

Shared Feelings: Understanding Facebook Reactions to Scholarly Articles

Curtis Cole Freeman
Northern Illinois University
DeKalb, Illinois
cole.freeman9@gmail.com

Michele Fattoruso
Northern Illinois University
DeKalb, Illinois
z1840898@students.niu.edu

Mrinal Kanti Roy
Northern Illinois University
DeKalb, Illinois
mkantiroy@niu.edu

Hamed Alhoori
Northern Illinois University
DeKalb, Illinois
alhoori@niu.edu

ABSTRACT

Research on social-media platforms has tended to rely on text-based tools to perform research tasks. While these tools have significantly increased our understanding of online behavior and social dynamics, they too often overlook features on these platforms that have grown in prominence in the past few years: click-based responses to content. In this paper, we present a new dataset of Facebook Reactions to scholarly content. We give an overview of its structure, analyze some of the statistical trends in the data, and use it to train and test two supervised learning algorithms. Our preliminary tests suggest the presence of stratification in the number of users following pages, divisions that seem to fall in line with distinctions in the subject matter of those pages.

CCS CONCEPTS

• **Information systems** → **Data mining; Social networks; Data analytics;**

KEYWORDS

Facebook Reactions, Altmetrics, Data Collection, Social Clicks, Research Community, Social Media, Supervised Learning

ACM Reference Format:

Curtis Cole Freeman, Mrinal Kanti Roy, Michele Fattoruso, and Hamed Alhoori. 2019. Shared Feelings: Understanding Facebook Reactions to Scholarly Articles. In *Proceedings of Joint Conference on Digital Libraries (ACM/IEEE)*. ACM, New York, NY, USA, Article 4, 4 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

As the prevalence of social media in the world around us increases and the number of users on these online platforms grows, so too grows the rate at which scholarly content is being proliferated and discussed in these venues. More and more, academics are finding it rewarding to look to these platforms for the insight they provide into research problems.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ACM/IEEE, June 2019, Urbana-Champaign, Illinois USA
© 2019 Copyright held by the owner/author(s).
ACM ISBN 123-4567-24-567/08/06.
https://doi.org/10.475/123_4

One reason scholars have turned to social media is to measure the influence their work is having in those spaces; this has become known as alternative metrics, or *altmetrics* [1]; another reason is for the knowledge online platforms provide about human behavior—an area of research known as *social-media analytics* [5]. Studies in social-media analytics tend to focus either on text, using tools such as Natural Language Processing (NLP), sentiment analysis, or opinion mining to arrive at and support research conclusions [8], or on the proliferation of content through online communities [4]. These approaches have proved effective for understanding or predicting many aspects of human behavior; but they leave a number of other expressive signals unexamined.

Click-based reactions, on the other hand, are a relatively under-utilized resource in social-media research. Examples of quick-draw, ready-made expressive features are becoming increasingly prevalent across many platforms, and as such have attracted some amount of attention from researchers in the past few years [2, 3].

In this paper, we present a new dataset of click-based reactions to scholarly articles on Facebook and use it to gain insight into how users are interacting with scholarly articles on that platform. In addition to information about the articles themselves, our dataset records the count of each click-based feature we could access through Facebook’s Graph API. We use our newly developed dataset to train and test two machine learning algorithms, and our analysis of the results shines some light on surprising relationships between features.

2 BUILDING THE DATASET

Before going any further, it will be useful to define a few terms and features:

- **Click-based reactions** - non-textual user interactions with shared content—sometimes referred to simply as *reactions*; includes Facebook Likes and Reactions, Re-shares, and Page visibility (definitions for these last two are below).
- **Reactions** - the five click-based reactions: Love, Amazed, Laughing, Sad, and Angry; will be distinguished from the common term “reaction” by capitalization.
- **Page visibility** - the number of followers a Facebook page has.
- **Re-shares** - the number of times users have re-shared a public post of an article into another location.

The roots of our dataset lie in the online resource Altmetric¹, which tracks the impact scholarly articles have across a variety of social media platforms. We used Altmetric as a “jumping-off point”, querying their API² for information on articles we were interested in and for the public pages onto which they had been shared. It gave us access to the titles, publication dates, subjects, and the URLs of Facebook shares for nearly 1.5 million scholarly articles.

We targeted reactions on Facebook as opposed to other social-media platforms because of the variety of click-based interactions available to users on that platform. Facebook’s enormous population of active users was another reason we directed our attention there: it has about 2.27 billion active monthly users, while Twitter only has roughly 67 million. With our list of Facebook URLs for article shares, we queried Facebook’s Graph API³ for the reaction counts on each post. Our dataset records their responses, and was collected during the period of December 1-13, 2018. Constraints in the number of queries allowed by Facebook’s API (200/hour) determined the rate at which we could work. The resulting dataset is publicly available on OSF⁴ as a comma-separated-value file (CSV).

We limited our collection efforts only toward scholarly articles published in 2017. Choosing this year accomplished three goals: (i) Reactions were released by Facebook in February 2016 [6], so any articles we looked at had to be published after that time to have meaningful data on this feature; (ii) any time a new feature is unrolled, it takes some amount of time for users to learn how to use it; Prah [7] finds that use of Reactions increased from 2.4% of all interactions in April 2016 to 5.8% by June 2016, and up to 12.8% of all interactions by June 2018; by the time of our data collection in December 2018 a large enough subset of users were comfortable expressing themselves with the feature to warrant more scholarly attention; and (iii) by the time we began our data collection, a sufficient interval of time had passed for articles to be widely shared and reacted to (between 11 and 23 months).

Of all the articles tracked by Altmetric, we found 296,052 were published in 2017 and had been shared on Facebook at least once. We eliminated entries that were missing data on the pages to which the articles had been shared; this reduced our set to 135,635 articles. We further limited the scope by selecting only articles with Scopus⁵ subjects in the scientific domain. We chose to focus only on articles in the Health Sciences, Physical Sciences, Social Sciences, and Environmental Science. Figure 1 shows that these four categories, article counts fall within one standard deviation of the mean number of articles, as do the total number of Facebook shares (Health sciences is the only exception, exceeding one standard deviation greater than the mean of article counts). It also displays the full list of subjects in all the 2017 articles and gives a sense of their distribution. The mean and two standard deviations are indicated there with blue lines for both axes, and the four subjects we target are indicated with arrows in the plot. Limiting the scope of subjects reduced the number of articles needed to process to just over 31,000. When we dropped articles that had missing records for features




Figure 1: Count of articles published in 2017 categorized by subject plotted against number of shares. Both features are shown on logarithmic axes.

such as abstract and title, we were left with 11,474 articles: these are the articles recorded in our dataset.

In our data collection process, we took the utmost care to respect Altmetric’s and Facebook’s specifications for *how* and *why* their data can be accessed and used and to protect the personal information of social-media users. Our interests are only in the ways that people are interacting *in the aggregate* with scholarly content on social media platforms—not in specific ways users’ beliefs or opinions may influence their behavior. We recognize that identifying information could in some instances be inferred *a posteriori* from some of the data we collect; however, our method of data collection does not target anything that could be used to consistently identify individual users and avoids collecting identifying information about individuals.

3 DATA EXPLORATION

The click-based features of our dataset are displayed along the axes in Figure 2; also displayed are the Pearson r correlation coefficients for all feature pairs. Highly correlated pairs are: Like and Love ($r = 0.82$), Sad and Angry ($r = 0.81$), Like and Amazed ($r = 0.77$), Love and Re-shares ($r = 0.71$). We can infer that high positive correlation is a sign that users employ features in similar contexts, and that the emotional expressions represented by those features overlap. For example, a Like seems to have a meaning comparable to a Love or (to a lesser extent) an Amazed, or (to an even lesser extent) a Laughing reaction. These relationships may not surprise us because they are all positive emotional states; but other feature

¹<https://www.altmetric.com>.

²<http://api.altmetric.com/>.

³<https://developers.facebook.com/docs/graph-api/>.

⁴<https://osf.io/4kh7r/>.

⁵<https://www.scopus.com/>.

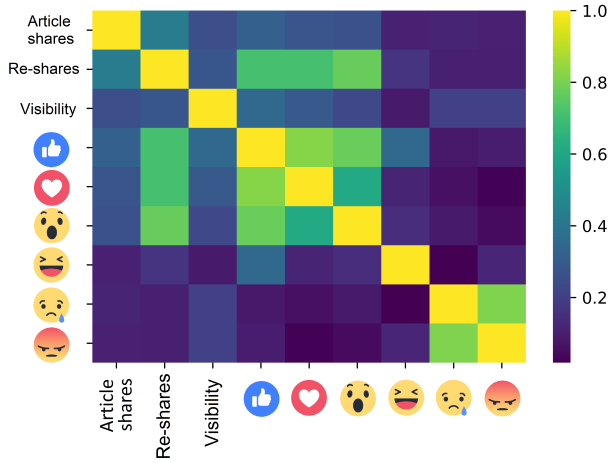


Figure 2: Correlation matrix of features using Pearson’s r .

pairs that have related expressive values in usage, such as Angry and Sad reactions, are not so intuitive.

Likes and Re-shares are correlated with the most other features; this might be explained by the fact that these two are the oldest reactions—but we also notice they are correlated with other *emotionally positive* reactions such as Love or Amazed and not with the negative emotions Sad or Angry. It follows that by Liking or Re-sharing a post, a user expresses a positive emotional reaction to its content. Looking at this from another angle, we infer that content that is more likely to inspire a negative reaction from users is less likely to be Re-shared or Liked.

High correlation between features can lead to increased variance in model results. To counter this, modelers often eliminate one of a pair of correlated features. Rather than removing features and losing data in our sparse dataset, we combined Love/Wow and Sad/Angry Reactions into two new composite features for our models.

Low correlation signifies that features have relatively distinct use values. Among the lowest r coefficients are Love/Angry ($r = 0.145$) and Laughing/Sad ($r = 0.224$); this makes intuitive sense, as these reactions nominally encompass opposite emotions. Laughing/Page visibility ($r = .229$) is another low-correlation pair, suggesting that articles that inspire humor are more likely to be posted to public pages with relatively low follower counts. It is likely that this relationship may be a result of our choice to limit the articles we include to those in the scientific domain, where humor is an under-utilized affect.

Our dataset also contains outliers in each feature category; to correct for these, we re-scaled the features to a range from 0 to 1, then took the cube root to these new values. Our root normalization function is demonstrated in Equation 1; it helped to smooth the distribution of values, raising the lower values by more than it raised the higher values. The result after combination/normalization is displayed in Figure 3.

$$rt_norm(F_i) = \sqrt[3]{\frac{F_i - F_{min}}{F_{max} - F_{min}}} \quad (1)$$




Figure 3: Feature values after transformation displayed by subject. Values only go up to 0.25 to better display the IQRs.

Set	Features	Model Accuracy/AUC
A	Likes, Re-shares, Love/Amazed, Laughing, Sad/Angry	Random Forest: 38.1% / 67% Decision Tree: 40.97% / 65%
B	Visibility, Love/Amazed, Laughing, Sad/Angry	Random Forest: 64.21% / 82% Decision Tree: 65.66% / 77%

Table 1: The two feature sets used in our classification models, with model accuracies and Areas Under the Curve (AUC).

Even after transformation, our dataset is still sparse—zeros are unchanged by the transformation; yet features with greater variance, such as Visibility or Likes, have less spread between the IQR and outliers. The median value of all Reactions is zero, and non-zero values in those features are all in the fourth quartile. Likes have the largest interquartile range (IQR), though the median is still close to zero. Page visibility and Likes have the highest median values of all features.

4 SUPERVISED-LEARNING MODELS

To explore the relationships in our dataset further, we isolated two feature subsets and trained two supervised learning classification algorithms with them. We used Decision Tree and Random Forest algorithms because of the insight they provide into the relationships between features, and our feature sets are detailed in Table 1. We were interested in gaining insight into the extent to which users’ interactions could be related to articles’ subject matter; and so we selected article subjects as the class labels for our models. This gave us four targets for our multiclass classification models to predict.

With the first set (A) we were interested in finding the extent to which click-based reactions that are immediately available to users on the post itself could be used to estimate an article’s subject. The second set (B) provides insight into how extended click-based




Figure 4: Various metrics used to score our two ML algorithms applied to the two feature sets. Baseline represents the results of random guesses at article subjects.

features such as Page Visibility can be used to approximate the subject matter of posts.

Table 1 records the accuracy and Area Under the Curve (AUC) of our models, and Figure 4 displays the results of our models using several different metrics. For reference, scores are shown against the baseline, which represents random guesses at which of the four class labels an article belongs to. Feature set B produced significantly better results than A with both algorithms. Average accuracy of models with feature set B is 160% greater than the baseline, while feature set A is only 58% greater.

Figure 5 shows the relative importance of each feature in our models. In feature set A, Likes have the greatest weight, accounting for 51% of the result on average between the two algorithms; the weight of Re-shares is the second highest importance, accounting for an average of 27% of the result. In feature set B, Visibility is the most important feature accounting for an average of 94% of the result; the remainder of the weight is spread relatively evenly among the remaining features.

5 DISCUSSION AND CONCLUSION

Our new dataset of click-based reactions to scholarly content on Facebook offers a wealth of possibilities for researchers interested in social media analytics. We have demonstrated how it can be used in the exploration of user interactions with scholarly content on Facebook, and how click-based reactions are an effective data source for investigating indicators of user emotional attitudes.

Results from the models trained and tested on our dataset suggest that the number of followers a page has (Visibility) may be predictive of article subject matter; this indicates that there may be patterns in the content shared on Facebook pages and the number of followers these pages have. It may prove useful for researchers to explore the ways in which Facebook page popularity is stratified by the type of content displayed on their pages.

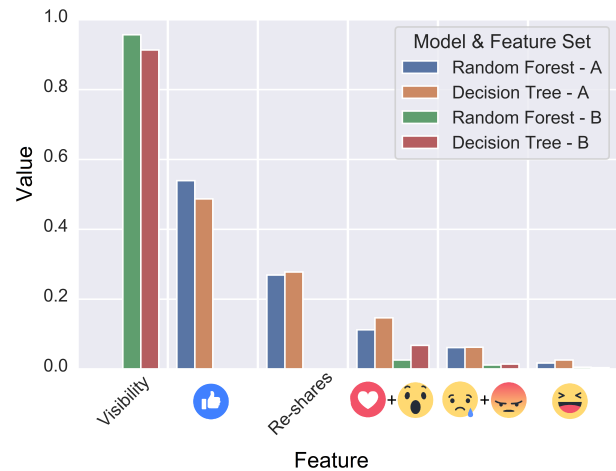


Figure 5: Feature weight for the four implemented models (two algorithms \times two feature sets).

We have also suggested some interpretation of Facebook click-based reactions that are not immediately apparent, notably that Re-shares convey an emotionally positive feelings toward content, and that Sad and Angry Reactions express similar affects. These relationships are not at all obvious, and give us insight into how these features are being used in practice.

REFERENCES

- [1] Hamed Alhoori and Richard Furuta. 2014. Do Altmetrics Follow the Crowd or Does the Crowd Follow Altmetrics?. In *Proceedings of the 14th ACM/IEEE-CS Joint Conference on Digital Libraries (JCDL '14)*. 375–378.
- [2] Angelo Basile, Tommaso Caselli, and Malvina Nissim. 2017. Predicting Controversial News Using Facebook Reactions. In *Proceedings of the Forth Italian Conference on Computational Linguistics (CLiC-it)*. 28–33.
- [3] Malvina Nissim Chris Pool. 2016. Distant supervision for emotion detection using Facebook reactions. *PEOPLES@COLING (2016)*, 30–39.
- [4] Maksym Gabielkov, Arthi Ramachandran, Augustin Chaintreau, and Arnaud Legout. 2016. Social Clicks: What and Who Gets Read on Twitter?. In *ACM SIGMETRICS / IFIP Performance 2016*. Antibes Juan-les-Pins, France.
- [5] Florian Krebs, Bruno Lubascher, Tobias Moers, Pieter Schaap, and Gerasimos Spanakis. 2018. Social Emotion Mining Techniques for Facebook Posts Reaction Prediction. *ICAART 2 (2018)*, 211–220.
- [6] Sammi Krug. 2016. Reactions Now Available Globally. <https://newsroom.fb.com/news/2016/02/reactions-now-available-globally/>. Accessed: Nov. 30, 2018.
- [7] Pritam Shah. 2018. Facebook's new Reactions are being used more - a lot more. <https://www.quintly.com/blog/new-facebook-reaction-study>. Accessed: Jan. 24, 2019.
- [8] Ye Tian, Thiago Galery, Giulio Dulcinati, Emilia Molimpakis, and Chao Sun. 2017. Facebook sentiment: Reactions and Emojis. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*. Association for Computational Linguistics, 11–16.