

Visual Narrative Flow: Exploring Factors Shaping Data Visualization Story Reading Experiences

S. McKenna^{1,2}, N. Henry Riche¹, B. Lee¹, J. Boy³, & M. Meyer²

¹ Microsoft Research ² University of Utah ³ United Nations Global Pulse
Redmond, WA, USA Salt Lake City, UT, USA New York, NY, USA

Abstract

Many factors can shape the flow of visual data-driven stories, and thereby the way readers experience those stories. Through the analysis of 80 existing stories found on popular websites, we systematically investigate and identify seven characteristics of these stories, which we name “flow-factors,” and we illustrate how they feed into the broader concept of “visual narrative flow.” These flow-factors are navigation input, level of control, navigation progress, story layout, role of visualization, story progression, and navigation feedback. We also describe a series of studies we conducted, which shed initial light on how different visual narrative flows impact the reading experience. We report on two exploratory studies, in which we gathered reactions and preferences of readers for stepper- vs. scroller-driven flows. We then report on a crowdsourced study with 240 participants, in which we explore the effect of the combination of different flow-factors on readers’ engagement. Our results indicate that visuals and navigation feedback (e.g., static vs. animated transitions) have an impact on readers’ engagement, while level of control (e.g., discrete vs. continuous) may not.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Information Interfaces and Presentation]: User Interfaces—Theory and methods

1. Introduction

Data-driven stories that tightly integrate visualizations have become a popular communication device in a variety of fields [SH10]. This has led the visualization research community to investigate the design factors that practitioners employ to craft narratives, from visual & interactive techniques [SH10, SLHS16] to specific genres [AHL*15, AHRL*17, BKH*16]. While specific knowledge on these factors is growing, there is still little understanding of which are predominant for, and how they may be combined to create effective *visual narrative flows*—which combine a reader’s input with story components and congruent visual feedback that tell the story matching the author’s intent and voice (Figure 1).

An ongoing informal debate on visual narrative flow centers around the effects of allowing readers to navigate through data-driven stories using either a click/tap input or a scroll input. We refer to this debate as the *stepper vs. scroller debate* (illustrated in Figure 2). Clicking to step through a story is like a slideshow, while scrolling is akin to panning up and down a long document. Practitioners from The New York Times recently advocate for scrolling because their readers tend not to fully consume stories that are delivered with steppers [Tse16, Ais16]. Others advocate for steppers, as they point out several potential issues pertaining to the use of scrolling (e.g., “scrolljacking”) [Kos16].




Figure 1: Here we illustrate high-level characteristics impacting visual narrative flow: a reader’s interactions with the story, the mechanisms tying the story components into a narrative, and the different forms of visual feedback perceived by the reader as they navigate, read, and interact with the visual data-driven story.

While navigation input may influence the flow and reading experience of a data-driven story, diverse examples from The Guardian [GMA*16] and The Wall Street Journal’s Custom Studios team [Stu16] suggest that there is more to shaping a visual narrative flow than just input. Furthermore, there is a limited understanding of how different flows might influence reading experiences. Here, we systematically examine what story design aspects are used by practitioners, which we name *flow-factors*, that encompass reader input, story components, and visual feedback as shown

in Figure 1. We then describe a series of studies we conducted to provide initial empirical evidence on how different combinations of flow-factors, i.e., how different visual narrative flows, can impact reading experiences. We focus specifically on readers' reactions, preferences, usability, and level of engagement with a visual, interactive, data-driven story because these aspects are important and known outcomes of readers' experiences with technology. As such, our main contributions are:

- a definition of visual narrative flow along with a characterization of seven flow-factors;
- a corpus of 80 visual data stories illustrating different combinations of these factors;
- two exploratory studies that shed initial light on the effect of different visual narrative flows on readers' reactions and preferences; and
- a crowdsourced study with 240 participants that provides an early assessment on the impact of different combinations of flow-factors on reader-perceived engagement.

2. Related Work

Few studies [HDH*13, BDF15] have focused on measuring the impact of different design factors on the reading experience of visual, interactive, data-driven stories. However, a number of tools have been developed to facilitate the creation of such stories. Ellipsis enables authors to create stories using a combination of scenes, visualization parameters, annotations, and triggers [SH14] to construct stepper-like sequences between views. VisJockey supports linking visualizations and animated transitions to corresponding text [KSJL14], which acts similarly to a stepper. Data-driven videos created with DataClips use eight types of animations [AHRL*17] and showcase automatic temporal sequencing, distinct from steppers or scrollers. While these tools effectively allow authors to create data-driven stories, the options they provide for fine-tuning the flow of a story are limited. Here, we discuss related work for the three characteristics in Figure 1 that can help shape the flow of a story: reader input, story components, and visual feedback. We also present previous work on measuring user engagement.

2.1. Reader Input & Navigation

The HCI research community has studied the role and trade-offs regarding reader input for navigating documents. The history of document navigation can be traced back to paradigms such as a scrolling document window with hypertext links and HyperCard, an adaptive layout with stepper-like navigation. Research has shown that scrolling can cause disruption while reading text [KMÅ02]. Recently, web developers introduced a new technique called scrolljacking, where scrolling input is disrupted on a webpage, and this becomes a usability issue as it disrupts navigation and can confuse or frustrate readers [Bos14, Kos16]. However, there are benefits for using animation of text down the page when readers scroll; they read the document faster and with fewer errors [KB05]. Similarly, direct manipulation [Shn82] and fluid interaction [EMJ*11] highlight the benefit of input which is rapid, reversible, and incremental. For steppers, Boy et al. showed that

readers engage with all sections on a page but that this may not occur linearly [BDF15], whereas scrolling linearly confines navigation. Graphics editors from The New York Times have advocated for flows that use scrolling input since readers did not reach all of their content using tabs, buttons, or sliders [Tse16]. Our work explores these input trade-offs for visual data-driven stories by observing reader preference and measuring engagement.

2.2. Story Components & Narrative Visualization

Previously, Segel and Heer's framework for narrative visualization presented a design space of story components and story genres [SH10]. However, we cannot fully characterize different visual narrative flows using this framework. For example, the scroller in Figure 2a contains animated transitions and motion which are linked directly to the reader's scrolling position on the page. Aspects like animated transitions are in Segel and Heer's design space, but there is a disconnect of how these forms of visual feedback connect to different types of reader input. Furthermore, other aspects like progress bars or linear slideshow stories are introduced, but the level of detail is insufficient to fully capture the range of visual narrative flows we have observed in stories on the web. In our work, we build upon this design space to capture properties that relate to how a reader navigates and consumes a visual data-driven story.

2.3. Visual Feedback & Animated Transitions

As readers navigate a story, visual feedback helps them by guiding their attention and showing change. Animation can play a role in visual discourse and storytelling [CRP16], and animated transitions help show state changes between two different graphics, such as changing the viewport and showing dynamic timesteps [HR07]. In studies on animated transitions, Heer and Robertson show animation's effectiveness for tracking objects [HR07], Feng et al. highlight how motion changes a person's affect [FBR14], Chang and Ungar argue that animation can increase user engagement for interfaces [CU95], and Chevalier et al. found that animating text vertically helps with reading more than page-flipping animations [CDBF10]. Animations can fail due to the apprehension principle [TMB02], and these detriments may be offset using interaction [TMB02], such as by linking animation to clicking or scrolling. Our work aims to explore how authors utilize these forms of visual feedback and measure the effect on reader engagement.

2.4. Engagement of a Data-Driven Story

Just like in web analytics, engagement plays an important role in determining the effectiveness and impact of a visual data-driven story. Engagement is a complex topic which lacks a unified definition in the community [MKK15], though some definitions include interdisciplinary aspects such as emotional, cognitive, and behavioral states between a person and an object [SES16, OT10]. For measuring engagement, the visualization community has primarily investigated time spent or number of interactions [BDF15, SES16] though others have explored subjective reaction cards to capture user feelings [SES16]. O'Brien and Toms break apart, define, refine, and validate a subjective questionnaire for measuring engagement based on a set of attributes: focused attention, per-

ceived usability, aesthetics, durability, novelty, and felt involvement [OT10]. For this work, we define *engagement* as the combination of a reader’s subjectively reported levels for each of these six attributes set forth by O’Brien and Toms. To build off of this rigorously evaluated scale on engagement [OT10], we selected a subset of questions to tailor the questionnaire to our domain for the crowdsourced study discussed in Section 5.

3. Visual Narrative Flow Design Space

Media experts generally agree that story or *narrative flow* is a difficult concept to define [Dig11]. According to Campo [Cam14] and Hill [Hil12], flow is what captures a reader’s attention and carries her beyond the written text into the story itself. In line with this, Brechman has proposed a model for narrative flow based on Csikszentmihalyi’s Flow theory [Bre10]. Brechman argues that there are four unique dimensions in narrative processing: focus, cognitive arousal, affective arousal, and absorption. This phenomenological perspective indicates that stories that “flow” are likely to be more engaging, since readers are likely to lose track of the activity of *reading* itself. To induce readers in this state of flow, the language used to convey the story should be as “transparent” as possible, so that the reader can connect directly with the author’s narrative voice [Sha09].

In the case of visual, interactive, data-driven stories, the use of language is not limited to text. It extends to the types of interactions the reader has with the content, e.g., how her input can influence the pacing of the story; the type of graphics and visual cues that are used; and the congruence between the story and the feedback she receives when navigating through the story. Based on this, we define **visual narrative flow** as the congruence between flow-factors, i.e., 1) the way a reader navigates the story, 2) the visual components of the story, and 3) the type of visual feedback the reader receives; along with the nature of the data and facts that the author wants to communicate. For this work, we focus specifically on the congruence of flow-factors. We introduce a design space for visual narrative flow that describes seven flow-factors: navigation input, level of control, navigation progress, story layout, role of visualization, story progression, and navigation feedback.

3.1. Methodology

To identify flow-factors of visual narrative flow, we constructed a corpus of 80 recent stories on the web (Table 1). The majority (54/80) of this corpus is stories published by leading media outlets and posted on websites and blogs such as Gapminder, EagerEyes, and FlowingData. The remaining (26/80) comes from heterogeneous sources such as links posted by researchers and practitioners in the field via social media and other visually interactive stories. We selected stories that contained visualizations and text as well as different kinds of visual narrative flow, excluding data videos to scope this work. We utilized additional selection criteria to capture a broad range of stories with different intents (e.g., to inform, to enjoy, to educate) and animations (e.g., text motion, highlighting, animated transitions).


To extract meaningful flow-factors, we performed a series of sessions of individual consumption for a set of stories and discussed characteristics emerging from them. We started with 10 stories, labeled *S-1* through *S-10* in Table 1, and we identified codes based on these different kinds of visual narrative flow. In addition to consuming these stories ourselves, we also implemented our own story to experiment with changing these flow-factors interactively. Over six months, we extended the corpus up to 80 stories to test the robustness of the codes, and the primary author iteratively adapted these codes into the formalized flow-factors and properties shown in Table 1. For further details, links to the individual stories, and to explore an interactive version of this corpus, please see Supplemental Materials at <https://narrative-flow.github.io>.

3.2. Factors of Visual Narrative Flow

Here we introduce seven factors that contribute to visual narrative flow along with illustrations of the various properties for each flow-factor. This design space framework captures aspects of flow like a reader’s input, connection of story components, and the visual feedback. Previous work by Segel and Heer focused on high-level story components (e.g., animation, progress bars) along with story genres and narrative approaches [SH10]. However, these flow-factors build upon their work by breaking down these properties to characterize and explore a broader range of visual data-driven stories than otherwise initially possible.



Navigation input is how a reader interacts to progress through a narrative visualization. For example, an author can choose to use scrolling input to move down a document, like in *S-51* in Table 1. Another input mechanism is a button, which corresponds to a click, tap, or keyboard press e.g., *S-4*. Another element, though rare, is a slider, where a reader can select and drag to choose where they wish to be in the story. A common theme across all flow-factors is that multiple properties can be combined in a story. For example, *S-79* combines both buttons with a slider in a timeline chart. Such a hybrid approach shows how the line between a stepper and scroller does not need to be rigid in terms of reader input.



Level of control corresponds to how much control a reader has over the motion or animated transitions of story components. For these levels of control, a reader can have *discrete* control if they trigger motion playback like the scroller in *S-2*, or *continuous* control if they can play through the keyframes or time points of that motion like the machine learning scroller *S-1* and in Figure 2. It is also possible for a *hybrid* style to combine or support aspects of




Figure 2: Two examples of different visual narrative flow in a data visualization story. The scroller shown in (a), by Yee and Chu, walks through a story to teach a basic concept of machine learning [YC15], where scrolling not only moves down the page but moves visualizations and continuously controls their linked animated transitions. We transformed this into a stepper narrative flow, shown in (b) which uses buttons for navigating the story across the story text with timed animated transitions.

both, such as the scrolling story in *S-8* with a timeline plot where points can be clicked to navigate. We break down levels of control based on the following categories: text, visualizations, and animated transitions. Text and visualizations can move or fade in or out within the page, and this motion is described by level of control for those elements. An animated transition is defined here as more specific, data-relevant motion that preserves data context across or within visualizations.



Navigation progress describes how the reader perceives their placement within the entire story. Not all stories may show navigation progress, relying on the implied progress of a scrollbar. Otherwise, stories may showcase this progress in a variety of ways. A common way is to represent steps with dots like the stepper in *S-4*, and another method utilizes numbers or text for story steps like in *S-6*. Authors also use visualization to convey story progress, such as a path on a small multiple map as utilized in *S-18*. These progress widgets can also be combined with button input for navigation.



Story layout captures both the type of layout model and the number of columns used in the story. Stories commonly utilize a single- or two-column approach, but these can be mixed as in *S-44*, which changes the layout across sections. The two kinds of layout models are either a document (e.g., *S-1*) or a slideshow (e.g., *S-4*), but hybrids (e.g., *S-3*) also occur where this example looks like a slideshow stepper but uses different animations and scrolling for input. These layout models are similar to steppers and scrollers.



Role of visualization examines the purpose and the part that visualizations serve with respect to the entire story. There are aspects of this flow-factor which may get decided when first creating a story based on the author's intent, but the role of visualization helps determine the visual narrative flow based on how the visualizations and text interact with each other. From the visual, interactive stories we surveyed, 29 stories have text and visualizations playing an equal role in telling the stories, like the interplay between the text and maps in *S-11*. However, other stories have visualizations serve as a figure to help convey part of the story, like the progress map used in *S-18*. Lastly, visualizations can drive telling a story, with text primarily annotating the visual story, such as the animated chart in *S-16*.



Story progression categorizes the possible story paths which can occur in a story, such as linear story points, skipping between, or more complex paths. The most common story type we observed is linear, where a reader is guided through each story point in order like in *S-21*. However, a variation of this approach is a linear skip, where readers can jump backwards or ahead – often included if a navigation progress widget is present such as the stepper in *S-9*. A story can also contain more complex story paths, such as a tree style or graph to include cycles or loops, and a good example is *S-29*, a visual, interactive story about how neurons work.



Navigation feedback combines animated transitions with additional animations of story text or other components, such as fading or movement. This factor is all about showing to readers that their input affects the story. For example, it is possible for both the text and visualizations to transition or move on the page simultaneously, or in sync, such as *S-5* where visualizations and text move down the page together. However, these animations can also occur one before the other, just the text or just the visualizations, so they swap their order like in *S-11* which swaps between maps and the text. Animated transitions that are not tied to data can show change using motion or fading, and these animations can occur in different parts of the story interface: the text, the progress widget, or the visualizations. The story in *S-80* uniquely combines all three kinds of feedback as a reader advances.

3.3. Expressiveness of the Design Space

The expressivity of the design space can be evaluated by looking at the model's descriptive and generative power [BLM04]. One type of visual narrative flow is the **stepper**: linear skip progression with button or swipe input, discrete control over elements, slideshow layout, and a progress widget. Conversely, **scrollers** commonly have linear progression, continuous control over elements, document layout, and no progress widget. There are also different kinds of scrollers, some of which use discrete control to trigger animations (e.g., *S-2*, *S-3*, *S-5*) while others do so continuously based on the scroll position (e.g., *S-1*, *S-13*, *S-18*). Thus, the design space characterizes differences between these discrete and continuous scrollers that were previously called one category in the community, demonstrating the framework's descriptive power.

Furthermore, there are a variety of hybrid properties of visual narrative flow that this model is able to capture. For example, Figure 3a shows a story (*S-51*) with a document model that scrolls text down the page, but animated transitions are triggered discretely and only for parts of the story. Despite using scrolling input, scrolling moves just the text and visualizations since there is discrete control over the animated transitions of the data visualizations. Another example (*S-79*) is shown in Figure 3b where a reader can click buttons to navigate to annotated sections of a temporal line chart but can also interactively move along the slider to explore the chart and skip over parts of the story. Simply calling either story a stepper or a scroller, or one of Segel and Heer's genres, would fail to capture many of the subtle and novel factors that lie within. Thus, this design space is a first step towards building knowledge on what factors shape the reading experience for visual data-driven stories. While it may certainly be expanded and flow-factors added as novel experiences appear, each factor we identified can be studied independently and inform future design.

To comprehend how flow-factors shape the reading experience, we implemented different visual narrative flows for a simple story conveying how a numerical table can be represented as a bar chart. An interactive demo is included in Supplemental Materials.[†] This demo enables dynamic switching between different combinations of flow-factors, thus allowing a reader to experience how each factor impacts the story delivery. For example, we explore varying the

level of control over the animated transitions in the story when the reader scrolls. In the "continuous" case, scrolling the pages enables the reader to control every step of an animated transition. In the "discrete" case, scrolling over a certain location in the page triggers an animated transition playback. By clicking a button in the interface, one can experience each of these. This demo supports interactive, visual changes to the location and duration of these triggers on the left. By constructing this example story and implementing variations of specific factors of the visual narrative flow such as level of control, we showcase the generative power of the design space to create different reading experiences.

4. Exploratory Studies

We conducted two exploratory studies to understand how visual narrative flow impacts the reader experience as well as to investigate how to capture this impact. We focused on two specific kinds of flows: steppers and scrollers, to mimic real-world reading experiences. For this exploration, our goal was to observe usability issues and reader preferences between the two different flows. Two studies employed different protocols: an online pilot survey and in-person observations and interviews.

For these exploratory studies, we utilized the story, "The Visual Introduction to Machine Learning" by Yee and Chu [YC15]. As shown in Figure 2, this story is a scroller, and we adapted the story and its various story points into another type of flow: a stepper. We selected this story because it won several awards, had many views, and was generally discussed as a good example of a scroller. In addition, we believe that it could transfer well to a stepper (slideshow) model. This story features technical insights and complex data visualizations in a longer format. Thus we hypothesized that the visual narrative flow could substantially impact the story usability and reader preference. To experience this story and our stepper variant, please see Supplemental Materials.[†]

4.1. Pilot Survey

We recruited eight participants in our lab to read through the machine learning story and instructed them to "explore" different reading experiences by toggling between the two types of narrative flow. After reading through the story, participants answered an online questionnaire on the usability of each approach and individual preference questions, selected from the system usability survey (SUS) [Bro96]. Survey materials used are included in Supplemental Materials.[†] The participants all have data visualization knowledge, different operating systems (Linux, Mac OS, Windows), various browsers (Firefox, Chromium, Chrome), display resolutions (1373x735 to 2560x1464), only three typically read visually animated stories, and five had read this story before.

The results from this study showed a mixed set of preferences for each participant. Three participants preferred the stepper while five preferred the scroller. Both visual narrative flows were also considered moderately usable, scoring 62 and 60 out of 100, respectively.

[†]. Supplemental Materials can be accessed via: <https://narrative-flow.github.io>

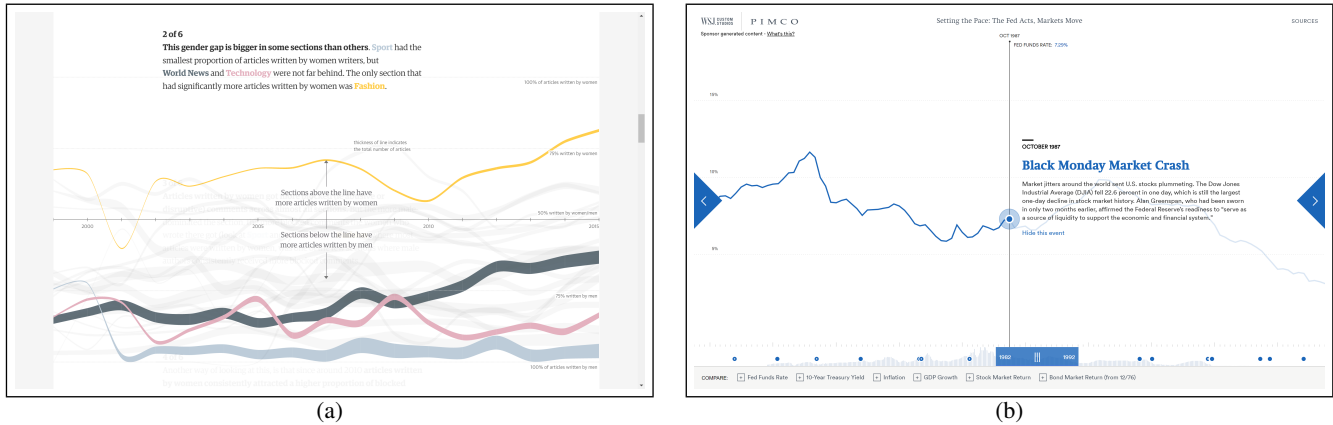


Figure 3: The stories shown here are a scrolling (a) story which contains animations in the story that trigger discretely [GMA*16] and a stepper (b) which contains a slider for navigating through time points [Stu16].

Thus, the differences we collected between both approaches varied mostly by subjective preference. For example, responses to open-ended questions pointed that three readers preferred the progress bar or the arrow keys enabling efficient story progression, while two others reported that scrolling required less effort. By analyzing interaction logs from the stories, we found that participants generally spent more time and interacted more in the flow they had rated higher.

4.2. Observations and Interviews

To gather deeper insights on the impact of visual narrative flows on preferences and usability, we conducted observations and semi-structured interviews with 10 participants. The study lasted 30 minutes on average, and we gave a \$10 compensation. To allow participants to experience both conditions, we broke the story into three different chapters. Participants experienced the two conditions in a random order, and were asked to pick a flow for the final chapter based on their preference. We improved the usability and breakpoints of the stepper version from the previous study.

Participants read the story on the same touchscreen tablet device. After each chapter, participants completed a longer usability questionnaire based on SUS [Bro96]. After reading the entire story, we conducted a semi-structured interview to gather preferences and feedback that the primary author analyzed iteratively with 3-5 passes on the interview notes along with informal coding techniques. Participants did not have formal training in data visualization, six participants had machine learning knowledge, and three participants had seen this story before. Materials used in the study are included in Supplemental Materials.

The results of this study shed more light on reader preference. For usability, the stepper scored 76 and scroller scored 80 out of 100, so both types of narrative flow were rated strongly usable, likely due to improvements to our implementations. Contrary to the previous study, only two (out of 10) readers preferred the stepper; the remaining eight preferring the scroller version, although two of them were not a strong preference. The interviews revealed

that the touch interface may have impacted the experience. Two participants commented that if swiping gestures (rather than tapping arrows) were implemented in the stepper version, they would have likely favored the stepper. These comments hint that different properties of visual narrative flow can affect reader preference.

All but one participant agreed that these visual narrative flows impacted their story reading experience. We observed that all participants in the scroller condition, would scroll back and forth to align text or replay animated transitions in certain positions on the screen. Three participants commented that it was hard to find the “perfect” view while reading the story. In contrast, we observed less interaction with the stepper. The participants generally hovered their fingers over the buttons to facilitate navigation and did not replay animation back and forth as much. We noted more expressions of joy and pleasure, such as laughter, from three participants in the scrolling experience. As participants stated: “scrolling allowed me to see the transformations more fluidly” (P4), “scrolling is sexier in some ways, it is unique ... and just more fun” (P10), and “my preference would be a stepper with slides ... but scrolling goes over these expectations, since it ... helps reasoning and understanding” (P5). Six participants expressed that the continuous control over animations afforded by scrolling was very engaging, though this could be an effect of novelty or self-reporting on engagement.

5. Crowdsourced Study on Engagement

To study how different flow-factors of visual narrative flow affect readers’ engagement, we conducted a larger scale crowdsourced study with 240 participants using Amazon Mechanical Turk.

5.1. Conditions and Hypotheses

We selected four conditions to study in detail. We identified two baseline conditions, one with only text and another with only static visuals, to first see if there is a measurable benefit to scrolling stories with visualization or with animation. Lastly, we included a stepper narrative flow to explore measuring the difference in en-

agement we witnessed in the exploratory studies. Specifically, our conditions were:

- *text*: a text-only story (baseline 1)
- *visual*: text paired with static visual images (baseline 2)
- *stepper*: text paired with visualizations and animated transitions via a stepper
- *scroller*: text paired with visualizations and animated transitions via a continuous scroller

For the 2 baseline conditions, we hypothesized that the inclusion of visualizations (**H1**) and animation (**H2**) would increase the visual appeal, attention, novelty, and felt involvement (all attributes of engagement) for readers. Furthermore, from our observations of readers, we hypothesized that transitions of dynamic data are more engaging using continuous control than discrete one (**H3**). In other words, readers from previous studies expressed that continuous scrolling was more of a gimmick until they experienced the final story chapter which uses continuous scrolling to show the timesteps of an algorithm. Thus our hypotheses were:

- **H1**: Visualizations contribute to make the data-driven story more engaging.
- **H2**: Animated transitions contribute to make the data-driven story more engaging.
- **H3**: Pairing dynamic transitions with continuous control contributes to make the data-driven story more engaging.

5.2. Study Design

We selected the machine learning story [YC15] as in previous studies because of its length and the various types of visualizations and transitions it included. Both the text and visual conditions used a document layout and scrolling navigation input. In the text condition, textual descriptions replaced the visualizations while the visual condition utilizes static screenshots of the visualizations. The scroller is the original story, and the stepper is our implemented version, which contains a progress widget, does not animate text, and controls animations by discrete button presses.

We broke apart the story into two chapters based on types of animated transitions used, and each participant went through each chapter and then filled out a survey at the end. Each participant received only two unique conditions, one for each story chapter. Questions were reserved until the end of reading to avoid breaking the reader's concentration and issues of memory recall of responses that would result in less effective comparison results. The study is balanced in design, with the four conditions crossed with the two chapters of the story resulting in 12 combinations.

For measuring engagement, we adapted a validated questionnaire from O'Brien and Toms [OT10] containing 14 questions on reader-perceived engagement across attributes such as usability to attention to aesthetics to novelty. Each statement then received a 5-point Likert response from the participant. These statements were duplicated across the two conditions (or chapters) that the participant had read. We randomized the order of questions, but questions per chapter were adjacent to support effective comparison. For analysis, several of the questions were negative attributes for

engagement, so their responses were flipped accordingly. All materials and conditions used for the study are included in Supplemental Materials.

5.3. Participants

In total, 240 participants read through the story, and they were compensated \$2.31 for their time. There were 20 participants per pair of conditions, or 40 per balanced set. Participants were recruited via Amazon's Mechanical Turk (at least 98% HIT approval rate, at least 100 approved HITs, and English-speaking countries only). We asked a series of questions to ensure that no one had seen this particular story before and that participants were actively following instructions and paying attention to the content of the story, both in the text and the visuals. Participants have a varied educational background (46.2% high school or some college, 39.2% with a bachelor's degree, 14.6% with masters or beyond), 58.3% have never taken a Computer Science course, 87.0% were unfamiliar with machine learning, 74.6% read stories on the web multiple times weekly or more, and 12.9% contain subscriptions to a popular news site. They used a variety of input devices (75.4% mouse, 22.9% trackpad, 2 touch devices, 1 pen, and 1 trackball).

5.4. Results

We performed a linear mixed effects analysis using R [R D08] and lme4 [BMW15] to study the relationship between different types of narrative flows and reader-perceived engagement (all 14 questions). As fixed effects, we used both the four narrative flows and chapter (without the interaction term) in the model. For random effects, we incorporated intercepts for participants and engagement questions, as well as a by-participant and by-question random slope model for the effect of engagement. Residual plots were visually inspected and no extreme deviations from homoscedasticity or normality were found to violate model assumptions. Despite the effect of engagement being captured using a Likert scale, the underlying concepts are likely continuous in nature, and it is accepted in several fields to utilize linear models for such ordinal data without succumbing to a negative bias as with other approaches [BS11]. The p -values reported here were obtained through likelihood ratio tests of the full effects model to one without the effect of different visual narrative flows.

Figure 4 shows the results of the model, which contains the average engagement score for all 14 questions. According to the model, the different conditions tested for this story affected the engagement score significantly ($p < .001$, $\lambda^2(3) = 30.71$), supporting **H1**. In other words, readers ranked the engagement of stories with visualizations higher than the first baseline, text-only story. It is important to note that these effects are small, since they are averages of the 14 questions on engagement. Certain questions had a stronger effect across the visual narrative flows, such as visualizations scoring higher on an aesthetics question: "This reading experience appealed to my visual senses." We provide an engagement model for each question and condition in Supplemental Materials.

The animated transition conditions scored significantly higher on engagement than static visualizations ($p < .001$, $\lambda^2(2) = 18.04$), supporting **H2** and suggesting that animated transitions increase




Figure 4: This mixed-effects model represents the average engagement score and 95% confidence interval of all 14 questions for 240 participants across the conditions. The model shows increased engagement when using visuals and especially when using animated transitions, but the effect of the other visual narrative flows, stepper and scroller, were not significantly different.

the reader-perceived engagement. Interaction effects of the engagement per chapter of the story were tested, but none were found to be significant so were left out of the model. Additionally, we found a significant effect on the chapter of the story ($p = .017$, $\lambda^2(1) = 5.72$), where the second chapter received, on average, a higher engagement score (0.10). In other words, the animated conditions scored even higher in engagement for the second chapter of the story, which contained dynamic transitions. A question on novelty scored higher for flows with animations: “The reading experience was different from a typical online reading experience.” Subsequently, animated transitions, such as navigation feedback, showed a measurable benefit for reader-perceived engagement.

Regarding the stepper versus scroller debate, we did not find a significant difference in engagement via our questionnaire, failing to support **H3**. While steppers scored higher for engagement on average, the difference over scrollers was not significant in the model. Thus, we are unable to conclude if continuous control over dynamic transitions, via scrolling, improves engagement measured in the second chapter of the story. We note that the difference did vary by reader preference and only for certain questions, such as those regarding usability. While our findings do not support **H3**, a carefully controlled user study may be able to investigate and measure this effect.

Furthermore, we conducted an analysis on the preferred conditions selected by the participants, shown in Figure 5. Note that each participant saw two conditions, the order shown in the table. They ranked which of the two they preferred for the story, or possibly none. We found that the stepper and scroller were largely preferred over other conditions by almost twice as many participants. However, a large portion of the participants overall did not have a preference between the two conditions they experienced. Lastly, Figure 5 shows split in preference across participants for both stepper and scroller. By inspecting the detailed breakdown of preferences, more participants preferred animated transitions and scrollers for the second chapter of the story.

In addition to questions on engagement, we asked participants to complete five comprehension questions, which varied in difficulty from terminology or concept recall to complex application of a concept to a new problem. Overall, participants comprehended the story well, scoring on average 4 out of 5. We did not find any major differences in comprehension across condition pairs.

6. Discussion and Limitations

We learned that individual differences may play a key role for different forms of visual narrative flow. While a navigation progress widget generally adds more information, only 2 out of 10 interviewees preferred knowing these steps of the story to help them manage their time and attention. Additionally, three interviewees expressed frustration with aligning transitions and text when reading a continuous scroller. When piloting the crowdsourced study, we first gave readers one story condition and compared engagement between participants, but we discovered that the subjective nature of the questions made results extremely variable. Therefore, we decided to adopt a within-subjects design for two conditions at a time. While these individual differences in subjectivity added a layer of complexity, we believe that it is important to discover novel ways to effectively measure engagement despite these challenges.

While individual differences affect reading experiences, there are still best practices recognized by the community for improving a story’s visual narrative flow. Navigation feedback can help guide readers by showing immediate change, and it is often detrimental to have navigation input feel laggy, unresponsive, or overridden (scrolljacking) without appropriate feedback. Supporting different kinds of reader input is complex, and this is further complicated by the variety of devices used to consume stories nowadays, from mobile phones to tablets. From our interviews using a tablet device, two readers recommended modifying the stepper for touch and to utilize swiping between steps. On the other hand, Mike Bostock suggests that “rapid, incremental, reversible scrolls are more usable than slow, animated swipes” [Bos14], so there may be some additional trade-offs to consider across different types of devices. Advice from the design community suggests that clicking is often about making decisions (e.g., to navigate through search results to find a target), whereas scrolling is more ideal for exploration and discoverability [Por14]. Choosing an input for the narrative flow may depend on the author’s intent or even the length or complexity of a story, since these inputs can be combined in a variety of ways.

A limitation of the results of our crowdsourced study is a lack of generalizability. While a measurable benefit for engagement was found for using animated narrative flows, this could be localized to this particular story. Another limitation of this study is the variability of the conditions, in particular the stepper and scroller, which vary in many flow-factors (progress widget, level of control, input, and layout model). Additionally, there are still factors of visual narrative flow which could be measured for their impact on engagement. Furthermore, the 14 questions we selected for engagement may not be a sufficiently detailed metric to fully capture reader engagement, and other measures of engagement exist beyond participant reported comparisons, such as time taken or number of user interactions [BDF15, SES16].

While the results of our crowdsourced study on reader-perceived engagement indicate a measurable benefit for stories that incorporate visuals and animated transitions, comparing flows of a stepper and scroller are inconclusive and could be studied further. For example, researchers could expand the questions we utilized in this study to encompass more attributes of engagement and validate this refined scale more rigorously [OT10]. Another interesting aspect to consider is the type of story, to see if these hypotheses and ques-

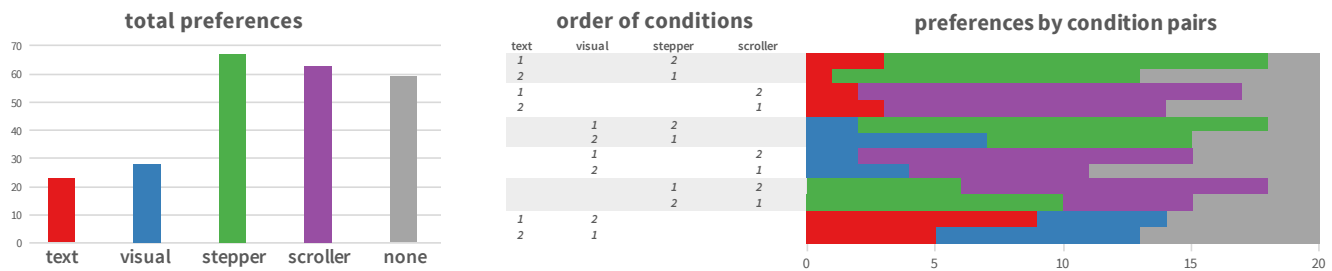


Figure 5: Participant preferences across all of the pairs of conditions ($N = 240$). There were 20 participants per pair, and balanced based on which condition was first or second. On the left, preference totals across all conditions emphasize that participants largely preferred conditions with visualizations and animation (stepper, scroller), otherwise they had no preference.

tions for engagement apply broadly. Another study could systematically break apart the flow-factors, exploring the factors individually, such as a discrete scroller compared with a continuous one. These studies could be conducted in a crowdsourced environment, in a laboratory setting to reduce variability, or embedded within real-world stories to increase the ecological validity.

7. Conclusion and Future Work

In this work, we shed light on *visual narrative flow* by introducing a novel design space describing the flow-factors that shape the reading experiences of data visualization stories. This design space contains seven factors we identified through analyzing a corpus of 80 stories. Visual narrative flow encompasses aspects of the reading experience such as how a reader navigates a story and how this navigation is coupled with visual feedback like animated transitions.

To build initial knowledge towards understanding what makes data stories with visualization compelling and engaging, we conducted a series of studies comparing different visual narrative flows. The first two exploratory studies we conducted support the claim that flow has an impact on readers' preference, and qualitative comments suggested that it impacts engagement as well. The third crowdsourced study measured reader engagement across different flows. Results from this study indicated that visualizations and animated transitions positively impact reader-perceived engagement, although other results of the study are inconclusive and require further study between steppers and scrollers.

Our studies focused on specific combinations of flow-factors for one story, but further research could generalize these results across stories and systematically analyze focused subsets of the visual narrative flow design space. Future studies in this space could explore other metrics for measuring engagement to more rigorously examine and measure the role engagement plays in these data visualization stories, validating and expanding upon our findings here. We believe that systematically identifying factors that increase engagement and make reading experiences more compelling is crucial to expand our knowledge on data-driven visual storytelling and to inform future story design and authoring tools.

8. Acknowledgments

We thank the Visualization Design Lab at the University of Utah for their feedback. This work is sponsored in part by the Air Force Research Laboratory and the DARPA XDATA program.

References

- [AHL*15] AMINI F., HENRY RICHE N., LEE B., HURTER C., IRANI P.: Understanding data videos: Looking at narrative visualization through the cinematography lens. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)* (2015), ACM Press, pp. 1459–1468. 1
- [AHL*17] AMINI F., HENRY RICHE N., LEE B., MONROY-HERNANDEZ A., IRANI P.: Authoring Data-Driven Videos with DataClips. *IEEE Transactions on Visualization and Computer Graphics* (2017), 501–510. 1, 2
- [Ais16] AISCH G.: Information+ Conference: Data visualization and the news. <https://vimeo.com/182590214>, 2016. 1
- [BDF15] BOY J., DETIENNE F., FEKETE J.-D.: Storytelling in information visualizations: Does it engage users to explore data? In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)* (2015), ACM Press, pp. 1449–1458. 2, 9
- [BKH*16] BACH B., KERRACHER N., HALL K. W., CARPENDALE S., KENNEDY J., RICHE N. H.: Telling stories about dynamic networks with graph comics. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)* (2016), ACM Press, pp. 3670–3682. 1
- [BLM04] BEAUDOUIN-LAFON M., MICHEL: Designing interaction, not interfaces. In *Working Conference on Advanced Visual Interfaces (AVI)* (2004), ACM Press, pp. 15–22. 5
- [BMBW15] BATES D., MÄCHLER M., BOLKER B., WALKER S.: Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* 67, 1 (2015), 1–48. 8
- [Bos14] BOSTOCK M.: How To Scroll. <https://bost.ocks.org/mike/scroll/>, 2014. 2, 9
- [Bre10] BRECHMAN J. M.: Narrative “flow”: A model of narrative processing and its impact on information processing, knowledge acquisition and persuasion. 3
- [Bro96] BROOKE J.: Sus-a quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7. 5, 7
- [BS11] BAUER D. J., STERBA S. K.: Fitting multilevel models with ordinal outcomes: Performance of alternative specifications and methods of estimation. *Psychological Methods* 16, 4 (2011), 373–390. 8
- [Cam14] CAMPO A.: Creating and Maintaining Narrative Flow: A Key Element of Story-Telling at Trial.

- <http://www.thefederation.org/documents/02.Narrative%20Flow%20Story-Telling.pdf>, 2014. 3
- [CDBF10] CHEVALIER F., DRAGICEVIC P., BEZERIANOS A., FEKETE J.-D.: Using text animated transitions to support navigation in document histories. In *Proceedings of the Conference on Human Factors in Computing Systems (CHI)* (2010), ACM Press, pp. 683–692. 2
- [CRP16] CHEVALIER F., RICHE N., PLAISANT C.: Animations 25 Years Later: New Roles and Opportunities. In *International Working Conference on Advanced Visual Interfaces (AVI)* (2016), pp. 280–287. 2
- [CU95] CHANG B.-W., UNGAR D.: *Animation: From cartoons to the user interface*. Tech. rep., 1995. 2
- [Dig11] DIGEST.COM W.: What Writers Mean by “Flow”. <http://www.writersdigest.com/qp7-migration-books/on-writing-fiction-excerpt>, 2011. 3
- [EMJ*11] ELMQVIST N., MOERE A. V., JETTER H.-C., CERNEA D., REITERER H., JANKUN-KELLY T.: Fluid interaction for information visualization. *Information Visualization* 10, 4 (2011), 327–340. 2
- [FBR14] FENG C., BARTRAM L., RIECKE B. E.: Evaluating affective features of 3D motionscapes. In *Proceedings of the Symposium on Applied Perception (SAP)* (2014), ACM Press, pp. 23–30. 2
- [GMA*16] GARDINER B., MANSFIELD M., ANDERSON I., HOLDER J., LOUTER D., ULMANU M.: The dark side of Guardian comments. <https://www.theguardian.com/technology/2016/apr/12/the-dark-side-of-guardian-comments>, 2016. 1, 7
- [HDH*13] HULLMAN J., DRUCKER S., HENRY RICHE N., LEE B., FISHER D., ADAR E.: A deeper understanding of sequence in narrative visualization. *IEEE Transactions on Visualization and Computer Graphics* 19, 12 (2013), 2406–15. 2
- [Hil12] HILL B.: How Goes the Flow in Your Story? <http://theeditorsblog.net/2011/10/01/hows-goes-the-flow-in-your-story/>, 2012. 3
- [HR07] HEER J., ROBERTSON G.: Animated transitions in statistical data graphics. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1240–7. 2
- [KB05] KLEIN C., BEDERSON B. B.: Benefits of animated scrolling. In *Extended abstracts on Human Factors in Computing Systems (CHI)* (2005), ACM Press, pp. 1965–1968. 2
- [KMÅ02] KAPTELININ V., MÄNTYLÄ T., ÅSTRÖM J.: Transient visual cues for scrolling. In *Extended abstracts on Human Factors in Computing Systems (CHI)* (2002), ACM Press, pp. 620–621. 2
- [Kos16] KOSARA R.: The Scrollytelling Scourge. <https://eagereyes.org/blog/2016/the-scrollytelling-scourge>, 2016. 1, 2
- [KSJL14] KWON B. C., STOFFEL F., JÄCKLE D., LEE B.: VisJockey: Enriching Data Stories through Orchestrated Interactive Visualization. *Computation+Journalism* (2014). 2
- [MKK15] MAHYAR N., KIM S., KWON B. C.: Towards a taxonomy for evaluating user engagement in information visualization. In *Workshop on Personal Visualization: Exploring Everyday Life* (2015), vol. 3. 2
- [OT10] O’BRIEN H. L., TOMS E. G.: The development and evaluation of a survey to measure user engagement. *Journal of the American Society for Information Science and Technology* 61, 1 (2010), 50–69. 2, 3, 8, 9
- [Por14] PORTER J.: Scrolling is easier than clicking. <http://bokardo.com/archives/scrolling-easier-clicking/>, 2014. 9
- [RD08] R DEVELOPMENT CORE TEAM: *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2008. ISBN 3-900051-07-0. 8
- [SES16] SAKET B., ENDERT A., STASKO J.: Beyond usability and performance: A review of user experience-focused evaluations in visualization. In *Proceedings of the Beyond Time and Errors on Novel Evaluation Methods for Visualization (BELIV)* (2016), ACM Press, pp. 133–142. 2, 9
- [SH10] SEGEL E., HEER J.: Narrative visualization: Telling stories with data. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (2010), 1139–48. 1, 2, 3
- [SH14] SATYANARAYAN A., HEER J.: Authoring narrative visualizations with Ellipsis. *Computer Graphics Forum* 33, 3 (2014), 361–370. 2
- [Sha09] SHADOWWOLF: The Importance of Narrative Flow. http://shifti.org/wiki/The_Importance_of_Narrative_Flow, 2009. 3
- [Shn82] SHNEIDERMAN B.: Direct manipulation: A step beyond programming languages. In *ACM SIGSOC Bulletin* (1982), vol. 13, ACM Press, p. 143. 2
- [SLHS16] STOLPER C. D., LEE B., HENRY RICHE N., STASKO J.: *Emerging and Recurring Data-Driven Storytelling Techniques: Analysis of a Curated Collection of Recent Stories*. Tech. Rep. MSR-TR-2016-14, Microsoft Research, 2016. 1
- [Stu16] STUDIOS W.: Setting the Pace: The Fed Acts, Markets Move. <http://www.wsj.com/ad/pimco-rateandreaction.html>, 2016. 1, 7
- [TMB02] TVERSKY B., MORRISON J. B., BETRANCOURT M.: Animation: Can it facilitate? *International Journal of Human-Computer Studies* 57, 4 (2002), 247–262. 2
- [Tse16] TSE A.: Why We Are Doing Fewer Interactives. <https://github.com/archietse/malofiej-2016/blob/master/tse-malofiej-2016-slides.pdf>, 2016. 1, 2
- [YC15] YEE S., CHU T.: A Visual Introduction to Machine Learning. <http://www.r2d3.us/visual-intro-to-machine-learning-part-1/>, 2015. 4, 5, 8