

# Storytelling in Argumentative Discussions: Exploring the Use of Narratives in ChangeMyView

Sara Nabhani   Khalid Al-Khatib   Federico Pianzola   Malvina Nissim

University of Groningen  
{s.nabhani,khalid.alkhatib,f.pianzola,m.nissim}@rug.nl

## Abstract

Psychological research has long suggested that storytelling can shape beliefs and behaviors by fostering emotional engagement and narrative transportation. However, it remains unclear whether these effects extend to online argumentative discourse. In this paper, we examine the role of narrative in real-world argumentation using discussions from the *ChangeMyView* subreddit. Leveraging an automatic story detection model, we analyze how narrative use varies across persuasive comments, user types, discussion outcomes, and the kinds of change being sought. While narrative appears more frequently in some contexts, it is not consistently linked to successful persuasion. Notably, highly persuasive users tend to use narrative less, and storytelling does not demonstrate increased effectiveness for any specific type of persuasive goals. These findings suggest that narrative may play a limited and context-dependent role in online discussions, highlighting the need for computational models of argumentation to account for rhetorical diversity.

## 1 Introduction

Argumentation is a key form of communication in online spaces, where people often try to justify their beliefs, challenge others' opinions, or persuade readers to reconsider their views. Research in computational argumentation has predominantly focused on identifying argumentative components (Lawrence and Reed, 2019), evaluating argument quality (Wachsmuth et al., 2024), and generating arguments automatically (Saha and Srihari, 2023). However, effective persuasion, arguably one of the primary goals of argumentation, often extends beyond logical structure, relying on rhetorical strategies such as framing, emotional appeal, and narrative (Nabi and Green, 2015). Among these strategies, narrative remains relatively underexplored in existing computational approaches.

In social sciences, narrative has been extensively studied as a persuasive strategy. A substantial body of research highlights the role of narrative transportation in enhancing persuasive effects (Fitzgerald and Green, 2017). Narrative transportation refers to the psychological process through which individuals become deeply immersed in a story, experiencing it as if it were unfolding around them. This immersion has been shown to reduce counter-arguing, increase identification with characters, and render arguments more natural and emotionally resonant. As a result, narrative emerges as a powerful tool for influencing beliefs, attitudes, intentions, and behaviors. While these effects are well-documented in controlled experimental settings (de Graaf et al., 2016; van Laer et al., 2013) and across applied domains such as education (Green, 2004), law (Mazzocco, 2011), and public health (Huang and Green, 2022; de Graaf et al., 2016), it remains unclear whether similar persuasive mechanisms operate in everyday online discourse, where users engage in informal and often unstructured debates across a wide range of topics.

Building on these insights from the social sciences, this paper investigates the role of *narrative* in online argumentation through a large-scale empirical analysis of discussions from the *ChangeMyView* (CMV) subreddit. Our goal is to examine whether the use of narrative in this context reflects the persuasive patterns established in prior research, and to better understand its association with successful persuasion. We adopt a computational approach, leveraging predictions from a story detection model trained on Reddit data to identify narrative across thousands of CMV discussions. We then conduct a series of analyses to explore how narrative presence correlates with *persuasion outcomes* (e.g., successful), *user argumentation skill* (e.g., average), and *targeted change type* (e.g., attitude). The analyses showed that while narrative is not strongly predictive of persuasive success at

the comment or user level, its frequent appearance suggests that it still plays a meaningful role in how users frame and express arguments.

This type of analysis is important for bridging theoretical and applied perspectives on narrative persuasion. It enables us to empirically test assumptions from narrative theory within real-world, user-generated discourse, offering insights into how storytelling functions as a rhetorical strategy in everyday argumentative practice. More broadly, this work contributes to the advancement of argument mining by extending its scope beyond logical structure to include richer rhetorical dimensions, such as narrative, that shape how arguments are constructed and received.

## 2 Related Work

Narratives have been widely studied in social science as a tool for persuasion, but their role in computational models of argumentation has received less attention. Recent efforts, however, have begun to bridge this gap by modeling narrative in argumentative contexts.

Falk and Lapesa (2023) introduce *StoryARG*, a multi-layer annotated corpus that captures both narrative and argumentative dimensions of stories. The corpus includes annotations for narrative elements (e.g., protagonist type, narrative proximity) and argumentative properties (e.g., claim, stance, effectiveness). Stories were collected from *ChangeMyView* (Egawa et al., 2019), *RegulatingRooms* (Park and Cardie, 2018), *Europolis* (Gerber et al., 2018), and *NYT* comments. The argumentative effectiveness of the stories was evaluated through four annotators’ ratings of how persuasive each story was. Their findings suggest that stories proposing solutions tend to be rated as more effective, and that narratives with certain structural properties (e.g. longer text, first-person perspective) are perceived as more persuasive. However, the persuasiveness here is measured via the annotators’ perception rather than real-world behavioral outcomes, and the dataset remains relatively small.

In a follow-up to their earlier work, Falk and Lapesa (2024) study the role of storytelling in online discourse during the COVID-19 pandemic, focusing on how users use personal experiences and narratives in socially tense discussions. Using Reddit data from multiple subreddits, including *ChangeMyView*, they apply the *StoryARG* framework to automatically detect storytelling spans and

classify them by narrative type and argumentative function. The study finds that different types of stories are used in distinct argumentative contexts: personal narratives are common in discussions around social distancing, harm disclosure stories often appear in environmental and social issues threads, solution-oriented stories dominate in homeschooling debates, and background-setting narratives are especially prevalent in conversations involving conspiracy theories. However, their study does not evaluate the persuasive impact of the stories.

Antoniak et al. (2024) present *StorySeeker*, a toolkit and classifier for detecting narrative spans in online discourse. Using a RoBERTa-based model (Liu et al., 2019) fine-tuned on Reddit data, they predict storytelling presence across 1,000 samples from 291 subreddits, showing that narrative is a widespread communicative strategy with over 50% of texts containing stories, with higher prevalence in subreddits centered on personal experience and lower rates in abstract or analytical domains. In a focused case study on *ChangeMyView*, they apply their model alongside topic modeling to examine where and how stories are used in persuasive discourse. They find that topics related to lifestyle and personal decisions tend to receive more storytelling, while abstract topics such as economics receive less. However, they also observe that the presence of narrative does not strongly correlate with persuasive success across topics. Our work complements these studies by taking a broader perspective, examining how narrative use relates to persuasion outcomes at scale using behavioral signals like CMV delta awards.

## 3 Data

Exploring the role of narrative in argumentation requires a source where narrative use is both likely and observable within argumentative discourse. One prominent source of publicly available data for such analysis is the *ChangeMyView* (CMV) subreddit.<sup>1</sup> CMV is an online discussion forum where users post a viewpoint on a given topic and invite others to challenge it. CMV is designed to promote open, thoughtful conversations in which participants are encouraged to reconsider their stances in response to sound counterarguments.

To indicate when persuasion has occurred, CMV employs a delta system. When a commenter successfully changes the mind of the original poster

<sup>1</sup><https://www.reddit.com/r/changemyview>

Statistic	Value
Total posts	20,436
Unique original posters (OPs)	13,704
Total comments	1,017,724
Awarded deltas	11,643
Avg. deltas per post	0.5697
Avg. comments per post	49.9619
Avg. unique participants per post	23.7417

Table 1: Summary statistics of the *ChangeMyView* dataset after preprocessing.

(OP) or another participant, the person whose view was changed responds with a delta symbol ( $\Delta$ ). We treat the presence of a delta as an explicit signal that a comment was persuasive. An example of a CMV thread and delta interaction is shown in Figure 1.

CMV hosts a diverse range of discussion topics from everyday issues and personal decisions to political, ethical, and philosophical debates. This diversity makes it a valuable resource for studying how people, often without formal training in debate, engage in argumentation, persuasion, and opinion expression through various rhetorical strategies, including the use of narrative.

For this study, we use the CMV corpus built by Tan et al. (2016), which comprises all CMV threads posted between January 2013 and August 2015. Although the corpus is split into training and held-out subsets, we include both in our analysis.

Prior to the analysis (see Section 5), we applied several preprocessing steps to clean the dataset. We removed all empty comments and those marked as deleted by the system, as they lack textual content. We also excluded threads created by moderators or system accounts, which typically serve administrative or community-oriented purposes such as gathering feedback, announcing new features, or outlining policy changes, and do not represent authentic attempts to present or defend a personal view. Summary statistics of the final dataset are provided in Table 1.

## 4 Analysis Method

We conduct our study to analyze the use of narrative in CMV discussions by following a structured methodology. First, we employ a computational model to detect the presence of narrative elements in CMV posts and comments. Next, we categorize these texts based on key criteria such as persuasion

outcome, the argumentation skills of the author, and the type of change being targeted. We then apply our identification model to each group to examine the presence, frequency, and intensity of narrative use across these different contexts. Finally, we synthesize our findings to draw informed conclusions about the role and patterns of narrative in CMV interactions.

### 4.1 Narrative Identification Method

To identify the presence of narrative in CMV posts and comments, we use the story detection model introduced as part of the *StorySeeker* toolkit<sup>2</sup> by Antoniak et al. (2024).

The model is based on a RoBERTa-base language model (Liu et al., 2019), fine-tuned on a dataset of Reddit posts sampled from the Webis-TLDR-17 Reddit corpus (Völske et al., 2017). The training data was annotated by two expert annotators, who labeled whether each post contained storytelling, as well as the specific spans that formed the narrative, using the definition of a story as “a sequence of events involving one or more people” provided in the annotation guideline (Antoniak et al., 2024).

The model was trained as a binary classifier to distinguish between texts that contain storytelling and those that do not. On the expert-annotated test set, it achieved strong performance: an F1 score of 0.86 for the *story* class and 0.88 for the *non-story* class. These scores were averaged across five cross-validation folds and demonstrated low variance, indicating the model’s stability and reliability. In addition to Reddit-based evaluation, the model was tested on several non-Reddit datasets. Despite domain differences, it maintained strong performance, showing good generalizability across various text types and topics.

We selected this model due to its high performance and close alignment with our data domain. Since both the training corpus and our target dataset for analysis are drawn from Reddit, the language style, tone, and structure are closely matched. This domain similarity, combined with the model’s demonstrated effectiveness, makes it particularly well-suited for detecting narrative content in *ChangeMyView* arguments.

<sup>2</sup><https://huggingface.co/mariaantonik/storyseeker>

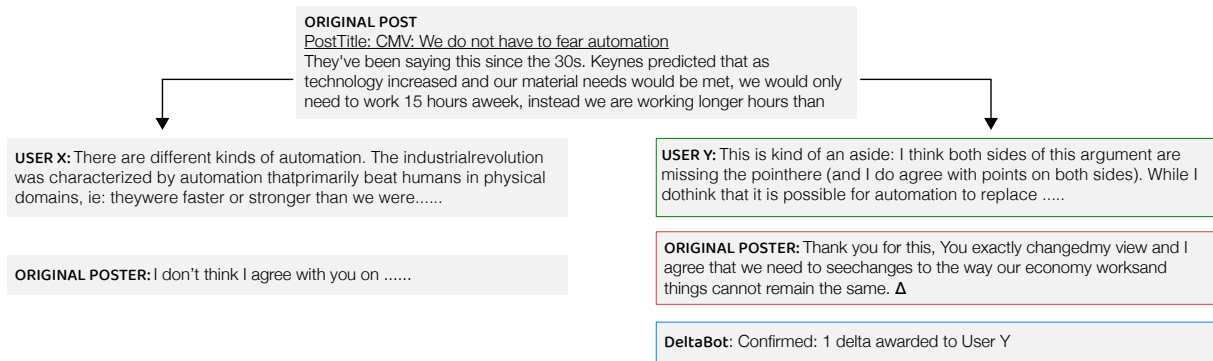


Figure 1: An example CMV discussion thread. The original poster (OP) expresses a belief, and other users respond. The green comment is the one that received a delta, indicating it successfully changed the OP's view. The red comment is the OP's response awarding the delta. The blue comment is an automatic confirmation from DeltaBot, Reddit's system for tracking awarded deltas.

## 5 Analyses and Results

Using the narrative predictions generated by the *StorySeeker* model, we conduct an analysis of how narrative is employed in argumentation on *Change-MyView*. Our investigation spans multiple levels including the comment level, user level, and discussion level to uncover broader patterns in narrative usage. Specifically, we examine: (1) the relationship between narrative and comment-level persuasiveness, (2) variations in narrative use based on debater effectiveness, and (3) shifts in narrative use across different types of targeted change. The following subsections detail our analytical approach and provide a discussion of the findings.

### 5.1 Narrative Use in Persuasive Comments

Our first analysis is conducted at the comment level, investigating whether persuasive comments (i.e., those that received a delta) are more likely to contain narrative elements than non-persuasive ones. This provides a direct means of assessing the potential relationship between storytelling and persuasive success in CMV argumentation. To address this question, we employ the binary predictions of the *StorySeeker* model to classify each comment in the CMV discussions as either narrative or non-narrative. We then compare the proportion of narrative comments between two groups: those that received a delta and those that did not. In addition, we examine the degree of narrativity among comments labeled as *story*, using the model's confidence scores as a proxy for narrative strength.

**Results:** Our findings reveal that 5.06% of delta-awarded comments are labeled as narrative, compared to 4.63% of non-delta comments. This indi-




Figure 2: Narrative confidence scores for comments labeled as *story*, grouped by whether the comment received a delta.

cates that storytelling is slightly more prevalent in persuasive arguments. To assess the significance of this difference, we conduct a chi-square test of independence, which yields a statistically significant result ( $p = 0.0384 < 0.05$ ). However, the effect size, measured using Cramér's  $V$ , is very small ( $V = 0.0025$ ), suggesting a weak association between narrative presence and persuasive outcome. Table 2 presents the distribution of both groups across the full dataset.

The results also show that non-delta comments have a slightly higher average narrative score ( $M = 0.8048$ ,  $SD = 0.1444$ ) than delta-awarded comments ( $M = 0.7914$ ,  $SD = 0.1459$ ). A Mann-Whitney U test indicated that this difference is statistically significant ( $p = 0.0183 < 0.05$ ), with a very small effect size ( $Rank - Biserical Correlation = 0.0598$ ). Figure 2 shows the distribution of narrative confidence scores across the two groups.



Comment Type	Narrative		Non-narrative		Total	
	Count	%	Count	%	Count	%
<b>Delta-awarded</b>	525	0.05%	9,844	0.97%	10,369	1.02%
<b>Non-delta</b>	46,628	4.58%	960,727	94.40%	1,007,355	98.98%
<b>Total</b>	47,153	4.63%	970,571	95.37%	1,017,724	100.00%

Table 2: Distribution of narrative and non-narrative comments across delta-awarded and non-delta comments. Percentages reflect proportions of the full dataset.

## 5.2 Narrative Use in Delta-Awarded Discussions

For this analysis, we move beyond individual comments to the level of full discussion threads, aiming to assess whether narrative use is more prevalent in persuasive conversations overall. Specifically, we compare threads where at least one delta was awarded (*delta threads*) to those without any deltas (*non-delta threads*) in terms of both narrative density and narrativity degree. For each thread, we compute the following metrics:

- **Narrative density:** the proportion of comments within the thread classified as *story* by the *StorySeeker* model.
- **Average degree of narrativity:** the mean narrative confidence score across all comments labeled as *story* within the thread (Steg et al., 2022).

**Results:** Threads that received deltas exhibit a higher average narrative density (0.0486) than those that did not (0.0407), and this difference is statistically significant ( $p < 0.001$ , *Rank – Biserial Correlation* = 0.1340). Figure 3 shows the distribution of narrative density across both groups, highlighting a greater concentration of narrative-heavy threads among delta threads.

As for the degree of narrativity, delta threads also display a slightly higher average confidence score (0.7900) compared to non-delta threads (0.7884). However, this difference is not statistically significant ( $p > 0.05$ ). The distribution of narrativity degree across groups is illustrated in Figure 4.

## 5.3 Narrative Use by Delta-Awarded vs. Other Participants

In this analysis, we investigate whether users who successfully persuaded someone in a thread (i.e., those who received a delta) tend to use more narrative than other participants within the same discussion. Focusing exclusively on delta-awarded




Figure 3: Narrative density per thread, grouped by whether a delta was awarded in the thread.




Figure 4: Degree of narrativity per thread, measured as the average confidence score for comments labeled as *story*, grouped by whether a delta was awarded in the thread.

threads, we compute two metrics for each user per thread: narrative density and degree of narrativity. We then compare these metrics between delta recipients and non-recipients within each thread.

**Results:** Interestingly, users who received a delta used slightly less narrative than their peers in the same persuasive threads. The average narrative density for delta recipients was 0.0544 ( $SD = 0.1703$ ), compared to 0.0560 ( $SD = 0.2064$ ) for non-recipients. This difference is statistically significant ( $p < 0.001$ , *Rank – Biserial Correlation* = 0.0474).

Similarly, the degree of narrativity was lower




Figure 5: Degree of narrativity per user in delta-awarded threads, among users who wrote at least one narrative comment, grouped by whether a user was awarded a delta in the post

for delta recipients ( $M = 0.7866$ ) than for non-recipients ( $M = 0.8037$ ), and this difference was also statistically significant ( $p < 0.001$ , *Rank – Biserial Correlation* = 0.0862). Figure 5 shows the distribution of narrativity degree across both user groups.

#### 5.4 Narrative and User Persuasiveness

Prior work on persuasion in *ChangeMyView* has categorized users based on their overall effectiveness in persuasion, measured by the frequency with which their comments receive deltas (Wiegmann et al., 2022). Following this classification, we examine the relationship between a user’s persuasive skill and their use of narrative. Users are grouped into three categories:

- **Good debaters:** at least 5% of their comments received deltas.
- **Average debaters:** between 0% and 5%.
- **Poor debaters:** 0% of their comments received deltas.

For each user, we compute two metrics: narrative density and degree of narrativity. These metrics are then compared across the debater categories.

**Results:** Narrative density varies significantly across debater groups, as indicated by a Kruskal–Wallis H test ( $p < 0.001$ ,  $\epsilon^2 = 0.0487$ ). As shown in Table 3, poor debaters exhibit the highest average narrative density (0.0714), followed by good debaters (0.0673), and average debaters (0.0445). All pairwise differences are statistically significant ( $p < 0.001$ ). However, the distributional patterns point to a more nuanced interpretation: the median narrative density for both good and poor debaters

Debater Type	Count	Density	Narr. Deg.
Good	2063	0.0673	0.7915
Average	2373	0.0445	0.7921
Poor	68040	0.0714	0.8081

Table 3: Narrative use statistics by debater quality.




Figure 6: Narrative density by debater quality.

is 0.0, suggesting that many users in both groups rarely employ narrative. Poor debaters also show the greatest variability ( $SD = 0.2058$ ), with a long tail of users who employ narrative heavily, thereby skewing the group mean upward. These patterns are visualized in Figure 6, which highlights the presence of heavy tails and outliers, particularly among poor debaters.

Narrativity degree also differs significantly among debater groups, again confirmed by a Kruskal–Wallis H test ( $p < 0.001$ ,  $\epsilon^2 = 0.0071$ ). Poor debaters have the highest average narrativity degree (0.8081), followed by average debaters (0.7921), and good debaters (0.7915). Pairwise comparisons show that poor debaters use significantly stronger narrative than both good and average debaters ( $p < 0.001$ ), while the difference between good and average debaters is not statistically significant ( $p = 0.019$ ). Figure 7 illustrates the distribution of narrativity degree across the three groups.

#### 5.5 Narrative Use and Effectiveness Across Persuasion Goals

Prior work in narrative persuasion suggests that the effectiveness of storytelling may vary based on the type of change being targeted, such as shifts in beliefs, attitudes, intentions, and behaviors (Zebregs et al., 2015; Green and Appel, 2024). To investigate whether similar patterns hold in the context of *ChangeMyView*, we analyze narrative use across different persuasion goals. We begin by classifying each original post according to the type of change




Figure 7: Narrativity degree by debater quality. Poor debaters show higher narrative confidence scores overall.

it aims to achieve: *belief*, *attitude*, *intention*, or *behavior*. For each post, we compute two key metrics: narrative density and degree of narrativity. These values are compared across the four goal types to examine how narrative use differs by persuasion context. Finally, to assess the effectiveness of narrative in each goal category, we compare its usage in persuasive (delta-awarded) versus non-persuasive (non-delta) comments within each group.

For the classification task, we employ a zero-shot classifier<sup>3</sup> (Laurer et al., 2024), trained as a universal classifier using Natural Language Inference (NLI). The model evaluates whether a given post (the premise) entails a hypothesis representing each potential goal type. We use the hypothesis template: “This post is primarily about changing someone’s [LABEL]”, with the placeholder replaced by one of the candidate labels. The label with the highest entailment score is assigned as the persuasive goal for the post. To ensure label reliability, we include only posts where the top prediction exceeds a confidence threshold of 0.7. Appendix A presents two *ChangeMyView* posts after classification, including the persuasive goals assigned by the model and example comments from the discussions.

**Results:** Narrative density varies significantly across persuasion goals ( $p < 0.001$ ,  $\epsilon^2 = 0.0167$ ). Posts targeting behavior and belief change exhibit the highest average narrative density (0.0549 and 0.0532, respectively), while posts targeting attitude and intention show lower values (0.0341 and 0.0405, respectively). Pairwise comparisons confirm that behavior and belief posts include significantly more narrative than intention posts ( $p < 0.001$ ); no other pairwise differences are statistically significant.

By contrast, narrativity degree does not differ

<sup>3</sup><https://huggingface.co/MoritzLaurer/deberta-v3-base-zeroshot-v1.1-all-33>

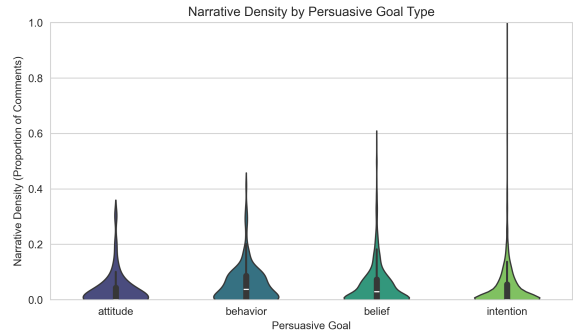


Figure 8: Narrative density by persuasive goal type.

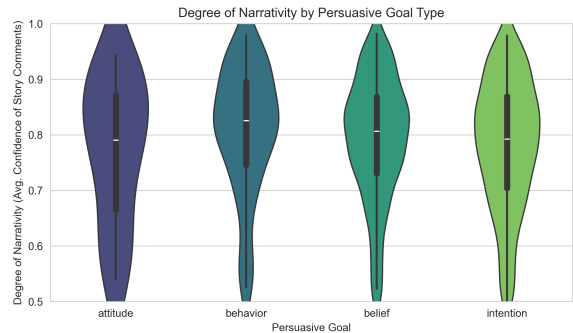


Figure 9: Narrativity degree by persuasive goal type.

significantly across persuasion goals ( $p > 0.05$ ), and none of the pairwise comparisons reach significance. Figures 8 and 9 show the distributions of narrative density and degree of narrativity, respectively. Summary statistics for narrative use across persuasion goal types are shown in Table 4.

Regarding narrative effectiveness, we find no significant differences in the proportion of narrative use between delta-awarded and non-delta comments within any of the goal categories ( $p > 0.05$ ). Likewise, the degree of narrativity does not differ significantly between persuasive and non-persuasive comments across all goal types. Full results are presented in Tables 5 and 6.

## 6 Discussion

The results of our analysis show that while narrative is present in online persuasive discourse, its association with success is weak and more limited than found in prior theory. We find a small but statistically significant increase in the proportion of narrative comments among delta-awarded comments, yet the effect size is negligible. Also, delta-awarded comments tend to have a slightly lower narrativity degree than the non-delta ones, suggesting that stronger storytelling, measured by model

Goal	Count	Density	Narrativity
Attitude	44	0.0341	0.7733
Behavior	212	0.0549	0.8073
Belief	1157	0.0532	0.7920
Intention	1149	0.0405	0.7826

Table 4: Narrative use by persuasive goal type.

Goal	Delta%	Non-Delta%
Attitude	0	0.0401
Behavior	0.0667	0.0625
Belief	0.0610	0.0550
Intention	0.0377	0.0432

Table 5: Proportion of narrative comments in *Delta* vs. *Non-delta* comments across goal types

confidence, does not guarantee persuasiveness.

At the thread level, posts where a delta was awarded tend to contain more narrative overall. However, when we zoom in to the user level within the same discussions, we find that users who received deltas used slightly less narrative than other participants. This finding is counterintuitive if narrative was a consistent indicator of persuasion success, and suggests that successful arguers may rely more on other rhetorical techniques, or that narrative alone is not sufficient to persuade.

The patterns observed at the user level support this idea. We found that poor debaters use narrative more than good ones, and use stories with higher narrativity. A potential explanation for this could be that relying on narrative may be a fallback strategy for users who struggle to persuade with reasoning, but this needs to be confirmed by additional research.

The analysis of the narrative use across different types of persuasive goals showed that stories are more frequently used in posts aiming to change beliefs or behaviors. Yet, this difference in usage does not translate into effectiveness: narrative comments are not more likely to succeed within any goal type, and narrativity degree does not differ between persuasive and non-persuasive comments in these contexts.

Overall, these findings suggest that narrative is a visible but not consistently persuasive part of online argumentation. Its effectiveness likely depends on how it is used, who the audience is, and what norms govern the discussion space.

Several limitations should be taken into account when interpreting these findings. First, our analysis

Goal	Delta	Non-Delta
Attitude	-	0.7761
Behavior	0.7435	0.8189
Belief	0.8220	0.8160
Intention	0.7903	0.7981

Table 6: Narrativity degree of comments in *Delta* vs. *Non-delta* comments across goal types

depends on the predictions of an automatic story detection model. Although the model was trained on Reddit data and demonstrates strong performance on held-out benchmarks, it may still overlook more subtle or implicit forms of storytelling. Second, the classification of persuasion goals relies on zero-shot predictions from an NLI-based model. To enhance reliability, we included only posts with high-confidence predictions; however, the absence of human validation may introduce noise. Third, we operationalize persuasion success through delta awards. While deltas serve as a useful community-driven signal of agreement, they do not capture all forms of influence. Some persuasive comments may go unrecognized, while others may receive deltas for reasons unrelated to argumentation quality. Finally, our findings are specific to the *Change-MyView* (CMV) platform, which has distinct rules, cultural norms, and moderation practices. These factors influence how arguments are constructed and what is considered persuasive. As such, the generalizability of our results to other platforms or domains remains an open question.

## 7 Conclusion

In this paper, we have examined the role of narrative in online argumentation through a large-scale analysis of persuasive interactions on *Change-MyView*. By applying an automatic story detection model to user comments, we have explored how narrative has been employed in argumentative contexts and whether it contributes to persuasive success. Our findings have shown that narrative is indeed used in argumentative discourse. However, its presence alone has not consistently predicted successful persuasion. These results suggest a more complex relationship between storytelling and persuasion than what has often been assumed in theoretical or experimental work.

Future research could build on these insights by examining how narrative interacts with other rhetorical strategies. Also, distinguishing between



types of stories, such as personal experiences versus hypothetical scenarios, and analyzing how they are integrated into arguments may offer a deeper understanding of when and how narrative enhances persuasion. Extending this analysis to other platforms may also help identify the broader conditions under which storytelling supports persuasiveness.

## References

- Maria Antoniak, Joel Mire, Maarten Sap, Elliott Ash, and Andrew Piper. 2024. [Where do people tell stories online? story detection across online communities](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7104–7130, Bangkok, Thailand. Association for Computational Linguistics.
- Anneke de Graaf, José Sanders, and J.A.L. Hoeken. 2016. [Characteristics of narrative interventions and health effects: A review of the content, form, and context of narratives in health-related narrative persuasion research](#). *Review of Communication Research*, 4:88–131.
- Ryo Egawa, Gaku Morio, and Katsuhide Fujita. 2019. [Annotating and analyzing semantic role of elementary units and relations in online persuasive arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 422–428, Florence, Italy. Association for Computational Linguistics.
- Neele Falk and Gabriella Lapesa. 2023. [StoryARG: a corpus of narratives and personal experiences in argumentative texts](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2350–2372, Toronto, Canada. Association for Computational Linguistics.
- Neele Falk and Gabriella Lapesa. 2024. [Stories and personal experiences in the COVID-19 discourse](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15320–15340, Torino, Italia. ELRA and ICCL.
- Kaitlin S. Fitzgerald and Melanie C. Green. 2017. [Chapter 3. Narrative persuasion: Effects of transporting stories on attitudes, beliefs, and behaviors](#), pages 49–67. John Benjamins Publishing Company.
- Marlène Gerber, André Bächtiger, Susumu Shikano, Simon Reber, and Samuel Rohr. 2018. [Deliberative abilities and influence in a transnational deliberative poll \(europolis\)](#). *British Journal of Political Science*, 48(4):1093–1118.
- Melanie C. Green. 2004. [Storytelling in teaching](#). *APS observer*, 17.
- Melanie C. Green and Markus Appel. 2024. [Chapter one - narrative transportation: How stories shape how we see ourselves and the world](#). volume 70 of *Advances in Experimental Social Psychology*, pages 1–82. Academic Press.
- Yan Huang and Melanie Green. 2022. [Reducing covid-19 vaccine hesitancy among african americans: The effects of narratives, character’s self-persuasion, and trust in science](#). *Journal of Behavioral Medicine*, 46.
- Moritz Laurer, Wouter van Atteveldt, Andreu Casas, and Kasper Welbers. 2024. [Building efficient universal classifiers with natural language inference](#). *Preprint*, arXiv:2312.17543.
- John Lawrence and Chris Reed. 2019. [Argument mining: A survey](#). *Computational Linguistics*, 45(4):765–818.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#). *Preprint*, arXiv:1907.11692.
- Philip J. Mazzocco. 2011. [Narrative persuasion in legal settings: What’s the story?](#)
- Robin L. Nabi and Melanie C. Green. 2015. [The role of a narrative’s emotional flow in promoting persuasive outcomes](#). *Media Psychology*, 18(2):137–162.
- Joonsuk Park and Claire Cardie. 2018. [A corpus of eRulemaking user comments for measuring evaluability of arguments](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Sougata Saha and Rohini Srihari. 2023. [ArgU: A controllable factual argument generator](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8373–8388, Toronto, Canada. Association for Computational Linguistics.
- Max Steg, Karlo Slot, and Federico Pianzola. 2022. [Computational detection of narrativity: A comparison using textual features and reader response](#). In *Proceedings of the 6th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, pages 105–114, Gyeongju, Republic of Korea. International Conference on Computational Linguistics.
- Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. 2016. [Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions](#). In *Proceedings of WWW*.
- Tom van Laer, Ko de Ruyter, Luca M. Visconti, and Martin Wetzels. 2013. [The extended transportation-imagery model: A meta-analysis of the antecedents and consequences of consumers’ narrative transportation](#). *Journal of Consumer Research*, 40(5):797–817.

Michael Völske, Martin Potthast, Shahbaz Syed, and Benno Stein. 2017. [TL;DR: Mining Reddit to learn automatic summarization](#). In *Proceedings of the Workshop on New Frontiers in Summarization*, pages 59–63, Copenhagen, Denmark. Association for Computational Linguistics.

Henning Wachsmuth, Gabriella Lapesa, Elena Cabrio, Anne Lauscher, Joonsuk Park, Eva Maria Vecchi, Serena Villata, and Timon Ziegenbein. 2024. [Argument quality assessment in the age of instruction-following large language models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 1519–1538, Torino, Italia. ELRA and ICCL.

Matti Wiegmann, Khalid Al Khatib, Vishal Khanna, and Benno Stein. 2022. [Analyzing persuasion strategies of debaters on social media](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6897–6905, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Simon Zebregs, Bas van den Putte, Peter Neijens, and Anneke de Graaf and. 2015. [The differential impact of statistical and narrative evidence on beliefs, attitude, and intention: A meta-analysis](#). *Health Communication*, 30(3):282–289. PMID: 24836931.

## A Example CMV Threads

Table 7 presents examples of *ChangeMyView* (CMV) threads targeting different persuasive goals, with comments annotated for narrative presence and persuasive success (delta).

Type	Content
<b>Title</b>	I'm financially stable, but I don't intend to give to charity until I'm older and have gathered even more money. CMV.
<b>Post</b>	I could: - Never give any money to charity - Give lots of my money to charity right now (i.e., give up my worldly possessions so that others may thrive) - Give a small amount of money now, and on a regular, continuing basis (e.g., every Christmas give out X% of my current net worth) - Invest now, and give lots of my money later (perhaps all of it, donated via my will) The last option appears wisest to me. It seems like it will get the most money out to charity in the long term. It also allows me a cushion in case hard times happen to hit me, which is likely enough. I think that people seem to encourage the 3rd option (continue giving, regularly). Why is my choice not the best? Is there a better choice, and I just didn't think of it to list here?
<b>Comment 1</b> <i>Delta</i> <i>Non-Story</i>	Charity works both ways. It helps the people you give it to, and it also helps you by changing your heart and making you more caring. Every time you put off good deeds, you are not just delaying the receipt of those good deeds, you are also becoming a less caring person. You are developing habits that make you less likely to help others around you as well as less likely to give to charity later. If you genuinely believe that your investment strategies are so much better than the strategies of any other charitable foundations (are you Warren Buffett?) then by all means invest your donations. Create a separate bank account (or foundation - if you're that good you'll need one) and donate to that account/foundation, acting as custodian of that money in terms of investment choices but knowing that you've already set it aside for others and that it is no longer yours. And by all means, recognize that most charity is not money. It is good deeds.
<b>Comment 2</b> <i>Non-Delta</i> <i>Non-Story</i>	I can give a defense of option 3, which is potentially the best one. Think of it this way- nonprofits have budgets, and they have to try and stay within them, like everyone else. Charities have a chronic problem of being unsure about what their budget will look like. Giving money regularly gives those charities budgeting power, meaning not only will they have more money, but they can use that money more effectively and consistently because they'll have a better understanding of what their finances will look like next month or next year. Big one-time donations are great, don't get me wrong, and it certainly helps, but it doesn't give them the same kind of budgetary power.

(a) Example of a thread targeting **intention change**

Type	Content
<b>Title</b>	CMV: 'Be yourself' is a totally useless piece of advice.
<b>Post</b>	What does it even mean? How can one not be one-self? Now I'm mostly talking about the everyday use of the phrase – where it is typically offered moments before an uncomfortable or unfamiliar social interaction, like a date. And in this scenario, I assume it means something like 'be yourself but without any of the debilitating neurosis and subsequent façade'. But this is just as useless. Don't you think if people could magically wish away such impediments, they would have already done so? It's possible that some people are able to do this – but then they presumably wouldn't be in need of such banal advice. But even if you are held back by anxiety / hiding behind a façade etc. – are these things not an integral part of the ego that constitutes the 'self'? ...
<b>Comment 1</b> <i>Delta</i> <i>Story</i>	So how can we change your view about this? You don't think there are people who need this advice, but I can anecdotally tell you that it took <i>several</i> people telling me <i>this exact advice</i> before it finally 'clicked' and I started acting more like my usual self on dates. For a long time I would take a girl out and then try to act like the person I thought she wanted me to be. It took my mother, father, and a few close friends, and then finally an actual <b>date</b> of mine, before I finally got the message that I really should just "be myself".
<b>Comment 2</b> <i>Non-Delta</i> <i>Story</i>	Oh sure, it's plainly obvious. Simply stating it in a nonchalant manner probably won't do anything. But I had a friend who was the opposite of me. We were very similar (humor, likes, dislikes, what we like to do for fun) except I would get anxious and this guy was fearless. By a kind of lead by example "be yourself" philosophy he lived by, he rubbed off on me slowly over time and I became less anxious. So yeah, I guess a simple statement won't do much, or maybe it's the slap on the face someone needs to get motivated. But the philosophy behind it definitely helped me personally gravitate towards confidence and normalcy.

(b) Example of a thread targeting **belief change**

Table 7: Examples of CMV threads targeting **intention change** and **belief change**.