Visualizing Article Similarities in Wikipedia

P. Riehmann, M. Potthast, H. Gruendl, J. Kiesel, D. Jürges, G. Castiglia, B. Ter-Akopyan and B. Froehlich

Bauhaus-Universität Weimar, Germany

Abstract

In this poster we present intermediate results regarding visual text analytics on Wikipedia. We implemented a visualization providing insight about similarities among Wikipedia articles in terms of structure as well as content. The presented data was gathered and processed via a pairwise comparison of all Wikipedia articles. Comparisons were appropriately pruned due to time and memory reasons when providing our in-memory database with the computed similarity values for visualization.

Categories and Subject Descriptors (according to ACM CCS): H.5.2 [Computer Graphics]: User Interfaces—Graphical user interfaces

1. Introduction and Related Work

Given its importance as an information source, and its free availability, Wikipedia has become a prominent subject for visualization. "See also" [LH] is a website that curates Wikipedia-related visualizations. Many of the visualizations found there capture various forms of article similarities: History Flow [VWD04] visualizes the editing process by revealing similarities of article revisions over time. "Seealsology" [Dm] uses the "See Also" links of Wikipedia pages to build a colored graph visualization that captures such user-generated article relations. "ClusterBall" [Har] visualizes Wikipedia's category pages as nodes-link diagram: inside a "ClusterBall" of categories arranged with links to a common parent category. Articles linked to the categories are arranged on the outer ring. The "Local Wikipedia Map" [Kre] presents another node-link diagram; starting with 3-5 articles, it recursively presents linked pages. Similarly, "WikiGalaxy" [Cora] provides an interactive three-dimensional visualization. Our approach, however, differs: instead of exploiting information already found on Wikipedia, we compute article similarities ourselves by comparing all (important) articles to all others using a text similarity measure.

| Similarity	Number of comparisons	Percentage	Size
[1.0, 0.9)	215,221	0.00033	1.35 MB
[0.9, 0.8)	1,125,858	0.00173	7.09 MB
[0.8, 0.7)	4,762,437	0.00732	29.98 MB
[0.7, 0.6)	16,598,675	0.0255	104.57 MB
[0.6, 0.5)	71,085,794	0.1093	1.17 GB
[0.5, 0.4)	329,595,171	0.5069	2.02 GB
[0.4, 0.3)	1,442,942,719	2.2	8.80 GB
[0.3, 0.2)	6,918,854,819	10.6	42.40 GB
[0.2, 0.1)	56,558,065,056	86.9	347.60 GB
> 0.1	65,344,422,345	100	402.14 GB

Table 1: *Distribution of Wikipedia article similarities under the vector space model.*

2. Data Acquisition

From a recent dump of the English Wikipedia, all articles were extracted (excluding meta pages such as user pages, help pages, and talk pages, etc.) and transpiled to plain text using Sweble [DR11]. Under the vector space model [MRS*08] with a *tf·idf* weighting scheme, stemming, and stop word removal, each text document d was represented as a high-dimensional vector **d**. Given two vectors $\mathbf{d}_1$ and $\mathbf{d}_2$, the semantic similarity of their original documents d_1 and d_2 can be estimated using the cosine similarity measure where a score of 0 means no similarity, and 1 maximum similarity. All articles' vectors combined form a term-document matrix which was distributed column-wise across four graphics cards (NVidia GTX 480) using CUDA. This way, computing the cosine similarity of a given article vector to all of Wikipedia took about 0.5 milliseconds. We compared every Wikipedia article in turn in decreasing order of length (where length served as an indicator of importance) to all others, recording the computed cosine similarities. To save space, we discarded all similarities of 0.1 or less as well as all similarities beyond the top 100,000 per article. To save time, we stopped computations once articles of length 50 terms or less were reached, since these articles are mostly superficial. Note in this respect, that the short articles still formed part of the aforementioned term-document matrix, so that all longer articles have been compared to the short ones, whereas the short ones have not been compared among themselves. Table 1 gives an overview of the similarity distribution; a total of 65 billion similarities have been computed in about 22 days, stored in 402 GB of disk space for a total 3.8 million articles disregarding the short ones.

3. Visualization and Interaction

Our visualization depicts articles as vertices and similarity relations as edges. Since an edge is only constituted if the similarity of the two adjacent articles' vertices exceeds a given threshold, the re-

Figure 1: Top left: *The radial layout places the subgraph containing the most articles in the center and orders the subgraphs with decreasing numbers of vertices more distally by growing the diameter level-wise.* **Top right**: *Additional edges appear within as well as across the subgraphs by lowering the global similarity threshold. Interaction capabilities like zooming, panning and dragging allow paying attention to particular subgraphs.* **Bottom left**: *A rectangular placement arranges the subgraphs (in local circle layout) row-wise from left to right starting with the one that contains the the most articles in the lower left corner.* **Bottom right**: *Local ring layout.*

sulting graph consists of a large number of isolated subgraphs (aka graph components). At first, the vertices of each subgraph are arranged locally circle-wise (Fig. 1 bottom left) or as ring metaphor by placing the articles that are in sum incident to more than 50% of the edges on an additional inner circle (Fig. 1 bottom right). Globally, the subgraphs are arranged by number of articles contained; either in rectangular or radial layout (see Figures 1 bottom left and 1 top left). Besides common interactions like zooming and panning most interactions are subgraph-centered such as dragging, dropping and highlighting subgraphs (Fig. 1 top right). Semantically zooming on particular subgraphs provides article titles and similarity values, whereas subgraph merging gathers all subgraphs that become connected to the current one when lowering the similarity threshold (Fig. 1 top right). Due to the increasing number of edges going along with lowering the similarity threshold, recreating the entire layout instantly for billions of edges was infeasible. Thus, a base drawing consisting of all edges exceeding the initial threshold is generated. Changing the threshold entails the appearance and disappearance of edges (Fig. 1 top right), however, some nodes may not emerge until the base layout is recreated.

Implementation was done completely in C++ and OpenGL by conveniently using gloost [WB], imgui [Corb] and FastDB [Kni] (an in-memory database for C++). We evaluated access times of custom-made, tailored data structures against FastDB, the former

being faster by a factor of 2. We still opted in favor of FastDB for its additional object-relational features, especially, since practical render performance was barely effected.

4. Discussion and Future Work

Providing threshold-adaptive layout; less rigid and static than before is crucial. A dynamic layout should at least manage slight changes of the threshold at run time by rearranging the node positions appropriately in local regions without distracting the larger picture and without loosing the user's mental map. For larger changes, however, we intend to create a preprocessed acceleration structure that provides initial positions or proximity information about clusters and articles for the entire threshold range.

Providing a proper subgraph labeling derived from the content of the articles assigned to a particular one will ease inspecting striking subgraphs. In our explorations we found that our similarity measures tends to focus more on structural similarity such as subgraphs of cities, countries, years, events, etc.; maybe due to the fact that Wikipedia articles discussing such entities tend to be structured homogeneously. In order to ameliorate the recognition of similarities in the actual running text, the data cleaning has to be improved as well as the similarity measures' weighting scheme. A comparison between different similarity measures that may be superimposed on top of each other is another interesting possibility.

References

[Cora] CORNEC O.: WikiGalaxy. http://wiki.polyfra.me. [Online; last accessed 07-April-2016]. 1

[Corb] CORNUT O.: imgui – bloat-free graphical user interface library for C++. https://github.com/ocornut/imgui. [Online; last accessed 04-April-2016]. 2

[Dm] DENSITYDESIGN, MÉDIALAB SCIENCES PO: Seealsology. http://tools.medialab.sciences-po.fr/seealsology. [Online; last accessed 07-April-2016]. 1

[DR11] DOHRN H., RIEHLE D.: Design and implementation of the sweble wikitext parser: unlocking the structured data of wikipedia. In *Proceedings of the 7th International Symposium on Wikis and Open Collaboration* (2011), ACM, pp. 72–81. 1

[Har] HARRISON C.: ClusterBall. http://www.chrisharrison.net/index.php/Visualizations/ClusterBall. [Online; last accessed 07-April-2016]. 1

[Kni] KNIZHNIK K.: FastDB. http://www.garret.ru/fastdb.html. [Online; last accessed 04-April-2016]. 2

[Kre] KREMPEL R.: Local Wikipedia Map. http://lwmap.uni-koeln.de. [Online; last accessed 07-April-2016]. 1

[LH] LAPORTE S., HASHEMI. M.: See also – A collection of our favorite visualizations built on Wikipedia data, curated by Hatnote. http://seealso.org. [Online; last accessed 04-April-2016]. 1

[MRS*08] MANNING C. D., RAGHAVAN P., SCHÜTZE H., ET AL.: *Introduction to information retrieval*, vol. 1. Cambridge university press Cambridge, 2008. 1

[VWD04] VIÉGAS F. B., WATTENBERG M., DAVE K.: Studying cooperation and conflict between authors with history flow visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2004), CHI '04, ACM, pp. 575–582. URL: http://doi.acm.org/10.1145/985692.985765, doi:10.1145/985692.985765. 1

[WB] WEISSIG F., BECK S.: gloost – handy wrappers for OpenGL and OpenCL entities. https://sourceforge.net/projects/gloost/. [Online; last accessed 04-April-2016]. 2