SANTA CLARA UNIVERSITY

DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

Date: June 9, 2021

I HEREBY RECOMMEND THAT THE THESIS PREPARED UNDER MY SUPERVISION BY

I Chang
Orion Sun
Jasper Sang Ahn

ENTITLED

Is Social Diversity Related to Misinformation Resistance? An Empirical Study on Social Communities

BE ACCEPTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

BACHELOR OF SCIENCE IN COMPUTER SCIENCE AND ENGINEERING

Yuhong Liu

Thesis Advisor

Nam Ling
Nam Ling (Jun 10, 2021 10:29 PDT)

Department Chair

Is Social Diversity Related to Misinformation Resistance? An Empirical Study on Social Communities

by

I Chang
Orion Sun
Jasper Sang Ahn

Submitted in partial fulfillment of the requirements
for the degree of
Bachelor of Science in Computer Science and Engineering
School of Engineering
Santa Clara University

Santa Clara, California
June 9, 2021

Is Social Diversity Related to Misinformation Resistance? An Empirical Study on Social Communities

I Chang
Orion Sun
Jasper Sang Ahn

Department of Computer Science and Engineering
Santa Clara University
June 9, 2021

ABSTRACT

Misinformation has become a pervasive issue on online social media. Whether it is spread purposefully with ill intent or accidentally through ignorance, misinformation can be dangerous and create confusion in those who are affected by it. The investigation presented in this paper found that no research has been performed that directly examine the correlation between a social community's misinformation resistance and diversity. This project utilizes CrowdTangle, a tool to gather data from Facebook groups, along with machine learning models to determine if this correlation can be drawn. We were able to find correlations between some diversity metrics with misinformation resistance and discuss potential rationales for such correlations.

Table of Contents

List of Figures

Chapter 1

Introduction

1.1 Motivation

Misinformation is a pervasive issue on social media. In the context of the current day, misinformation spreads false and harmful ideas regarding topics such as the 2020 election and the COVID-19 virus and vaccine [1]. This misinformation can result in drastic consequences such as the loss of human lives. The propagation of misinformation must be prevented as much as possible — especially in social media environments where it is able to thrive and be consumed by a significant portion of the population.

Virtually every person is susceptible to misinformation to a certain degree [2]. However, in a social media environment, groups of individuals exhibit highly diverse levels of resistance to misinformation. This suggests that there may be characteristics of group environments that contribute to a greater level of misinformation propagation within the group. If these characteristics are examined and a significant link to misinformation resistance is found, it would provide useful insight into the spread of misinformation and how we can potentially work to reduce this spread.

1.2 Solution

The goal of this project is to determine how likely a community (i.e. group) will be susceptible to misinformation based on its behavior. The core methodology for achieving this goal is to propose diversity metrics for social media groups and to determine which forms of diversity is related to the group's resistance level to misinformation. The belief is that a group with more diversity would tend to consume a greater variety of information and is more frequently faced with the problem of distinguishing reliable sources from unreliable ones. Groups which do not face such an issue would be less likely to be able to identify misinformation and as a result be less resistant to misinformation. Additionally, it is important to note that diversity comes in many forms. This research takes the opportunity to develop a variety of diversity metrics based on group behavior.

The following are the major contributions presented in this paper:

1. We propose a set of quantitative metrics to capture the several types of diversity which can exist within a social media group. The proposed quantitative metrics relate the concept of entropy with diversity and include a novel network-based metric which we refer to as Mutual Network Analysis (MNA).

2. We perform analysis to connect existing credibility scores of news sources with posts shared by social media groups in order to establish a ground truth data set containing a list of social media groups and quantitative values of their misinformation resistance.

3. We utilize machine learning algorithms and discover significant correlation between the proposed set of diversity metrics and group resistances to misinformation.

Chapter 2

Related Works

The investigation into the background of this project includes a few major components: approaches to quantifying diversity, sources for establishing ground truth, and methods for analyzing the collected data. Additionally, some of the works we examined provided inspiration for additional objectives regarding possible future extensions of this project and the potential applications for the conclusions.

This project requires a method of measuring diversity quantitatively and consistently across the analysis of different types of features of social media groups. Prior research [3, 4] has provided groundwork which establishes information entropy as a suitable way of measuring the variety of any particular characteristic. The research also investigates characteristics of social media information with different contexts such as political bias and information relevance, which are features we can use to examine in relation to misinformation resistance.

For other potential measures of diversity, Zheng et al. [5] presents an analysis of the posting patterns and vocabulary variety to identify anomalous Twitter users, providing an indication that examining vocabulary diversity through entropy scores is a method that can potentially distinguish groups of online users with distinct behavioral patterns. Harell et al. [6] presents a discussion of a metric relating to network diversity which serves as an inspiration for a proposed metric which we named Mutual Network Analysis.

To evaluate the metrics we generate, we need to compare them with an established ground truth of social media groups and their misinformation resistance. Prior research [7, 8] have proposed methods for measuring social media credibility through analysis of user interaction patterns and user perception of credibility. In contrast, we will take existing factual ratings of news sources and track social media groups which have shared the news sources. In another research related to identifying source credibility, Broniatowski et al. [9] used a data set of news source credibility found from the website "Media Bias/Fact Check" [10], which we will also make use of to establish social media groups with misinformation resistance scores.

For the method of analyzing the correlation between misinformation resistance scores and metrics, we decided to take on several machine learning algorithms. Several research papers [11, 12] have been able to take advantage of machine learning to recognize human patterns in social media. We also can look at Menczer [13] whose research specifically was designed to use machine learning algorithms to detect misinformation spread through fact checking to create a theoretical model for finding truth.

Finally, it is important to be mindful of the current limitations and the possible future extensions of this research. This paper primarily focuses on the consumption of misinformation from known news sources. Kulshrestha et al. [14] indicates that the production, recommendation, and spread of misinformation are also significant contributors to the misinformation resistance of social media groups. Another thing of note is that most research also treats all types of misinformation with homogeneity while Brennen et al. [15] suggests that placing misinformation into categories such as 'fabricated vs. reconfigured' and 'bottom-up vs. top-down' can be beneficial for analysis. Previous research [16] has suggested that the common beliefs regarding the idea of echo chambers and polarization on social media could be a misconception, meaning that we should be cautious of any preconceptions regarding this topic in current and future work.

Chapter 3

Methodology

3.1 Data Collection

Throughout the length of this research, the CrowdTangle API — a tool provided by Facebook which is built for analysing the virality and trends of social media content — was used for data collection. This tool provides several features for data analysis which was used to fetch the post information of Facebook groups and the Facebook posts with specific link mentions. Specifically, the information can be retrieved by sending requests to endpoints. The following is an elaboration of the parameters and data associated with the endpoints used: */links* and */posts*.

The */links* endpoint accepts a link parameter and returns a set of Facebook posts which share the link:

- If the link given is a domain name, the request will return posts that share links which are equally specific or more specific (e.g. requesting *www.cnbc.com/* will return posts mentioning the link *www.cnbc.com/2021/03/31/biden-infrastructure-plan* as well).

- For this project, each */links* request retrieved at most 1000 posts made on or after November 1st, 2019, prioritizing posts with the most interactions. These constraints were set around the limitations of the CrowdTangle API. However, the constraints also ensured that the data collected will be timely and relevant to relatively recent events.

- This endpoint returns information about posts from Facebook profiles and pages in addition to Facebook groups. As this research is focused around the behavior of groups, the information of posts from non-group sources were discarded after the request, meaning that there may be less than 1000 posts collected for each request.

- The endpoint returns information for each Facebook post such as post text, referenced links, and engagement (likes, comments, heart reactions, etc). It is important to note that the CrowdTangle API places an emphasis on user privacy and does not allow data regarding specific users to be collected. The impact of this restriction is a limitation on the diversity metrics generated in the future sections of this paper: the diversity of group characteristics such as gender and race distribution cannot be measured.

The */posts* endpoint accepts a list of Facebook groups and returns a set of Facebook posts from those groups:

- For this project, this endpoint was used to fetch the most interacted posts from a single group. This was done by using a feature of the endpoint to filter by specific group IDs in a manually created list of Facebook groups.
- Additionally, for this project, each */posts* request retrieved at most 100 posts made on or after June 1st, 2020.
- The returned information for each post is the same as the information returned by the */links* endpoint.

These two endpoints were used throughout the project to retrieve data for establishing the ground truth and for generating values for our diversity metrics.

3.2 Establishing Ground Truth

A ground truth is necessary for evaluating the characteristics of social media groups. In this research, the ground truth is a set of Facebook groups with predefined values of misinformation resistance. To establish this ground truth, data from the source "Media Bias/Fact Check" (MBFC) was used. MBFC contains several lists of news sources with a manual review of each of the news source's overall factual accuracy. The review contains a categorical factual reporting score with the possible values of "very high", "high", "mostly factual", "mixed", "low", and "very low". For this research, these categories were converted into numerical factual accuracy scores with the values of 100, 80, 60, 40, 20, and 0, respectively. For the base data, we created a list containing the domain name of every news source covered by MBFC along with the converted factual accuracy score. In total, there are around 3000 pairs of domain names and factual accuracy scores.

Next, a */link* request was performed for every domain name of every news source in the collected list. This process retrieved the top 1000 post mentions for every news source in the list. Afterward, the occurrence of all Facebook groups associated with each retrieved post was recorded and the most commonly occurring groups across the entire list of news sources was found. Each of these groups were given a 'misinformation resistance score' equal to the average factual accuracy score of all the website domains it is associated with (For example, if a group is found to have posted a link with 60 factual rating and another link with 80 factual rating, the group is given a misinformation resistance score of 70). This process was used to generate misinformation resistance scores since a group which propagates more factually inaccurate information should be more likely to consume misinformation.

Additionally, if a group posted information from a single website multiple times, the score of the website was only counted once. The goal of this exclusivity is to lower the data count for Facebook groups made by news sources that only post articles from their own news source, which effectively function as Facebook pages and are not made for public discourse in a group environment.

In the end, the groups which are associated with the highest amount of MBFC links are chosen as ground truth data. For this project, groups with 50 or more associated links were selected, which resulted a ground truth data set of 719 Facebook groups with misinformation resistance scores.

3.3 General Analysis Metrics

After establishing the ground truth of a set of Facebook groups and their misinformation resistance, quantitative descriptions of each group's features can be developed in order to further analyze how the features are related to the misinformation resistance. This section provides a description of the metrics generated for each group and the generation methods.

All the metrics generated use data collected using the */posts* endpoint, meaning that the top 100 most interacted posts of each Facebook group were analyzed. For clarification, the set of metrics were measured once for each group in the ground truth data for a total of 719 sets of quantitative values. The metrics we generate can be split into two groups, general metrics and diversity metrics:

3.3.1 General Metrics

General metrics are metrics collected which do not directly relate to diversity. Despite this, these metrics are kept for the possibility that they may provide more insight into the characteristics of misinformation resistance. The general metrics collected are as follows:

- **Share and Comment Count:** Users can share and comment on Facebook posts. The number of shares and comments were totalled individually across all posts. The result is recorded as two separate metrics.
- **Total Reactions:** Users can also engage with Facebook posts through reactions (like, love, wow, haha, sad, angry, thankful, care). Total reactions is the sum of the count of each reaction. Each user may only use one reaction, meaning that this metric measures the total number of reacting users.
- **Reaction Proportions:** The amount of each of the reactions were counted across all posts and divided by the number of total reactions. The result is recorded as eight separate metrics, one for the proportion of each reaction.
- **Top Level Domain (TLD) Proportion:** For each link in the posts of a group, the top level domain is recorded for proportion calculation if it is one of the seven original top level domains (com, org, net, int, edu, gov, mil). The result is recorded seven separate metrics.

3.3.2 Diversity Metrics

All diversity metrics used entropy calculations in order to quantitatively describe the variance of the examined data and describe its diversity. The entropy value for any set of data was calculated by placing the data into categories. The entropy value was then generated with the following formula, where n represents the number of categories and each p_i value represents the proportion of the number of items in each category i:

$$Entropy = -(\sum(p_i * log(p_i))/log(n)) \tag{3.1}$$

The formula converts the data into a diversity value ranging from 0 to 1, with 0 being the least diverse and 1 the most. If an examined metric does not have a predefined amount of categories, the amount of categories is assumed to be the number of unique items collected for the metric. The diversity metrics collected are as follows:

- **Word Entropy:** Facebook posts often contain a message from the poster. For every analyzed post with a message, the included message was parsed and processed such that it is reduced to a list of words in its base form (for example, words such as 'running' and 'ran' in the original message would be transformed into the base form 'run'). The proportion of the appearance of every word across all posts was calculated to generate the overall entropy value. In the context of entropy equation, the n value would be the number of unique words while each p value would be their proportions. Thus, the max entropy is achieved if every word in all posts are different.

- **Post Type Entropy:** CrowdTangle categorizes Facebook posts into 12 types (album, igtv, link, live_video, live_video_complete, live_video_scheduled, native_video, photo, status, video, vine, youtube). The frequency and type of all of the sampled posts are collected and used to generate entropy over 12 categories.

- **Average/Total Reaction Entropy:** As previously mentioned, users can add reactions to Facebook posts to express their opinion. Each of the reaction types was counted as a category. The total reaction entropy was generated using the total reaction counts across all posts. The average reaction entropy calculates the average of the reaction entropy values of each individual post.

- **Domain Entropy/Top Level Domain Entropy:** Facebook posts have the option to contain a direct link. If any post contains a link, the domain name (ex. cnbc.com) and top level domains (.com, .org) are stored. Both entropy values are calculated using as many categories as the amount of unique entries.

- **Mutual Network Analysis (MNA) Entropy:** The mutual network analysis entropy is generated by analyzing other Facebook groups which share the same links as the target group. For each of the posts in the target group which contains a link, the */links* endpoint was used to find other groups sharing the link. Each group found was recorded in a list and the resulting entropy value of the list is the MNA entropy. The highest MNA entropy is achieved if no individual group has shared 2 or more links shared by the target group.

3.4 Machine Learning

The collective of all of the collected general metrics and calculated diversity metrics were used as independent variables for the machine learning algorithms. The calculated misinformation resistance, derived from the ground truth scores, was the dependent variable for the machine learning algorithms. For training, k-fold cross validation was utilized to establish training and testing subsets from the overall data [17].

As for the specific machine learning model, any algorithm can be utilized as long as it can take in any number of non-normalized independent variables and output a binary classification for the misinformation resistance. The rationale behind this is to focus more on the data and recognize if a pattern could be gathered from the data set rather than the specific machine learning models. Thus we leave the details of the method open for further research opportunities in different machine learning models on this data.

3.5 Statistical Analysis

As a rudimentary test to see whether the metrics have any statistical significance when correlated with a group's misinformation resistance score, we ran a t-test, a commonly used statistical test for hypothesis testing in machine learning. By running the t-test on each metric, the probability of a null hypothesis being false was calculated. The null hypothesis used in this implementation was that changes in a certain metric have no effect on the outcome of the misinformation resistance score. A list of the metrics which were statistically significant was compiled and used as a reference when implementing various machine learning algorithms [18].

Chapter 4

Results

4.1 Data Characteristics

After collecting the ground truth and metric data, a few noteworthy patterns emerged. Regarding the misinformation resistance scores collected for the 719 groups in the ground truth data, a significant clustering behavior can observed as shown in the following histogram.

Instead of a normal distribution of misinformation resistance scores, there are two peaks which suggests that there are Facebook groups with at least two distinct types of behaviors. A possible explanation of this pattern may be that there exists a feedback loop effect which causes groups with less misinformation resistance to become even less resistant over time and vice versa.

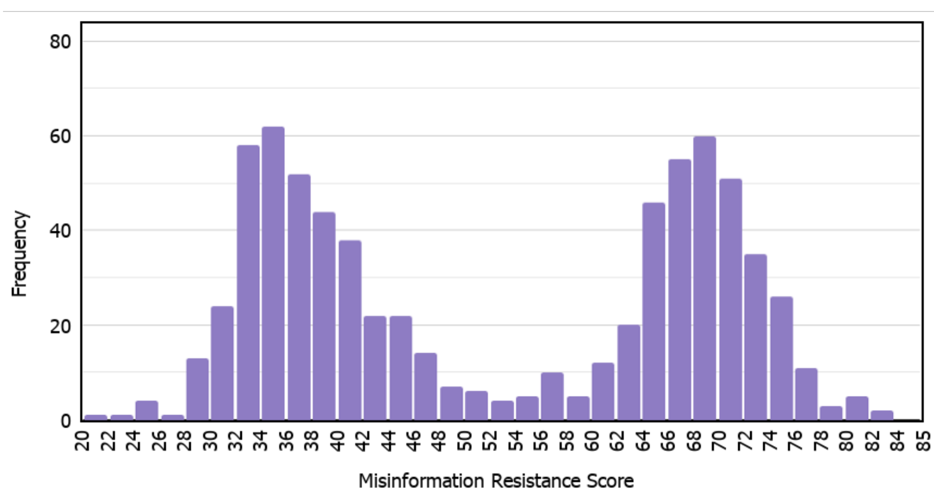

Figure 4.1: Histogram of Misinformation Resistance Scores across 719 Ground Truth Groups

Regarding the collected diversity metrics, the following is table of the range of values measured for each of the metrics (which are all entropy values). The lower and upper bound capture 95% of the measured values to exclude outliers.

Diversity Metric	Lower Bound	Upper Bound
Word Entropy	0.77	0.91
Post Type Entropy	0.00	0.00
Average Reaction Entropy	0.30	0.44
Total Reaction Entropy	0.50	0.73
Top Level Domain Entropy	0.17	0.69
Domain Entropy	0.67	0.90
MNA Entropy	0.84	0.95

Table 4.1: Table of Ranges of 95% Diversity Metric Values

This table provides an overview of the diversity metric values. The resulting post type entropy for most of the ground truth groups were 0 as most of the groups exclusively shared posts which were classified as link posts. This is likely due to the fact that the ground truth generation method was aimed at finding groups which share information through news sources — an action typically done by sharing links. As the focus of this paper is around the propagation and consumption of information in these types of environments, such a result is appropriate.

4.2 Machine Learning Model Comparison

For testing, we specifically utilized six different machine learning models implemented in Python utilizing the Scikit-Learn libraries. Three were classification models and three were regression models. Classification models are designed to sort data points into groups and allow the model to predict which group an entry will fall under based on the values of the features. Classification is useful for this research as being able to determine whether a group is resistant to misinformation or not would be a classification problem. Regression models output a continuous value from a range. This was useful in our data set as the computed misinformation resistance is a continuous value [19, 20, 21].

The classification models we utilized were:

- Decision Tree Classifier
- Ridge Regression Classifier
- Logistic Regression Classifier

The regression models we utilized were:

- Decision Tree Regressor
- Ridge Regressor
- Lasso Regressor

We worked on ensuring that each model was able to output a binary classification in order to normalize the results of the machine learning models for comparison. To accomplish this, we needed to develop a baseline classification for the output data. The basic classification methodology was to split the misinformation resistance score, which had a range of [0,100] into resistant groups (classified by a 1) and nonresistant groups (classified by a 0). When statistically analyzing the data, we noticed that clear groups had formed with a low percentage of groups having a misinformation resistance score around 50, so the data was split at a score of 50. If a group had a misinformation score above 50, they would be classified as a resistant group. If a group had a misinformation score below or equal to 50, they would be classified as a nonresistant group.

The above classification was run on the data set at different steps in the procedure depending on if it was a classification model or a regression model. On classification models, the groups were classified prior to running the model on the data. On regression models the groups were classified after running the model on the data.

For comparing the output data between each of these machine learning models, the F-1 score of the confusion matrix was utilized. The F-1 scores we calculated are subject to small variance (±0.05) due to the random nature of some of the machine learning models. The table below shows the overall F-1 scores for each model across all folds.

Machine Learning Model	**F-1 Score**
Decision Tree Regressor	0.81276
Decision Tree Classifier	0.82870
Ridge Regressor	0.86107
Ridge Regression Classifier	0.85457
Lasso Regressor	0.79733
Logistic Regressor	0.46642

Table 4.2: F-1 Scores of Corresponding Machine Learning Models

From the table, it can be gathered that there is a correlation between the proposed metrics and misinformation resistance as most models yielded F-1 scores around the 0.80 to 0.85 range. The one exception to this is the logistic regression model that yielded an F-1 score below 0.50. This is most likely due to the large number of features or the classification method not translating well for the logistic regression model.

4.3 Feature Analysis

To perform feature analysis, two approaches were taken; statistical analysis and feature importance from the machine learning algorithms.

For the statistical analysis, the p-value threshold — the probability that the results could be attributed to random change — was set to .05. Based on this boundary, it was determined which of the metrics were statistically significant. Among the diversity metrics, Average Reaction Entropy (p-value = 1.43e-13), Top Level Domain Entropy (8.50e-12), MNA Entropy (.00116), Word Entropy (.00133), and Domain Entropy (.0307) were statistically significant. It was also gathered that Post Type Entropy (.05766) and total Reaction Entropy (.07345) were not statistically significant.

For general metrics, out of the 21 that were included in our model, 10 of them were determined to be statistically significant. With these results, we are able to reinforce any findings determined by a machine learning model. Although some of the metrics were determined to be statistically insignificant, they were included in the model to test if they have any effects on the reliability of the model's ability to predict a group's misinformation resistance score.

The decision tree models were utilized to confirm the results of the statistical significance. This was decided due to the ability to not only determine feature importance from the model, but also because it allows the computation of the information gain of each metric [22, 23].

This process involved starting with decision trees that could fully expand to determine which metrics were utilized earlier in the decision tree and sorted more data points. On those decision trees we can gather the importance of each feature through the model. The earlier in the tree the feature is utilized for a decision and the more data entries that are filtered by a feature, the more important that feature is.

In order to ensure that the decision trees were not over fitted to the data, we performed two operations on the decision trees. The first was to filter out features from the overall feature list that were not deemed to have a high importance. As this did not drastically decrease the F-1 score of the decision trees, the features that were removed did not have a close correlation with misinformation resistance.

The other operation to ensure that the decision trees were not over fitted to the data was to reduce the max depth of the trees. When performing this operation, the F-1 score increased. This suggested that the full size tree was over fitted to the data, and the tree with reduced max depth is a better model for predicting the misinformation resistance of new data points. This would be an important characteristic to note for predicting an unknown group's misinformation resistance utilizing the decision tree model.

With the statistical significance of the features and the feature importance gathered from the decision trees, we were able to gather which features had the most relevant correlations with misinformation resistance. The graph below displays some of the features with higher correlation to misinformation resistance.

We observed several patterns amongst all of the experiments in terms of important features between the general metrics and diversity metrics.

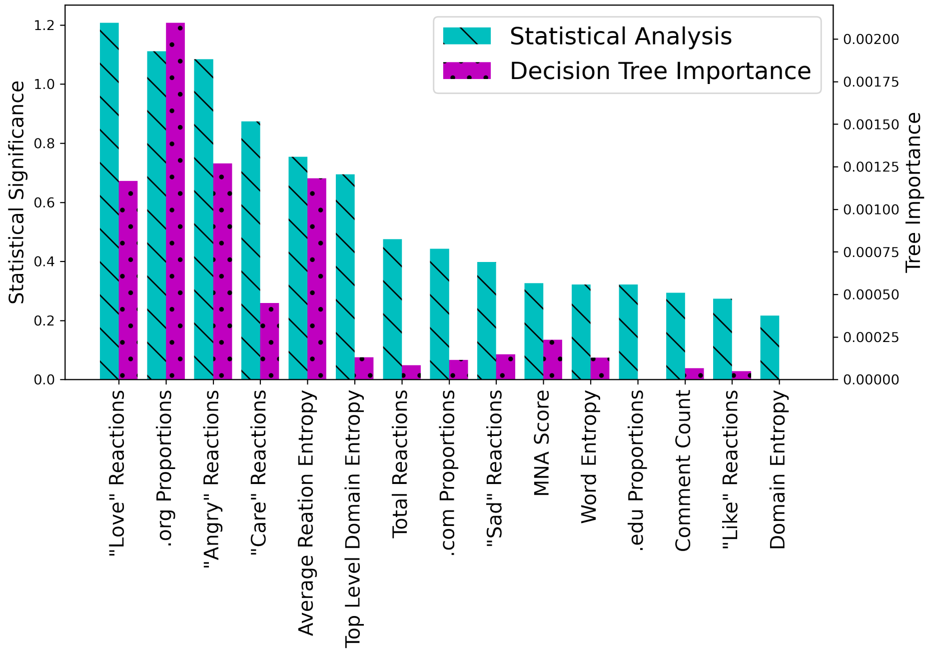

Figure 4.2: Comparison Between Statistical Significance and Feature Importance

4.3.1 Top Level Domains

While .com was the most commonly shared top level domain from the data set, there was not strong correlation between the sharing of sites under the .com domain with misinformation resistance. The domain .org — the second most commonly shared top level domain — had a positive correlation with misinformation resistance. This likely is due to the idea that as groups share more posts from organizations, they would be less likely to be exposed to solely misinformation, which allows them to be more resistant to misinformation.

4.3.2 Reactions

We had found two reactions that had a strong correlation with misinformation resistance: "love" and "angry." The love reaction had a positive correlation with misinformation resistance, while the angry reaction had a negative correlation with misinformation resistance. We recognized that within the list of all reactions, these were subjectively the most extreme emotions displayed among the reactions. A conclusion drawn from this was the notion that the more misinformation resistant a group was, the more likely that they would respond more positively to new information. On the other hand, the less misinformation resistant a group was, the more likely they would display hostility toward new information.

4.3.3 Diversity Metrics

Between all of the calculated diversity metrics, the two metrics that displayed the most significant correlation with misinformation resistance were average reaction entropy and MNA score. Both of these metrics had a positive correlation with misinformation resistance. For the average reaction entropy, we were able to draw the conclusion that groups that would have more diversity in their reactions would be more likely to critically analyze and discuss new information from several perspectives. If a group of people reacts in different ways to new information, they may be more likely to discuss the information prior to accepting it, which would help minimize the misinformation from spreading within the group. For the MNA score, a conclusion can be drawn that the posts shared by resistant groups were more likely to be accepted as factual by a diverse set of groups. This indicates that the group is likely sharing information that other groups with high diversity has also shared, resulting in a lower likelihood of misinformation being spread within the group.

Chapter 5

Conclusion

In this work, we tried to investigate the correlation between the diversity of a group and their resistance to misinformation. We proposed varying diversity metrics based on the behaviors of Facebook groups' that could be analyzed through the CrowdTangle API. Utilizing these tools, we were able to collect a real data set to perform machine learning algorithms to determine that some of the proposed metrics have a correlation with misinformation resistance.

As our research primarily focused on link sharing within Facebook groups, our data collection methods were focused on that aspect. Due to this, a potential research opportunity would be through extending the methodology to other social media platforms such as Twitter or Instagram. Through the application on other social media platforms, there is the opportunity to expand or refine the current diversity metrics to better suit the characteristics of the specific social media platform. There is also the opportunity to compute new diversity metrics to determine if they have a correlation with misinformation resistance. There is also a possibility to analyze our current data through another perspective. For example, if we were able to collect the temporal data of posts, we could run analysis on the diversity of post types over time. This opens up many possibilities for varying analyses of misinformation in social media.

Bibliography

[1] Emily Chen, Herbert Chang, Ashwin Rao, Kristina Lerman, Geoffrey Cowan, and Emilio Ferrara. Covid-19 misinformation and the 2020 us presidential election. *The Harvard Kennedy School Misinformation Review*, 2021.

[2] Rachel Leigh Greenspan and Elizabeth F Loftus. Pandemics and infodemics: Research on the effects of misinformation on memory. *Human Behavior and Emerging Technologies*, 3(1):8–12, 2021.

[3] Engin Bozdag, Qi Gao, Geert-Jan Houben, and Martijn Warnier. Does offline political segregation affect the filter bubble? an empirical analysis of information diversity for dutch and turkish twitter users. *Computers in human behavior*, 41:405–415, 2014.

[4] Munmun De Choudhury, Scott Counts, and Mary Czerwinski. Identifying relevant social media content: leveraging information diversity and user cognition. In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, pages 161–170, 2011.

[5] Xiuwen Zheng and Amarnath Gupta. Social network of extreme tweeters: A case study. In *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 302–306, 2019.

[6] Allison Harell, Dietlind Stolle, and Ellen Quintelier. Network diversity and political participation: A complication or asset. In *Conference on Comparative Perspectives on Political Socialization, Bruges, Belgium*, 2008.

[7] Mohammad-Ali Abbasi and Huan Liu. Measuring user credibility in social media. In *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction*, pages 441–448. Springer, 2013.

[8] Minjeong Kang. Measuring social media credibility: A study on a measure of blog credibility. *Institute for Public Relations*, pages 59–68, 2010.

[9] David A Broniatowski, D Kerchner, F Farooq, X Huang, AM Jamison, M Dredze, and SC Quinn. Twitter and facebook posts about covid-19 are less likely to spread false and low-credibility content compared to other health topics. 2021.

[10] Media bias/fact check, 2015. URL `https://mediabiasfactcheck.com/`.

[11] Vikas S Chavan and SS Shylaja. Machine learning approach for detection of cyber-aggressive comments by peers on social media network. In *2015 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 2354–2358. IEEE, 2015.

[12] Qijin Cheng, Tim MH Li, Chi-Leung Kwok, Tingshao Zhu, and Paul SF Yip. Assessing suicide risk and emotional distress in chinese social media: a text mining and machine learning study. *Journal of medical internet research*, 19(7):e243, 2017.

[13] Filippo Menczer. The spread of misinformation in social media. In *Proceedings of the 25th International Conference Companion on World Wide Web*, pages 717–717, 2016.

[14] Juhi Kulshrestha, Muhammad Zafar, Lisette Noboa, Krishna Gummadi, and Saptarshi Ghosh. Characterizing information diets of social media users. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 9, 2015.

[15] S Brennen, F Simon, PN Howard, and RK Nielsen. Types, sources, and claims of covid-19 misinformation. 2020.

[16] Jesse Shore, Jiye Baek, and Chrysanthos Dellarocas. Network structure and patterns of information diversity on twitter. *arXiv preprint arXiv:1607.06795*, 2016.

[17] Jason Brownlee. A gentle introduction to k-fold cross-validation. *Machine Learning Mastery*, 2018.

[18] Nathalie Japkowicz. Why question machine learning evaluation methods. In *AAAI workshop on evaluation methods for machine learning*, pages 6–11, 2006.

[19] Sebastian Raschka and Vahid Mirjalili. Python machine learning: Machine learning and deep learning with python. *Scikit-Learn, and TensorFlow. Second edition ed*, 2017.

[20] Luca Massaron and Alberto Boschetti. *Regression Analysis with Python*. Packt Publishing Ltd, 2016.

[21] scikit-learn: Machine learning in python. URL `https://scikit-learn.org/stable/index.html`.

[22] Jason Brownlee. How to calculate feature importance with python. *Machine Learning Mastery*, 2020.

[23] Jason Brownlee. Information gain and mutual information for machine learning. *Machine Learning Mastery*, 2019.

Is_Social_Diversity_Related_to_Misinformation_Resistance_An_Empirical_Study_on_Social_Communities

Final Audit Report 2021-06-10

Created:	2021-06-10
By:	Bioengineering Department (bioengineering@scu.edu)
Status:	Signed
Transaction ID:	CBJCHBCAABAAnG5OZZ4MlzOP0JX6mpJo7HsF-14DRLAG

"Is_Social_Diversity_Related_to_Misinformation_Resistance_An_Empirical_Study_on_Social_Communities" History

Document created by Bioengineering Department (bioengineering@scu.edu)
2021-06-10 - 5:25:23 PM GMT- IP address: 24.6.105.116

Document emailed to Nam Ling (nling@scu.edu) for signature
2021-06-10 - 5:25:55 PM GMT

Email viewed by Nam Ling (nling@scu.edu)
2021-06-10 - 5:26:04 PM GMT- IP address: 72.14.199.27

Document e-signed by Nam Ling (nling@scu.edu)
Signature Date: 2021-06-10 - 5:29:00 PM GMT - Time Source: server- IP address: 75.4.202.62

Agreement completed.
2021-06-10 - 5:29:00 PM GMT