

# Explainable Visualization of Collaborative Vandal Behaviors in Wikipedia

Siva Sandeep Subramanian\*

Parab Pushparaj†

Zerong Liu‡

Aidong Lu§

Department of Computer Science  
University of North Carolina at Charlotte

## ABSTRACT

Online social networks are prone to be targeted by various frauds and attacks, which are difficult to detect due to their complexity and variations. The challenge is to make sense of all information with suitable exploration tools for different groups of users. This project focuses on an explainable visualization approach to study collaborative behaviors of vandal users on Wikipedia. Our approach creates visualization with commonly used techniques from cartography and statistical graphics that are familiar to the general public for effectiveness and explainability. We build a large-scale visualization system which supports an illustrative interface with multiple data query, filtering, analysis, and interactive exploration functions. Examples and case studies are provided to demonstrate that our approach can be used effectively for a set of Wikipedia behavior analysis tasks.

**Index Terms:** Human-centered computing—Visualization—Visualization techniques; Human-centered computing—Visualization—Visualization design and evaluation methods

## 1 INTRODUCTION

Online social networks, such as Wikipedia, have become an irreplaceable component of our every day lives. They are however very likely to be targeted by various frauds and attacks, especially as Wikipedia can be edited by anyone in the world. Successful vandal detection is crucial to maintain the normal operations and it is extremely challenging to achieve due to the complexity of heterogeneous data and variances of attacks. Many detection approaches have been developed to support administrator tasks and they often target at making sense of heterogeneous data, identifying vandal users at early stages, or providing effective interactive analysis functions [20, 29, 33]; however the tools for the general public to understand such malicious attacks are lacking.

Explainable visualization is designed to address this issue by presenting data in visual formats that can be easily understood by users. It is useful for applications like Wikipedia which involves of the general public as the main users. Previous visualization techniques have demonstrated a variety of interactive analysis functions for many applications. While complex visual interfaces can match the needs of challenging applications, the potential drawback is on the usability of visualization system, especially for users without professional experiences [19, 23]. Advanced interactive analysis functions can be achieved by integrating visual interface with suitable algorithms. For explainable visualization, we target at self-explanatory approaches by limiting the complexity of interface and adopting concepts that are already familiar to the public or can be easily understood.

In this project, we present an explainable visualization approach for an important and challenging task – analyzing various collabora-




Figure 1: Minard's flow map illustrates Napoleon's disastrous Russian campaign of 1812 in a succinct but effective style.

tive vandal behaviors in Wikipedia. We design our visualization with commonly used techniques from cartography and statistical graphics [1], which is a proven technique to improve the efficiency, accuracy, and memory of data for various users (including the general public without much training) [18, 22]. As shown in Figure 1, flow maps in cartography are a mix of maps and flow charts [27]. It successfully combines two components: the overlaying of the main event (invasion of Russia) through the flow on the map, and the associated factor (temperature). Our goal is to inherit the familiar concepts and maintain the succinct styles from example techniques for self-explanatory features.

Our visualization approach transforms the vandal exploration problem onto the flow map domain through synthesizing wiki pages as semantic maps and temporal user behaviors as line charts and sankey diagrams overlaid on top. To support interactive analysis for large-scale data, we have developed a system with the support of multiple data query, filtering, semantic analysis, and interactive exploration functions. The metaphor of flow map is used to capture important statistical features of users vs. pages relationships, and integrated with a temporal visualization to represent collaborative editing patterns intuitively. We provide examples of interactive exploration of various collaborative and vandal behaviors across different time durations. Our system demonstrates how an explainable interface can be designed to tackle large-scale and complex data analysis applications. Our approach of creating explainable interfaces through visual metaphors can be extended other temporal and relational data tasks.

## 2 RELATED WORK

We describe the related work from the aspects of visual metaphor for explainable visualization and wikipedia vandal detection.

### 2.1 Visual Metaphor for Explainable Visualization

Visual metaphor is one important method of visualization design [18, 22]. As the choices of metaphors are often familiar objects or concepts in real-life, these approaches are shown to be effective for a variety of users on different applications. Here we list several work on the topics of space-time sense-making and anomaly behaviors which are most related to this paper. Spaces and map-like visualizations have been widely used to visualize topics from common topics among multiple sources [21], anomalous information spreading on social media [34], text stream [11], diffusion patterns

\*e-mail: ssubra18@uncc.edu

†e-mail: pparab1@uncc.edu

‡e-mail: zliu17@uncc.edu

§e-mail: aidong.lu@uncc.edu

and user behavior in social media [5]. In this work, our visualization design simulates not only the appearance of metaphor, but also the succinct style for explainable interface.

Specifically, we use flow chart and map in our approach for visualizing temporal and spatial patterns in wikipedia. Flow maps [27] are important techniques in cartography and geospatial visualization. Maps are intuitive to show geospatial data, which often contains the location information directly. The map as a metaphor has also been used to visualize relationships beyond geospatial datasets [3, 10, 16]. Different from the previous approaches, we try to find 2D spaces embedded with semantic relationships for explainability.

## 2.2 Wikipedia Vandal Detection and Visualization

A number of vandal detection methods have been developed for Wikipedia data [25]. For example, Kumar et al. [20] used three machine learning models for early detection of vandal users. Their user behavior identification is based on statistical analysis performed on already known vandal and benign users.

There are also previous work of visual analytics of social network data [6], including visualizations of page ranks [32] and behaviors of participants [3]. Our work is also related to the topic of visualization for network security [30]. For example, network host behavior is visualized through positional changes in a two dimensional space using a force-directed graph layout algorithm [2]. Suh et al. [31] proposed revert graph where a user is denoted as a node and a revert relationship as a link. Hecking and Hoppe [13] proposed a node-link approach for visualizing dynamically evolving edits collaboration based on the article content. Flock and Acosta [9] presented a system that combines various visualization techniques to facilitated the understanding of intra-article disagreement between editors. Holloway et al. [14] proposed a system to analyze and visualize the semantic coverage of Wikipedia and its authors. Pang and Biuk-Aghai [26] introduced map-like Wiki Visualization to highlight the semantics coverage of Wikipedia. Chevalier et al. [7] integrated five visual indicators to Wikipedia layout in order to highlight the maturity and quality assessment of an article for the readers. Borra et al. [4] investigated the analysis and visualization of controversies in Wikipedia articles. Different from previous methods, this work focuses on an explainable visualization approach that is easy for the general public to understand.

## 3 EXPLAINABLE MAP VISUALIZATION

We use the UMDWikipedia dataset which contains a host of information about 34M pages and 23M users and their behaviors during a year long duration [20]. This section describes the design goals and our approach for visualizing temporal and collaborative behavior patterns with two components.

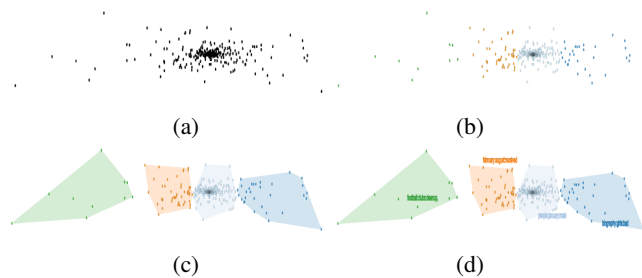


Figure 2: The procedure to generate Wiki map visualization: (a) data projection from editing records to visualize the relationships of pages as a map, (b, c) clustering pages as regions with similar topics, and (d) labeling for regions. The four labels from left to right are: football, received, people, and biography. They provide a quick hint of the regions on the map.

## 3.1 Design Goals and Overview

A complex Wikipedia visualization system can tackle multiple challenges, while our work focuses on the following design goals. (*DG1*) Facilitate intuitive understanding of multiple data elements, including a large number of users and pages, page tags, and editing records (reverted or not) over a time period. (*DG2*) Facilitate intelligent interaction to explore vandal user characteristics and collaborative editing behaviors. Different from the previous Wikipedia visualization approaches [15, 24], this work explores how to apply a classical cartography technique, flow map, with very simple and familiar concepts to visualize Wikipedia data (*DG1*). This design is flexible to support a number of analysis tasks (*DG2*), especially for studying the similarities and differences among behavior patterns of benign and vandal users.

Figure 1 shows the famous illustration of Napoleon’s invasion of Russia, which suggests the correlation of the low winter temperatures and the horrendous casualties in Napoleon’s army. Minard’s famous illustration shows the decreasing size of the Grande Armee as it marched to Moscow (brown line, from left to right) and back (black line, from right to left) with the size of the army equal to the width of the line. Temperature is plotted on the lower graph for the return journey, indicating the correlation between the attribute and retreat of the army. Similarly, our visual interface uses the map and line charts to simulate the technique of flow map. We describe the details of each component in the following respectively.

Our design of Wikipedia visualization includes two components: summary map visualization and behavior visualization with line charts. Figure 3 provides two examples of our design, which uses summary visualization to show Wikipedia pages and behavior visualization to show users. This design demonstrates the editing behaviors of selected users across time, features of editing patterns, as well as page categories. For example, we can easily identify users who have edited pages from a single or multiple categories; or users whose edits are reverted for different pages.

## 3.2 Generating Maps for Wiki Pages

The maps embed a rich set of information and therefore selected to visualize the most complex data component – 34M Wiki pages. Our approach uses map as a metaphor to visualize page relations from different measurements, time ranges, and user interaction. Specifically, we use the spatial dimensions to visualize page relationships and tags of grouped pages to label regions on the map with the following two steps: data projection for map foundation and topic clustering for labeled regions.

The main challenge of summary visualization is to provide meaningful context in the map like format for reflecting a large number of data. Different from the illustration examples, the datasets from Wikipedia do not have the intuitivity for users to understand. For wiki pages, the data also contains different aspects of information according to the relationships with a subset of users.

### 3.2.1 Data Projection for Map Foundation

The first step is to identify important data statistical features from a time duration. As the data contains mainly user-page editing patterns, we adopt the latent dimension generated by a dimension-reduction technique – singular vector decomposition (SVD) [17].

Specifically, for a user-page matrix  $M$  with  $m$  users and  $n$  pages, the SVD algorithm factorizes  $M$  into three matrices such that  $M = USV^T$ . It is common to truncate these matrices to yield  $U_k$ ,  $S_k$ , and  $V_k$ , in order to decrease the dimensionality of the vector space, and only leave the strongest effects by dropping dimensions with small singular values [8]. As the input to the user vs. page matrix, we assign value 1 for good edits and  $-1$  for reverted edits.

Each SVD dimension captures a combination of statistical features of data. To demonstrate the correlation between users and pages, it is ideal to select dimensions that contain features related to




Figure 3: Flow maps in cartography are a mix of maps and flow charts that “show the movement of objects from one location to another”. This example Visualization of collaborative editing patterns organize pages as a map from SVD and the behaviors of users as glyphs on the bottom bars for simulating the explainable features of flow maps.



Figure 4: Our behavior visualization for line charts: multiple features are shown as glyph patterns.

selected users. Therefore, we choose the top two dimensions as the default setting that maximizes the distributions of the edges which have been edited by selected users.

### 3.2.2 Topic Clustering for Labeled Regions

To better demonstrate the page topics on the map, we apply the  $k$ -means clustering algorithm to divide the space into multiple regions. The space can be the 2D subspace selected from the SVD or a larger range such as the first  $k$ -dimensional subspace of SVD. Since we only use topics as a suggestion of page contents which are very hard to cluster, we use the subspace from the first 2 dimensions of SVD for acceleration. Each cluster includes a set of closely distributed pages with similar features such as titles, categories, or edited by the same group of users. The regions are visualized with convex hulls automatically generated to provide the map background.

Once the map is generated, the user has an option to generate labels for the regions. The labels are based on the categories of pages included in the cluster. Our system analyzes the words in all categories, removes generic words like “of” and “the”, and returns the top three words with the highest occurrences. Better NLP algorithms can be applied to improve the results. This way, user can get a general idea of what kind of pages are present in the cluster rather than hovering over each page individually.

### 3.3 Line Charts for Temporal Behaviors

Line charts have been used to show sequential data for centuries [1]. Our design is multiple line stripes for representing several important features of users’ editing patterns that have been identified useful for vandal detection [20]. We have the top lane to represent one important attribute of pages, regular or meta pages. Since we focus on studying editing patterns, we have another three rows to represent the editing speed. Here we measure the time difference between two continuous edits of the same user on the same page as follows: very fast (i.e. re-edit less than 3 minutes), fast (i.e. re-edit within 3-15 minutes) and slow (i.e. re-edit longer than 15 minutes). Both the page attribute (meta or regular) and editing speed have been identified as useful features for vandal detection [20].

To connect the map and line charts as in the flow maps, we add curved stripes in between to visualize the editing patterns. The curves are generated by sankey diagrams [12], which are especially useful for depicting a flow from one set of values to another and show a many-to-many mapping between two domains. Figure 3 demonstrates combined examples.

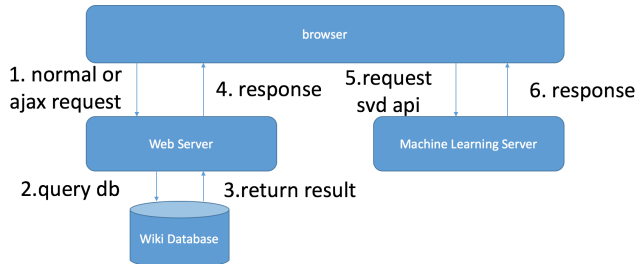


Figure 5: The architecture of our system. There are three core components, a web server, a database and a machine learning server.

## 4 VISUAL ANALYTICS SYSTEM FOR WIKI EXPLORATION

### 4.1 System Architecture and Implementation

As shown in Figure 5, our system contains three main components: a database by MySQL for storing the wiki data and providing high performance for querying, a web server for supporting web browser interaction and visualization, and a machine learning server for running computationally expensive algorithms. Specifically, **the web server** is set up to provide the visualization interface. In this component, we choose LAMP (Linux+Apache Server+MySQL+PHP) as our framework for its light weight and high development efficiency. We design our application base on the Model-View-Controller architecture (MVC) [28] and separate our static resources, include Html, CSS, JavaScript, with our PHP script. Finally, we run it on the web development platform known as WampServer. **The machine learning server** is set up using Flask, a microweb-framework of python. We provide an API for truncated SVD base on Scikit Library.

Overall, the web server is the heaviest component in our system, since it provides rich interaction and visualization functions implemented by D3.js and responds to various requests from clients. We retrieve the input data for SVD algorithm and send request to call SVD API on machine learning server to process data and return result. Finally, we use D3.js to manipulate and visualize the data.

### 4.2 Interactive Exploration Functions

We add two additions to the system interface to support interactive exploration functions: a timeline for selecting durations and scatter plots for visualizing user statistics, both are commonly used statistical graphics existed for years [1]. This helps our system to match all data elements onto the simple interface (*DGI*) by providing an overview of all data elements, users or pages, in a time range based on user-page editing relationships.

For scatter plots, multiple features have been extracted for detecting vandal users from the rich set of editing records [20]. We choose the most significant factors, including consecutive re-edit frequency (fast or slow), meta or non-meta page, number of reverts, and number of edits. Interactive functions are provided to select a group of users

with a lasso drawing tool on any scatter plot. Similarly, brushing on the timeline triggers an event mechanism to re-generate the map visualization and update the scatter plots (DG2).

Apart from the three main visualization components, we have also added the following interactive exploration functions (DG2). We have provided an option for the users to select the dimensions from the SVD for the map. Several filtering functions are also provided to select users based on different combinations on the scatter plots, such as filtering users by selected pages in a given time range.

## 5 RESULTS AND CASE STUDIES

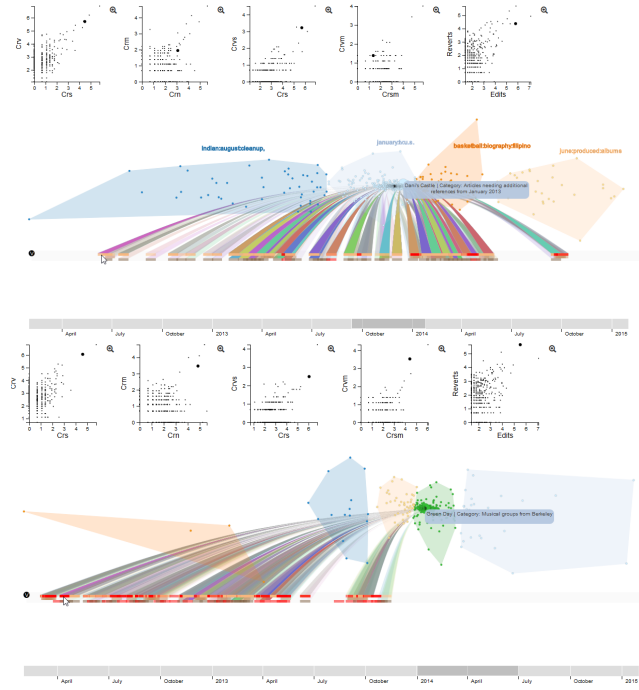


Figure 6: Examples of collaborative vandal behavior visualizations.

To start interactive exploration, we often filter the data by selecting a time range. We can then choose any one or multiple features with the scatterplots, such as outliers distributed on the side or popular behaviors on the dense regions. We can also fine-tune data with the provided filtering functions. Vandal behaviors can be identified by number of reverted edits shown in red and lack of edits on meta pages (missing the first row). The following demonstrates our approach with different behaviors and collaboration patterns.

Figure 6 (top row) shows an example of a mixed user behavior. During this period, multiple wiki pages related to a TV show “Dani’s castle” were edited, covering a variety of topics of the show. We can use the region labels or hover over pages to view their tags and find the map regions range from discussions of actors to the stories of the show from different seasons (left to right on the map). Among all users who edited the show pages, a user “Dragondevil22” had edited multiple pages but not a single meta-page, a strong indicator of vandal users. The visualization demonstrates how the user vandalized the show pages during this time period. Almost 95% of his edits were related to this show. His attack was centered around the time when the second season of TV show was scheduled to be aired.

Figure 6 (bottom row) shows that a wiki page titled “Green day” is edited several times from only two users over a long time. We can easily observe one of the users named “TreCoolGuy” vandalizing the page during two separate durations. Another example of vandal behavior is shown in Figure 3 (right), where two users edited several

pages related to the TV show “Bad girls club” in just two months. They targeted the pages related to the show at a time when the show was popular to mislead readers to wrong information.

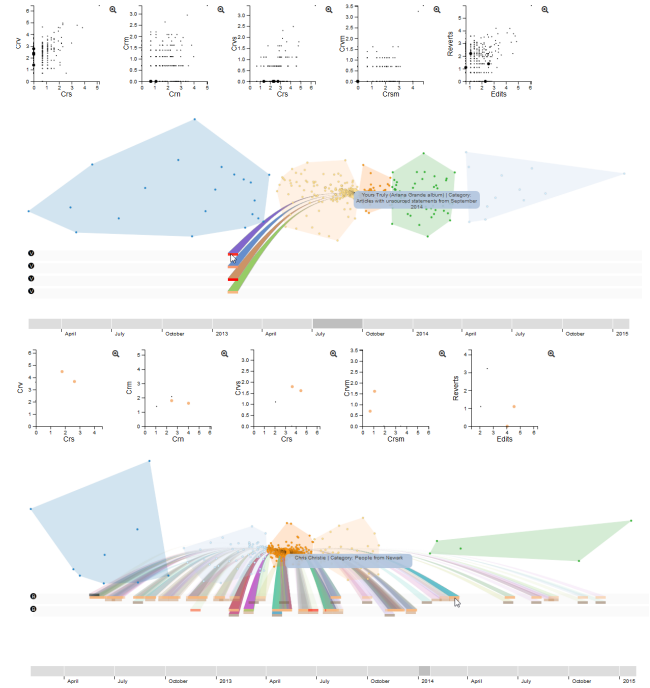


Figure 7: Examples of collaboration behaviors: (top) short-term group attack, (down) long-term positive collaboration.

The collaboration patterns can be positive for multiple benign or vandal users; or negative between benign and vandal users. Figure 7 (top) demonstrates a group attack where four vandal users edited the same page during a short period. The users are identified by filtering functions. They are hard to detect as their behavior statistics shown on the scatter plots are very different. Figure 7 (bottom) shows a positive collaboration where benign users contribute content to the same page alternatively during a long period. We use filter functions to show only users editing the same page, narrowing down the data to find collaborations.

## 6 CONCLUSION AND FUTURE WORK

This work presents an explainable visualization approach for exploring collaborative vandal behaviors in the Wikipedia editing records. Toward the goal of intuitive interfaces to improve the system usability for the general public, our design simulates flow maps and several commonly used statistical graphics techniques and integrates components to extract useful information from the data. The results demonstrate that our simple interface can effectively support a set of interactive visualization and exploration functions for studying various vandal collaboration patterns.

Our future work includes evaluating explainable interface for visualization of vandal behaviors to different groups of users, such as security experts and the general public. We will improve the system with scalable functions to visualize more users. We are also interested in exploring other design options for similar applications that can significantly benefit from explainable interfaces.

## ACKNOWLEDGMENTS

We thank for the detailed comments from our reviewers. This work was supported by the National Science Foundation under Grant No. 1564039.

## REFERENCES

- [1] J. R. Beniger and D. L. Robyn. Quantitative graphics in statistics: A brief history. *The American Statistician*, 32(9):111, 1978.
- [2] D. M. Best, S. Bohn, D. Love, A. Wynne, and W. A. Pike. Real-time visualization of network behaviors for situational awareness. In *Proceedings of the Seventh International Symposium on Visualization for Cyber Security, VizSec '10*, pp. 79–90. ACM, New York, NY, USA, 2010. doi: 10.1145/1850795.1850805
- [3] R. P. Biuk-Aghai, C.-I. Pang, and Y.-W. Si. Visualizing large-scale human collaboration in wikipedia. *Future Generation Computer Systems*, 31:120 – 133, 2014. Special Section: Advances in Computer Supported Collaboration: Systems and Technologies.
- [4] E. Borra, E. Weltevrede, P. Ciuccarelli, A. Kaltenbrunner, D. Laniado, G. Magni, M. Mauri, R. Rogers, T. Venturini, et al. Contropedia—the analysis and visualization of controversies in wikipedia articles. In *OpenSym*, pp. 34–1, 2014.
- [5] S. Chen, S. Chen, Z. Wang, J. Liang, Y. Wu, and X. Yuan. D-map+: Interactive visual analysis and exploration of ego-centric and event-centric information diffusion patterns in social media. *ACM Trans. Intell. Syst. Technol.*, 10(1):11:1–11:26, Nov. 2018. doi: 10.1145/3183347
- [6] S. Chen, L. Lin, and X. Yuan. Social media visual analytics. *Comput. Graph. Forum*, 36(3):563–587, June 2017. doi: 10.1111/cgf.13211
- [7] F. Chevalier, S. Huot, and J.-D. Fekete. Wikipediaviz: Conveying article quality for casual wikipedia readers. In *2010 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 49–56. IEEE, 2010.
- [8] M. D. Ekstrand, J. T. Riedl, and J. A. Konstan. Collaborative filtering recommender systems. *Found. Trends Hum.-Comput. Interact.*, 4(2):81–173, Feb. 2011.
- [9] F. Flöck and M. Acosta. whovis: Visualizing editor interactions and dynamics in collaborative writing over time. In *Proceedings of the 24th International Conference on World Wide Web*, pp. 191–194. ACM, 2015.
- [10] J. Fowler, M. Schneider, T. Johnson, C. Acedo, L. Lazos, P. Simonetto, and S. Kobourov. Imap: Visualizing network activity over internet maps. In *VizSec 2014 - Proceedings of the 11th Workshop on Visualization for Cyber Security*, pp. 80–87. Association for Computing Machinery, France, 2014. doi: 10.1145/2671491.2671501
- [11] E. R. Gansner, Y. Hu, and S. North. Interactive visualization of streaming text data with dynamic maps. *J. Graph Algorithms Appl.*, 17(4):515–540, 2013.
- [12] R. L. Harris. *Information Graphics: A Comprehensive Illustrated Reference*. Publisher: Oxford University Press (January 6, 2000), New York, 2000.
- [13] T. Hecking and H. U. Hoppe. A network based approach for the visualization and analysis of collaboratively edited texts. In *Proceedings of the Visual Approaches to Learning Analytics (VISLA15) workshop co-held with 5th International Learning analytics and knowledge conference*, 2015.
- [14] T. Holloway, M. Bozicevic, and K. Börner. Analyzing and visualizing the semantic coverage of wikipedia and its authors. *Complexity*, 12(3):30–40, 2007.
- [15] T. Holloway, M. Bozicevic, and K. Börner. Analyzing and visualizing the semantic coverage of wikipedia and its authors: Research articles. *Complex.*, 12(3):30–40, Jan. 2007. doi: 10.1002/cplx.v12:3
- [16] Y. Hu, S. G. Kobourov, and S. Veeramoni. Embedding, clustering and coloring for dynamic maps. In *2012 IEEE Pacific Visualization Symposium*, pp. 33–40. IEEE, Songdo, South Korea, Feb 2012. doi: 10.1109/PacificVis.2012.6183571
- [17] D. Kalman. A singularly valuable decomposition: the svd of a matrix. *The college mathematics journal*, 27(1):2–23, 1996.
- [18] N. W. Kim, E. Schweickart, Z. Liu, M. Dontcheva, W. Li, J. Popovic, and H. Pfister. Data-driven guides: Supporting expressive design for information graphics. *IEEE Transactions on Visualization and Computer Graphics*, PP(99):1–1, Jan 2017 2017.
- [19] M. Kirby, D. Laidlaw, R. S. Laramée, K. Mueller, H.-W. Shen, and A. Ynnerman. On the death of scientific visualization, 2016.
- [20] S. Kumar, F. Spezzano, and V. Subrahmanian. Vews: A wikipedia vandal early warning system. In *Proceedings of the 21th ACM SIGKDD international conference on Knowledge discovery and data mining, KDD '15*, pp. 607–616. ACM, Sydney, 2015.
- [21] S. Liu, X. Wang, J. Chen, J. Zhu, and B. Guo. Topicpanorama: A full picture of relevant topics. In *2014 IEEE Conference on Visual Analytics Science and Technology (VAST)*, pp. 183–192, Oct 2014. doi: 10.1109/VAST.2014.7042494
- [22] M. H. Loorak, C. Perin, C. Collins, and S. Carpendale. Exploring the possibilities of embedding heterogeneous data attributes in familiar visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):581–590, Jan. 2017. doi: 10.1109/TVCG.2016.2598586
- [23] B. Lorensen. On the death of visualization, 2004.
- [24] M. McKeon. Harnessing the information ecosystem with wiki-based visualization dashboards. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1081–1088, Nov 2009. doi: 10.1109/TVCG.2009.148
- [25] S. Moiss Mola-Velasco. Wikipedia vandalism detection. pp. 391–396, 01 2011. doi: 10.1145/1963192.1963349
- [26] C.-I. Pang and R. P. Biuk-Aghai. Wikipedia world map: method and application of map-like wiki visualization. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration*, pp. 124–133. ACM, 2011.
- [27] D. Phan, L. Xiao, R. Yeh, P. Hanrahan, and T. Winograd. Flow map layout. In *Proceedings of the Proceedings of the 2005 IEEE Symposium on Information Visualization, INFOVIS '05*, pp. 29–. IEEE Computer Society, Washington, DC, USA, 2005. doi: 10.1109/INFOVIS.2005.13
- [28] D.-P. Pop\* and A. Altar. Designing an mvc model for rapid web application development. In *Procedia Engineering*, pp. 1172–1179. ELSEVIER, Amsterdam, Netherlands, 2014.
- [29] M. Potthast, B. Stein, and R. Gerling. Automatic vandalism detection in wikipedia. In C. Macdonald, I. Ounis, V. Plachouras, I. Ruthven, and R. W. White, eds., *Advances in Information Retrieval*, pp. 663–668. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.
- [30] H. Shiravi, A. Shiravi, and A. A. Ghorbani. A survey of visualization systems for network security. *IEEE Transactions on Visualization and Computer Graphics*, 18(8):1313–1329, Aug 2012. doi: 10.1109/TVCG.2011.144
- [31] B. Suh, E. H. Chi, B. A. Pendleton, and A. Kittur. Us vs. them: Understanding social dynamics in wikipedia with revert graph visualizations. In *Visual Analytics Science and Technology, 2007. VAST 2007. IEEE Symposium on*, pp. 163–170. IEEE, 2007.
- [32] J. Xia, Y. Hou, Y. V. Chen, Z. C. Qian, D. S. Ebert, and W. Chen. Visualizing rank time series of wikipedia top-viewed pages. *IEEE Computer Graphics and Applications*, 37(2):42–53, 2017.
- [33] S. Yuan, P. Zheng, X. Wu, and Y. Xiang. Wikipedia vandal early detection: from user behavior to user embedding. In *PKDD*, pp. 832–846. ECML, Skopje Macedonia, 2017.
- [34] J. Zhao, N. Cao, Z. Wen, Y. Song, Y. Lin, and C. Collins. #fluxflow: Visual analysis of anomalous information spreading on social media. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1773–1782, Dec 2014. doi: 10.1109/TVCG.2014.2346922