



ARL-TR-8448 • AUG 2018



The Investigation of Social Media Data Thresholds for Opinion Formation

by DE Asher, J Caylor, M Mittrick M, J Richardson, E Heilman, E
Bowman, G Korniss, and B Szymanski

Approved for public release; distribution is unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



The Investigation of Social Media Data Thresholds for Opinion Formation

**by DE Asher, J Caylor, M Mittrick M, J Richardson, E Heilman,
and E Bowman**

Computational and Information Sciences Directorate, ARL

G Korniss and B Szymanski

Rensselaer Polytechnic Institute (RPI), Troy, NY

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) August 2018		2. REPORT TYPE Technical Report		3. DATES COVERED (From - To) March 27, 2017–October 5, 2017	
4. TITLE AND SUBTITLE The Investigation of Social Media Data Thresholds for Opinion Formation				5a. CONTRACT NUMBER W911NF-09-2-0053	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) DE Asher, J Caylor, M Mittrick M, J Richardson, E Heilman, E Bowman, G Korniss, and B Szymanski				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) US Army Research Laboratory Computational and Information Sciences Directorate (ATTN: RDRL-CII-T) 2800 Powder Mill Road Adelphi, MD 20783-1138				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-8448	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The pervasive use of social media has grown to over two billion users to date, and it is commonly utilized as a means to share information and shape world events. Evidence suggests that passive social media usage (i.e., viewing without taking action) has an impact on the user's perspective. This empirical influence over perspective could have significant impact on social events. Therefore, it is important to understand how social media contributes to the formation of an individual's perspective. A set of experimental tasks were designed to investigate empirically derived thresholds for opinion formation as a result of passive interactions with different social media data types (i.e., videos, images, and messages). With a better understanding of how humans passively interact with social media information, a paradigm can be developed that allows the exploitation of this interaction and plays a significant role in future military plans and operations.					
15. SUBJECT TERMS opinion dynamics, human subjects experiment, social media, opinion threshold, media types					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 24	19a. NAME OF RESPONSIBLE PERSON DE Asher
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code) (410) 278-3011

Contents

List of Figures	iv
1. Introduction	1
2. Methods	2
2.1 Human Subjects	3
2.2 Data Types	3
2.3 Data-Type Orientations	3
2.4 Contexts	4
2.5 Data Collection	5
2.6 Data Analysis Techniques	6
3. Results	7
3.1 Jarque-Bera (JB) Tests	7
3.2 Wilcoxon Rank Sum Tests	8
3.3 Kruskal-Wallis Tests	8
3.4 Opinion Formation Threshold Estimates	10
4. Discussion and Conclusion	12
5. References	14
List of Symbols, Abbreviations, and Acronyms	16
Distribution List	17

List of Figures

Fig. 1	Data-type presentations: The different orientations of the representative data types (images, videos, and messages) were used in the study to determine if data-type interactions needed to be considered for opinion formation threshold estimates.....	4
Fig. 2	Contexts: Four different contexts (None, Low, Medium, and High) with three different questions. Subjects were provided with tangible examples to help introduce context into the experiment.	5
Fig. 3	Condition distributions: The image shows the distribution of 235 subjects across 28 conditions. The y-axis represents the seven distinct data-type presentations (see Fig. 1) for images (I), videos (V), and messages (M). The x-axis shows the different contexts (see Fig. 2). The color legend shows the number of subjects corresponding to the binned color group. The figure illuminates the number of subjects per condition.	6
Fig. 4	Kruskal-Wallis p-values across questions: The graphs show p-values calculated from Kruskal-Wallis tests for null hypothesis (same underlying distribution) across questions, per data-type presentation (y-axis) and per context (x-axis). There were a total of 16 p-values evaluated, corresponding to the four data-type presentations by the four contexts. The color legends show the p-values, truncated at 0.50. All p-values for a–c were above 0.01, which indicates that the questions did not result in the same underlying distribution for any of the data-type presentation/context comparisons. P-values shown are for a) images, b) videos, and C) messages.....	8
Fig. 5	Kruskal-Wallis p-values across contexts: The graphs show p-values calculated from Kruskal-Wallis tests for null hypothesis (same underlying distribution) across contexts, per question (y-axis) and per data-type presentation (x-axis). A total of 12 p-values were computed corresponding to the three questions by the four data-type presentations. The color legends show the p-values, truncated at 0.50. All p-values for a–c were above 0.01, which indicates that the contexts did not result in the same underlying distribution for any of the question/data-type presentation comparisons. P-values shown are for a) images, b) videos, and C) messages.	9
Fig. 6	Images data-type medians: The figures show medians calculated from the varying number of subjects per condition (see Fig. 3). The y-axes represent the four different contexts, the x-axes represent the four relevant data-type presentations, and the color bars show the full range of median values for the images data type across the three questions (see Section 2.5).....	10
Fig. 7	Videos data-type medians: The figures show medians calculated from the varying number of subjects per condition (see Fig. 3). The y-axes represent the four different contexts, the x-axes represent the four	

relevant data-type presentations, and the color bars show the full range of median values for the videos data type across the three questions (see Section 2.5)..... 11

Fig. 8 Messages data-type medians: The figures show medians calculated from the varying number of subjects per condition (see Fig. 3). The y-axes represent the four different contexts, the x-axes represent the four relevant data-type presentations, and the color bars show the full range of median values for the messages data type across the three questions (see Section 2.5)..... 12

1. Introduction

Social media viewing or passive social media consumption has been shown to shape an individual's perspective (i.e., opinion).¹ Given the influence social media can have in shaping the individual's opinion, it is important to understand opinion formation, the factors that contribute to the shifts in opinion, and the dynamics associated with the formation and shifting of opinions. Much theoretical work has focused on opinion dynamics through the use of computational models and simulation experiments.^{2-6,8-10} Empirical studies have investigated opinion polarization¹¹ and opinion evolution;¹² however, empirical research is still needed to understand how passive social media consumption of different data types (e.g., images, videos, and messages of pure text) contributes to opinion changes. Specifically, in exploring how people form and change opinions, the various forms of bias are critical considerations.

Content bias has been extensively investigated within the context of media¹³⁻¹⁷ and is in line with the findings associated with social media content.¹² The evidence suggests that a priori perspective and personal experience result in polarization from the consumption of content.¹² From this evidence, it would be of interest to understand how opinions form without the built-in bias of experience related to the content. Our approach aims to investigate how opinions form without the bias of content.

Information media sources are pervasive within most areas where military operations occur. Military information operation analysts manipulate and control the information environment to provide commanders with a decisive advantage over adversaries, threats, and enemies.¹⁸ Determining the types and number of media to present to a populous with intentions of altering opinions is highly relevant to military information operations. The results from this experiment represent a first step toward identifying the effort needed to achieve opinion formation in a given population. Alternatively, commanders are interested in understanding existing opinions held by individuals and social groups within an area of operations to develop appropriate courses of action.

Social influence has been shown to have a strong effect over the formation of an individual's opinion.^{19,20} The term can be interpreted as conformity, peer pressure, and compliance resulting from the expressed opinions or perspectives of other individuals. Within the context of this report, social influence is meant to capture the contribution that a like-minded group makes to opinion formation versus a group holding an opposing perspective.

Context provides a framework for pieces of information, essentially bonding these pieces into a set and resulting in reduced variance over the interpretation and increased coherence of the information as a whole. In the absence of context, pieces of information are subject to individual interpretation, resulting in increased variance of interpretation and a potential general lack of coherence over the whole set. In the study, we address context using levels of controversy. Different levels of controversy helped subjects estimate how pieces of information contribute to the formation of an opinion.

Crowdsourcing is a term used to describe the use of crowds to answer difficult questions or solve difficult problems.²¹ This phenomenon is based on the “wisdom of the crowd” concept, in that a large group of aggregated answers has generally been found to be as good as or often better than a single answer by any person in the group²² or an expert.²⁵ Recently, the method of crowdsourcing was used to help researchers discover the most likely folding patterns for proteins by turning the problem into an online game called Foldit.^{23,24} An explanation for this phenomenon is that noise inherently exists in estimates and an average over a large amount of these noisy estimates results in a reduction in the overall noise, which abides by the law of large numbers in probability theory.²⁶ This makes crowdsourcing an ideal method for estimating empirical thresholds of passive social media consumption for opinion formation.

The aim of this study was to illuminate opinion formation thresholds from passive social media consumption. To eliminate content bias while maintaining information coherence, a general and generic content-free framework with differing contexts (represented by levels of controversy) was established. To the best of our knowledge, no prior work has examined how passive social media data consumption contributes to opinion formation in the absence of content. Furthermore, we used levels of controversy and social influence to understand how opinion formation depends on these factors.

2. Methods

Amazon Mechanical Turk (MTurk) was used to collect the human subject data from 235 participants. A simple computerized task required subjects to enter numbers in boxes that represented their estimates of the amount of data types expected to be viewed in a static time frame (one day) before formulating an opinion, given a specified context. There were 235 subjects randomly assigned to 1 of 28 conditions. A condition consisted of one data-type presentation within one context and with three questions related to the cued information source.

2.1 Human Subjects

The study falls under the US Army Research Laboratory (ARL) internal review board (IRB) Exempt Research Determination for Protocol ARL 17-087, which exempts the study from regulation 32 CFR 219. The research falls into the exemption criteria defined by the Common Rule, which states that human subjects cannot be identified by the collected data; and the responses provided by the subjects place them at no risk of criminal or civil liability, nor could they be damaging to their financial standing, employment, or reputation.

Upon selecting to participate in the study for a quarter (\$0.25), subjects were notified that it would require approximately 3 min to complete, and no personally identifiable information would be collected. Subjects needed to complete all questions in order to be compensated. Twelve demographic questions focused primarily on social media usage, and three task-related questions established the data per participant. Exclusionary criteria consisted of the subjects' general use of social media. If a participant indicated that they did not use social media, they were thanked for their interest in the study and their participation was ended without collecting any data.

User bias was minimized by allowing each subject to participate in the study only once. The MTurk account name was used solely to determine if a subject had attempted to participate previously. If a subject attempted to participate a second time, the program informed them that they were no longer eligible.

2.2 Data Types

Subjects were asked to estimate their opinion formation thresholds for three distinct data types: 1) images, 2) videos, and 3) messages. These data types were selected for their easily identifiable differences. Subjects were shown the following descriptions corresponding to the data types:

- **Messages:** Data type includes text, a tweet, or a post on Facebook.
- **Videos:** Data type includes any moving pictures, animations, GIFs, and videos.
- **Images:** Data type includes still pictures, images, and drawings.

2.3 Data-Type Orientations

To account for potential data-type interactions, we tested seven arrangements or presentations of the data types (Fig. 1). Our general hypothesis regarding data-type interactions is that the population of subjects will not have significantly adjusted

their estimates for the different data-type presentations. This suggests that the estimated threshold for opinion formation from data-type presentation I (images alone) should not be significantly different from the estimates for images in presentations IV, V, or VII (Fig. 1). Therefore, the different data-type presentations were selected to provide sufficient evidence for only utilizing a single presentation in future studies (i.e., data-type presentation VII; Fig. 1).




Fig. 1 Data-type presentations: The different orientations of the representative data types (images, videos, and messages) were used in the study to determine if data-type interactions needed to be considered for opinion formation threshold estimates.

2.4 Contexts

To counterbalance for the abstraction introduced by excluding content from the experiment, context was introduced as levels of controversy. For the purposes of this report, “controversy” and “context” are used interchangeably. The four levels of controversy (None, Low, Medium, and High) were used to investigate influence over opinion formation thresholds. To summarize, the four contexts were the following:

- **None:** No reference to controversy
- **Low:** Minimal controversy (some people would have an opinion)
- **Medium:** Controversial (many people would have an opinion)
- **High:** Highly controversial (most or all people would have an opinion)

With an exception for the “None” case, the different levels of controversy were introduced to the subjects with an example and a color-coded word (Fig. 2). These levels were selected to try and shift the opinion formation threshold estimates of the subjects, essentially allowing us to observe how different contexts influence opinion formation thresholds.

NONE	→ no context referenced
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 US allied country (e.g., France).

Fig. 2 Contexts: Four different contexts (None, Low, Medium, and High) with three different questions. Subjects were provided with tangible examples to help introduce context into the experiment.

2.5 Data Collection

Each MTurk subject was randomly assigned one condition and asked to answer three questions investigating the influence of different sources. This implies that each condition has independent data. The only dependencies between the data are across the three questions a subject answered. The experiment contained seven data-type presentations (see Fig. 1) and four contexts, which resulted in 28 conditions distributed over 235 subjects participating in the study (Fig. 3).

The three questions were used to measure the influence of the source of the social media data type being estimated and included the following sources:

- **Question 1:** Before you FORM an OPINION, how many data types listed below would you expect to view in a day?
- **Question 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?
- **Question 3:** 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 with different viewpoints?

The first question did not specify a source of the social media posts and it was used as a control case. The second question emphasized that the social media data were posted by like-minded people, aiming to measure the influence that in-group posts have on a subject's estimate of opinion formation threshold. The third question emphasized that the social media data were posted by multiple groups with different

perspectives, capturing the combination of in-group and out-group influence. Together, the three questions allow us to measure influence from different sources over opinion formation threshold.

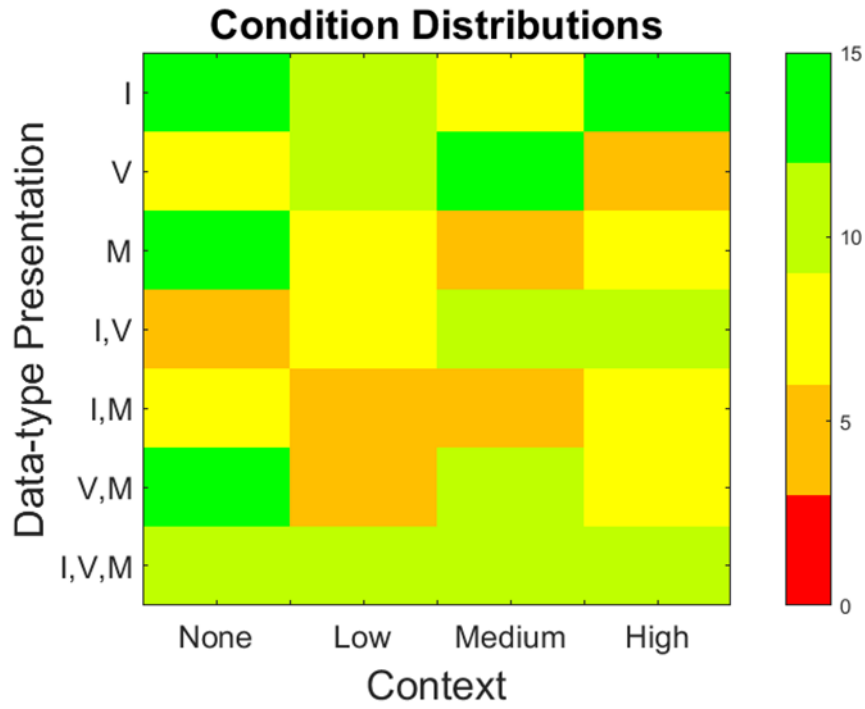


Fig. 3 Condition distributions: The image shows the distribution of 235 subjects across 28 conditions. The y-axis represents the seven distinct data-type presentations (see Fig. 1) for images (I), videos (V), and messages (M). The x-axis shows the different contexts (see Fig. 2). The color legend shows the number of subjects corresponding to the binned color group. The figure illuminates the number of subjects per condition.

2.6 Data Analysis Techniques

Nonparametric tests are needed if the data are not distributed normally. The Jarque-Bera (JB) goodness-of-fit test indicates if a data sample came from an unspecified normal distribution. Therefore, the JB test was used to determine if the data from the respective conditions, questions, data-type presentations, and data types were normally distributed.

The Wilcoxon rank-sum test or the Mann-Whitney U test is a nonparametric test. The null hypothesis we used here for the Wilcoxon rank-sum test states that the distributions of the compared samples are equal. Small p-values indicate that the null should be rejected and the distributions are not equal. The two compared samples are assumed to come from identical and continuous distributions with a possible shift.

To quantify differences between the conditions, the Kruskal-Wallis test was used. This is a nonparametric method similar to a classical one-way analysis of variance (ANOVA). The Kruskal-Wallis test is equivalent to the Wilcoxon rank sum test for equal medians but allows comparisons with more than two groups. Medians of the subject data from the different conditions were compared to determine if the samples are from the same underlying distribution. This test assumes that all samples come from populations having the same continuous distribution and that all observations are mutually independent. In comparison to the Kruskal-Wallis test, a classical one-way ANOVA replaces the “same continuous distribution” assumption with a stronger assumption that the populations are normally distributed. However, in the study, normally distributed data were not observed.

The Kruskal-Wallis test uses ranks of the subject data, rather than numeric values, to compute the test statistics. It finds ranks by ordering the data from smallest to largest across all compared groups, taking the numeric index of this ordering. The rank for a tied observation is equal to the average rank of all observations tied with it. The F-statistic used in classical one-way ANOVA is replaced by a chi-square statistic, and the p-value measures the significance of the chi-square statistic.

The p-values generated from the Kruskal-Wallis test represent the statistical significance associated with all compared data from the respective conditions originating from the same distribution (i.e., the null hypothesis). P-values are considered significant if below 0.01 and corrections for multiple comparisons were not necessary.

3. Results

To assess the factors contributing to opinion formation estimates across the different conditions, the analyses necessitated the segmentation of data by data type (see Section 2.2). With this segmentation, data were measured under four data-type presentations. For example, the image data type (I) was compared to images and videos (I, V), images and messages (I, M), and images, videos, and messages (I, V, M). However, the analysis of data-type presentations for images did not include videos (V), messages (M), or videos and messages (V, M). In contrast, all three data types were compared across the four contexts (see Section 2.4) and three questions (see Section 2.5).

3.1 Jarque-Bera (JB) Tests

The JB test for normally distributed data revealed that less than 50% of the data per data type were normally distributed (46% images; 46% videos; 44% messages),

indicating that nonparametric tests such as Kruskal-Wallis and Wilcoxon rank sum tests are appropriate for further analysis.

3.2 Wilcoxon Rank Sum Tests

Pair-wise Wilcoxon rank sum tests were performed between each data-type presentation (six comparisons) per context, question, and data type. The Bonferroni corrected significance level alpha initially set to 0.01 was corrected to 0.0017. At this significance level, no pairwise compared medians reached significance when correcting for multiple comparisons. This result indicates that thresholds across the four data-type presentations per data type (e.g., data type: images = four data-type presentations: 1. I; 2. I, V; 3. I, M; 4. I, V, M) were not statistically different. Therefore, the medians (i.e., opinion formation thresholds) were not found to be significantly different across the data-type presentations within a data type, context, and question (e.g., data type: videos, context: Low, question: Q1).

3.3 Kruskal-Wallis Tests

The Kruskal-Wallis tests were first conducted across the three questions. This indicates that the subject responses for a data-type presentation (e.g., I, V) and a context (e.g., Medium) were compared across the three questions. The results suggest that the Kruskal-Wallis test for the same underlying distribution was not significant when comparing across contexts (Fig. 4a–c). This indicates that we cannot conclude that the samples are from the same distribution; however, it does not indicate that the samples are significantly different.

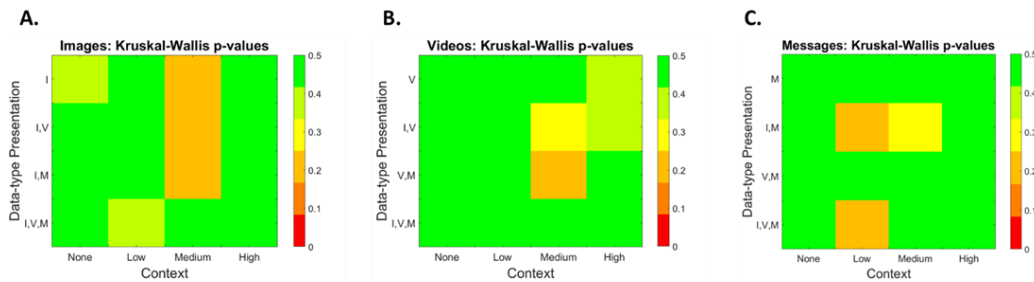


Fig. 4 Kruskal-Wallis p-values across questions: The graphs show p-values calculated from Kruskal-Wallis tests for null hypothesis (same underlying distribution) across questions, per data-type presentation (y-axis) and per context (x-axis). There were a total of 16 p-values evaluated, corresponding to the four data-type presentations by the four contexts. The color legends show the p-values, truncated at 0.50. All p-values for a–c were above 0.01, which indicates that the questions did not result in the same underlying distribution for any of the data-type presentation/context comparisons. P-values shown are for a) images, b) videos, and C) messages.

The p-value comparisons shown in Fig. 4 clearly illustrate the absence of a significant null hypothesis, which states that the compared samples come from the same underlying distribution. A significant null result would indicate that the three questions did not result in different underlying distributions, suggesting that the questions may have had a significant influence over the subjects' opinion formation threshold estimates.

Although none of the comparisons in Fig. 5 reached significance, six of the comparisons did reach marginal significance (see red-colored comparisons in Figs. 5b and c). For the videos data type, the p-values across contexts show that one data-type presentation, “I, V”, reached marginal significance (Fig. 5b) for both No Source Specified ($p = 0.0345$; Question 1) and Multiple Different Perspectives ($p = 0.0217$; Question 3). Similarly, the messages data type showed four marginally significant p-values at data-type presentation “M” for No Source Specified ($p = 0.0356$) and Multiple Different Perspectives ($p = 0.0184$), and for data-type presentation “V, M” for Like-Minded ($p = 0.0700$), and Many Different Perspectives ($p = 0.0225$). Given that none of the p-values were less than 0.01, these results suggest that context may not have an influence over opinion formation thresholds in the aforementioned cases.

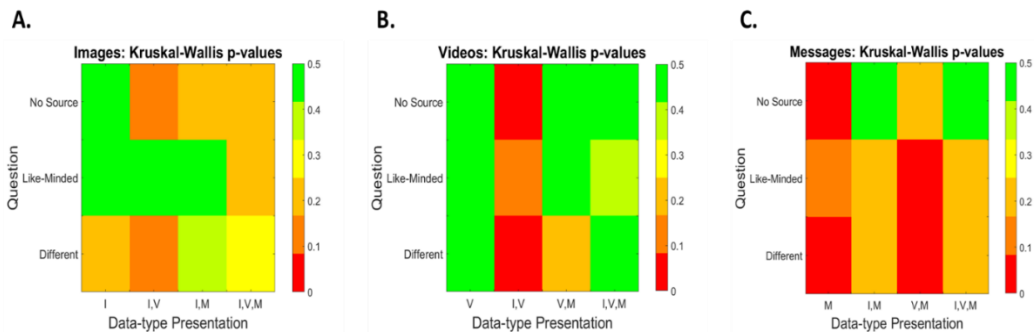


Fig. 5 Kruskal-Wallis p-values across contexts: The graphs show p-values calculated from Kruskal-Wallis tests for null hypothesis (same underlying distribution) across contexts, per question (y-axis) and per data-type presentation (x-axis). A total of 12 p-values were computed corresponding to the three questions by the four data-type presentations. The color legends show the p-values, truncated at 0.50. All p-values for a–c were above 0.01, which indicates that the contexts did not result in the same underlying distribution for any of the question/data-type presentation comparisons. P-values shown are for a) images, b) videos, and C) messages.

Together, the Kruskal-Wallis tests showed that the opinion formation thresholds may have been influenced by source (Questions 1–3) and context (see Fig. 2). In addition, several cases reached marginal significance (see Figs. 5b and c). To better understand the marginally significant cases, additional data could help, given that the number of subjects per condition had large differences (see Fig. 3).

3.4 Opinion Formation Threshold Estimates

Due to the uneven number of data points per condition (see Fig. 3) and the lack of normally distributed data, medians per data type (images, videos, and messages) from the respective conditions were used to determine the opinion formation thresholds. Although possible, no calculated median was zero and across all data types, contexts, questions, and data-type presentations the range is [1, 54]. Overall, the images data type tended to have higher threshold values compared to videos and messages (compare Fig. 6a–c to Figs. 7a–c, and 8a–c). This result indicates that subjects estimated that a larger number of images would be needed to form an opinion.

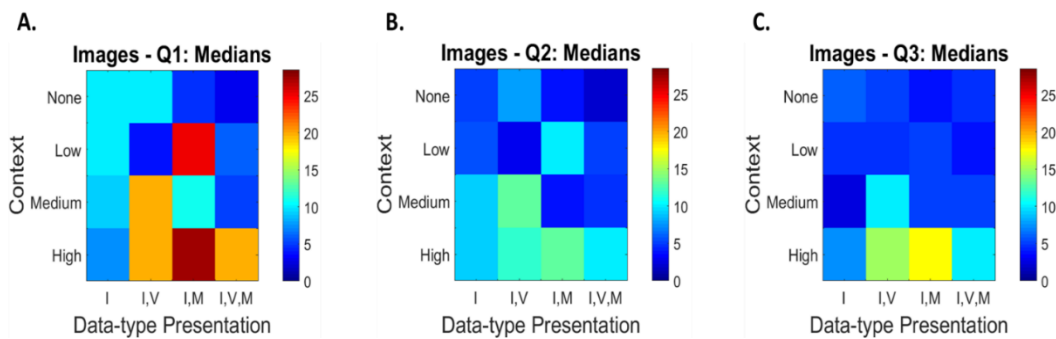


Fig. 6 Images data-type medians: The figures show medians calculated from the varying number of subjects per condition (see Fig. 3). The y-axes represent the four different contexts, the x-axes represent the four relevant data-type presentations, and the color bars show the full range of median values for the images data type across the three questions (see Section 2.5).

Select trends can be observed in the threshold data for the images data type. For example, the context High had a clear increase in median values (compare High row to other contexts in Fig. 6a–c). This result suggests that the context with high controversy shifted the opinion formation thresholds up when compared to the other contexts. Although these values are higher, more data are needed to confirm statistical significance. Similarly, the data in Fig. 6 show that the data-type presentations “I, V” and “I, M” appear to have overall higher thresholds when comparing across contexts and questions. Finally, the thresholds tend to decrease when comparing across questions (compare Fig. 6a to b and b to c). These results, though not necessarily significantly different, show trends that may become significant with additional data.

Surprisingly, some of the trends for the videos data type are similar to the images data type. The data show that the context High has greater threshold values compared to other contexts within a question (compare Fig. 7 High rows to other three contexts), and the context High decreases across questions (compare Fig. 7a to b and b to c). Notably, the maximum median for the videos data type is 10,

indicating that subjects generally estimate lower opinion formation thresholds for the videos data type versus the images data type (compare Fig. 6 to Fig. 7). Furthermore, this result is intuitive since a person would generally expect a video to take more time to view than an image, which implies that someone can see more images versus videos in the same time frame.

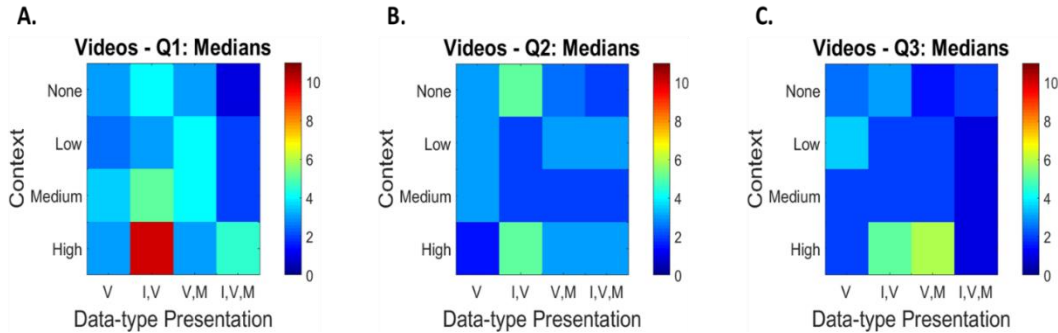


Fig. 7 Videos data-type medians: The figures show medians calculated from the varying number of subjects per condition (see Fig. 3). The y-axes represent the four different contexts, the x-axes represent the four relevant data-type presentations, and the color bars show the full range of median values for the videos data type across the three questions (see Section 2.5).

In contrast to both video and image data types, the messages data type shows different trends. The messages data type shows consistencies between Questions 1 and 3 (compare Fig. 8a to c), whereas Question 2 appears to have a different pattern (compare Fig. 8b to a and c). In addition, the threshold for context Low and data-type presentation “V, M” is approximately five times higher than all other conditions. This can be explained by a small number of subjects ($n = 4$), and by the fact that half the data points could potentially be outliers. The trend persists across the questions because subjects were asked to provide estimates to all three questions, which continued to capture their potentially unreasonable opinion threshold estimates. As with the other data types, additional data are likely to produce results without outliers.

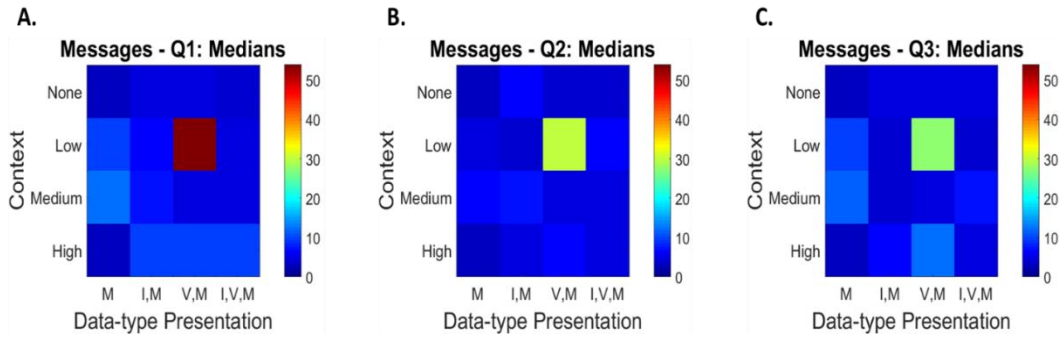


Fig. 8 Messages data-type medians: The figures show medians calculated from the varying number of subjects per condition (see Fig. 3). The y-axes represent the four different contexts, the x-axes represent the four relevant data-type presentations, and the color bars show the full range of median values for the messages data type across the three questions (see Section 2.5).

Together, the medians show tentative opinion formation thresholds derived from the distributions of subject responses to the various conditions. The results from the Wilcoxon rank-sum tests indicate that the data-type presentations did not have a significant effect on subjects’ opinion formation thresholds. However, additional data will be collected to resolve potential outliers within the respective distributions and illustrate potential significant differences between opinion formation thresholds.

4. Discussion and Conclusion

The aim of this study was to illuminate opinion formation thresholds from passive social media consumption. To eliminate content bias while maintaining information coherence, a general and generic content-free framework with differing contexts was established. To the best of our knowledge, no prior work has examined how passive social media data consumption contributes to opinion formation thresholds in the absence of content. Furthermore, we used levels of controversy (see Fig. 2) and social influence (see Section 2.5) to understand how opinion formation depends on these factors.

Evidence suggests that passive social media usage (i.e., viewing without taking action) has an impact on the user’s perspective.¹ This influence over the individual’s perspective could have a significant impact on social events. For this reason, an experimental task was developed to determine thresholds for opinion formation from passive interactions with different social media data types (i.e., images, videos, and messages). The results suggest 1) different presentations of the data-types (see Fig. 1) did not have a significant effect over opinion formation threshold (see Section 3.2), 2) contexts and source information (Questions 1–3) appear to have a differential impact on opinion formation (see Figs. 4 and 5), and

3) opinion formation thresholds do appear to exist and are data-type-dependent (see Figs. 6–8). More data are needed to substantiate conclusions made for results 2 and 3.

Crowdsourcing is an ideal method for understanding phenomena that is inherently noisy like opinion formation threshold estimates. In our experimental paradigm, 235 subjects participated in one of 28 conditions, leaving several conditions relatively underrepresented (see Fig. 3). This underrepresentation was primarily due to randomizing the condition selection criteria and the lack of subjects. In light of this, the 235 subjects were sufficient to test the impact of data-type presentations (see Fig. 1), finding them nonsignificant (see Section 3.2). This implies that future experiments regarding opinion thresholds will not require different data-type presentations.

Some work has been done to understand how small groups of individuals' opinions can influence public opinion,²⁷ but little attention has been focused on the factors that drive an individual toward the formation of an opinion based solely on passive social media consumption. Given the novelty of this work, it is reasonable to remain skeptical about the existence of opinion formation thresholds without additional experimentation. However, the results from this experiment suggest that these opinion formation thresholds do exist and are independent of data-type presentation but dependent upon the data type. Additional data will need to be collected to confirm these results.

5. References

1. Romero DM, Galuba W, Asur S, Huberman BA. Influence and passivity in social media. Heidelberg (Germany): Springer Science+Business Media; 2011; p. 18–33. (Lecture Notes in Computer Science, vol. 6913)
2. Battiston F, Cairoli A, Nicosia V, Baule A, Latora V. Interplay between consensus and coherence in a model of interacting opinions. *Phys D Nonlinear Phenom.* 2016;323–324:12–19.
3. Hegselmann R, Krause U. Opinion dynamics and bounded confidence: Models, analysis and simulation. *JASSS.* 2002;5(3).
4. Deffuant G, Neau D, Amblard F, Weisbuch G. Mixing beliefs among interacting agents. *Adv Complex Syst.* 2000;3(01n04):87–98.
5. Martins ACR. Continuous opinions and discrete actions in opinion dynamics problems. *Int J Mod Phys C.* 2007;19(4):10.
6. Weisbuch G, Deffuant G, Amblard F, Nadal JP. Interacting agents and continuous opinions dynamics. *Lect Notes Econ Math Syst.* 2003;521:225–242.
7. Hegselmann R, Krause U. Opinion dynamics and bounded confidence: models, analysis and simulation. *JASSS.* 2002;5(3):1–15.
8. Duggins P. A psychologically-motivated model of opinion change with applications to American politics. *JASSS.* 2016;20(1).
9. Ramirez-Cano D, Pitt J. Follow the leader: Profiling agents in an opinion formation model of dynamic confidence. *Proceedings of the IEEE/WIC/ACM International Conference on Intelligent Agent Technology; 2006; Hong Kong, China.* p. 660–667.
10. Sznajd-Weron K, Sznajd J. Opinion evolution in closed community. *Int J Mod Phys C.* 2000;11(6):1157–1165.
11. Lee JK, Choi J, Kim C, Kim Y. Social media, network heterogeneity, and opinion polarization. *J Commun.* 2014;64(4):702–722.
12. Xiong F, Liu Y. Opinion formation on social media: an empirical approach. *Chaos.* 2014;24(1):13130.
13. Covert TJA, Wasburn PC. Measuring media bias: A content analysis of time and Newsweek coverage of domestic social issues, 1975–2000. *Soc Sci Q.* 2007;88(3):690–706.

14. DellaVigna S, Kaplan E. The Fox News effect: media bias and voting. *Q J Econ.* 2007;122(3):1187–1234.
15. Baron DP. Persistent media bias. *J Public Econ.* 2006;90(1–2):1–36.
16. Gentzkow M, Shapiro JM. Media bias and reputation. *J Polit Econ.* 2006;114(2):280–316.
17. Gentzkow M, Shapiro N, Stone N. Media bias in the marketplace: theory. NBER Work Pap. 2013:1–28.
18. Headquarters, Department of the Army. Operations. Washington (DC); Headquarters, Department of the Army; 2008 Feb. Field Manual No.: FM 3-0.
19. Cialdini RB, Goldstein NJ. Social influence: compliance and conformity. *Annu Rev Psychol.* 2004;55(1):591–621.
20. Cialdini R, Trost M. Social influence: Social norms, conformity and compliance. New York (NY): McGraw-Hill Companies, Inc.; 1998. p. 151–192. (The Handbook of Social Psychology, vol. 2.)
21. Estellés-Arolas E, González-Ladrón-de-Guevara F. Towards an integrated crowdsourcing definition. *J Inf Sci.* 2012;38(2):189–200.
22. McClellan JE III. The Académie Royale des Sciences, 1699–1793: a statistical portrait. *Isis.* 1981;72(4):541–567.
23. Markoff J. Video gamers tackle protein-folding puzzles in Foldit. *The New York Times.* 2010 Aug 4.
24. Cooper S, et al. Predicting protein structures with a multiplayer online game. *Nature.* 2010;466(7307):756–760.
25. Yi SKM, Steyvers M, Lee MD, Dry MJ. The wisdom of the crowd in combinatorial problems. *Cogn Sci.* 2012;36(3):452–470.
26. Kunda Z, Nisbett RE. Prediction and the partial understanding of the law of large numbers. *J Exp Soc Psychol.* 1986;22(4):339–354.
27. Xie J, Emenheiser J, Kirby M, Sreenivasan S, Szymanski BK, Korniss G. Evolution of opinions on social networks in the presence of competing committed groups. *PLoS One.* 2012;7(3).

List of Symbols, Abbreviations, and Acronyms

ARL	US Army Research Laboratory
IRB	internal review board
I	images
JB	Jarque-Bera
M	messages
MTurk	Amazon Mechanical Turk
V	videos

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

2 DIR ARL
(PDF) IMAL HRA
RECORDS MGMT
RDRL DCL
TECH LIB

1 GOVT PRINTG OFC
(PDF) A MALHOTRA

6 ARL
(PDF) RDRL CII T
D ASHER
J CAYLOR
M MITTRICK
J RICHARDSON
E HEILMAN
E BOWMAN