

Visualizing Information Diffusion and Polarization with Key Statements

Andrew Salway¹, Nicholas Diakopoulos² and Dag Elgesem²

¹Computational Language Unit, Uni Research, Bergen. ²Dept. of Information and Media Studies, University of Bergen.
andrew.salway@uni.no, nicholas.diakopoulos@gmail.com, dag.elgesem@infomedia.uib.no

Abstract

This paper reports ongoing work in the “Networks of Texts and People” project, which is developing methods to visualize the social and epistemological contexts of information contained in blogs. Here, we propose an approach to visualize information diffusion and polarization in the blogosphere, with two novel characteristics. Firstly, we demonstrate how text content can be analyzed and visualized as key statements, rather than as keywords. Secondly, we sketch and discuss ideas for a visual analytic tool that integrates data about blog networks with data about the occurrence of related key statements in blog posts.

Introduction

Blogs and other social media are playing an increasingly central role in the public sphere as sources of information. However online discussions around important topics, like climate change, are vast and highly complex, and so social media visualization has a crucial role to play in helping people to interpret and interact with them. Current tools for navigating the blogosphere provide little help to users who want to assess information quality and credibility, the diversity of perspectives, how information nuggets fit together, how controversial certain statements are, who holds which viewpoints on a topic, and what the original source or provenance of a statement was. These features are not only relevant for people engaging in online discussions but they also inform the study of phenomena such as information diffusion (Katz and Lazarsfeld 1955, Hirshleifer 1995) and polarization (Sunstein 2006).

We are working to develop a visual analytics system for both social scientists and media monitors (to help their investigations into information diffusion and polarization), and for people accessing blogs (to help them to assess information). As a general point, we argue that the use of keywords to represent text content is a limiting factor for

many of the tools and techniques currently used to access, analyze, and visualize the blogosphere. Furthermore, we note that whilst some tools visualize explicit linkage networks, and some visualize implicit textual relationships, there is an underexplored opportunity to develop visualizations that fuse and allow for the simultaneous exploration of these two kinds of networks.

With regards to the representation of text content, this paper demonstrates a novel technique that extracts a set of key statements about a given topic from the blogosphere. In effect, these key statements characterize different points of view and different aspects of a debate. We then sketch and discuss the design of the interface for an interactive visual analytic tool that integrates data about the occurrence and relations of key statements in blog posts, with data about links between blogs. Finally, we show how this design supports the elucidation of information diffusion and polarization within a topical blog corpus.

Related Work

Current online tools for blog search efficiently deliver blog posts matching a query, in some cases with structural information, e.g. incoming links from other blogs, or authority. However, these services provide little help in answering questions about the structure of the discussion around topics or the pattern of its distribution in the blogosphere. Elsewhere, a common technique is to extract keywords to produce word clouds and trend graphs that depict the topics of blogs, but these do not reveal what in particular is being said about the topics. Other tools can be used to map and analyze networks in the blogosphere but they provide only a static picture of the network and nothing about the development of interactions over time. Because there is no specific information about the textual context of the links, the interpretation of the links is difficult. It is hard to know, for example, if a particular blogger has linked to another blog because she agrees with it, criticizes it, or expands on a statement in the other blog.

Visualization research has addressed some of these limitations by investigating alternatives for representing and exploring blog-related information. This research can be understood along three dimensions: unit of interest, analytic scheme, and analytic unit. In terms of unit of interest, some related work concerns itself with the analysis of individual blogs (e.g. Indratmo and Gutwin 2008; Bross, Schilf, and Meinel 2010). However, in our work we are instead interested in developing methods to understand a topically oriented corpus of blog texts. Various analytic schemes have been explored with respect to blog analysis, including temporal schemes (Leskovec Backstrom, and Kleinberg 2009; Indratmo and Gutwin 2008), as well as spatial text similarity schemes (e.g. Michelle et al. 2007; Gretarsson et al. 2012). Our work incorporates both temporal and spatial similarity schemes to more effectively support answering questions relating to information diffusion and polarization in the same unified interface. Finally, we can see a variety of analytic units in the literature on document and blog visualization. In particular, keyword based visualizations (Bansal and Koudas 2007; Indratmo and Gutwin 2008) and topic based visualizations (Gretarsson et al. 2012; Michelle et al. 2007) are the most often used representations. More recently, quotations have also been used as the analytic unit for blog exploration (Leskovec et al. 2009; Simmons et al. 2011). However, there has been little, if any, research that we are aware of that has incorporated the standard textual units of analysis (keywords and topics, let alone statements) together with explicit linking information. In our work we aim to close this gap and tie together a statement-based unit of analysis with a visual representation depicting topical and network based relationships in a blog corpus.

Automatic extraction of key statements

If we can extract key statements, rather than keywords, from blog posts then this will provide a more informative view on the content of individual posts and of online debates. If these statements, and hence the posts containing them, can be related then this will provide new data about networks of texts from which we can learn more about information diffusion and associated social networks.

Previous research has shown how statements and meta statements can be extracted from the blogosphere in the form of relations between entities, e.g. people, organization, places and dates (Mesquita and Barbosa 2011). Whilst impressive in its own right, this approach is not directly relevant for our main goal: (1) we need to capture more diverse and detailed kinds of information, but we do not require a formal knowledge representation as output, and (2) we want an approach that can be adapted easily to multiple languages and domains, which is not a

feature of most information extraction approaches. Instead, we use a technique described and evaluated by Salway et al (2010) to extract a list of semi-structured key statements about a given topic from the web, ranked by a keyness metric; the statements comprise the things typically written about the topic. To port the technique to different domains and languages, it is only necessary to specify a sentence template, e.g. "TOPIC CUE FILLER", and some cues.

Figure 1 shows a selection of the statements extracted from the English-language blogosphere for the topic "climate change", in January 2012. Over 1,000 different statements about climate change were returned using one simple sentence template and nine cues. These results have yet to be evaluated, but an initial inspection indicates that

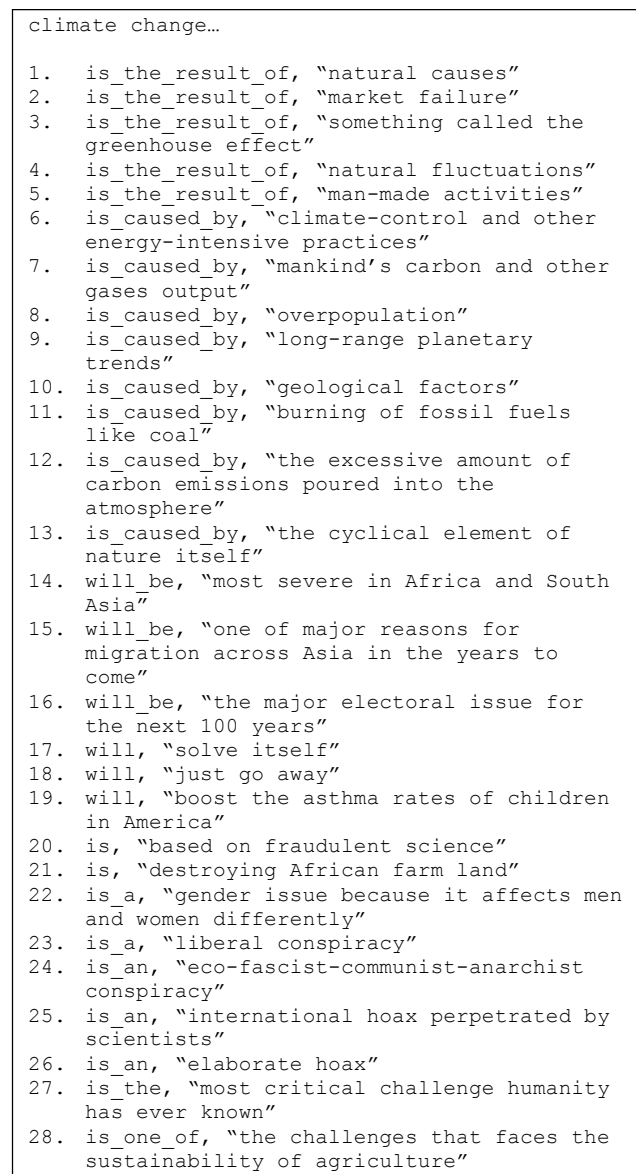


Figure 1. A small selection of statements about climate change extracted automatically from the English-language blogosphere, January 2012.

the key statements capture many diverse aspects of the climate change debate. Figure 1 begins to both demonstrate and address the challenge of grappling with a vast and complex online discourse. It shows statements made from a variety of perspectives – scientific (e.g. statements 3, 9, 10, 11), political (16, 23, 24) and social (22, 27). Some statements are apparently global in their reference, whilst others take more localized viewpoints (14, 19). Just in the small sample presented here it is possible to see examples of various ways in which statements could be related. Some could be considered equivalent (17, 18), (1,4) and (20, 25, 26). There are some pairs of statements that both share the main point, but may augment each other (7, 12) and (21, 28). Then there are the contradictory statements (1/4/9/13, 5/6/7/8). Other statements assume agreement on the basic point that climate change is caused by CO₂-emissions but focus on different mechanisms (2, 8). Other statements assume that global warming is happening but emphasize different consequences of it (15, 19, 21, 22).

Thus it seems that the identification of key statements can reveal the diverse content of an online discussion. However, to explore diffusion and polarization, it will be necessary to instantiate relations between statements. Current meme-tracking techniques rely on matching exact

or nearly-exact forms of a phrase. Such techniques could be enhanced with data about instances where memes appear in different forms, e.g. equivalent statements. Conversely, to show polarization requires knowing about contradictory statements. To infer such relations automatically, we will draw on work from the field of paraphrase and textual entailment detection (Androutsopoulos and Malakasiotis 2010).

Visualization design

As a first step towards working with large-scale and complex online discourses, we are developing a visual analytics system that operates on data about statements and their relations, together with blog network data that relates blogs and blog posts via blog rolls and hyperlinks. Our users are social scientists, media monitors, and members of the public who want to understand online discussions more fully. We wish to support the exploration of three analytic use-cases or questions: (1) How do statements diffuse across the blog network over time? (2) How polarized are different statements within the blog network? (3) How can we characterize statements and related statements?

Figure 2 shows a mockup of our interface depicting a

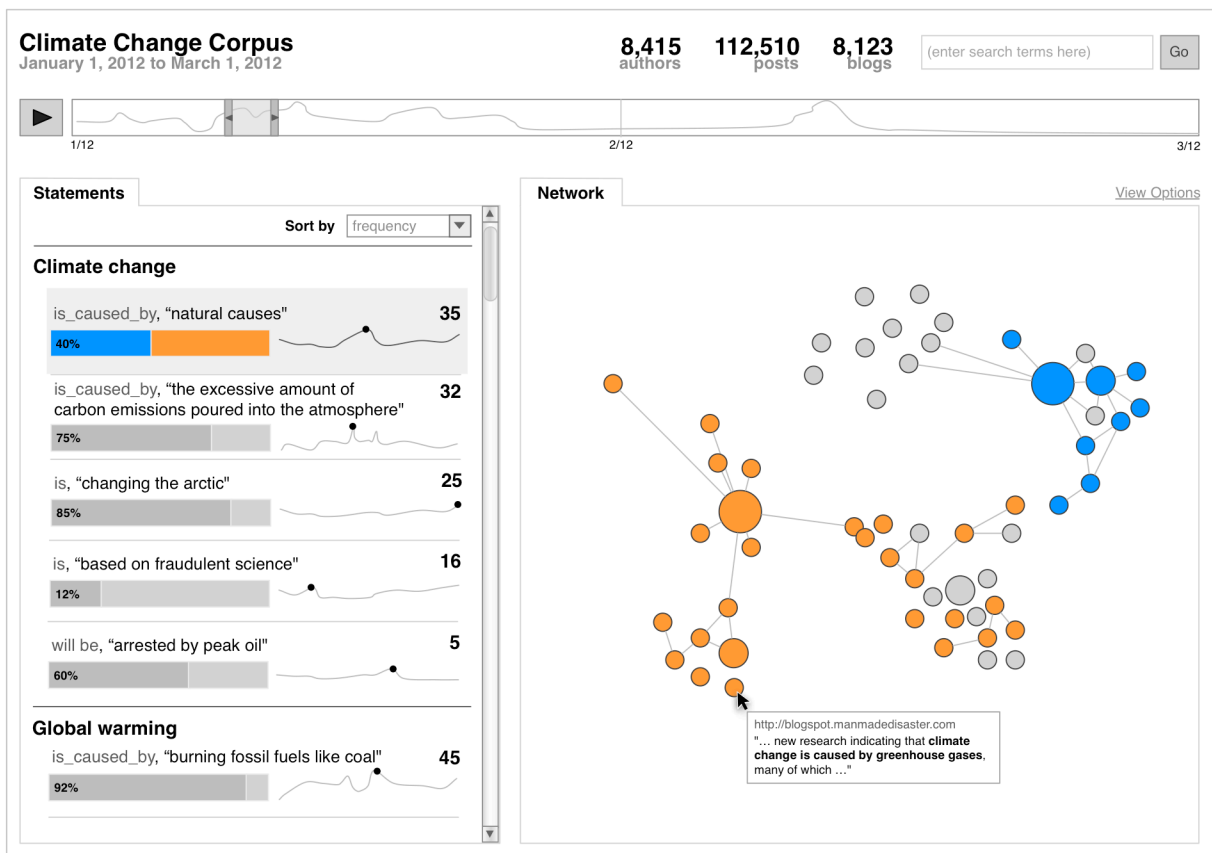


Figure 2. A mockup of an interface depicting a corpus of blog posts collected around the topic of climate change. Extracted statements are listed in the panel at left and the network is shown in the panel at right. The selected statement is shown in color with blue (similar statements) and orange (dissimilar or disagreeing statements).

blog corpus collected around the topic of climate change. The interface has three primary areas: the temporal navigator at the top, the statement browser (shown left), and the network view (shown right). The temporal navigator depicts the volume of blog activity over time and the user can filter for a temporal subset of the corpus by dragging the slider. Hitting the play button starts an animation with each frame depicting a time-slice of data.

The statement browser area allows the user to browse the statements that have been extracted from the corpus. Each topic is shown prominently with all cues and filler text ordered below it; frequency in the corpus is given to the right. Each statement has two visualizations associated with it: a stacked bar chart, and a sparkline. The stacked bar chart shows the percentage of statements that are clustered both similarly (darker gray) and dissimilarly (lighter gray). The sparkline shows the relative frequency of use in the corpus over time, with the peak indicated by a black dot. A statement is selected by clicking on it, which causes it to colorize and highlight the network display. Clicking the spark graph expands the cluster and lists all statement variants sorted by time.

The network display maps each blog in the corpus to a node, with the in-degree of that blog representing the size of the node. Blue nodes are blogs that use the selected statement or a similar statement: orange nodes are blogs that use a disagreeing statement. Based on the view settings, edges in the graph depict either the network defined by blogroll links or the network defined by links between posts. In order to reduce visual clutter, edges are only drawn if they connect to one of the orange or blue nodes. The nodes in the graph are spatially arranged by statement similarity using multi-dimensional scaling or principal components analysis (Gretarsson et al. 2012). Using these methods, a feature vector is constructed for each blog based on the statements it uses. These high dimensional vectors are projected to reflect high dimensional similarities in the 2D layout. Thus blogs with similar content appear near each other on the graph, and, with edges depicting explicit citations, this allows the user to see both content similarity and network connectivity at the same time. To avoid spatial inconsistency based on varying filters, one layout will be done for the network based on an entire corpus.

To look at diffusion the user can first select a statement and then, using the temporal navigator, scrub through the data to see how the network lights up with blue and orange as the statement diffuses. The edges depicted indicate likely paths for the color (i.e. statement) to spread. Another view will show small multiples of the network as an alternative to animation. Selecting a statement colorizes the network according to the similar and dissimilar or disagreeing statements. This will show if there is a large separation in the statement space in terms of where the

statement and its antithesis is used – one possible indication of polarization. Statements are characterized in the statement panel by indicating the cues for each statement as well as the proportion of similar and dissimilar statements that have been identified, and the volume of the use of the statement over time. Statement clusters can be drilled into to show additional details, e.g. individual statements within a cluster, timestamps, etc.

Closing Remarks

Current tools for analyzing blog networks are limited in several ways. The use of keywords does not reveal what is said about topics, nor how arguments about topics are formulated. Furthermore, while textual information is highly relevant to the interpretation of links, it is not well integrated with the structural information about the network of blogs. Approaches that overcome these limitations should provide the basis for tools that would be useful to social scientists, media analysts and general users.

Acknowledgements. This research was supported by a grant from the Norwegian Research Council's VERDIKT program.

References

- Androustopoulos, I. and Malakasiotis, P. 2010. A Survey of Paraphrasing and Textual Entailment Methods. *Journal of Artificial Intelligence Research* 38:135-187.
- Bansal, N. and Koudas, N. 2007. BlogScope: a system for online analysis of high volume text streams. In *Proc. VLDB '07*.
- Bross, J.; Schilf, P.; and Meinel, C. 2010. Visualizing Blog Archives to Explore Content- and Context-Related Interdependencies. In *Conf. Web Intelligence and Intelligent Agent Technology*.
- Gretarsson, B. et al. 2012. TopicNets: Visual Analysis of Large Text Corpora with Topic Modeling. *ACM Trans. on Intelligent Systems and Technology* 3(2).
- Hirshleifer, D. 1996. The Blind Leading the Blind: social influence, fads, and information cascades. In *The New Economics of Human Behavior*. Cambridge: CUP.
- Indratmo; Vassileva, J.; and Gutwin, C. 2008. Exploring Blog Archives with Interactive Visualization. In *Proc. AVI '08*.
- Katz, E. and Lazarsfeld, P. 1955. *Personal Influence*. Free Press.
- Leskovec, J.; Backstrom, L.; Kleinberg, J. 2009. Meme-tracking and the Dynamics of the News Cycle. In *Proc. KDD 2009*.
- Mesquita, F. and Barbosa, D. 2011. Extracting Meta Statements from the Blogosphere. In *Proc. ICWSM 2011*.
- Michelle, G.; Payne, D.; McColgin, D.; Cramer, N.; and Love, D. 2007. Visual Analysis of Weblog Content. In *Proc. ICWSM 2007*.
- Salway, A.; Kelly, L.; Skadiņa, I.; and Jones, G. 2010. Portable Extraction of Partially Structured Facts from the Web. In *Proc. ICETAL 2010*, LNAI 6233, 345-356. Heidelberg, Springer.
- Simmons, M.; Adamic, L.A.; and Adar, E. 2011. Memes Online: Extracted, Subtracted, Injected and Recollected. *ICWSM'11*.
- Sunstein C., R. 2006. *Infotopia*. Oxford: OUP.