

Anonymous FreeSpeech

Carl Vogel

Computational Linguistics Group
Centre for Computing and Language Studies
School of Computer Science and Statistics
Trinity College Dublin
Dublin 2
Ireland
vogel@tcd.ie

Abstract

In public and private discourse, some may be heard to express disquiet about the supposed dangers of anonymity. Anonymous suggestion boxes may be classed with anonymous accusation of crime with the accusation forming the basis of legal proceedings. In some contexts, anonymity does appear to create danger. However, other contexts reveal that important benefits accrue from having the possibility of anonymous expression. Some of the literature on behavioral impacts of anonymity is reviewed with the aim of analyzing social media systems in light of their support for anonymous contribution. A new system is described. The new system supports anonymous communication, while thwarting some of the obvious risks that anonymity affords.

Keywords: social media, anonymity, free speech, sentiment analysis, meme analysis, ephemera, dynamic social groups

1. Introduction

Popular thinking is conflicted on the value of anonymity. The US Constitution's First Amendment protects anonymous free speech, and accordingly the US courts have established guidelines for discovering the identity of anonymous internet posters (Den, 2001; Mob, 2007; Kri, 2008). It appears to be a coherent position to hold that anonymity has no social value, since it is an asocial concept. Davenport (2002, p.33) wrote "By allowing anonymous Net communication, the fabric of our society is at risk." However, it will be seen (§2.) that if anonymous communication is asocial, it is paradoxically so, since there are conditions in which group identity is more strongly present for anonymous rather than identified communicators. Neumann (1996) describes the uneasy availability of anonymity and argues that accountability must always compromise anonymity. The CEO of Facebook has called for the end of internet anonymity;¹ however, he has a commercial interest in the users of his company's system being identifiable. The UK Defamation Act 2013 provides means to preserve the anonymity of posters, but passes accountability to hosting website operators.²

It is a reasonable principle of law that if one is to be tried, one should know what the accusation is and who the accuser is, and this is contrary to the possibility of anonymity. On the other hand, society has not yet established effective means of protecting "whistle-blowers", who are in many cases neither directly harmed by the accused nor have any ground to benefit from a successful claim, from retribution by the accused or the system in which the accused operates.

Social media have been attributed a role in enabling societal reform (Khondker, 2011; Eltantawy and Wiest, 2011), and prohibitions of anonymity in social media have been cited as an obstacle to reforms (Youmans and York, 2012). However, one mob's societal reform is another's treasonous rebellion. Nonetheless, a number of social media platforms have emerged in support of anonymous communication. While each of these has a niche, and acknowledging that the actual utility of any tool frequently diverges from its intended use, the architecture of a new platform for anonymous communication is detailed.

This article reviews some of the literature on the potential positive or negative impacts of anonymity (§2.). The impacts are mixed whether the interactions are online or in person, both with respect to likelihood of aggression and generosity. It also describes extant systems that support anonymous communication to varying degrees. The design of a new system is specified (§3.): it is a system that supports anonymous communication in its outward facing function, and it provides a testing ground for methods of text analytics and analytics of multi-modal content within its system architecture.

2. Background

2.1. Effects of anonymity

In a lab based setting Connolly et al. (1990) showed that anonymous groups working jointly on idea generation produced more ideas than identified groups, particularly when a critical mode of interaction was induced by a confederate. Postmes et al. (2001) demonstrated that anonymous groups would exhibit behaviors consistent with introduced norms, while identified groups did not – this finding alone is a powerful rebuttal to the notion that anonymity is inherently asocial. Smith et al. (2007) obtained results which showed that participating individuals with strong group self-identification behaved according to group norms even in anonymous conditions. Diener (1976) showed in a lab

¹See a 2011 staff journalist article in *The Daily Mail* – www.dailymail.co.uk/news/article-2019544/Facebook-director-Randi-Zuckerberg-calls-end-internet-anonymity.html – last verified February 2014

²<http://www.legislation.gov.uk/ukpga/2013/26/section/5/enacted> – last verified February 2014.

setting involving male undergraduate students that group presence yielded decreased aggression, but that providing anonymity had no effect on aggressive behaviors. Pizazza and Bering (2008) conducted a mixed gender study in which economic punishment behaviors were more pronounced in identified conditions than where anonymity was preserved.

On the other hand, Donnerstein et al. (1972) found aggression in lab situations invoking race conflict more likely in situations of preserved anonymity than where the aggressor could be identified. More recently, Ellison-Potter et al. (2001) found that people drive automobiles more aggressively when anonymous than when identified. In a natural field experiment Alpizar et al. (2008) found public donors to be 25% more generous than anonymous benefactors. Rule-breaking has been found in lab settings to be more likely where participants are not identifiable and not accountable for their actions (Nogami and Takai, 2008).

Lapidot-Lefler and Barak (2012) found that lack of eye contact was a greater contributor to toxic online communications than anonymity in computer mediated communications in which visibility and eye-contact were mediated by web-cams, and anonymity, by random online aliases.

Stone et al. (1977) showed that in student evaluations of lecturers, signed statements were more positive than anonymous feedback.

Kidder et al. (1977) demonstrated gender differences in behavior where decisions regarding generosity would become public or remain anonymous, with men choosing a more generous response than women when the decision would remain anonymous (the reverse of when decisions were made public). Durant et al. (2002) also noticed that anonymity guarantee yielded higher quality data in medical self reporting than where identification was supported.

The literature mentioned above demonstrates both support for anonymity as a social construct and potentially negative impacts of anonymity. Popular opinion on anonymity is perhaps conflicted because the value of anonymity is context dependent.

2.2. Extant systems supporting anonymity

I distinguish between systems that, perhaps despite stated terms and conditions of use, do not force identification of users by verifying the identity of the person registering but support the identification of posters or creation of online persona (e.g. Facebook, Twitter, Reddit, Tumblr), and those systems that purport to maintain the anonymity of posters. Kwikdesk³ provides an outwardly anonymous messaging service like Twitter, with messages that last only a very short amount of time, but requires user registration. Self-created account names have been analyzed to assess the likelihood that account names that appear in sentiment lexica (e.g. SentiWS (Remus et al., 2010)) are likely to host content that anyone wishing to filter offensive content might wish to exclude, and found a correlation between negative sentiment user identifiers and such content (Vogel, 2013). While self-created account names do not reveal identity, they do lend well to establishing distinctive online

persona. This latter category, systems that do not record details of posters, are described further below.

Some sites allow posters to post content via random identification numbers into threads of communication, for example 4chan.org.⁴ From third party perspectives, unless posters reveal themselves, they are anonymous. To the system managers, posters must reveal an effective email address, and presumably system logs record connecting IP address information — the terms and conditions suggest that the system managers have the means of identifying and permanently banning posters, which suggests a relatively high level of identification. Postings to this service have been studied in light of the fact that a high-volume of postings exist, despite the very ephemeral nature of some of the threads and estimated 90% anonymous posting (Bernstein et al., 2011). This demonstrates that a social niche is filled through the possibility of such a facility existing to support outward facing anonymity, in spite of the frequently offensive content (Knuttila, 2011). Schoenebeck (2013) points out that YouBeMom.com provides a positive means of expression of negative emotions through the dis-inhibition enabled by anonymity.

A service provided by smalltime.com⁵ allows users to post messages anonymously to the next person who posts a message. The system functions something like an ephemeral message-in-a-bottle, not directly supporting a general view of sentiments expressed by many. A poster need not post self-identifying information, nor information that reveals who intended readers are. In contrast, hadtosay.com⁶ allows users to send email messages to named individuals, but from an anonymous address, with the proviso that the content of the message is posted to a public forum. This can have potential negative consequences for recipients whose email contact details are deposited with the system; however, it is curious to see the preponderance of sentiment that emerges for occasions like Valentine's Day.

2.3. Observations

The review here has attempted to be suggestive rather than exhaustive. The literature reveals both benefits and risks associated with anonymous communication. Social media systems support anonymity to varying degrees, most depending on the construction of online social personae, if not actually identifying individual posters. Some that support greater levels of public anonymity still require system registration, and nonetheless have become a clearinghouse for content of dubious merit.

3. Design of a new system

This section provides a specification of FreeSpeech, a system for anonymous expression and perusal of ideas. The system is currently under prototyping.

3.1. Desiderata

As in any system design, some principles are formulated in positive terms and others are proscriptions.

³See kwikdesk.com — last verified February 2014.

⁴Last verified February 2014.

⁵See <http://www.smalltime.com/anon.html> — last verified Feb 2014.

⁶See <http://hadtosay.com/> — last verified Feb 2014.

As a matter of principle the system should afford users the possibility of posting text and reading texts. The system should provide comparable facilities for other modalities; however the primary focus should be the articulation of thoughts. Users may wish to tag the text in various ways to indicate perceived categories appropriate to the text. However, the system should, in the background, supply continuous indexing that re-categorizes submissions in additional, alternative manners.

The system should be built on a principle that posters and readers may retain anonymity. It should be possible for a user to express thoughts without registering an email, IP address or other identifying signals, and similarly it should allow other users to examine expressed thoughts, via an assortment of possible searches and indexing mechanisms.

The system should also respect the rule of law in the jurisdiction in which it is hosted, and to the extent possible given inconsistencies, internationally. It should facilitate the expression of thought and estimation of current thinking. Content that is not legal should be rejected by the system; content that is merely offensive or inappropriate should be quarantined.

Preservation of anonymity may prevent the system from allowing the geo-location of current thinking, but should facilitate analysis of the time course of popular thought.

The system should not provide support for the maintenance of persona, whether of actual individuals or constructed for an online presence. The primary function of the system is not maintenance of synchronous or asynchronous online dialog.

The system should support multi-lingual views on the content recorded.

Through the totality of its public facing functions as supported by the system-interior functions, the system should provide a testing ground for research associated with the system-interior functions, content analytics, in particular. It should be possible to easily adapt the system to include alternative classifiers and test the relative differences and similarities obtained by integrating alternative methods.

3.2. Public features

The individual posting content may do so into a window amidst extant content, and doing so will automatically tag the new content in relation to the extant visible content, or posting may be made from a top level portal that comes with no *a priori* tags. In either case, the poster may add additional tags to the content.

An individual seeking to monitor thought expressed in the system should have the possibility of seeing raw content or content aggregation summaries. Content should be visible in relation to the most recent postings or in relation to categories. The categories may be supplied as a query term from the monitor or dominant themes. Category related views should also be possible with respect to past temporal intervals.

It should be possible for users to register support or dissent with respect to ideas hosted by the system, and the time course of support and dissent should be available to monitors. Support or dissent may be registered in categorical terms or supplemented with textual comment. Similarly,

search term use will also be aggregated and reported for the benefit of monitors wishing to sample global thought.

While it will not be possible in general to sample content relative to geo-location parameters, because of the multi-lingual support, it should be possible to view content with language of input as a parameter. Equally, it should be possible to view (via automatic translation), content made available cross-linguistically.

The system requires active monitoring because as a matter of principle it does not record email addresses of individuals wishing to be informed of updates.

The system should not declare which system-interior classifiers are at work in clustering postings together at any time of use.

3.3. System-interior features

On the system side, there are facilities for multi-dimensional indexing of content. From the perspective of research in content analytics and machine learning. The indexing of the content comes from the tags supplied by the user, or supplied contextually at the point of posting, or from background indexing processes.

Interfacing to machine translation will be necessary possibly as a remote online service. Language detection services may be used in order to tag input textual content (Cavnar and Trenkle, 1994; van Noord, 1998).

Background analysis of posted content is intended to support filtering out spam, defamatory claims, pornography and so on.⁷ At the same time, it should be possible to classify and index posts according to multiple criteria: terms used, user-supplied tags, system detection of similarity to other posts, system detected anomaly of content. The same general techniques that are used for detecting effects of source language on translation and stylistic quality can be used to detect content based anomalies (Vogel et al., 2013; Moreau and Vogel, 2013). In fact, it is a design principle of the system that it should provide a scenario for testing out theories and methods of content classification.

The system will thwart posting of links to external content. Thus, pre-processing of posts will remove external links. Content beyond text that is uploaded will be filtered. While the system will support the posting of visual content, this will be secondary to text.

Indexing will be based on user-supplied key-words as well as system discovered relevant terms. Index expansion will be facilitated with reference to machine readable lexical resources like WordNet.

Tools for detection of named individuals and automated anonymization may be necessary in order to avoid compromising the system (or supporting defamation).

A database suitable for large scale data sets will accommodate the data and its multiple indices.

3.4. Interaction types

3.4.1. Perusal

A person may approach the site from its front end web presence, and through a default language setting will be presented with a few alternative views of the data recorded in

⁷See O'Brien and Vogel (2003), for example.

the back-end, the n most: recent posts; endorsed posts; denounced posts; anomalous posts; active user tagged topics; active system categorized topics. A reader may therefore browse posts through any of those mechanisms for ranking past postings, or a user may search for posts, which will be indexed by their content and also by tags (user supplied, and system generated, as described above). The possibility of posting will also be possible at this point (see §3.4.2.), and an alternative access point will support posting only, without a view of current topics.

In selecting a post to view, an additional interface context is created in which other posts related by topic, recency, etc. are also visible. Posts will have unique identifiers, but this will not include any system-generated signal of who the user is. Thus, readers should be able to quickly view linked (and linking) posts and their contexts. Exploring a post should give the impression that there is an unending potential for expanding links from posts to related posts, even if those links are not user-supplied links, but links based on shared index terms.

Tools for aggregate summary analytics will also be available. That is, it may be interesting to some users not to read posts directly, but to study the volume of postings on topics over time, or as a function of language of posting, etc. Tools for visualizing trend analysis will be supplied.

3.4.2. Posting

A person may post content into the context supplied by postings perused. This may, if attending to immediacy and certain contextual features, have the effect of sustaining an interaction between contemporaneous posters. However, postings all without user identifiers, will make it difficult to sustain a dialogue as in other online social media – direct quotation may be the easiest means of directly signalling such interactions, but post identifiers may link posts, too. Rather, postings will be amid other postings, creating the effect of ideas interacting with other ideas, rather than personalities interacting with personalities. On entering a contribution, a user posting will have the option of including multi-media files in the post (knowing that these will not appear with the post immediately, as the media file filtering is expected to take longer than text filtering) as well as key phrases that contribute to the index of the post. A post may link to the post-id of another post, but not outside the system.

A user will not have a facility to edit or delete any post after accepting submission. Some edit facilities, including draft preview, are supported. Additionally, a feature allows a user to decline to submit a draft, and in that case it is not recorded in the system except as a count of aborted drafts. Once the draft is committed to the system, the post ceases to be controlled by the user posting content.

Once a post passes text filters, it is available for perusal to others in the original context, and in other contexts linked by time, user-supplied key phrases, and key phrases generated from the text, and through index expansion. Thus, a post may appear to a user to be part of one topic of discussion, but may, through indexing be visible to others in relation to topics of discussion that the original poster perhaps did not consider.

3.5. Anticipated use

Given that social media exist for interactive social communication, it is not anticipated that the FreeSpeech system will occupy this dimension of interaction. Rather, it is foreseen that the system will be a living meme monitor, serving to digest the global pulse in a way that extensively supplements information about trending topics or top search terms. It is imagined that the system will attract use more general than the automatic mood detector supplied by We-FeelFine.org which trawls the internet for content that contains the expressions “I feel” and “I am feeling”.⁸ Yet neither is the system likely to emerge as a crowd-sourced compendium of knowledge in competition with Wikipedia (not least because wiki functions such as the potential for user editing or deleting of posts will not be supplied – not for the others’ posts, nor their own).

Some individuals may try to subvert the spirit of anonymity either by naming themselves in their posts or by creating trademark identifiers that effectively tag their contributions and establish a public persona. Nonetheless, it is hoped that the system will facilitate debate of ideas in their own sake, and without reference to their proponents. Within this particular social medium, the role of “key movers” will be diminished through anonymity, so that ideas may be applauded or detracted because of the ideas themselves rather than because of the personalities expounding them or endorsing them.

It is foreseen that in order to overcome system obstacles to identifying users will make it cumbersome to use the system for social chat. Therefore, while it might be feasible to use the system to assemble a “flash mob” (Gore, 2010), it would be very difficult to foster a pen-friendship through the system.

Background filter that remove obnoxious content are meant to have the effect of deterring the posting of such content and use of the site for sharing such. Similarly, it is not expected to provide affordances for copyright violation.

The system is not proposed with any model of economic exploitation in mind as a driving principle; rather, the driving principles are in the provision of a tool for expressing thoughts, and for testing out methods of content analytics. However, it might be noted that the system, as described, is compatible with having a side panel that displays relevant advertizing. Moreover, it is compatible with the goal of the system in providing a testing bed for research into methods of content analytics that one might explore study of relations of compatibility (as opposed to inappropriateness) between content displayed in the primary content areas and any ads displayed in such an advertizing area, as an instance of the general problem of monitoring compatibility of message across types of media (Pastra, 2006; Pastra, 2008). Such ads would presumably have to respect the constraints imposed on regular posts.

4. Evaluation

The interactive elements of the system will be evaluated through its uptake: the system may be deemed successful if it demonstrates growth in usage. By construction, it will

⁸See <http://www.wefeelfine.org> – last verified February 2014.

not be possible to measure the number of distinct users.⁹ Rather it will be necessary to measure the quantity of posts, diversity of topics, and depth within topics.¹⁰ Empirical hypotheses underlying the system design may be simultaneously tested. It is a hypothesis of this work that sufficiently many alternative systems exist to support fostering relationships or hosting particular kinds of multi-modal content. It is not imagined that this system will compete with those on the same ground, because this system is not designed to facilitate that sort of content management. Nonetheless, it will be possible to evaluate the quantity of postings that are attempted, but filtered through the various possibilities of excluded contexts. The ratio of desirable to undesirable content in this system will be interesting to monitor in itself and in relation to the same ratios for other social network systems.

The capacity of the system to support research into analysis of the competition of ideas as a dynamic system, meme analysis, is also to be evaluated. This will require study once sufficient content accumulates within the system.

The back-end facilities are designed in order to provide a sandbox for developing and evaluating methods of content analytics. Evaluation of the proposed system in this dimension relates to the ease with which alternative classifiers may be correlated with each other and differentiated from each other, for example.

5. Final remarks

This article has defended the role of anonymity in online communication and expression of ideas. It has suggested a new system for facilitating anonymous contribution of ideas into public debate. The system provides features for users to express their ideas and for others to monitor ideas that are current (or which have subsided in general interest). The system itself provides a locus for research in multi-lingual multi-modal content analytics. It may be argued that the design described here is ambivalent in approaching content beyond the textual mode. It is true that the prohibition of external links and severe filtering of other content provides a greater level of visual filtering than is applied to text. These features are seen as limiting the potential for an anonymity preserving communication device to be used in ways that other online facilities already appear to support with ample resources. That should leave interest in the system presented here primarily for expressing thoughts that require consideration.

The value of anonymous communication during disaster management remains for further consideration. Manifestly, during disaster management, the provenance of messages matters to citizens a great deal, and trust is accorded to messages seemingly as a direct function of the trust accorded to the public persona of the source. Without automatic geo-location of postings in the system described, it would be challenging to use the system in order to assess the situation “on the ground” as reported by citizen observers, and

⁹Arguably, systems like Facebook are not in a position to provide reliable methods for estimating unique user numbers accurately, either; however, here, it is in principle not possible.

¹⁰Developing metrics for topic diversity and depth presents an interesting scientific challenge.

it would be difficult to convey messages of urgency to the public one would hope to respond, if one were using the proposed system as a channel for such communications. However, in the case of emergencies that are more politicized, it is easy to see the value that could be provided by anonymous communication mediated through FreeSpeech.

6. Acknowledgements

I am grateful to Eamonn Lawlor for prototyping some of the features of FreeSpeech suggested here in the context of his final year project on the undergraduate of honors course in Computer Science, Linguistics and German at Trinity College Dublin; he also alerted me to the existence of 4chan and Kwikdesk.

7. References

- Francisco Alpizar, Fredrik Carlsson, and Olof Johansson-Stenman. 2008. Anonymity, reciprocity, and conformity: Evidence from voluntary contributions to a national park in costa rica. *Journal of Public Economics*, 92(56):1047 – 1060.
- Michael Bernstein, Andrs Monroy-Hernandez, Drew Harry, Paul Andr, Katrina Panovich, and Greg Vargas. 2011. 4chan and /b/: An analysis of anonymity and ephemerality in a large online community. In *Fifth International AAAI Conference on Weblogs and Social Media*.
- W. B. Cavnar and J. M. Trenkle. 1994. N-gram-based text categorization. In *Proceedings of Third Annual Symposium on Document Analysis and Information Retrieval*, pages 161–175. Las Vegas, NV, UNLV Publications/Reprographics.
- Terry Connolly, Leonard M. Jessup, and Joseph S. Valacich. 1990. Effects of anonymity and evaluative tone on idea generation in computer-mediated groups. *Management Science*, 36(6):689–703.
- David Davenport. 2002. Anonymity on the internet: Why the price may be too high. *Communications of the ACM*, 45(4):33–35.
2001. *Dendrite International v. Doe*, 775 A2d 756 (N.J. App. Div.).
- Edward Diener. 1976. Effects of prior destructive behavior, anonymity, and group presence on deindividuation and aggression. *Journal of Personality and Social Psychology*, 33(5):497 – 507.
- Edward Donnerstein, Marcia Donnerstein, Seymore Simon, and Raymond Ditrachs. 1972. Variables in interracial aggression: Anonymity, expected retaliation, and a riot. *Journal of Personality and Social Psychology*, 22(2):236 – 245.
- Lauren E. Durant, Michael P. Carey, and Kerstin E.E. Schroder. 2002. Effects of anonymity, gender, and erotophilia on the quality of data obtained from self-reports of socially sensitive behaviors. *Journal of Behavioral Medicine*, 25(5):439–467.
- Patricia Ellison-Potter, Paul Bell, and Jerry Deffenbacher. 2001. The effects of trait driving anger, anonymity, and aggressive stimuli on aggressive driving behavior. *Journal of Applied Social Psychology*, 31(2):431–443.
- Nahed Eltantawy and Julie Wiest. 2011. The arab spring—social media in the egyptian revolution: Reconsidering

- resource mobilization theory. *International Journal of Communication*, 5(0).
- Georgiana Gore. 2010. Flash mob dance and the territorialisation of urban movement. *Anthropological Notebooks*, 16(3):125–131.
- Habibul Haque Khondker. 2011. Role of the new media in the arab spring. *Globalizations*, 8(5):675–679.
- Louise H. Kidder, Gerald Belletirrie, and Ellen S Cohn. 1977. Secret ambitions and public performances: The effects of anonymity on reward allocations made by men and women. *Journal of Experimental Social Psychology*, 13(1):70–80.
- Lee Knuttila. 2011. User unknown: 4chan, anonymity and contingency. *First Monday*, 16(10).
2008. Krinsky v. Doe 6, 159 Cal.App. 4th 1154 (Cal. Ct. App.).
- Noam Lapidot-Lefler and Azy Barak. 2012. Effects of anonymity, invisibility, and lack of eye-contact on toxic online disinhibition. *Computers in Human Behavior*, 28(2):434 – 443.
2007. Mobilisa v. Doe, 170 P.3d 712, (Ariz. Ct. App.).
- Erwan Moreau and Carl Vogel. 2013. Weakly supervised approaches for quality estimation. *Machine Translation*, 27:257–280.
- Peter G. Neumann. 1996. Inside risks: Risks of anonymity. *Commun. ACM*, 39(12):162–, December.
- Tatsuya Nogami and Jiro Takai. 2008. Effects of anonymity on antisocial behavior committed by individuals. *Psychological Reports*, 102:119–130.
- Cormac O’Brien and Carl Vogel. 2003. Spam filters: Bayes vs. chi-squared; letters vs. words. In Markus Alesky et al., editor, *Proceedings of the International Symposium on Information and Communication Technologies*, pages 298–303.
- Katerina Pastra. 2006. Beyond multimedia integration: Corpora and annotations for cross-media decision mechanisms. In *Proceedings of the 5th Language Resources and Evaluation Conference (LREC)*, pages 499–504.
- Katerina Pastra. 2008. Cosmoroe: A cross-media relations framework for modelling multimedia dialectics. *Multimedia Systems Journal*, 14(5):299–323.
- Jared Piazza and Jesse Bering. 2008. The effects of perceived anonymity on altruistic punishment. *Evolutionary Psychology*, 6(3):487–501.
- Tom Postmes, Russell Spears, Khaled Sakhel, and Daphne de Groot. 2001. Social influence in computer-mediated communication: The effects of anonymity on group behavior. *Personality and Social Psychology Bulletin*, 27(10):1243–1254.
- R. Remus, U. Quasthoff, and G. Heyer. 2010. SentiWS - a publicly available German-language resource for sentiment analysis. In *Proceedings of the 7th International Language Resources and Evaluation (LREC’10)*.
- Sarita Yardi Schoenebeck. 2013. The secret life of online moms: Anonymity and disinhibition on youbemom.com. In *Seventh International AAAI Conference on Weblogs and Social Media*.
- Joanne R. Smith, Deborah J. Terry, and Michael A. Hogg. 2007. Social identity and the attitudebehaviour relationship: effects of anonymity and accountability. *European Journal of Social Psychology*, 37(2):239–257.
- Eugene F. Stone, Mark D. Spool, and Samuel Rabinowitz. 1977. Effects of anonymity and retaliatory potential on student evaluations of faculty performance. *Research in Higher Education*, 6(4):313–325.
- Gertjan van Noord. 1998. <http://odur.let.rug.nl/vannoord/TextCat/>. last verified 25 June, 2001.
- Carl Vogel, Ger Lynch, Erwan Moreau, Liliana Mamani Sanchez, and Phil Ritchie. 2013. Found in translation: Computational discovery of translation effects. *Translation Spaces*, 2(1):81–104.
- Carl Vogel. 2013. Multimodal conformity of expression between blog names and content. In Péter Baranyi, Anna Esposito, Mihoko Niitsuma, and Bjorn Sølvang, editors, *4th IEEE International Conference on Cognitive Informatics*, pages 23–28.
- William Lafi Youmans and Jillian C. York. 2012. Social media and the activist toolkit: User agreements, corporate interests, and the information infrastructure of modern social movements. *Journal of Communication*, 62(2):315–329.