

Compact Phase Histograms for Guided Exploration of Periodicity

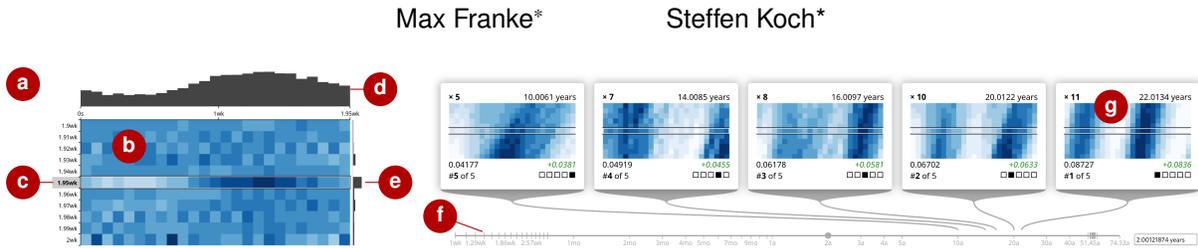


Figure 1: Our approach is centered around the phase histogram widget (a), which indicates the von Mises distribution [36] for given period lengths as a row-wise heat map (b). The duration of the period can be adjusted by vertically scrolling the heat map. The current period length is indicated by a frame (c) and visualized additionally as a bar chart (d). The quality measures for period lengths are visualized as a vertical bar chart (e). Bigger changes of the period length or picking a dedicated one is possible with the time slider (f). Other interesting period lengths that are fractional multiples of the current one are suggested above (g) the slider.

ABSTRACT

Periodically occurring accumulations of events or measured values are present in many time-dependent datasets and can be of interest for analyses. The frequency of such periodic behavior is often not known in advance, making it difficult to detect and tedious to explore. Automated analysis methods exist, but can be too costly for smooth, interactive analysis. We propose a compact visual representation that reveals periodicity by showing a phase histogram for a given period length that can be used standalone or in combination with other linked visualizations. Our approach supports guided, interactive analyses by suggesting other period lengths to explore, which are ranked based on two quality measures. We further describe how the phase can be mapped to visual representations in other views to reveal periodicity there.

Index Terms: Human-centered computing—Visualization—Visualization techniques

1 INTRODUCTION

Many datasets contain temporal information. For many years, visualization techniques have been used to understand developments, trends, and other temporal patterns in such time data series [1]. Such patterns include periodically occurring accumulations of events, measured values, or frequencies. In time series data, one or more attributes depend on time: $a(t)$. Accordingly, periodic behavior of period length τ can be seen as a similarity of the characteristics of $a(t)$ in the intervals $[t_0 + k\tau, t_0 + k\tau + \tau]$, $k \in \mathbb{Z}$. A special case are event data, where the temporal distribution of data points is of interest, rather than the distribution of values over time. Periodic behavior can be interesting for analyses; for instance, to predict future values of a data attribute, or to find hidden dependencies in the data. Several well-known automated procedures for detecting periodic effects exist, such as Fourier transforms [6], seasonal-trend decomposition based on Loess (STL) [10], and dynamic mode decomposition (DMD) [20, 27]. However, some of these methods require the period length of interest as an input parameter, or can be expensive regarding the compute time.

In certain situations, the *interactive* visual identification and exploration of periodic effects has benefits over automatic procedures and static visual approaches; for example, if the interplay of periodicity in events with their geographical location or other data at-

tributes is of interest. Different periodicities can be explored without making presumptions on interesting periods. In addition, important facets of the time-dependent data become visible immediately. Various approaches [5, 19, 30, 35, 37] map phase and time to positions in a grid- or spiral-based layout, revealing periodicity in an intuitive manner. We extend previous work with a compact, aggregated, and interactive representation (Fig. 1a) of time series data that reveals periodic behavior even for datasets with larger temporal extent. We offer guidance [9] in the form of suggestions (Fig. 1g) and quality measures [28, 36]. Since exploring the interplay with other data attributes can be interesting, our approach can be integrated within a larger, multiple-view visualization.

The contributions of the approach presented in this work comprise the introduction of a new composite widget that (i) helps users detect and explore periodic occurrences of aggregated quantitative information in long time series data, (ii) offers guidance to point users to potentially interesting periods exhibiting seasonal effects, and (iii) lets users pick suitable mappings including glyphs and color scales to understand periodicity in other views.

2 RELATED WORK

We briefly discuss related automatic methods for analyzing periodic behavior, but focus on interactive visual approaches. Various possibilities for representing cyclical temporal data exist [1, 5]. Hence, we limit our discussion here to approaches closely related to ours.

Automatic Periodicity Analysis. Fourier analysis [6] converts an input signal to the frequency domain. This requires a high sampling rate on the input signal, which can be expensive for event data, where it must be generated as a long, sparse histogram of the events. Fourier analysis also returns many false positives for non-sinusoidal periodicity. Cleveland et al. [10] introduced STL, which splits a time-dependent data signal into a seasonally recurring component, a linear trend, and a remainder. Constrained DMD (cDMD) [20] produces similar results using a different method. Cycle plots [7, 11] produce visualizations that also reveal such seasonal characteristics. While these methods produce good results even for noisy data, they are usually expensive to compute, and require prior knowledge on the frequency of the periodic behavior. STL also requires the period length to be an integer multiple of the sampling rate. Unconstrained DMD [27] does not require the period length as input, but has similar computation costs as cDMD [20], and does not always output the expected period length. Fourier analysis [6], DMD [27], and cDMD [20] find sinusoidal signal components; but struggle with other signal characteristics that could occur with event data. In contrast to these methods, we focus on interactive and exploratory analysis of periodic behavior in event data, often in the context of other data attributes.

*Both authors are with the University of Stuttgart, Germany.
E-mail: {Max.Franke,Steffen.Koch}@vis.uni-stuttgart.de

Visualization of Periodicity. Various works have explored non-linear layouts of data to reveal periodicity. Carlis et al. [8] and Weber et al. [37] proposed Archimedean spirals, where one turn along the spiral represents one period. Periodic behavior of that period length would then appear as lines or cones going radially outward from the center (Fig. 2i). The same effect can be reached with a line-wise, rectangular representation (Figs. 2e to 2g), where each line represents one period [19]. In both representations, spiral and rectangular, the period length can be adjusted interactively, by tightening the spiral or changing the aspect ratio of the rectangle, to find periodicities. Our approach focuses particularly on scalability, as well as guidance of users to interesting period lengths. We show these representations as detail views on demand. A special case of the rectangular representations are calendar-based layouts [21, 29, 35], which reveal periodic behavior in human-made time concepts on multiple scales, but are limited to fixed time concepts. Several visualizations utilizing concentric circles [2, 3, 16, 23] or stacks [22] also fit this category. Frey et al. [14] present a matrix representation which reveals self-similar patterns in time series data, but requires considerable screen estate.

Pinus view [30] shows an aggregated visual summary of a time series in a triangular fashion, where each position in the triangle corresponds with an interval and an offset in the temporal domain of the dataset. The residuality model by Van de Weghe et al. [34] describes a similar concept. These concepts, like ours, show a visual summary of the temporal domain. However, they are orthogonal to ours in that they consider aggregated intervals, rather than looking at repeating patterns. Closer to our approach is the work by Suschnigg et al. [31], who compare characteristics between periods via glyphs, or a matrix of anomaly scores. The approach by Ishii and Misue [17] is also closely related to our quality metrics, as they essentially visualize the von Mises [36] vector strength and direction for different data attributes. Recurrence quantification analysis [24, 25] also share similarities to our work. In contrast to these works, we focus on interactive discovery of interesting period lengths, and offer a more compact visual representation.

3 APPROACH

Our approach is centered around an aggregated view on all periods at once for a given period length, which we call the *phase histogram*. Periodic behavior for that period length is visible by non-uniform distribution of values (Figs. 1b, 1d, 2b, and 2c; Table 1). We calculate quality measures for different period lengths to guide towards interesting ones, focusing on fractional multiples of the current period length to balance out the challenges discussed in Section 3.2.

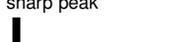
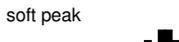
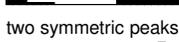
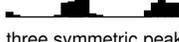
3.1 Phase Histogram

Our approach considers a time series of events S , $|S| = n$. The time series $S = \{t_i \mid i \in \{1, \dots, n\} \subset \mathbb{N}\}$ consists of points in time t_i . It has an extent from t_0 to t_1 : $t_i \in [t_0, t_1]$, where $t_0, t_1 \in \mathbb{R}$. For a user-selectable *period length* τ , each point in time t has a *phase* $\varphi(t, \tau) \in [0, 2\pi)$. The phase is the offset from the start of the period relative to the period length, where the first period starts at t_0 :

$$\varphi(t, \tau) = 2\pi \cdot \frac{(t - t_0) \bmod \tau}{\tau}$$

For a given τ , we then calculate a histogram with N bins over the phases of all time series events (Fig. 2a). The choice of N depends mainly on available space. It affects the granularity of the resulting phase histogram, as well as the scaling of the entropy estimation values. In the subsequent examples, we use $N = 25$. The shape of the histogram then reveals potential periodic behavior if one or some of the bins contain considerably more items than the others. Figures 2b, c, e to k, and Table 1 (rows 2–5) show some examples of histograms which represent periodic behavior. Multiples and fractions of the actual periodic behavior’s period length will also produce interesting-looking patterns. We consider such

Table 1: Example phase histograms alongside the two quality measures used: Shannon entropy [28] and von Mises vector strength [36]. For the vector strength, the distribution of the data points on the unit circle (blue) and their barycenter (red) are also visualized.

Phase histogram	Distribution	Vector strength	Entropy
uniform random 		0.079	4.5 Sh
sharp peak 		1	0 Sh
soft peak 		0.78	3.5 Sh
two symmetric peaks 		0.0037	3.3 Sh
three symmetric peaks 		0.01	4 Sh

multiples and fractions to evaluate whether they show clearer periodic behavior than the current period length τ (see Section 3.2). While our approach focuses on event data, it can also be employed to show periodic behavior of time-dependent data attributes (Section 3.3). Our representation can be understood as a vertical aggregation of the rectangular binned representations [19, 21, 29, 35], which reveals overarching patterns in all rows (see Fig. 2a).

3.2 Pre-calculation and Guidance

Interesting period lengths are those that match periodically reoccurring events in the data. In these cases, the phases of a larger-than-average number of events are similar, and the phase histogram is not uniform. Our approach calculates two quality measures that help decide which period lengths show promising patterns. The first measure is the Shannon entropy [28] of a phase histogram. Shannon entropy is high for more uniformly distributed histograms, and low for histograms with clear peak outliers (Table 1). The second measure is von Mises vector strength [36], which is a measure between 0 and 1, and is highest when the phases of all data points are the same. Von Mises vector strength is approximated by the distance from the center of a unit circle to the barycenter of all data points after these have been projected onto the unit circle at the angle of their respective phase. Both measures are suited for the task, since they generally indicate when the phase histograms are not uniformly distributed. They are also fast to compute, and so can be used as part of a highly interactive exploration approach. They have individual downsides, which is why we calculate both for a more nuanced view on the data: For Shannon entropy, the values quickly rise with only little noise. Von Mises vector strength cannot recognize whole multiples of the period length of periodic behavior, as the barycenter then moves to the center of the unit circle (Table 1, last two rows).

For a dataset, our approach initially pre-calculates the histograms and quality measures for a set of period lengths between a lower bound (e.g., 1 min) and an upper bound (the temporal extent of the dataset). Within this range, multiples of time units are sampled (1...59 min, 1...23 h, etc.). In addition, exponentially increasing period lengths are included. Data for additional period lengths can then be requested during analysis, and are calculated ad-hoc.

We also offer suggestions based on the current period length as a guidance. A key challenge here is to determine the actual frequency or period length of periodic behavior: A repeating pattern of period length τ looks repeating also for multiples of τ (Table 1 and Fig. 2f), but also for fractions thereof, if only an aggregation of all periods is considered (Fig. 2g). Furthermore, fractions of multiples (such as $4/3^{\text{rds}}$ of the actual period length) can look interesting as well in

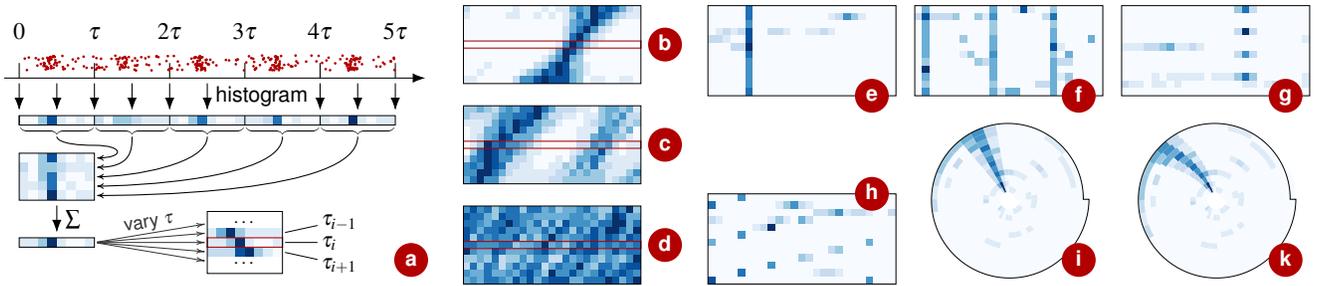


Figure 2: A schematic explanation (a) of our approach: Event data gets binned over the temporal domain. For the Cartesian representations [19, 21, 29, 35] the bins then get placed so that each row represents one period. Our approach shows an aggregated view on this, and varies the period durations slightly from row to row. The current period length is framed in red in the figure. Example patterns for periodic behavior in our approach for a sharp (b) and a less sharp (c) periodic pattern, and for uniform noise (d). As a comparison, periodic behavior in the Cartesian binned representation is shown for the actual signal period length, and for multiples (f), integer fractions (h, here $\frac{6}{5}$) thereof. Here, the periodicity manifests as vertical lines. Patterns appear as straight lines emanating from the center in the Archimedean binned representation (i) [8, 37], but are also visible if the period length is nearly right (k).

an aggregated view on the data (Fig. 2h). To facilitate discovery of better period lengths that are located at such factors, our approach samples a small set of fractions k/n , where n is a small natural number ≥ 2 and $k \in \{1, \dots, 2n - 1\}$. Small natural-numbered multiples of the current period length are sampled as well, and a selection of the best-ranking suggestions are returned based on available space. This sampling happens on the fly for the current period length τ .

3.3 Visual Representation

We propose a compact widget (Fig. 1a) that can be included as part of a larger, coordinated-views visualization. The widget shows the aggregated phase distribution for the currently selected period (Fig. 1d) length τ , as well as for the immediate neighborhood period lengths. This context (Fig. 1b) is shown as a pixel-based [18] visualization or heat map, where each row shows a period length. Each row visualizes the histogram as a line of colored rectangles, where the bin’s respective value is mapped to color. The period lengths are shown in ascending order from the pre-calculated data (Section 3.2 and Fig. 2a), with a context of n rows above and below the current row, which is located at the center of the heat map (Fig. 1c). By scrolling on the heat map with the mouse wheel, or clicking on rows, users can browse the rows, changing the current period length accordingly. By default, the color of the heat map cells represents the number of data items in that histogram bin. It is also possible to map this to other measures calculated on these subsets of data, such as the mean value or variance of another data attribute (shown in Fig. 1g for the number of sun spots per day).

To the right of the heat map, an interest measure is indicated for each row by a vertical bar chart (Fig. 1e). For the vector strength (Section 3.2), which is shown by default, the measure is mapped directly. For the entropy, the highest possible entropy (least interesting) is mapped to an empty bar, and the lowest possible entropy (most interesting) is mapped to a full bar. At the top of the widget, the current period length’s histogram is visualized again, as a bar chart (Fig. 1d). In applications where the phase is mapped to visual attributes in other views (see Section 3.4), this visual mapping is displayed in a legend (Figs. 3d and 3e) below the bar chart.

Our approach also contains a time slider (Fig. 1f), which represents the domain of possible period lengths on a logarithmic scale and can be used for larger adjustments to the current period length τ . Period lengths with the highest-ranked quality measure values are indicated as longer tick marks to guide users. During exploration, multiples and fractions for the current period length are sampled. The sampled period lengths are ranked by the selected quality measure, and the most promising ones are suggested as small thumbnails (Fig. 1g) above the time slider. Clicking on one such thumbnail changes the current period length to that of the thumbnail.

3.4 Visually Mapping the Phase

The presence of periodicity in the context of other data attributes can be of interest as well. For example, users might be interested in understanding whether certain events displayed on a map occur at recurring times (see Section 4). We have tested this mapping with position-based visualizations in our approach, but it could be extended to other visualization types. Our prototype implementation contains a scatter plot (Fig. 3a), which represents the spatial attribute of the data. To reveal spatio-temporal periodicity patterns, we propose two different types of mapping that can be applied to the markers of the scatter plot: color and shape (Fig. 3).

The first variant maps the phase of the data item to a position on a continuous color scale. Cyclical color scales, such as a rainbow color scale, can be a reasonable choice in this case. However, some datasets might favor color scales with a clear cut between the end and the beginning. The second variant maps to a visual mark with a parameterized shape. For our prototype, we offer a mapping to a moon phase-like mapping, or the angular rotation within $[0, \pi)$ of a rectangle, both cyclical. As acyclic mappings, we offer star shapes that morph into circles. Previous works [4, 12, 26] on glyph design offer additional options here. Future work might explore suitable shapes for representing cyclical data attributes.

Within our central widget, we show a legend (Figs. 3d and 3e) of this mapping. Especially for the acyclic mappings, it can be beneficial to change the offset of the mapping. Hence, the offset can be adjusted from within the legend.

4 CASE STUDY

We tested our approach with various synthetic and real-world datasets. Here, we demonstrate its applicability on tidal sea level data obtained from NOAA [32], with two specific datasets. Such data contains various periodicities pertaining to the rotation of the earth in relation to the moon, the orbit of the moon around the earth and the earth around the sun, but also to annual or perennial patterns stemming from meteorological phenomena. The interplay of these factors is not always obvious, demonstrating the need for exploratory analysis of the resulting patterns in the data.

The first dataset was generated from hourly measurements of the mean sea level (MSL) at one station located in Honolulu, Hawai’i [33]. The station measurements cover 118 years, from 1905 to the present day. We reduced this data to the measurements that exceeded a threshold value (+0.5m), leaving us with 5060 events. Our approach guided us to find the spring and neap tide periodicity at half a sidereal month (13.66d, Fig. 1a). We also could see and verify the 18.61-year nodal cycle discussed by Haigh et al. [15]. By mapping time of year to the x and year to the y coordinate in a scatter plot (Fig. 3c, days increase to the right, years to

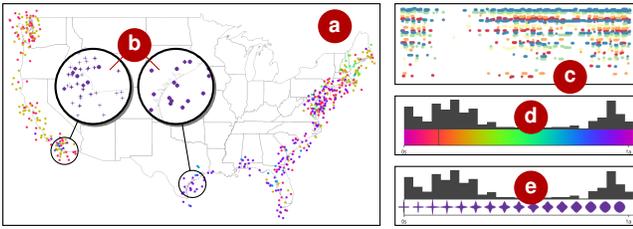


Figure 3: Example mappings of phase to color (a, c) or shape of visual marks (b) in a scatter plot. Periodic behavior related to the spatial aspect of the data is revealed by uniform areas. The mapping is shown in our widget (Fig. 1a) as a legend (d, e) that can be interactively adjusted to change the mapping of glyph or color to phase.

the top), we can also see an increase of events over the years, and a seasonal yearly component (denser and less dense vertical areas).

The second dataset contains the dates with highest extreme water levels for all observation stations, as provided by NOAA [32]. Over a record interval of 122.5 years, this data only contains 571 events, distributed over 74 unique stations. Hence, the observed patterns are not as expressive; additional preprocessing of raw data with more domain knowledge might yield better results here. Still, we can see some periodicity at the 18.61-year period length [15]. Furthermore, seasonal differences in events between the US east and west coast are clearly visible in the mapping of phase to color in a geographical scatter plot (Fig. 3a) for a one-year period length.

5 DISCUSSION

We implemented a web-based prototype for our approach. Its backend calculates the required data, as well as additional detail data on demand. We discuss initial results, and benefits and disadvantages compared to other visual representations and automated methods.

Suitable Datasets and Use Cases. Our approach can be applied to use cases where the periodic re-occurrence of events is of interest. Such events can also be generated from general time series data by determining points in time where time-dependent data values match some criteria. We demonstrated the viability of this strategy for finding periodic behavior in real-world data in our case study (Section 4). We offer a few sample datasets in our prototype, but also allow for arbitrary datasets to be loaded by the user.

Implementation Details and Scalability. Our web-based prototype visualizes pre-calculated data, and additional data is calculated on demand asynchronously in a backend. With this strategy, even for larger datasets with tens of thousands of events, interaction with the user interface is very responsive and smooth, with no or very few frame drops. Suggestions are calculated ad-hoc by the backend and typically appear within half a second after the last interaction. The thumbnails at the top (Fig. 1g) work well to guide users to interesting period lengths. We have found that the two measures supplement each others' drawbacks (Section 3.2) quite well. The mapping of phase to color or shape in the scatter plot has proven useful to detect periodic behavior that is local in other data attributes (Fig. 3). For performance, redrawing the scatter plot is the bottleneck, but we have found datasets with thousands or tens of thousands of data points to still render interactively. The implementation is available on GitHub and Zenodo [13].

Comparison to automated analyses. STL [10] and related methods [20] allow for the automated analysis of periodic behavior, and handle noisy data better. These methods require prior knowledge of the period length for which the periodic behavior should be analyzed. We compared mean computation time over ten runs of our method, STL, and DMD for two datasets: For a smaller synthetic dataset, STL took 227.2ms to calculate for *one* period length

of interest, and 3.45 s for the Hawai'i tide dataset [33]. DMD took 2.53 s and 1.63 s, respectively, to produce decompositions with six components and a sensible delay parameter, but did not find the expected periodic behavior. In contrast, our method took 470 ms and 589.5 ms, respectively, to calculate the phase histograms and quality measures for over 1800 sampled period lengths. The fairly high computation costs of methods like STL or cDMD prohibit extensive pre-calculations, or on-the-fly computation during smooth interaction. We also see that DMD struggles to produce good results for non-sinusoidal periodic behavior. However, these methods could still be utilized as a second step to verify results found interactively with our approach on demand. Fourier analysis [6] is another option, but not always suitable to real-world periodicity patterns. We argue that automated analysis is most useful for targeted use cases, but that exploratory analysis, especially in the context of other data attributes (Section 3.4), still requires more interactive methods.

Comparison to other visual representations. Cycle plots [7, 11] are a powerful way to visualize periodicity, but again presume knowledge of the periodic behavior. The calendar-based representations [21, 35] presume specific period lengths as well, but are often suitable for data based on from human-specified time concepts, such as months. We think other rectangle- [19] and spiral-based [8, 37] visualizations show periodic behavior in an intuitive manner, especially when period length can be adjusted. However; these representations do not scale well for larger time spans; and interactive use can often lead to large shifts, bad aspect ratios, and flickering effects. Our approach condenses these two-dimensional representations down to a one-dimensional aggregation. This opens up space for visualizing different data attributes to understand relations between them, and allows showing close-by period lengths as a context to the current one. This supports users in finding the best local period for periodic events. The aggregation hides periodicity that only appears for a part of the temporal extent, and multiples or fractions of the actual periodicity are not clearly identifiable as such (Section 3.2). Hence, we offer guidance, as well as detailed views as a tooltip to closer examine interesting period lengths (Fig. 2).

Open challenges. We demonstrated the general applicability of our approach with a prototype and different datasets, and plan to explore its integration into larger visualization systems. Understanding the best choices with respect to phase mapping (Section 3.4) in different scenarios, and determining the most appropriate visual mapping in terms of color scales and glyph designs to represent periodicity, is another interesting research direction we plan to pursue. The integration of other quality measures such as variance of the phase, domain-specific measures, or the combination of multiple measures, could be an interesting extension. Incorporating more complex automatic analysis methods, such as STL, could be another good addition to verify promising findings on demand.

6 CONCLUSION

We have presented a novel, aggregated visual representation of temporal data to explore periodicity. The time domain is mapped to the phase or remainder for a given and adjustable period length and visualized in a binned manner. Patterns that repeat with the same period length appear clearly in our representation. Our approach offers a more compact, interactive alternative to existing representations. It can be employed as a part of visualization systems with many views to let users configure the mapping of phases for depicting periodicity in other visual contexts.

ACKNOWLEDGMENTS

This work has been partially funded by the German Research Foundation (DFG) project # 314647693. We thank Tim Krake in particular; as well as Frank Heyen, Daniel Klötzl, and Kuno Kurzhals; for their valuable feedback and assistance.

REFERENCES

- [1] W. Aigner, S. Miksch, H. Schumann, and C. Tominski. *Visualization of Time-oriented Data*. Springer, London, UK, 2011.
- [2] E. N. Argyriou and A. Symvonis. Detecting periodicity in serial data through visualization. In *Advances in Visual Computing*, pp. 295–304. Springer Berlin Heidelberg, Jul 2012. doi: 10.1007/978-3-642-33191-6_29
- [3] K. Bale, P. Chapman, N. Barraclough, J. Purdy, N. Aydin, and P. Dark. Kaleidomaps: A new technique for the visualization of multivariate time-series data. *Information Visualization*, 6(2):155–167, Jan 2007. doi: 10.1057/palgrave.ivs.9500154
- [4] R. Borgo, J. Kehler, D. H. S. Chung, E. Maguire, R. S. Laramée, H. Hauser, M. Ward, and M. Chen. Glyph-based visualization: Foundations, design guidelines, techniques and applications. In *Eurographics STARs*, pp. 39–63. The Eurographics Association, Aug 2013. doi: 10.2312/conf/eg2013/stars/039-063
- [5] M. Brehmer, B. Lee, B. Bach, N. H. Riche, and T. Munzner. Time-lines revisited: A design space and considerations for expressive storytelling. *IEEE Transactions on Visualization and Computer Graphics*, 23(9):2151–2164, Sep 2017. doi: 10.1109/tvcg.2016.2614803
- [6] E. O. Brigham and R. E. Morrow. The fast Fourier transform. *IEEE Spectrum*, 4(12):63–70, Dec 1967. doi: 10.1109/mspec.1967.5217220
- [7] M. Bögl, P. Filzmoser, T. Gschwandtner, T. Lammarsch, R. A. Leite, S. Miksch, and A. Rind. Cycle plot revisited: Multivariate outlier detection using a distance-based abstraction. *Computer Graphics Forum*, 36(3):227–238, Jun 2017. doi: 10.1111/cgf.13182
- [8] J. V. Carlis and J. A. Konstan. Interactive visualization of serial periodic data. In *Proc. UIST*, pp. 29–38. ACM, Nov 1998.
- [9] D. Ceneda, T. Gschwandtner, and S. Miksch. A review of guidance approaches in visual data analysis: A multifocal perspective. *Computer Graphics Forum*, 38(3):861–879, Jun 2019. doi: 10.1111/cgf.13730
- [10] R. B. Cleveland, W. S. Cleveland, J. E. McRae, and I. Terpenning. STL: A seasonal-trend decomposition procedure based on Loess. *Journal of Official Statistics*, 7(1):3–73, 1990.
- [11] W. S. Cleveland. *Visualizing data*. Hobart Press, Sep 1993.
- [12] D. S. Ebert, R. M. Rohrer, C. D. Shaw, P. Panda, J. M. Kukla, and D. A. Roberts. Procedural shape generation for multi-dimensional data visualization. *Computers & Graphics*, 24(3):375–384, Jun 2000. doi: 10.1016/s0097-8493(00)00033-9
- [13] M. Franke. UniStuttgart-VISUS/periodic-time-vis. Zenodo, Jul 2023. v1.0.0. doi: 10.5281/zenodo.8164734
- [14] S. Frey, F. Sadlo, and T. Ertl. Visualization of temporal similarity in field data. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2023–2032, Aug 2012. doi: 10.1109/tvcg.2012.284
- [15] I. D. Haigh, M. Eliot, and C. Pattiaratchi. Global influences of the 18.61 year nodal cycle and 8.85 year cycle of lunar perigee on high tidal levels. *Journal of Geophysical Research*, 116(C6):C06025, Jun 2011. doi: 10.1029/2010jc006645
- [16] M. Hao, M. Marwah, S. Mittelstädt, H. Janetzko, D. Keim, U. Dayal, C. Bash, C. Felix, C. Patel, M. Hsu, Y. Chen, and M. Hund. Visual analytics of cyber physical data streams using spatio-temporal radial pixel visualization. In P. C. Wong, D. L. Kao, M. C. Hao, C. Chen, and C. G. Healey, eds., *SPIE Proceedings*. SPIE, Feb 2013. doi: 10.1117/12.2002948
- [17] T. Ishii and K. Misue. 2.5D extension of ChronoView for exploring periodic features of temporal data. In *Proc. IV. IEEE*, Jul 2018. doi: 10.1109/iv.2018.00014
- [18] D. Keim. Designing pixel-oriented visualization techniques: Theory and applications. *IEEE Transactions on Visualization and Computer Graphics*, 6(1):59–78, 2000. doi: 10.1109/2945.841121
- [19] R. Kosara. Spirals for periodic data. <https://eagereyes.org/techniques/spirals>, Aug 2011. Accessed: 2023-02-22.
- [20] T. Krake, D. Klötzl, B. Eberhardt, and D. Weiskopf. Constrained dynamic mode decomposition. *IEEE Transactions on Visualization and Computer Graphics*, pp. 182–192, Sep 2022. doi: 10.1109/tvcg.2022.3209437
- [21] T. Lammarsch, W. Aigner, A. Bertone, J. Gärtner, E. Mayr, S. Miksch, and M. Smuc. Hierarchical temporal patterns and interactive aggregated views for pixel-based visualizations. In *Proc. IV. IEEE*, Jul 2009. doi: 10.1109/iv.2009.52
- [22] T.-Y. Lee, A. Chaudhuri, F. Porikli, and H.-W. Shen. CycleStack: Inferring periodic behavior via temporal sequence visualization in ultrasound video. In *Proc. PacificVis. IEEE*, Mar 2010. doi: 10.1109/pacificvis.2010.5429602
- [23] G. C. Mariano, V. G. Staggemeier, L. P. Cerdeira Morellato, and R. da Silva Torres. Multivariate cyclical data visualization using radial visual rhythms: A case study in phenology analysis. *Ecological Informatics*, 46:19–35, Jul 2018. doi: 10.1016/j.ecoinf.2018.05.003
- [24] N. Marwan, M. C. Romano, M. Thiel, and J. Kurths. Recurrence plots for the analysis of complex systems. *Physics Reports*, 438(5–6):237–329, Jan 2007. doi: 10.1016/j.physrep.2006.11.001
- [25] T. Rawald, M. Sips, and N. Marwan. PyRQA—Conducting recurrence quantification analysis on very long time series efficiently. *Computers & Geosciences*, 104:101–108, Jul 2017. doi: 10.1016/j.cageo.2016.11.016
- [26] T. Ropinski, S. Oeltze, and B. Preim. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Computers & Graphics*, 35(2):392–401, Apr 2011. doi: 10.1016/j.cag.2011.01.011
- [27] P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656:5–28, Jul 2010. doi: 10.1017/s0022112010001217
- [28] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(4):623–656, Oct 1948. doi: 10.1002/j.1538-7305.1948.tb00917.x
- [29] P. Silva, C. Macas, E. Polisciuc, and P. Machado. Visualisation tool to support fraud detection. In *Proc. IV. IEEE*, Jul 2021. doi: 10.1109/iv53921.2021.00022
- [30] M. Sips, P. Kothur, A. Unger, H.-C. Hege, and D. Dransch. A visual analytics approach to multiscale exploration of environmental time series. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2899–2907, Dec 2012. doi: 10.1109/tvcg.2012.191
- [31] J. Suschnigg, B. Mutlu, G. Koutroulis, V. Sabol, S. Thalmann, and T. Schreck. Visual exploration of anomalies in cyclic time series data with matrix and glyph representations. *Big Data Research*, 26:100251, Nov 2021. doi: 10.1016/j.bdr.2021.100251
- [32] Tides & Currents, National Oceanic and Atmospheric Administration, United States of America. Extreme water levels. <https://tidesandcurrents.noaa.gov/est/>. Accessed: 2023-04-25.
- [33] Tides & Currents, National Oceanic and Atmospheric Administration, United States of America. Station home page of Honolulu, HI. <https://tidesandcurrents.noaa.gov/stationhome.html?id=1612340>. Accessed: 2023-04-26.
- [34] N. Van de Weghe, R. Docter, P. De Maeyer, B. Bechtold, and K. Ryckbosch. The triangular model as an instrument for visualising and analysing residuality. *Journal of Archaeological Science*, 34(4):649–655, Apr 2007. doi: 10.1016/j.jas.2006.07.007
- [35] J. J. van Wijk and E. R. van Selow. Cluster and calendar based visualization of time series data. In *Proc. InfoVis. IEEE*, Oct 1999. doi: 10.1109/infvis.1999.801851
- [36] R. von Mises. Über die “Ganzzahligkeit” der Atomgewichte und verwandte Fragen. *Physikalische Zeitschrift*, 19:490–500, 1918.
- [37] M. Weber, M. Alexa, and W. Müller. Visualizing time-series on spirals. In *Proc. InfoVis. IEEE*, Oct 2001. doi: 10.1109/infvis.2001.963273