Narrative Design Patterns for Data-Driven Storytelling

Benjamin Bach Moritz Stefaner Jeremy Boy
Steven Drucker Lyn Bartram Jo Wood
Paolo Ciuccarelli Yuri Engelhardt
Ulrike Koeppen Barbara Tversky

February 13, 2017

Abstract

This chapter introduces the concept of *narrative design patterns*, which aim to facilitate the shaping of compelling data-driven stories. There are many different ways storytellers can narrate the same story, depending on their intentions and their audience. Here, we define and describe a set of these narrative design patterns that can be used on their own or in combination to tell data stories in a myriad of ways. We analyze eighteen of them, and illustrate how they can help storytellers think about the stories they want to tell and the best ways to narrate them. Each pattern has a specific purpose, for example, engaging the audience, evoking empathy, or creating flow and rhythm in the story. We assume storytellers already know what story they want to tell, why they want to tell it, and who they want to tell it to. These patterns should not only facilitate the process of creating compelling narratives, they should also stimulate a wider discussion on techniques and practices for data-driven storytelling.

1 Introduction

The terms "story" and "narrative" are often used interchangeably. They also harbor different meanings in different communities (see Chapter 1 of this volume). Here we are guided by the definitions in Gerard Genette's *The Narrative Discourse* [1], Bruner's *The Actual Mind* [2] and which are summarized in the Media Student's Book [3]. A **story**

is defined as *"all of the events in a narrative, those presented directly to an audience and those which might be inferred,"* [3]. In a wider sense the story are the facts—the data—including characters, places, times, and actions important to the specific plot. A **narrative**, on the other side, can be described as *"the 'telling' of a sequence of [those] events [...] [which] shape the events, characters, arrangement of time, etc. in very particular ways so as to invite particular positions towards the story on the part of the audience."* [3].

Thus, a narrative gives shape to the unfolding events in a story, with the goal of making them clear and compelling to an audience. The narrative can be linear, e.g., representing the actual order of events (the story) chronologically; or non-linear, e.g., circular, where the beginning and end mirror themselves. Non-linear storytelling has been used in highly acclaimed movies, such as Pulp Fiction and Memento, where the narration either jumps between scenes (Pulp Fiction), or tells the story backwards (Memento, following the main character gradually remembering his personal past).

Consequently, a single story can be narrated in an endless number of ways. Beyond the temporal ordering of events, it can be told at a varying pace to create emphasis: slowing down the narrative pace can give the audience time to think about what is being 'told' and to increase the importance of this moment in the story; speeding up the narrative pace can produce an overwhelming effect of information onslaught and reduce the importance of moments in the story. Arguments can be made explicit or kept implicit. The narrative can emphasize comparison and contrasting points of view; it can convey abstract concepts through the use of analogies, making them more concrete; it can address and engage the audience directly with abrupt questions, while interrupting the flow of the story; or it can draw the audience in so that it becomes a part of the story, releasing it with a call-to-action. Each variation is likely to have different effects on audiences' reactions and emotions, as well as on their understanding of the story.

Creating compelling narratives is a labor-intensive process that requires expertise, creativity, iterations, and feedback. It calls for a deep understanding and creative use of media, as well as a sense for explaining, convincing, persuading, and especially engaging audiences. There is a wealth of guidelines and best-practices (e.g., [4, 5]) for creating narratives in the more established disciplines that involve storytelling, such as literature, the performing arts, comics, radio drama, theater,

and cinema. However, these fall short for stories based on data that use graphic representations [6, 7, 8] (visualizations) to explain facts and develop evidence-based arguments. The goal of this chapter—in a wider sense—is to explore such guidelines for data-driven storytelling.

This chapter presents a set of narrative design patterns (*narrative patterns*) that we believe can help storytellers design data-driven narratives that rely heavily, but not exclusively, on data visualization. Narrative patterns help connect the form of the narration with the intent of a story. They are meant for journalists, web and visualization designers, presenters, public speakers, and others trying to shape compelling data-driven stories, and engaging interactive environments. The only requirements for using narrative patterns are to have:

1. a story,
2. an idea of who your audience is, and
3. to know the effect your story and narration should have on that audience. This can include sympathy or distaste, encouragement for action, information, explanation, and so on.

It may further be useful to have an understanding of the benefits and limitations of using different media (e.g., text, numbers, images, visualization, audio etc.), as well as the culture of the audience, as these factors usually drive the ways in which stories are told: different audiences might react very differently to the choice of media, or to the shaping of the story.

As we believe narrative patterns are best explained through examples, we have created an interactive browseable collection, available online at `http://napa-cards.com` (Figure 1). This collection is inspired by *method cards*, i.e., design cards used to stimulate creativity and facilitate ideation, that are commonly used by design practitioners and visual artists (see e.g., IDEO cards [9], oblique strategies [10], design-with-Intent [11], Creativity Cards [12], Service Design tools [13]), or the more recent VizItCards for teaching visualization techniques [14]. We firmly believe in the power of data-driven storytelling, and our experience as academics, designers, and journalists leads us to claim that narrative patterns will undoubtedly facilitate further structuring and discussion of this emerging field.

In the rest of this chapter, we describe our set of narrative patterns, and introduce a higher level structure for grouping them. We then discuss possible combinations of patterns, using examples of data-driven

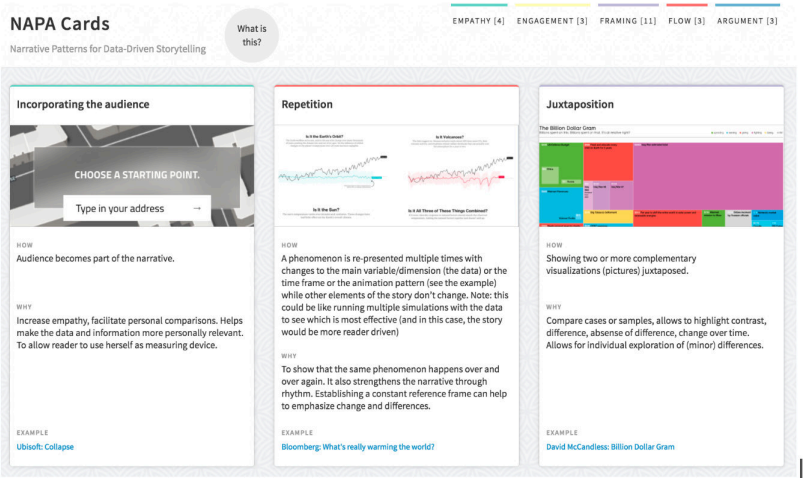


Figure 1: Three example of narrative pattern "cards", browsable online at http://napa-cards.net.

narratives. After that, we detail certain aspects related to the application of patterns. We conclude with an outlook of open questions and directions of future work that results from our discussions.

2 Narrative Patterns

We define narrative patterns as follows:

A narrative pattern is a low-level narrative device that serves a specific intent. A pattern can be used individually or in combination with others to give form to a story.

The specific intents and possible combinations of patterns should be defined by the narrator (or storyteller). These are usually influenced by e.g., the data, the formal setting, or the particular audience and its assumed background knowledge. The narrator may try different unique patterns and/or combnations to see how they fit the intent(s), or use the concept of patterns to analyze existing stories and reflect on their

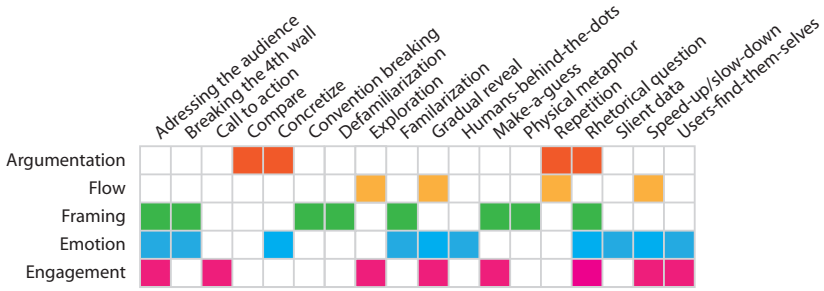


Figure 2: 18 Narrative patterns grouped into 5 major pattern groups.

specific intent(s). Examples of narrative intents range from enlightening audiences, to evoking empathic response, to engaging them to take action, or to questioning their beliefs and behavior. The goal may be to trigger dialog, or simply to immerse people in an enjoyable experience. These are all rather traditional intents, and are common to all kinds of storytelling disciplines (e.g. cinema, theater, etc.). They are a good starting point for data-driven storytelling. Other, more specific intents for data-driven storytelling may include delivering convincing arguments backed with data, explaining a type of data and sensitizing people to their existence and power, or simply educating an easily targeted online audience.

In the context of data-driven storytelling, narrative patterns should not be considered direct implementations of visualization or interactions techniques, nor should they be bound to any specific implementation. In most cases, they are decoupled from the final medium and visual presentation. Patterns can overlap, they can be fuzzy, and in some cases, they may only be applicable to a single type of story.

We initially identified around 40 distinct narrative patterns, based on a great number of data-driven stories that we found in the literature and on the web. Through extensive discussions, we were able to narrow down this initial number to a core of 18 patterns. While we focus mainly on these 18, we stress that any existing story can inspire new patterns, and that our collection is by no means definitive. To better understand these patterns, and to facilitate their discussion and application, we categorize them in five different groups: *argumentation*, *narrative flow*, *framing*, *empathy and emotion*, and *engagement*.

This typology (Table 2) was derived from an initial set of open-

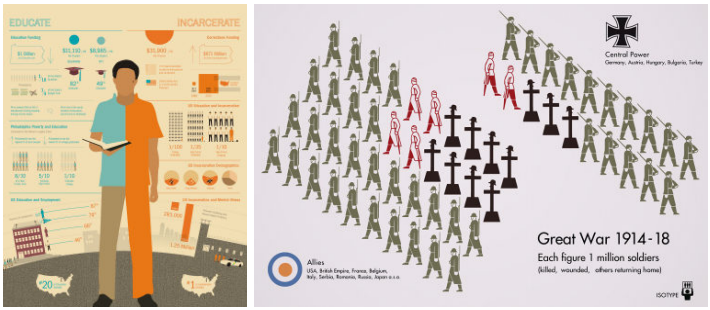


Figure 3: Examples for comparison (left, [15]) and concretize (right, [16]).

ended questions, such as *What types of pattern exist? What is their specific purpose? How are they employed? How can they be applied in a story? Can a story be told with a different set of patterns?*. Note that individual patterns can serve multiple intents, and that our typology is not meant to be exclusive. It merely presents an initial framework to think about narrative patterns, and aims to facilitate their application. In our online collection, patterns can be looked up according to their main intent(s), and each is illustrated with an existing, published example.

2.1 Patterns for Argumentation

Following the Oxford dictionary, argumentation is *"the action or process of reasoning systematically in support of an idea, action, or theory."* Data-driven storytelling is often about argumentation, so it is important to be aware that there are different ways to present, support, reinforce, contradict or discuss a given message. Patterns for argumentation serve the intent of persuading and convincing audiences.

The classical devices of rhetoric involve *logos* (reason, word), *ethos* (character, ideal), and *pathos* (experience, emotion). Though we believe the ultimate goal of data-driven storytelling is to communicate truth (most closely to logos), there are traces of both pathos (see Section 2.4) and ethos in every story, which help connect the narrator with the audience.

One very common pattern of argumentation is **Compare** (Figures 3). The narrator presents two or more datasets, and draws respective

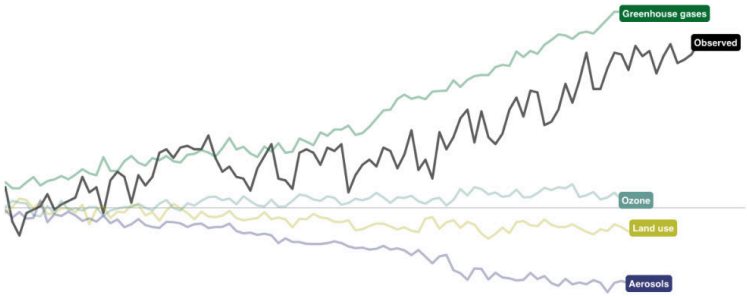


Figure 4: Example for repetition: What really warms the world [23].

conclusions from them. Datasets can refer to categories (e.g., young people, old people), elements in the data (e.g., countries), or temporal evolution. Visually, the comparison can be made through the juxtaposition of graphics (through side by side presentation), or through changes of a single graphic over time (e.g., animation). Comparison allows the narrator to to make the point about equality of both data sets, to explicitly highlight differences and similarities, or to give reasons for their difference. If the graphic is rich enough, the audience can also be given the opportunity to explore differences and similarities for themselves.

Concretize is another pattern that can help build an argument by illustrating abstract concepts with concrete objects. Concretization usually implies that each data point is represented by an individual visual object (e.g., a point or shape), making them less abstract than aggregated statistics. Perhaps the most common example for concretization is the ISOTYPE language [16, 17]. ISOTYPE links a specific pictogram (human, animal, cars, Figure 3-right) to quantities (people fought in war, animals eaten per year, cars produced, water consumed [18]). Concretization has recently been used in interactive visualizations [19, 20, 21], and its effects studied in [**?**, 22].

Repetition can also help argumentation, as it can increase a message's importance and memorability, and can help tie together different arguments about a given dataset. Repetition can be employed as a means to search for an answer in the data. For example, the Bloomberg

piece *"What's really warming the world?"* [23] (Figure 4) shows different measures of possible culprits over time, such as volcanic or solar activity, deforestation, or ozone pollution. While none of these measures correlates with the increase in global temperature, the last example, greenhouse gases, reveals a high correlation.

2.2 Patterns for Flow

Patterns for flow are devices that help structure the sequencing of messages and arguments. They are essential to every story, as they set the order, rhythm, and pace, and help build up the climax and resolution.

One common pattern for flow is **Reveal**, in which elements of the data are progressively revealed, eventually leading to the whole picture and the final argument. Famous implementations of this pattern include the data videos on the U.S. dept [24], and on *The Fallen of World War II* [25]. In this last piece, the narrator guides the audience through various sections of the data, starting with the fallen soldiers from the US, then Poland, then France, etc., while sticking to the temporal order of events and battles. The narrative reaches its climax when the huge number of fallen Soviet soldiers is revealed; as more and more dead soldiers are stacked on top of each other—in the form of pictograms—the camera moves up in a seemingly endless vertical panning. The number of deaths seems infinite. By doing this, the narrator provokes strong emotional responses in viewers. It should be noted though that this particular example uses Reveal in conjunction with Concretize to achieve this highly effective outcome.

Reveal has also been used in Bloomberg's *"What is really warming the world?"* (mentioned above—see Figure 5). Morever, *gradual* reveal and **repetition** are closely related, and both can be seen to create rhythm and flow. Yet, repetition can happen without revealing new insights or data, and gradual reveal can happen without being repetitive. An example is the rhythm of slow-downs and breaks in Hans Rosling's narration of the world development data over 200 years [26]. Rossling slows down the pace of his animations as he focuses on dramatic changes during specific periods (here, WWII), and completely stops the animation to be able to explain the data in detail.

Narrative time in general can be a powerful means. In the example of *The Fallen of World War II* [25], presentation of data elements happens at a constant pace which supports the impression of the number of data elements. The counting of data elements in this particular ex-

ample also provides a **slowing down** effect. The narration is slowed down and focussed on a single aspect of the data, while leaving the user alone with the data. In classical storytelling, a slow-down of the narrative pace can also happen through brakes.

The opposite to slow-downs and breaks is **speeding-up** the narration or the presentation of data, which can be used for a similar effect. The example *U.S. Gun Deaths* [27] presents the observer with over 10.000 data points (individuals) and speeds up their reveal after an initial slow animation that explains how to read the visualization. The visualization uses the metaphor of bullets and ballistic arcs to represent people's expected lifespans, and the age at which they were shot. Due to both the metaphor and the effect of a sped-up, gun-fire like presentation, the result is a strong emotional ensemble that creates an overwhelming effect based on data.

2.3 Patterns for Framing the Narrative

Framing builds the way facts and events in a story are perceived and understood through narration. It feeds on the audience's expectations, but it can also play with those expectations and go against them to create surprise. Patterns for framing can be used to integrate the audience to the story, or conversely to keep it at a distance. They range from elements of the story itself that help the audience relate to the content, and interaction techniques that give the audience control over the way the story unfolds, to purely narrative devices that challenge the audience and the way it is addressed.

Creating a **familiar setting** can help audiences relate to the content; it sets an entry point in the story that the audience can identify with. For example, the *OECD's Regional Well-Being* website [28] allows viewers to first inspect well-being in their region of the world, before moving on to explore and compare the quality of life in other regions. Similarly, inviting viewers to **make-a-guess** about the data or the outcome of a series of events presented in the narrative can enable them to take control over the way the story unfolds. This usually relies on an intelligent and original use of interaction. In the *"You draw it yourself"* example from the NYTimes [29], readers are explicitely asked to draw the trend they expect to see in the data. This can engage them in causal reasoning about the issue addressed, which can stimulate their curiosity, and possibly lead to a game-lke experience.

Conversely, **defamiliarization** challenges audiences' expectations,

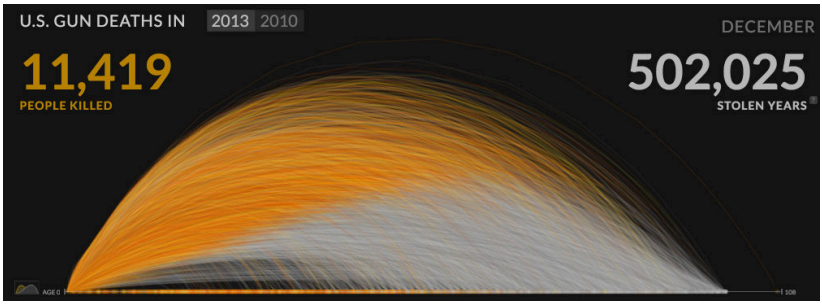


Figure 5: Example for **speed-up**; the way the individual traces are displayed, giving the impression of arms firing [27].

by presenting something familiar in an unexpected way. A typical example is to show a Mercator projected map (e.g. a Google Map) upside-down. This challenges reading habbits, and encourages exploring the map in new ways. This pattern can be used to highlight and question implicit assumptions, and to force thinking differently about well-known facts. **Convention breaking** can also be used to similar effects, by first establishing a graphical convention through e.g., color, scales, or visual rhythm in the narrative, and then disrupting it at a specific point. This can highlight unordinary patterns or subsets in the data presented, and can create a surprise effect in the audience. Finally, introducing **silent data**, i.e., deliberately hidden data, can help create rhythm, like pauses in a piece music. The audience can use these pauses to try to infer the missing data, or to stitch different pieces of information back together.

From a purely visual perspective, a meaningful use of space and **physical metaphors** can help reinforce positive or negative feelings. Frequently used metaphors are: up=good, down=bad, left=retreat, and right=progress [30]. These take advantage of conventions to ease understanding, and can convey cultural or embodied messages.

2.4 Patterns for Empathy and Emotion

Emotion and empathy are critical in storytelling. In fiction, emotional responses help create what the author Guy Gavriel Kay [31] terms *"imaginative empathy"*: resituating and reorienting the reader's perspective while enhancing our ability to understand and share the feel-

ings and experiences important in the story. Empathy engages us with the story content: we pay attention, and we perceive and reflect on the message. In fictional stories, empathy and emotions are evoked through realistic characters and situations; technology has provided us with immersive experiences through sound and stereo-vision. Technology and multimedia can enrich the audience's connection to the story through pure sensory experience, for focus, for compassion, for joy, for excitement. However, data-driven stories—in the ideal case—do not rely on fictional content, or on ludic engagement, but are meant for education and to convey insight. What, then, are the specific applications and purposes of emotions and empathy in data-driven storytelling? How can they be used to enhance understanding, provoke a call-for-action, direct the audience to reflect on the data, or, more generally, enrich the persuasive message of the data story?

Some of the previously mentioned patterns, e.g. **gradual reveal** or **slowing down**, can contribute to emotional responses like surprise or the feeling of being lost. Managing flow can affect a sense of urgency, intensity, optimism, or seriousness. For example, the **speed-up** in the onset of trails in the *U.S. Gun Deaths* piece [27] (Figure 5) heightens the sense of urgency as the visualization progresses. Similarly, the measured pace of transitions in *"What's actually Warming the World?"* [23] suggests a deliberate and methodical tone, invoking a sense of gravity and seriousness. These emotions can help engage the audience with the story's communicative intent.

Concretize, used in combination with other patterns [**?**], is another example that can strongly evoke emotions, grounding an empathic response within that emotional context. Consider the example of the human icons in the NY times visualization of casualties in Iraq [32] (Figure 6). Using a sober black avatar of a mother and child to represent civilians moves the perception from a numerical assessment to the stark comparison of familial comfort, love and security with the horror of war. Another intense example for a story provoking emotion is *Netwars* [34] (Figure 7-right). This example uses a pattern known from comics and movies; in **directly addressing the audience** and talking about his personal computer and his data, the narrator is **breaking-the-fourth-wall**. Originating from theater—the 4th wall is the wall between the stage and the audience, while the other 3 walls frame the scene (Figure 7-left)—breaking-the-fourth-wall can be a strong device for immersion into the story.

Another common pattern used to increase empathic ties between

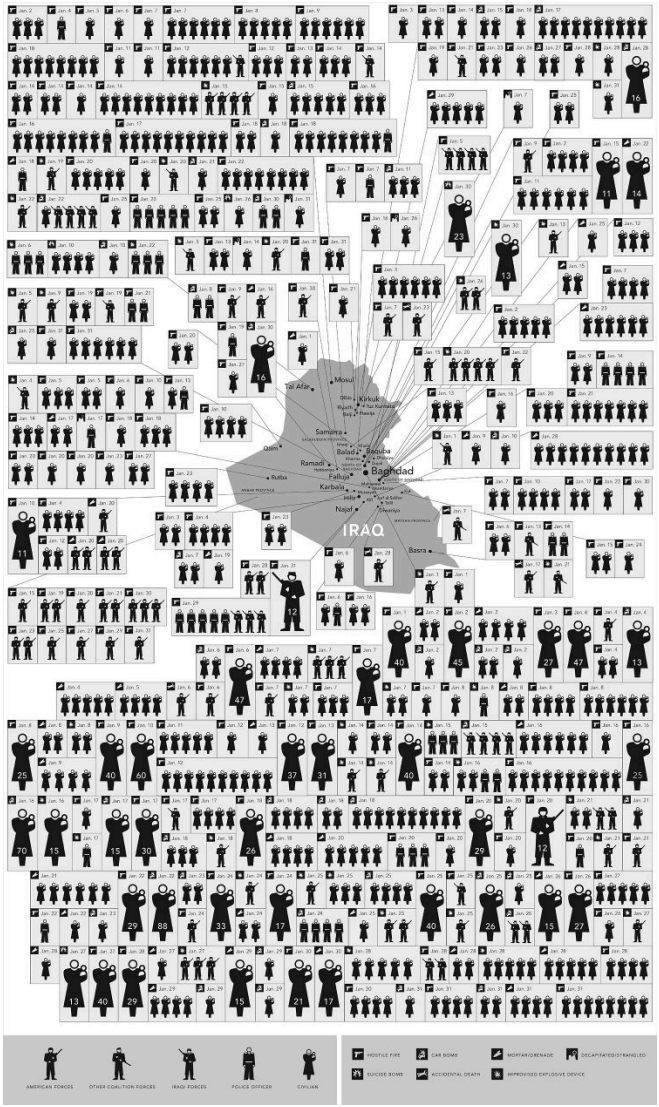


Figure 6: Example for **concretize** to provoke a emotional response: *31 Days in Iraq*, representing war causalities, 2007 [32].



Figure 7: Breaking-the-fourth wall: The Three Pigs by David Wiesner (left) [33], Netwars (right) [34].

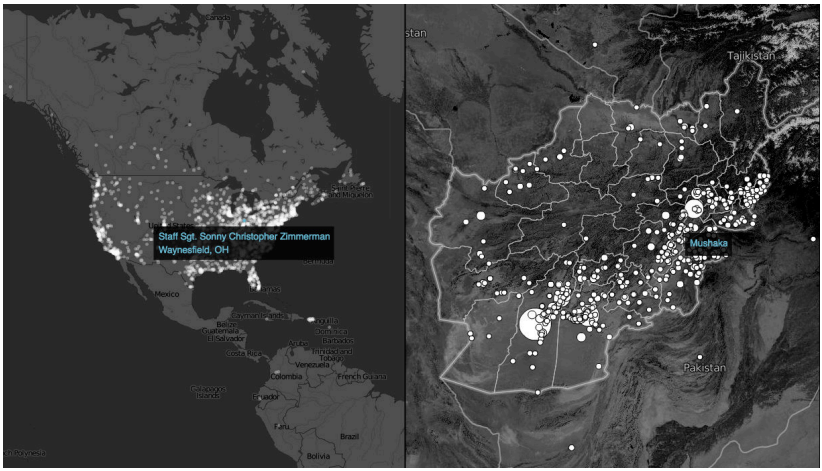


Figure 8: Example for **humans behind the dots**: War Casualties [35].

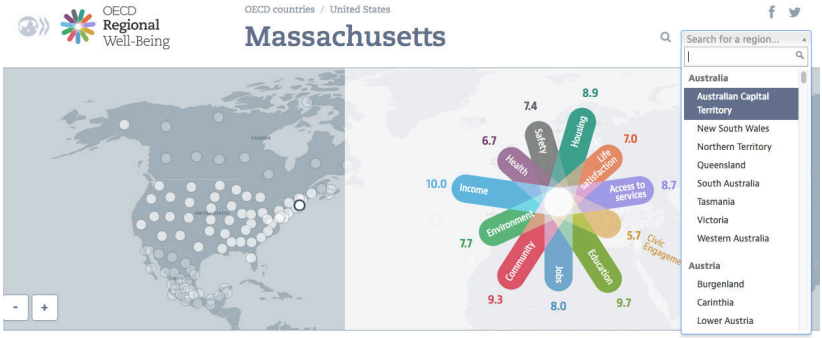


Figure 9: Example for the **familiarize** pattern: OECD Better Life Index [28].

audiences and the story is **humans-behind-the-dots**. This pattern is similar to concretize, and is often used in journalism, where a concrete person or setup is introduced as an example for the matter of the story. Humans-behind-the-dots consists in presenting individual stories through detail-on-demand in a visualization. For example, the CNN story on US casualties in the wars in Afghanistan and Iraq [35] (Figure 8) shows casualties as dots on two maps. Clicking on the dots reveals the identities of each person, including their name, photo, hometown, as well as a description of what exactly happened to them. While pictograms or shapes can be used to concretize numbers, these details can highlight the uniqueness of each individual character and their life's story.

Related to humans-behind-the-dots, is a pattern we call **familiarize**. This patterns helps create a familiar setting for audiences, so that they can relate pieces of the data to what they already know (e.g. the situations in which they live). In publishing and explaining data on human development (education, health, income, etc.) for different geographic regions, the *OECD's Regional Well-Being* [28] website starts by asking viewers for their current location; it then presents the development data for that location, and shows regions with similar development profiles.



Figure 10: Examples for patterns **rhetorical question** (left) and **call-to-action** (right). Both pictures from [18].

2.5 Patterns for Engagement

Engagement can be seen as the feeling of being part of the story, of being connected to it, and being in control over the interactions with the story's content.

One way to provoke (passive) engagement can be through emotion. Another way is to start with a **rhetorical question**. By directly addressing the audience (see **Breaking-the-4th-wall**), a rhetorical question can connect the narrator with the audience, and can trigger its will to answer that question, even if only for a moment. In this sense, a rhetorical question can set the context for a subsequent exploration of the reasons behind the answer, which the narration can then reveal and explain. Examples for rhetorical questions in data-driven stories are widespread, e.g. *What If I told you: you eat 3496 liters of water?* in Angela Morelli's Virtual Water [18], or *Can you live on the minimum wage?* [36].

Sometimes, a rhetorical question can also imply a **call-to-action**. After having learned much about how much water the production of certain food requires, Morelli [18] finishes with the words "Can we help?" and a list of action items for reducing personal water consumption, e.g. by regarding which food to eat, and how often. While many data-driven stories may intend for a call-to-action on different scales (e.g. *"What's Really Warming the World"* [23], or to reflect on war casualties [35]) these are not always explicit, nor are they necessarily coupled with concrete behavioral change takeaway (as is the case in Morelli's example).

Besides passive engagement in the form or emotion, empathy, or questioning the audience's own behavior, *active* engagement happens through audiences' action; it can shape the way the story is told, and

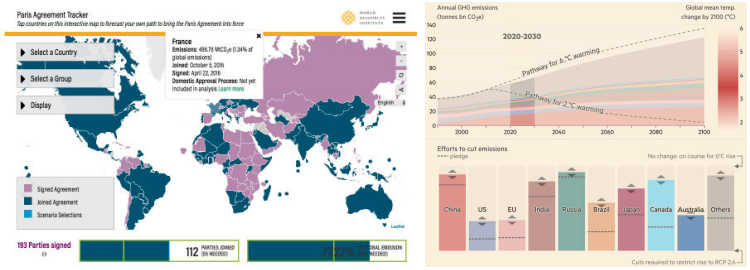


Figure 11: Examples for exploration and simulation: The Paris Agreement [39], and the Climate Calculator [40] (right).

which data is revealed (first).

Active engagement usually requires interactive features in a story presentation, if not a live presentation. For example, other than rhetorical questions, actual questions can ask for an audience's opinion [34], which may potentially alter the narration of the story. The inclusion of personal data, such as location [28, 37], or body weight [38] can also help interactively tailor the narration for the audience.

Alternatively to asking for actual data, the audience can be asked to **make-a-guess** about the provided information (e.g. about values or distributions). One example for this pattern is the NY Times' story on education [29]; the viewer is asked to draw a line inside a graph, representing the distribution between household income and percentage of children with a college degree. Make-a-guess questions can stimulate the audiences' curiosity to know the answer, and can engage them in causal reasoning about the phenomenon, while providing a motivation for following the narrative. On the one hand, audiences are asked to question their own perception of reality, and on the other hand, they are asked to question why this perception may have been wrong—in the case of a mismatch between their perception of reality and the actual data. Eventually, the narrative can also reveal what other audiences have guessed (a device used in the example of the college degrees, as well).

Make-a-guess questions can be posed in both remote and co-located presentation setups. In a co-located setup, the narration can be interrupted to address the audience directly. In a remote setup, interactive featurs can enable the audience to input their perception and understanding into the content of the narrative.



Figure 12: Case 1: US Debt Ceiling Visualized as $100 bills [24].

A last form of active engagement is **exploration**. Exploration can stand at the end of a linear narration (without user input) [7, 41], such as in [39] (Figure 11-left). Probably the highest degree of active engagement can be achieved through narrative that include simulations, games, and puzzles; rather than simply being asked to make a guess, the audience is confronted with a complex interface, where parameters of the data can be set freely (e.g. in the *Climate Change Calculator* [40]—Figure 11-right—or in the *Budget Hero* piece [42]).

3 Use Cases

We now cover three examples of data-driven narratives to illustrate how different narrative patterns can be used in combination.

3.1 Case 1: US Debt Visualized

Our first example talks about the scale of large numbers and how to communicate huge amounts of money. Demonocracy.info created a story to visualize the total debt of the United States [24]. This piece combines several narrative design patterns, woven into a static website that requires scrolling to follow the story. Each step consists of an image and explanatory text (Figure 12). The same story with the same narration also exists as an online video [24].

The example starts with an image of a $100 bill, **concretizing** a certain amount of money into a tangible object of a certain size that everyone is familiar with. The next stage in the story shows a stack of one hundred $100 dollar notes, bundled into a package that represents $10,000. The explanatory text states that this is the amount needed to buy a used car, or the equivalent of what the average human on earth

earns in a year. The amount of $10,000, or a car, or one year worth of work, is concretized into a small package of paper easily fitting the pocket of a jacket. The next pictures scale the amount to $100 million, using the analogy of money stocked on a common pallet that usually is used to carry and lift goods. After another few pictures and increases in scale, the money shown grows to $1 trillion, which largely exceeds the surfaces of an American football field and of the white house in pallets of $1 million dollar. Eventually, the number of $122 trillion dollars amounts into multiple skyscrapers, exceeding the height of the most important buildings in New York.

Thus, while **concretize** is predominent in this piece, it is not the only patterns used in this story. *Familiarization* is used through the different scale comparisons, juxtaposing and combining the number of bills and their volume with common spatial structures or material objects (e.g., pallets, a football field, and skyscrapers) and other immaterial values (e.g. that of a car, or of one year worth of work). This helps relate the numbers to things we can easily imagine, and gives a reference to other amounts of money.

The narrative also uses **repetition** and **gradual reeal**, as it is divided into individual stages that share the same strucutre, and present an increasing amount of money. This introduces a sense of flow to the narration.

3.2 Case 2: Can You Live on the Minimum Wage?

This example is an interactive piece from the New York Times that reflects upon living on the minimum wage [36]. The piece briefly begins with some context, mentioning the number of workers currently living on the lowest legal pay in the United States. It then moves on to invite readers to see whether, given their current lifestyle, they could make ends meet with such an income. The hypothetical yearly income is **concretized**, using a unit visualization in which each unit represents $1. Units are displayed as small green squares piled from bottom to top. To the left of this visual representation are a number of input fields, in which readers are asked to indicate how much they spend on basic necessities like rent, transportation, food, etc. Every time a new amount is entered, the units in the visualization are removed to show how much money is left. If no money is left, a negative amount of units appears—shown with red squares piled from top to bottom—indicating that the reader would end up in debt by the end of the year.

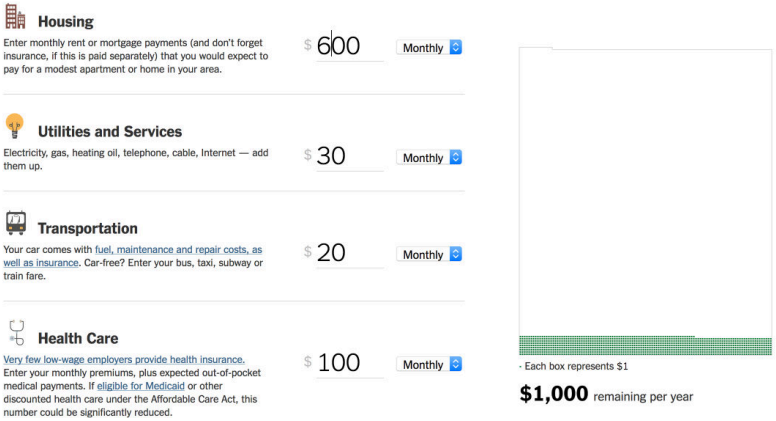


Figure 13: Case 2: Interface to enter and alter expenses for visualizing what can be afforded with a minimum wage [36].

This example also combines several narrative patterns. The first two are in the title of the piece itself: it **directly addresses the audience** by using a **rhetorical question**. The title asks whether "you" can live on the minimum wage, encouraging the reader to enter his personal expenses. The piece then serves to determine to what extent readers could not live with a minimum wage income. In addition, as the piece adapts to readers' input, it helps them **familiarize** with the amount of money that the minimum wage represents, by enabling them to compare it to the amount they spend on everyday necessities. This helps readers immediately identify with the story and the issue it addresses, i.e., the difficulties that 4.8 million Americans who live on the minimum wage have. Finally, readers can modify their entered values, in order to *explore* through *repetition*.

3.3 Case 3: What's Really Warming the World?

This last example is another interactive web piece, created by the Bloomberg Graphics team [23] (Figure 4). It sequentially interrogates possible explanations for global warning, in an attempt to debunk skeptics' arguments against manmade climate change. The piece starts with a **rhetorical question**: "What's Really Warming the World?," and

confronts it with a timeseries linegraph that shows the observed general increase in land-ocean temperature between 1880 and 2014. Scrolling down the page reveals different trends for possible explanations, and relates them to the general increase in temperature. The question also updates to e.g. "Is It the Earth's Orbit?," "Is it the Sun?," or "Is it Volcanoes?," implying the reader should **make a guess**, based on the correlation s/he might see in the graphic. Of course, there are none, or a best, they are very weak. After going through several possible *natural* explanations, the questions and data switch to *manmade* explanations. This use of **repetition** is very powerfull, as it pushes readers to try and **find answers for themselves.** While deforestation, ozone pollution, and aerosol pollution still show no real correlation, greenhouse gases show a very strong correlation. The question then changes to a direct statement, **breaking the fourth wall**, telling the reader (and skeptics): "No, it Really is Greenhouse Gases." The piece finishes by aggregating data for all possible manmade explanations, and the correlation becomes even stronger.

4 Discussion

In this chapter, we have described a set 18 prominent narrative patterns in Section 2, which we have categorized in 5 groups. These groups can help inform us as to when to use a certain pattern, i.e. what might be a potential purpose for the pattern. Section 3 then provided a closer look at how different narrative patterns can be woven into a single narrative. While we believe this is an interesting start, we are also convinced that several open questions remain, e.g. How can we employ our collection to create our own narratives? or, How are patterns affected by the presentation medium or the audience?

4.1 Using Narrative Patterns

We think narrative patterns can be used in both a generative and analytical fashion by a wide variety of sotrytellers: journalists, teachers, students, designers, visualization researchers, scientists, etc.

Generative use: Our main intention is to provide a pool of possibilities and ideas that help authors create data stories. Like other examples of design cards [9, 10], narrative patterns can provide a starting point and initial ideas to assist and augment the creative process.

In a typical use case, cards can be randomly selected, and their respective narrative patterns can be used to shape a tentative narrative for a given story. Other cards can then be selected to produce alternative narratives, creating diversity and potentially triggering discussion.

Analytical use: Narrative patterns can also help to analyze narratives of existing data-driven stories, by encouraging researchers and critics to think about the specific intents behind the use of different pattern. The patterns could typically be used in education, to teach data-driven storytelling, by analyzing which patterns go well together in a specific story? Which patterns exclude each other? or, Which are the favorites?

Eventually, the process of creating new narratives may lead to novel patterns or novel variants of existing patterns. We envision our corpus of narrative design patterns will grow organically, enriching both the creation of new narrative experiences as well as our understanding of the use of narrative techniques in data-driven storytelling.

4.2 Storytelling Techniques

Narrative patterns are not storytelling "techniques" themselves. This means, our patterns are not technical blue prints or code that can be downloaded and applied in a straight forward manner such as User Interface Patterns [43], Visualization Design Patterns [44], or Software Design Patterns [45]. Narrative patterns do not imply any visual (or other sensory) form, or even interaction techniques; the eventual visual, auditory, and interaction design remains to be created for each story individually, and benefits from (and requires) the creativity and expertise of the author.

4.3 Presentation Medium

Since narrative patterns are independent from the eventual presentation, they are also generally independent from the presentation medium; web, oral presentation, print, video, and so forth. Some patterns may not apply to specific media. For example, patterns for engagement might require a back channel between the audience and the storyteller. Emotion can potentially be transported better through certain media (e.g. movie, sound as in *Out of CTRL* [34]). We have consistently talked about the storyteller and the audience. Of course, depending on the presentation medium, the storyteller can be colocated with the

audience, or having authored the story content, the audience may consume it remotely and asynchronously. We think there are vast possibilities for future research here.

4.4 Different Notions of Time

Talking about the presentation medium and the potential colocation of storyteller and audience, we necessarily have to talk about the notion of time in (data-driven) storytelling. We refer to three notions of time: *authoring-time*, *presentation-time*, and *data-time*.

Authoring time: Authoring time refers to the time when the story and its narration is created by the author. In some cases, story and narration are detached in time, i.e. the story is usually written before the narration and can be written by another person. In any case, We consider both as authoring-time. During authoring-time, patterns are "selected" and/or "implemented" into the final narration.

Presentation time: Presentation-time specifies the time at which a presenter presents the story to an audience, or the audience consumes the story, depending on the presentation medium. Generally, patterns are not changed during presentation time. However, there may be cases where the audience may chose its own narration through the story, through interaction and exploration. Also, patterns such as repetition can be implied in an exploratory narration; for example *Can You Live on the Minimum Wage?* [36], as users play with different budgets, expenses, and geographical regions.

Data time: Finally, the third notion of time is the time indicated in the data itself. For the sake of communication, presentation-time can be related to data-time. One such example is the counting of soldiers in *The fallen of WWII* [25]; time of people being killed in the war is implicitly mapped to the time as the video continues while more and more data points (soldiers) are added to the screen. A similar example relating presentation-time and data-time is presented by Hans Rosling talking about the development of nations [26]; Rosling slows-down, speeds-up and stops the pace of the data as it is animated through the changing years. Interestingly, his voice reflects the same pace as the data changes, slow when there are small changes and speeding up as the data begins to change more rapidly.

Besides animation, visual techniques for presenting time are well studied in the literature; Aigner et al. [46] surveys general visualization for temporal data, Brehmer et al. [47] overview timeline visualizations,

Beck et al. [48] compile a an overview over visualizations for time-chancing networks, and Bach et al. [49] describe spatio-temporal data such as geographical data, videos, networks, multidimensional data, and others. Visualizations and systems specific for presenting temporal data are only in the beginning of gaining greater attention [47, 50, 51].

4.5 Audience and General Intention

Similar to the presentation medium, the target audience can vary; children, students, experts, car owners, consumers. Certain patterns may work for a specific audience (e.g. *rhetorical questions, make-a-guess, exploration*) while not for other types of audiences. Audiences differ in homogeneity, in background knowledge, in their willingness to interact, in cultural behaviors, and in individual views on the world. Depending on the audience, a storyteller may have different motivations for communication (informing, educating, confirming, questioning, moving to act, etc.) and hence change the narrative patterns that they choose. Sometimes an author may have the luxury of preparing multiple versions for different audiences (or media) while at other times, they may need to fall back on a lowest common denominator. See the Chapter 9: "Audience" for more details.

5 Conclusions

This chapter introduced narrative patterns for data-driven storytelling. We described and illustrated a core collection of 18 patterns with examples, discussed relations between patterns, and gave scenarios on their usage.

We are only at the very beginning to understand how data-driven storytelling works and how narrative patterns can benefit the communication of a message. Beyond creating data-driven stories, we hope that narrative patterns will help authors think more broadly about data-driven storytelling. For example, the discussion of what can be called data-driven storytelling is far from settled (hence the motivation for this book), and the transition from explanations based on illustrations and graphics towards data-driven argumentations is gradually sinking in. Also, we do not think (all) narrative patterns are limited to data-driven storytelling. Rather, we got inspiration from common storytelling mediums and looked for common patterns across the domains, i.e. where the rather novel medium of data-driven stories is related to the classical mediums. We also do not see narrative patterns bound to a specific technical medium such as videos, static or interactive graphics, interactive presentations, games, news articles etc. We hope that the narrative patterns presented here can contribute to the discussion on what data-driven storytelling is, what its mediums are, and what are good or bad practices.

Many questions remain to be answered: What are the specific rules a specific narrative pattern should adopt or implement? Which patterns go well with each other and which exclude each other? Though we do not aim to restrict and over-quantify narrative patterns, such questions can serve to establish a common ground to create effective stories.

Other questions remain on which patterns are suitable for which audience and how particular patterns apply to specific media; how does the presentation medium change the implementation of the narrative pattern? This is closely related to the specific techniques that are used in the implementation of a narrative pattern, and how those are related to the presentation medium, e.g. animation, interactive highlighting, multiple choice questions, interactive simulations, etc. So far, the links between narrative pattern and presentation technique remain unclear. We are just beginning to think about the respective implications and purposes of each pattern. How persuasive is a narrative pattern? Can we measure effectiveness? Is a pattern more emotional or more ra-

tional? Does it aim for objectiveness or for provoking an emotion? Potentially, there are alternative groupings than the ones proposed in Section 2. Eventually, there are many open questions related to the higher-level structure of data-stories, including rhythm, and pace. Are there specific patterns for story-structures, similar to the structures of classic greek drama, comic books and graphic novels, documenteries, or scientific articles?

While there are an infinite number of possibilities to combine, adapt, extend, and implement patterns, we think our initial collection is a good starting point to discuss all these questions and to start better understanding data-driven storytelling. We hope for a growing collection of patterns, as well as for a creative discussion in the near future.

References

[1] G. Genette, *Narrative discourse: An essay in method*. Cornell University Press, 1983.

[2] J. S. Bruner, *Actual minds, possible worlds*. Harvard University Press, 1985.

[3] G. Branston and R. Stafford, *The media student's book*. Psychology Press, 2003.

[4] S. McCloud, "Understanding comics: The invisible art," *Northampton, Mass*, 1993.

[5] B. Alexander, *The New Digital Storytelling: Creating Narratives with New Media: Creating Narratives with New Media*. ABC-CLIO, 2011.

[6] N. Gershon and W. Page, "What storytelling can do for information visualization," *Communications of the ACM*, vol. 44, no. 8, pp. 31–37, 2001.

[7] E. Segel and J. Heer, "Narrative visualization: Telling stories with data," *IEEE transactions on visualization and computer graphics*, vol. 16, no. 6, pp. 1139–1148, 2010.

[8] K.-L. Ma, I. Liao, J. Frazier, H. Hauser, and H.-N. Kostis, "Scientific storytelling using visualization," *IEEE Computer Graphics and Applications*, vol. 32, no. 1, pp. 12–19, 2012.

[9] "Ideo Cards." online: https://www.ideo.com/work/method-cards. [last visited: Nov. 29, 2016].

[10] "Oblique Strategies." online: https://en.wikipedia.org/wiki/Oblique_Strategies. [last visited: Nov. 29, 2016].

[11] "Design with Intend." online: http://designwithintent.co.uk/docs/designwithintent_cards_1.0_draft_rev_sm.pdf. [last visited: Nov. 29, 2016].

[12] "Creativity Cards." online: http://ja-ye.atom2.cz/form/download.ashx?FileId=52. [last visited: Nov. 29, 2016].

[13] "Service Design Tools." online: http://www.servicedesigntools.org/frontpage?page=3. [last visited: Nov. 29, 2016].

[14] S. He and E. Adar, "Vizitcards: A card-based toolkit for infovis design education," *IEEE Transactions on Visualization & Computer Graphics*, no. 1, pp. 561–570, 2017.

[15] "Education vs. Incarceration." online: http://visual.ly/education-vs-incarceration. [last visited: Nov. 29, 2016].

[16] O. Neurath, *International Picture Language; the First Rules of Isotype: With Isotype Pictures*. K. Paul, Trench, Trubner & Company, 1936.

[17] "Gerd Arntz Web Archive." online: http://gerdarntz.org/isotype. [last visited: Nov. 29, 2016].

[18] "Virtual Water." online: http://www.angelamorelli.com/water. [last visited: Nov. 29, 2016].

[19] F. Chevalier, R. Vuillemot, and G. Gali, "Using concrete scales: a practical framework for effective visual depiction of complex measures," *IEEE transactions on visualization and computer graphics*, vol. 19, no. 12, pp. 2426–2435, 2013.

[20] "Sand Dance." online: https://www.sanddance.ms. [last visited: Nov. 29, 2016].

[21] "The Globe." online: http://globe.cid.harvard.edu/. [last visited: Nov. 29, 2016].

[22] S. Haroz, R. Kosara, and S. L. Franconeri, "Isotype visualization: Working memory, performance, and engagement with pictographs," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, (New York, NY, USA), pp. 1191–1200, ACM, 2015.

[23] "What's Really Warming the World?." online: http://www.bloomberg.com/graphics/2015-whats-warming-the-world/. [last visited: Nov. 29, 2016].

[24] "US Debt Ceiling Visualized in $100 Bills." online:http://demonocracy.info/infographics/usa/us_debt/us_debt.html, note=[last visited: Nov. 29, 2016].

[25] "The Fallen of World War II." online:https://vimeo.com/128373915. [last visited: Nov. 29, 2016].

[26] "Hans Rosling's 200 Countries, 200 Years, 4 Minutes - The Joy of Stats - BBC Four." online:https://www.youtube.com/watch?v=jbkSRLYSojo. [last visited: Nov. 29, 2016].

[27] "U.s. gun deaths in 2013." online:http://guns.periscopic.com, note=[last visited: Nov. 29, 2016].

[28] "OECD Better Life Index." online: httpssolo://www.oecdbetterlifeindex.org. [last visited: Nov. 29, 2016].

[29] "You Draw It: How Family Income Predicts Children's College Chances." online: http://www.nytimes.com/interactive/2015/05/28/upshot/you-draw-it-how-family-income-affects-childrens-college-chances.html. [last visited: Nov. 29, 2016].

[30] "The Globe." online: https://www.theguardian.com/world/ng-interactive/2015/aug/26/china-economic-slowdown-world-imports. [last visited: Nov. 29, 2016].

[31] "A conversation with guy gavriel kay." online:http://www.fantasy-magazine.com/non-fiction/articles/a-conversation-with-guy-gavriel-kay. [last visited: Nov. 29, 2016].

[32] A. L. de Albuquerque and A. Cheng, "31 days in iraq." online:http://www.formfollowsbehavior.com/2007/09/17/artistic-data-based-visualization/31days_iraq_20061/. [last visited: Nov. 29, 2016].

[33] D. Wiesner, *The three pigs*. Random House, 2014.

[34] "Out of CTRL: die Web-Doku." online: http://future.arte.tv/de/netwars-krieg-im-netz-1/out-ctrl-die-web-doku. [last visited: Nov. 29, 2016].

[35] "Casualties." online: http://edition.cnn.com/SPECIALS/war.casualties/index.html. [last visited: Nov. 29, 2016].

[36] "Can You Live on the Minimum Wage?." online:http://www.nytimes.com/interactive/2014/02/09/opinion/minimum-wage.html. [last visited: Nov. 29, 2016].

[37] "The Division game." online: http://collapse-thedivisiongame.ubi.com. [last visited: Nov. 29, 2016].

[38] "Your Olympic athlete body match." online: http://www.bbc.com/news/uk-19050139. [last visited: Nov. 29, 2016].

[39] "When Could the Paris Agreement Take Effect? Interactive Map Sheds Light." online: http://www.wri.org/blog/2016/04/when-could-paris-agreement-take-effect-interactive-map-sheds-ligh [last visited: Nov. 29, 2016].

[40] "Climate Change Calculator." online: http://ig.ft.com/sites/climate-change-calculator. [last visited: Nov. 29, 2016].

[41] J. Boy, F. Detienne, and J.-D. Fekete, "Storytelling in Information Visualizations: Does It Engage Users to Explore Data?," in *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, (New York, NY, USA), pp. 1449–1458, ACM, 2015.

[42] "Budget Hero." online: https://www.wilsoncenter.org/budget-hero. [last visited: Nov. 29, 2016].

[43] "User Interface Patterns." online: http://ui-patterns.com/patterns. [last visited: Nov. 29, 2016].

[44] "Visualization Design Patterns." online: http://www.infovis-wiki.net/index.php?title=Visualization_Design_Patterns. [last visited: Nov. 29, 2016].

[45] E. Gamma, *Design patterns: elements of reusable object-oriented software*. Pearson Education India, 1995.

[46] W. Aigner, S. Miksch, W. Müller, H. Schumann, and C. Tominski, "Visualizing time-oriented data—a systematic view," *Computers & Graphics*, vol. 31, no. 3, pp. 401–409, 2007.

[47] M. Brehmer, B. Lee, B. Bach, N. H. Riche, and T. Munzner, "Timelines revisited: A design space and considerations for expressive storytelling," *IEEE Transactions on Visualization and Computer Graphics*, 2016.

[48] F. Beck, M. Burch, S. Diehl, and D. Weiskopf, "The state of the art in visualizing dynamic graphs," *EuroVis STAR*, 2014.

[49] B. Bach, P. Dragicevic, D. Archambault, C. Hurter, and S. Carpendale, "A descriptive framework for temporal data visualizations based on generalized space-time cubes," in *Computer Graphics Forum*, Wiley Online Library, 2016.

[50] B. Bach, N. Kerracher, K. W. Hall, S. Carpendale, J. Kennedy, and N. H. Riche, "Telling stories about dynamic networks with graph comics," in *Proceedings of the Conference on Human Factors in Information Systems (CHI)*, ACM, New York, United States, 2016.

[51] B. Bach, C. Shi, N. Heulot, T. Madhyastha, T. Grabowski, and P. Dragicevic, "Time curves: Folding time to visualize patterns of temporal evolution in data," *IEEE transactions on visualization and computer graphics*, vol. 22, no. 1, pp. 559–568, 2016.