

Challenges in Logging Interactive Visualizations and Visualizing Interaction Logs

Romain Vuillemot*
École Centrale de Lyon

Jeremy Boy†
UN Global Pulse

Aurélien Tabard‡
Université Lyon 1

Charles Perin§
University of Calgary

Jean-Daniel Fekete¶
INRIA

ABSTRACT

A growing number of visualization tools are now publicly released on the Web. While this has many benefits, such as reaching more users without any installation time or procedure, it is often unclear how those tools are being used. The most common method to remotely observe usage is remote logging through a web server. Analyzing recorded logs has already been successful to improve the usability of tools, assess the performance of users and even to enrich the user interface with histories or logs visualizations. However, from our own practice of recording and analyzing logs, we have found a lack of methodology to support this process and use the results consistently. Our goal is to raise awareness of the potential of logging to improve visualization tools and their evaluation, as well as paving the way for a long term research agenda on the use of logs in Information visualization (Infovis).

1 CONTEXT AND MOTIVATION

Logging is a mechanism for automatically capturing the behavior of a program or of a user. It is usually invisible, non-obtrusive, and can be set up remotely for long periods of time [1]. Logging can be particularly useful for information visualization research, as it can serve to debug a visualization, to test its usability, or to evaluate a user's behavior while interacting with it. Although logs usually aim at capturing system events resulting from e. g., user interactions, they can also record other valuable information like a visualization's state at specific moments of a user-session—typically what data is being used, what window layout shows up on the user's UI, etc. These recordings can be set up explicitly (e. g., using a log tracker), or can be indirectly generated using web server logs or proxies [7].

We, the authors of this article, have used logging mechanisms for almost a decade now. We have mainly conducted system evaluation and user behavior analysis using logs, but we have also started to explore novel ways of visualizing logs themselves to facilitate their analysis. We have developed a variety of tools to track user-activity, which we have deployed in various online visualizations and tools, some of which have reached great masses of users (+100 000). This experience has led us to appreciate the need for developing structured ways of making sense of logs, and is what drives the questions and discussions we raise in this proposal. As so far we have failed to find proper documentation on best practices in this area in the Infovis literature, we hereby intend to encourage a community effort to share best practices, resources, and outline promising directions for future research and developments.

Logging is difficult because it provides only a partial view of users' behavior. This is the trade-off to accept in order to remotely track them in their own settings, such as computer, desktop and

real collaboration environment. Logs cannot capture everything as most users are distracted by other applications, emails and social network notifications, and coffee breaks with colleagues. Recording every single event the user generates is also not reasonable as the volume of logs will be too important and the signal drowned in the noise especially as we have said since users are often multi-tasking. Finally, if logging spreads over long periods of times, Infovis software may have been upgraded during the period, the user may have worked offline and her environment has changed (new input or output device such as mouse or screen). For all those reasons, logging is a non-trivial problem but have a huge potential if done properly.

As far as we know, there hasn't been any attempt of tackling Infovis logging research and technical questions head-on. The workshop BELIV (Beyond Time And Errors: Novel Evaluation Methods For Visualization) has been running every 2 years for 10 years now, and a series of articles [1, 6, 4] investigate logging as an evaluation mechanism. Over the same period of time, the VAST challenge also released many datasets related to logs. For instance in 2011 the challenge contained firewall, IDS (Intrusion Detection System) and syslog (System) logs. Last year's workshop on Personal Infovis at IEEE Vis gathered researchers analyzing and visualizing human behavior data. Other research communities have organized workshops focused on logging user activity for specific contexts, such as such as WWW [3] for Web browsing. None of those workshops address the characteristics of Infovis interaction techniques and evaluation procedures.

2 RESEARCH AGENDA PROPOSAL

Our agenda focuses primarily on five issues associated with logging: 1) defining logging format(s); 2) reporting and analyzing logs; 3) setting up logging infrastructures; 4) reflecting on the legal issues and necessary ethical practices associated with logging; and 5) applications related to logs, such as their visual representations. In the following subsections, we briefly develop on each of these issues, and we propose a series of open questions intended as 'food for thought' for future research directions.

2.1 A Standard Logging Format

The first step when setting up a logging process is to ask what should be recorded, when it should be recorded, and how (by the web server, by the application itself, etc.). For example, even a simple and ubiquitous interaction, like a mouse dragging, requires carefully considerations as it can generate a lot of noisy events resulting in very large and thus difficult to interpret log files.

- How to record low-level interactions (mouse moves, keystroke, ...) and data-intensive interactions (dynamic queries, brushing and linking, ...) efficiently?
- How to track multiple and coordinated views? How to track the view the user currently focuses on?
- What is the scope of the context that should be recorded beyond user's interaction? Desktop UI configuration? Computer and office setup?
- How to record collaborative and multi-device activities?

*e-mail: romain.vuillemot@gmail.com

†e-mail: myjyby@gmail.com

‡e-mail: aurelien.tabard@univ-lyon1.fr

§e-mail: charles.perin@ucalgary.ca

¶e-mail: jean-daniel.fekete@inria.fr

- How to reduce the size of data intensive interactions? Should there be a low frequency / interaction sampling, filtering and/or aggregation to shrink log files? Should there be any buffering strategy?
- Is the Common Log Format (IP, User ID, Timestamp, etc.) generated by web servers, expressive enough to be the standard for Infovis? What are the related and upcoming standards (W3C, others)? Should Infovis define its own log format?

2.2 Logs Reporting and Analysis

Logs reporting in academic papers varies with high discrepancies. In PivotSlice [14] authors report "interaction logs were recorded by the software". While in À Table [11], authors provide a detailed "Participation Logs" analysis of the 185636 interaction from 648 visitors. This raises the need to improve logs analysis reporting to allow sound conclusions, and reproducibility of the evaluation.

- What relevance have vanity metrics (# users, # visits) to assess the success of an Infovis tool/technique?
- What should be the standard procedure or the best practices in logs reporting, for applications ranging from usability testing to evaluation?
- How to improve the reproducibility of research results and interoperability between logging tools and techniques?
- What are the specifics of logs for controlled experiments versus in the wild ones?
- How do user behavior framework like the HEART framework translate into logs? (and vice versa)

2.3 Logging Infrastructure

As we have mentioned earlier, a series of tools log users by default (e.g. proxies, web servers). However, from the authors' practical experience, it is oftentimes necessary to build its own tools for the sake of control over the logging format and flexibility in types of events to tracks.

- What is a simple and affordable setup for logging in Infovis?
- How to deal with offline tracking, synchronization? How to merge collected logs with other data sources, e.g. to clean, validate or enrich them with more contextual information?
- How updating an Infovis technique impacts previously collected/legacy logs?
- How existing APIs (Google Analytics, KissMetrics) can be used to track Infovis techniques? And perform tests such as A/B testing, perform cohort analysis, and real time monitoring?
- Beyond remote servers: what logging device or tracker can be used for logging? Can logging be manual and self-reported by users, instead of automated?

2.4 Legal and Ethical questions

As log collection and analysis is related to behavioral research involving humans, it requires approval from researchers' employer.

- How to make logging comply with IRB applications? How those applications shape the logging collection and evaluation procedure?
- What are the disclaimers best practices to notify users of logging activity?
- What would be the design of logging respecting privacy (e.g. logging that doesn't enable to reveal people's identity)?

2.5 Application Related to Logging

Finally, we think that agreeing upon a logging format and infrastructure, would have spillovers such as data interoperability and allow more applications building upon logs. Letting users visualize logs, whether it is their own or others, is a rich and promising area to

identify patterns [10, 12], insights [13] of large logs collections [9]. Logs may also enrich the user experience with enhanced history navigation [2], browsing [5] and monitoring [8]. More Infovis and Visual Analytics application already make sense of logs and further research is need to tackle challenges with the growing complexity of data types, user tasks, and the need for scalable solutions as logs volume increases exponentially.

REFERENCES

- [1] Ben Shneiderman and C. Plaisant. Strategies for evaluating information visualization tools: multi-dimensional in-depth long-term case studies. In *Proceedings of the 2006 AVI workshop on Beyond time and errors: novel evaluation methods for information visualization*, pages 1–7. ACM, 2006.
- [2] F. Chevalier, P. Dragicevic, A. Bezerianos, and J.-D. Fekete. Using Text Animated Transitions to Support Navigation in Document Histories. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 683–692, New York, NY, USA, 2010. ACM.
- [3] A. Edmonds, K. Hawkey, M. Kellar, and D. Turnbull. Logging traces of Web activity: The mechanics of data collection. In *Workshop presented at the 15th International Conference on World Wide Web (WWW 2006)*, 2006.
- [4] J. Gerken, P. Bak, C. Jetter, D. Klinkhammer, and H. Reiterer. How to use interaction logs effectively for usability evaluation. 2008.
- [5] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala. Graphical Histories for Visualization: Supporting Analysis, Communication, and Evaluation. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1189–1196, Nov. 2008.
- [6] N. Henry, N. Elmqvist, and J.-D. Fekete. A Methodological Note on Setting-up Logging and Replay Mechanisms in InfoVis Systems. In *BELIV'08*, Florence, Iran, Apr. 2008. ACM.
- [7] J. I. Hong, J. Heer, S. Waterson, and J. A. Landay. WebQuilt: A Proxy-based Approach to Remote Web Usability Testing. *ACM Transactions on Information Systems*, 19(3):263–285, July 2001.
- [8] S. Huron, R. Vuillemot, and J. D. Fekete. Visual Sedimentation. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2446–2455, Dec. 2013.
- [9] K. E. Isaacs, P. T. Bremer, I. Jusufi, T. Gamblin, A. Bhatlele, M. Schulz, and B. Hamann. Combing the Communication Hairball: Visualizing Parallel Execution Traces using Logical Time. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2349–2358, Dec. 2014.
- [10] A. Makanju, S. Brooks, A. N. Zincir-Heywood, and E. E. Milios. LogView: Visualizing Event Log Clusters. In *Sixth Annual Conference on Privacy, Security and Trust, 2008. PST '08*, pages 99–108, Oct. 2008.
- [11] C. Perin, R. Vuillemot, and J.-D. Fekete. A Table!: Improving Temporal Navigation in Soccer Ranking Tables. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems*, CHI '14, pages 887–896, New York, NY, USA, 2014. ACM.
- [12] R. Vuillemot and C. Perin. Sports Tournament Predictions Using Direct Manipulation. *IEEE Computer Graphics and Applications*, 36(5):62–71, Sept. 2016.
- [13] K. Wongsuphasawat and J. Lin. Using visualizations to monitor changes and harvest insights from a global-scale logging infrastructure at Twitter. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pages 113–122, Oct. 2014.
- [14] J. Zhao, C. Collins, F. Chevalier, and R. Balakrishnan. Interactive Exploration of Implicit and Explicit Relations in Faceted Datasets. *IEEE Transactions on Visualization and Computer Graphics*, 19(12):2080–2089, Dec. 2013.