The Line of Time: Backgrounds and Usage of Event-Timelines

Johanna Fulda*
Department of Computer Science, University of British Columbia

ABSTRACT

One of the most prevalent parameters in visualized data is time. A large part of information visualization concerns time-series data, where a progress or shift is shown over time. In this survey paper we want to look at the special case of event timelines that aim for giving an overview over a sequence of events that all belong to a superior topic. Timelines can for example show happenings in the news, the biography of a person or the history of a company, the development of a certain technology or a review of events that led to or influenced a historic event. In most cases events are aligned along a horizontal line, which represents the progression in time from left to right. It seems like a commonly accepted representation, but how much insight do we really get through timelines and how do people read and understand them? Moreover there are some issues arising during the design process of a timeline, e.g. when the dates are distributed irregularly, when there are long periods of time where we don't know the density of events or when a date is not defined precisely. We will look at those challenges and how several different tools and techniques try to solve them, especially in the domains of science and journalism - even though time constraints and a general technophobia showed up to be reasons for not using them too often.

1 INTRODUCTION

> With one dimension marching along to the regular rhythm of seconds, minutes, hours, days, weeks, months, years, centuries, or millennia, the natural ordering of the time scales gives this design a strength and efficiency of interpretation found in no other graphic arrangement. (Edward Tufte)

In "The Visual Display of Quantitative Information"[38] Tufte, one of the big pioneers in modern information visualization wrote that, emphasizing the strength of the human understanding of temporal order and thus the tangibility of aligning data chronologically along a path. For a better understanding of the time factor Section 2 describes how time is transferred linguistically in our daily businesses - conscious or unconscious. It can happen through different channels, we look at the written or spoken word on the one hand, computers on the other and how Natural Language Processing offers translation between humans and machines. Section 3 then goes deeper into the cognition process. Experiments showed that humans are fairly well in using spatial metaphors when thinking and talking about time. But when bringing those images to paper they still might end up differently. Reasons can be general differences in visual perception, handedness and the different directions of writing. Still the most widespread timeline-technique in English speaking environments is a horizontal line, read from left to right. That representation seems to be commonly accepted and understood for quite some time now, but it was not always like that, as we will see in Section 4, where the historic development as well as today's usage are elaborated in more detail. Examples from science and newspapers show the usage and usefulness of timelines but also highlight some challenges the creators have to face. Those spring from the temporal data itself, as well as from time restrictions, technical problems and reservations towards the techniques behind it or just the lack of knowledge. But as freely available tools are spreading out all over the internet, we looked at some of them in Section 5 and picked out some examples for every level of technical expertise.

*e-mail: jfulda@cs.ubc.ca

2 TRANSMISSION OF TIME

Time is the 4th dimension[1], time is inseparable from space[2], "Time is what keeps everything from happening at once"[3], Time is "one of the greatest mysteries in all of nature"[4]. But we don't want to delve deep into the underlying structure of physical reality, psychology or behavioral science, and just inspire some thoughts on what time actually is (or rather seems to be). For the purpose of this paper we first look at the hard facts of how language deals with temporal information, how time can be used as data, and how we should format it to make it understandable for a machine.

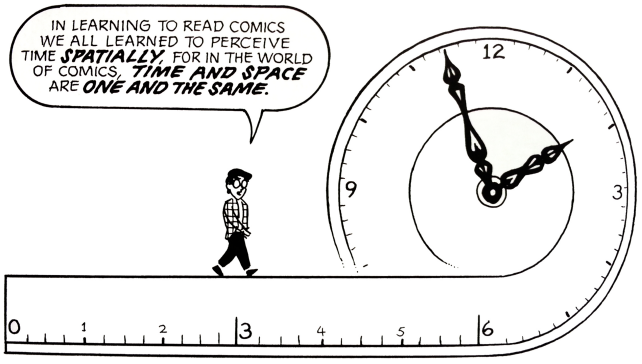

Figure 1: From Scott McCloud's "Understanding Comics". Chapter Four "Time Frames."

2.1 Language

Time is ubiquitous and thus also a very present topic when we speak or write. Either to indicate when something happened, or to place it related to something else. First we will quickly name the components for formulating temporal expressions (most languages use that kind of representation, the primarily represented language here will be English) and then how that knowledge can be used to translate it into a more functional and machine-readable format.

[1] In addition to the three dimensions of space; introduced by Sir Isaac Newton in his "Mathematical Principles of Natural Philosophy"

[2] That is at least in Comics - see Figure 1, but also in the common language use, see Section 4

[3] From Ray Cummings' novel "The Girl in the Golden Atom", 1922

[4] Said Michio Kaku, an American theoretical physicist, futurist, creator of a four-part BBC documentary about time.

2.1.1 Grammar

In our language, written or spoken, *tense* and *aspect* give us information about time. *Tense* (the grammatical form of a verb, e.g. past, present, future) lets us "localize" events or states in time, whereas *aspect* tells us more about the flow of the event. Aspect can either be *perfective* or *imperfective*. *Perfective* means a bounded and self-contained event, where the verb indicates an action that took place at one point in time, beginning and ending included ("I ate ice cream"). Whereas *imperfective* indicates a continuous or repeating event ("I was eating ice cream" or "I used to eat ice cream") where it is implied that the event had a certain duration. In "How time is encoded" Klein describes four more types of devices that encode time in language. Those are *lexical aspect* (where the meaning of the word implies more about its temporal feature, e.g. sleep vs. crack), *temporal adverbials* (temporal expressions like "soon", "now", "rarely"), *temporal particles* (only present in few languages, e.g. Chinese, where a suffix of a verb can influence its temporal meaning), and *discourse principles* (positioning inside whole story, assuming that the default order is chronological) [14].

Thus we can set three crucial subsystems of linguistic time [22]:

- **Tense**
  "Is it already over, enduring or pending?"
- **Aspect** (perfective, imperfective, lexical)
  "How long did/will it endure?"
- **Temporal sequencing** (including *temporal adverbials*, *discourse principles*) lets us order events either to our viewpoint or to one another

To sequence an event we need a reference date, which is the time of speaking or the time a written document was created. With that we are able to get information about the order of events, compared to the reference date or in relation to another event (was it before, at the same time or after something else).

2.1.2 Translation to machine readability

When we speak or write about events, we rarely use concrete dates, that we could register in a calendar, like *"on November 25th 2014 xy happened"*. Of course the style of writing or speaking varies according to different purposes, but if we don't write formal reports, it's rather unlikely that we annotate every event with a concrete date. Most often we rather give vague directions, like some point in the future or past or a reference time, like *"at the time of xy"*. Also the duration rarely gets indexed, sometimes we can just tell if it was rather short or long. And for understanding the context that might be just enough information and we can already order events in time and get the chronology straight in our heads. But if we want to automate the process of extracting temporal information inside sentences, for example because we have a vast amount of documents and don't want to read them all manually, this issues some challenges - and in the field of Natural Language Processing scientists devoted a lot of attention to tackle those. There are two different approaches: The **Rule-based approach** and the **Machine-Learning (ML) approach**. A machine can either just stubbornly follow some rules, use search strings and regular expressions to identify temporal expressions, or Machine-Learning algorithms can be applied to identify them based on previous "experience" gained from an annotated example document collection. (There are also unsupervised learning approaches - where the example data doesn't have to be annotated as much, but we won't go deeper into that).

Machine-Learning approaches are much more dominant these days, still sometimes the heuristic rule-based approach is more appropriate. It is very powerful when it comes to pattern matching and the big advantage compared to ML is, that there is no test data set needed. All the rules can be written by just knowing very well about the domain, which spares the often quite expensive training cycle and lets the creator understand more clearly why the system actually decided doing what it did. Disadvantages are that writing the rules can become a very tedious process, also the targeted expressions have to be assessable and non-varying, to match with the provided patterns and dictionaries. Performance-wise it depends on the number of rules. Rule engineering doesn't scale too well, the more rules you have, the more possibilities of them getting into each others way appear and it can easily end up in a chaotic system. When the raw data gets bigger, more chaotic and noisy Machine-Learning approaches become more suitable. Those algorithms can for example create decision trees to enables a decision in case of ambiguity. Based on experience they are able to go for the more likely case. Because ML models are robust also to unfamiliar input (like typos or unknown words), they can handle huge unorganized datasets much better than the rule-based approach. And to improve the ML systems the creators can just provide them with more annotated test data - which is also time consuming, but much less complex than expanding rule-based systems. Still in NLP tasks like "Named Entity Recognition" (to recognize names e.g. of persons or organizations) or "Morphological Segmentation" (to identify the class of a word) and our needed "Temporal Expression Recognition" (which can be seen as special case of "Named Entity Resolution") hand written rules often obtain better precision - but as mentioned require hardworking and persistent computational linguists to create them and are most likely not perfect.
But for now let's assume the NLP for recognizing temporal expressions works perfectly, and when analyzing a written document the machine manages to annotate all words that somehow refer to time - the machine still can't do too much with that information, so we need another process, where temporal expressions become "normalized". That means anchored to a concrete date, and we can do that with the help of a reference date, that being the time of speaking or the time when the document containing the information was created. With those two steps a machine can nowadays fairly well interpret the time of concern of any written document (also spoken words can be interpreted like that, but we first need to translate the audio into text - which also works very well these days but is another topic). [18]

2.2 Data

Because a computer handles information differently than humans, a translation from one format to the other is needed, as seen in the previous paragraph. Annotating events brings us halfway there, but the data still has to be transformed into a structured format, a table or one of many structured file types - like JSON, XML, CSV etc..

2.2.1 Time-varying data

There are different use cases and data sets which have time as their key value, to make a clear distinction, we define two kinds of data: One where the time component is rather predictable, most often in regular distances and the task exploratory: **Time-series data**; and one unpredictable, irregular data set created for explanatory tasks: **timeline data** (which is the underlying dataset for our timelines). Figure 2 illustrates the different data sets.

Explore: Time-series data With today's computational power and the omnipresence of sensors and smart machines, it is very easy to collect huge amounts of data. Because very often data is collected over time, with varying values at every measuring point, we produce a lot of time-series data where time is the key-value of the data set. The intervals of measurement are most often regular, thus the rhythm of the time values is consistent and predictable. Questions like: "Does a data element exist at a certain point in time?", "Is there periodicity in the data elements?", "How

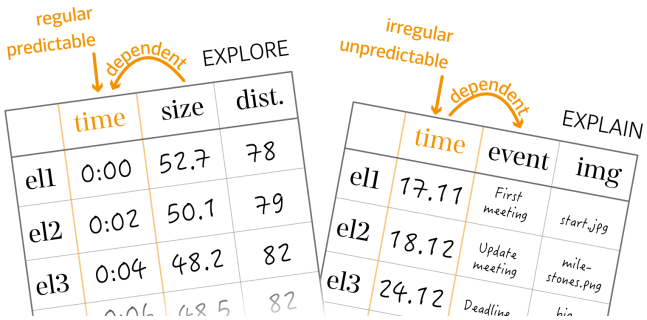

Figure 2: Paradigm data set for both defined kinds: time-series data (left) and time-line data (right)

fast does change happen?" or "When did something start and end?" can be answered using those data sets [23]. Time-series data is multi-dimensional and usually has several quantitative values, time being just one of them. When analyzing them, we either want to confirm something we are already expecting, or we want to discover things like patterns, peaks, lows, trends or irregularities. When we are on the course of discovery we face the challenge of finding the most meaningful parameters - sometimes the combination of specific parameters brings up interesting insights. It is undisputable that visualizing the data facilitates answering those questions immensely or even makes it possible in the first place. But still a lot of preprocessing (like filtering, aggregating, ordering) is usually needed, especially when it comes to data sets with thousands of elements and values. For science, engineering or business tasks this analysis is widespread and very valuable - for the context of event timelines that we are considering here, we need another kind of data set.

Explain: Timeline data One big difference between those two data sets is the predictability. While time-series data is rather predictable (due to regular measurements), timeline data in general is unpredictable. When creating time-series data, values are measured at certain times, so they are dependent on their measuring time. In timeline data the time is dependent on the event. The latter case is the underlying data set for the event timelines considered in this paper - we use the term *timeline data*. It is usually much smaller than a time-series data set. Instead of automatically created by sensor measurements it often is manually created by an author or journalist or at least refined manually. It consists of only one or two quantitative values, which is the time when an event began and often the duration (or the ending time - resulting in the same outcome). The starting time is the value that predicts the ordering, other parameters can be seen as its companions and are most likely text or other associated elements (like hyperlinks or media). In contrast to time-series data, the intervals between the data points are rather irregular. A data point should only be created, when the date is relevant to the discussed topic (which of course is a rather subjective decision) to not overwhelm the reader with too much information and rather give an overview than very detailed background information. The mantra should be "less is more" in contrast to "more is more" for time-series data (where more measurements over time usually results in more precise insight). When it comes to visualizing timeline data we basically only need one dimension for aligning the data points according to their starting date. If there are ordered categorical values they could be aligned along the second dimension in a 2D space, but that is already a modification from the very basic one-dimensional event timeline.

2.3 Generating events

To create a timeline data set, we first have to divide a story into single events and then choose the ones of which we think are most relevant to understanding the superior topic. That can either (and is most often) be done by a human, who probably has quite some domain knowledge, or (to be a little more futuristic) could be done by a computer (which nowadays mostly still requires some human post-processing, but we're getting there). The input data for a machine can either be a document collection, a long text or the like. Let's say we have a document collection where all the single documents cover the same superior topic, the goal is then to split it into single events, get their date (the quantitative value) and summarize them (e.g. title and one sentence explanation as accompanying parameters). The research field of Multi-Document Summarization (MDS) has methods to do so. One of them is to first divide all documents into single sentences and cluster those that contain similar words and are interpreted to belong to the same subtopic. To only get the most relevant events, those topics are weighted according to their importance or "salience", as it is called in that context. The computed salience often depends on the number of appearance of certain terms throughout the whole document collection (also synonyms can be considered with the help of lexical databases like *WordNet*). The computer then filters out only the x highest ranked events and presents one summary sentence for the topic, which can for example be generated through alignment trees [2]). If no temporal expression can be found inside the sentence, the date is set to be the document creation time (DCT) in the multi-document case. (That is just one of many different approaches.) It seems like a promising approach, nevertheless studies showed MDS attempts can't compete yet with the human assessment of importance and especially when summarizing an event we still preferably have a human administrator in the loop, who decides what information is needed to best explain it [5, 6].

2.4 Format of time

To specify the date value in timeline data sets it has to be in a format that is readable for the computer. But dates can be defined in many different ways. We can either just tell the century in which something happened, or define it much more precise, by naming the hour of the day or even the minutes, seconds or nanoseconds. All those definitions are legit phrasings for dates. The hierarchy in time is "deeply multiscale"[24] and we can't say that any granularity is better than another one in general. There is also no evenness in the hierarchy. One hour has 60 minutes, a day has 24 hours, one month has something between four an four and a half weeks, some years have 357 days instead of 356. Thus also calculating dates is not as easy as doing the math with regular numbers and arises some challenges and exacerbates transformations and aggregations. And an event often is not just one date but a time period - a duration with a start date and an end date. Or it can be periodic, repeating every hour, or year or maybe even irregularly. To bring some order into this confusing matter, the International Organization for Standardization published the *ISO 8601 standard* (Information interchange - Representation of dates and times). It is there to prevent misinterpretation between different conventions, because different countries write dates in different formats or even inside one environment people tend to use different ways of writing. So since 1988 there is at least an officially right way of writing dates and times. We start from the left, with the year (largest temporal term) and go into more detail to the right (month, day, hour, second), depending on how much information we have. In the end it could look like this: **2014-11-08** (Format: YYYY-MM-DD). We are only allowed to use numbers and some special characters, like W if the date should be expressed by the year's week number, or Z for indicating the time zone at the end of the expression.

3 Cognition of Time & Space

Now that we got the linguistics and formal matters straight, let's look at how well people actually understand temporal information and how spatial imagination is used for it. A lot of expressions in our language combine time and space, and also our imagination of time is closely connected to space. Or as Mitchell brings it together [22]:

> Time is space [16], in which we think of time as an imaginary place or path along which we walk. In this metaphor, the future lies before us, the past behind, and we stand at now [37] (Marilyn Mitchell)

3.1 Time & Space in our minds

In the paper "Metaphoric structuring: understanding time through spatial metaphors" Boroditsky looks at how metaphors of space influence our way of thinking about temporal information [4]. Because we know that time is continuous and nonrecurring, we often use one-dimensional terms to describe time. We say for example "I am looking *forward* to something" or "when I'm looking *back*". Those spatial schemas that we create in our minds provide us information about how to organize events in the continuous flow of time. Boroditsky looks at two metaphors in more detail. The "ego-moving metaphor" (Figure 3 a) and the "time-moving metaphor" (Figure 3 b). In the first case, the person speaking is inside the stream and everything that is *in front* (in the observer's direction of movement) is future. In the "time-moving" case events are moving relative to the observer who is standing outside - *front* relates to earlier events, in the case of Figure 3 to those who are on the left side [19]. Mixing up these two representations can lead to misunderstandings. For example if we say: "The Wednesday meeting has been moved forward two days" it depends on the condition of the recipient if he will show up on Monday or on Friday. (Using the ego-moving condition it will be Friday, because "forward" suggests it is later in time, the time-moving condition concludes Monday because in that case, "forward" makes the event move to an earlier time - the perspective of the recipient is most likely implied by prior context) [19].

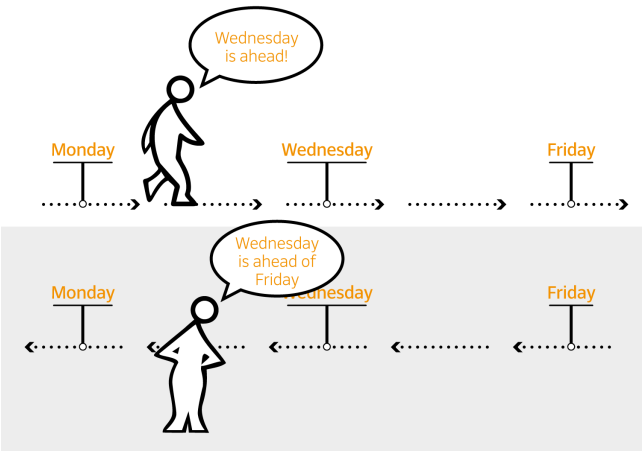

Figure 3: (a) Ego-moving schema. (b) Time-moving schema. Redrawn from L. Boroditsky: Fig. 2 in "Metaphoric structuring: understanding time through spatial metaphors"

Testing those different interpretations McGlone and Harding prompted that there is a correlation between language and spatial thinking [19], Boroditsky performed another experiment to confirm that there are conceptual similarities between space and time, by influencing her test users with spatial primes and then asking them questions about time [4]. Based on the results she concluded that spatial schemas can be useful, but are not required when thinking about time. (She also tested how it might work the other way around - if people can make use of temporal information when thinking about space - and the result is: they can't).

If we are able to use space to support understanding temporal order in our head, why not concretize it by putting that mental image on paper? We are most likely able to use the connection to make information accessible more quickly by drawing a chronological overview of a collection of events. That means a reader doesn't initially have to understand all details to build his own mental image, but can begin with using one that was already created by somebody who knows more about that topic. That lets the reader access the topic more quickly, thus lowers the barrier of entry. Because our language assumes time to be something one-dimensional, a line might be the right choice for representation (more about that in Section 4).

3.2 Time & Space on paper

Facilitating the understanding of something more complex can be done by visualizing it on paper (or today rather digitally on computer screens). Pictures have the power to easily communicate complex information, thus our visual system is "the major player in visualization"[32]. That's why the creators of visualizations should be aware of the principles of visual perception, including that people tend to perceive things differently. That can either be due to physical causes (like dyschromatopsia, color blindness or other visual impairments) or influenced by prior knowledge, the motivation of the user or the degree of visual literacy [39]. In the previous section we saw that people imagine time as something spatial, and tend to see the future "before", and the past "behind" them. So what would be a suitable representation of that imagination to paper?

3.2.1 The line

Joseph Priestley (the inventor of the modern timeline - more about that in the history-section 4.1) claims that even though time is something ungraspapal it is somehow quantitative and "admits of a natural and easy representation in our minds by the idea of a measurable space, and particularly that of a LINE; which, like time, may be extended in length [...] and thus a longer or shorter space of time may be most commodiously and advantageously represented by a longer or shorter line."[31]. The one-dimensionality of time suggests using a one-dimensional representation, but why we chose the horizontal line and the alignment from left to right, we will see in the following section. In the book "Catographies of Time" Rosenberg and Grafton demonstrate that "the line is a much more complex and colorful figure than is usually thought" [35], also lines can appear in many different forms, they can for example also be circular, like the trace of the hour hand on an analog clock - still our focus will be on the line as in its straight form.

3.2.2 Alignment

If we draw visualizations in the two dimensional space, we have two axes along which we can align our elements. The horizontal and the vertical axis. Both very often indicate increase in value, either to the top or to the right.

Vertical The reasons for why a top position is interpreted bigger, more or better than a lower one comes from our ubiquitous gravity and thus how we and everything around us is oriented [36]. Physically speaking: something at the top has more potential energy than something lying on the ground. The vertical position is also prevalent in our language: "being on the top", "feeling down".

Horizontal But why is it that right is accepted to be bigger and more than left - it seems that there are no natural reasons for that ordering, because on the horizontal plane most things are generally pretty symmetric. Including our body. Except that most people have one dominant hand. The handedness influences not only our ability to act but according to a study from 1988 by Ldavas also our judgment of the referential term on the horizontal dimension [15]. But even more influential are the different directions of writing. A study by Tversky et al. from 1991 tested how that influences spatial perception [40]. With around 1200 children and adults from different language cultures, with different writing directions (English, Hebrew, Arabic) they looked at how they would use space to represent relations that are non-spatial. They let their study participants place elements on a paper, that either represented temporal information, quantitative information or their preferences. The result for the temporal alignment was a strong tendency to left-to-right for English speakers, and the other way around for Arabic speakers (Hebrew speakers were less coherent - probably because they are more likely to learn European languages, also arithmetic operations are made from left to right in Hebrew). Very young participants tended to align the temporal information from top to bottom [40].

So concluding from those experiments we can say that a horizontal timeline, that shows the progress of time from left to right seems to be the most accepted representation, at least for a left-to-right writing audience.

3.2.3 Marking events

The single events that are aligned along that line can have different properties. They can vary in duration and importance or belong to different categories. So we need to think about how to indicate these differences. In today's timeline tools it seems established to use a dot for an instantaneous event, a line for a period of time and either coloring them accordingly to their categories or use vertical alignment to indicate their affinity. Weighting events according to their importance is not very common, but could be done by changing the radius of the circle and increasing the line width or for example through saturation of color. Those are the most common methods of marking events on a timeline, but there are no rules or studies (that we know of) that prove them to be the most effective and favorable ones.

4 VISUALIZATION OF TIME

As seen in the previous section the horizontal line is a very intuitive way of understanding chronology and nowadays it seems like a horizontal timeline is just the most obvious thing to draw, when wanting to represent events belonging to one bigger topic. But as Rosenberg pointed out it is "only quite recently that scholars first thought to represent chronological relationships among historical events by placing them on a measured timeline"[34]. He nominates the 1750s and 60s as the time of birth of the modern timeline. After discovering how that happened we will look at how timelines are used these days, by exemplary looking at their application in science and journalism. And finally there are some noteworthy challenges designers have to face to create a meaningful and accepted timeline.

4.1 Short historic jaunt

We could go back to ancient cave paintings and claim those drawings to be very early timelines, but let's leave those out in the cold and start with "modern" timelines from the 17th century. ("A Timeline of Timelines" by Sasha Archibald and Daniel Rosenberg shows an elaborate listing with representative examples, beginning with the very first discovered timeline-like records, referring back to the second century A.D. [1]).

> "The timeline seems among the most inescapable metaphors we have. And yet, in its modern form, with a single axis and a regular, measured distribution of dates, it is a relatively recent invention."
>
> (Rosenberg and Grafton)

In the book "Cartographies of Time: A History of the Timeline" Rosenberg and Grafton give an in-depth review of graphic representations of time in Europe and North America from 1450 until today. They entitle Joseph Priestley, an Anglo-American theologist, to be the inventor of the modern timeline as we know it today. 1765 he published "A Chart of Biography" (see Fig. 4), where the life spans of 2,000 famous men between 1750 A.D. and 1200 B.C. were shown.

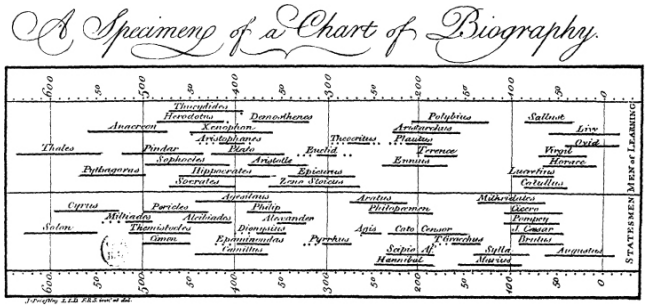

Figure 4: "A Specimen of a Chart of Biography" from Joseph Priestley, 1765 [35]

Because it was the first chart that included all the necessary visual vocabulary to map time, it was seen as the worthy successor of matrices - the standard representation of chronological data up to that point. He for example established the idea of showing the lifespan of a person with a horizontal line (starting at birth date, ending at its death) and to indicate dates that are uncertain he used dots. Still that kind of representation needed explanation to be understood by people and also had opponents. Laurence Sterne for example, an English writer didn't like the idea of showing history in a straight line and preferred telling stories with digressions and deviations. Presenting time as something linear seemed like a construction to him. Nevertheless since the 18th century the timeline became more familiar and understood. It even played a crucial role in the "*modern* understanding of historical time as linear, chronological, and flowing from future to past"[28], and it enabled people to compare different events in history. The broad dissemination let new questions arise, like how can you fit all important events on a horizontal line without resulting in a 54 feet long scroll (like Dubourg's Chronological Chart from 1753) or how can you combine the time data with additional information. In our modern digital times those questions can luckily be answered by using interactive tools, like zooming and details-on-demand. [33, 35, 28].

4.2 Benefits of Timelines

Visualizing data can massively improve the understanding process, we don't have to prove that anymore. The four most beneficial aspects of timelines we defined in the following:

- **Overview**
  Dividing a big story into several single events lets the reader better understand relationships and interferences between them. Every defined story started at some point and either ended at a certain time or is still ongoing. It is up to the writer of the story to define those borders (subjectively, so it is very likely that this decision is debatable). But it is helpful to have that defined time period to gain overview and get a first quick

and general understanding of the whole. Even if the reader doesn't understand everything in detail at that point, the time frame is set and thus the stage for further investigation. It also gives an idea of how much is going on in general - is it one thousand events or only a few, did they all happen at the same time or spread over the whole time frame and so on. When used in education, for example in history class, "they translate wonderfully from weighty analytic history books to thrilling narrative ones"[35]. One example that gives an overview of big game-changing events since the Big Bang (but mostly attempts to show the tremendous difference in scale between human history and the history of the universe) is the ChronoZoom project in Section 4.4.1.

- **Chronological Understanding**
  When we learn about something that evolved over time, it is very likely that we arrange single events belonging to that topic in chronological order in our heads - much like the timeline representation, as seen in Section 3.1. So the timeline basically represents that mental construct on paper und thus facilitates a quick entry into a story. We don't have to have read everything to already get some information out of the ordering. What happend before something else and could have been influencial? What was the aftermath of some other event? Which events took place at the same time? (see *The Guardian* example in Section 4.4.2)

- **Comparison**
  If we have two timelines with for example biographies of two people who lived around the same time and share let's say a scientific field, we could align them, allowing comparison. Meeting points or possibilities of influence could be revealed. Who was first with having an idea? Or who met whom before the other one? Another use case is the comparison of many timelines as in the electronic health records tool *LifeLines* in Section 4.4.1. In that case the timelines are not all aligned absolute to time, but relative to one joint event. That can for example reveal patterns and disease indicators.

- **Guidance**
  Another supportive task of a timeline could be to indicate the reader at what point in time he or she is right now. So when reading about one event in more detail, it could be of interest how that fits into the bigger story. What happend before that and what afterwards? Was that already something happening close to the end or rather at the very beginning of a process? Like an index it gives orientation and makes sure that the reader doesn't get lost (see the *New York Times* example in Section 4.4.2).

4.3 Potential drawback

Even though timelines seem to offer a lot of benefits, we also want to show one counterexample: An experiment by Law et al. in 2005 revealed that sometimes, despite all the glory and simplicity of visual representations, textual summaries are to favor [17]. In their study they let 40 medical professionals make treatment decisions based on either a graphical representation or a textual summary. The graphical representation was a trend graph generated from measurements indicating the physiological state of a patient over a time period of about 40 minutes. The textual summary was generated by experts and just described the course of the measurements in words (without any interpretation). Even though the staff was familiar with that kind of visual representation, the result was that more correct actions were chosen after reading the textual summary of the particular medical scenario (even though most of the physicians liked the graphs better). Also interesting to mention, that there was no difference in speed of response [17]. The evaluation of *The BabyTalk project* by Portet et al. also showed that although graphical representation of clinical data is very valuable and well accepted by clinicians, when it comes to decision-making textual summarizations are more suitable for making correct choices [30]. Their attempt was to instead of letting experts write summaries manually, using NLP techniques to generate them (so called "data-to-text generation"). This approach would make it much less expensive and actually practical in medical environments. The text generation turned out to still have some flaws, e.g. the generated text contains too many irrelevant events, important events are not highlighted, because needed context is missing etc. But the evaluation showed that it already improves the decision-making process compared to visual representation. (Still these approaches are not exploitable at present, because the input data would be too noisy in the real-world setting, but the approach does point in promising direction [30].)

Those studies show that of course we can't say that visual representation is always to be preferred to textual representation. And in most cases they are also not meant to replace text, but rather to be combined with it or as a supplemental offering for all the benefits as outlined above. The designer just has to be careful about not conveying the impression that a visual overview (that is only meant as supplement) already says it all.

4.4 Domains for timelines

Many research areas use timelines for investigation and representation of events. Medicine, Engineering, History only being a few of them. And there are some who would like to use them, but are lacking the tools or the knowledge, as for example in Computer Forensics - Olsson and Boldt claim that evidence plotted on a timeline would help investigators solve crimes faster and more intuitively [26].

4.4.1 Timelines in science

We want to look at two examples, one from medicine, meant to facilitate discovering courses of disease and improve medical treatment - which is used for exploratory tasks - and one from geology or human science, which tackels the timeline scalability problem with interactivity and can be seen as explanatory.

LifeLines This project that was first presented 1996 at CHI by Plaisant et al. and is being improved and refined until today (LifeLines2, EventFlow[41]), visualizes personal medical histories generated from electronic health records. Different incidents from medical records like reported problems, diagnoses, test results or medications are aligned on a timeline to enable the discovery of patterns and "reveal important cause-and-effect phenomena"[43]. Temporal patterns can be highlighted, and it is possible to align several records from different patients one above the other according to one central event, e.g. a heart attack. That enables a comparison of symptoms and treatments between patients, because in that case the exact calendar date is less important than the elapsed time between incidents and the course of the disease itself[29]. The University of Maryland research group continued improving the tool, and now also developed EventFlow, which can not only handle point event data but interval data [41].

ChronoZoom The second example approaches the problem of scalability in timelines. ChronoZoom started in 2009 at the University of California Berkeley. The idea was to develop something that gives an understanding of the very profound difference between the scale of humanity (which starts with the Ancient Egypt, ca. 5,000 years ago), geologic history (the earth shaped around 4.54 billion years ago), and the history of the universe (the big bang, which is supposed to be around 13.8 billion years ago). Doing the math shows: Humanity could fit 2,760,000,000 times into the time of evolution of the universe. Because it is impossible to show such

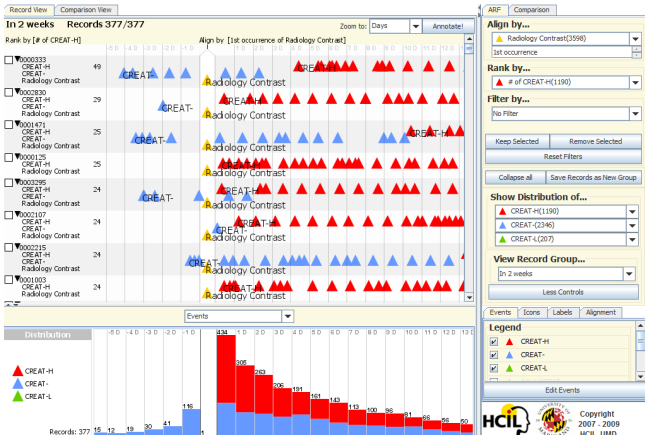

Figure 5: Screenshot of the Lifelines2 project to discovering temporal categorical patterns across multiple patient records. From http://www.cs.umd.edu/hcil/lifelines2

scales on a paper or even a monumental poster, they assumed interactive tools could remedy this problem. But the obvious challenge in scalability needed new deep zooming technologies. Microsoft Research, Moscow State University and the Outercurve Foundation joined in and (after some rather less performant attempts) an interactive, browser-based beta version of ChronoZoom was released in 2012. The tool enables the user to navigate through 13.8 billion years, offering an all-in-one overview and interactive zooming into human history even to the level of selected single day events (which corresponds to a five trillion to one zoom ratio) - important events can be selected to get more information. Animations between the states give an idea of the distances that lie in between. All code is open-source and since 2013 there are also authoring tools available, enabling students, educators and scientists to create their own timelines, with other scientific, personal, or statistical data. According to the authors it offers the "opportunity to algorithmically generate timelines and exhibits that can help researchers in various fields, as well as the general public"[42]. A further aim of the project is now to establish some lesson plans for teaching historical thinking in schools and universities. [42, 27]

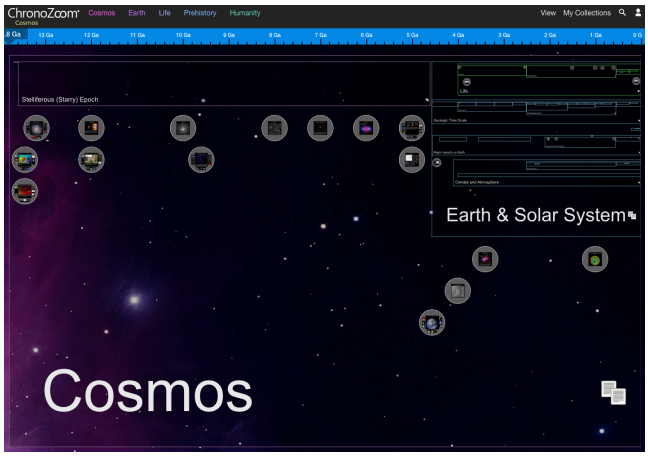

Figure 6: Screenshot of the project ChronoZoom, which dedicates itself to "visualizing the history of everything" [27]

4.4.2 Timelines in journalism

Even though journalism seems to be one of the most obvious areas of application, we don't see the usage of timelines too often. We could imagine them in several scenarios:

- **Summary of previous events**
  As an entry point into a topic, summing up previous events that belong to or led to the current situation - or are just happening at the same time (current news stories often evolve pretty quickly, in that case a timeline could help the reader keeping up pace) - Example 2 *"The path of protest"*

- **Single aspect background**
  Offering background information about a subject inside an article. E.g. when referencing to a politician, scientist or writer, the reader might quickly want to know, when he lived, what his influence was etc.

- **Index**
  The timeline as guide through a chronological story, first to give an overview of the whole story and second to indicate where something belongs to inside the overall structure - Example 1 *"The Higgs, From Theory to Reality"*

We will have a look at two glamorous examples in the following. One from the always pioneering *New York Times*, and one rather unconventional one from *The Guardian*.

The Higgs, From Theory to Reality The most simple and elegant timeline, that gives a quick overview of a news story can be found at the *New York Time*'s website - Figure 7 shows an example. NYT authors used that framework for several stories and it guides the reader through the development of news stories, life stories, scientific explorations or political issues. The accompanying texts and pictures of the single events are listed as normal text underneath each other, the timeline offers navigation through the page on the one hand and indicates the position of the currently looked at event inside the whole story on the other hand. The timeline is so simple and unobtrusive but still offers a very convenient visual guidance through the story.

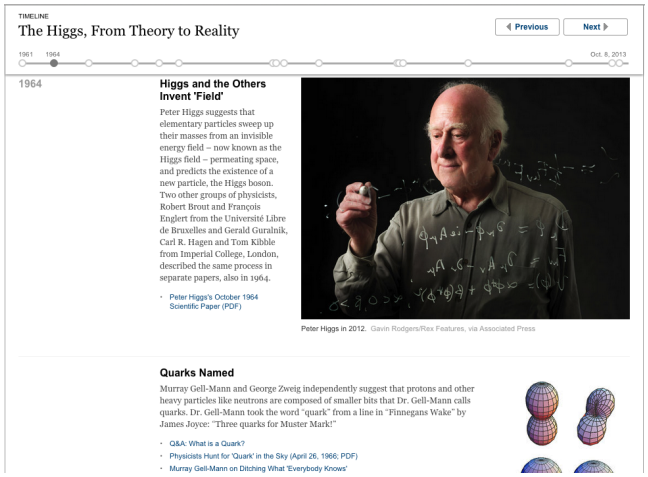

Figure 7: Screenshot of the timeline about the Higgs boson from the *New York Times*, March 2013 [7]

The path of protest One more colorful and ludic approach was done by *The Guardian* for their timeline about the Arab spring, see Figure 8. The reader can navigate through a huge data set of events. The data points are generated from all articles that covered

that topic. All involved countries have their own track on a long three-dimensional uphill road towards the most recent date. Different icons indicate the different categories of events (protest, political move, regime change and international response). Hovering over an event gives a quick summary of the data point, clicking on it guides the reader to its accompanying article. A horizontal timeline at the top lets the reader keep track of the position inside the whole timeframe. It's a rather unusual type of visualization, bringing together horizontal and vertical navigation and it has a huge underlying data set. The interactive graphic can be used for getting a general idea of how much happened during the Arabic Spring, for comparing the course of it throughout different countries or also to look up a specific date of interest and discover what happened around that time.

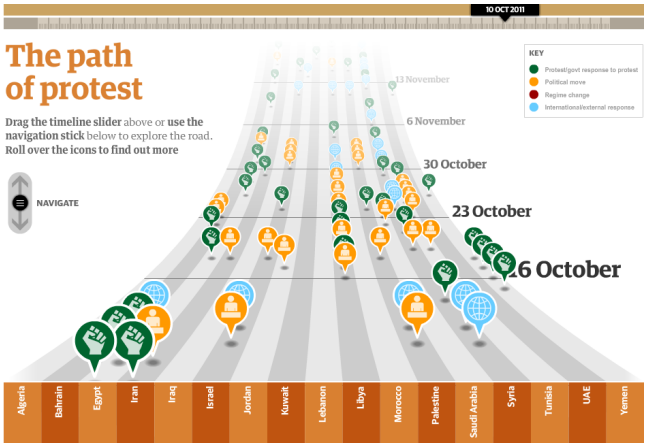

Figure 8: Screenshot of the interactive timeline of Middle East protests from *The Guardian*, January 2012 [3]

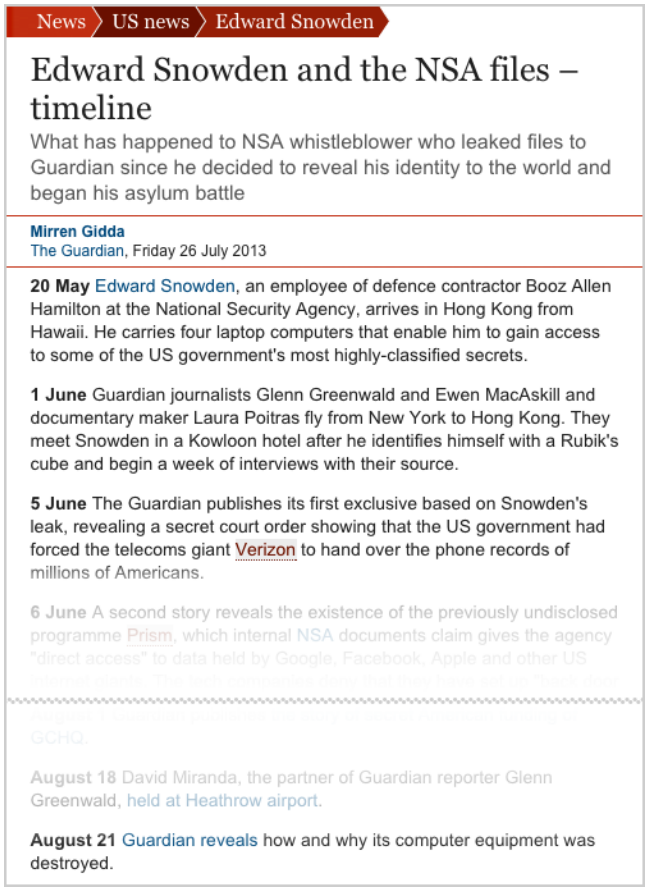

Figure 9: Example of a non-visual timeline from *The Guardian*, July 2013 [9]

4.5 Challenges for the timeline creators

Those two examples are rare appearances in news reporting. The idea of bringing together all important events inside a superior topic can be found quite often, but it is more likely to find those lists without any kind of visualization - and thus initially without any idea of the size and time frame of the whole topic, Figure [**?**] shows an example. The reason why journalists don't make usage of timelines too often are that there are quite some challenges that have to be faced when creating them - time constraints and technophobia only being two of them.

Scalability As seen in previous paragraphs, there are problems of scalability, density and widespread distribution - sometimes it just doesn't work out to fit the whole time frame into one view. Nowadays' interactive exploration tools enable the user to zoom and scale inside the timeline, to either have a complete overview or to see segments in more detail. Already existing timeline tools (which we go into more detail in Section 5) mostly offer an overview first, where elements that are too close together can be aggregated, others let the user rank single events and offer a semantic zoom - with only showing the highest ranked events in the initial overview.

How and what Let's assume a designer doesn't have to resolve those fundamental problems of scale and can just use already existing tools, which solve that issue. But of course these tools still need meaningful input, which has to be provided by the creator. If we assume the content springs from a human author (rather than from a machine as discussed above), he or she has to spend some time thinking about what time frame the story has, which events are interesting and relevant for the general story, when did they happen. Even if we assume the author has a good knowledge about the topic, this process can take some time - which leads us to the time issue.

Time In news rooms especially, but in many other surroundings as well, time is always scarce. Tight deadlines prevent eagerness to experiment. Also it is not very common that writing journalists think a lot about how they could use graphical elements or visualizations for their stories. And those who do usually don't see it as a primary task. That results in them tackling that issues pretty late in the publishing process and this results in not having a lot of time for it. In big publishing companies it is also not common that journalists make the graphics themselves. There is a graphical department that creates them together with, or after scribbles or notes of the journalists. The quite short time that is calculated for that process leads to using only a limited amount of variation. And because of the outsourcing process journalists are normally not familiar with tools and possibilities for visualizations (of course there are many exceptions, especially when it comes to journalists who deal with the evolving topic of data journalism). But the general situation makes it very unlikely that journalists are willing to spend a lot of time with generating graphics - even if they might be aware of them adding value to their stories.

Accessibility Surprisingly though that we sometimes see lists of events (with title, date and summary, as in Figure 9) as "article". This would be the perfect starting situation for a timeline, because the content and structure of the needed data set is already created and it would be just a question of the right tool to throw it in to generate a visual overview. And there are a lot of tools that make

it really easy to create timelines (more about that in Section 5). The user can either just upload a table, or create the events inside a comfortable graphical user interface. Of course the journalists have to know about them and the results vary strongly in visual appeal and usability. Also news sites generally have their very own styleguide, so an author cannot just throw in any form of a generated timeline, which leads us to the question of integration.

Integration News sites work with huge and complex Content Management Systems (CMS). For traditional print authors it might already have been a big commitment to occupy themselves with those systems at all and they don't feel the need to spend any more time with other technical issues. CMS offer different layouts for different formats of articles. News articles, report, commentary, image gallery and so on. But they are not made for people going all creative. Integrating interactive graphics originally is something rather exotic. Nowadays most CMS have the possibility to include iFrames that load in an external link containing e.g. an interactive infographic. But just too often it's not integrated very well, because the size of the article doesn't match the one of the graphic itself and resizing the page or placing the element at a wrong spot ends up in a huge mess. And if that was not enough, nowadays all websites that want to call themselves modern, have to be responsive, which means have to look good on every device, whether it be the tiny little Smartphone or the huge 84-inch screen on the family's living room wall. High visited news sites (like the shining examples *New York Times* or *The Guardian*) are probably very well equipped for all occasions, but less huge websites still seem to struggle with formats that are different to the established, but often outdated ones. So if we want to integrate a timeline into a news article, first an integration feature has to be available inside the CMS, second the style has to adjust to the general styleguide of the website, and third the timeline has to be responsive (at least as much as the website itself) and adjust to different screen sizes to be enjoyable on as many devices as possible. (Things like usability and performance are left to be concerns of the tool itself).

4.6 More exotic timelines

To look a little outside the box of a straight horizontal timeline we want to show two examples of different kinds of timelines. They are both originally hand drawn and can rather be seen as artworks, more than infographics. It is probably necessary to see them in poster size to acknowledge their full beauty.

The Temple of Time Figure 10 shows a graphic created by Emma Willard in 1846. It projects important figures and historical events onto a 3D temple. The intention behind that representation was to facilitate remembering historical facts due to spatial imagination and architectural details. As Willard wrote herself onto the chart: "The attempt to understand chronology by merely committing dates to memory, is not only painful, but [...] useless [...]. The relation which any given event bears to others constitutes the only useful knowledge." She suggests that the viewer should (mentally) locate himself inside the temple and look around, to see the characters of the time. To that time, this kind of illustration was very unconventional and must have been pretty mind-blowing.

The History of U.S. Political Parties That illustration devotes itself to the history of political parties in the United States from the birth of the political system in 1789 until today (or the day of creation of the second part of the timeline). It is a two-part work made by designers living in two different centuries. Around 1894 the first part was created by an unknown designer, 115 years later Larry Gormley and Bill Younker continued the graphic with surely way more modern design tools, but modeled on the look and structure of the original piece. The graphic shows the serpentinous course of popularity of the single parties, combined with details about political and historical events. All the reigning presidents and their

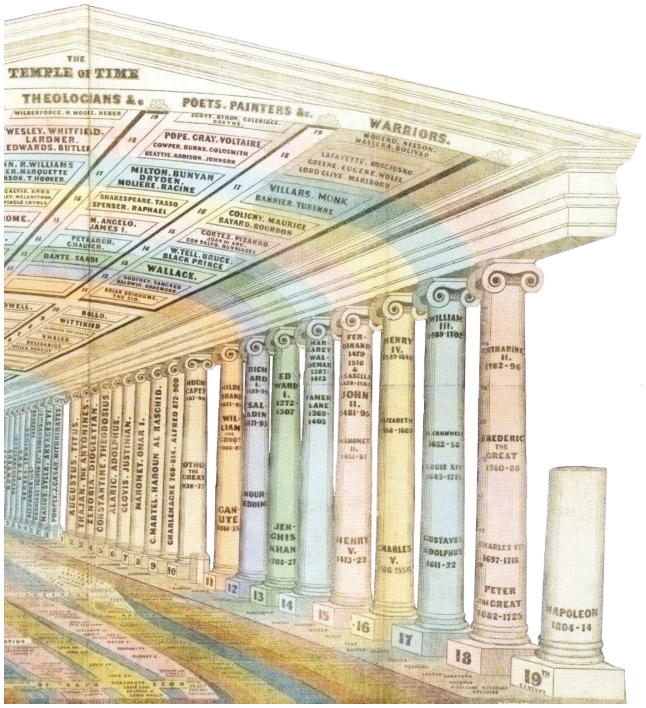

Figure 10: Cutout of "The Temple of Time" by Emma Willard, 1846. The columns represent the centuries, the floor shows historical events, the ceiling biographies of statesmen, philosophers and discoverers, theologians, artists and warriors. The whole graphic can be found on DataVis.ca [8])

cabinet members are listed to give American history enthusiasts a beautiful visual summary to hang on a long wall. The horizontal line indicating time is estranged here to also indicate the popularity, and also uses the possibilities of annotating it with summaries or headlines of events.

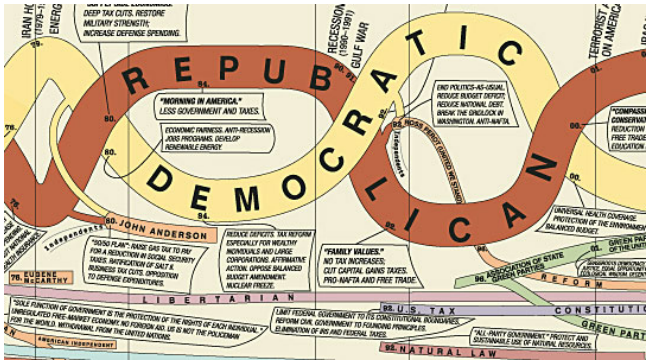

Figure 11: Cut-out of "The History of the Political Parties", the complete two-part poster can be seen at HistoryShots website [11]

5 Tools for Timelines

Overcoming the obstacles of creating a timeline seems only worth the time for big topics, like elections, huge sports events, or as seen above for great scientific achievements or big political issues. But maybe it is just the lack of knowledge and techniques that keeps people from creating them. We want to show that there are tools for every level of technical expertise to create timelines. They are

all freely available online and steadily improved by the enthusiastic web community. The better the technical understanding the more modifications can be made by the creator.

5.1 For non-programmers

Programming languages are becoming more popular, not least because they are becoming more understandable. Nowadays they are taught in school and the digital natives are often familiar at least with the fundamental principles of them. Still there are many people that have a very skeptical attitude towards writing code or are just not familiar with it. But also for those fellows there are tools. Their advantages are they easy accessibility, most often it is quite intuitive to handle and they deliver a quick outcome. But of course that brings along disadvantages, e.g. that the users don't have influence on how it will look like - and they often look rather clunky. Also the possibilities of integration into external sites vary and (at least the freely available ones) are not always reliable in performance. Still we want to mention two of them, which do a fairly good job and are easy and comfortable to use.

Timerime Also available as commercial software for professionals, it is freely available for the online community. Signing up and getting something started is very easy, but it is also possible to just look at publicly available timelines, which were created by others. The backend side of the application also allows people to work together on a timeline from different computers. When creating a new timeline, the dates for the single events can either be year only or exact to the second and it can be a data point or a time period (with start and end date). A very nice feature here is, that the user can set the level of importance. So you end up having a semantic zoom, where only the highest rated events are shown at the zoomed out view. Every event can have text, links, external multimedia (like YouTube videos or Google Maps), pictures, even music or sounds. A big drawback is, that it is based on Flash, which is pretty outdated nowadays and restricts the availability on some devices. Also the loading times sometimes take very long and it happens that the application crashes when creating new events (maybe something that works better in the professional solution). The visual appeal can be judged individually - all in all it's a nice tool, not only for creating own timelines but also to discover what the community created.

Figure 12: A timeline about Steve Jobs and Bill Gates, created by TimeRime. www.timerime.com

Dipity The attempt of this tool is to organize web content by its time. Timelines can automatically be populated via platforms like YouTube, Twitter, Flickr, Google News or Tumblr. The creator enters a search term and the top search results will be inserted into the timeline (according to their creation time). There is even an API that connects to a SMILIE timeline XML file (more about SIMILE in 5.3), but unfortunately no option to upload an excel sheet or the like, which would make it even more useful. Nevertheless the imported events can then be edited or deleted very easily and combined with other search queries or manually generated input. In this tool however there are only data points available, no time periods. Populating the timeline with dates seems rather randomly, but is a very quick way to start. The purpose behind those timelines is mainly ordering social media content or Internet memes (see Figure 13), and for that it is a very good and functional tool. Because it is all HTML it can be integrated into other websites without restrictions, and of course it can easily be shared via social media platforms. Again the beauty of the appearance is subjective, but seems kind of suitable for its clientele.

Figure 13: A timeline about Internet Memes, created by Dipity. www.dipity.com

5.2 For technophiliac semi-programmers

If the user wants to have a little more control over appeal, behavior and field of application and has a basic knowledge about web-programming, there are some more sophisticated approaches. The first example, TimelineJS is even suitable for non-programmers, but because it offers additional possibilities of customization, we decided placing it in here.

TimelineJS The Northwestern University Knight Lab is dedicated to news media innovation and education. 2012 they released this open-source tool for journalists, which is probably the most used tool for timeline creators and can be found on various news sites around the world, from *Aljazeera* to *RadioLab* to the *TIME magazine* website, it has even been part of a *Pulitzer Prize*-winning story by *The Denver Post* about the shooting in the Aurora theater in 2012, which is shown in Figure 14. The creators even claim it is "among the most widely used interactive storytelling tools on the web"[25]. The Aurora theater-shooting story demonstrates that dates can either be in the format of the exact minute of an event or on a day level (month- or year-level are not designated). TimelineJS is suitable for non-programmers as well as for advanced users, because the timelines can either be created using Google spreadsheet (and just exporting it to the web, using their servers) or it can be manually configured and implemented into the user's personal code (by downloading the code and using JSON as data format). Many different media types can be included very elegantly, even PDF documents (which is pretty unique), excerpts from Wikipedia articles and any website, by just entering the URL (it shows a snapshot of the site and links to the URL when clicking on it). Because of its simplicity and responsiveness it fits well into all kinds of websites, and the community contributed in translating and improving the tool all the time.

Google Charts Timeline Less dedicated to news stories, but for example to visualizing project schedules or conference timetables as in Figure 15, this approach uses *Gantt charts* to illustrate the start and finish of events. It is based on script (JavaScript), but the implementation is fairly simple. Users can just copy and paste the code from the many examples in the documentation, and modify the values - the scale is adjusted according to the whole time span of the data and fit into the assigned div container in the HTML code. More advanced programmers can make adjustments to their

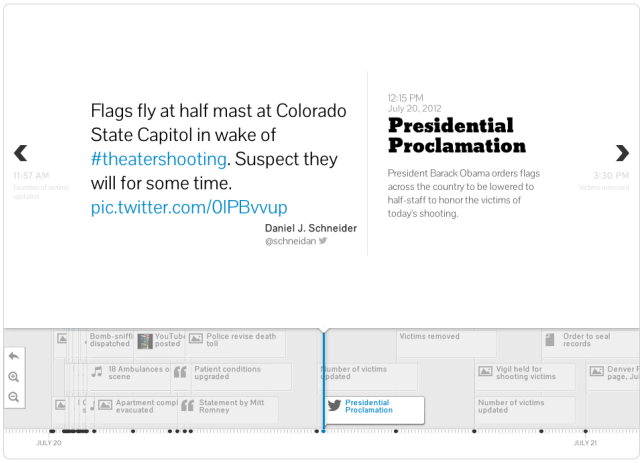

Figure 14: *Pulitzer Prize*-winning story about the movie theater shooting in Aurora, Colorado on July 20, 2012 [20].

liking, change styles and behaviors. It doesn't offer a graphical user interface, and some basic code skills are needed, but it is very simple and elegant and can serve as a solid foundation for further development.

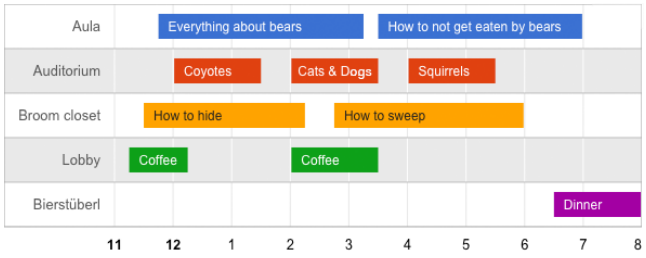

Figure 15: Example schedule created with Google Chart's Timeline tool [10].

5.3 For programmers

The previous presented tools are suitable either for inexperienced programmers or for making a quick draft. The following two are rather appropriate for bigger projects, because they are highly customizable. For experienced programmers everything is possible, but the creators have to have a clear idea of what should be shown and know if they have enough time to implement it.

SIMILE The MIT project on **S**emantic **I**nteroperability of **M**etadata and **I**nformation in un**L**ike **E**nvironments developed several open-source tools, enabling users to access, manage, visualize and reuse data. One of them is *Timeline* a "Web Widget for Visualizing Temporal Data"[21]. It started in 2006, the latest release was in September 2008 (version 2.2.0), but the community is still active and discussions are going on. To tackle the scalability problem, this tool offers to use several "bands" with different fixed scales. This enables to have an overview sight and an enlarged detail view at the same time. But it can also just have one band and a magnifier to zoom in and out - a lot of options allow many variations, through customization even more is possible. The code is in JavaScript, the events are implemented through XML, JSON or SPARCL files. Events can have six different formats, with all combinations of being a duration or an instantaneous event and having precise or imprecise starting (and ending) dates. Impreciseness is marked through a transparent line over the indicated possible time (something we haven't seen in any of the previous tools). It is also possible to let the user search for certain events (through keywords and highlighting) or filter out events to only show what's interesting for the user (as shown in Fig. 16). Even though it is already a little older (in the Internet scale of old), it tackles a lot of the issues that appear when working with temporal data - like imprecise dates, durations, overview and context (not only through zooming but different scaled bands). It takes a while to get familiar with it, and it doesn't appear very modern, but e.g. for big historic events it does a really good job of giving an overview.

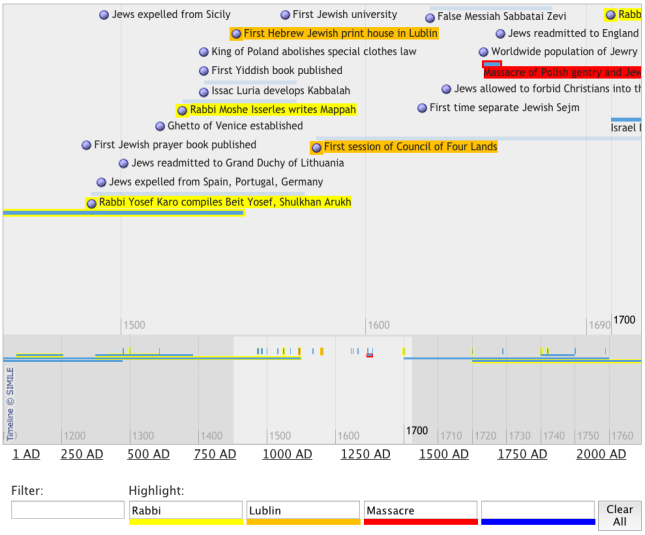

Figure 16: Timeline about Jewish History over 2000 years. With the oportunitiy for filtering and highlighting [13].

d3-timeline D3 is a JavaScript library for **D**ata-**D**riven **D**ocuments, springing from the Visualization Group at Stanford University. It is a powerful tool for creating dynamic and interactive data visualizations using HTML, CSS, and SVG. The underlying data can come from various sources - almost everything is possible as long as it is structured. The most common data format is JSON or CSV. Transitions, popups and any other dynamic effects can be included easily through pre-built JavaScript functions. It has a huge Internet community and new plug-ins and extensions are built every day. One of these is the d3-timeline by Jia Huang [12]. It takes temporal data as input and can display it either as dots (for instantaneous event) or bar (for durations), it can be all aligned next to each other on one horizontal line or vertically separated by categories, labels can have different appearances and some callback methods are available. Also the scale automatically matches the input data. It basically offers some nice features but is completely customizable by anybody who is familiar with writing code in D3. Figure 17 shows a very basic example of how a very small data set could be visualized.

6 Conclusion

We looked at various aspects of timelines: How temporal information is transferred in language as well as in machine readable form, how we perceive temporal information mentally as well as visually, in what context timelines are beneficial and what tools can be used for creating timelines.

Looking at the transmission of time we pointed out certain grammatical rules that on the one hand give us a direction of when something happened and on the other hand can be used to translate temporal expressions into machine-understandable data. Natural Language Processing offers opportunities to extract information about

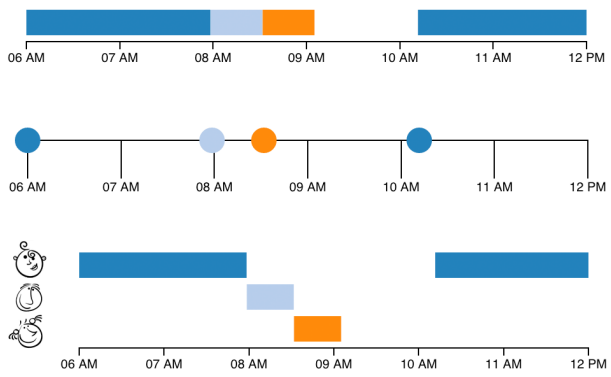

Figure 17: Screenshot of example application of d3-timeline [12]

time inside written text and normalizes it into a format that is valid for a machine and matches defined standards. Looking at the human cognition of time, we saw that combining space and time is not only present in our language, but also in our minds. Studies showed that we use our spatial imagination to think about time and that it is very intuitive to order events on a straight line according to their appearance. The direction of alignment depends on physical and cultural factors - in English speaking environments it is very common to use a left-to-right alignment.

For visual literate people visualizing a sequence of events on a timeline seems like something very obvious. Still we discovered that historically as well as in non-visual trained surroundings it is not too common to use that representation when demonstrating single events inside a story. But the possibilities could and should be exploited more often, because there are some obvious benefits. They can for example serve as a quick introduction or update of a story, superseding the initial mental efforts of ordering events chronologically in our heads, and thus quicken the understanding process. Also timelines offer a general chronological overview, indicating when something began and ended, when a lot of things happened in parallel and how events could have influenced each other. And they can serve as a temporal orientation inside a story.

We showed some flagship examples from science as well as from journalism, which demonstrated that timelines can have a positive impact on medical treatment or historical understanding, even though there might be situations where a raw textual representation of events is to favor.

But we think possibilities are not exhausted yet and we would like to see more timelines, especially in the news. We presented tools for all levels of technical expertise of users that enable a fast and easy creation. Because of the growing popularity of data-driven journalism and the possibilities of computer generated (or at least computer-aided) visual representation of data we are positively looking forward to further development of interactive visual elements inside news reporting and many other domains.

REFERENCES

[1] S. Archibald and D. Rosenberg. A Timeline of Timelines. *Cabinet Magazine*, (http://www.cabinetmagazine.org/issues/13/timelines.php), 2004.

[2] R. Barzilay and K. McKeown. Sentence fusion for multidocument news summarization. *Computational Linguistics*, 31(3):297–328, 2005.

[3] G. Blight, S. Pulham, and P. Torpey. Arab spring: an interactive timeline of Middle East protests. *The Guardian*, (http://www.theguardian.com/world/interactive/2011/mar/22/middle-east-protest-interactive-timeline), 2012.

[4] L. Boroditsky. Metaphoric structuring: understanding time through spatial metaphors. *Cognition*, 75(1):1–28, Apr. 2000.

[5] J. Christensen, Mausam, S. Soderland, and O. Etzioni. Towards Coherent Multi-Document Summarization. In *HLT-NAACL*, number Section 3, pages 1163—-1173, 2013.

[6] J. Christensen and S. Soderland. Hierarchical Summarization: Scaling Up Multi-Document Summarization. *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, 2014.

[7] J. DelViscio and D. Overbye. The Higgs, From Theory to Reality. *New York Times*, http://www, 2013.

[8] M. Friendly. Timelines and Visual Histories. *DataVis.ca*, (http://www.datavis.ca/gallery/timelines.php), 2010.

[9] M. Gidda. Edward Snowden and the NSA files timeline. *The Guardian*, (http://www.theguardian.com/world/2013/jun/23/edward-snowden-nsa-files-timeline), 2013.

[10] Google Developers. Google Charts Timeline. *Google Developers*, (https://developers.google.com/chart/interactive/docs/gallery/timeline).

[11] L. Gormley and B. Younker. History of the Political Parties. *HistoryShots*, (http://www.historyshots.com/partiescombined), 2012.

[12] J. Huang. D3-timeline. *GitHub*, (http://github.com/jiahuang/d3-timeline), 2012.

[13] D. Huynh. SIMILE Timeline Jewish History. (http://www.simile-widgets.org/timeline/examples/religions/jewish-history.html).

[14] W. Klein. How time is encoded. *The expression of time*, 3:1–43, 2009.

[15] E. Làdavas. Asymmetries in processing horizontal and vertical dimensions. *Memory & cognition*, 16(4):377–382, 1988.

[16] G. Lakoff and M. Johnson. *Metaphors we live by*. The University of Chicago Press, 1980.

[17] A. S. Law, Y. Freer, J. Hunter, R. H. Logie, N. McIntosh, and J. Quinn. A comparison of graphical and textual presentations of time series data to support medical decision making in the neonatal intensive care unit. *Journal of clinical monitoring and computing*, 19(3):183–94, June 2005.

[18] I. Mani. Computational Modeling of Narrative. *Synthesis Lectures on Human Language Technologies*, 5(3):1–142, Dec. 2012.

[19] M. S. McGlone and J. L. Harding. Back (or forward?) to the future: The role of perscpective in temporal language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(5):1211, 1998.

[20] C. Minshew and D. J. Schneider. Interactive Timeline: Shooting at Aurora theater. *The Denver Post*, (www.denverpost.com/breakingnews/ci_21119904), 2012.

[21] MIT. SIMILE Timeline. (http://www.simile-widgets.org/timeline/).

[22] M. Mitchell. The Visual Representation Of Time In Timelines, Graphs, And Charts. *Conference paper delivered to the Australian & New Zealand Communication Association Conference 2004.*, 2004.

[23] W. Müller and H. Schumann. Visualization Methods For Time-Dependent Data - An Overview. In *Simulation Conference, 2003. Proceedings of the 2003 Winter*, pages 737—-745. IEEE, 2003.

[24] T. Munzner. Visualization Analysis and Design (Chapter 2). A K Peters Visualization Series, 2014.

[25] Northwestern University Knight Lab. TimelineJS. *Knight Lab Projects*, (http://projects.knightlab.com/projects/TimelineJS).

[26] J. Olsson and M. Boldt. Computer forensic timeline visualization tool. *Digital Investigation*, 6:78–87, Sept. 2009.

[27] Outercurve Foundation. ChronoZoom - Zoom through all of time. (http://chronozoom.tumblr.com/about), 2013.

[28] R. Pasts. Why Time is so Now. *Rethinking History: The Journal of Theory and Practice*, 16(2):303–317, 2012.

[29] C. Plaisant, R. Mushlin, and A. Snyder. LifeLines: using visualization to enhance navigation and analysis of patient records. *Proceedings of the AMIA Symposium. American Medical Informatics Association*, 1998.

[30] F. Portet, E. Reiter, A. Gatt, J. Hunter, and S. Sripada. Automatic generation of textual summaries from neonatal intensive care data. *Artificial Intelligence*, 2009.

[31] J. Priestley. *A description of a chart of biography*. Printed at Warrington, 1765.

[32] L. H. Reuter. Human Perception and Visualization - Introduction. In *VIS '90 Proceedings of the 1st conference on Visualization '90*, page

401, 1990.

[33] D. Rosenberg. The Trouble with Timelines. *Cabinet Magazine*, (http://www.cabinetmagazine.org/issues/13/timelineIntro.php), 2004.

[34] D. Rosenberg. Joseph Priestley and the Graphic Invention of Modern Time. *Studies in Eighteenth Century Culture*, 36:55–103, 2007.

[35] D. Rosenberg and A. Grafton. *Cartographies of time: A history of the timeline*. Princeton Architectural Press, 2013.

[36] R. N. Shepard and S. Hurwitz. Upward direction, mental rotation, and discrimination of left and right turns in maps. *Cognition*, 18(1-3):161–193, Dec. 1984.

[37] E. C. Traugott. Spatial expressions of tense and temporal sequencing: A contribution to the study of semantic felds. *Semiotica*, 15(3):207–230, 1975.

[38] E. R. Tufte. *The Visual Display Of Quantitative Information*, volume 2. Graphics press Cheshire, CT, 1983.

[39] P. Tukey. Human Perception and Visualization - Visual Perception and Scientific Data Display. In *VIS '90 Proceedings of the 1st conference on Visualization '90*, pages 401–403, 1990.

[40] B. Tversky, S. Kugelmass, and A. Winter. Cross-Cultural and Developmental Trends in Graphic Productions. *Cognitive Psychology*, 23(4):515–557, 1991.

[41] University of Maryland. EventFlow: Visual Analysis of Temporal Event Sequences and Advanced Strategies for Healthcare Discovery. *Human-Computer Interaction Lab*, (http://www.cs.umd.edu/hcil/eventflow/), 2014.

[42] R. L. Walter, S. Berezin, and A. Teredesai. ChronoZoom : Travel Through Time for Education , Exploration , and Information Technology Research. *Proceedings of the 2nd annual conference on Research in information technology*, pages 31–36, 2013.

[43] T. D. Wang, C. Plaisant, A. J. Quinn, R. Stanchak, S. Murphy, and B. Shneiderman. Aligning temporal data by sentinel events. In *Proceeding of the twenty-sixth annual CHI conference on Human factors in computing systems - CHI '08*, page 457, New York, New York, USA, Apr. 2008. ACM Press.