

Mining personal media thresholds for opinion dynamics and social influence

Casey Doyle^{*†}, Alex Meandzija[†], Gyorgy Korniss[†], Boleslaw Szymanski^{†‡}, Derrik Asher[§], and Elizabeth Bowman[§]

*Email: doylec3@rpi.edu

[†]Department of Physics, Applied Physics, and Astronomy

[‡]Department Computer Science

Rensselaer Polytechnic Institute, Troy, NY, USA

[§]Computational & Information Sciences Directorate

US Army Research Laboratory, Aberdeen, MD, USA

Abstract—To study the detailed effects of social media consumption on personal opinion dynamics, we gather self reported survey data on the volume of different media types an individual must consume before forming or changing their opinion on a subject. We then use frequent pattern mining to analyze the data for common groupings of responses with respect to various media types, sources, and contexts. We show that in general individuals tend to perceive their behavior to be consistent across many variations in these parameters, while further detail shows various common parameter groupings that indicate response changes as well as small groups of individuals that tend to be consistently more easily swayed than the average participant.

Index Terms—opinion formation, personal influence thresholds, social media, crowdsourcing, data mining, association rules

I. INTRODUCTION

The importance of understanding how individuals interact with and are influenced by media has been growing as people continue to get increasingly large amounts of their news via various social networking websites and media exchange platforms [1]–[3]. The problem of knowing exactly how an individual takes in and processes this information is difficult to accurately study, and further the variation among different people makes it even worse when scaling up to understand how the preferences and tendencies of an individual inform on the behavior of the larger population as a whole. In this work we attempt to bridge some of these difficulties, extending the work done in [4], [5]. In those studies, a group of individuals was given a survey on the number of media items they would need to consume given various parameters such as the media type (format), source (general like mindedness) and context (general level of controversy of the subject). With this

information, general thresholds were established for various media types as well as their interaction with the source and context of the media. Yet, each individual was asked a very limited set of questions to establish the general behavior of a large population. Here we broaden the scope of the study such that each individual is asked a larger number questions covering many of the combinations between media types, sources, and contexts in order to establish a better profile for how the individual responses change for each person. Additionally, we add a new field for some participants where they are asked about the number of media items required for “shifting” their opinion instead of “forming” it. Using this data, we mine for various frequent patterns, attempting to pick up on different types of people whose behavior deviates from the average while still being common enough to not be considered outliers.

While data mining techniques are popular in the realm of opinion formation, they are mostly used to understand the evolution of opinions as they change in empirical networks [6]–[8]. Other applications of data mining to the field have been applying mining techniques to extract personality types and behavior from social network and cell phone data [9], [10]. Here we study a more direct data set via an online survey and use frequent pattern detection tools to search for categories of people that relate directly to their social media use and behavior. We first investigate the basic statistics corresponding to the data that we present, then analyze the pattern lists to search for strong relationships between items. Finally we present the combinations that represent the most interesting relationships as well as general conclusions that can be drawn from the overall shape of the patterns found.

II. DATA COLLECTION

A. Platform and Participant Selection

The data in this experiment was collected using Amazon Mechanical Turk (MTurk), which is a survey hosting platform that connects researchers to diverse pools of participants [11]–[13]. In total, the experiment involves 1431 participants.

This work was supported in part by the Army Research Laboratory (ARL) under Cooperative Agreement Number W911NF-09-2-0053 (NS-CTA), and by the Office of Naval Research (ONR) Grant No. N00014-15-1-2640. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies either expressed or implied of the Army Research Laboratory or the U.S. Government.

IEEE/ACM ASONAM 2018, August 28-31, 2018, Barcelona, Spain
978-1-5386-6051-5/18/\$31.00 © 2018 IEEE

Before beginning the survey, participants were asked if they use social media; only those that responded 'yes' and had not already participated were admitted to the study. In addition, only users aged eighteen years or older and located in the US were accepted for participation in this experiment. More details on the setup and questions asked of the individuals are given in the Appendix.

B. Survey and Data Description

Upon entering the survey, the participants were assigned randomly to one of two groups, the "fixed-source" group or the "fixed-context" group. Users were then asked to fill out a brief demographic questionnaire that contained items such as their age, occupation, and average level of social media usage. Once completed, the subjects were asked to report on the number of social media items they would need to see before forming an opinion on the subject of the media. These questions include three parameters that describe the media: *type*, *context*, and *source*. We identify three media types:

- 1) Images: for still photos and drawings
- 2) Videos: for any animations or moving pictures
- 3) Messages: for text, tweets, and Facebook posts

Four media controversy levels defined:

- 1) Low: minimal (some people would form an opinion)
- 2) Medium: generally controversial (most would form an opinion)
- 3) High: very controversial (most or all would form an opinion)
- 4) None: no reference to controversy

And three media sources:

- 1) Unknown: individual has no knowledge of the source
- 2) Like-minded: the source of the media generally thinks similarly to the recipient
- 3) Different-minded: the source of the media generally thinks differently from the recipient

To limit survey fatigue, we decreased the number of questions asked. Instead of asking each participants all possible combinations of the three parameters, we used the groups as follows. The participants in each group ("fixed-context" and "fixed-source") were given surveys that corresponded to each possible combination of the above parameters with the exception of their group parameter. For instance, individuals in the "fixed-context" group were assigned a random context at the start of questioning, and asked about every combination of type and source with that one context level. Similarly, the "fixed-source" group was given a single source and asked about all combinations of type and context.

Finally, the "fixed-source" group was given an extended questionnaire that included questions on how many media items they would need to see to *shift* their opinion as well as the original set that asked about *forming* their opinion. From here on out, responses to the opinion shifting questions will be referred to as their own group: "shifting", and the "fixed-source" group will be implied to mean data only from the opinion formation answers. When discussing the full suite

of answers provided by those individuals, the data will be referred to as the "shifting-formation" group. The final sizes of each group are 616 users in the "fixed-source" group (and consequently 616 users for the "shifting" and "shifting-formation" groups) and 815 users in the "fixed context group.

III. ANALYSIS

A. Prior Analysis and Binning

The prior results on similar data sets revealed some key features of how individuals respond to these types of questions [4], [5]. First, reported thresholds are shown to be distributed log-normally, and thus a log transform can easily be performed to normalize the data for analysis. Second, both the source and context each have a significant effect on the requisite number of media items viewed dependent on the value of other parameters. In general, images are less susceptible to change depending on other parameters, while videos are more sensitive to source and messages are more sensitive to context. Average values for each media type are 4 – 7 images, 2 – 5 videos, and 3 – 6 messages required to form opinions. We do not repeat much of this analysis in the current paper as it is beyond our scope, but the overall average number of media items to form an opinion (4.5 items) is commonly seen throughout our analysis.

Using the prior results to inform on the current analysis, we clean and preprocess the data for mining by coding reported thresholds into bins. Since the data has been shown to be normalized via a log transform, we utilize a logarithmic binning scheme to sort the data (i.e., bin one contains responses of threshold 1, bin two contains responses with thresholds 2-3, bin three responses of thresholds 4-7, bin four responses of thresholds 8-15, etc.). From here on in the paper all 'response' values will correspond to the log-binned value of the reported thresholds, not the raw threshold given. This allows us to better speak in generalities about behavior, as threshold values represent the category instead of specific values that have less meaning in a self-reported data set like the one here. Additionally, the binning allows for a fuzzy look at thresholds for the sake of finding more representative patterns by grouping together similar thresholds instead of attempting to pick out only patterns that contain the exact same values.

B. Association Rule Mining and Processing

In order to identify interesting subsets and trends within the overall population, we perform *frequent pattern* and *association rule* mining on the dataset. The frequent patterns are simply itemsets that appear commonly (identified via a minimum support defined as the fraction of all transactions that contain that pattern), while the rules contain directional information on the implications of the frequent itemsets, i.e., individuals that give responses *A* and *B* are also likely to give response *C* [14]. Rules are defined as frequent via both a minimum support and confidence (defined as the probability that a transaction contains the consequent given that it contains the antecedent). The frequent patterns are mined using the Apriori algorithm within the `arules` package in R [15]–[17].

For the "fixed-context" group, the minimum support value $sup = 0.01$ is used with a confidence of $conf = 0.6$, yielding 3263 rules with a minimum absolute count of 8. Finally Similarly, the "fixed-source" group is mined with a minimum support of $sup = 0.015$ and a confidence of $conf = 0.6$, yielding 4716 rules with an absolute minimum count of 9. In order to remove redundancies they are then filtered to remove any rule belonging to an itemset that is not maximal (itemsets where none of their supersets are considered frequent) [14]. Further, all rules are tested for statistical significance via the Fisher Exact Test (using $p > 0.01$), and insignificant rules are also removed. This yields a final ruleset of 1484 in the "fixed-context" group and 2212 in the "fixed-source" group.

The mining on opinion shifting is much the same. The "shifting" group is mined with a minimum support of $sup = 0.015$ and minimum confidence of $conf = 0.6$, yielding 2767 rules with a minimum absolute count of 9. Due to its larger size, the "shifting-formation" group is mined with a minimum support of $sup = 0.025$ and confidence $conf = 0.6$, yielding 2346 rules and a minimum absolute count of 15. After pruning for maximal and significant rules, the sets become 1790 and 1957 rules long, respectively.

C. Response Statistics

In this section we look at a few different cross sections of the data. The first is general user data, getting a general view of how the individuals behaved as a whole. Then we look at the mined rules, to see if there are any patterns or statistics among the frequent itemsets that differ from the user statistics. Finally, we rank the rules based on their lift value. The lift score is a measure of the surprise of the rule, defined as the ratio of the percentage of times the pattern appeared in the dataset divided by the probability that the pattern would arise randomly if the items within were independent, represented as

$$lift(XY) = sup(XY) / sup(X)sup(Y)$$

. We rank the rules based on their lift, then study the top 100 rules more closely to see how the most unlikely patterns behave in relation to the rest.

The main focus of this analysis is to study how different parameters affect the response scores in a more group-focused approach than the prior analysis. By calculating the average and standard deviation for each rule we can observe some initial trends in the rules, shown in Table I. Immediately it is clear that the rules have a generally very low average standard deviation, meaning that most rules contain consistent thresholds for varying parameter values. This effect is strong for all rules when compared to the thresholds for the users as a whole, indicating the rules are picking up on a tendency outside of simply general user behavior, an idea further reinforced by the often large deviations that responses have from their corresponding user's average as shown in Fig. 1. Further, the effect is increased for the top 100 rules by lift. In both cases the effect is true even if, to account for the many short rules with no changes at all, we include only rules with changes. Even with that reduced data set the standard deviation of rules

is noticeably lower than that for the users as a whole. Another important difference between the lists is the average response for the rules is lower than the responses in general, another effect that is stronger among rules with high lift scores. These findings show two important features of the frequent patterns within the data sets: there are significant groups within the populations that tend to have consistent parameter subsets that lead to lower general thresholds for forming and shifting their opinion.

Some of the details causing this effect can be seen in Fig 2, where the distributions show a large difference in the tail end of the responses. The averages for all users tend to drop off after bin four (reported thresholds ranging from eight to fifteen), but there is a size-able tail to the distribution that contains many more extreme responses. When studying the averages across the rules, the tail disappears, and after the drop-off at bin four there are nearly no responses with thresholds in the higher bins. In this case it is not necessarily that the rules have a higher volume of low values, it's that users that respond with high values don't tend to be consistent enough to form significant groups despite the logarithmic providing more flexible boundaries for group inclusion at those values. The top ranking rules by lift show that this is no mere statistical effect, either, as the rules that contain responses with the lowest thresholds dominate the high lift rules. Interestingly this is not absolute, as despite the head dominance of the high lift values there are still significant groups of individuals that respond with the higher threshold values. So while it is possible that the volume of rules with average thresholds similar to those of the population as a whole could be a relic of the high number of responses with thresholds in that range, the high lift rules show that there are statistically interesting groups that respond mostly with low thresholds.

D. Contents of Rules

With a basic understanding of what the general responses are and how they are represented when mining for rules, we now turn our attention to what tends to be the content of the rules. In Fig. 3, we can visually inspect the top 100 rules and see how the individual rules tend to interact with each other. In this case, the rules are highly modular, dominated by a few very large, tight clusters. As would be expected from the prior analysis, these clusters contain many rules that match together similar responses. Similarly, Fig. 4 shows a graphical view of the top 100 for the "shifting" and "shifting formation" groups. Here we see that the "shifting" group still maintains some of this structure, although there are fewer large clusters than with the formation groups. Moreover in the "shifting formation" group the the large structures are largely gone and replaced by many smaller, isolated structures. Again, this mirrors the results in Table I, which shows a much lower standard deviation for this group than for the others arising from the lack of rules bridging the gaps to make larger structures, and instead staying more within the same threshold values.

TABLE I
AVERAGE THRESHOLDS FOR EACH RULE

Group	Rules	Average Threshold	Average SD	Average SD w/ Change
Fixed Source	All	3.021	0.217	0.568
	Top 100	2.257	0.061	0.544
Fixed Context	All	2.774	0.254	0.578
	Top 100	2.620	0.063	0.608
Shifting	All	3.064	0.204	0.578
	Top 100	2.717	0.115	0.566
Shifting Formation	All	2.749	0.079	0.586
	Top 100	2.419	0.012	0.577

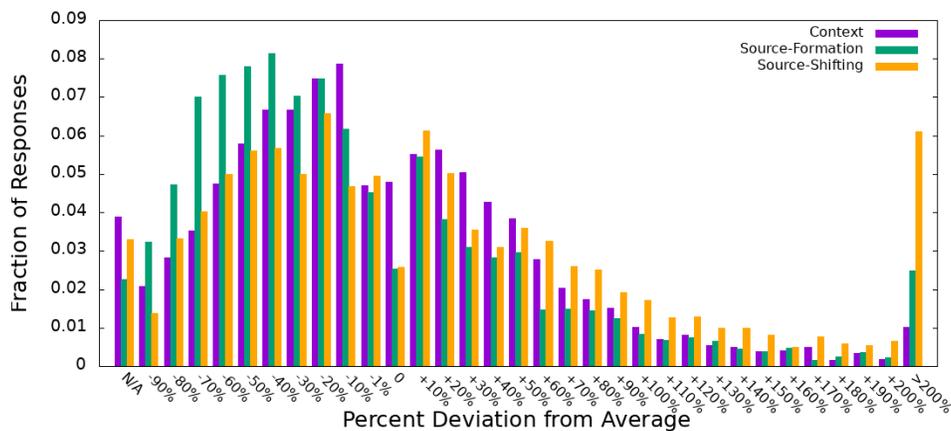


Fig. 1. The proportion of responses as a function of deviation relative to the corresponding user's average response.. This shows the consistency of responses for each user.

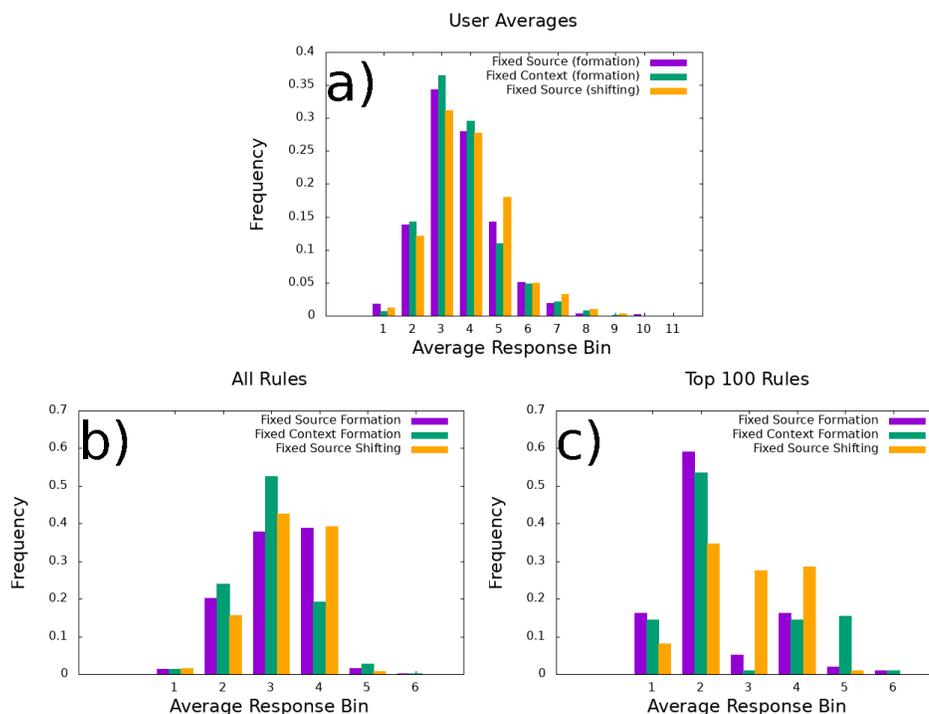


Fig. 2. (a) Distribution of the average thresholds reported by each user. Thresholds are binned logarithmically as described in Sec. III-A. (b) Distribution of the average thresholds within each calculated rule. (c) Distribution of average thresholds from only the top 100 rules by lift.

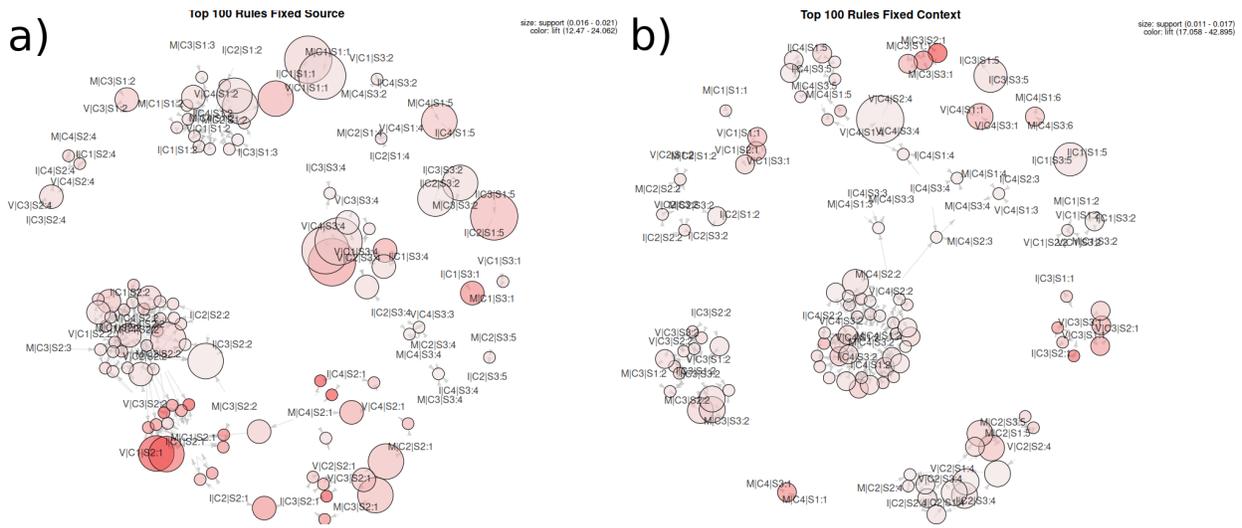


Fig. 3. Network representation of the top 100 rules by lift for the (a) "fixed-source" and (b) "fixed-context" groups. Each circle represents a rule with the connections being the items within those rules. The size of the circle is scales with the support of the rule, while the color represents the lift score (darker red corresponds to higher lift).

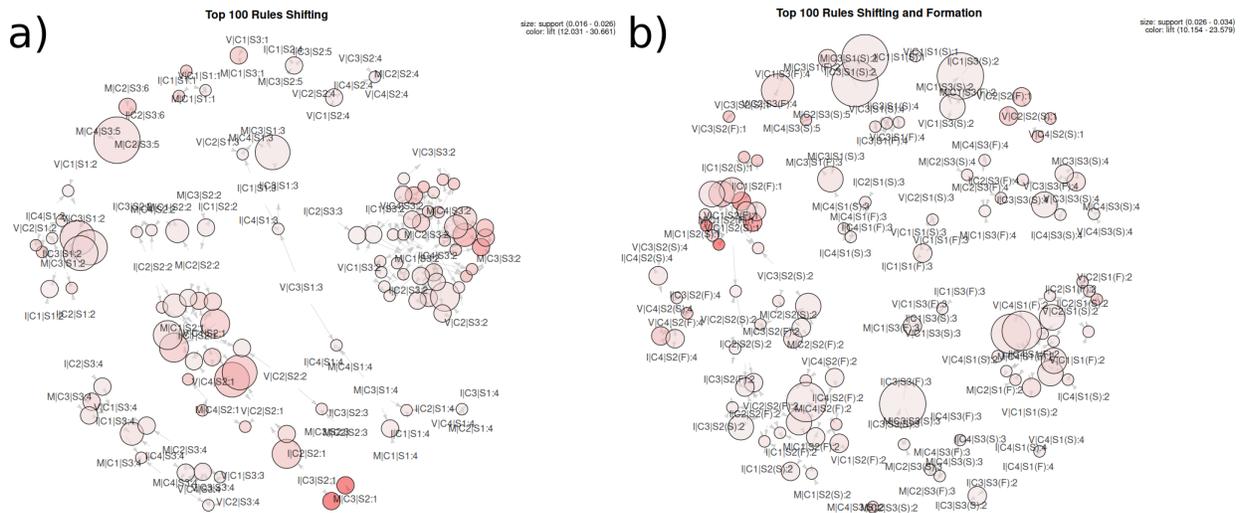


Fig. 4. Network representation of the top 100 rules by lift for the (a) "shifting" and (b) "shifting-formation" groups. Each circle represents a rule with the connections being the items within those rules. The size of the circle is scales with the support of the rule, while the color represents the lift score (darker red corresponds to higher lift).

TABLE II
AVERAGE THRESHOLDS FOR EACH USER

Group	Average Threshold	Average SD	Average SD w/ Change
Fixed Source	3.228	0.826	0.884
Fixed Context	3.183	0.870	0.892
Shifting	3.392	0.880	0.909

To look deeper into what makes the rules (and in particular what drives changes in threshold values), we look into the statistics of the parameter makeups of each rule. This is somewhat difficult due to the deep interplay between the three parameters in driving changes to user reported thresholds, as nearly every rule has at least one change in parameter. For

instance, in the "fixed-source" group, 94% of rules have at least one change in media type, while 97% have at least one change in context level. This remains true for each of the other groups as well, as each have > 90% of their rules changing in media type and context (or source for the "fixed-context" group). However, looking more deeply into the makeup of each of these rules we observe some common trends.

Fig. 5 shows that for the formation groups, the media types are relatively stable across combinations. There are very few rules with only a single type, and no clear preference for any pairing of two media types. There is, however, a slight difference in the percentage of rules in each category for the "fixed-source" group when looking at rules with messages and images versus those with messages and videos. This effect is

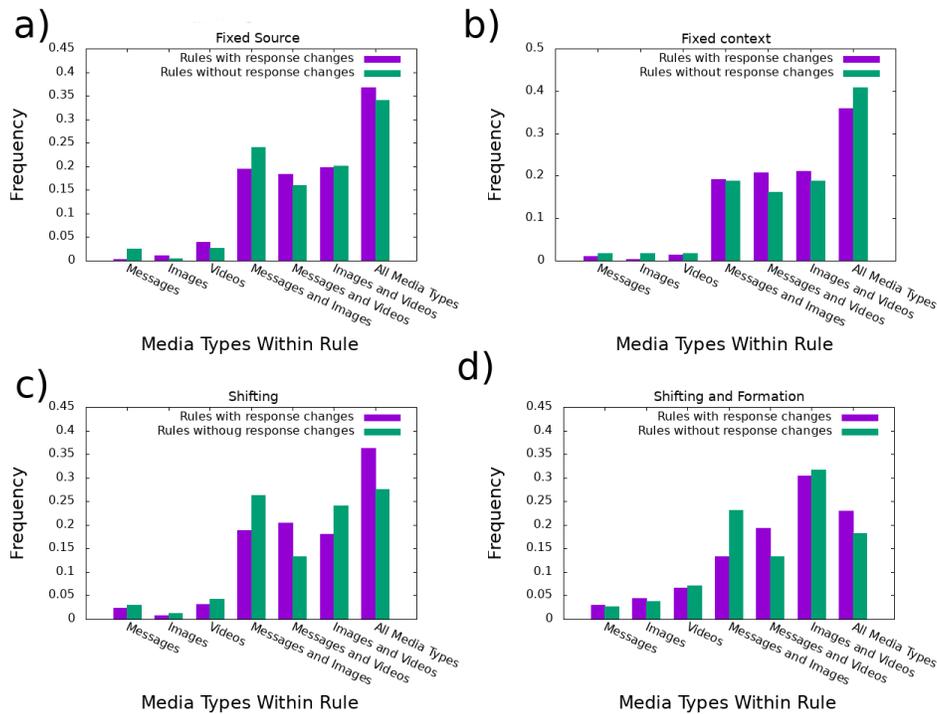


Fig. 5. Percentage of rules that contain each possible combination of media types, separated by whether that rule also contains a change in response. Groups corresponding to the plots are (a) "fixed-source", (b) "fixed-context", (c) "shifting", and (d) "shifting-formation"

even more pronounced in the groups where shifting responses are included, and the percentage of rules that produce no change is much higher for rules that contain messages and images than those that contain messages and videos. Similarly, the percentage of rules that contain a change in response and include both messages and videos in them is higher than that which includes messages and videos but no change in response. From this, we can conclude that users tend to think of messages and images as being more similar in their potential to shift opinions than messages and videos.

In Fig. 6, we extend this analysis to rules that include response changes to investigate the source and context levels within those rules. The average values for the source (in the "fixed-context" group) and context (in the "fixed-source" group) are generally very close to the expected mean for a random sample, but when separated based on whether the rules contain a change in response or not some trends begin to emerge. First, the source level is slightly higher in the rules that produce change, indicating more changes happen when the source is differently minded, although the differences are too small to draw any concrete conclusions. A similarly small effect can be seen in the "fixed-source" group indicating that lower context values (less controversy) are slightly more common in rules that produce changes. This effect is also present in the "shifting" and "shifting-formation" groups, with the "shifting" effect being the most pronounced. Similarly, the coefficient of variation (a normalized version of the standard deviation $C_v = \sigma/\mu$ [18]) shows a trend towards a higher variation in context implicating a higher likelihood that the

rule also contains a response change. Conversely, it shows the opposite for the source. This indicates a stronger relationship between the context and response than the source, especially in the case of opinion shifting, where the effect is far more pronounced.

IV. CONCLUSIONS

While studies that focus on mining data from large datasets such as those collected by social networking websites can present many interesting and unbiased findings on the nature of how users interact with data online, they are limited in the types of data that can be collected in this way. More detailed questions relating to exactly when users reach conclusions, what pushes them in one direction or another, and how staunchly they stick to the opinions they form are difficult to answer from these types of data collection. In order solve this hurdle, we use large scale surveys to crowdsource a solution, asking individuals in a more direct manner how they interact with media online. From these surveys, we can not only begin to understand the basic thresholds of individual opinion formation [4], [5], but also gain an idea of what kinds of popular trends various groups of people exhibit by mining the survey sets for interesting features and patterns.

The most predominant feature discovered in this way is the high degree of consistency within itemsets. Users tend to behave more similarly to each other when describing the pieces of media they feel are equivalent, and thus is far easier to pick up on the aspects of media that make them similar than what makes them different. Additionally, the frequent itemsets

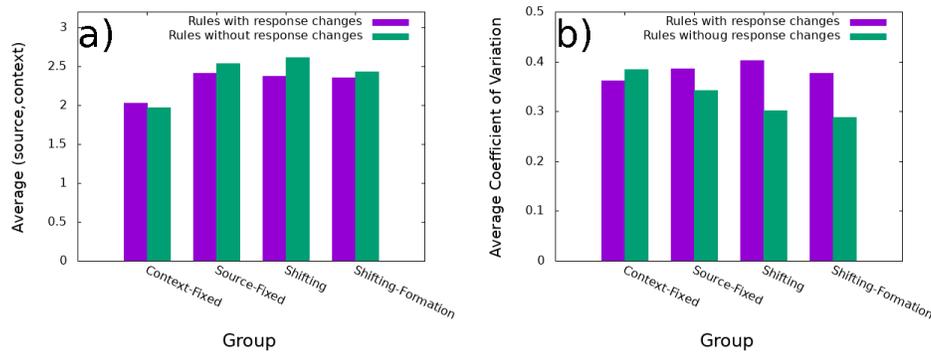


Fig. 6. (a) The average source (for "fixed-context" group) or context (for all other groups) level for each rule averaged over all rules within that group, separated by whether the rule produced a change in response. (b) The average coefficient of variance (σ/μ) of the source (for "fixed-context" group) or context (for all other groups) level for all rules within that group, separated by whether the rule produced a change in response.

have a lower average threshold value than the thresholds of the populations as a whole, meaning that users are far more consistent when describing media that they feel is more convincing, and thus it is easier to group people based on what they are partial to rather than what they dislike. Further, when looking at rules that do contain response changes, some trends in the parameter values become apparent. For instance, media types such as messages and images have very similar swaying power over individuals, while others such as messages and videos tend to be far more different. Additionally, the context (controversy level) of the media appears to be more important in predicting response changes than other parameters, as the variation in the context values within rules is far higher when that rule also contains a response change.

These conclusions are a start towards a deeper understanding of the size and behavior of different groups of individuals, but further analysis is necessary. First, the surveys conducted recorded basic demographic data, social media consumption statistics, and favorite social media websites and news sources for each participant. While not studied in depth in this work, this information has potential for future examination into the links between demographic groups, popular social media communities, and how individuals respond to and consume data. The results presented here could be extended via additional surveys to include more study groups that are asked about opinion shifting as well as formation, in particular in the case of a fixed context and varying source. Further studies could also extend into the differences even among single media types could be valuable in differentiating how users respond to videos or articles of different lengths for example.

As information and opinions spread through societies at ever greater rates, the subject of how social influence occurs becomes increasingly centered around the media with which people interact online. Greater insight into the how different groups behave and how large they are in the first place is necessary in order to understand and model these processes to further our understanding of how people influence each other.

REFERENCES

- [1] H. Kwak, C. Lee, H. Park, and S. Moon, "What is Twitter, a social network or a news media?" in *Proceedings of the 19th international conference on World wide web - WWW '10*. New York, New York, USA: ACM Press, 2010, p. 591.
- [2] H. Gil de Zúñiga, N. Jung, and S. Valenzuela, "Social Media Use for News and Individuals' Social Capital, Civic Engagement and Political Participation," *Journal of Computer-Mediated Communication*, vol. 17, no. 3, pp. 319–336, Apr. 2012.
- [3] A. Hermida, F. Fletcher, D. Korell, and D. Logan, "SHARE, LIKE, RECOMMEND," *Journalism Studies*, vol. 13, no. 5-6, pp. 815–824, Oct. 2012.
- [4] D. Asher, J. Caylor, M. Mittrick, J. Richardson, E. Heilman, E. Bowman, G. Korniss, and B. Szymanski, "The Investigation of Social Media Data Thresholds for Opinion Formation," in *ICCRTS 2017: 22nd International Command and Control Research & Technology Symposium*, Los Angeles, CA, 2017, p. 27.
- [5] A. Derrik, E. Bowman, C. Doyle, G. Korniss, and B. K. Szymanski, "An experimental approach to identify perceived opinion formation thresholds from social media computers in human behavior," *Computers in Human Behavior*, submitted 2017.
- [6] M. Adedoyin-Olowe, M. M. Gaber, and F. Stahl, "A Survey of Data Mining Techniques for Social Media Analysis," Dec. 2013.
- [7] M. Kaschesky, P. Sobkowicz, and G. Bouchard, "Opinion mining in social media," in *Proceedings of the 12th Annual International Digital Government Research Conference on Digital Government Innovation in Challenging Times - dg.o '11*. New York, New York, USA: ACM Press, 2011, p. 317.
- [8] P. Sobkowicz, M. Kaschesky, and G. Bouchard, "Opinion mining in social media: Modeling, simulating, and forecasting political opinions in the web," *Government Information Quarterly*, vol. 29, no. 4, pp. 470–479, Oct. 2012.
- [9] D. Markovikj, S. Gievska, M. Kosinski, and D. Stillwell, "Mining facebook data for predictive personality modeling," in *Proceedings of the 7th international AAAI conference on Weblogs and Social Media (ICWSM 2013)*, Boston, MA, USA, 2013, pp. 23–26.
- [10] G. Chittaranjan, J. Blom, and D. Gatica-Perez, "Mining large-scale smartphone data for personality studies," *Personal and Ubiquitous Computing*, vol. 17, no. 3, pp. 433–450, Mar. 2013. [Online]. Available: <http://link.springer.com/10.1007/s00779-011-0490-1>
- [11] M. Buhrmester, T. Kwang, and S. D. Gosling, "Amazon's Mechanical Turk," *Perspectives on Psychological Science*, vol. 6, no. 1, pp. 3–5, Jan. 2011.
- [12] G. Paolacci, J. Chandler, and P. G. Ipeirotis, "Running Experiments on Amazon Mechanical Turk," *Judgment and Decision Making*, vol. 5, no. 5, pp. 411–419, June 2010.
- [13] K. Crowston, "Amazon Mechanical Turk: A Research Tool for Organizations and Information Systems Scholars." Springer, Berlin, Heidelberg, 2012, pp. 210–221.
- [14] M. J. Zaki, W. Meira Jr, and W. Meira, *Data mining and analysis: fundamental concepts and algorithms*. Cambridge University Press, 2014.

Medium Controversy

2) **Before you FORM an OPINION** how many data types listed below would you expect to view in a day , given that the data type(s) were **posted by people who think like you?**

Videos

Next

Fig. 7. Sample question asked of participants. Each question is presented on its own page, and participants are given the context at the top, then asked a question containing the source, and given the media type in question at the bottom.

- [15] M. Hahsler, C. Buchta, B. Gruen, and K. Hornik, *arules: Mining Association Rules and Frequent Itemsets*, 2018, r package version 1.6-1. [Online]. Available: <https://CRAN.R-project.org/package=arules>
- [16] M. Hahsler, B. Gruen, and K. Hornik, "arules – computational environment for mining association rules and frequent item sets," *Journal of Statistical Software*, vol. 14, no. 15, pp. 1–25, Oct. 2005. [Online]. Available: <http://dx.doi.org/10.18637/jss.v014.i15>
- [17] M. Hahsler, S. Chelluboina, K. Hornik, and C. Buchta, "The arules r-package ecosystem: Analyzing interesting patterns from large transaction datasets," *Journal of Machine Learning Research*, vol. 12, pp. 1977–1981, 2011. [Online]. Available: <http://jmlr.csail.mit.edu/papers/v12/hahsler11a.html>
- [18] B. Everitt and A. Skrondal, *The Cambridge dictionary of statistics*. Cambridge University Press Cambridge, 2002, vol. 106.

APPENDIX QUESTIONNAIRE

The questionnaire used in this experiment is simply an extension to that used in prior work [4], [5]. All participants accessed the study via MTurk, where they were immediately asked whether they use social media. If the user responded 'no', they were thanked for their time and not allowed to proceed with the study. If they responded 'yes', they were asked to fill out a short demographic form with general information about the user as well as more detailed questions on their social media usage (favorite websites, amount of social media consumption, and main news sources). This demographic and usage information was not used in the current study, but could be the source of future work in identifying different consumption behaviors within different demographic groups or social media communities.

After the demographic form, participants were asked about how much of a given data type they would need to see to form an opinion, as shown in Fig. 7. Each question was given its own page, and participants were only asked for one response at a time. The controversy level and media type were listed above and below the question, respectively, and the source was incorporated into the question. Participants were required to

LOW **✉ an example of a LOW level of controversy is:**
A car company introduces a new standard car color in hot pink.

MEDIUM **✉ an example of a MEDIUM level of controversy is:**
A typically conservative state (e.g., Texas) approves a liberal law (e.g., recreational marijuana).

HIGH **✉ an example of a HIGH level of controversy is:**
A dictator-run country (e.g., North Korea) fires a chemical weapon into a U.S. allied country (e.g., France).

NONE **✉ no context referenced**

Fig. 8. Definitions of the context values used throughout the paper, as given to participants in the study.

answer each question to move onto the next question in the survey, and their prior responses were not visible to them when answering the following questions. This process continued until the participant was asked every combination of media type, source, and context for their randomly chosen grouping and the survey was complete. Upon completing the survey, the participants were paid a small sum and thanked for their time.

Users were given the definitions of different controversy levels to be used in the study with examples of each controversy level on the instructions page to better orientate each participant to the same approximate level of controversy for each question (Fig. 8). In this paper, context is equivalent to controversy, so the examples given for controversy here also serve as the definitions for the different context values used throughout this work. In the cases that participants were asked about *shifting* their opinion instead of *forming* it, the question was the same as shown in Fig. 7 only with 'shift' in place of the word 'form' in the given example. A single account was only eligible to take the survey once, and was not permitted to repeat it after completion.