IEEE, 2017.
- [30] A. C. Gleiss, R. P. Wilson, and E. L. Shepard. Making overall dynamic body acceleration work: on the theory of acceleration as a proxy for energy expenditure. *Methods in Ecology and Evolution*, 2(1):23–33, 2011.
- [31] D. Gotz and H. Stavropoulos. Decisionflow: Visual analytics for high-dimensional temporal event sequence data. *IEEE transactions on visualization and computer graphics*, 20(12):1783–1792, 2014.
- [32] B. A. Han, A. M. Kramer, and J. M. Drake. Global patterns of zoonotic disease in mammals. *Trends in parasitology*, 32(7):565–577, 2016.
- [33] M. Harrower and C. A. Brewer. Colorbrewer.org: an online tool for selecting colour schemes for maps. *The Cartographic Journal*, 40(1):27–37, 2003.
- [34] J. Hochenbaum, O. S. Vallis, and A. Kejarawal. Automatic anomaly detection in the cloud via statistical learning. *arXiv preprint arXiv:1704.07706*, 2017.
- [35] J. A. Horn, N. Mateus-Pinilla, R. E. Warner, and E. J. Heske. Home range, habitat use, and activity patterns of free-roaming domestic cats. *The Journal of Wildlife Management*, 75(5):1177–1185, 2011.
- [36] I. Jonsen, M. Basson, S. Bestley, M. V. Bravington, T. Patterson, M. W. Pedersen, R. B. Thomson, U. H. Thygesen, and S. J. Wotherspoon. State-space models for bio-loggers: A methodological road map. *Deep Sea Research Part II: Topical Studies in Oceanography*, 88:34–46, 2013.
- [37] Y. Kang, D. Belušić, and K. Smith-Miles. Detecting and classifying events in noisy time series. *Journal of the Atmospheric Sciences*, 71(3):1090–1104, 2014.
- [38] D. A. Keim, G. Andrienko, J.-D. Fekete, C. Görg, J. Kohlhammer, and G. Melançon. Visual analytics: Definition, process, and challenges. In *Information visualization*, pp. 154–175. Springer, 2008.
- [39] D. A. Keim, F. Mansmann, J. Schneidewind, J. Thomas, and H. Ziegler. Visual analytics: Scope and challenges. In *Visual data mining*, pp. 76–90. Springer, 2008.
- [40] E. Keogh, S. Chu, D. Hart, and M. Pazzani. Segmenting time series: A survey and novel approach. In *Data mining in time series databases*, pp. 1–21. World Scientific, 2004.
- [41] J. Krause, A. Perer, and H. Stavropoulos. Supporting iterative cohort construction with visual temporal queries. *IEEE transactions on visualization and computer graphics*, 22(1):91–100, 2015.
- [42] D. Krauze-Gryz, J. Gryz, J. Goszczyński, P. Chylarecki, and M. Zmihorski. The good, the bad, and the ugly: space use and intraguild interactions among three opportunistic predators—cat (*Felis catus*), dog (*Canis lupus familiaris*), and red fox (*Vulpes vulpes*)—under human pressure. *Canadian Journal of Zoology*, 90(12):1402–1413, 2012.
- [43] J. B. Kruskal and J. M. Landwehr. Icicle plots: Better displays for hierarchical clustering. *The American Statistician*, 37(2):162–168, 1983.
- [44] S. A. Levin. The problem of pattern and scale in ecology: the Robert H. MacArthur award lecture. *Ecology*, 73(6):1943–1967, 1992.
- [45] C. Leys, C. Ley, O. Klein, P. Bernard, and L. Licata. Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, 49(4):764–766, 2013.
- [46] B. Liu, J. Li, C. Chen, W. Tan, Q. Chen, and M. C. Zhou. Efficient motif discovery for large-scale time series in healthcare. *IEEE Transactions on Industrial Informatics*, 11(3):583–590, 2015.
- [47] S. R. Loss, T. Will, and P. P. Marra. The impact of free-ranging domestic cats on wildlife of the United States. *Nature communications*, 4(1):1–8, 2013.
- [48] J. Meng, J. Yuan, M. Hans, and Y. Wu. Mining motifs from human motion. In *Eurographics (Short Papers)*, pp. 71–74, 2008.
- [49] G. A. Miller. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological review*,

- 63(2):81, 1956.
- [50] A. Mueen. Time series motif discovery: dimensions and applications. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 4(2):152–159, 2014.
- [51] K. Nazemi, D. Burkhardt, E. Ginters, and J. Kohlhammer. Semantics visualization—definition, approaches and challenges. *Procedia Computer Science*, 75:75–83, 2015.
- [52] O. Ovaskainen, H. J. de Knecht, and M. M. Delgado. *Quantitative Ecology and Evolutionary Biology: Integrating models with data*. Oxford Series in Ecology and Evolution. OUP Oxford, 2016.
- [53] G. T. Pecl, M. B. Araújo, J. D. Bell, J. Blanchard, T. C. Bonebrake, I.-C. Chen, T. D. Clark, R. K. Colwell, F. Danielsen, B. Evengård, et al. Biodiversity redistribution under climate change: Impacts on ecosystems and human well-being. *Science*, 355(6332):eaai9214, 2017.
- [54] D. Pinaud. Quantifying search effort of moving animals at several spatial scales using first-passage time analysis: effect of the structure of environment and tracking systems. *Journal of Applied Ecology*, 45(1):91–99, 2008.
- [55] L. R. Rabiner. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.
- [56] T. Rakthanmanon, B. Campana, A. Mueen, G. Batista, B. Westover, Q. Zhu, J. Zakaria, and E. Keogh. Searching and mining trillions of time series subsequences under dynamic time warping. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 262–270. ACM, 2012.
- [57] M. Röhlig, M. Luboschik, F. Krüger, T. Kirste, H. Schumann, M. Bögl, B. Alsallakh, and S. Miksch. Supporting activity recognition by visual analytics. In *Visual Analytics Science and Technology (VAST), 2015 IEEE Conference on*, pp. 41–48. IEEE, 2015.
- [58] S. Sharma, S. Couturier, and S. D. Cote. Impacts of climate change on the seasonal distribution of migratory caribou. *Global Change Biology*, 15(10):2549–2562, 2009.
- [59] E. L. Shepard, R. P. Wilson, F. Quintana, A. G. Laich, N. Liebsch, D. A. Albareda, L. G. Halsey, A. Gleiss, D. T. Morgan, A. E. Myers, et al. Identification of animal movement patterns using tri-axial accelerometry. *Endangered species research*, 10:47–60, 2008.
- [60] S. Sherub, G. Bohrer, M. Wikelski, and R. Weinzierl. Behavioural adaptations to flight into thin air. *Biology letters*, 12(10):20160432, 2016.
- [61] S. Sherub, W. Fiedler, O. Duriez, and M. Wikelski. Bio-logging, new technologies to study conservation physiology on the move: a case study on annual survival of himalayan vultures. *Journal of Comparative Physiology A*, 203(6-7):531–542, 2017.
- [62] D. Spretke, P. Bak, H. Janetzko, B. Kranstauber, F. Mansmann, and S. Davidson. Exploration through enrichment: a visual analytics approach for animal movement. In *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pp. 421–424. ACM, 2011.
- [63] M. Stein, D. Seebacher, T. Karge, T. Polk, M. Grossniklaus, and D. A. Keim. From movement to events: Improving soccer match annotations. In *International Conference on Multimedia Modeling*, pp. 130–142. Springer, 2019.
- [64] G. Swan, V. Naidoo, R. Cuthbert, R. E. Green, D. J. Pain, D. Swarup, V. Prakash, M. Taggart, L. Bekker, D. Das, et al. Removing the threat of diclofenac to critically endangered asian vultures. *PLoS biology*, 4(3):e66, 2006.
- [65] A. Thiebault and Y. Tremblay. Splitting animal trajectories into fine-scale behaviorally consistent movement units: breaking points relate to external stimuli in a foraging seabird. *Behavioral Ecology and Sociobiology*, 67(6):1013–1026, 2013.
- [66] C. Truong, L. Oudre, and N. Vayatis. A review of change point detection methods. *arXiv preprint arXiv:1801.00718*, 2018.
- [67] C. Vamoş. Multiscale structure of time series revealed by the monotony spectrum. *Physical Review E*, 95(3):033310, 2017.
- [68] J. S. Walker, M. W. Jones, R. S. Laramee, O. R. Bidder, H. J. Williams, R. Scott, E. L. C. Shepard, and R. P. Wilson. Timeclassifier: a visual analytic system for the classification of multi-dimensional time series data. *The Visual Computer*, 31(6-8):1067–1078, 2015.
- [69] Y.-C. Yao. Estimating the number of change-points via schwarz' criterion. *Statistics & Probability Letters*, 6(3):181–189, 1988.
- [70] C.-C. M. Yeh, N. Kavantzias, and E. Keogh. Matrix profile vi: Meaningful multidimensional motif discovery. pp. 565–574, 11 2017. doi: 10.1109/ICDM.2017.66
- [71] C.-C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, and E. Keogh. Matrix profile i: all pairs similarity joins for time series: a unifying view that includes motifs, discords and shapelets. In *Data Mining (ICDM), 2016 IEEE 16th International Conference on*, pp. 1317–1322. IEEE, 2016.
- [72] J. Zhao, F. Chevalier, and R. Balakrishnan. Kronominer: using multi-foci navigation for the visual exploration of time-series data. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1737–1746, 2011.
- [73] Y. Zhu, Z. Zimmerman, N. S. Senobari, C.-C. M. Yeh, G. Funning, A. Mueen, P. Brisk, and E. Keogh. Matrix profile ii: Exploiting a novel algorithm and gpus to break the one hundred million barrier for time series motifs and joins. In *2016 IEEE 16th international conference on data mining (ICDM)*, pp. 739–748. IEEE, 2016.
- [74] W. Zucchini, I. L. MacDonald, and R. Langrock. *Hidden Markov Models for Time Series: An Introduction Using R, Second Edition*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. CRC Press, 2017.
- [75] Unicorn Systems. Entso-e transparency platform. <https://transparency.entsoe.eu/>. Accessed: 2020-04-28.
- [76] Icarus Initiative. Icarus - global monitoring with animals. <https://www.icarus.mpg.de/en>. Accessed: 2020-04-11.
- [77] M. Wikelski, S. C. Davidson, and R. Kays. Movebank: archive, analysis and sharing of animal movement data. hosted by the max planck institute of animal behavior. <http://www.movebank.org>. Accessed: 2020-04-11.
- [78] S. Sherub and M. Wikelski. Movebank: e-obs gprs himalayan griffon, sherub, bhutan. https://www.movebank.org/cms/webapp?gw_t_fragment=page=studies,path=study20202974. Accessed: 2020-04-29.