

TE_xVis: An Interactive Visual Tool to Explore Twitter Data

Shah Rukh Humayoun¹, Saman Ardalan¹, Ragaad AlTarawneh², and Achim Ebert¹

¹Computer Graphics and HCI Group, University of Kaiserslautern, Germany

²Information Technology Department, Mutah University, Jordan

Abstract

*Exploring tweets enables us understanding people's reaction and feedback regarding any particular event or product. Many tools have been developed to visualize Twitter data based on some criteria, e.g., keyword frequency or evolution of topics. Visualizing the relations between the keywords of the underlying Twitter data opens another window to analyze the people's reaction towards a particular event/product. Targeting this concern, our developed tool, called **TE_xVis** (**T**weets **E**xplorer and **V**isualizer), visualizes important keywords (e.g., hashtags, nouns, verbs) from the underlying tweets based on their frequency and shows the relations between them based on some criteria (e.g., the common tweets), using an extended Chord diagram. TE_xVis also visualizes the sentimental polarity for a better understanding of the keywords associated tweets. Further, the provided interaction, multi-level navigation, and filtering options help the users in better exploration of the underlying tweets. A user study with 16 participants shows a high acceptance towards the tool and our approach in general.*

Categories and Subject Descriptors (according to ACM CCS): H.5 [Information interfaces and presentation]: - [-]: —H.5.2[User Interface]: Graphical user interfaces, interaction styles

1. Introduction

Current social media platforms (e.g., Twitter, Facebook, etc.) are important mediums for people to express their feelings or to provide their feedback towards some currently happening or recently happened events. The exploration and the analysis of these feelings and feedbacks help us understanding the overall behavior or the reaction of the community about the underlying event. Nowadays, even companies are interested in analyzing their customers' feedback obtained from social media to better understand trends and the attitude towards their products. The 140-characters limitation on Twitter enforces users to write their tweets in a shortened and more precise way, which could be more useful for the exploration and the analysis compared to other social media platforms.

Many visualization tools have been proposed in the past to explore Twitter data, e.g., Nokia Internet Pulse [KLJ*12], Spark-Clouds [LRKC10], TopicFlow [MSH*13], Conference Monitor (CM) [SRBS12], TweetViz [SDM14], etc. Most of these tools visualize Twitter data based on either high number of keywords frequency (e.g., [KLJ*12]), a timeline of keywords (e.g., [SRBS12]), or evolution of topics (e.g., [MSH*13] or [CLT*11]). However, the relations between keywords based on some criteria (e.g., a *co-occurrence* relation that occurs between any two keywords if both are in the same tweet) opens another window to explore and analyze Twitter data. Visualizing such relations also helps users exploring and understanding people's feelings and feedback towards the related event(s) of the underlying tweets. Further, such visualization support can be useful for many application domains, e.g., exploring

users' feedback towards a product for marketing purpose or analyzing users' desired functionalities from the underlying tweets.

Targeting this concern, we developed a visual tool called **TE_xVis** (**T**weets **E**xplorer and **V**isualizer). It visualizes not only the keywords from the underlying tweets based on their frequency but also shows the relations, based on the selected criteria between these keywords (e.g., if two keywords occur in the same tweet then they have a direct relation). TE_xVis provides the resulting visualization through an extended Chord diagram that limits the extra chords cluttering, which might appear due to multiple relations' associating to the underlying nodes (i.e., keywords). Further, it visualizes the sentimental polarities of the keywords associated tweets, which helps understanding the associated tweets' subject. The provided multi-level navigation facility, the intuitive interaction and the filtering options help users in better exploring of the underlying tweets. We conducted a user study with 16 participants to see how they analyze people's feedback towards the "Brexit" event using the extracted tweets of the first ten days of July 2016 from Twitter. The participants showed high interest in the exploratory tasks and provided positive feedback towards the provided visual approach.

2. Related Work

One of the earlier work in visualizing text was done by Havre et al. [HHN00] in their famous ThemeRiver tool that visualizes the themes' variations over time for a collection of documents. Later, Don et al. [DZG*07] developed FeatureLens tool that visualizes

text collection at several levels of granularity to enable users exploring interesting text patterns based on length and frequency, while Cui et al. [CLT*11] in their TextFlow work used a semi-supervised clustering technique for the topic creation and represented the topic convergence and divergence using the river metaphor. As social media platforms (e.g., Twitter) provide large volume of real-time data; therefore, many researchers have focused on finding several techniques to visualize Twitter stream and hashtags. Few examples are: Kaye et al. [KLJ*12] developed a tool to visualize the evolution of Twitter discussions with a time series of stacked tag clouds. TopicFlow [MSH*13] [NMM11] visualizes the evolution of tweets to help understanding statistical topic modeling. Sopan et al. [SRBS12] provided an analysis of academic conferences hashtags over time to analyze conferences' trends. SparkClouds tool, developed by Lee et al. [LRKC10], integrates spark lines into the cloud tags to convey the trends between the multiple tag clouds. Stojanovski et al. [SDM14] developed the TweetViz tool that represents topic distribution in a set of tweets to allow users searching for any hashtag or keywords in the proposed interface. Thom et al. used ScatterBlogs to visualize geo-located Twitter messages [TBK*12] and to study crisis intelligence [TKE*15], e.g., using sentiment volumes of geo-located tweets. While Dork et al. [DGWC10] visualized Twitter data in three modes: topics over time through Topics Streams layout, people and their activity through spiral layout, and popularity of event photos through Image Cloud. Most of the previous work focused on keywords frequency or evolution of topics. None of them investigated the impact of different relations between the keywords (based on some relation criteria). Visualizing these relations could help the users to explore and understand people's feelings and feedback towards a particular event, product, or term.

3. The Enhanced Chord Diagram

The Chord diagram is a radial layout, initially popularized by The New York Times to show the relations between Genomes using the Circos package [KSB*09] [TNYT]. Radial (circular) layouts produce compact visualizations and use space efficiently, as they support a larger data domain on a squared area than rectangular or square layouts provide [KSB*09] [Krz]. They encourage the eye movement to proceed along the curved lines rather than a zig-zag fashion in a square or rectangular figure, which helps viewers to better understand and explore the underlying data [Krz]. Also, they can show the flow on relations between pairs more intuitively [KSB*09]. Due to these reasons, we selected Chord diagram rather than any rectangular/square layout (e.g., matrix layout).

In standard Chord diagrams, data elements (also called *nodes* or *arcs*) are arranged in a circular fashion and relations (also called *chords*) are drawn between the nodes. Mostly, chords associated to a node are mutually exclusive due to their association with different data in the underlying data-element; hence, no chords overlapping happens at the node side (see Fig. 1.a). However, sometimes chords associated to a node may not be mutually exclusive, which creates chords overlapping at the associated node side. This can cause a cluttering issue in the resulted visualization. Handling cluttering resulted from many-to-many relations in the visualization is a challenging task, which has been targeted by some researchers in the past for different visualization techniques (e.g., for matrix-base vi-

ualization [YDGM17] or for correlation map [ZMZM15]). However, as per our knowledge no one handled in the past the cluttering issue in Chord diagram resulting from non-mutually exclusive many-to-many relations.




Figure 1: (a) A standard Chord diagram taken from [Bos], (b) relations associated to nodes in TExVis extended Chord diagram.

To deal with the cluttering issue in our case, we propose an extension to the standard Chord diagram (see Fig. 1.b), which we use in our TExVis tool. In our case, the width of a node represents the weight value of the data element (e.g., frequency of a keyword), while the height of a node represents the number of associated chords or relations (e.g., co-occurrence relations between this keyword and other keywords in the underlying tweets). The width of a chord represents the weight value of this relation (e.g., frequency of co-occurrence relation between two keywords). In order to avoid extra chords cluttering, we order the chords based on their weights, such that the chord with highest weight value is placed at the bottom, while the second next one is placed above the previous one, and so on till the chord with the least weight value. However, the chord with the highest weight value starts a little below the upper/outer side of the specified node, the second one starts few pixels downwards, and so on (see Fig. 1.b). In this case, the node with the least weight value starts just above the inner boundary of the node. In this proposed solution, no chord is hidden behind the other chords; therefore, it provides more readable visualization.

4. TExVis: Tweets Explorer and Visualizer

Our developed TExVis tool visualizes the frequent keywords in the underlying tweets, the relations between these keywords using some particular criteria, and the sentimental polarities of the associated tweets. The web-based client side was developed using HTML, CSS, and JavaScript to provide the visual view as well as the interaction and filtering options, while the server side was developed in C#.Net to manage and process the data.

TExVis fetches the tweets from Twitter using the *Tweetinvi* and *Twitter REST* APIs. It fetches the data per user's request using a requesting loop. Then it separates interesting words (we call them *keywords* or *tokens*) based on hashtags, nouns and verbs using the *Apache OpenNLP*, which is a Natural Language Processing (NLP) library. If a non-noun or a non-verb hashtag is used frequently in the previous retrieved tweets then it is marked as a global hashtag, called *g-hashtag*, and the corresponding tokens are also separated in the current extracted Twitter data. Sentiment classifier is a term used to classify the text based on the contained sentimental polarities (e.g., *positive*, *negative*, or *neutral*) [LLC*10] [PLV02],

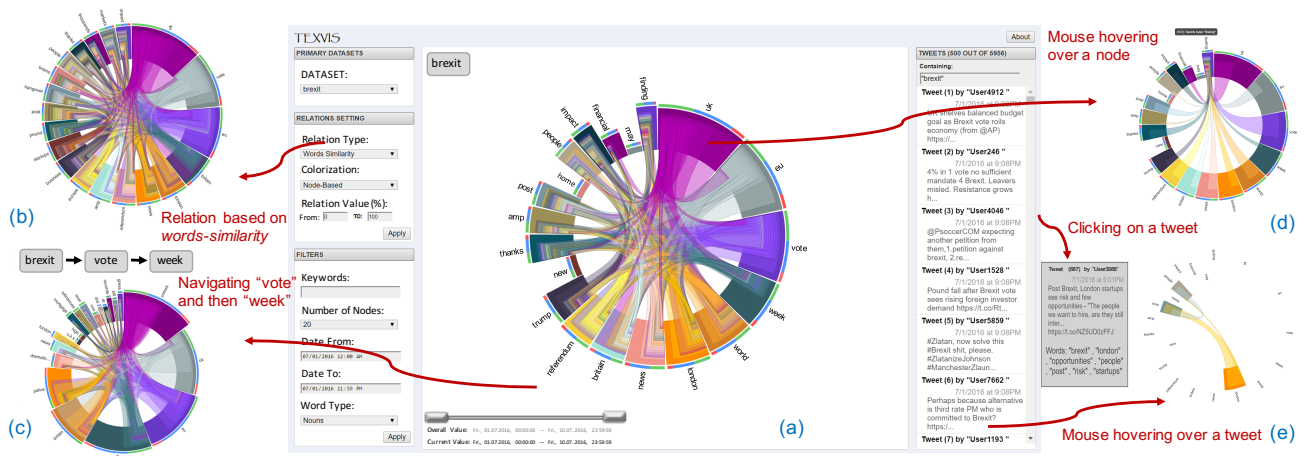


Figure 2: (a) TExVis tool with the left filter panel, the central visualization panel, and the right tweet panel, (b) the Chord diagram based on the words similarity relations, (c) navigating the keyword “vote” and then keyword “week”, (d) showing only the relations of the keyword “uk”, and (e) showing only the associated nodes and relations of a selected tweet.

which is useful for getting an overall opinion towards the text subject [PLV02]. TExVis uses Aylien.TextApi [ATAA] for each retrieved tweet, which returns the sentimental polarity value of the tweet, along with the polarity confidence value (between 0% to 100%) to show the confidence level of the stated polarity value.

Figure 2.a shows the overall view of TExVis client-side. It consists of three parts: the *filter panel* at the left side, the *visualization panel* in the center containing the extended Chord diagram built using the *d3.js* library, and the *tweet panel* at the right side. For a proof of concept, here we use the tweets extracted from Twitter using the keyword “brexit”, that was a popular hashtag in July 2016 referring to the UK referendum about its quitting from EU. Our extracted data consists of 41,199 tweets (with 56,701 distinct keywords in all categories) between July 01 to July 10, 2016. We assigned random ID numbers in order to make users’ IDs anonymous.

In order to explore and analyze the underlying Twitter data in visual form, the central part of the visualization panel is dedicated to our extended Chord diagram. In this diagram, the node’s width represents the keyword’s frequency while the node’s height is based on the number of associated relations to this node. The relations between the nodes are decided through the given criteria. Currently, TExVis supports relations based on two criteria: the *co-occurrence* criteria in which a relation between two nodes (keywords) occurs if both belong to the same tweet, and the chord (relation) width depends on the frequency of this co-occurrences in the underlying tweets (see Fig. 2.a); and the *words similarity* criteria in which the chord (relation) width indicates the value of the words similarity (range between 0 to 1) of the connected nodes (see Fig. 2.b) that is calculated in TExVis using the WordNet.Net [SC] library, which acts as a semantic dictionary for English lexical tokens. The colors of chords associated to a node ranges from darker to lighter from the wider to the thinner chords respectively with the same node color. Also, the chord color between two nodes depends on the larger associated node (e.g., a chord between “uk” and “vote” in Fig. 2.a has the same color as of the “uk” node). However, mouse

hovering over a particular node fades all other nodes’ relations and changes the colors of this node’s relations according to the opposite associated nodes (see Fig. 2.d). Mouse hovering over a particular node or a relation also brings a tooltip to show further details (e.g., no. of associated tweets, percentage, etc.). Further, we provide arcs outer-side of nodes to show the sentimental polarity (here, green represents *positive*, blue represents *neutral* and red represents *negative*) of all the associated tweets to each node. The color opacity shows the average confidence level for each polarity, i.e., 100% is the darkest color it gets. The length of each color in the arc represents the frequency of this polarity in the associated tweets.

Users can navigate the Chord diagram on-demand by selecting the navigation option from the menu bar (appears by clicking on a particular node), which results a new Chord diagram as a next level of details. For example, navigating “vote” and then “week” visualizes the data in the resulting Chord diagram related to only those tweets that have “brexit”, “vote”, and “week” together in them (see Fig. 2.c). This helps users exploring the tweets based on intersection of keywords. A navigation path is also shown at the top, which is used for going back to any previous level of details. TExVis also provides the navigation option through a chord; however, in this case user goes two levels down, e.g., Fig. 2.c navigation can be achieved by selecting the option from the chord between the “vote” and “week” nodes.

The right-side tweet panel shows tweets associated to the current Chord diagram. User can filter it to see tweets only related to a particular node or chord. Clicking a particular tweet opens a down space to show full text and all the associated keywords, while mouse hovering over a particular tweet fades all non-associated nodes and relations in the central Chord diagram (see Fig. 2.e).

TExVis provides a number of filtering options in the left side *filter panel*, for example: selecting the main extracted dataset based on the extracted keyword (e.g., “brexit”), selecting the relation type (e.g., *co-occurrence* or *words similarity*), navigating the visualization based on a given keyword, visualizing the Chord diagram

with the selected number of the most frequent nodes (e.g., Fig. 2 shows the Chord diagram with 20 high frequent keywords), filtering through date and time, and filtering based on the keyword-type (e.g., *nouns*, *verbs*, or *hashtags*). A time-bar is also provided at the bottom side of the central Chord diagram to filter the current view based on time interval.

5. The User Study

We conducted a user study with 16 participants (6 females), aged 24-33 ($M = 28.9$). The used dataset was the earlier described “brexit” dataset. The study goal was to check how users analyze and understand people feedback regarding the “brexit” event through exploring the provided visualization. We were also interested to see how different relationship types between the keywords influence on users in understanding and exploring the underlying tweets. Finally, we wanted to know users’ reaction towards our approach and the visualization. Based on this, we defined the study with a total of six tasks where the purpose of the first four tasks (i.e., getting information, navigating the visualization, data filtering, and understanding keywords’ relations) was to make the participants well aware of the TExVis visualization and the provided interaction and filtering options. The last two tasks were exploratory in nature where participants were asked to analyze people’s feedback and reaction towards the “brexit” event, first from the perspective of navigating to “eu” and then from the perspective of further navigating to “vote”. The study was done *between-subjects* manner, where eight of the randomly selected participants performed the experiment using the *co-occurrence* relation type while the other eight performed it using the *words similarity* relation type. At the end of the experiment, participants were asked to answer a 6 closed-ended questionnaire (using Likert-scale from 1 to 5) and their feedback in general. Each experiment lasted no more than 1 hour.

In the first four tasks, both groups performed approximately the same in terms of accuracy, i.e., 90.63%, 98.94%, 100%, and 100% by the co-occurrence group compared to 95.75%, 100%, 100%, and 100% by the words similarity group. However, the co-occurrence group’s performance time was overall better in these tasks, i.e., 208.75, 171.88, 214.38, and 98.75 seconds compared to 278.75, 185.63, 246.25, and 88.13 seconds by the words similarity group. In the last two exploratory-natured tasks, participants in both groups showed high interests and provided interesting feedback. Few examples of feedback from the co-occurrence group are: “*People are still in shock of the referendum event and do not have clear idea about the consequences*”, “*There is interest in signing a petition to protest the results*”, and “*Many people are linking the situation back to WW2*”. Two examples of the words similarity group feedback are: “*After navigating to vote, people are talking about ‘hate’, ‘fear’ and ‘crime’*” and “*It is interesting to see ‘nato’, ‘crimes’, ‘hate’, ‘police’, ‘fear’, ‘Youtube’, and ‘Brussels’. Something related to terrorism had happened at the same time?*”

We found out that the co-occurrence group highly relied on the co-occurrence relation for the analysis, especially where they were relating two events. Most of the participants in this group had a similar approach, i.e., they found the most occurring keyword pairs and then tried to find out a reason behind it. Few of them used navigation option to go further to their own decided next level of

details in order to focus on some specific topics. From the feedback, it is clear that they used features like sentiment polarity, co-occurrence relation, words frequency, and navigation to understand people’s behavior and to give the answer. Overall, they were able to find out various aspects of the “brexit” event. Some were even shocked to see the sentiment polarity of tweets related to particular keywords. Finding the relation of “brexit” with some other events (e.g., “WW2”) helped them understanding people’s reaction and feedback. While the second group had some other perspectives. Although they used the sentiment polarity, the navigation, and the words frequency; however, it seems that they hardly made any conclusion out of the similarity relation. It is because the words similarity relation shows that how much two words are semantically or lexically similar and this does not help the users to relate it to an event. However, the similarity relation can be useful in some other scenarios, e.g., clarifying two confusing words and knowing which ones in a certain context could be used by people (e.g., the case of ‘Geek’ vs. ‘Nerd’ by Settles [Set13]).

In closed-ended questions, most of participants from both groups either agreed or strongly agreed with the statements. In the case of intuitiveness of the visualization, 7 agreed and 8 strongly agreed. In the case of clearness and understandability, 7 agreed and 6 strongly agreed. When asked about visualization support for the analysis purpose, 5 agreed and 9 strongly agreed with it. However, 5 agreed and 1 strongly agreed in the co-occurrence group towards the usage of co-occurrence relation compared to 3 agreed and 2 strongly agreed the second group towards the usage of words similarity relation. There is some disagree of usage of word similarity relation that we can understand from the given scenario. This and the exploratory-natured tasks’ feedbacks indicate that different relation types suit different exploration scenarios, e.g., co-occurrence relation type suited more to our experiment scenario compared to the words similarity relation type. Therefore, it is recommended to first investigate whether a particular relation type could help in exploring and analyzing the underlying scenario. Most of the participants (5 agreed and 8 strongly agreed) favored the semantics polarity option as well as high positive feedback towards using the tool in future (4 agreed and 10 strongly agreed). In the open-ended feedback, few participants suggested to show initially only the important relations and then the remaining ones on demand. However, all of them provided high positive feedback about the tool, the visualization, and the approach. Few also asked to make the tool public.

6. Conclusion and Future Work

The presented TExVis tool enables the visual exploration of Twitter data through keywords frequency, keywords relation, and associated tweets’ sentimental polarities using the proposed extended Chord diagram. The conducted user study indicates that the used keywords’ relation type is useful when it supports the analysis of the underlying events or scenarios. In the future, we aim to provide additional facilities in TExVis, such as: selecting the relation based on other criteria, support of other social media platforms, showing only the important relations initially and further other relations only on demand, and visual comparison of people reaction about two or more events. We also plan to conduct detailed user studies on larger scale to generalize our findings.

References

- [ATAA] AYLIEN-TEXT-ANALYSIS-API. URL: <http://docs.aylien.com/docs/introduction.3>
- [Bos] BOSTOCK M.: Chord diagram. URL: <https://bl.ocks.org/mbostock/4062006.2>
- [CLT*11] CUI W., LIU S., TAN L., SHI C., SONG Y., GAO Z., QU H., TONG X.: Textflow: Towards better understanding of evolving topics in text. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (Dec. 2011), 2412–2421. URL: <http://dx.doi.org/10.1109/TVCG.2011.239>, doi:10.1109/TVCG.2011.239. 1, 2
- [DGWC10] DORK M., GRUEN D., WILLIAMSON C., CARPENDALE S.: A visual backchannel for large-scale events. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov. 2010), 1129–1138. URL: <http://dx.doi.org/10.1109/TVCG.2010.129>, doi:10.1109/TVCG.2010.129. 2
- [DZG*07] DON A., ZHELEVA E., GREGORY M., TARKAN S., AUVIL L., CLEMENT T., SHNEIDERMAN B., PLAISANT C.: Discovering interesting usage patterns in text collections: Integrating text mining with visualization. In *Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management* (New York, NY, USA, 2007), CIKM '07, ACM, pp. 213–222. URL: <http://doi.acm.org/10.1145/1321440.1321473>, doi:10.1145/1321440.1321473. 1
- [HHN00] HAVRE S., HETZLER B., NOWELL L.: Themeriver: visualizing theme changes over time. In *IEEE Symposium on Information Visualization 2000. INFOVIS 2000. Proceedings* (2000), pp. 115–123. doi:10.1109/INFVIS.2000.885098. 1
- [KLJ*12] KAYE J. J., LILLIE A., JAGDISH D., WALKUP J., PARADA R., MORI K.: Nokia internet pulse: A long term deployment and iteration of a twitter visualization. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI EA '12, ACM, pp. 829–844. URL: <http://doi.acm.org/10.1145/2212776.2212856>, doi:10.1145/2212776.2212856. 1, 2
- [Krz] KRZYWINSKI M. I.: Benefits of a circular layout. *circos*. URL: http://circos.ca/intro/circular_approach/. 2
- [KSB*09] KRZYWINSKI M., SCHEIN J., BIROL I., CONNORS J., GASCOYNE R., HORSMAN D., JONES S. J., MARRA M. A.: Circos: an information aesthetic for comparative genomics. *Genome research* 19, 9 (2009), 1639–1645. 2
- [LLC*10] LI S., LEE S. Y. M., CHEN Y., HUANG C.-R., ZHOU G.: Sentiment classification and polarity shifting. In *Proceedings of the 23rd International Conference on Computational Linguistics* (Stroudsburg, PA, USA, 2010), COLING '10, Association for Computational Linguistics, pp. 635–643. URL: <http://dl.acm.org/citation.cfm?id=1873781.1873853.2>
- [LRKC10] LEE B., RICHE N. H., KARLSON A. K., CARPENDALE S.: Sparkclouds: Visualizing trends in tag clouds. *IEEE Transactions on Visualization and Computer Graphics* 16, 6 (Nov. 2010), 1182–1189. URL: <http://dx.doi.org/10.1109/TVCG.2010.194>, doi:10.1109/TVCG.2010.194. 1, 2
- [MSH*13] MALIK S., SMITH A., HAWES T., PAPADATOS P., LI J., DUNNE C., SHNEIDERMAN B.: Topicflow: Visualizing topic alignment of twitter data over time. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining* (New York, NY, USA, 2013), ASONAM '13, ACM, pp. 720–726. URL: <http://doi.acm.org/10.1145/2492517.2492639>, doi:10.1145/2492517.2492639. 1, 2
- [NMM11] NALLAPATI R., MCFARLAND D. A., MANNING C. D.: Topicflow model: Unsupervised learning of topic-specific influences of hyperlinked documents. In *AISTATS* (2011), pp. 543–551. 2
- [PLV02] PANG B., LEE L., VAITHYANATHAN S.: Thumbs up?: Sentiment classification using machine learning techniques. In *Proceedings of the ACL-02 Conference on Empirical Methods in Natural Language Processing - Volume 10* (Stroudsburg, PA, USA, 2002), EMNLP '02, Association for Computational Linguistics, pp. 79–86. URL: <https://doi.org/10.3115/1118693.1118704>, doi:10.3115/1118693.1118704. 2, 3
- [SC] SIMPSON T., CROWE M.: Wordnet.net. URL: <http://www.ebswift.com/wordnetnet.html.3>
- [SDM14] STOJANOVSKI D., DIMITROVSKI I., MADJAROV G.: Tweetviz: Twitter data visualization. *Proceedings of the Data Mining and Data Warehouses* (2014), 1, 2
- [Set13] SETTLES B.: On “geek” versus “nerd”, 2013. URL: <https://slackprop.wordpress.com/2013/06/03/on-geek-versus-nerd/.4>
- [SRBS12] SOPAN A., REY P. J., BUTLER B., SHNEIDERMAN B.: Monitoring academic conferences: Real-time visualization and retrospective analysis of backchannel conversations. In *Proceedings of the 2012 International Conference on Social Informatics* (Washington, DC, USA, 2012), SOCIALINFORMATICS '12, IEEE Computer Society, pp. 62–69. URL: <http://dx.doi.org/10.1109/SocialInformatics.2012.20>, doi:10.1109/SocialInformatics.2012.20. 1, 2
- [TBK*12] THOM D., BOSCH H., KOCH S., WORNOR M., ERTL T.: Spatiotemporal anomaly detection through visual analysis of geolocated twitter messages. In *Proceedings of the 2012 IEEE Pacific Visualization Symposium* (Washington, DC, USA, 2012), PACIFICVIS '12, IEEE Computer Society, pp. 41–48. URL: <http://dx.doi.org/10.1109/PacificVis.2012.6183572>, doi:10.1109/PacificVis.2012.6183572. 2
- [TKE*15] THOM D., KRÄJGER R., ERTL T., BECHSTEDT U., PLATZ A., ZISGEN J., VOLLAND B.: Can twitter really save your life? a case study of visual social media analytics for situation awareness. In *2015 IEEE Pacific Visualization Symposium (PacificVis)* (April 2015), pp. 183–190. doi:10.1109/PACIFICVIS.2015.7156376. 2
- [TNYT] THE-NEW-YORK-TIMES: Close-ups of the genome, species by species by species. *IEEE Transactions on Visualization and Computer Graphics*. URL: http://www.nytimes.com/imagepages/2007/01/22/science/20070123_SCI_ILLO.html.2
- [YDGM17] YANG Y., DWYER T., GOODWIN S., MARRIOTT K.: Many-to-many geographically-embedded flow visualisation: An evaluation. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 411–420. doi:10.1109/TVCG.2016.2598885. 2
- [ZMZM15] ZHANG Z., MCDONNELL K. T., ZADOK E., MUELLER K.: Visual correlation analysis of numerical and categorical data on the correlation map. *IEEE Transactions on Visualization and Computer Graphics* 21, 2 (Feb 2015), 289–303. doi:10.1109/TVCG.2014.2350494. 2