Evaluating Alignment Approaches in Superimposed Time-Series and Temporal Event-Sequence Visualizations

Yixuan Zhang ⓘD*   Sara Di Bartolomeo ⓘD   Fangfang Sheng ⓘD   Holly Jimison ⓘD   Cody Dunne ⓘD
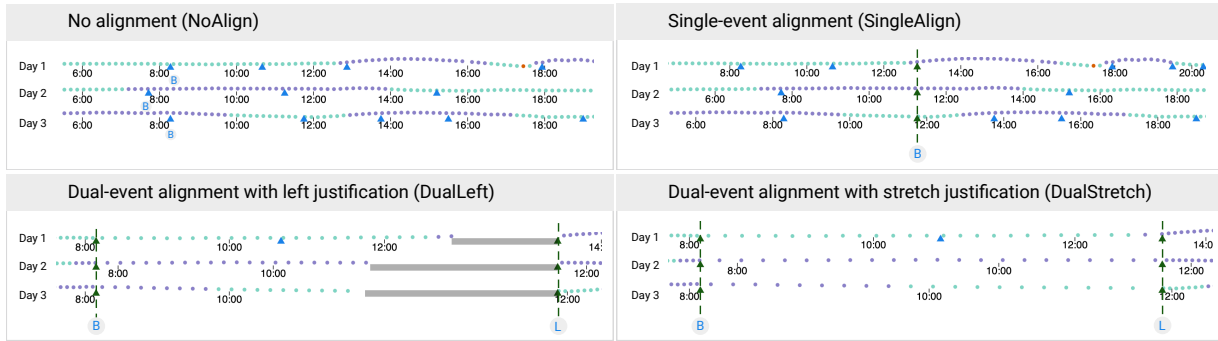
Northeastern University

Figure 1: User task performance was evaluated on four temporal event sequence alignment visualization approaches: no sentinel event alignment (NoAlign), single-event alignment (SingleAlign), dual-event alignment with left justification (DualLeft), and dual-event alignment with stretch justification (DualStretch). Stimuli were created using IDMVis [15] and Type 1 diabetes treatment data. Measures were task completion time, correctness, and error rate.

ABSTRACT

Composite temporal event sequence visualizations have included sentinel event alignment techniques to cope with data volume and variety. Prior work has demonstrated the utility of using single-event alignment for understanding the precursor, co-occurring, and aftereffect events surrounding a sentinel event. However, the usefulness of single-event alignment has not been sufficiently evaluated in composite visualizations. Furthermore, recently proposed dual-event alignment techniques have not been empirically evaluated. In this work, we designed tasks around temporal event sequence and timing analysis and conducted a controlled experiment on Amazon Mechanical Turk to examine four sentinel event alignment approaches: no sentinel event alignment (NoAlign), single-event alignment (SingleAlign), dual-event alignment with left justification (DualLeft), and dual-event alignment with stretch justification (DualStretch). Differences between approaches were most pronounced with more rows of data. For understanding intermediate events between two sentinel events, dual-event alignment was the clear winner for correctness—71% vs. 18% for NoAlign and SingleAlign. For understanding the duration between two sentinel events, NoAlign was the clear winner: correctness—88% vs. 36% for DualStretch—completion time—55 seconds vs. 101 seconds for DualLeft—and error—1.5% vs. 8.4% for DualStretch. For understanding precursor and aftereffect events, there was no significant difference among approaches. A free copy of this paper, the evaluation stimuli and data, and source code are available at osf.io/78fs5

Index Terms: Human-centered computing—Visualization—Empirical studies in visualization

---

*Corresponding author. E-mails: [zhang.yixua | dibartolomeo.s | sheng.f]@husky.neu.edu, [h.jimison | c.dunne]@northeastern.edu

1 INTRODUCTION

Time-series data visualization has been widely used in a variety of fields, including healthcare, business, engineering, and social media because of its ability to show how variables change over time. In recent years, visualizing discrete events on a time axis (temporal event sequence visualization) is also gaining popularity as this approach can reveal event patterns over time. Recent studies have started to composite these two types of visualization by overlaying the event sequence on top of time-series data visualization. As the superimposed view highlights the relationship between the events and the time-series data, the superimposition demonstrates the great potential to facilitate the process of making sense of the data [8].

Among various types of strategies to tackle the growing volume and complexity of data, alignment is a commonly adopted approach. With alignment, a visualization tool can put multiple instances of the same type of events at the same vertical or horizontal position, so that users can easily see the interaction between the aligned events and other events. In superimposed time-series and event-sequence visualizations, alignment can help discover the precursor (what triggers the event) and the aftereffect (what the event triggers) of the event, in the same visual space. Appropriately using superimposition would help ease comparison across the time-series data and related attributes displayed in the event-sequence visualization. However, applying alignment techniques in superimposed views might also increase more visual clutter and cognitive load. The usefulness of aligning by one sentinel event has been explored by Wang et al. [13]. However, previous evaluations focus on temporal event-sequence visualizations, while alignment techniques in superimposed visualizations are not sufficiently studied.

Moreover, the visualization community is no longer satisfied with single-event alignment (e.g., [13]), and dual-event alignment [15] has been proposed to explore the more complex temporal data interaction. IDMVis [15] exploits three different types of dual-event alignment to visualize the interaction of more than one event with the underlying value; the tool has received positive feedback from domain experts. However, the evaluation with domain experts cannot be easily generalized. A formal study that compares the effective-

ness of single-event and dual-event alignment is required before widely accepting the single-event alignment and dual-event alignment in superimposed time-series and event-sequence visualizations. Therefore, we designed a series of tasks that are tailored for the superimposed time-series and temporal event-sequence visualization and a controlled experiment to evaluate different temporal event alignment techniques, including no alignment, single-event alignment, and dual-event alignment.

2 RELATED WORK

2.1 Temporal Event Sequence Visualization

Visualization researchers have proposed a variety of techniques to deal with temporal event sequences, such as temporal folding and aligning. Temporal event alignment is a strategy to reduce pattern complexity by "aligning all records by the occurrence of a selected alignment" [4]. The purpose of applying aligning techniques is to help identify precursor, co-occurring, and aftereffect events [13]. Alignment techniques have been widely used in different domains, such as healthcare [6, 13], application log analysis [3], and sports [11]. For example, LifeLines2 [13] allowed users to align, rank, and filter medical events, such as first heart attack. In the application log analysis case study [3], alignment enabled analysts to identify the antecedents of the aligned events when the error messages appeared. Alignment can also be used to support cause-and-effect relation analysis. A similar case of analyzing Twitter user activity logs allowed users to pick a customized alignment point of interest—by default the beginning of events—to identify the sequence before and after the alignment point [14]. In these prior works, event alignment can be interacted with using a single view.

Other researchers explored the design space of applying alignment in composite views. Using alignment techniques, Care-Cruiser [6] supported comparison between the effects of the selected treatment plan among different patients as well as evaluating alternative treatment plans for a single patient. In our previous work by Zhang et al., we introduced the dual-event alignment in IDMVis [15] to help clinicians make treatment decisions for patients with Type 1 diabetes. Though our evaluation with domain experts was positive and gave us ample feedback, there has been a lack of thorough evaluation of the usefulness of using dual-event alignment versus single-event alignment. In this work, we aim to understand the effectiveness of different event alignment approaches in superimposed time-series and temporal event-sequence visualizations.

2.2 Tasks in Understanding Temporal Event Sequences

Temporal event-sequence visualizations can be used to support a wide range of sequence and timing analysis tasks. Along with popularizing the idea of sentinel event alignment for event-sequence analysis, the authors of Lifelines2 [13] described a set of six low-level domain-specific tasks. Here we re-frame these as three general tasks and present the results of their controlled experiment. Task 1: Given an event of interest and a subsequent time window, count aftereffect events of given event type within that window. In this case, participants viewing 5 records performed similarly in the non-aligned and aligned conditions. However, when viewing 20 records participants were on average 65% faster, more accurate, and had less error. Task 2: Given an event of interest and a subsequent time window, explore the most frequent aftereffect or co-occurrence events within that window. In one case, participants were on average 62% faster but there were no differences in accuracy or error. In others, the results were mixed. Task 3: Given *two* events of interest, count intermediate events that match a given event type. In this case participants were possibly more accurate but there were no differences in time or error.

The authors of DecisionFlow [5] described 12 high-level tasks focusing on three dimensions: interpretation of the temporal flow, understanding the properties of intermediate events, and reasoning
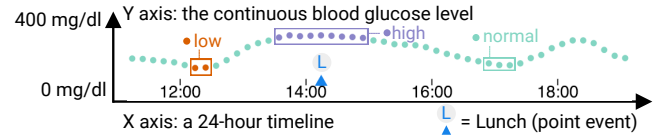

Figure 2: Legend of the main visual encodings used in the study.

about the changes of event sequences. Among these 12 tasks, those most related to sequence and timing analysis were (1) understanding temporal flow relationships (e.g., A before B) and (2) sentinel event (milestone) time comparisons. Plaisant and Shneiderman [12] further developed eight high-level tasks for event analytics across three high-level categories: heightening awareness, identifying data for further study, and understanding the impact of event patterns. This last category includes a particularly relevant task for sequence and timing analysis that harkens back to Lifelines2 [13] tasks: given an event of interest and optionally a surrounding time window, characterize the precursor (antecedent) or aftereffect (sequelae) events. However, most of the tasks prior researchers have considered have been too high-level, domain-specific, or unrelated to sequence and timing analysis to be used as guidance for developing tasks for our evaluation. Therefore, we first must build domain-independent task abstractions for temporal event sequence and timing visualization.

3 HYPOTHESES

The purpose of this study is to understand the effectiveness of event-alignment in superimposed time-series and temporal event-sequence visualizations. We form the hypotheses as follows:

H1: *The benefit of using single event alignment will be greater than no alignment for identifying precursor and aftereffect events and groups of events.* Rationale: This hypothesis is grounded in prior works (e.g., Lifelines2 [13]) which show that aligning by a sentinel event improved completion time and accuracy. We need to check if the previous results still hold in superimposed visualizations.

H2: *There is no benefit of using single- vs. dual-event alignment for identifying precursor and aftereffect events.* Rationale: Since dual-event alignment builds upon single-event alignment and is similar for a single sentinel event, we assume they perform similarly.

H3: *Dual-event alignment will be more useful than single-event alignment or no alignment for understanding intermediate events.* Rationale: This hypothesis is based on the IDMVis qualitative evaluation with domain experts which indicates that using dual-event alignment will increase performance [15].

4 COMPARATIVE EVALUATION

We conducted a controlled experiment on Amazon Mechanical Turk (MTurk) to assess the effect of sentinel event alignment approaches on participant task performance.

4.1 Stimuli and Tasks

We selected stimuli from our open source tool IDMVis [15] using its de-identified Type 1 diabetes treatment data. Example stimuli are shown in Fig. 1 for the four alignment approaches. The main visual encodings are described in Fig. 2, including point events (e.g., lunches) shown as labeled blue triangles ▲; time series data for a continuous glucose monitor shown as color-coded dots ● ● ● ; and the time and value axes. We selected tasks (detailed in Fig. 3) guided by our hypotheses, a synthesis of prior work in temporal event sequences discussed in Sect. 2.2, and taskonomies for Type 1 diabetes management [1, 15]. We designed a set of low-level tasks for evaluation considering the insufficiency of high-level tasks for sensemaking in the clinical context. In chronic illness management, sensemaking requires iterative data exploration and feature selection before formalizing hypothesis formulation and evaluation [10].

To examine co-occurrence of events, we categorize co-occurrence of events into ▲ disconnected and ● ● ● contiguous co-occurrence of events. Contiguous co-occurrence of events are events that are

Precursor	Aftereffect	Intermediate	Co-occurrence	Task abstraction	Mapped tasks in this study
X			●●●	Identify the contiguous co-occurrence of events **BEFORE/ AFTER** the aligned event	**Task 1**. How many days did the patient have <u>only high blood glucose values</u> within the 2 hours **BEFORE lunch**?
	X		●●●		**Task 2**. How many days did the patient have <u>at least one high blood glucose value</u> within the 2 hours **AFTER breakfast**?
	X		●●●		**Task 3**. How many days did the patient have within 2 hours **AFTER lunch** <u>at least three consecutive high blood glucose values</u>?
		X		Identify the duration **BETWEEN aligned events**	**Task 4**. Among the days shown, what is the length of the longest gap **BETWEEN lunch and dinner** (in hours)?
		X	●●●	Identify <u>the contiguous co-occurrence of events</u> **BETWEEN aligned events**	**Task 5**. How many days did the patient have <u>at least three consecutive high blood glucose values</u> **BETWEEN breakfast and lunch**?
		X	●●● ▲	Identify <u>the contiguous & *disconnected co-occurrence of events*</u> **BETWEEN aligned events**	**Task 6**. Examine the interval **BETWEEN lunch and dinner**. How many <u>*snack events coinciding with*</u> at least three consecutive low blood glucose values?

●●● Contiguous co-occurrence of events   ▲ Disconnected co-occurrence of events

Figure 3: Tasks for evaluating precursor, aftereffect, intermediate, and co-occurrence of events in superimposed visualizations.

adjacent to each other. Such events can include "interval events" (e.g., exercise events with duration logged) but also time series data (e.g., continuous blood glucose level) as the data points are contiguous. Disconnected co-occurrence of events are events that do not have constraints on adjacency. They are similar to the concept of "point events" (e.g., insulin injection). We designed questions to understand both ▲ disconnected (Task 6) and ●●● contiguous co-occurrence of events (Task 1,2,3,5,6).

We are also interested in assessing participant task performance at various visualization scales. Our visualization stimuli consist of one row for every day of data. We use a small scale of three days (as shown in Fig. 1) as well as a larger scale of 14 days. We asked each participant to answer 12 questions: 6 composite tasks×2 scales. Questions 1–6 used 3-day data visualizations and Questions 7–12 used 14-day data visualizations. In total, we selected a disjoint subset of the temporal event sequence for each question and created 48 visualization stimuli: 6 composite tasks×2 scales×4 conditions. To avoid ordering effects we randomized the order in which we presented the tasks to each participant.

4.2 Procedure

In this study we used a between-subjects approach to separate participants randomly but equally between four experimental conditions (shown in Fig. 1): no sentinel event alignment (NoAlign), single-event alignment (SingleAlign), dual-event alignment with left justification (DualLeft), and dual-event alignment with stretch justification (DualStretch). We did not test dual-event alignment with right justification as we believe that participant performance would mirror dual-event alignment with left justification.

Upon approval from the Institutional Review Board of our institution, we recruited participants from MTurk. Informed consent was gained before they accepted the HIT. Participants were asked to read through a brief introduction to the domain background (Type 1 diabetes) and a tutorial with example questions, answer 12 questions about alignment visualizations, and answer a questionnaire about their experience using alignment approaches and their demographics. Participants completed the study in 12 minutes 59 seconds on average. They were compensated with a base pay of $4 USD—average hourly rate $18.48—with a bonus of $0.60 for every correct answer after the sixth one. Participants were given a 10% and 20% bonus if they completed the study in less than 12 minutes and 10 minutes, respectively. We provide bonus rewards to encourage both fast and correct performance, as discussed in section 6. To ensure honest participation and sufficient understanding of the study and tasks, based on the pilot results we set the rejection criteria as participants (1) having less than or equal to one correct answer AND spending less than eight minutes in total, OR (2) less than one minute on the tutorial, OR (3) less than four minutes on questions.

4.3 Participants

Based on mean and variance estimates from a pilot study, we conducted a power analysis—Type I error rate $\alpha = 0.05$, power $1 - \beta = 0.8$—from which we determined we needed 108 participants. In total we recruited 123 U.S. participants and accepted 108. After applying a k-means clustering of the time spent on the questions and the numbers of correct answers to all accepted participants, we discovered that there was a clear clustering effect and identified 46 participants as "speeding" with the bound as four answers out of 12 questions and at least 8 minutes were spent on the questions. Therefore, we filtered out 46 participants leaving us 62 (16 for NoAlign, 18 for SingleAlign, 14 for DualLeft, and 14 for DualRight) for analysis. Among the 62 participants, the median age was 34 ($IQR = 15.5$).

4.4 Measures and Analysis

Due to the composite nature of our questions, we analyzed participant performance separately for each question. For each question, we measured the completion time, correctness, and error rate. The completion time is the time, in minutes, a participant spent to complete the question. The correctness is a binary measure, measured by whether or not the participant got the right answer. If the participant answered the question correctly, correctness was recorded as 1, otherwise treated as 0. We report error rate (ER) using $ER = |O_p - O_t|/O_t$, where the O_p and O_t represent the option that the participant selected and the correct answer, respectively. To select the appropriate statistical analysis plan, we first tested for normality with two methods qualitatively and quantitatively by using Q-Q plots and the Shapiro-Wilk test respectively. Our results show that our data does not have a normal distribution. Therefore, we used chi-square test for independence to analyze correctness and the Kruskal-Wallis H test as the nonparametric alternative to the one-way ANOVA to analyze time and error rate. We used the Bonferroni Adjustment for post hoc tests.

5 RESULTS AND DISCUSSION

We present the results using the filtered data ($n = 62$) which removes likely "speeding" actors as described in Sect. 4.3. Fig. 4 shows correctness and completion time for each combination of scale, task, and approach. For space considerations results with the unfiltered data and charts of error are relegated to supplemental material.

5.1 Precursor and Aftereffect Events

Tasks 1–3: We did not observe any significant differences ($p < 0.05$) between groups in terms of correctness and error rate. Participants were significantly slower using SingleAlign (45s) than NoAlign (31s) for the 3-day scale ($p = .02$) —but not for the 14-day scale. Contrary to our hypothesis, our results indicate that NoAlign was equal or better than SingleAlign for identifying precursor and after-
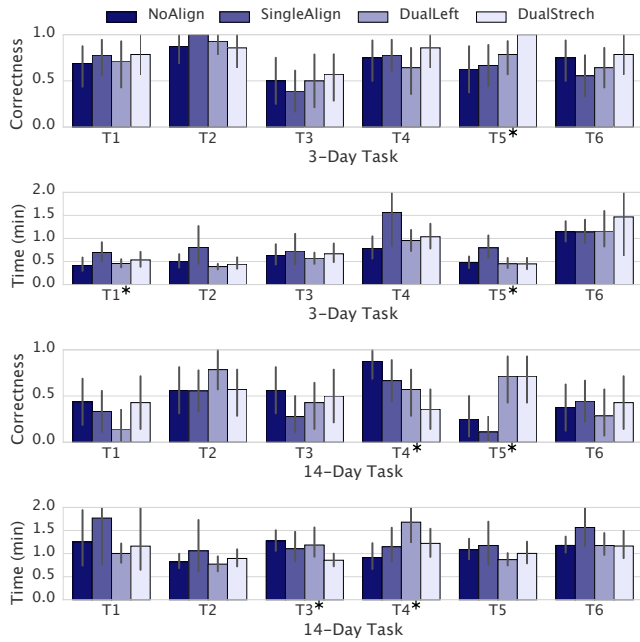
Figure 4: Correctness and time by scale, task, and approach with bootstrapped 95% confidence intervals, with asterisk (*) representing the results having significant ($p < .05$) differences.

effect events. These results seem to contradict the Lifelines2 [13] evaluation which showed an improvement in completion time.

However, there are several differences we must point out. The Lifelines2 [13] study explored point events, whereas we examine more complex superimposed visualizations of multiple data sources. Superimposed visualizations allow direct comparison in the same visual space, but may also cause high visual clutter and thus sacrifice readability [8]. We also visualize event types differently than LifeLines2. We display point events (e.g., breakfast) similarly using triangles. However, blood glucose levels are also treated as point events shown by dots that encode quantitative values using color and position. The complex Type 1 diabetes data and associated domain tasks [15] may have also contributed to task difficulty. Tasks that involve understanding contiguous co-occurrence events were complex, but have parallels with the tasks used to evaluate LifeLines2.

5.2 Intermediate Events

Task 4: For the 3-day scale we did not observe any significant differences in terms of completion time, correctness, and error rate. For the 14-day scale however, participants were more correct using NoAlign (88%) than DualStretch (36%) ($\chi^2 = 6.54$, $p = .01$), were faster using NoAlign (55s) than DualLeft (101s) ($p = .02$), and had less error using NoAlign (2%) than DualStretch (8%) ($p = .00$). This suggests that dual-event alignment is not an appropriate solution for identifying interval duration.

Task 5: For the 3-day scale, participants were faster using NoAlign (47s) than SingleAlign (94s) ($p = .05$), more correct using DualStretch (100%) than NoAlign (62%) ($\chi^2 = 4.43$, $p = .04$), and had less error using DualStretch (0%) than NoAlign (14%) ($p = .01$). For the 14-day scale, we did not observe any significant difference in terms of completion time. However, participants were more correct using DualStretch (71%) and DualLeft (71%) than NoAlign (25%) ($\chi^2 = 4.74$, $p = .03$) or SingleAlign (11%) ($\chi^2 = 9.79$, $p = .00$). These results support our hypothesis that dual-event alignment will be beneficial for understanding intermediate events, as long as they are not based on interval duration (Task 4).

Task 6: We did not observe any significant difference in terms of completion time, correctness, and error rate. One possible explanation from the participants was that the wording of Task 6 was

unclear. For example, one participant commented that *"[the] wording was confusing, because I thought it meant that both of those things happened between the two meals, not simultaneously, so I got that one wrong unfortunately."* However, we had not observed any difficulties in interpretation of the tasks in our pilot studies. One important take-away from this study is that the design of user study tasks requires substantial design and usability iteration.

6 REFLECTIONS AND DESIGN IMPLICATIONS

In this controlled experiment on MTurk, we measured task completion time, correctness, and error rate to evaluate the effect of different sentinel event alignment techniques in superimposed time-series and temporal event-sequence visualizations. We highlight several design implications for using alignment techniques, as well as design considerations when designing relevant tasks for evaluation.

First, consider using dual-event alignment for exploring intermediate events, regardless of the visualization scale (e.g., number of days). In most cases, DualStretch performed slightly better than DualLeft. However, dual-event alignment might not be most suitable for identifying the duration between two events.

Second, when applying alignment techniques, the nature of the composite visualizations should be taken into consideration. The complexity of a superimposed time-series and event-sequence visualizations is likely to add cognitive burdens to users. We need to further evaluate the design space of superimposed time-series and temporal event-sequence visualizations such as showing categorical tags or types of events through alignment alone, using categorical shape, color, and position. Also, we believe it is necessary to perform a more systemic evaluation on temporal event sequence in composite visualization that also includes juxtaposed, integrated, overloaded, and nested views—in addition to superimposed visualizations [8].

Third, we suggest a call for action for the design of low-level tasks. We noted that the complexity of the low-level tasks related to superimposed visualizations might affect users' performance. When designing tasks for non-experts on crowdsourcing platforms, the appropriate balance between workload, ease of tasks, and domain knowledge remains unclear and thus requires further study.

Fourth, participant compensation requires thoughtful consideration. The monetary reward is the primary motivator for many workers [2] who may have faster response times than in lab studies due to optimizing money/time [9]. Performance-based payments, e.g. a correctness bonus, may increase work quality but also increase the amount of time workers spend [7]. In order to use time as one of our measures, we use a speed bonus to help counteract any increase in time spent. Separate bonuses could lead to a situation in which a participant favors optimizing one over the other. We did not detect any negative impact of our bonus setup but it merits further study.

7 CONCLUSIONS

We designed low-level tasks for superimposed time-series and event-sequence visualizations, and conducted a comparative evaluation on MTurk to understand the effect of different sentinel event alignment approaches. We found that dual-event alignment allows users to perform tasks of intermediate events more accurately and quickly, except for identifying the duration. The benefits of using single-event alignment were not significant, compared with no alignment. There is no significant difference in using single- vs. dual-event alignment for identifying precursor and aftereffect events in terms of task completion time, correctness, and error rate. We encourage visualization researchers to conduct further studies to assess the effect of alignment techniques in composite visualizations.

REFERENCES

[1] R. M. Bergenstal, A. J. Ahmann, T. Bailey, R. W. Beck, J. Bissen, B. Buckingham, L. Deeb, R. H. Dolin, S. K. Garg, R. Goland, I. B. Hirsch, D. C. Klonoff, D. F. Kruger, G. Matfin, R. S. Mazze, B. A. Olson, C. Parkin, A. Peters, M. A. Powers, H. Rodriguez, P. Southerland, E. S. Strock, W. Tamborlane, and D. M. Wesley. Recommendations for standardizing glucose reporting and analysis to optimize clinical decision making in diabetes: The Ambulatory Glucose Profile (AGP). *Diabetes Technology & Therapeutics*, 15(3):198–211, 2013. doi: 10.1089/dia.2013.0051

[2] J. J. Chen, N. J. Menezes, A. D. Bradley, and T. North. Opportunities for crowdsourcing research on Amazon Mechanical Turk. *Interfaces*, 5(3):1, 2011.

[3] Y. Chen, P. Xu, and L. Ren. Sequence synopsis: Optimize visual summary of temporal event data. *TVCG*, 24(1):45–55, 2018. doi: 10.1109/TVCG.2017.2745083

[4] F. Du, B. Shneiderman, C. Plaisant, S. Malik, and A. Perer. Coping with volume and variety in temporal event sequences: Strategies for sharpening analytic focus. *TVCG*, 23(6):1636–1649, 2017. doi: 10.1109/TVCG.2016.2539960

[5] D. Gotz and H. Stavropoulos. DecisionFlow: Visual analytics for high-dimensional temporal event sequence data. *TVCG*, 20(12):1783–1792, 2014. doi: 10.1109/TVCG.2014.2346682

[6] T. Gschwandtner, W. Aigner, K. Kaiser, S. Miksch, and A. Seyfang. CareCruiser: exploring and visualizing plans, events, and effects interactively. In *Proc. PacificVis*, pp. 43–50, 2011. doi: 10.1109/PACIFICVIS.2011.5742371

[7] C.-J. Ho, A. Slivkins, S. Suri, and J. W. Vaughan. Incentivizing high quality crowdwork. In *Proc. 24th International Conference on World Wide Web*, WWW '15, pp. 419–429, 2015. doi: 10.1145/2736277.2741102

[8] W. Javed and N. Elmqvist. Exploring the design space of composite visualization. In *Proc. PacificVis*, pp. 1–8, 2012. doi: 10.1109/PacificVis.2012.6183556

[9] R. Kosara and C. Ziemkiewicz. Do mechanical turks dream of square pie charts? In *Proc. BELIV Workshop: BEyond Time and Errors: Novel evaLuation Methods for Information Visualization*, pp. 63–70, 2010. doi: 10.1145/2110192.2110202

[10] L. Mamykina, E. M. Heitkemper, A. M. Smaldone, R. Kukafka, H. J. Cole-Lewis, P. G. Davidson, E. D. Mynatt, A. Cassells, J. N. Tobin, and G. Hripcsak. Personal discovery in diabetes self-management: discovering cause and effect using self-monitoring data. *Journal of biomedical informatics*, 76:1–8, 2017. doi: 10.1016/j.jbi.2017.09.013

[11] M. Monroe, R. Lan, H. Lee, C. Plaisant, and B. Shneiderman. Temporal event sequence simplification. *TVCG*, 19(12):2227–2236, 2013. doi: 10.1109/TVCG.2013.200

[12] C. Plaisant and B. Shneiderman. The diversity of data and tasks in event analytics. In *Proc. IEEE VIS 2016 Workshop on Temporal & Sequential Event Analysis*, 2016.

[13] T. D. Wang, C. Plaisant, A. J. Quinn, R. Stanchak, S. Murphy, and B. Shneiderman. Aligning temporal data by sentinel events: discovering patterns in electronic health records. In *Proc. CHI*, pp. 457–466, 2008. doi: 10.1145/1357054.1357129

[14] K. Wongsuphasawat and J. Lin. Using visualizations to monitor changes and harvest insights from a global-scale logging infrastructure at Twitter. In *Proc. VAST*, pp. 113–122, 2014. doi: 10.1109/VAST.2014.7042487

[15] Y. Zhang, K. Chanana, and C. Dunne. IDMVis: Temporal event sequence visualization for type 1 diabetes treatment decision support. *TVCG*, 25(1):512–522, 2019. doi: 10.1109/TVCG.2018.2865076