

# Now You See It, Now You Don't! A Study of Content Modification Behavior in Facebook

Fuxiang Chen  
Singapore Management University  
Singapore  
fxchen@smu.edu.sg

Ee-Peng Lim  
Singapore Management University  
Singapore  
eplim@smu.edu.sg

## ABSTRACT

Social media, as a major platform to disseminate information, has changed the way users and communities contribute content. In this paper, we aim to study content modifications on public Facebook pages operated by news media, community groups, and bloggers. We also study the possible reasons behind them, and their effects on user interaction. We conducted a detailed study of Content Censorship (CC) and Content Edit (CE) in Facebook using a detailed longitudinal dataset consisting of 57 public Facebook pages over 3 weeks covering 145,955 posts and 9,379,200 comments. We detected many CC and CE activities between 28% and 56% of these pages (in both Facebook Posts and Comments). Manual judgements on these post/comment removals and edits show that majority of the content censorship is related to negative reports on events and personal grouses, and content edit is mainly performed to improve content quality and correctness. Furthermore, recency effect is also observed as part of Facebook content modification behavior.

## Keywords

Content modification; content censorship; content edit; user interaction; Facebook

## 1. INTRODUCTION

The success of social media is fueled by rich content contributed by large user population, including both individuals and organizations. Most research in the past focused on determining the content topics, sentiments and opinions in the social media content. The popular social media content are where users pay attention to are likely to be further shared by users. For example, Kwak, et. al, found that most tweeted topics are related to news [1]. Separately, Suh et. al reported that tweets have a higher chance of getting retweeted (shared) if the twitter user account has a larger number of followers and followees (more popular) [2].

Content posts and comments can be edited or removed. We therefore define two types of **content modification**,

namely: (a) *content censorship (CC)*, and (b) *content edit (CE)*. Content censorship refers to complete deletion of some content post or comment. Content edit refers to edits made to a content post or a comment. Content modification, like content contribution, may be performed for different reasons. Content modification can have major impact to user interaction. When not treated carefully, content modification can turn into disputes, controversies or even crises.

For example, in 2010, Nestle's Facebook page was flooded by criticism comments by environmentalists when the company attempted to delete comments on its Facebook page that carries altered Nestle's logos [3]. The saga suddenly turned Nestle's Facebook page into a platform for public outrage instead of for harnessing public goodwill.

In this paper, we examine the different types of content modifications that can be observed in 57 public Facebook pages over a period of three weeks (tracking period). These pages are selected as they are well known to offer interesting content and have been frequently visited by many users. We select Facebook as our target of study due to its immense popularity as well as its not-so-well understood content modification mechanism.

Past studies have focused on analysing the reasons behind content *self-censorship* which actually refers to the case of users not sharing any censored content at all. Our study in contrast focuses on censorship that comes *after* the censored content is shared. We also investigate into content edits. These two types of content modification have not been studied before although they can be similar to content self-censorship.

**Our contributions.** Our paper carefully classifies content modification into content censorship and content edit. We aim to understand content modifications on public Facebook pages and reasons behind these activities. We also compare and contrast related existing studies on Facebook content modification.

Our research approach consists of a data driven study which constructs the required Facebook dataset with longitudinal changes for analysing content modification. The dataset covers a three weeks tracking period from which removals and edits of posts and comments can be derived. Manual judgements on these post/comment removals and edits show that majority of the content censorship is related to negative reports on events and personal grouses. On the other hand, content edit is mainly performed to improve content quality and correctness. Furthermore, recency effect is also observed as part of Facebook content modification behavior.

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License. WWW'17 Companion, April 3–7, 2017, Perth, Australia. ACM 978-1-4503-4914-7/17/04. <http://dx.doi.org/10.1145/3041021.3054763>



## 2. RELATED WORK

Several studies have been conducted on censoring content in social media [4] [5] [6] [7] [8] [9] [10]. Chancellor et al. reported that Instagram users remove their own posts when they feel that they are inappropriate. Instagram was also reported to remove inappropriate content such as sexual photos, illegal behaviors, spam, self-harm and pro-eating disorders content. Zhu et al. analysed Weibo, a Chinese Twitter variant, and observed that much of the censorship in Weibo occur within the first 24 hours. They also reported that blacklisted keywords have been used to automatically remove content in the Chinese social media. Almuhammedi et al. reported that a lot of tweets have been deleted due to spelling errors, rephrasing and spam. Study has also shown that there are more deleted tweets than what Twitter has claimed [10].

There are much less content modification research on Facebook. Due to the heavy use of Facebook for sharing information, researchers have begun to study content changes in Facebook. Wang et al. [7] interviewed participants and studied their regrets on Facebook. They found that a comment could lead to controversial debate and severed online friendships. On the topic of self-censorship, Sleeper et al. [11] conducted user studies on 18 participants and observed that Facebook users are likely to self-censor entertainment content and content related to personal updates and opinions. Last-minute censorship refers to initial content creation by Facebook users but was never shared. Das et al. conducted an exploratory study on Facebook users examining last-minute self-censorship. They found that posts published on the user’s page are censored more than those on other pages, and that comments on photos are censored more than other types of comments.

Our work differs from existing studies in that we distinctly delineate content censorship and content edit within the Facebook environment. We focus on textual modifications of Facebook posts and comments. We also adopt a data driven study which include the acquisition of content snapshots over time and detection of different types of content changes

## 3. RESEARCH APPROACH

To understand content modification in Facebook, reasons behind these activities and their effects on user interaction, we first carefully assemble a set of Facebook posts and comments which allow us to detect changes in them. We then conduct analysis of the types of content modification.

### 3.1 Dataset

We selected 57 public pages of three different regions (Hong Kong, Singapore and United States) ranging from News to Community, Event and Group pages. All these pages contain content mainly in English language. These pages are selected as they are well known to offer interesting content and have been frequently visited by many users. We focus on detecting textual content modification of posts and comments created during the period from **1 January 2016 to 23 August 2016 (Study Period)**. These pages contain a total of 145,955 posts and 9,379,200 comments.

Specifically, to download the pages’ posts and comments, we made use of Facebook’s Graph API [12]. We downloaded posts and comments from the study period from the 57 public Facebook pages. As Facebook imposes rate limit mecha-

**Table 1: Number of CC & CE found from posts-/comments during the study & tracking periods**

Study Period	# of CC	# of CE
145,955 Posts	65 (0.04%)	463 (0.3%)
9,379,200 Comments	191,260 (2%)	4,573 (0.05%)
Tracking Period	# of CC	# of CE
9,828 Posts	36 (0.4%)	438 (4.5%)
749,559 Comments	72,343 (9.6%)	4,432 (0.6%)

**Table 2: Number of Facebook pages affected by CC & CE during the study and tracking periods**

Study Period	# of Pages (CC)	# of Pages (CE)
Posts	16 (28%)	28 (49%)
Comments	26 (46%)	23 (40%)
Tracking Period	# of Pages (CC)	# of Pages (CE)
Posts	18 (32%)	28 (49%)
Comments	32 (56%)	23 (40%)

nisms [13], we made special efforts in downloading the complete set of posts and comments (from the study period) from the 57 pages.

We also repeatedly download the pages’ posts and comments for a period of three weeks, from **8 August 2016 to 23 August 2016 (Tracking Period)**. Each cycle of downloading was scheduled to start as the previous cycle completed. Depending on the time taken in each downloading cycle, we were able to obtain every 1 to 5 hours a new version of the same Facebook page. This repeated downloading procedure obtains many versions of the Facebook pages. From the successive versions, we finally obtain modifications to the posts and comments in these pages.

### 3.2 Empirical Data Analysis

Every downloaded post/comment contains both an unique identifier and its content. To identify content modification of a Facebook post or comment, we need to compare both the identifier and the content from its two consecutive versions  $\langle P, C \rangle_t$  where  $P$  denotes the previous version and  $C$  denotes the current version.  $t$  is the time where we have downloaded  $C$ . This comparison is initiated in every cycle of download.

To detect a removed post/comment, we compare the  $P$  &  $C$ ’s identifiers. If  $P$  contains an identifier from a post/comment, and is missing in  $C$ , then the post/comment is deemed to be removed. When a comment is removed together with its associated post, we call it a *propagated comment removal*. Otherwise, it is an *organic comment removal*. There are a total of 3,822 propagated removed comments residing in the removed posts. This number is relatively small as compared to the number of organic removed comments from the tracking period (749,559) Table 1. Furthermore, we could not identify any differences between propagated removed comments and organic removed comments. Thus, in our study, we focus on organic comment removal.

Similarly, to detect an edited content in a post/comment, we compare both its identifier and its content. When we detect the same identifier in the  $\langle P, C \rangle$  pair, we compare their content to detect for changes. A change in the post/comment content can be a removal or replacement of some partial content, or it can be additional content.




Figure 1: Recency Effect in CE (Tracking Period)

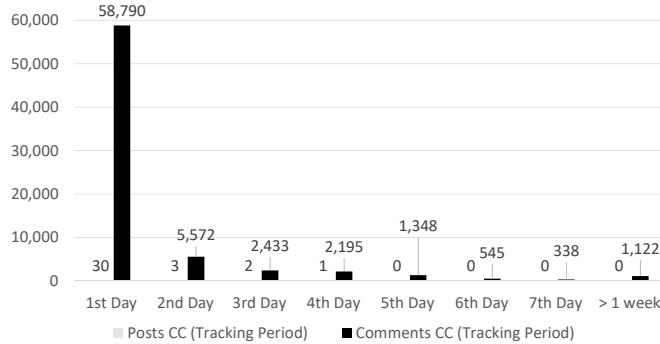


Figure 2: Recency Effect in CC (Tracking Period)

As shown in Table 1, a total of 145,955 posts and 9,379,200 comments were created from the 57 Facebook pages during the study period, and a total of 9,828 posts and 749,559 comments were created from the 57 Facebook pages during the tracking period.

We observed that posts are more likely to be edited than removed. We believed that Facebook users generally spend more time on crafting and writing a post, and thus the posts are not easily removed. In contrast, comments are more likely to be removed than edited. We believed that Facebook users generally spend much less time to write comments, and they are more likely to remove the comments, rather than to edit them.

Recency effect has been studied in different areas such as making use of recent tweets to improve search results on twitters and recent tags made to improve the accuracy of tweet recommendation [14] [15]. However, little study on recency effect is conducted in the context of content modification in Facebook. Thus, we investigate the recency effects in two aspects. We first analyse the content censorship and edit actions performed on posts and comments created during the tracking period. For each censorship and edit action, we determine the number of days between the content creation date and the action date. We then bin each action by its number of days, and count the number of censorship and edit actions in each bin. Figures 1 and 2 show the edit and censorship action count for different day bins respectively. Figure 1 shows that the number of posts and comments edits achieves the highest volume in the first day, and then decreases exponentially until the 7th day. Similarly, Figure 2 shows the same exponential decreasing trend in the number of posts and comment censorship. Thus, this suggests that




Figure 3: Content Edit Count on Pre-Tracking Period Posts/Comments

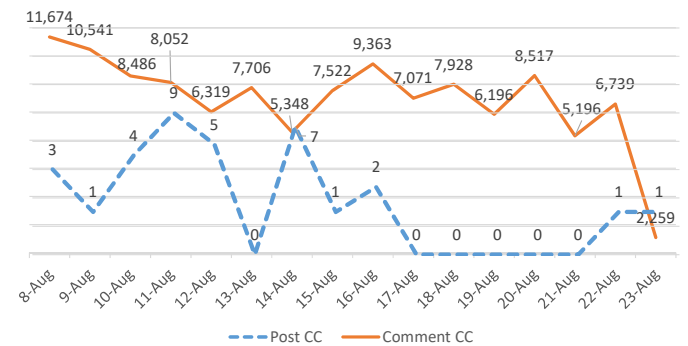


Figure 4: Content Censorship Count on Pre-Tracking Period Posts/Comments

users are more likely to perform content modification on the more recent posts/comments.

We next analyse the number of posts and comments during the **pre-tracking period (1 Jan - 7 Aug 2016)** getting edited and censored each day during the tracking period. We expect the amount of modification to decrease as time advances. Figures 3 and 4 show the rate of decrease for content edit and censorship respectively. Figure 3 shows that the numbers of posts and comments edits decrease sharply after the first day and these numbers do not significantly increase in the subsequent days. This is in line with our expectation, and we believe that recency effect plays a role here as users are more likely to perform edit operations on recent posts and comments. Figure 4 shows that the number of comments getting censored decreases over time. This matches our expectation. However, we observed that for posts censorship, there is no obvious decreasing pattern.

As shown in Table 2, we also noted that during the tracking period, 18 (32%) pages and 32 (56%) pages contained CC instances in posts and comments respectively, and 28 (49%) pages and 23 (40%) pages contained CE instances in posts and comments respectively. Similarly, we noted that during the study period, 16 (28%) pages and 26 (46%) pages contained CC instances in posts and comments respectively,

The five suspects, who from a terror cell, whose leader plotted to fire a rocket into Marina Bay, were brought to Jakarta under tight guard.

Figure 5: An example of a visualised edited content using Myer’s general purpose diff algorithm.

Table 3: Classification of Possible Reasons behind Content Edit

Types of Edited Posts	# of Instances
Supplementary Information	227 (52%)
Grammatical/Typographical Errors	174 (40%)
Substantial Removal of Content	37 (8%)
Types of Edited Comments	# of Instances
Grammatical/Typographical Errors	270 (68%)
Supplementary Information	125 (31%)
Substantial Removal of Content	5 (1%)

and 28 (49%) pages and 23 (40%) pages contained CE instances in posts and comments respectively. Overall, content modification can be seen affecting many pages.

## 4. CC AND CE ANNOTATION

### 4.1 CE Annotation

After detecting the edited posts and comments, we then seek to understand the reasons behind these post/comment modifications as a manual annotation task. We made use of the state-of-the-art Myer’s general purpose diff algorithm [16] to help an annotator visualise the changes in the edited posts/comments. Figure 5 shows an example of a modified content. The characters highlighted in red with strikethrough are the content that have been removed while the underlined characters highlighted in green are the newly added content.

With the visual highlights, we manually inspected all the changes to all the 438 edited posts from the tracking period (see Table 1) and categorized them. As there are too many edited comments, we randomly sampled 400 edited comments out of the 4,432 edited comments found during the tracking period. We manually categorized the different types of content edit reasons through an iterative annotate and label creation process. We systematically analyzed each of the visualised edited change and assigned one of the available category labels to the change. If the available labels do not apply, we create and assign a new category label. Finally, we have every edited posts and comments assigned with category labels. The final labels for edited posts and comments consists of three categories, namely: 1) Supplementary Information (users edit the posts/comments to include new content), 2) Grammatical/Typographical errors (users edit the posts/comments to fix language errors), and 3) Substantial Removal of Content (users edit the posts/comments to remove a large percentage of the content from the original version).

#### 4.1.1 Edited Posts

As shown in Table 3, we observed that 227 (52%) of the edited posts are due to addition of supplementary information. We found that many of these supplementary infor-

mation are website links to direct users to other web pages for more information. We also found that some of these supplementary information are updates to the original post content.

It is further observed that 174 (40%) of the edited posts are related to typographical or grammatical errors in the posts content by the post owners.

Furthermore, we noted that 37 (8%) of the posts have substantial content being removed and replaced. We observed that several of these removed content are supplementary information from previous versions of the post and we do not observe any controversial content in them.

#### 4.1.2 Edited Comments

As shown in Table 3, we observed that 270 (68%) of the edited comments are due to typographical or grammatical errors, and 125 (31%) of the edited comments contain supplementary information. The supplementary information includes newly added information, clarifications of previously created content, and expression of discontent. We also observed that only 5 (1%) of these edited comments have a large percentage of their original content removed. These removed content contained more offensive statements and was replaced by mildly toned statements.

## 4.2 CC Annotation

To understand the reasons behind the removed posts, we also manually assigned them with different types of reasons through a annotate and label creation process.

**Post Censorship** As the number of censored posts is small (36), we manually inspect them for annotation. We further separate the censored posts into two categories, *Deleted Posts not from Page Owner*, and *Deleted Posts from Page Owner*. A post is not contributed by the page owner if the post’s contributor id is different from the page owner’s user id. Otherwise, the post is contributed by the page owner. For deleted posts not from page owner, there are a total of three types of annotation: 1) *Frustration* (the content of the posts contain frustration vented by the user), 2) *Spam* (the content of the posts contain advertising information), and 3) *Personal Casual Remarks* (the content of the posts contain personal remarks made by the user). For deleted posts from page owner, there are a total of two types of annotation: 1) *Negative Reports on Events* (the content of the posts contain reports on adverse events such as kidnapping of children), and 2) *Controversial Remarks Messages* (the content of the posts contain controversial remarks on countries, races or politicians).

**Comment Censorship** We recruited five human annotators to perform annotation on 500 randomly sampled removed comments. These sampled removed comments are from existing posts, instead of removed posts. Each removed comment (together with it’s corresponding post) was read, analysed and annotated by two independent humans. The annotators were first trained for the first two hours before embarking on the annotation tasks individually. They were allowed to search for more information online to clarify the context of the removed comments.

A common set of predefined category labels was given to the human annotators as annotation options. This common set of predefined labels were derived by the first round of categorization of 100 random sampled removed comments. The predefined category labels are: 1) *Neutral Sentiments*

**Table 4: Annotated Deleted Posts**

Types of deleted Posts (not from Page Owners)	# of Instances
Frustration	7 (19%)
Spam	3 (8%)
Personal Casual Remarks	2 (5%)
Types of deleted Posts (from Page Owners)	# of Instances
Negative Reports on Events	17 (47%)
Controversial Remarks Messages	7 (19%)

**Table 5: Annotated Removed Comments**

Types of Removed Comments	# of Instances
Neutral Sentiments Comments	44 (38%)
Comments that Tagged Other Users	43 (37%)
Nuisance Comments	20 (17%)
Negative Sentiments Comments	10 (9%)

*Comments* (the content of the comments do not have any relationship to the post content, or that there is no observable topic of interest in the comment’s content), 2) *Comments That Tagged Other Users* (the content of the comments contain only other tagged users), 3) *Nuisance Comments* (spam content or the content of the comments is about asking others to add or follow the commenters), and 4) *Negative Sentiments* (the content of the comments contain more offensive statements made towards other users). The human annotators could select an appropriate label from the pre-defined list or suggest a new label. No time restriction is being imposed for each of the annotated removed comment as we do not want the annotators to rush through the tasks, which could jeopardised the output quality.

If the two human annotators could not assign a common label for a removed comment, we would rope in another independent human annotator to perform the same annotation on that comment. If we could not find two annotators agreeing a label for a removed comment, we would exclude the comment from our study. In the final list of category labels, there are no new category labels suggested by the annotators for removed comments, and it is the same as the initial version.

#### 4.2.1 Removed Posts

We manually inspected the 36 removed posts and observed that 12 (33%) of them are not posted by the page owners themselves.

For the deleted posts which are not published by the page owners, 7 (19%) of them are about users venting their frustration about their daily lives. We observed that 3 (8%) of them are spam messages and 2 (5%) of them are about personal casual remarks (e.g. The user posted a message saying he likes a particular video.).

For those page owners’ deleted posts, we observed that 17 (47%) of them (Table 4) cover negative events such as a competitor swimmer losing a competition and reported encounters of kidnapping of young children. Also, we noted that 7 (19%) of the posts contains controversial remarks about certain countries, races, or politicians.

#### 4.2.2 Removed Comments

Table 5 shows the deleted comments annotated by the five humans. A total of 117 out of 500 randomly sampled deleted comments were annotated and have unanimous agreement on their labels by the human annotators. We observed that 44 (38%) of the deleted comments contain neutral sentiments, the highest among the other three categories. Comments with neutral sentiments generally contain content that do not have any relationship to the post content, or that there is no observable topic of interest in the comment’s content. This may explain why they are removed. The second highest category, Comments that Tagged Other Users, constituted 43 (37%) of the sampled deleted comments. Comments in this category contain other tagged Facebook users accounts. The third highest deleted comments category, Nuisance Comments, involved 20 (17%) of the entire sampled population. These comments are marked as nuisance comments and are generally spam messages or messages asking others to add or follow the commenters. The last category belongs to comments that contain negative sentiments. This category constituted 10 (9%) of the entire sampled population. These comments generally contain offensive content aimed at attacking government, politician or other people. It also contained comments that excessively displayed discontentment on the post content.

## 5. DISCUSSIONS

### 5.1 User Attention in Content Modification

**Recency Effect.** As Facebook users perform content modification, they are more likely to focus on recent content: posts or comments. As shown in Tables 1, our data driven study saw most of the post censorship and edits performed on the recent posts and comments. For example, 4.5% of the posts created in the tracking period were edited during the same period compared with only 0.3% of the posts created in the entire study period getting edited during the tracking period.

**Multiple edits.** For the edited posts, we observed that some Facebook users made several rounds of edits to their posts. The changes can be an entirely new content added to supplement existing post, or removal of some previously added post content, in replacement of new content. However, in the case of edited comments, we found that the changes were mainly due to typographical or grammatical mistakes. This suggests that Facebook users generally are careful in writing. This explains some cases of multiple comment changes to fix previous typographical or grammatical mistakes.

### 5.2 Self-Censorship

*Self-Censorship* refers to the act of preventing oneself from content sharing. Usually, self-censorship leaves no trace of the content to be shared. In cases where content has been created but not shared, we call it the *Last-Minute Censorship* [8] [9] [7]. As last-minute censorship involves content that may have been entered in some text box but are subsequently abandoned, it is very similar to the post-sharing content censorship in this paper. In the following, we relate our findings with those of previous self-censorship and last-minute self-censorship studies on Facebook reported by Sleeper et al. and Das et al respectively.

### 5.2.1 Comparison with Self-Censorship on Facebook

Sleeper et al. [8] in their user study reported that users would self-censor more on content related to entertainment, personal updates and personal opinions. Many of the reported unshared content were negative in general. One of the reasons for censoring the personal updates is because the participant thinks that it might be too negative. Results from our data driven study concurred with the study report from Sleeper et al. Based on the annotated deleted posts from page owners, majority of them are related to negative reports on events.

### 5.2.2 Comparison with Last-Minute Self-Censorship on Facebook

Das et al. reported that 51% of the users performed last-minute self-censorship in at least one post while 44% of the users performed last-minute self-censorship in at least one comment. Results from our data driven study however suggested that comments are more frequently removed than posts although we do not know exactly if the comments were removed by the users themselves.

### 5.2.3 Comparison with Regrets on Facebook

Wang et al. [7], based on their user study, reported that participants regret posting three main types of content: (a) sensitive content (e.g., religion, politics, personal and family issues), (b) content with strong sentiment (e.g. negative or offensive comments), and (c) lies and secrets. Interestingly, we noted that some of the removed comments detected by our data driven study share several reasons of regrets reported by Wang et al. Specifically, among the removed comments, 115 (23%) of them are related to negative sentiments comments. These negative sentiments comments contain both negative and offensive remarks on religion and politicians.

## 6. CONCLUSION

This paper investigates into content modification activities in Facebook which are motivated by different reasons but complicated by the inter-connections among content pieces. This research turns out to be challenging due to unavailable data sets and a lack of user provided ground truth. In this work, we downloaded posts and comments from a set of Facebook public pages tracking their changes very closely over a three week period. We devised a data driven study which revealed the amount of post and comment removal and edit in the target pages which are usually not noticeable to Facebook users. The data driven study also showed the different reasons of post/comment removal/edits as annotated manually.

## 7. ACKNOWLEDGEMENTS

This research is supported by the National Research Foundation, Prime Minister's Office, Singapore under its International Research Centres in Singapore Funding Initiative.

## 8. REFERENCES

- [1] Haewoon Kwak, Changhyun Lee, Hosung Park, and Sue Moon. What is twitter, a social network or a news media? In *WWW*, pages 591–600, 2010.
- [2] Bongwon Suh, Lichan Hong, Peter Pirolli, and Ed H Chi. Want to be retweeted? large scale analytics on factors impacting retweet in twitter network. In *Social computing (socialcom), 2010 IEEE second international conference on*, pages 177–184. IEEE, 2010.
- [3] Caroline McCarthy. Nestle mess shows sticky side of facebook pages. CNet News, March 2010. Retrieved September 17, 2016 from <http://www.cnet.com/news/nestle-mess-shows-sticky-side-of-facebook-pages>.
- [4] Stevie Chancellor, Zhiyuan Jerry Lin, and Munmun De Choudhury. "this post will just get taken down": Characterizing removed pro-eating disorder social media content. In *CHI*, pages 1157–1162, 2016.
- [5] Tao Zhu, David Phipps, Adam Pridgen, Jedidiah R Crandall, and Dan S Wallach. The velocity of censorship: High-fidelity detection of microblog post deletions. In *USENIX Security 13*, pages 227–240, 2013.
- [6] Hazim Almuhammedi, Shomir Wilson, Bin Liu, Norman Sadeh, and Alessandro Acquisti. Tweets are forever: a large-scale quantitative analysis of deleted tweets. In *CSCW*, pages 897–908, 2013.
- [7] Yang Wang, Gregory Norcie, Saranga Komanduri, Alessandro Acquisti, Pedro Giovanni Leon, and Lorrie Faith Cranor. I regretted the minute i pressed share: A qualitative study of regrets on facebook. In *Proceedings of the Seventh Symposium on Usable Privacy and Security*, page 10, 2011.
- [8] Manya Sleeper, Rebecca Balebako, Sauvik Das, Amber Lynn McConahy, Jason Wiese, and Lorrie Faith Cranor. The post that wasn't: Exploring self-censorship on facebook. In *CSCW*, pages 793–802, 2013.
- [9] Sauvik Das and Adam D.I. Kramer. Self-censorship on facebook. In *ICWSM*, 2013.
- [10] Rima S Tanash, Zhouhan Chen, Tanmay Thakur, Dan S Wallach, and Devika Subramanian. Known unknowns: An analysis of twitter censorship in turkey. In *Proceedings of the 14th ACM Workshop on Privacy in the Electronic Society*, pages 11–20. ACM, 2015.
- [11] Manya Sleeper, Justin Cranshaw, Patrick Gage Kelley, Blase Ur, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. I read my twitter the next morning and was astonished: A conversational perspective on twitter regrets. In *CHI*, pages 3277–3286, 2013.
- [12] Facebook. Graph api - documentation - facebook for developers, 2016. Retrieved September 17, 2016 from <https://developers.facebook.com/docs/graph-api>.
- [13] Facebook. Rate limiting - graph api - documentation - facebook for developers, 2016. Retrieved September 17, 2016 from <https://developers.facebook.com/docs/graph-api/advanced/rate-limiting>.
- [14] Jimmy Lin and Miles Efron. Temporal relevance profiles for tweet search. In *SIGIR Workshop on Time-aware Information Access*, 2013.
- [15] Santiago Larrain, Christoph Trattner, Denis Parra, Eduardo Graells-Garrido, and Kjetil Nørnvåg. Good times bad times: A study on recency effects in collaborative filtering for social tagging. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 269–272. ACM, 2015.
- [16] Eugene W Myers. Ano (nd) difference algorithm and its variations. *Algorithmica*, 1(1-4):251–266, 1986.