Graphical Abstract

Guided Spiral Visualization for Periodic Time Series and Residual Analysis

Julian Rakuschek, Helwig Hauser, Tobias Schreck

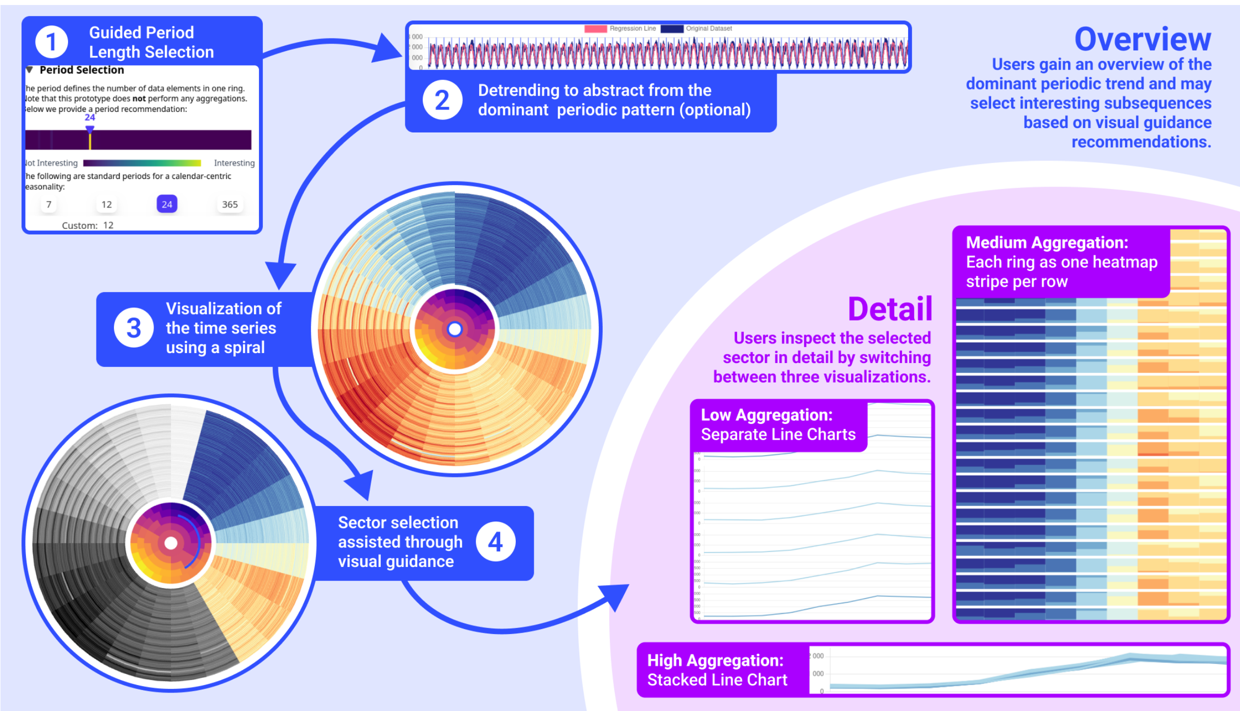

Guided Spiral Visualization for Periodic Time Series and Residual Analysis

Julian Rakuschek[a], Helwig Hauser[b], Tobias Schreck[a]

[a]*Institute of Visual Computing, TU Graz, Austria*
[b]*University of Bergen, Norway*

Abstract

Time series in domains such as climate, traffic, and energy often contain multiple, overlapping periodic patterns. Spiral visualizations can support the exploration of such data, but their effectiveness is limited in practice. Outliers and global trends skew the color mapping, dominant periodic components can hide weaker patterns, selecting a meaningful period length is challenging, and comparing subsequences within large datasets remains cumbersome. To address these challenges, we present a guided analytical workflow centered on an enhanced time series spiral visualization. A regression model tailored to periodic data helps identify suitable period lengths and exposes secondary patterns through its residuals. Visual guidance mitigates issues caused by skewed color mappings and highlights relevant spiral sectors even when global trends or outliers are present. Users can interactively select and compare sectors based on measures of average, trend, and similarity, and examine them in linked views or a provenance dashboard, which maintains a record of all user interactions and allows comparing multiple spirals with each other. Application examples demonstrate use cases where the visual sector selection guidance together with the exploration of model residuals leads to insights. In traffic data, for instance, removing the dominant day–night rhythm reveals rush-hour effects that become visible through exploration of the residuals.

Keywords: Periodic Time Series, Time Series Analysis, Spiral Visualizations, Guidance, Regression, Residual Analysis

1. Introduction

Periodic patterns are at the core of many time series. For example, when observing regional average daily temperature records, a seasonal effect is apparent when comparing records acquired from summer versus winter months. Time series spirals [2] are a common visualization technique for such data because they maintain continuity while aligning recurring cycles, enabling direct comparison of periods such as winter months across years. Despite their utility, spirals have two key limitations: choosing a meaningful period length can be difficult, and trends or outliers can dominate the view and occlude seasonal behavior. Even with clear seasonal patterns, secondary periodic structures may remain hidden under dominant cycles. This raises an analytical question: Once the main seasonal component is understood or removed, can a spiral visualization reveal periodic occurrences of second-level patterns? Existing spiral-based visualizations have largely focused on depicting global periodic structure and have not addressed the difficulties of subsequence comparison or the analysis of second-level patterns that emerge beyond dominant seasonal trends.

This paper introduces an analytical workflow (Figure 1) for exploring periodic time series using the time series spiral as the central representation. The workflow supports identifying and comparing informative subsequences using an overview+detail strategy [1] in which users locate interesting regions in the spiral and then examine them through coordinated detailed views. Relative to our earlier work [3], this work provides three key extensions. First, we refine the visual guidance mechanism for subsequence recommendation by introducing the *guidance donut*, a circular indicator placed at the center of the spiral. In contrast to the previous isolated guidance view, the guidance donut establishes a direct spatial correspondence between recommendation scores and spiral segments, enabling guidance even when trends or outliers affect the

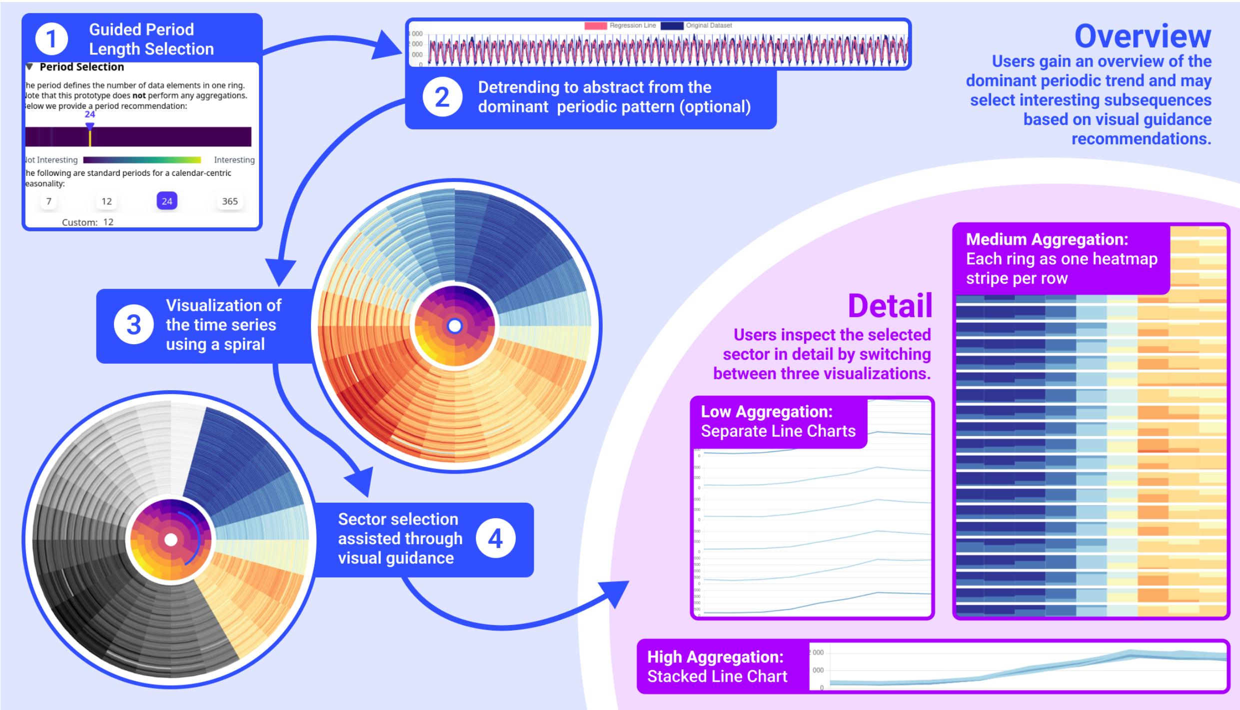

Figure 1: The design of our analytical workflow follows the **overview**+**detail** strategy [1]. In the overview, analysts select a period length (step 1), optionally apply a regression model to remove the dominant periodic trend (step 2) and visualize the data through a time series spiral (step 3). Users select periodic subsequences using visual guidance (step 4), which can be inspected in the detail view. Analysts may switch between three visualizations to control the required visual aggregation.

appearance of the spiral. Second, we integrate a regression model into the workflow to subtract the dominant seasonal component. The resulting residuals allow the analysis to focus on second-level periodic structures that may be obscured in the original data. Third, the regression model also supports period selection by providing cues for well-aligned spirals, thereby improving the reliability of subsequence comparison.

We demonstrate the utility of our approach through several application examples drawn from diverse domains, including environmental sensing, transportation, and energy consumption. These examples illustrate how the workflow supports three central analysis tasks: selecting an appropriate period length, identifying informative subsequences through visual guidance, and revealing second-level periodic patterns by analyzing regression residuals. Each case demonstrates situations in which specific components of the workflow, such as the guidance donut, the different interestingness measures, or the detrending step, provide insights that are difficult to obtain with a standard spiral visualization alone. In addition to these examples, we conducted an expert interview with two facility-management experts who applied the workflow to energy-consumption data from their domain. Their insights and observations complement the application examples and further demonstrate the practical value of the approach.

2. Related Work

We first provide background on time series visualization with emphasis on periodic data. Afterwards, we discuss existing approaches for extending the time series spiral visualization, specifically regarding guidance.

2.1. Visualization of Periodic Time Series

Our work builds on previous research in time series visualization, which has been subject to extensive research [6]. Several glyph-based approaches have been proposed to compare periods [7, 8, 9], especially to detect anomalous periods. Periodicity is also visualized through matrix-based or

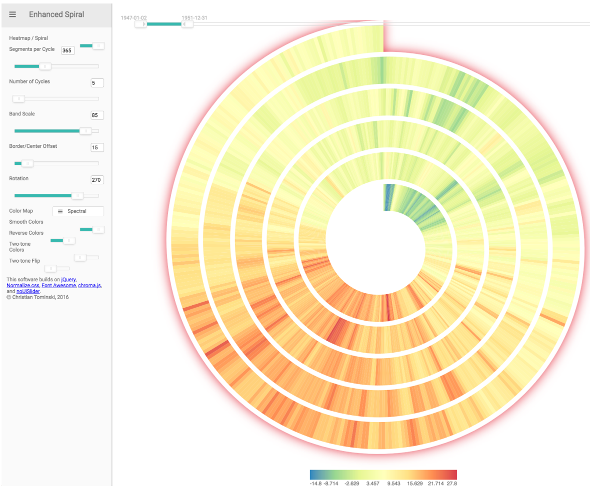

Figure 2: The enhanced interactive spiral visualization for time series data proposed by Tominski et al. [4] and later extended by Ceneda et al. [5] utilizes two-tone coloring to represent data values and provides visual guidance indicators to users to acquire well-suited period lengths. The intensity of the red glow surrounding the spiral indicates the fitness of a selected period length. The figure has been kindly provided by Davide Ceneda.

cycle-based layouts. Similar to our approach, cycle plots [10, 11] have been used to identify trends across periods through an element-wise comparison between periods. Time series are also frequently visualized through pixel-based techniques [12, 13], which encode data values similar to heatmaps. Circular layouts [14] arrange pixel-based visualizations in circular arrangements. Examples include Solar Plots [15], Kaleidomaps [16] and Ring Maps [17], which leverage repeated intervals for comparison and often incorporate additional context in the center region, which is a technique we also employ for our guidance donut approach that we place in the center of the spiral.

In the technique introduced by Weber et al. [2], the time series is mapped onto an Archimedean spiral in which each period corresponds to one turn. Each value of the time series is represented through a colored annular segment within the spiral. Tominski et al. later proposed to replace the coloring by a two-tone pseudo coloring [4] as seen in Figure 2. The resulting spiral is at the core of this work and is explained in Section 3. Time series spirals have further been explored in the 3D space [18, 19, 20] and in spatiotemporal context [21, 22]. Visualizing spirals in 3D comes with the drawback of occlusion, which requires additional interaction techniques as mitigation strategies. Such interaction techniques are often more complicated than the 2D version, which is why we opted for a 2D visualization in our approach.

The majority of the mentioned techniques can be explored in the TimeViz browser, an interactive collection of time series visualization techniques [23].

2.2. Exploration of Periodic Time Series

The interactive visual exploration of time series datasets is an integral part of many analysis tasks, such as discovering similar time series [24], constructing preprocessing pipelines [25], selecting features [26], detecting anomalies [27] or analyzing clusters [28]. In this broad context, our approach situates itself as a method to explore and discover groups of periodic subsequences following certain interestingness measures. Many of the mentioned related approaches follow the well-established overview+detail strategy or the related focus+context paradigm [29]. A well-known best practice approach in this regard is the cluster and calendar visualization [30].

When exploring spiral-based visualizations of time series data, a focus+context approach [31] is effective, as shown in the TimeSpiral system [32], which adaptively increases the width of rings in the user's focus while narrowing those farther from attention. This is similar to the SignalLens approach [33] and the ChronoLens technique [34], which display a dense time series through a line chart and provides the focus+context strategy in the form of a magnifying glass. The Lin-spiration system [35] employs two spirals for the context while showing the selected part as a line chart, following the analogy of an audio cassette.

2.3. Guidance in Time Series Applications

Although interactive visual exploration of datasets can provide valuable insights, a limitation is encountered with large datasets. The field of guidance seeks to provide visual cues to users to explore potentially interesting parts in a dataset [36, 37, 38]. Interestingness measures have been proposed to highlight and rank interesting patterns through visual cues [39, 40, 41, 42]. Our approach further builds on this idea by introducing the guidance donut, which visualizes several interestingness measures to guide users towards interesting spiral sector selections. The previous version of this work [3] introduced several interestingness measures, which we refine and revise in this work. Interestingness measures are not only used to highlight relevant sectors, but also to assist users in the configuration of the spiral by recommending suitable period lengths [5, 43, 44]. By contrast to previous work, our tool utilizes the regression model employed in our residual analysis workflow to recommend suitable period lengths. Finally, provenance is another key aspect in guidance, as it provides a history of previous analysis steps, which can be used as a reference to novice users [45]. Our tool also provides a provenance overview through which users can compare previous selections, which may act as guidance to novice users.

2.4. Residual Analysis of Regression Models

Analyzing residuals of a fitted regression model is an established practice in model diagnostics [46] to analyze secondary patterns [47]. The concept of *Interactive Model Prototyping* [48] combines residual analysis with interaction techniques. Users sketch the distribution, subtract the resulting estimates and subsequently explore the residuals, which can reveal secondary data features obfuscated by the dominant distribution. This concept has been further explored in the time series domain, in which the concept is especially useful to detect anomalies [49]. Our work utilizes the concept of Interactive Model Prototyping by providing the users a second level in the exploration of periodic time series data such that model residuals can be analyzed instead of the original time series to reveal secondary patterns. Regression-based modeling has been applied in further areas of Visual Analytics, such as simulation refinement using surrogate models [50], optimization of cooling systems [51] and understanding prediction biases of regression models [52]. Our approach uses a generic regression model, which is not bound to a specific use case and which can reveal secondary patterns in the data.

3. Design of the Analytical Workflow

In this section, we describe the technical details of the proposed analytical workflow. The following requirements were derived from a task analysis of analytical activities identified in the literature on periodic time series visualization. The resulting requirements align with the analytical workflow illustrated in Figure 1.

R1: Segmentation of Periodic Data: The system should support dividing a time series into equally sized intervals of length p, where p is a user-defined period length, to enable the alignment and analysis of the periods [2, 4].

R2: Removal of Dominant Periodic Component: The system should allow users to subtract the dominant periodic pattern from the data to reveal residual variations, anomalies, or secondary structures [48].

R3: Visual Comparison of Periods: The system should facilitate the side-by-side or overlaid comparison of the periods to support the identification of similarities, differences, and temporal shifts [11, 3].

R4: Efficient Subsequence Selection: The system should allow users to interactively select relevant sectors of interest in an efficient manner [3].

R5: Configurable Detail View: The system should offer three detail view modes that differ in aggregation level: an overview mode that aggregates all selected subsequences, a mid-level mode that displays each subsequence as one aligned spiral row, and a detail mode showing each subsequence as an individual line chart for fine-grained inspection [3].

R6: Provenance and Analytical Traceability: The system should track user interactions and selections to maintain a history of findings, enabling users to revisit, compare, and synthesize analytical outcomes [53].

3.1. Overall Workflow Design

The resulting workflow based on the described requirements can be seen in Figure 1. In the following we provide a high level overview before discussing the underlying technical details. This section introduces the overall conceptual workflow design.

In the first step, users select a suitable period length assisted through visual guidance clues (**R1**), which are based on the fitness of a regression model. The period length is required to properly align the spiral visualization such that each turn in the spiral corresponds to one period. This allows users to compare periods and detect patterns in the subsequences (**R3**), which corresponds to the *overview* component of the information seeking mantra [1]. In cases where the spiral does not provide sufficient indicators to discern patterns, the guidance

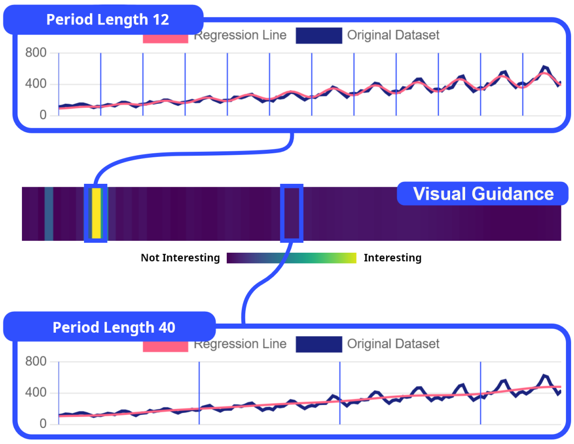

Figure 3: The visual guidance for the period length selection utilized model fitting to determine the fitness of each period. The result is a heatmap stripe seen in the middle, which indicates appropriate period selections to the user. Users select a period length by clicking on the heatmap, a hover effect provides an additional numeric indicator of the hovering period.

donut in the center of the spiral can assist users by providing interestingness measures, for example the average value per sector. Once the overall dominant periodic behavior of the time series has been understood, it can be subtracted through a fitted regression model to explore secondary patterns (**R2**). Users can switch anytime between analyzing the original time series (primary patterns) and the model residuals (secondary patterns).

The *detail* component is implemented by providing users an interaction technique to select interesting sectors. Users can hover the guidance donut to select segments within the spiral, which reflect a certain interestingness measure (**R4**). The selection is visualized in a linked view side-by-side with the spiral, called detail view. The segment is visualized in three visualization aggregation modes (**R5**) such that overall trends can be inspected (high aggregation) and individual elements (low aggregation).

Finally, each interaction with the spiral is logged and can be reviewed in a provenance view, which further allows users to compare previously made selections with each other (**R6**).

The next sections present technical details of the conceptual workflow starting with the period length selection guidance.

3.2. Selecting an Appropriate Period Length

As a first step, a period length p must be defined to segment the time series into equal intervals of length p. Choosing the period length manually can be challenging if the time series does not come with a natural or otherwise known periodicity. To assist users in the selection process, we provide a guidance approach by recommending suitable period lengths, which result in a well-aligned spiral. The basic idea is to fit a periodic model to the time series with a fixed frequency component given through a selected period and to use the resulting residual error as a guidance measure to recommend suitable periods. A basic model of periodicity is given as follows:

$$m(t) = \underbrace{\left(o_0 + o_1 \frac{t}{p}\right)}_{\text{Linear trend}} + \underbrace{\left(a_0 + a_1 \frac{t}{p}\right)}_{\substack{\text{Amplitude}\\ \text{modulation}}} \cdot \underbrace{\cos\left(2\pi \frac{t}{p} + \phi\right)}_{\text{Oscillation}}$$

The model consists of 3 components with 5 parameters in total. First, an overall linear trend of the time series is described through a basic linear function $o_0 + o_1 \frac{t}{p}$ with parameters o_0 and o_1. Next, the dominant periodic pattern is captured by a simple cosine function with parameter ϕ as the phase shift. The third component modulates the amplitude of the periodic pattern, also modeled as a linear function $a_0 + a_1 \frac{t}{p}$ with parameters a_0 and a_1. In our implementation, we fit this model with the Levenberg-Marquardt algorithm [54].

We evaluate the quality of the fit by computing the root-mean-squared error between the fitted model and the time series. We conducted experiments with other evaluation metrics as well, for example the R^2 score, which all yielded similar results. The result is a fitness score for each considered period length p, which we then visualize as a heatmap stripe as seen in Figure 3, also closely related to previous work by Ceneda et al. [5]. Users select a period by hovering and clicking on the heatmap stripe. Since computing the fitness scores can be expensive for large time series with over $10,000$ data points, the fitness scores are computed incrementally, starting with the smallest possible period and incrementing the period over time. As a result, users are not forced to wait until all possible periods have been examined, usually the first suitable period length is already sufficient.

This visual guidance method for period selection assistance demonstrates that data modeling in this scenario is a valid alternative to the Discrete Fourier Transform and the Chi-Square Periodogram seen in previous work by Ceneda et al. [5]. As an alternative, users are also able to select the period

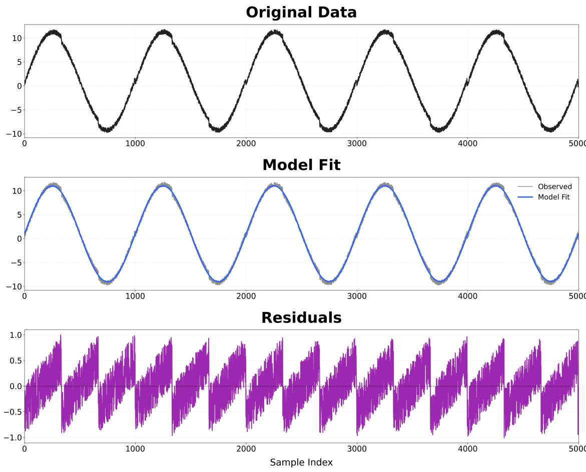

Figure 4: This example illustrates the model prototyping approach for a synthetic time series, which is a noisy sine wave with a secondary sawtooth pattern. To reveal the sawtooth pattern, model $m(t)$ seen in Section 3.2 can be fitted to the data and subsequently the residuals are visualized, which reveal another periodic pattern within the time series.

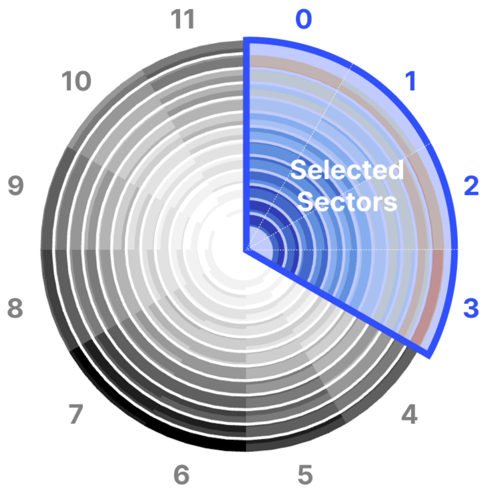

Figure 5: This didactic example illustrates our definition of a sector and their indexing. Each revolution in the spiral contains p data points, where p is the selected period length (in this example $p = 12$). Each element in one revolution is indexed clockwise, shown through the circular arranged integers. A sector is then the collection of elements having the same index within one revolution. Therefore, sectors are indexed clockwise, which enables a simple sector selection.

manually by defining a custom period length. This is beneficial in cases when the period length is already known, for example when the periodic time series exhibits a calendar-centric seasonality. In both cases, guided and manual, **R1** is fulfilled sufficiently. The resulting model $m(t)$ can be further used in a model prototyping scenario as will be shown in the next subsection.

3.3. Modeling the Dominant Periodic Trend

With a suitable period length selected, the dominant periodic trend is easily understood and modeled. However, the periodic trend may obstruct secondary patterns within the time series. Lampe and Hauser illustrated this scenario with a Gaussian point distribution [48]. The authors demonstrate that secondary features can be made visible by fitting a model and subsequently subtracting the model estimates. The remaining residuals reveal data features that were not captured by the dominant model, which can result in originally obstructed patterns becoming visible. This approach is called *Interactive Model Prototyping* and can be applied to time series as well – an illustrative example is shown in Figure 4. Our analytical workflow provides analysts the option to remove the dominant trend from the time series and use the result in all further exploration steps, which fulfills **R2**. The fitted model is additionally externalized by displaying the fitted model parameters.

3.4. Construction of the Time Series Spiral

The original time series or the model residuals from the previous step can be visualized appropriately through a time series spiral [2] to fulfill **R3**. As seen in Figure 6, the spiral is mapped on an Archimedean spiral with one revolution in the spiral corresponding to one period in the time series. Our implementation closely follows the design extensions to the original spiral proposed by Tominski et al. [4] using two-tone coloring [55] instead of a continuous color scale. The two-tone method colors each data point with a smooth split between two adjacent interval colors, where the boundary height within the slice precisely encodes the data value. A comparison between a continuous and a two-tone coloring is provided in Figure 6. According to Tominski et al. the two-tone coloring fosters an overview+detail inspection of the data within the spiral making it ideal for data exploration. However, the detail inspection can be hindered by large amounts of data, therefore we implement a sector selection to compare details within a periodic subset in a linked view seen in Figure 1. Additionally, the spiral does not start in the center, instead, it starts at a certain distance from the center. This has two advantages: First, the distortion close to the center is reduced and second, the resulting empty area can be elegantly utilized to provide visual guidance for

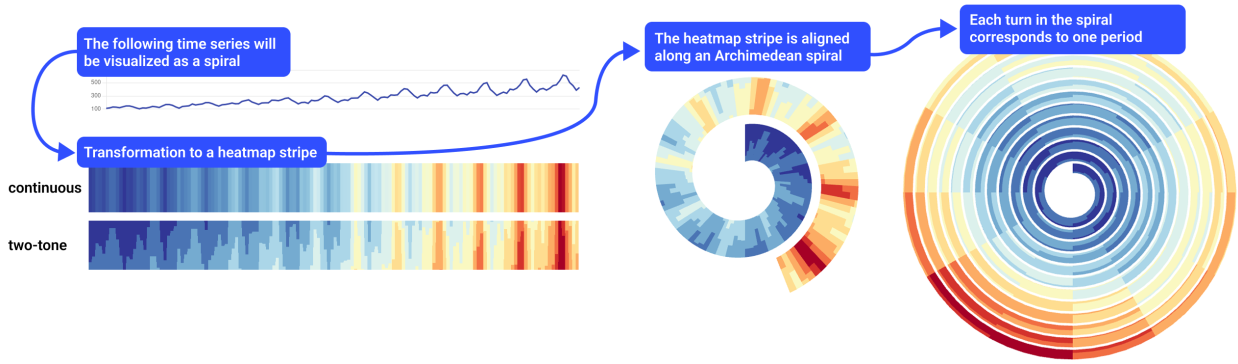

Figure 6: The spiral is constructed by first converting the time series into a heatmap stripe utilizing two-tone coloring [55] and afterwards bending the heatmap stripe into an Archimedean spiral. The number of data points within one revolution corresponds to the user-defined period length, such that periodic subsequences are aligned and can be easily compared.

the sector selection, which will be elaborated in the subsequent section.

3.5. Interactive Sector Selection with Visual Guidance

Selecting subsets within the data is a key requirement (**R4**). Since each turn in the spiral corresponds to one period, we can select subsets through a sector selection as seen in Figure 5. In the following, we introduce our design for a visual guidance element to assist users in the selection of sectors. The visual guidance is heatmap-based where each cell coloring corresponds to the interestingness of each possible sector selection. Each interestingness measure is a real-valued function of two variables: the sector's start index and end index, representing the sector's interestingness. Several examples are explained in Section 3.6. The following steps (also seen in Figure 7) elucidate our design process in constructing the guidance component to visualize the interestingness measures and provide reasoning for our design decisions. Users only see the final step of this process in the system.

Each data point within one period is assigned a sector index, allowing subsets to be selected via sector index ranges instead of radiant measurements. While selections can be characterized by a start index and end index, the large number of possibilities necessitates visual guidance. We designed this guidance by visualizing the selection space through a heatmap. Initially, using sector index ranges (Step 1, Figure 7) resulted in a discontinuous visualization due to the sharp break along the diagonal. Replacing the sector end index with sector width (Step 2) improved this, but traversing the start index still required users to manually account for the spiral's circular nature (modular arithmetic). Since the visualization should behave like a cylinder with a cyclic sector start index and bounded sector width, we transformed the heatmap into a toroidal (donut-shaped) visualization (Step 3).

The result is a direct mapping of each sector selection to the spiral, which we call the *guidance donut*. Each cell in the heatmap represents a possible sector selection and may indicate the interestingness of the sector through coloring. Possible choices for interestingness measures will be explained in the next subsection.

The guidance donut is placed in the center of the spiral as seen in Figure 8. Hovering the guidance donut highlights the corresponding selection in the time series spiral. Users control the sector width selection by moving along the radius (larger sectors close to the center and smaller sectors close to the border) and control the sector location in the spiral by moving along the arc of the guidance donut. The selection is further unaffected by the aggregation in the guidance donut such that every possible sector selection can be made. We provide a short demo video clip of this interaction technique in the supplemental material. Since the spiral is in the peripheral field of vision of the user when inspecting the guidance donut, changes in the sector selection due to hovering can be traced without needing to change context. This effectively removes the need of superpositioning another coloring layer on top of the spiral visualization. Overall, the proposed guidance donut provides a mapping of interestingness measures for sector selections to the spiral.

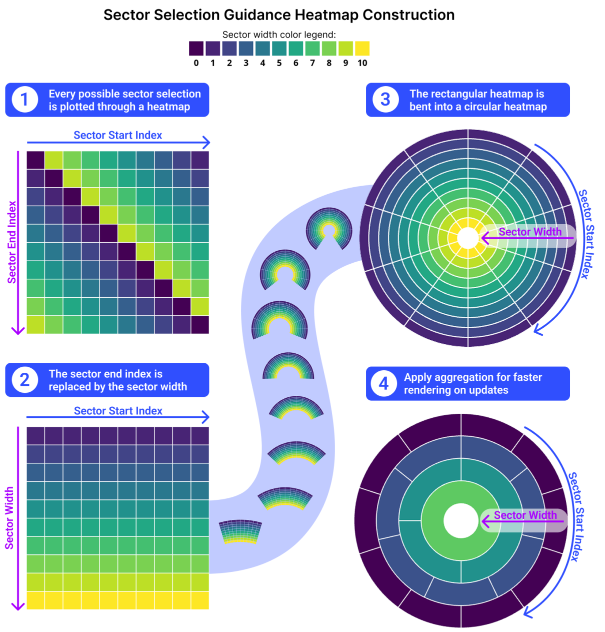

Figure 7: Construction of the guidance donut: sector selections are transformed from index-based coordinates (Step 1) to width-based (Step 2), and finally to a toroidal layout (Step 3) to establish a direct mapping to the spiral. The guidance donut is aggregated in Step 4 to increase rendering efficiency for large periods. The cells in this figure are colored by sector width, which serves as didactical placeholder for the interestingness measures shown in the remaining figures.

3.6. Interestingness Measures for Sector Selections

For each possible sector selection, we can compute several data characteristics, which can be visualized within the guidance donut and help provide cues on the interestingness of different sectors. Our approach is compatible with any interestingness measure following the definition introduced in the previous section. In the following, we describe three suitable interestingness measures we found useful for the data exploration task: average, trend and similarity. The interestingness measures are not only computed for the original time series, but also for model residuals. They are based on common analysis tasks faced by data analysts.

Sector Mean: As the name suggests, this interestingness measure is computed by computing the mean of all data points within the selected sector. The main usage of showing the average is to depict differences in sectors when the color distribution of the spiral is skewed through outliers or a global trend. Examples can be seen in Figure 11. The average can also be replaced by any other standard statistical measurements, for example median, variance or entropy.

Trend: The trend interestingness measure is computed by fitting a linear model on the data contained within a sector selection. We fit a regression model on the time series under a masking scheme that excludes observations outside the periodic segments. This can be viewed as regression with missing or masked data, where the mask corresponds to the selected periodic subsequences. The resulting gradient of the fitted model is used as the interestingness measure.

Similarity: Determining whether a sector in a time series is unique or recurring is a common time series analysis problem and is frequently solved using similarity search. In the case of a selected sector, the similarity interestingness measure indicates whether at least one of the subsequences is a recurring pattern. The interestingness measure is computed by using each subsequence as a query for a similarity search across the full time series, which is accomplished through *Mueen's algorithm for similarity search (MASS)* [56]. MASS uses the Euclidean distance and outputs a distance profile, indicating the distance to the query for each index. The distance profile is computed for each subsequence in the sector selection, i.e. the subset within each ring, and z-normalized for comparability reasons. The similarity interestingness measure is the minimum distance value of all distance profiles. Although the resulting interestingness measure is tailored towards recurring patterns, it can be adapted to anomalies as well by using the maximum instead of the minimum.

The three proposed interestingness measures are visualized through the guidance donut introduced in the previous subsection and provide visual cues for analysts during the exploration process. We demonstrate several use cases with each interestingness measure in Section 4.

3.7. Scalable Visual Guidance with Incremental Computation and Aggregation

Computing the interestingness measures within the guidance donut for a large number of period lengths can quickly result in infeasible running times since the number of possible sector selections grows quadratically with p^2. However, sector selections close to each other in the guidance donut often do not rapidly change and thus not every sector selection needs to be computed immediately. Therefore, the guidance donut can be quickly approximated

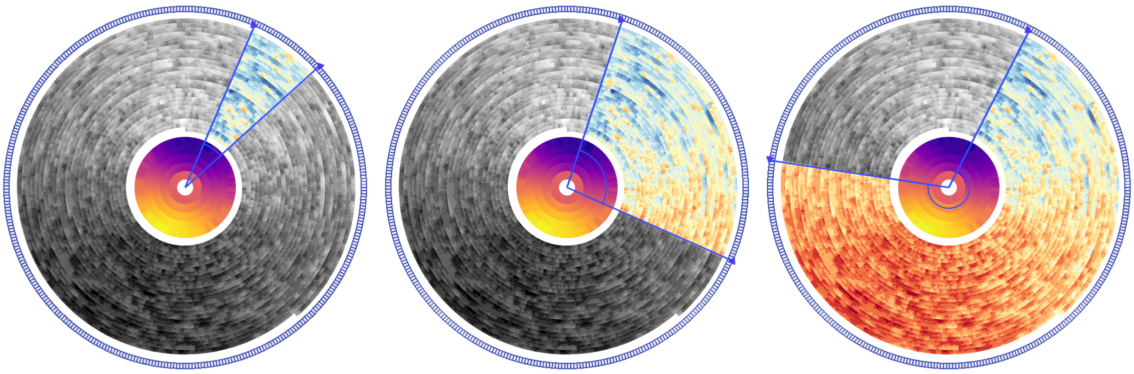

Figure 8: The three spirals demonstrate the sector-selection interaction. The guidance donut is placed in the center of the time series spiral to support the selection of sectors. Users hover the guidance donut to select a sector and refine it with sliders on the outer arc. Elements excluded from the selection are grayed out.

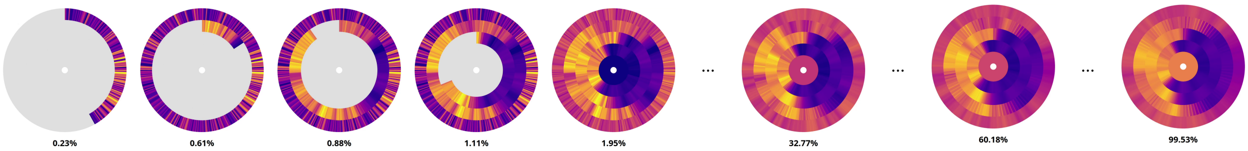

Figure 9: The guidance donut is computed incrementally following a round-robin sampling strategy. The first five donuts from the left demonstrate the round-robin strategy, which results in every guidance donut cell being sampled once and resulting in a reasonable approximation with a fraction of all possible sector selections. The remaining three donuts (from the right) show that with increased sampling progress, the approximation becomes smoother, but the overall pattern remains stable.

with only a fraction of all possible sector selections. Computing such approximations with only a fraction of the data is the core idea of *Progressive Data Analysis* (PDA) [57]. Ulmer et al. provide a taxonomy on progressive visualizations within their survey [58]. Following that, we will adapt the guidance donut to the PDA paradigm by processing the list of possible sector selections in chunks and updating the guidance donut through overwriting.

First, the guidance donut is aggregated as seen in Figure 7, which increases rendering efficiency on updates. The aggregation of the donut is uneven with more cells for smaller sectors and only one cell for sectors nearly encompassing the full spiral. Computing interestingness measures for wider sectors tends to exhibit less variation, therefore less cells are needed to adequately represent them. The zones in the aggregated guidance donut are visited in a “round-robin” strategy, which ensures that each zone has at least one computed sector selection after one iteration and thus, providing a quick approximation. During each iteration of zone visiting, one possible sector selection is chosen for each zone by traversing the space of possible sector selections in each zone. As a result, the guidance donut is incrementally computed, an example is shown in Figure 9. It shows that approximately 2% of possible sector selections are sufficient to provide a reasonable approximation to users. Afterward, the guidance donut remains stable in most zones.

3.8. Comparing Periodic Subsequences with the Detail View

Comparing sectors in a detail view is an important requirement (**R5**) and a natural interaction step after having selected a sector. Therefore, we provide a linked view consisting of three visualizations for a selected sector seen in Figure 1. Only one of the three visualizations is shown at a time, allowing users to switch between them. The three views show the selected subsequences at three aggregation levels. A stacked line chart offers the densest aggregation, allowing analysts to quickly identify an overall dominant pattern or trend. By contrast, separating all subsequences into separate line charts offers the lowest aggregation. Analysts may utilize it to inspect specific subsequences in detail. The middle ground is offered by a stacked heatmap visualization, which is essentially the selected sector of the spiral as a rectangular heatmap.

3.9. Comparing Spiral States Through the History

The final user task in the analytical workflow is concerned with the comparison of different spiral states (**R6**) such that important findings are recorded. To fulfill **R6**, we record the configuration of the analytical workflow on each user interaction. Each history element, as seen in Figures 12 and 13, shows the spiral with its selection state and the corresponding selected subsequences as a stacked line chart. Further, a fitted regression model shows the trend within the sector whose parameters are externalized for a quantitative comparison between selections.

4. Application Examples

We demonstrate the usefulness of our approach through several application examples, which focus on the period length selection, the sector selection guidance and the detection of secondary periodic patterns. The subsequent cases demonstrate scenarios for each interestingness measure in which the visual guidance provides beneficial assistance to users.

4.1. When do we need a guided period length selection?

In many cases, the most suitable period length is not known upfront. As an illustrative scenario, an analyst acquires a dataset containing IoT traffic data [59], which usually exhibit periodic patterns and finding a suitable period length is therefore crucial for further analysis steps. [1] To acquire an initial starting point, the analyst utilizes the visual period length guidance shown in Figure 10 and selects 1452 by clicking on a part of heatmap stripe indicating a suitable period fitness value. The resulting spiral is already well aligned, but not perfectly matching. The offset within the neighboring rings provides the analyst with the idea, that 1440 is a reasonable period length candidate as it resembles the total minutes for one day. The analyst therefore proceeds with a manual adjustment, yielding a perfectly aligned spiral.

Although the described scenario is hypothetical, it demonstrates the interplay between guided and manual period selection. The data-driven period length selection provides initial hints to experts, which can be used to refine the period length and incorporate semantic knowledge.

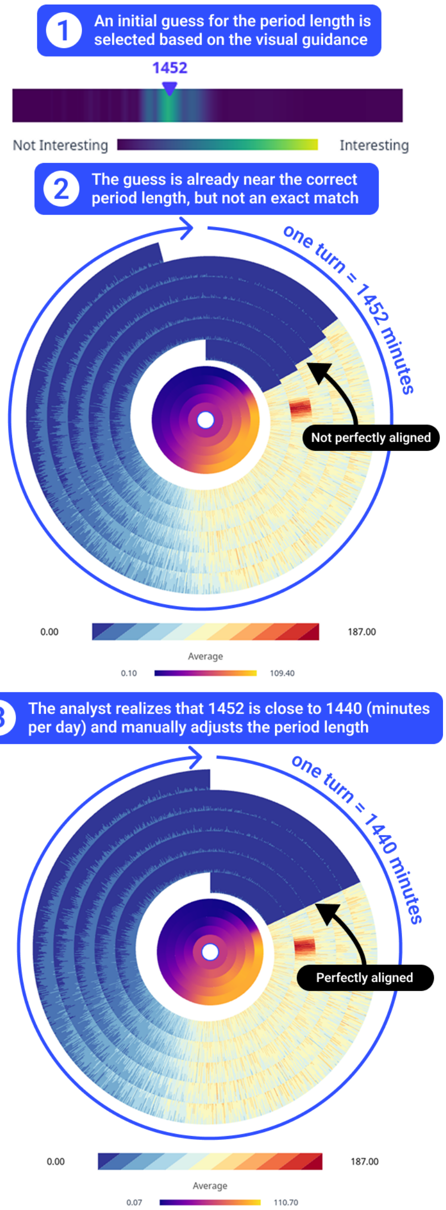

Figure 10: Visual guidance for selecting and refining a suitable period length: The heatmap-based period length fitness indicator (step 1) highlights candidate period lengths and enables the analyst to select an initial value (here: 1452). The corresponding spiral visualization reveals remaining misalignment (step 2), signaling that the chosen period length is close but not optimal. This offset suggests a semantically meaningful alternative (1440 minutes per day), which the analyst tests and confirms through manual adjustment (step 3). The figure illustrates how automated period-length guidance and manual interventions complement each other to achieve a well-aligned spiral visualization.

[1] Although the correct period length is stated on the source web page of the dataset, we assume the period length

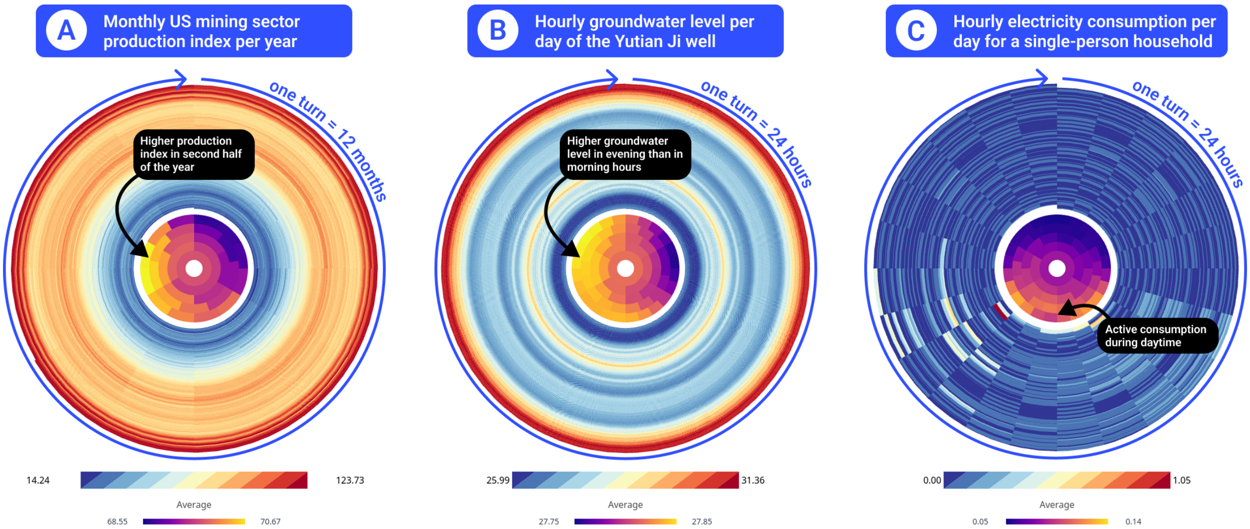

Figure 11: Three examples of the sector mean indicator highlight the benefit of guidance when differences between sectors are hard to interpret through the spiral. In all three cases, the average indicates subtle differences between sectors, which is easy to see due to the mapping between guidance donut and the spiral.

4.2. When do we need which guidance interestingness measure?

Whenever the difference between sectors within one period is only marginal, the spiral is hard to interpret since the contrast in coloring is reduced. This can happen in two common cases: Either the data contains outliers, which skews the color distribution, or the global trend causes a wide domain range such that the color space rather depicts the global trend instead of changes within periods.

4.2.1. Sector Mean Indicator

Figure 11 provides examples of cases with difficult sector comparison through the spiral visualization and shows three scenarios in which the sector mean indicator highlights differences between sectors. Figure 11, Scenario A shows the US mining sector production index, which has been explored as an example by Ramsay and Silverman [60]. As the authors discovered, the index exhibits a seasonal fluctuation. Although the global trend increases the difficulty in discovering this seasonal component in the spiral visualization, the average interestingness guidance highlights this seasonal fluctuation since the index is on average lower during winter months and higher in summer months. Scenario B in Figure 11 depicts the groundwater level of the Yutian Ji well [61]. Again, the difference between sectors is difficult to see due to the global trend, however, the guidance donut shows subtle differences in groundwater level between morning and evening hours. Scenario C in Figure 11 presents the electricity consumption of a single-person household within one month. The time series is dominated by discrete spikes in consumption rather than smooth, continuous variations. Although the spikes are well visible within the spiral, they cause secondary patterns to be occluded. However, the guidance donut shows the difference between active and inactive periods by averaging subsequences of the time series.

4.2.2. Trend Guidance

Although the spiral can point users towards sectors with differing patterns, estimating common time series properties such as trends can be challenging. To illustrate this challenge, Figure 12 shows the well-known air passengers dataset, which is included in the standard R dataset package and has been originally introduced by Box et al. [62]. When inspecting the spiral visualization, a higher passenger volume can be observed for summer months, however, estimating the growth rate can be difficult. Using the trend guidance, users get an impression of sectors with higher growth rates. By hovering and clicking on outstanding areas within

is not known for demonstration purposes.

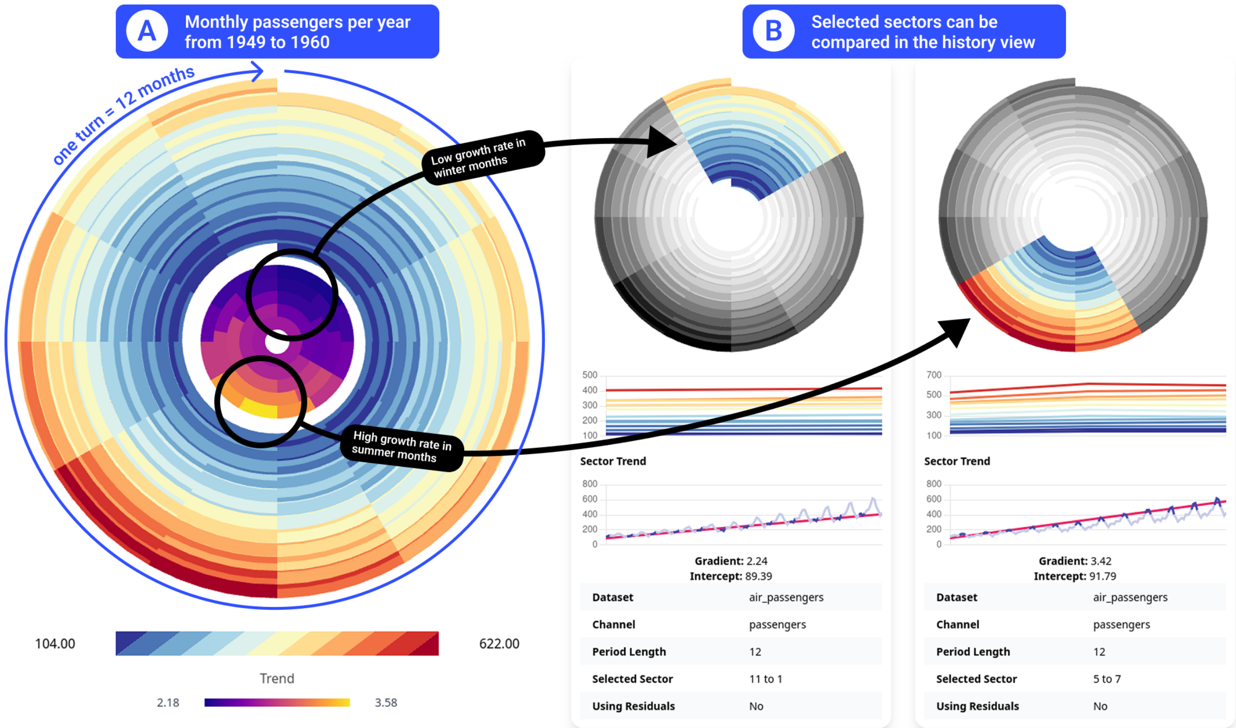

Figure 12: The air passengers time series exhibits a steeper trend in summer months compared to winter months, which is shown through the guidance donut with the trend guidance measure. The history view provides users with a comparison of these sectors together with an externalization of the trend parameters.

the guidance donut, users can inspect the sectors in the detail view and additionally inspect the linear regression fit, which serves the basis for the trend guidance. The linked view supports detail inspection while comparing selections with each other can be accomplished through the history view seen in Figure 12. Further, the parameters of each regression fit are exposed such that differing trends can be compared quantitatively as well and further externalized. In this case, the trend guidance shows that summer months are gaining popularity faster than winter months. Gaining insights in trend differences can be a valuable insight for enterprises to adapt strategies accordingly.

4.2.3. Similarity Guidance

Finding recurring patterns is a common time series analysis task in domains such as traffic and electricity consumption. Identifying such patterns through the spiral visualization can be hindered by challenges discussed in the beginning of Section 4.2, i.e. outliers skewing the color distribution or a global trend obfuscating secondary patterns. Figure 13 revisits the electricity consumption data of a single-person household and demonstrates the similarity guidance. The guidance donut seen in Figure 13A points users towards recurring patterns during noon and unique patterns during the night. In other words, the electricity consumption during the night exhibits a higher degree of randomness compared to consumption data during daytime. The same effect can be seen with traffic data shown in Figure 14, which shows the average vehicle speed at a busy highway. To ensure that the analysis captures consistent commuting behavior, only weekdays are considered in the dataset (excluding weekends and holidays). The dataset has been obtained from the *Caltrans Freeway Performance Measurement System* [63]. By inspecting the spiral visualization, two speed regimes can be observed. During daytime, the vehicle speed tends to be lower due to high traffic volumes while increasing during the night as the flow increases. The similarity guidance reflects this hypothesis as it points towards recurring patterns during nighttime and unique patterns during daytime, i.e. the average vehicle speed exhibits more randomness during daytime and less variance during nighttime.

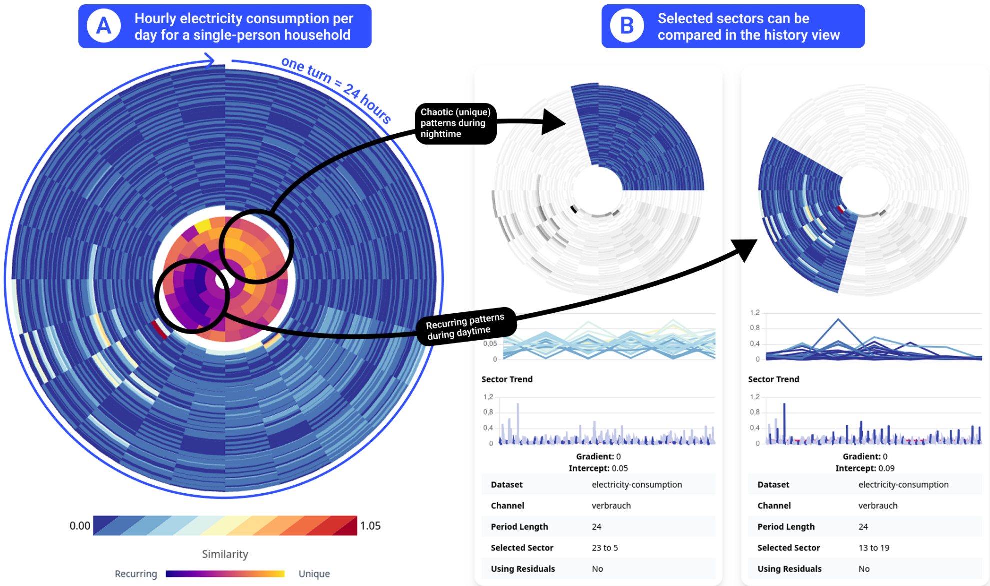

Figure 13: The similarity guidance highlights active and inactive electricity consumption periods in a single-person household. The history view allows users to compare these subsequences. Based on the stacked line charts, users identify consistent spikes in the active consumption period while the inactive period exhibits more randomness.

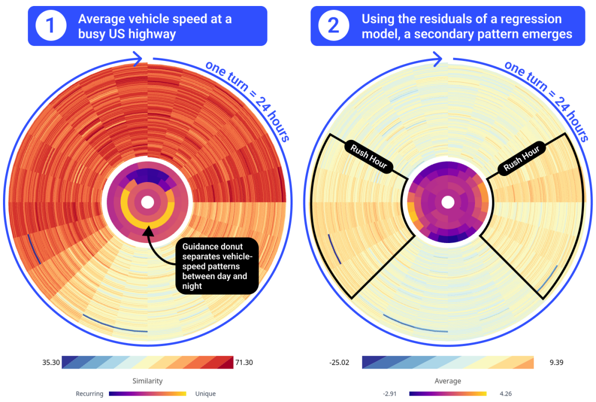

Figure 14: The average vehicle speed at a busy highway can be analyzed in two levels: (1) The similarity guidance highlights two speed regimes in the dominant periodic trend. (2) When analyzing the residuals of a fitted model, the effect of rush hours becomes apparent, which is also shown through the sector mean indicator.

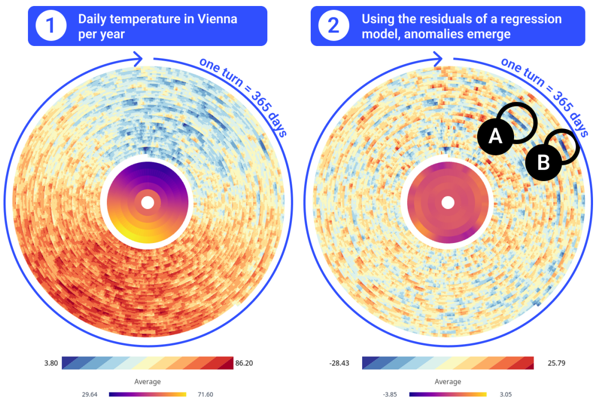

Figure 15: The spirals show the daily temperature in Vienna per year. Analyzing residuals of a fitted model can unveil anomalies in the data. Markers A and B correspond to cold waves, which have been overestimated by the model.

4.3. When do we need to analyze the residuals of a fitted regression model?

The insights of the previous subsections have been acquired by analyzing the dominant periodic trend. Although the main trend can unveil valuable information, it can obfuscate a secondary trend as well. Figure 4 presents an example to illustrate this scenario. Revisiting the traffic domain from the previous subsection, we observe a stable periodic trend as the result of the interplay between daytime and nighttime traffic. Due to the stability of this trend, regression model $m(t)$ introduced in Section 3.2 can be fitted to the data. Figure 14 shows the resulting residuals, which unveil a secondary periodic pattern. By inspecting the spiral visualization, it can be seen that the model consistently underestimates the vehicle speed at the transition of the two speed regimes observed in the dominant periodic trend.

These underestimates of the model might be caused by the daily rush hours and thus provide a predictable pattern.

Analyzing the residuals of a fitted regression model not only provides insight into secondary trends, but also in anomalies. Figure 15 shows the daily temperature in Vienna per year from 2000 to 2025. The data has been acquired from the GeoSphere Austria platform [64]. The markers A and B in Figure 15 highlight subsequences in the time series which have been greatly overestimated by the regression model, indicating unexpected patterns. Both events correspond to cold waves, which have been reported in newspaper articles.

5. Expert Interview

We conducted one joint expert interview with two domain experts from the field of Building Information Modeling (BIM) at *CAFM Systems GmbH*, both based in Austria. The experts' domain area includes acquiring and analyzing time series related to facility management, for example energy consumption rates, to detect potential saving opportunities.

Procedure. We conducted a semi-structured interview with the domain experts that lasted 90 minutes. First, we introduced the overall problem space, provided an explanation of the time series spiral together with our guidance donut indicator and gave a tour through the prototype. The experts were then tasked to explore several use cases demonstrated in Section 4.2 to acquire an initial understanding of the sector selection guidance mechanism. Afterwards, the experts were tasked to explore their own datasets, which include room temperature measurements and energy consumption rates in several building areas. The scope of the interview was limited to assess the use of the guidance donut.

Overall Assessment. After an initial onboarding, the experts found the spiral and the guidance donut straightforward to use. The detail view together with the three visual aggregation modes were helpful for the experts to inspect certain parts of a sector selection in detail. The experts used the stacked line charts (highest aggregation) to gain an overview of the selected sector and to quickly spot overall trends and outliers while the other two aggregation modes were mainly used to locate spotted phenomena. However, an extension to foster comparisons across multiple variables is necessary for the adoption of the guided spiral visualization in their field.

Findings and Lessons Learned. The guidance donut was particularly helpful to the experts to identify differences: *"For me, the guidance donut is meaningful. Thanks to the donut, I know in which area I have recurring consumption, so I can narrow down the sector. So the guidance donut is understandable for me."* The experts analyzed the dataset shown in Figure 13 and confirmed that the similarity guidance is reasonable and understandable: *"During nighttime, the consumption is significantly lower, you can already see that in the donut. You also know from your own consumption that there will be points [related to nighttime behavior] in there that are rather unique."* The experts afterwards examined their own datasets, which consisted of measurements acquired from a seminar room. In the case of fan coil activity to control the room temperature, they found the similarity measure helpful to provide a meaningful segmentation between daily patterns: *"The logical thing is that the least happens at night, and then during the day when the students are there, you always have unique control [patterns] of the fan coils."*

The visual guidance further helped in a specific scenario when the experts inspected one of their own datasets, which contained outliers. The spiral did not provide any insights, while the guidance donut with the sector mean indicator gave an immediate clue of the outlier location in the daily patterns. Although the accompanying line chart immediately reveals the presence of overall outliers, the question remains whether the outliers occur in certain hours of a day. The experts used the guidance donut to select the sector containing the outlier and through the detail view, they were able to locate it. The stacked line chart (highest visual aggregation) showed the overall presence of the outlier in the sector, while the experts found the stacked heatmap (low visual aggregation) particularly helpful to inspect the sector rings in detail, which shows that having multiple visual aggregation modes within the detail view is a useful enhancement in the workflow.

During the exploration of the datasets, the experts mentioned several opportunities for improvement. First, the spiral and the donut did not communicate time scales and units sufficiently, which created confusion during sensemaking. The experts further raised the need for sector comparison across datasets. Although the history view was found to be a useful first start, the experts need to compare sectors in detail across several datasets: *"This [radiator usage] would be interesting to compare with the measured temperatures to see why so much energy was consumed here."* The lesson learned in particular is the need for multivariate time series visualizations, which is an important requirement for a further adoption of the workflow in the field of building management. Finally, the experts mentioned a need for visualization onboarding as the visualization is complex and requires an explanation prior to using it.

6. Limitations and Future Work

Although the proposed approach is applicable to a broad spectrum of problems and datasets, several limitations exist, providing opportunities for future works.

6.1. Technical Limitations

The workflow assumes that the underlying time series exhibits a reasonably stable dominant periodic component. Time series with drifting or irregular periodicity may not align well in the spiral, limiting interpretability. The underlying regression model further assumes equal sized intervals defined by the period length, which can result in meaningless secondary patterns if this assumption is violated. Future works could therefore investigate advanced models to account for complex periodic patterns with irregular intervals.

The workflow is further limited to univariate time series. Weber et al. [2] proposed to plot each variable in a separate lane within the spiral to account for multivariate time series, however, this approach is only applicable when using monochrome coloring to ensure a stark contrast between variables within the spiral. Additional guidance measures tailored towards multivariate relations between variables would be required as well, for example a correlation-based measure. Providing additional guidance measures for univariate time series can also be a valuable extension in future works to support various analytical tasks, such as identifying clusters or correlations across subsequences.

Figure 11C presents a scenario in which the sector mean indicator is beneficial whenever outliers skew the color space. Despite the guidance donut proving a useful assistance in this case, outliers present an inherent limitation to the spiral visualization. Therefore, future works may investigate the use of model prototyping to remove understood outliers, such as traffic jams shown in Figure 14, to decrease the imposed skewness by outliers.

6.2. Evaluation Scope

The evaluation is limited to a small set of application examples and an expert interview. Broader studies across domains, user groups, and task types, as well as quantitative assessments of usability and task performance, are needed to better understand the workflow's general utility.

6.3. Visualization Onboarding

Although the sector selection presents a direct mapping between guidance donut and the spiral, the two concepts are both still complex for practitioners unfamiliar with spiral-based visualizations, which has been seen in the expert interview. Techniques from visualization onboarding [65] aim to educate users about the visualization, for example by animating a folding animation to achieve a spiral as seen in Figure 6. Future works could further investigate onboarding techniques for guidance visualizations, such as the guidance donut introduced in this work.

7. Conclusion

We presented a visual analytics workflow for exploring periodic time series and model residuals that combined spiral-based visualization, visual guidance, regression-based modeling, and progressive computation. We found the integration of these techniques particularly beneficial, as they compensated for each other's limitations: regression modeling reduced the visual dominance of primary cycles and enabled residual analysis, visual guidance lowered the interaction cost of identifying informative subsequences and suitable period lengths, and progressive computation alleviated the computational cost of guidance by providing early and stable approximations. The guidance donut exemplified this integration by enabling scalable and responsive exploration despite the large space of possible selections. We concluded that coupling guidance, modeling, and progressive computation is a promising design strategy for visual analytics systems operating on large periodic time series.

Supplemental Material

We provide the source code of our prototype through a public GitHub repository under MIT license: https://github.com/julian-rakuschek/interactive-time-series-spiral. The repository further contains a link to a live demo.

Funding

The work by Julian Rakuschek and Tobias Schreck was partially funded by the Austrian Research Promotion Agency (FFG) within the framework of the flagship project ICT of the Future PRESENT, grant FO999899544. Parts of this work have been done in the context of CEDAS, i.e., the Center of Data Science at the University of Bergen, Norway.

Acknowledgment

Our thanks go to Stefanie Stoppacher (TU Graz) for contributing to the initial version of this work [3]; to Kurt Battisti (A-NULL Development GmbH, Austria) for supplying the electricity consumption data used in Figures 11 and 13; to the domain experts at CAFM Systems GmbH (based in Austria) for their time and insights during the expert interview; and to the anonymous reviewers for their thoughtful and constructive feedback, which substantially strengthened the manuscript.

References

[1] B. Shneiderman, The eyes have it: a task by data type taxonomy for information visualizations, in: Proceedings 1996 IEEE Symposium on Visual Languages, 1996, pp. 336–343. doi:10.1109/VL.1996.545307.

[2] M. Weber, M. Alexa, W. Müller, Visualizing time-series on spirals, in: IEEE Symposium on Information Visualization, 2001. INFOVIS 2001., 2001, pp. 7–13. doi:10.1109/INFVIS.2001.963273.

[3] S. Stoppacher, J. Rakuschek, T. Schreck, Guided Visual Analysis of Time Series Data with Spiral Views and View Quality Measures, in: H.-J. Schulz, A. Villanova (Eds.), EuroVis Workshop on Visual Analytics (EuroVA), The Eurographics Association, 2025. doi:10.2312/eurova.20251097.

[4] C. Tominski, H. Schumann, Enhanced Interactive Spiral Display, in: Proceedings of the Annual SIGRAD Conference, Special Theme: Interactivity, Vol. 34, Linköping University Electronic Press, 2008, pp. 53–56.
URL https://ep.liu.se/ecp/034/ecp08034.pdf

[5] D. Ceneda, T. Gschwandtner, S. Miksch, C. Tominski, Guided visual exploration of cyclical patterns in time-series, in: Proceedings of the IEEE Symposium on Visualization in Data Science (VDS). IEEE Computer Society, 2018.
URL https://www.cvast.tuwien.ac.at/sites/default/files/bibcite/485/publik_272789.pdf

[6] W. Aigner, S. Miksch, H. Schumann, C. Tominski, Visualization of Time-Oriented Data, Second Edition, Human–Computer Interaction Series, Springer London, 2023. doi:10.1007/978-1-4471-7527-8.

[7] H. Tang, S. Wei, Z. Zhou, Z. C. Qian, Y. V. Chen, TreeRoses: outlier-centric monitoring and analysis of periodic time series data, Journal of Visualization 22 (5) (2019) 1005–1019. doi:10.1007/s12650-019-00586-1.

[8] F. Fischer, J. Fuchs, F. Mansmann, ClockMap: Enhancing Circular Treemaps with Temporal Glyphs for Time-Series Data, in: M. Meyer, T. Weinkaufs (Eds.), EuroVis - Short Papers, The Eurographics Association, 2012. doi:10.2312/PE/EuroVisShort/EuroVisShort2012/097-101.

[9] J. Suschnigg, B. Mutlu, G. Koutroulis, V. Sabol, S. Thalmann, T. Schreck, Visual Exploration of Anomalies in Cyclic Time Series Data with Matrix and Glyph Representations, Big Data Research 26 (2021) 100251. doi:10.1016/j.bdr.2021.100251.

[10] W. S. Cleveland, Visualizing Data, Hobart Press, Summit, NJ, USA, 1993.

[11] M. Bögl, P. Filzmoser, T. Gschwandtner, T. Lammarsch, R. Leite, S. Miksch, A. Rind, Cycle Plot Revisited: Multivariate Outlier Detection Using a Distance-Based Abstraction, Computer Graphics Forum 36 (3) (2017) 227–238. doi:10.1111/cgf.13182.

[12] T. Lammarsch, W. Aigner, A. Bertone, J. Gärtner, E. Mayr, S. Miksch, M. Smuc, Hierarchical Temporal Patterns and Interactive Aggregated Views for Pixel-Based Visualizations, in: 2009 13th International Conference Information Visualisation, 2009, pp. 44–50. doi:10.1109/IV.2009.52.

[13] M. Hao, U. Dayal, D. Keim, T. Schreck, Multi-Resolution Techniques for Visual Exploration of Large Time-Series Data, in: K. Museth, T. Moeller, A. Ynnerman (Eds.), Eurographics/ IEEE-VGTC Symposium on Visualization, The Eurographics Association, 2007. doi:10.2312/VisSym/EuroVis07/027-034.

[14] G. M. Draper, Y. Livnat, R. F. Riesenfeld, A Survey of Radial Methods for Information Visualization, IEEE Transactions on Visualization and Computer Graphics 15 (5) (2009) 759–776. doi:10.1109/TVCG.2009.23.

[15] M. Chuah, Dynamic aggregation with circular visual designs, in: Proceedings IEEE Symposium on Information Visualization (Cat. No.98TB100258), 1998, pp. 35–43. doi:10.1109/INFVIS.1998.729557.

[16] K. Bale, P. Chapman, N. Barraclough, J. Purdy, N. Aydin, P. Dark, Kaleidomaps: A New Technique for the Visualization of Multivariate Time-Series Data, Information Visualization 6 (2) (2007) 155–167. doi:10.1057/palgrave.ivs.9500154.

[17] J. Zhao, P. Forer, A. Harvey, Activities, ringmaps and geovisualization of large human movement fields, Information Visualization 7 (3-4) (2008) 198–209. doi:10.1057/palgrave.ivs.9500184.

[18] Z. Heliang, F. Sugimoto, C. Kato, K. Tsuchida, An application of 3D spiral visualization to the Uchida-Kraepelin psychodiagnostic test, in: 15th

IEEE/ACIS International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD), 2014, pp. 1–6. doi:10.1109/SNPD.2014.6888734.

[19] C. Tominski, P. Schulze-Wollgast, H. Schumann, 3D information visualization for time dependent data on maps, in: International Conference on Information Visualisation, 2005, pp. 175–181. doi:10.1109/IV.2005.3.

[20] J. V. Carlis, J. A. Konstan, Interactive visualization of serial periodic data, in: Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology, UIST '98, Association for Computing Machinery, New York, NY, USA, 1998, p. 29–38. doi:10.1145/288392.288399.

[21] J. Struhár, P. Rapant, Spatiotemporal Visualisation of PS InSAR Generated Space–Time Series Describing Large Areal Land Deformations Using Diagram Map with Spiral Graph, Remote Sensing 14 (9) (2022) 2184. doi:10.3390/rs14092184.

[22] C. Tominski, H. Schumann, G. Andrienko, N. Andrienko, Stacking-Based Visualization of Trajectory Attribute Data, IEEE Transactions on Visualization and Computer Graphics 18 (12) (2012) 2565–2574. doi:10.1109/TVCG.2012.265.

[23] C. Tominski, W. Aigner, The TimeViz Browser – A Visual Survey of Visualization Techniques for Time-Oriented Data, https://browser.timeviz.net, version 2.0 (2023).
URL https://browser.timeviz.net

[24] M. Sachdeva, J. Burmeister, J. Kohlhammer, J. Bernard, LFPeers: Temporal similarity search and result exploration, Computers & Graphics 115 (2023) 81–95. doi:10.1016/j.cag.2023.06.009.

[25] J. Bernard, M. Hutter, H. Reinemuth, H. Pfeifer, C. Bors, J. Kohlhammer, Visual-Interactive Preprocessing of Multivariate Time Series Data, Computer Graphics Forum 38 (3) (2019). doi:10.1111/cgf.13698.

[26] L. Cibulski, T. May, B. Preim, J. Bernard, J. Kohlhammer, Visualizing time series consistency for feature selection, Journal of WSCG 27 (2) (2019) 93 – 102. doi:10.24132/JWSCG.2019.27.2.2.

[27] M. Steiger, J. Bernard, S. Mittelstädt, H. Lücke-Tieke, D. A. Keim, T. May, J. Kohlhammer, Visual Analysis of Time-Series Similarities for Anomaly Detection in Sensor Networks, Computer Graphics Forum 33 (3) (2014). doi:10.1111/cgf.12396.

[28] T. Schreck, J. Bernard, T. von Landesberger, J. Kohlhammer, Visual Cluster Analysis of Trajectory Data with Interactive Kohonen Maps, Information Visualization 8 (1) (2009) 14–29. doi:10.1057/ivs.2008.29.

[29] H. Hauser, Generalizing Focus+Context Visualization, in: Scientific Visualization: The Visual Extraction of Knowledge from Data, Springer, Berlin, Heidelberg, 2006, pp. 305–327. doi:10.1007/3-540-30790-7_18.

[30] J. Van Wijk, E. Van Selow, Cluster and calendar based visualization of time series data, in: Proceedings 1999 IEEE Symposium on Information Visualization (InfoVis'99), 1999, pp. 4–9. doi:10.1109/INFVIS.1999.801851.

[31] G. W. Furnas, A fisheye follow-up: further reflections on focus + context, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06, Association for Computing Machinery, 2006, p. 999–1008. doi:10.1145/1124772.1124921.

[32] D. Zhang, L. Zhu, C. Wang, L. Zhang, TimeSpiral, an enhanced interactive visual system for time series data, in: 2016 2nd International Conference on Information Management (ICIM), 2016, pp. 127–133. doi:10.1109/INFOMAN.2016.7477546.

[33] R. Kincaid, SignalLens: Focus+Context Applied to Electronic Time Series, IEEE Transactions on Visualization and Computer Graphics 16 (2010) 900–907. doi:10.1109/TVCG.2010.193.

[34] J. Zhao, F. Chevalier, E. Pietriga, R. Balakrishnan, Exploratory Analysis of Time-Series with ChronoLenses, IEEE Transactions on Visualization and Computer Graphics 17 (12) (2011) 2422–2431. doi:10.1109/TVCG.2011.195.

[35] E. Graells-Garrido, A. Jaimes, Lin-spiration: using a mixture of spiral and linear visualization layouts to explore time series, in: Proceedings of the 2012 ACM International Conference on Intelligent User Interfaces, IUI '12, Association for Computing Machinery, 2012, p. 237–240. doi:10.1145/2166966.2167006.

[36] D. Ceneda, N. Andrienko, G. Andrienko, T. Gschwandtner, S. Miksch, N. Piccolotto, T. Schreck, M. Streit, J. Suschnigg, C. Tominski, Guide Me in Analysis: A Framework for Guidance Designers, Computer Graphics Forum 39 (6) (2020) 269–288. doi:10.1111/cgf.14017.

[37] D. Ceneda, T. Gschwandtner, T. May, S. Miksch, H.-J. Schulz, M. Streit, C. Tominski, Characterizing Guidance in Visual Analytics, IEEE Transactions on Visualization and Computer Graphics 23 (1) (2017) 111–120. doi:10.1109/TVCG.2016.2598468.

[38] W. Han, H.-J. Schulz, Providing visual analytics guidance through decision support, Information Visualization 22 (2) (2023) 140–165. doi:10.1177/14738716221147289.

[39] D. Sacha, M. Kraus, J. Bernard, M. Behrisch, T. Schreck, Y. Asano, D. Keim, SOMFlow: Guided Exploratory Cluster Analysis with Self-Organizing Maps and Analytic Provenance, IEEE Transactions on Visualization and Computer Graphics 24 (1) (2018) 120–130. doi:10.1109/TVCG.2017.2744805.

[40] J. Bernard, T. Ruppert, M. Scherer, T. Schreck, J. Kohlhammer, Guided discovery of interesting relationships between time series clusters and metadata properties, in: Proceedings of the 12th International Conference on Knowledge Management and Knowledge Technologies, i-KNOW '12, Association for Computing Machinery, New York, NY, USA, 2012. doi:10.1145/2362456.2362485.

[41] M. Luboschik, C. Maus, H.-J. Schulz, H. Schumann, A. Uhrmacher, Heterogeneity-based guidance for exploring multiscale data in systems biology, in: 2012 IEEE Symposium on Biological Data Visualization (BioVis), 2012, pp. 33–40. doi:10.1109/BioVis.2012.6378590.

[42] J. Suschnigg, B. Mutlu, M. Burgholzer, M. Bauer, T. Schreck, Cluster-Based Approach for Visual Anomaly Detection in Multivariate Welding Process Data Supported by User Guidance, in: Proceedings of the 30th International Conference on Intelligent User Interfaces, IUI '25, Association for Computing Machinery, New York, NY, USA, 2025, p. 325–340. doi:10.1145/3708359.3712076.

[43] F. de Luz Palomino Valdivia, H. A. H. Baca, I. S. Solis, M. A. Cruz, A. M. C. Valdivia, Guided Interactive Visu-

alization for Detecting Cyclical Patterns in Time Series, in: K. Arai (Ed.), Advances in Information and Communication, Springer Nature Switzerland, Cham, 2023, pp. 391–402. doi:10.1007/978-3-031-28076-4_29.
[44] M. Franke, S. Koch, Compact Phase Histograms for Guided Exploration of Periodicity, in: 2023 IEEE Visualization and Visual Analytics (VIS), 2023, pp. 191–195. doi:10.1109/VIS54172.2023.00047.
[45] T. Langer, R. Meyes, T. Meisen, Guided Exploration of Industrial Sensor Data, Computer Graphics Forum 43 (1) (2024). doi:10.1111/cgf.15003.
[46] R. D. Cook, S. Weisberg, Residuals and Influence in Regression, Chapman and Hall, New York, 1982.
URL https://hdl.handle.net/11299/37076
[47] B. W. Johnson, R. E. McCulloch, Added-Variable Plots in Linear Regression, Technometrics 29 (4) (1987) 427–433. doi:10.2307/1269453.
[48] O. D. Lampe, H. Hauser, Interactive Model Prototyping in Visualization Space, in: Proceedings of SIGRAD 2011, 2011, pp. 43–51.
URL https://ep.liu.se/konferensartikel.aspx?series=ecp&issue=65&Article_No=7
[49] N. V. Andrienko, G. L. Andrienko, A visual analytics framework for spatio-temporal analysis and modelling, Data Mining and Knowledge Discovery 27 (1) (2013) 55 – 83. doi:10.1007/s10618-012-0285-7.
[50] K. Matković, D. Gračanin, R. Splechtna, M. Jelović, B. Stehno, H. Hauser, W. Purgathofer, Visual Analytics for Complex Engineering Systems: Hybrid Visual Steering of Simulation Ensembles, IEEE Transactions on Visualization and Computer Graphics 20 (12) (2014) 1803–1812. doi:10.1109/TVCG.2014.2346744.
[51] R. Splechtna, M. Behravan, M. Jelović, D. Gračanin, H. Hauser, K. Matković, Interactive Design-of-Experiments: Optimizing a Cooling System, IEEE Transactions on Visualization and Computer Graphics 31 (1) (2025) 44–53. doi:10.1109/TVCG.2024.3456356.
[52] T. Mühlbacher, H. Piringer, A Partition-Based Framework for Building and Validating Regression Models, IEEE Transactions on Visualization and Computer Graphics 19 (12) (2013) 1962–1971. doi:10.1109/TVCG.2013.125.
[53] E. D. Ragan, A. Endert, J. Sanyal, J. Chen, Characterizing Provenance in Visualization and Data Analysis: An Organizational Framework of Provenance Types and Purposes, IEEE Transactions on Visualization and Computer Graphics 22 (1) (2016) 31–40. doi:10.1109/TVCG.2015.2467551.
[54] J. J. Moré, The Levenberg-Marquardt algorithm: Implementation and theory, in: G. A. Watson (Ed.), Numerical Analysis, Springer Berlin Heidelberg, Berlin, Heidelberg, 1978, pp. 105–116.
[55] T. Saito, H. Miyamura, M. Yamamoto, H. Saito, Y. Hoshiya, T. Kaseda, Two-tone pseudo coloring: compact visualization for one-dimensional data, in: IEEE Symposium on Information Visualization, 2005. INFOVIS 2005., 2005, pp. 173–180. doi:10.1109/INFVIS.2005.1532144.
[56] C.-C. M. Yeh, Y. Zhu, L. Ulanova, N. Begum, Y. Ding, H. A. Dau, D. F. Silva, A. Mueen, E. Keogh, Matrix Profile I: All Pairs Similarity Joins for Time Series: A Unifying View That Includes Motifs, Discords and Shapelets, in: 2016 IEEE 16th International Conference on Data Mining (ICDM), 2016, pp. 1317–1322. doi:10.1109/ICDM.2016.0179.
[57] J.-D. Fekete, D. Fisher, M. Sedlmair, Progressive Data Analysis, The Eurographics Association, 2024. doi:10.2312/PDA.20242707.
[58] A. Ulmer, M. Angelini, J.-D. Fekete, J. Kohlhammer, T. May, A Survey on Progressive Visualization, IEEE Transactions on Visualization and Computer Graphics 30 (9) (2023) 6447–6467. doi:10.1109/TVCG.2023.3346641.
[59] D. Hajhussein, M. Ibnkahla, IoT Time-Series Traffic Data: Smart City, eHealth, and Smart Factory (2024). doi:10.21227/5hyq-sw82.
[60] J. O. Ramsay, B. W. Silverman, Functional Data Analysis, 2nd Edition, Springer Series in Statistics, Springer, New York, NY, 2005. doi:10.1007/b98888.
[61] L. Yang, Groundwater level data minute value (2024). doi:10.21227/3hs2-3c27.
[62] G. E. P. Box, G. M. Jenkins, G. C. Reinsel, Time Series Analysis: Forecasting and Control, 3rd Edition, Prentice Hall, 1994.
[63] Caltrans / PeMS Development Group, Freeway Performance Measurement System (PeMS), Traffic performance database / data service (2025).
URL https://pems.dot.ca.gov/
[64] GeoSphere Austria, Messstationen Tagesdaten v2, Dataset (2024). doi:10.60669/gs6w-jd70.
[65] C. Stoiber, M. Wagner, F. Grassinger, M. Pohl, H. Stitz, M. Streit, B. Potzmann, W. Aigner, Visualization Onboarding Grounded in Educational Theories, Springer International Publishing, Cham, 2023, pp. 139–164. doi:10.1007/978-3-031-34738-2_6.