Visualising Data in Web Observatories: A Proposal for Visual Analytics Development & Evaluation

Paul Booth
Electronics and Computer Science
University of Southampton
pmb1g11@soton.ac.uk

Wendy Hall
Electronics and Computer Science
University of Southampton
wh@ecs.soton.ac.uk

Nicholas Gibbins
Electronics and Computer Science
University of Southampton
nmg@ecs.soton.ac.uk

Spyros Galanis
Economics Division,
School of Social Sciences
University of Southampton
s.galanis@soton.ac.uk

ABSTRACT

Web Observatories use innovative analytic processes to gather insights from observed data and use the Web as a platform for publishing interactive data visualisations. Recordable events associated with interactivity on the Web provide an opportunity to openly evaluate the utility of these artefacts, assessing fitness for purpose and observing their use. The three principles presented in this paper propose a community evaluation approach to innovation in visual analytics and visualisation for Web Observatories through code sharing, the capturing of semantically enriched interaction data and by openly stating the intended goals of all visualisation work. The potential of this approach is exampled with a set of front-end tools suitable for adoption by the majority of Web Observatories as a means of visualising data on the Web as part the shared, open, and community-driven developmental process. The paper outlines the method for capturing user interaction data as a series of semantic events, which can be used to identify improvements in both the structure and functionality of visualisations. Such refinements in user behaviour are proposed as part of a new methodology that introduces Economics as an evaluation tool for visual analytics.

Categories and Subject Descriptors

A.0 [**Conference Proceedings**]

D.2.2 [**Software Engineering]**: Design Tools and Techniques – *software libraries*, *user interfaces*

General Terms

Management, Measurement, Documentation, Design, Economics, Experimentation, Human Factors, Standardization.

Keywords

Visual Analytics, Data Visualisation, Web Observatory, Economics, Evaluation.

1. VISUAL ANALYTICS & WEB OBSERVATORIES

The Web Science Observatory was envisioned as a 'global data resource and open analytics environment' [1]. In general terms a Web Observatory (WO) is repository of Web data 'structured such that observations can be made, activity can be monitored, and experiments may be performed' [2]. These concepts are well supported by the growing research in Visual analytics (VA), which is defined as 'the science of analytical reasoning facilitated by interactive visual interfaces' [3]. As a field of study VA is a tool of illumination for a range of industries that depend increasingly upon the valuable insights obtained through data analysis. Transforming data into valuable insights is a primary objective for Web Observatories and all observatories stand to benefit if a form of standardised for developing VA applications can be agreed upon. Standardising methods, frameworks and critical perspectives will contribute towards development and evaluation of VA tools - effective evaluation being a critical factor in iterative design and development. There is a unique opportunity for WOs to develop a common framework ahead of the general VA community, beginning with the proposed methods in this paper. This approach draws on the recommendations from the pivotal VA report by Cook and Thomas [3]. Each recommendation in the report highlights an area of necessary advancement for VA research as follows:

1. The science of analytical reasoning
2. Visual representations and interaction techniques
3. Data representations and transformations
4. Product, presentation and dissemination
5. Moving research into practice

The focus herein is with evaluation as a central requirement of a VA framework for WOs, which is expressed under the fifth heading of the report - *moving research into practice*. In particular the report calls for the development of an infrastructure to support common methods and measures for VA evaluation. These include the need to 'clearly articulate research hypotheses to be verified through evaluation', 'encouraging and challenging researchers in a particular area', 'comparing technical approaches' and 'determining progress through definable achievements' [3]. The following three principles have been devised from the report and are hereby proposed for use by WOs that are developing visualisations and VA tools:

1. Openly articulate the purpose / goals of any WO visualisation or VA application
2. Share the code (e.g., on GitHub) for community use, evaluation and development
3. Collect and publish interaction data in a meaningful way for community evaluation (measurable against the first principle) and analysis

2. CLASSIFYING VISUALISATION TYPES

By definition these principles immediately exclude the use of tools that do not facilitate code sharing for interactive artefacts (e.g., Flash), or would be created as a flat graphics file (e.g., with Adobe Illustrator, or as GIF, JPG, or PDF files). The terms

WWW'14 Companion, April 7–11, 2014, Seoul, Korea.
ACM 978-1-4503-2745-9/14/04.
http://dx.doi.org/10.1145/2567948.2579211

information and *data* are often used interchangeably, the meaning of either covering both interactive and non-interactive graphics that visually communicate data as information. The term *info-graphics* is a used for images that choose design and style over functionality and clarity in data communication. Furthermore, the 'depth' of interactivity can also define the term used for a visualisation, with some visualisations (and desktop visualisation tools such as Tableau[1]) offering more exploratory experiences falling under the term of Visual Analytics. While visualisations that lead users along a sequential and narrative path of interactions and insights have been analysed as having a 'visual rhetoric'[4] and also categorised extensively into genre, narrative type and structural elements [5]. The concept of narrative in visualisation and VA is something of a current trend in visualisation [6], supporting new questions about measuring engagement and insight [7].

3. NARRATIVE VISUALISATION & ENGAGEMENT

Measuring a user's engagement with online content is a key objective for online publishers. Building a story around data presented on the Web has driven *narrative visualisation* to be an emerging research area ([8], [5], [9], [10], [7]). Although narrative may help to engage users in a story supported by data, a key question remains as to whether or not engagement through narrative is more relevant to the types of journalistic visualisations that seek to disseminate known insights and formulated opinions in the media than those that facilitate more objective insights and decision-making. Supporting the argument for the conceptual bisection of engagement and insight is the difference between users being led towards known insights (e.g., narrative visualisation) and the alternative of providing analysts with the tools to uncover new insights and to inform decision-making (e.g., visual analytics). By comparison, the broad subject of Design is a unifying factor in both VA and narrative visualisation that positively or negatively affects user engagement and also facilitates or impedes insight [4].

4. FINDING INSPIRATION

Having followed the first principle ('Openly articulate the purpose / goals of any WO visualisation or VA application'), Observatories have an opportunity to share previous experience and resources to support a decision about how best to visualise the data at hand. For example, if the goal to share known insights with users or perhaps provide a data-centric interactive educational resource then simple charts, then interactive info-graphics or the use of a journalistic narrative visualisation approach may be most suitable. To facilitate new insights through visual analysis, to enable data discovery, or to support decision making for a given task, then a VA application would be more appropriate. There are numerous existing examples of all these that can already be found online.

The use of visualisation in terms of decision support now extends into (and beyond) public safety and security, finance, insurance and climate science (Sedig, Parsons and Babanski 2012:1). There are well known examples of interactive data visualisations, some with narrative, and others that are more akin to VA, that have the potential to inspire solutions for WOs with less practical knowledge of this area. Highlights would include the work of Hans Rosling[2] whose visualisations have been seen on TED.com[3] and the BBC[4]. David McCandless[5] whose work has been featured in the Guardian newspaper online, and Mike Bostock[6] who is one of the original and *the* continuing lead author for D3.js currently working for the New York Times. Looking across only these three will open up a world of developmental tools, methods, and styles used to present data interactively (or statically) on the Web. No one tool can claim to be suitable for the requirements and resources of every WO, however we need to start somewhere if we are to share best practice in VA and visualisation development. A major goal must be to move forward together with evaluation at the heart of any framework in order to improve visualisations by better understanding the utility of this work in the future. The following sections focus on delivering the proposal for a visualisation framework to support the creation of VA applications, narrative visualisation and data visualisation component graphs and charts. These recommendations are born from professional work completed at the University of Southampton with two organisational partners.

5. DATA VISUALISATION WITH JAVASCRIPT

D3.js (Data-Driven Documents) is the JavaScript library that is a central recommendation for any WO looking to visualise data for the Web. D3 looks to have been influenced from Bertin's [11] work on visual encoding, Wilkinson's Grammar of Graphics [12] and from the previous tools created by some of the same developers [13]. It has been influential on, and arguably has also been influenced by Wickham's ggPlot2 data visualisation package for the widely adopted and open-source R statistics tool [14]. In their supporting paper, Bostock, Ogievetsky and Heer [13] outline specific objectives for D3.js that should be considered by WOs:

Compatibility. Tools do not exist in isolation, but within an ecosystem of related components. Technology reuse utilizes prior knowledge and reference materials, improving accessibility.

Debugging. Trial and error is a fundamental part of development and the learning process; accessible tools must be designed to support debugging when the inevitable occurs.

Although set out as low-level objectives for the D3.js library, these points highlight the limitations for compatibility when closed technologies such as Flash are used, as they have been for some commendable WO projects already[7]. As such, sharing visualisations that are not transparent between observatories and the iterative improvements that could be made in a 'trial and

[1] See http://www.tableausoftware.com/

[2] Rosling's work can be seen on http://www.gapminder.org/

[3] Rosling has appeared in five TED videos in total, all currently available for viewing on www.ted.com

[4] Rosling presented a one-hour documentary titled 'The Joy of Stats' on BBC Four, last transmission at time of writing was Wednesday 16th 2013.

[5] Self described as a 'data journalist', McCandless' work can be found on www.davidmccandless.com and www.informationisbeautiful.net, he has also delivered a lecture viewable on www.ted.com

[6] Bostock's official site (http://bost.ocks.org/mike/) shows a mixture of personal and New York Times visualisations

[7] Existing work at TWC LOGD Rensselaer has employed Google Charts and Flash technologies to visualize their impressive work in linked data. http://logd.tw.rpi.edu/demo/trends_in_smoking_prevalence_tobacco_policy_coverage_and_tobacco_prices

error' process are not rendered as community issue. By using Flash and generating charts with Google Charts API it is not possible to track events in a meaningful and open way as it is with D3.js.

5.1 Compliant & Transparent

D3.js is set apart from Flash and other visualisation tools by being openly viewable in the DOM, by the use of W3C API standards as well and for maintaining up-to-date and extensive documentation with strong online community support. What is rendered to the browser is viewable to anyone who chooses to view the source code in his or her browser. D3 wraps the W3C Selectors API to identify document elements, where selections made are divided into *enter, update* and *exit* sub-selections for visual encoding. Using a wrapper for the W3C DOM API, a number of operators can be applied within the 'd3' namespace to programmatically get or set attributes, styles and properties once the data operator binds input data to selected visual elements from the DOM. D3 also enables simplified animated transitions between selection states (e.g., dots on a scatterplot can fade out or move to a new location following a new *update* selection using new data). At every stage of this process the resulting code and changes in data are transparent to the user and the browser. This fundamentally enables the event tracking proposed in this paper; the collection of low level user actions that can be consolidated into higher level goals and tasks against which evaluations about visualisations can be made. With D3 at least, events can be tracked as inline code in code or preferably as part of a 3rd party script that can be shared and improved upon by the community along with other modular components the D3 library enables.

5.2 Reusable Components

Establishing a modular framework with open-source, and reusable components addresses the second principle - *share the code (e.g., on GitHub) for community use, evaluation and development*. The examples already on d3js.org, GitHub, on blogs and tutorial sites that openly share D3.js code that works well as 'scaffolding' for new projects. These convert easily from generalized examples to the specific task at hand. For the existing WO community the barriers for sharing any newly completed components (e.g., geographic maps, network graphs and charts) should be lower, perhaps more so if the code is centrally catalogued. Beyond singular charts and graphs though, there exists the potential for a WO visualisation framework of reusable charts and added functionality for interactivity between graphs that builds upon the D3 library.

A partnership between a leading electronics research company and a PhD researcher at the University of Southampton has led to development of a framework for reusable visualisation components using the Backbone.js MVC[8] JavaScript library and D3.js. The result is a set of reusable JavaScript objects that can be deployed to a HTML page as a simplified JavaScript object, the methods and properties of which are used to connect to data sources and set attributes for the visual encoding of data facilitated by D3. Styles are managed as a separate concern through external CSS files. Interactivity between components i.e., 'brushing', is also a feature of the framework made possible in prototype experiments using Crossfilter.js. As an example, a user can select one element in a social network graph by interacting with other linked chart components (say by clicking on a bar in a bar chart) that provides a powerful framework for further development in a WO context. There is no insistence here that Backbone.js is the best choice, but against the size and complexity of Angular.js and the limiting design patterns imposed by Ember.js it was felt that Backbone.js would offer the most flexibility while enabling rapid development and frequent iterative evaluations. Plans are already underway to experiment with Angular.js as an alternative to using Backbone.js and to also move away from static CSV and JSON by extending the JavaScript stack with Express.js and MongoDB. The plans for addressing the solution for the final principal - *collect and publish interaction data in a meaningful way for community evaluation* – is outlined in section 5 below with the collection of semantically enriched interaction data.

6. COLLECTING DATA FOR EVALUATION

Simply publishing VA tools and sharing single page applications or any visualisation for that matter is no longer enough. The tools exist to capture and analyse data about how Web artefacts are being used in the form of Web analytics, yet the practice of doing this in the context of visualisation is not widespread or standardised. Currently insights into basic interactivity help in general terms to know more about page views, mouse movement and click activity. However, the proposal made in this paper is that evaluating the utility of data visualisations and VA tools on the Web requires that the information be sharable by using semantic event tracking. For example, by making use of DOM events, actions captured as a *click* event can be given an analytic context of *delete, undo, filter, zoom, sort* or *inspect*.

6.1 Collecting Semantic Events

The Semantic Action Taxonomy (SAT) from Gotz and Zhou [19] provides a classification of actions that supports the evaluation of visualisation. Emerging areas of VA research have focused on extending the utility of interaction data for use in analytic collaborations [15], [16], to add functionality and make usability improvements to software [17] and to dynamically adapt visualisations to the analyst's task [18]. The SAT is shown in Figure 1.

Using a common vocabulary for describing events means that

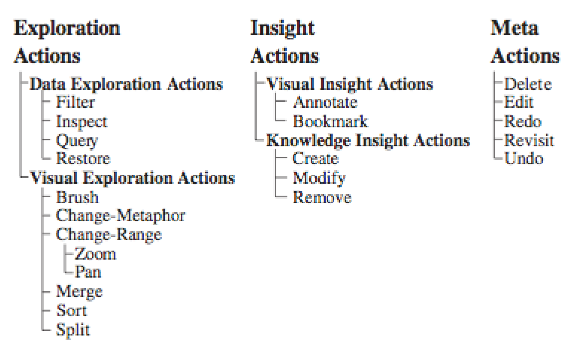

Figure 1: The Semantic Action Taxonomy. Gotz and Zhou, 2009

---

8 Backbone.js is considered a library more than a traditional MVC framework. An full explanation of the use of the View class as a controller can be seen at http://backbonejs.org/#FAQ-mvc

higher-level tasks can be determined and inefficiencies in the design of work can be identified. Indeed, the semantics based approach has already led to technology that 'supports automatic detection of user action patterns for better visualization recommendation' [18]. Examining interaction data in VA tools is an important step towards improving analysts' performance [15], and has been used to 'assist in building more effective analysis environments' [19]. These possibilities are all relevant to WOs, starting with logging data for evaluating the goals of a visualisation. WOs must move beyond sharing a list of tools and instead provide links to reusable resources and interaction logs that can be evaluated by the community. This would align with the recommendation from Cook and Thomas (2005) that data be openly shared amongst researchers in order to contribute to the progression of VA, from which the third principle proposed here was derived - *collect and publish interaction data in a meaningful way for community evaluation*.

6.2 Moving from Web Logs to Event Logs

Current web-log analytics tools enable publishers to see information on the user; geographic location, time (arrived and time spent on site), pages visited, as well as browser, OS and device breakdowns. Beyond the common user/agent segmentation is the potential to capture semantic actions and events. Segment.io has created several libraries for doing this in client-side JavaScript, Python, Node, Ruby and Java that are able to act as a wrapper for web analytics services.

These libraries allow developers to use a single code source to integrate with dozens of analytic services (e.g., Google Analytics, SalesForce). Core methods in the library are: *identify, track, group* and *alias*. *Identify* links a user to their actions to record traits (email, name, subscription plan), while the *track* method records an event in the application. Essentially any event listener in the DOM – created with a simple JQuery script or as part of a Backbone.js framework implementation for example – can be used to trigger an event with the segment.io libraries. As D3.js outputs to SVG to the DOM, actions and events can be tied directly to interactions with visualisations, making it possible to track the behaviour of an analyst or regular Web user.

6.3 A Vocabulary for the Structural Elements of Visualisation

Tracking events semantically is beneficial if interaction logs are to be shared with others, but without better knowledge of the original site how well will anyone else be able to comprehend the data in context; the layout, the types of graphs, maps or charts used? Developers can solve this using the 'special properties' permitted in the segment.io libraries for a given task. The JavaScript version of the library (Analytics.js, which is implemented dicrectly into Google Analytics and does not require WOs to necessarily invest in a Segment.io account) takes the following general form for tracking:

analytics.track (action, properties)

This can be applied using the following code structure:

analytics.track('Selected a bar', { value: 1500,
graph: 'bar-chart', rank: 1, type: 'bar', action: 'filter'
});

The value in the data - bound to the click event 'Selected a bar' - is retrievable, as is the class property ('bar-chart'). Rank is speculative here and untested currently, but examples the potential to record exactly which bar in a graph was clicked (the first, second… last etc.). 'Bar' and 'filter' in this example are existing properties generated in the D3 code and made accessible in the DOM while action is attached to the *analytics.track* event code directly. This method is still in the early stages of testing, but what is evident is that the use of a consistent language is vital if event data is to be shared. As mentioned earlier, D3 shows itself to be influence by Bertin's work on visual encoding [11], where commonalities between D3 and ggplot2 (a graphing library for R) vocabularies are also evident in the code. This must be continued in the visualisations and VA artefacts created in the WO community.

6.4 A Task is a Chain of Events

An individual event identified semantically can provide a level of insight above standard interaction-log analytics. Taking this further means successfully identifying groups of actions and events as high-level tasks. The recording of user events already allows the replay of interactions so 'users can evaluate the ways they and others have explored it in the past' [15]. For Lu *et al.* [16] 'a linear, logical sequence of user actions constitutes an analytic trail' that can used to support analytic provenance, asynchronous collaboration, and reuse of analyses in VA applications.

The notion of meaningful structures in user behavior in VA can be traced back to the extraordinary work achieved with the VISAGE system [20]. Following this with a look into the *branching history of user operations* [21] Gotz and Zhou's empirical study in 2008 concluded two major points. First, that trails are 'chains' of activity 'lead to points of insight', and secondly that these patterns can be used (by 'smart' VA systems) to 'proactively assist users in completing their desired goals' [22]. To do this effectively means understanding how a set, or 'chain', of low-level user events relates to high-level tasks and goals, which will be addressed in future work at the University of Southampton. However, even if consistent and meaningful interaction data can be collected from VA applications on the Web a method of evaluation is still required. It is suggested in the following section that Economics is a subject that offers new insights into user behavior in situations of decision-making and choice for Web Observatories where the effect of data on the user can be tested and modified to increase efficiency and reduce errors in interpretation and subsequent action.

7. VISUAL FRAMING WITH ECONOMICS

When individuals are given a decision problem they create a mental model [23] referred to as a frame. This model used to solve the problem and includes information about the problem and the context [24]. By changing an element of the frame such as the written or visual context while the information remains the same, individuals have been shown to evaluate the same information differently to others and to be divided in their choice of response [25]. Framing in Economics has already been applied to the study of data visualisation [26] to examine the affect of *visual* framing. Looking further at the relationship between visualisation and Economics it is noted that visual analytics has been used to facilitate decision support for financial portfolio selection [27], [28], [29].

7.1 Testing Visual Framing with Economics

Historically the experimental method used to test framing in Economics involved presenting two groups of people the same information framed differently [30], [31]. However, methods have since evolved to recognise the need to examine the process of decision *making* and not only the decision maker [24]. A recent experimental method to test visual framing was conducted using Amazon's Mechanical Turk (AMT) [32], effectively mimicking the historic approach from Economics on a larger scale. An extensive review of problems encountered in data reliability, technical issues and timescale using AMT for visualisation research [33] is taken into consideration for the methodological approach given in the next section.

7.2 The Process Model

There are four complementary levels in the process, starting with the *interaction data* level at the bottom. The collection of interaction data is continuous in the process shown in Figure 2 (below), while observing users in a 'think aloud' session is used to identify event chains and the tasks they relate to. The users event chain is expected to be different compared to how it would be in an unobserved task, but the general order of actions and event types are expected to be fundamentally similar.

Event chains in the interaction logs are mapped to think aloud observations (see #4 in Figure 2), marking the beginning at end of tasks undertaken by the user and examining the chain of events within. Any aspect of visual framing that may affect the user's interpretation of the data should be manually identifiable from the data. Following successful identification is the formation of a hypothesis about visual framing, which is tested with either (or both) A/B test(s) and a controlled visual experiment (#6 & #7 in the process model) more akin to the historical approach [32], [34], [25]. What remains then is to explore the potential to identify framing effects from just only the event chains found in the interaction logs.

8. CONCLUSIONS

This paper proposes three principles to guide the future of interaction visualisations and visual analytics tools created by WOs. These principles emphasise the need be open about the objectives of an application, to share source code and for the collection and publication of interaction data so applications can be evaluated against stated objectives. The well-established JavaScript visualisation library D3.js is suggested as part of a framework alongside Analytics.js and a semantic taxonomy for event logging. Further to this it is proposed that Economics could be a suitable resource and perspective to develop new and innovative methods of examining user behaviour captured through observation and event logs. The use of Economic theory in this way is preliminary, but builds on existing research to outline a new application area.

Visualisation methods and tools used in Web Observatory projects do not align to the three principles proposed in this paper, regardless of the technology used to produce them. If these existing examples shared the objectives of the work, the code used to create it and data on the user behaviour then other WOs could benefit in their own work and improve upon the design of the original. Greater collaboration and development in the publication of visualisation work can facilitate better insights and inform decisions for WO users.

Standardisation can stifle our innovation and creativity when the boundaries are limiting, but developing the right standards can also liberate real-world applications. There are limitations to be considered with the implementation of the D3.js library and the use of SVG to visualise data, as it cannot cope well when rendering 'big data' on the front-end. Developers must rely instead on the majority of data processing to be done on the back-end, which is a well-established process in WOs.

There are JavaScript libraries that can support a greater capacity for visualising large data, such is the case with Three.js - a JavaScript library that makes use of WebGL to work with processing power of a user's graphics card and not just working memory from a Web browser. However, Three.js has not been used extensively in visualisations or VA tools.

9. FUTURE WORK

Publishing Methods of Event Chain Capture: The current work undertaken with tracking interaction events is currently experimental. These must be formalised and made available to the WO community (and others) for feedback and further development across multiple domains.

Events into Tasks: Low-level events (clicks) can be attributed to semantic actions (filter, undo, delete). Chains of actions can be viewed as tasks that might be generalizable across Observatories. It is important therefore to implement event tracking across multiple WOs to support analysis of event types and sets of events to explore commonality in users' tasks for different visualisations.

Insights and Decisions: Further research is required to explore the potential to reliably identify the exact points of insight and decisions made by a user from interaction logs alone.

Universal analytics: Analytics.js has recently be adopted within Google Analytics as part of the *Universal Analytics* framework. The concept is to capture not only online events, but also offline events and actions by the same user. Further research could explore the potential for recording actions and events across multiple services both online and offline to trace chains of user events in VA in a more holistic organisational context.

Framing Effects: The proposed methodology for identifying framing effects in visualisation work and the application of Economic theories to increase the efficiency of decision making will be undertaken across 2014 at the University of Southampton in a WO context and externally with two industry partners.

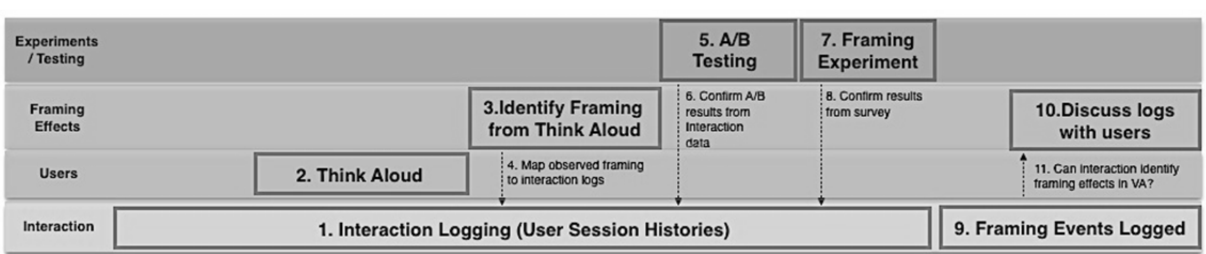

Figure 2: Process Model for Evaluating VA for Framing Effects

10. REFERENCES

[1] T. Tiropanis, W. Hall, N. Shadbolt, D. De Roure, N. Contractor, and J. Hendler, "The Web Science Observatory," *IEEE Intelligent. Systems*, vol. 28, no. 2, pp. 100–104, 2013.

[2] M. J. K. Gloria, J. S. Luciano, and D. L. McGuinness, "Building Semantically-Enriched Web Observatories," presented at the Web Science Workshop Building Web Observatories, Web Science 2013, Paris.

[3] K. A. Cook and J. J. Thomas, "Illuminating the path: The research and development agenda for visual analytics," Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2005.

[4] J. Hullman and N. Diakopoulos, "Visualization Rhetoric: Framing Effects in Narrative Visualization," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2231–2240, 2011.

[5] 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.

[6] T. M. Rhyne and M. Chen, "A Snapshot of Current Trends in Visualization," *Computing Now*, vol. 7, Jan. 2014.

[7] R. Kosara and J. Mackinlay, "Storytelling: The Next Step for Visualization," *Computer*, vol. 46, no. 5, pp. 44–50, 2013.

[8] R. Eccles, T. Kapler, R. Harper, and W. Wright, "Stories in GeoTime," presented at the VAST '07: Proceedings of the 2007 IEEE Symposium on Visual Analytics Science and Technology, 2007.

[9] C. Baber, D. Andrews, T. Duffy, and R. McMaster, "Sensemaking as Narrative: Visualization for Collaboration," *Proc. VAW2011, University London College*, pp. 7–11, 2011.

[10] G. Dove and S. Jones, "Narrative Visualization: Sharing Insights into Complex Data," 2012.

[11] J. Bertin, "Semiology of graphics: diagrams, networks, maps," *University of Wisconsin Press*, 1983.

[12] L. Wilkinson, *The Grammar of Graphics*. Springer, 2005.

[13] M. Bostock, V. Ogievetsky, and J. Heer, "D³ Data-Driven Documents," *IEEE Transactions on Visualization and Computer Graphics*, vol. 17, no. 12, pp. 2301–2309, Dec. 2011.

[14] H. Wickham, "A Layered Grammar of Graphics," *Journal of Computational and Graphical Statistics*, vol. 19, no. 1, pp. 3–28, Jan. 2010.

[15] H. Goodell, C.-H. Chiang, C. Kelleher, A. Baumann, and G. Grinstein, "Collecting and Harnessing Rich Session Histories," presented at the 10th International conference on Information Visualization, 2006. IV 2006. pp. 117–123.

[16] J. Lu, Z. Wen, S. Pan, and J. Lai, "Analytic Trails: Supporting Provenance, Collaboration, and Reuse for Visual Data Analysis by Business Users," vol. 6949, no. 22, pp. 256–273, 2011.

[17] J. Heer, J. D. Mackinlay, C. Stolte, and M. Agrawala, "Graphical Histories for Visualization: Supporting Analysis, Communication, and Evaluation," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 6, pp. 1189–1196, 2008.

[18] D. Gotz, Z. When, J. Lu, P. Kissa, N. Cao, W. H. Qian, S. X. Liu, and M. X. Zhou, "HARVEST: an intelligent visual analytic tool for the masses," *IVITA '10*, pp. 1–4, Feb. 2010.

[19] D. Gotz and M. X. Zhou, "Characterizing users' visual analytic activity for insight provenance," *InfoVis*, vol. 8, no. 1, pp. 42–55, 2009.

[20] M. Derthick, J. Kolojejchick, and S. F. Roth, "An interactive visualization environment for data exploration," *Proceedings of knowledge discovery in databases*, pp. 2–9, 1997.

[21] M. Derthick and S. F. Roth, "Enhancing data exploration with a branching history of user operations," *Knowledge-Based Systems*, vol. 14, no. 1, pp. 65–74, 2001.

[22] D. Gotz and M. X. Zhou, "An empirical study of user interaction behavior during visual analysis," 2008.

[23] P. N. Johnson-Laird, *Mental Models*. Harvard University Press, 1983.

[24] D. J. Koehler and N. Harvey, *Blackwell Handbook of Judgment and Decision Making*. John Wiley & Sons, 2008.

[25] A. Tversky and D. Kahneman, "The framing of decisions and the psychology of choice," *Science*, vol. 211, no. 4481, pp. 453–458, 1981.

[26] E. K. Choe, "Visual Framing: Nudging Toward Health Behavior Change," Jan. 2012.

[27] A. Savikhin, R. Maciejewski, and D. S. Ebert, "Applied Visual Analytics for Economic Decision-Making," *IEEE Symposium on Visual Analytics Science and Technology, 2008 VAST '08*, pp. 107–114, Jun. 2008.

[28] S. Rudolph, A. Savikhin, and D. S. Ebert, "FinVis: Applied visual analytics for personal financial planning," *IEEE Symposium on Visual Analytics Science and Technology, VAST 2009*, pp. 195–202, 2009.

[29] A. Savikhin, H. C. Lam, B. Fisher, and D. S. Ebert, "An Experimental Study of Financial Portfolio Selection with Visual Analytics for Decision Support," *44th Hawaii International Conference on System Sciences (HICSS), 2011*, pp. 1–10, 2011.

[30] A. Tversky and D. Kahneman, "Judgment under uncertainty: Heuristics and biases," *Science*, vol. 185, no. 4157, pp. 1124–1131, 1974.

[31] D. Kahneman and A. Tversky, "Prospect theory: An analysis of decision under risk," *Econometrica: Journal of the Econometric Society*, pp. 263–291, 1979.

[32] E. K. Choe, J. Jung, B. Lee, and K. Fisher, "Nudging People Away From Privacy-Invasive Mobile Apps Through Visual Framing," *Proceedings of Interact 2013*, Sep. 2013.

[33] R. Kosara and C. Ziemkiewicz, "Do Mechanical Turks dream of square pie charts?," Proceedings of the 3rd BELIV"10 Workshop: BEyond time and errors: novel evaLuation methods for Information Visualization, New York, New York, USA, 2010, pp. 63–70.

[34] A. Tversky and D. Kahneman, "Rational choice and the framing of decisions," *Journal of business*, pp. 251–278, 1986.