

Timeline Storyteller

The Design & Deployment of an Interactive Authoring Tool for Expressive Timeline Narratives

Matthew Brehmer, Bongshin Lee, Nathalie Henry Riche,
David Tittsworth, Kate Lytvynets, Darren Edge, Christopher White
Microsoft Research

{mabrehme,bongshin,nath,dattittsw,kalytv,daedge,chwh}@microsoft.com

ABSTRACT

Timeline Storyteller is an open-sourced interactive application for authoring expressive visualization-based narratives about event sequences. Its design reflects a recently proposed timeline design space, which itself is grounded in an extensive survey. In this paper, we highlight its capabilities and reflect upon its usage and adoption.

KEYWORDS

Visualization, storytelling, event sequence data, deployment.

Citation:

Matthew Brehmer, Bongshin Lee, Nathalie Henry Riche, David Tittsworth, Kate Lytvynets, Darren Edge, Christopher White. 2018. Timeline Storyteller: The Design & Deployment of an Interactive Authoring Tool for Expressive Timeline Narratives. In *Proceedings of Computation + Journalism Symposium (C+J'19)*.

INTRODUCTION

In recent years, there has been an increasing demand for tools that produce compelling narratives with information visualization. A recurring limitation of such tools is their expressive range, offering limited options for specifying visual encoding design choices or selecting chart templates. This is particularly true of timelines, a type of chart used to present information about event sequences in domains such as journalism. Timelines are not a standard chart type in many existing chart creation tools, and existing timeline-specific visualization tools have limited expressivity with respect to how time is represented, how time is scaled, and how a narrative is revealed to the viewer.

In this paper, we present Timeline Storyteller, a freely-available open-source tool for authoring and presenting expressive visual narratives about event sequence data. We profile its design and document its deployment and usage as a web application and as a free extension for Power BI, an established data visualization tool with a large international user community and an increasing level of adoption by news organizations. Following Brehmer et al.'s recently proposed timeline design space [2], our goal was to identify how people make use of an authoring interface that makes this design space accessible. We also discuss how Timeline Storyteller and the design choices it offers shaped how authors produced stories with their own event sequence data. Finally, we reflect upon the

dissemination of Timeline Storyteller as well as our methodology for studying its usage following its deployment.

Leveraging visualization in a narrative medium is an evolving art form and research area [8], with several established genres [15]. Research questions posed by those studying visual data-driven storytelling are notably different from those studying the use of visualization in data analysis. While the latter group is generally concerned with scalability, performance, and how analysts can spot patterns, the former group considers topics such as visualization literacy and interpretability by laypersons, best practices for sequentially revealing information [6], and the role of annotation [10].

While many existing interactive tools allow people to create charts [12], few tools provide simultaneous support for expressive visual encoding, storytelling, and annotation. For instance, it is possible with tools such as Tableau, Power BI, or Ellipsis [13] to produce a sequence of discrete scenes, where each scene can present a limited range of chart types. In contrast, tools such as Lyra [14] are more visually expressive but cannot produce sequential multi-scene narratives. Timeline Storyteller combines aspects of both classes of tools, albeit focusing on timeline data.

TIMELINE STORYTELLER

The aim of Timeline Storyteller is to enable people to author a visually expressive narrative: a succession of scenes incorporating information visualization and annotations connected by animated transitions. However, it specifically focuses on data that can be represented as a timeline: sequences of events, a form of data underserved by existing interactive chart creation tools.

The design choices in Timeline Storyteller reflect Brehmer et al.'s recently proposed design space [2], which itself is grounded in a survey of hundreds of timeline graphics, techniques published in the research literature [1], and existing timeline tools. The majority of existing timeline tools are intended primarily for exploratory data analysis [2] in specific application domains, such those involving electronic health records. Meanwhile, existing presentation tools used in online journalism such as the Knight Lab's TimelineJS [9] and ProPublica's TimelineSetter [16] are limited to linear representations of time and chronological time scales, whereas Brehmer et al.'s survey of timelines showed that timeline designers often make use of alternative representations of time (e.g., radial, grid), time scales (e.g., relative, logarithmic), and timeline layouts (e.g., categorically faceted timelines, or chronologically segmented timelines, such as segmenting a timeline spanning a century into decades). Before Timeline Storyteller, those who opted to incorporate any of these design choices resorted to time-intensive manual illustration or to custom programmatic implementations.

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(s).

C+J'19, Feb 2019, Miami, FL USA

© 2019 Copyright held by the owner/author(s).




Figure 1: The interface of Timeline Storyteller, where the timeline canvas spans the entire browser window. Interface components include: (A) toggles for alternative timeline representations, scales, and layouts; (B) a movable event legend; (C) widgets for loading data, annotating the canvas with captions and images, filtering and highlighting events, and exporting content; (D) scene recording and playback controls along with a scene list; and (E) tutorial and help widgets. This example displays a scene from a story about the performance record of professional golfers at the US Open golf tournament throughout their careers (see the interactive story at <https://aka.ms/timelinestoryteller-golf-story>). Golfer images: Wikimedia (CC).

The interface and authoring workflow reflects the intent of presenting a narrative to an audience. We assume that authors already have a sense of the purpose of their story, and accordingly Timeline Storyteller is designed to help them realize ways to present it.

Importing event data: Timeline Storyteller can ingest event sequence data in a variety of formats (e.g., CSV, JSON, a spreadsheet). Given an array of events, each event must have a start date/time, which itself can be specified in a variety of formats. Each event can also be augmented with an end date/time, a text description, a category, and a facet (a second categorical attribute).

Selecting representation, scale, and layout: Initially all events appear as a single chronological timeline drawn linearly from left to right across the canvas, allowing authors to select alternative combinations of representation, time scale, and layout (see Figure 1-A), choices corresponding to the dimensions of Brehmer et al.'s design space [2]. In Figure 1, timelines depicting the careers of professional golfers are represented radially along a relative time scale and are faceted by golfer; event categories such as tournament victories, placements, and disqualifications are shown in a legend that can be collapsed and repositioned (see Figure 1-B). We restricted the possible combinations of representation, scale, and layout to the 20

viable combinations identified by Brehmer et al. as being perceptually interpretable, generalizable across datasets, and purposeful with respect to a communicative or narrative intent.

Filtering and highlighting events: Authors can filter and restore events according to their category, facet, or when they occur in time (see Figure 1-C). Authors can also choose to emphasize filtered events by either hiding or reducing the opacity of events not meeting the filter criteria.

Annotating and captioning: Authors have the option to annotate the canvas with captions, event labels, and images. When hovering over an event, event descriptions appear as a label with a rectilinear pointer to the event; two event labels are shown in Figure 1. Clicking on an event pins the annotation to the canvas. Authors can also position and resize captions, images, and labels via dragging.

Presenting a story: Authors can record the current canvas as a scene, which is added to a scene array (see Figure 1-D), where clicking on a scene thumbnail triggers a transition to display that scene. Timeline Storyteller smoothly animates the position of events during scene transitions to provide the audience with a continuous context. For live presentations, authors can trigger a read-only playback mode, hiding everything except the canvas and a minimal scene navigation interface.




Figure 2: The Microsoft Power BI version of Timeline Storyteller; Power BI’s canvas (left) contains the web version interface. Shown here is a faceted linear timeline depicting the biographies of Henri Matisse and Pablo Picasso.

Sharing a story: When completed, authors can export the current timeline canvas as a PNG or SVG image, or they can export entire stories as animated GIFs or as interactive iFrames for embedding in a web page or news article, allowing stories to be consumed asynchronously¹. When viewing an embedded story, viewers cannot make changes to the content; instead, they are limited to navigating between the next and previous scenes in a story. Finally, the story can also be saved so that authors can resume editing later.

DEPLOYMENT

We were interested in expressivity and how people tell stories with their own event sequence data. Given this goal, we deemed a controlled laboratory study to be inappropriate, as these contexts tend to focus on metrics that are better suited for visual analysis tools, such as accuracy and task completion time [11]. We therefore deployed Timeline Storyteller as an open-source web application in early 2017², and as a free extension for Microsoft’s Power BI via its custom visual store in mid-2017 (see Figure 2). We instrumented the web version to store a copy of exported content in a secure storage repository, subject to the consent of the author. With this approach, our intent was to collect and classify the event data that authors imported into Timeline Storyteller: the cardinality of events, their chronological extent, the heterogeneity of categories, and the subject matter of the data. We also hoped to gain insight into the design choices that authors make when using Timeline Storyteller, including their coverage of Brehmer et al.’s timeline design space [2], their use of annotations and images, and their narrative sequencing when producing stories with multiple scenes.

Facilitating usage and adoption: As Timeline Storyteller is a novel authoring environment for a datatype not typically accommodated by existing tools, we incorporated an Intro.js [7] walkthrough for the interface (see Figure 1-E) to facilitate author onboarding. On our website, we also provided a data formatting guide, an example image gallery, and a collection of example story videos, with the latter serving to demonstrate all of Timeline Storyteller’s functionality. We also produced a 37-minute video interview and tutorial

for Power BI’s YouTube channel³, which demonstrated the entire workflow of producing a story and embedding it in a web page.

To further inspire authors to create a wide range of stories, we provided over 20 example event datasets that could be loaded from within Timeline Storyteller. These datasets differ from one another in terms of the number of events, their chronological extent and distribution, the number of categories and facets, and topic area, including politics, science, world history, and art.

Throughout 2017, we publicly demonstrated Timeline Storyteller to promote its usage at practitioner-oriented visualization conferences and at post-secondary educational institutions, which included a business school and two journalism schools. Finally, we held a contest in which we solicited stories produced using Timeline Storyteller from the Power BI user community over the course of two weeks. Unlike academic visualization contests such as the VAST Challenge or the SciVis contest, we did not designate a theme and we encouraged participants to use their own data to tell a story. Together with Microsoft’s Modern Journalism team, we judged submissions based on their use of Timeline Storyteller, their use of Power BI, and the overall quality and novelty of the story, announcing the results via social media and the Power BI blog.

Usage at public events: In the summer of 2017, we partnered with several presenters who used Timeline Storyteller in three live presentations at public events. First, we worked with a representative from the UK National Trust, a charitable organization that manages a portfolio of historical buildings and properties. In collaboration with this representative, we prepared a dataset and story depicting the history of eight National Trust properties, which he presented during a keynote presentation at the Data Insights Summit in Seattle⁴. This story featured multiple scene transitions, a faceted layout of timelines for eight historical properties, as well as a set of image and text annotations. Second, one of the authors of this article used Timeline Storyteller to present a history of artificial intelligence during a keynote talk at the Dublin Data Summit⁵. This story began with an unconventional spiral timeline of milestones in AI history and made substantial use of Timeline Storyteller’s filtering features to incrementally reveal different aspects of this history. Finally, members of Microsoft’s Modern Journalism team repurposed the AI history story for an exhibition at the Future of Storytelling Summit in New York City, where they demonstrated it to journalists and new media enthusiasts.

Content produced with Timeline Storyteller: As of Fall 2018, the Timeline Storyteller Power BI extension has been downloaded over 36,700 times; meanwhile, our YouTube tutorial has been viewed over 46,000 times. Given this response to our deployment, we now examine the content that people produced with the tool.

We recorded export events from the web version between March and December 2017, which included 223 exports of content by 41 authors visualizing their own data who opted into sharing their content with us. The Power BI version was not similarly instrumented, and thus we cannot comment on the content produced with it. Upon inspection of the content exported from the web version, we found that authors often exported incremental versions or

¹e.g., the story featured in Figure 1: <https://aka.ms/timelinestoryteller-golf-story>

²<https://github.com/Microsoft/timelinestoryteller>, <https://timelinestoryteller.com>

³<https://youtu.be/bwiMfwBVsq>

⁴<https://youtu.be/zaVGt-IbuhU?t=28m1s>

⁵<https://aka.ms/timelinestoryteller-ai-story>

duplicates of the same content in different formats (PNG, GIF, SVG, interactive stories); after removing these, we identified 53 unique instances of exported content. The event data in this content varied considerably and included timelines consisting of only 4 events as well as those consisting of hundreds of events, from timelines with a single event category to those with a dozen categories, and from timelines spanning a single day to those spanning 500 years. The topics of these timelines were also varied, though it appeared that political history, art history, and military history were particularly common topics, based on their annotations and captions; we can only infer these themes as a substantial amount of the exported timeline content was not in English.

The majority of this content featured a conventional linear representation, a chronological time scale, and a single unified timeline: the most prevalent design choices in Brehmer et al.'s survey of timelines [2]. However, we discovered that radial representations and faceted layouts were also relatively common among the content exported from Timeline Storyteller. We further determined that every design choice from Brehmer et al.'s design space was represented at least once within this corpus, except for the century grid representation and a logarithmic time scale, design choices that were also relatively uncommon in their survey. Similarly, all of Timeline Storyteller's annotation and filtering design choices were also represented in the exported content. While we were able to catalog this exported content, what remains unclear is how the authors subsequently used this content, or whether they were merely experimenting with the tool. It is therefore possible that some of the exported content was later integrated into print or online articles.

For our storytelling contest, we received seven submissions, with the qualifying submissions making use of several design choices within Timeline Storyteller. We selected one winning story and two honorable mentions, which were profiled and embedded in a blog post about the contest⁶. The winning story, which depicted the history of TV ratings, made use of many of Timeline Storyteller's design choices (See Figure 3-left). It contained seven scenes which featured a variety of image and caption annotations, a timeline segmented by decade, a radial representation, and use of the filtering controls to emphasize a subset of events. The author of this story subsequently posted a video tutorial describing his process of creating the story⁷. Another contest entry worth noting is one of the two honorable mentions, which depicted timelines of tropical cyclones (See Figure 3-right). In this instance, the spiral representation was a particularly fitting visual metaphor considering the subject matter. Additionally, the author manually configured a custom color palette corresponding to wind speed ranges, setting it apart from all previous content generated with Timeline Storyteller in that the events were distinguished according to a quantitative attribute rather than a categorical one.

DISCUSSION

Considering all the author-generated content that we collected, the storytelling contest submissions, and the stories presented at public events, we now have a sense of how people make use of the design choices offered by Timeline Storyteller. Because these design




Figure 3: Left: The winning entry of the Timeline Storyteller contest by Devin Knight depicted the history of top-rated TV programs on major networks. Right: An honorable mention in the contest by Manga Solutions depicted the length and wind speed of several tropical cyclones.

choices reflect Brehmer et al.'s design space and their proposed considerations for expressive storytelling [2], these results also serve as an ecological validation of that work.

Facilitating authoring with recommendations: Though only 20 combinations of timeline representation, scale, and layout are currently achievable in Timeline Storyteller, when combined with annotations and filter states, the scope of design choices in Timeline Storyteller is immense. This scope may be intimidating for prospective adopters, which perhaps explains why most of the exported content described above makes use of a conventional linear representation and a chronological time scale. A possible solution would be to replace or supplement the current design choice interface with a ranked list of recommended combinations of representation, scale, and layout that is based on properties of the data, such as the chronological distribution of events, the number of events, or the number of unique event categories. This would be similar to how tools like Voyager [17] recommend charts based on the statistical properties of tabular quantitative data.

Authors have remarked to us that it can be tedious to manually annotate events on the timeline canvas. Thus, another design recommendation would be to annotate a subset of events with their description and to automate the placement of these labels, like how Bryan et al. automated the annotation of stacked area charts [3], and to allow the author to manually adjust these labels.

Supporting the entire storytelling workflow: We had initially assumed that authors would already have a story in mind before using Timeline Storyteller, and prior to its release, these authors would have resorted to time-intensive illustration tools or to bespoke programmatic implementations. However, prospects for design choice and annotation recommendations along with requests from authors for exploratory interactions such as panning and zooming suggests that people want to use Timeline Storyteller as a data analysis tool as well. However, in its current form, it is primarily an expressive presentation tool. It would be therefore useful to investigate how Timeline Storyteller could be used in the context of the entire storytelling process, from analysis to presentation [4, 8]. This process begins with data collection and curation. For example, several of the sample timeline datasets available in Timeline Storyteller were generated using a tool called TimeLineCurator [5], which allows people to extract events from text documents and export timeline datasets in several formats; it is therefore worth considering ways to integrate such functionality into Timeline Storyteller's workflow. Following data curation, a data analysis process could be augmented

⁶<https://aka.ms/timelinestoryteller-contest>

⁷<https://youtu.be/fyyO2JmuNsg>

with recommendations and enhanced navigation capabilities, as well as an ability to view an interaction history or bookmark states of the application, where bookmarking could be a precursor to recording the state as a scene for presenting to an audience. In regards to data analysis in general, it is worth remarking on Timeline Storyteller's integration with Power BI, a tool that is used for both data analysis and presentation, and that news organizations are increasingly embedding Power BI content in their articles. Published content may reflect only a small part of the analysis and storytelling workflow, a workflow that invites further in-depth study.

Challenges of studying deployed tools: There is little consensus on how to assess the adoption and impact of tools for authoring data-driven stories [11]. Typically reported metrics of accuracy and completion time for graphical perception and data analysis tasks do not seem appropriate for visualization authoring tools such as Timeline Storyteller. In contrast, the constructs we were interested in studying included expressivity, the audience and community reception of content created with Timeline Storyteller, and barriers to adoption. With respect to the latter, we initially performed typical think-aloud usability tests with example datasets before its deployment, which did reveal low-level learnability and interface complexity issues that we resolved prior to deployment. However, these tests could not tell us anything about expressivity or content reception. Ultimately, we believed that an evaluation of our tool would require the study of people visualizing their own data. With the analysis reported in this paper, we have begun to shed light on the issue of expressivity, though more work is needed with respect to studying the audience response to content produced using it.

Our methodological approach is not without its limitations. Collecting exported content and contest submissions allowed us to study the expressive range of content produced with Timeline Storyteller, but these artifacts tell us little about the process of their creation and the difficulties their authors faced while producing them. Furthermore, we have little insight on instances in which a person discontinued their use of Timeline Storyteller before exporting content, except in cases where they contacted us directly to explain their motivations and the challenges they encountered.

Another issue to consider when collecting authored content from a deployed visualization tool such as Timeline Storyteller is the privacy of the author and potentially identifying content in their stories, which is particularly true of timelines containing biographical information; for instance, we noted one instance where an author exported a timeline of their curriculum vitae. When we initially released Timeline Storyteller as a web application in 2017, we collected the content that authors generated using Timeline Storyteller, allowing them to opt out sharing this information if they wished. It is important to stress that this form of logging requires compliance with data collection and storage laws, and the interpretation and sharing of the collected collected must be handled delicately, particularly in the absence of direct communication with content authors. Thus, we no longer monitor how authors use Timeline Storyteller's design choices. However, we continue to learn from authors who contact us directly.

The visualization research community therefore needs to devise new methods with which to study the expressiveness of deployed visualization tools. These methods must respect the privacy of

people using these tools with their personal datasets, and they should scale beyond a handful of case studies and short-term data collection events such as contests where people consciously opt to share their content with a public audience.

CONCLUSION

We introduced Timeline Storyteller, a realization of a recent timeline design space [2] and an open-source web- and Power BI-based tool for authoring expressive timeline stories in domains such as journalism. We reported on the deployment and adoption of Timeline Storyteller over the course of a year, reflected upon the expressivity of the tool and our methodology, and indicated opportunities for future research and development.

REFERENCES

- [1] Wolfgang Aigner, Silvia Miksch, Heidrun Schumann, and Christian Tominski. 2011. *Visualization of Time-Oriented Data*. Springer. <http://timeviz.net/>
- [2] Matthew Brehmer, Bongshin Lee, Benjamin Bach, Nathalie Henry Riche, and Tamara Munzner. 2017. Timelines revisited: A design space and considerations for expressive storytelling. *IEEE Trans. Visualization & Computer Graphics (TVCG)* 23, 9 (2017), 2151–2164. <https://doi.org/10.1109/TVCG.2016.2614803>
- [3] Chris Bryan, Kwan-Liu Ma, and Jonathan Woodring. 2017. Temporal summary images: An approach to narrative visualization via interactive annotation generation and placement. *IEEE Trans. Visualization & Computer Graphics (Proc. InfoVis 2016)* 23, 1 (2017), 511–520. <https://doi.org/10.1109/TVCG.2016.2598876>
- [4] Fanny Chevalier, Melanie Tory, Bongshin Lee, Marian Dörk, Jarke van Wijk, Giuseppe Santucci, and Jessica Hullman. 2018. From analysis to communication: Supporting the lifecycle of a story. In *Data-Driven Storytelling*, Nathalie Henry Riche, Christophe Hurter, Nicholas Diakopoulos, and Sheelagh Carpendale (Eds.). A K Peters/CRC Press. <https://www.crcpress.com/Data-Driven-Storytelling/Riche-Hurter-Diakopoulos-Carpendale/p/book/9781138197107>
- [5] Johanna Fulda, Matthew Brehmer, and Tamara Munzner. 2016. TimeLineCurator: Interactive authoring of visual timelines from unstructured text. *IEEE Trans. Visualization & Computer Graphics (Proc. VAST 2015)* 22, 1 (2016), 300–309.
- [6] Jessica Hullman, Steven Drucker, Nathalie Henry Riche, Bongshin Lee, Danyel Fisher, and Eytan Adar. 2013. A deeper understanding of sequence in narrative visualization. *IEEE Trans. Visualization & Computer Graphics (Proc. InfoVis)* 19, 12 (2013), 2406–2415. <http://dx.doi.org/10.1109/TVCG.2013.119>
- [7] Intro.js 2018. <https://introjs.com/>.
- [8] Bongshin Lee, Nathalie Henry Riche, Petra Isenberg, and Sheelagh Carpendale. 2015. More than telling a story: Transforming data into visually shared stories. *IEEE Computer Graphics and Applications (Visualization Viewpoints)* 35, 5 (2015). <http://dx.doi.org/10.1109/MCG.2015.99>
- [9] Northwestern University Knight Lab. 2013. TimelineJS. <http://timeline.knightlab.com/>
- [10] Donghao Ren, Matthew Brehmer, Bongshin Lee, Tobias Höllerer, and Eun Kyoung Choe. 2017. ChartAccent: Annotation for data-driven storytelling. In *Proc. IEEE Pacific Visualization Symp. (PacificVis)*. <https://doi.org/10.1109/PACIFICVIS.2017.8031599>
- [11] Donghao Ren, Bongshin Lee, Matthew Brehmer, and Nathalie Henry Riche. 2018. Reflecting on the evaluation of visualization authoring systems. In *Workshop Proc. Evaluation and Beyond - Methodological Approaches for Visualization (BELIV)*. <https://aka.ms/renbeliv18>.
- [12] Lisa Charlotte Rost. 2016. What I learned recreating one chart using 24 tools. *Source* (2016). <https://goo.gl/uGE5dc>.
- [13] Arvind Satyanarayan and Jeffrey Heer. 2014. Authoring narrative visualizations with Ellipsis. *Computer Graphics Forum (Proc. EuroVis)* 33, 3 (2014), 361–370. <http://dx.doi.org/10.1111/cgf.12392>
- [14] Arvind Satyanarayan and Jeffrey Heer. 2014. Lyra: An interactive visualization design environment. *Computer Graphics Forum (Proc. EuroVis)* 33, 3 (2014), 351–360. <http://doi.org/10.1111/cgf.12391>
- [15] Edward Segel and Jeffrey Heer. 2010. Narrative visualization: Telling stories with data. *IEEE Trans. Visualization & Computer Graphics (Proc. InfoVis)* 16, 6 (2010), 1139–1148. <http://dx.doi.org/10.1109/TVCG.2010.179>
- [16] Al Shaw, Jeff Larson, and Ben Welsh. 2011. TimelineSetter. <http://propublica.github.io/timeline-setter/>
- [17] Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, and Jeffrey Heer. 2016. Voyager: Exploratory analysis via faceted browsing of visualization recommendations. *IEEE Trans. Visualization & Computer Graphics (Proc. InfoVis 2015)* 22, 1 (2016), 649–658. <https://doi.org/10.1109/TVCG.2015.2467191>