
Simplifying Overviews of Temporal Event Sequences

Matthew Louis Mauriello
HCIL
Dept. of Computer Science
University of Maryland
mattm@cs.umd.edu

Ben Shneiderman
HCIL
Dept. of Computer Science
University of Maryland
ben@cs.umd.edu

Fan Du
HCIL
Dept. of Computer Science
University of Maryland
fan@cs.umd.edu

Sana Malik
HCIL
Dept. of Computer Science
University of Maryland
maliks@cs.umd.edu

Catherine Plaisant
HCIL
UMIACS
University of Maryland
plaisant@cs.umd.edu

Abstract

Beginning the analysis of new data is often difficult as modern datasets can be overwhelmingly large. With visual analytics in particular, displays of large datasets quickly become crowded and unclear. Through observing the practices of analysts working with the event sequence visualization tool EventFlow, we identified three techniques to reduce initial visual complexity by reducing the number of event categories resulting in a simplified overview. For novice users, we suggest an initial pair of event categories to display. For advanced users, we provide six ranking metrics and display all pairs in a ranked list. Finally, we present the Event Category Matrix (ECM), which simultaneously displays overviews of every event category pair. In this work, we report on the development of these techniques through two formative usability studies and the improvements made as a result. The goal of our work is to investigate strategies that help users overcome the challenges associated with initial visual complexity and to motivate the use of simplified overviews in temporal event sequence analysis.

Author Keywords

Information visualization; temporal event sequences.

ACM Classification Keywords

H.5.2. Information interfaces and presentation (e.g., HCI): User Interfaces.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).

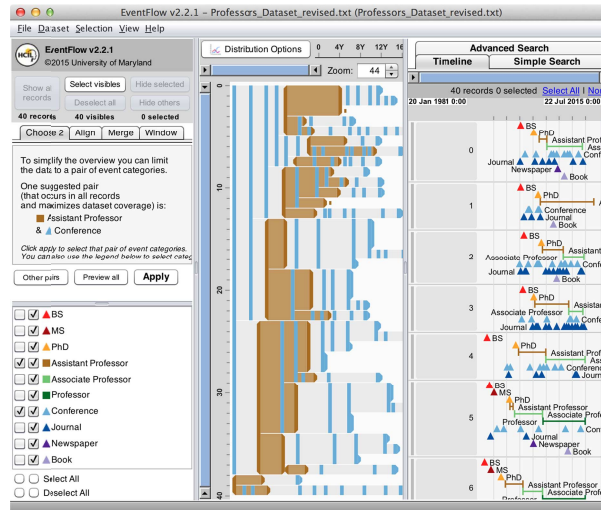
CHI'16 Extended Abstracts, May 07-12, 2016, San Jose, CA, USA

ACM 978-1-4503-4082-3/16/05.

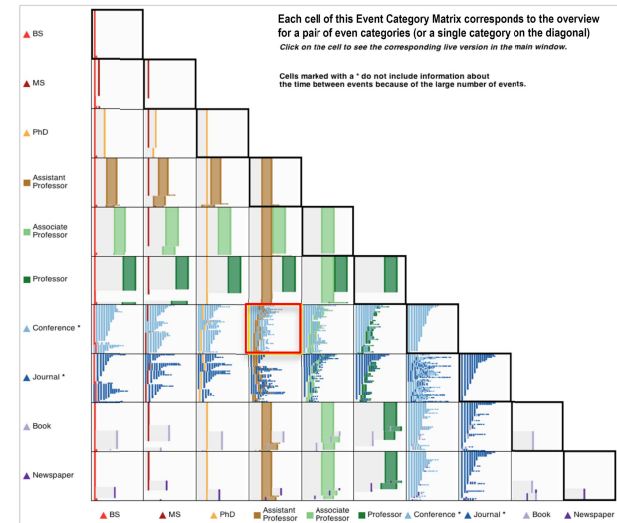
<http://dx.doi.org/10.1145/2851581.2892440>

Figure 1: The Choose2 panel (top left of 1a) as implemented in EventFlow recommends a pair of event categories to simplify the overview of all the records. On the right, 1b shows an Event Category Matrix with the overview of all pairs. In this example the user has applied the recommended simplification, which filtered out all event categories except Assistant Professor and Conference. This recommendation is shown in the Choose2 panel and highlighted in red in the Event Category Matrix.

See Appendix I below for a longer description of the Event Category Matrix (ECM).



(1a)



(1b)

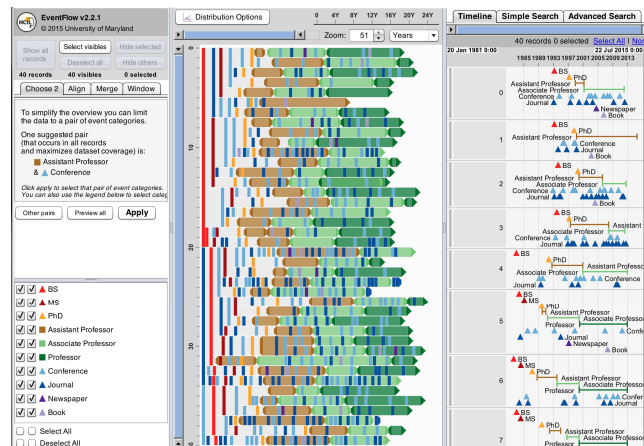
Introduction

Modern datasets are voluminous and varied in structure. This presents a challenge when trying to discover meaningful insights, which requires analysts to explore data, generate hypotheses, draw insights, and reach conclusions [19]. This paper focuses on the analysis of temporal event sequence data (e.g., electronic health records, activity logs) and how this data is visualized in aggregated display overviews. While early work on these displays demonstrated that analysts can discover meaningful relationships between temporal events [2,3,5,10,16,22], our observations from case studies [7] indicate that, regardless of skill level, analysts looking at a new dataset can have difficulty answering the question: "Where do I begin?"

When using aggregated display overviews for temporal event sequence analysis (e.g., EventFlow¹), analysts can become overwhelmed by the visual complexity of the overview because it summarizes all the patterns (i.e., includes all the event categories) found in the dataset. This tends to occur when the data is unfamiliar. To reduce this initial visual complexity in event analytics, we propose Choose2: a set of three techniques that provide simplified overviews based on the idea of starting with fewer event categories (i.e., 2 or even 1) while still representing a large number of records. The motivation for the techniques came from observing practicing analysts derive insights using simplification strategies in 18 case studies [7] and from Spotfire's automatic initialization of scatterplots [1].

¹ EventFlow: <http://www.cs.umd.edu/projects/hcil/eventflow/>

Figure 2: The professor dataset displayed in EventFlow might overwhelm novice and expert analysts alike when all event categories are shown. The dataset includes events in the careers of 40 professors: point events (e.g., time they received their bachelor or master degrees, published a journal or conference paper) and interval events (e.g., appointments as assistant or associate professor). Choose2 (top left) suggests starting with only two events.



The first Choose2 technique, designed for novice analysts who are new to event sequence analytics, automatically suggests a pair of event categories to begin analysis. The second technique suggests additional pairs using a set of 6 ranking metrics. The third technique was designed for expert analysts (*i.e.*, experienced analysts), and consists of an Event Category Matrix (ECM) displaying overviews of all event pairs and single event categories. This paper motivates these techniques and reports on their iterative development during two usability studies. The continuing goal of this work is to (i) encourage the use of simplified overviews in temporal event sequence analysis and (ii) inspire discussion of simplified overviews for other data types.

Related Work

After a brief overview of event sequence analytics and EventFlow, we briefly discuss topics influential to our techniques including: visual complexity, simplified overviews, quality metrics, and automated ranking.

We implemented Choose2 within EventFlow, which is a visual analytics tool that assists analysts with understanding complex temporal data composed of event sequences [11]. When exploring this type of data, analysts are trying to understand the relationships that exist among event sequences found in multiple records [12]. For example, Electronic Health Records (EHR) are studied to understand how differing procedures or medications lead to successful outcomes. Early work on EventFlow showed how analysts could use it to discover meaningful temporal relationships between point-based events using aggregated display overviews [20,21]. Subsequent revisions to EventFlow enabled the analysis of interval events [13,14] as well. However, despite the changes to EventFlow, case studies with practitioners indicated a consistent problem: How should analysts begin their analysis when they are overwhelmed by the initial visual complexity of their data? While analysts readily recognize their data in the detail record view (right side of Figure 2), they were often troubled by the visual complexity of the overview (center Figure 2).

Visual complexity refers to the level of detail in a visualization created by an analytics tool [8]. For example, in scatterplots the number of points, occlusion, and entropy has been used as measures of visual complexity. For aggregated event overviews, as used in EventFlow, Monroe *et al.* defined two types of visual complexity: number of visual elements (e.g., vertical bars) and average height of these elements [12]. A decrease in the former and an increase in the latter indicate a simpler display. Fine-grain measures (e.g., number of events categories) have also been noted to influence visual complexity [7].

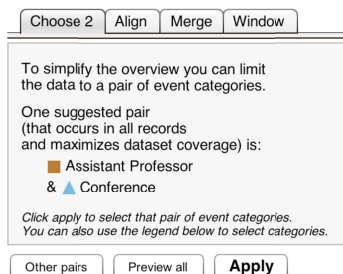


Figure 3: The Choose2 simple panel suggests a pair for novice analyst to start their analysis with and provides a brief rationale; it is a modal window within EventFlow.

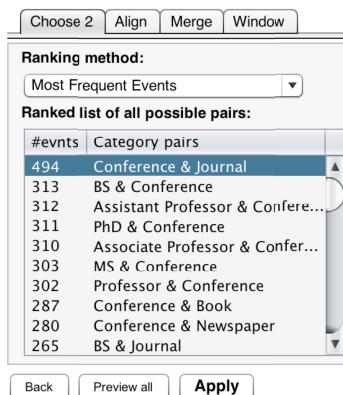


Figure 4: The ranked list of suggestions interface displays all pairs ranked by the selected metric.

When overviews are complex, skilled analysts can simplify them by manually filtering event categories [6]. An alternate approach could be to have the system pre-analyze the data and then present suggestions on how to proceed [9]. One method for generating these suggestions is through the use of quality metrics, which are calculated about the data to highlight potential areas of interest. These metrics can be used to generate sample visualizations, find meaningful abstractions of the data, and reduce clutter [4]. Once calculated, metrics allow suggestions to be ranked and communicated to analysts: the Hierarchical Clustering Explorer, with its Rank-by-Feature framework [18], provided similar rankings (e.g., correlation coefficient). Choose2 applies these ideas to event analytics and offers six metrics for measuring a pair’s complexity.

Description of the Three Techniques

We describe the three Choose2 techniques using the same synthetic professor dataset (Figure 2) used in the usability testing (see also the supplementary video).

Simple Panel with Single Suggestion

For most novice users, the suggestion of a single pair of event categories to start with helps speed their efforts (Figure 3) to explore new data. The default suggestion is the pair of events that appears across the largest number of records. In the professor dataset, the automatic suggestion is to look at “Assistant Professor” and “Conference” events. Clicking “Apply” executes the suggestion and simplifies the overview accordingly. This technique provides quick access to one simplified overview without having to understand the details of the pair selection process. A simple rationale for the suggestion is also provided. The suggestion is non-modal so users can ignore it and access other controls.

Ranked List of Suggestions

When analysts click “Other pairs”, they see advanced simplification options: a ranked list panel of six ranking metrics (Figure 4). This list is a dropdown menu, where pairs are ranked by the selected metric’s value. A popup explains the metrics’ calculation methods. See Appendix II for definitions of ranking metrics.

Event Category Matrix

While the simple panel seemed appropriate for novice users, and the ranked list could guide more advanced users toward several simplified overviews, it seemed equally important to provide rapid access to all simplified overviews for expert analysts. When analysts click on “Preview all” in either of the previously described panels a new window opens and displays the ECM, an organized set of small overviews for all pairs as well as for individual event categories (Figure 1b, Appendix I). Clicking on a cell of the matrix applies the event category pair simplification. While related to the Lower Triangular Scatterplot Matrix for multivariate data [15], this summarization of event sequence data using a matrix provides a novel way to review temporal relationships between events. There are $(n(n-1))/2$ pairwise event category overviews. The diagonal is used to show the temporal relationships among repeated events of a single category, for a total of $(n(n-1))/2 + n$ cells. With 10 event categories, there are $45 + 10 = 55$ simplified overviews in the ECM display. As a result, we believe the ECM enables analysts to extract insights difficult to see in the initial overview.

Evaluation and Evolution of the Design

The interfaces were designed iteratively using a process that included two usability studies with five and six participants respectively (4 Computer Science graduate

Appendix I: Expanded explanation of Figure 1b, the Event Category Matrix.

Starting along the diagonal and from the top analysts can review the data one event category at a time: a few professors have multiple Bachelor degrees, not everyone has a Masters. Analyst can also see that not everyone has reached full professor status. At a glance, analysts can see that professors all have many conference and journal papers; however, most have either one or no book or newspaper publications. The pairwise overviews show that those who had multiple bachelors received their PhD faster than those who had only one. Some professors started as Assistant professor before defending their PhD. We see that book and newspaper publications generally appear in later career with a few dramatic, easily visible, exceptions. In the middle portion, "Conference" and "Journal" events repeat frequently which adds complexity to the visualization. This complexity may indicate potential category pairs of interest and analysts can choose to review the details in EventFlow's main window.

students, 1 Information Science graduate student, and 1 professor of Urban Studies) with limited or no experience with visual analytics. The professor dataset was used because the context of the data was familiar to the participants and there were enough sequence variations to produce a visually complex overview at the start (*i.e.*, when all event categories are included). Using a simpler training dataset, participants learned how to read an EventFlow overview and filter by event categories using controls in the legend. Participants were also informed of our motivation to help users begin analysis when the overviews are complex. The professor dataset was then loaded and the observer pointed to the Choose2 panel asking participants to explore freely and provide feedback on that feature. No additional information was provided. Participants were asked to "think aloud" during their exploration, point at what they found unclear, and indicate what they learned from the data. Improvements were discussed at the end of each 45 minute session.

During the first study only the simple and ranked list panels were tested. The design of these panels was revised, the ECM was developed, and then all three designs were tested in a second study before being refined again. This paper summarizes a few of the lessons learned (*i.e.*, what didn't work) and what improvements were made. We organize the results of both studies by the three proposed techniques.

Review of Simple and Ranked List Panels

The first participants had difficulties with the early versions of explanatory text and the names and description of the metrics. Improvements were discussed and new explanations and labels were drafted with early participants and further refined with

later ones (overall about 7 or 8 versions of the text of the simple panel and ranked list were tested with participants and colleagues). Short popup explanations were added, which now appear when the cursor hovers over a method in the dropdown menu. Participants also asked for additional ranking metrics (*i.e.*, they had access to a subset of ranking metrics from the list presented in Appendix II) so three more were added.

The early version of the interface appeared as a one-time modal dialog that disappeared after the overview was simplified. Participants wanted to come back and try more pairs, so this dialog was refactored to be a permanent non-modal panel, visible after loading the dataset and always available. The 6 participants in the second study were content with the panel placement.

Originally the simple panel included a menu of the ranking metrics to choose from, but it was confusing for first time users causing us to push that menu to a separate "ranked list" panel (so that the more extensive exploration of the pairs and their rankings would take place in the ranked list panel). Participants had difficulties guessing what the ranking method did based solely on the labels, but had better results by inspecting the list, trying different pairs, and looking at the scores assigned to each pair. The original design of the ranked list panel colored each pair using a color gradient from green (high ranking) to red (low ranking). This was found confusing because color is used heavily in EventFlow (and mapped to event categories), so the use of color was discontinued.

Review of Event Category Matrix

The first study led to the design of the ECM as a method to systematically explore all pairs, and it was

Appendix II:

Six ranking metrics were implemented by Choose2's ranked list based on maximizing or minimizing the following: record coverage, event co-occurrence, and frequency of event occurrence. We define these metrics as followed:

Record Coverage. The cardinality of the union of records containing event category A and records containing event category B.

Event Co-occurrence. The cardinality of the intersection of records containing event category A and records containing event category B.

Frequency of Event Occurrence. The top two occurring event categories based on the total number of occurrences across all records.

tested in the second study. The "Preview all" and the ECM panel were immediately well received by 4 out of the 6 participants who felt at ease and said they understood what it was, stating "*I know it's a matrix*"; however, all took a few minutes to fully understand this new display. Two participants said at first that the diagonal was not needed and could be removed; however, they all eventually understood that the diagonal was useful to show one event category at a time. They commented that it was not really a pair, but still was useful to have. Out of the 4 participants who were immediately comfortable, 2 said they would not use the other panels but would go to the ECM directly. One participant said "*I like [the ECM]; it shows where the complexity is.*" Two participants took more time understanding what the ECM did (between 3 and 5 minutes). All participants discovered that they could click on cells of the ECM to select the event categories and see the overview in more detail in the main EventFlow window. The text and labels on the panels were also thoroughly discussed and revised.

Summary of Findings

Overall the placement and clarity of the three panels improved over the course of the two usability studies. All participants understood what had happened when the overview was simplified. However, it is likely that the ranking methods will remain difficult to grasp for some users. On the other hand the ECM seems very promising because it was understandable (without training, and only 2-4 minutes of exploration). Once they mastered the ECM participants were able to describe what they learned from it and seemed more focused on the data than when using other techniques.

Limitations and Future Work

The goal of Choose2 is to provide techniques for analysts of varying skill levels to help reduce overwhelming/initial visual complexity; however, there are many possible ways to do this (e.g., providing details-on-demand). Scaling the ECM is a challenge since the number of pair views increases quadratically with the number of event categories, which can be addressed by applying the rank by feature framework [17]. Additionally, our usability studies were formative and further testing is needed to investigate the usefulness with professional analysts. Future work will focus on exploring ranking metrics, solutions to scaling, and testing with analysts.

Conclusion

This work investigates three new techniques for providing simplified overviews that reduce the visual complexity of temporal event sequences in aggregated display overviews. We described how the designs evolved over time and what improvements were made. The display of a single pair of categories provides a simple way for novice users to start analysis of temporal event sequences. The ranked list of suggestions is slightly more complicated, but provides much greater assistance for novices. For more advanced users the ECM provides a way to systematically explore all event category pairs. Our usability studies demonstrated that all three techniques were learnable without training. Mitigating challenges with visual complexity is critical in event analytics. We believe that the proposed techniques will help event sequence analysts reduce initial visual complexity, thereby accelerating data exploration. Additionally, we believe that these techniques are applicable to other data types and should be further explored.

References

1. Christopher Ahlberg. 1996. Spotfire: An Information Exploration Environment. *SIGMOD Record* 25, 4: 25–29. Retrieved from <http://dblp.uni-trier.de/db/journals/sigmod/sigmod25.html>
2. Jurgen Bernard, David Sessler, Thorsten May, Thorsten Schlomm, Dirk Pehrke, and Jorn Kohlhammer. 2015. A Visual-Interactive System for Prostate Cancer Cohort Analysis. *Computer Graphics and Applications, IEEE* 35, 3: 44–55.
3. Jürgen Bernard, Nils Wilhelm, Bjorn Kruger, Torsten May, Tobias Schreck, and Jorn Kohlhammer. 2013. Motionexplorer: Exploratory search in human motion capture data based on hierarchical aggregation. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12: 2257–2266.
4. Enrico Bertini, Andrada Tatu, and Daniel Keim. 2011. Quality Metrics in High-Dimensional Data Visualization: An Overview and Systematization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12: 2203–2212. <http://doi.org/10.1109/TVCG.2011.229>
5. Jesus J Caban and David Gotz. 2015. Visual analytics in healthcare--opportunities and research challenges. *Journal of the American Medical Informatics Association* 22, 2: 260–262.
6. Stuart K Card, Jock D Mackinlay, and Ben Shneiderman. 1999. *Readings in information visualization: using vision to think*. Morgan Kaufmann Publishers Inc.
7. Fan Du, Ben Shneiderman, Catherine Plaisant, Sana Malik, and Adam Perer. Coping with Volume and Variety in Temporal Event Sequences: Strategies for Sharpening Analytic Focus. (*under review*).
8. Alexandra Forsythe. 2009. Visual Complexity: Is That All There Is? *Proceedings of the 8th International Conference on Engineering Psychology and Cognitive Ergonomics: Held As Part of HCI International 2009*, Springer-Verlag, 158–166. http://doi.org/10.1007/978-3-642-02728-4_17
9. Jerome H Friedman and John W Tukey. 1974. A Projection Pursuit Algorithm for Exploratory Data Analysis. *Computers, IEEE Transactions on* C-23, 9: 881–890. <http://doi.org/10.1109/T-C.1974.224051>
10. Tim Lammarsch, Wolfgang Aigner, Alessio Bertone, Silvia Miksch, and Alexander Rind. 2014. Mind the time: Unleashing temporal aspects in pattern discovery. *Computers & Graphics* 38: 38–50.
11. Tamra E Meyer, Megan Monroe, Catherine Plaisant, et al. 2013. Visualizing Patterns of Drug Prescriptions with EventFlow: A Pilot Study of Asthma Medications in the Military Health System US Army. *Workshop on Visual Analytics in HealthCare*, 1–4.
12. Megan Monroe, Rongjian Lan, Hanseung Lee, Catherine Plaisant, and Ben Shneiderman. 2013. Temporal Event Sequence Simplification. *Visualization and Computer Graphics, IEEE Transactions on* 19, 12: 2227–2236. <http://doi.org/10.1109/TVCG.2013.200>
13. Megan Monroe, Rongjian Lan, Juan del Olmo, Ben Shneiderman, Catherine Plaisant, and Jeff Millstein. 2013. The Challenges of Specifying Intervals and Absences in Temporal Queries: A Graphical Language Approach. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, ACM*, 2349–2358. <http://doi.org/10.1145/2470654.2481325>

14. Megan Monroe, Krist Wongsuphasawat, Catherine Plaisant, Ben Shneiderman, Jeff Millstein, and Sigfried Gold. 2012. Exploring point and interval event patterns: Display methods and interactive visual query. *University of Maryland Technical Report*.
15. Francois Nielsen. 2006. Lower Triangular Scatterplot Matrix for Multivariate Data. *University of North Carolina at Chapel Hill*. Retrieved from <http://www.unc.edu/~nielsen/soci709/m8/m8002.jpg>
16. Adam Perer and Fei Wang. 2014. Frequence: Interactive Mining and Visualization of Temporal Frequent Event Sequences. *Proceedings of the 19th International Conference on Intelligent User Interfaces*, ACM, 153–162. <http://doi.org/10.1145/2557500.2557508>
17. Jinwook Seo and Ben Shneiderman. 2005. A Rank-by-feature Framework for Interactive Exploration of Multidimensional Data. *Information Visualization* 4, 2: 96–113. <http://doi.org/10.1057/palgrave.ivs.9500091>
18. Jinwook Seo and Ben Shneiderman. 2006. Knowledge discovery in high-dimensional data: case studies and a user survey for the rank-by-feature framework. *Visualization and Computer Graphics, IEEE Transactions on* 12, 3: 311–322. <http://doi.org/10.1109/TVCG.2006.50>
19. Ben Shneiderman. 1996. The eyes have it: a task by data type taxonomy for information visualizations. *Visual Languages, 1996. Proceedings., IEEE Symposium on*, 336–343. <http://doi.org/10.1109/VL.1996.545307>
20. Taowei David Wang, Catherine Plaisant, Ben Shneiderman, et al. 2009. Temporal summaries: Supporting temporal categorical searching, aggregation and comparison. *IEEE Transactions on Visualization and Computer Graphics*, IEEE Computer Society, 1049–1056. <http://doi.org/10.1109/TVCG.2009.187>
21. Krist Wongsuphasawat, John Alexis Guerra Gómez, Catherine Plaisant, Taowei David Wang, Meirav Taieb-Maimon, and Ben Shneiderman. 2011. LifeFlow: visualizing an overview of event sequences. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 1747–1756.
22. Krist Wongsuphasawat and Jimmy Lin. 2014. Using visualizations to monitor changes and harvest insights from a global-scale logging infrastructure at Twitter. *Visual Analytics Science and Technology (VAST), 2014 IEEE Conference on*, 113–122.