Who is the hero and who is the villain?

Detection of roles in news articles by using sentiment analysis

Diego Gomez-Zara
Northwestern University
USA
dgomezara@u.northwestern.edu

Miriam Boon
Northwestern University
USA
MiriamBoon2012@u.northwestern.edu

Larry Birnbaum
Northwestern University
USA
l-birnbaum@northwestern.edu

ABSTRACT

Many news articles use tropes to present people, organizations, and facts. These narrative frames follow cultural archetypes, where readers can associate each one of the presented elements with familiar stereotypes, well-known characters, and recognizable outcomes. In this way, authors can cast real people or organizations as villains or heroes. We describe a system that uses sentiment analysis of words co-occurring with identified entities to determine when an entity is being cast as a villain or as a hero. Our hope is that by informing readers when an entity is cast in one of these roles, we can make implicit bias explicit, and assist readers in applying their media literacy skills.

CCS CONCEPTS

• **Computing methodologies** → Natural language processing • **Human-centered computing** → Natural language interfaces

KEYWORDS

Computational journalism; tropes; contextual search; sentiment analysis.

1 INTRODUCTION

A frequent problem in understanding news lies in comprehending the broader context of the particular news story or other informational content. One influential way that the media may shape public opinion is by *framing* people, events, and issues in particular ways. By using narrative structures and cultural references, framing involves a communication source presenting and defining an issue within a field of meaning [1]. Framing helps to introduce events by matching readers' prior knowledge, cultural narratives, and moral values [2]. Also, it presents the people involved in the story as heroes, villains, and victims, and with that, readers can anticipate and comprehend their attitudes, beliefs, decisions, and actions. In sum, narrative frames are a critical part of how readers understand new things in terms of older things, and make sense of the causes, events, and consequences reported by news articles [3].

When narrative frames are used in news articles, they can both clarify and entertain. But when misused, they can mislead and introduce bias that can be difficult for a reader to recognize. For example, by framing a story to match with a hero or a villain, the author can dispose readers towards positive or negative impressions of real people or organizations without ever explicitly criticizing them. As media contributes to public perceptions, the presence of frames can lead to misinformation and misperceptions.

Based on computational techniques, we propose an information system able to detect how the characters of news stories are presented by journalists, and classify the frame applied to them. By doing so, we aim to contribute by detecting the use of narrative frames, in order to distinguish potential bias in news articles. We view this as the first step towards research that leads to understanding the effects of these frames on the readers' perceptions.

2 RELATED WORK

Narrative frames are not only used in films or literary genres; they are also frequently used in news media. They respond to a pervasive cultural mode that in which narrative structures the presentation of political discourse and identities [4, 5]. Narrative frames are a mode of popular culture narrative that employ emotionality to provide an unambiguous distinction between good and evil through clear designations of victimization, heroism, and villainy [6]. Based on these emotional presentations, we took a sentiment analysis approach to detect who is designated as hero and villain in a particular news story. By using sentiment analysis, we aim to understand how the characters of the story are framed and presented to the readers.

Sentiment analysis have been used to detect emotions and attitudes relating to reviews [7], consumption [8], reputation [9], credibility [10], stock markets [11], and elections [12]. Most applications simply return a score for the event, entity or person in the study. These scores have little meaning to the average user, however. We seek to classify entities into the semantic categories of "heroes" and "villains," categories which are much more likely to be comprehensible to out typical user.

3 A MODEL FOR ROLE DETECTION

We have developed a framework for role detection. This framework comprises a platform that instantiates the following core procedures:

3.1 Entity recognition. In this step, we proceed to identify the main characters involved in a news story. The system takes news article headlines and text as inputs, and analyzes sentence by sentence detecting and tokenizing relevant entities. Characters can be a *Person*, an *Organization*, or only a *Position* (such as "the witness" or "the police officer"). As news articles have determined narrative styles, we used grammar structures to detect chunks, entities, and their relations. The relations associated with *Persons* and *Organizations* are demonyms (i.e. "The *French* president…"), countries (i.e. "The president of *Spain…*"), organizations (i.e. "The CEO of *Google…*"), and positions (i.e. "The *judge* Mark…"). We used the entity recognition methods provided by the package Natural Language Toolkit (NLTK) [13]. We saved all the relations found for each entity in a temporal database. Then, we analyze the content of the news article and its headline.

One difficulty in applying entity detection in this context is that some news articles work with previously unreferenced names that *a-priori* are novel for knowledge databases (i.e., the name of a victim), or using proper nouns that are not part of English dictionaries (i.e. "Brexit"). For this, we support the entity recognition task with an *n-gram* and proper nouns search. Based on frequencies, we check if there are any new concepts relevant to the news structure, and add them to our entity list.

3.2 Sentiment analysis. After entities are detected, we analyze the use of emotional words appearing near the identified entities. Culturally, heroes are idealized for courage, outstanding achievements, or noble qualities, while villains are idealized by evil intentions, plotting, and having unacceptable purposes. Narrative structures use positive words to describe heroes and negative words for villains. By using *TextBloob*[1], a Python library for processing textual data, we calculate the polarity and subjectivity of each sentence. While polarity score is a float within the range from -1.0 (very negative) to 1.0 (very positive), the subjectivity is a float within the range 0.0 (very objective) to 1.0 (very subjective). Scores come from a default classification defined in *TextBloob* package, where the use of certain words add more polarity and subjectivity to the sentence. For each one of the entities recognized, we assign to each entity (a) the polarity score of the entire sentence, and (b) the sentiment score of the closest verbs, adjectives, and adverbs to that entity.

3.3 Entity cut-off. Once we recognize all the entities and related words, we must detect and merge the multiple references to a given entity, and select the most relevant entities in the corpus. First, we merge different mentions under the same concept and select the longest entity having proper nouns (i.e. the best results for "Mr. Smith," "John Smith," and "Smith," would be "John Smith"). During this process, we count the number of times that all the entities are mentioned in the text, and save their locations in a list using the indexes of their corresponding sentences. Secondly, we calculate the relevance score of each merged entity as:

$$r_i = \alpha h_i + \frac{n_i}{\|s\| * f_i}$$

where r_i is the relevance score of the merged entity *i*, h_i is a dummy variable where 1 means that the merged entity is mentioned in the headline, n_i is the number of times that the merged entity is mentioned, $\|s\|$ is the number of sentences in the text, and f_i is the first location where the merged entity is located. In other words, we assume that entities that appear at the beginning of the text are more relevant. The coefficient (α) is to weight the relevance of the news article's headline and body.

3.4 Role assignment. Finally, we look who are the main characters of the story. At this point, we filter all the entities that are not Persons, Organizations, or Positions. Then, by taking the relevance and sentiment scores, we assign the hero as the most relevant entity from the top half of the most positive entities, and the villain as the most relevant entity from the bottom half of the most negative entities. Using these two metrics, we establish the most relevant positive character as the hero, and the most relevant negative character as the villain.

5 EXAMPLES

To test our model "in the small", we analyze four different news articles published between July 11th and 18th 2017. While the first two were published by the New York Times, the other two were covered by the Washington Post. The first article discloses the meeting between Donald Trump, Jr. and a Russian lawyer to attack Hillary Clinton's campaign. The second article corresponds to President Trump's visit to France, where he was received by the French president Emmanuel Macron. The third news article refers to the shooting of a 40-year-old Australian woman by a Minneapolis Police officer. Finally, the last article refers to a clash between Ann Coulter and Delta Airlines. Table 1 shows the classification results.

In the first case, the narrative structure did not follow clearly a hero-villain structure. Hillary Clinton was classified as a hero, but she would be better classified as a victim, which our system currently does not allow for. Between "the Russians" and Donald Trump Jr., the latter was considered as the most relevant entity. Although the villain classification would have been correct, *TextBloob* did not identify any sentiment-laden words related to these two entities. This does not indicate that the article did not cast these individuals as villains. It did. However, it did not do so by using obviously negative, sentiment-laden words. Instead, it used implicit language, and relied on the reader's understanding of the specific sequence of events and conversations described to communicate the harmful intentions of these meetings. These are pragmatic elements. That is, they are not communicated directly through the semantics–meaning of words–or syntactics–structure of phrases and sentences. Pragmatic aspects of meaning are notoriously difficult to detect computationally. Our work aims to develop a simple to compute heuristics that would allow us to recover some of this pragmatic content.

[1] http://textblob.readthedocs.io/en/dev/

For the second news article, the model selected the French president as the hero and the U.S. president as the villain. By reading the news article, we verified that Mr. Macron was presented as a positive world-leader aiming to continue working on the current environmental trades. In contrast, Mr. Trump was criticized for stepping back American efforts and research funding on global warming.

At time of publication of the third news article, the police officer's name had not yet been revealed by the Minneapolis Police Department. Nonetheless, our model was able to identify "the police officer" as the villain. Our model was able to reconcile references to "the Australian woman" with her name, Justine Damond. It also identified her as a hero, however, when again she would better be described as a victim.

Finally, in our fourth news article, Ann Coulter was identified as the villain. This is not entirely inaccurate. However, the article described a complex situation in which Ann Coulter began as a victim but rapidly transitioned into being the villain. Our system could not capture a situation where someone shifted from victim or hero to villain.

Table 1: *Classification results*

Headline	Hero	Villain
"Russian dirt on Clinton? 'I love it,' Donald Trump Jr. said"	Hillary Clinton (r: 0.49, s: 0.16)	D. Trump Jr. (r: 0.16, s: 0.0)
"Emmanuel Macron to welcome Trump, an unlikely partner, to France"	E. Macron (r: 0.75, s: 0.26)	Donald Trump (r: 0.66, s: -0.01)
"After police officer fatally shoots Australian woman, her relatives plead for answers"	J. Damond (r: 0.11, s: 0.24)	Police officer (r: 0.40, s: -0.20)
"Ann Coulter had to switch seats on a Delta flight. Then came the tirade."	Delta Airlines (r: 0.41, s: 0.07)	Ann Coulter (r: 0.82, s: -0.07)

5 FUTURE WORK

At present, we are leveraging standard sentiment analysis, which as a simple first approach yielded some promising results. However, because it situates our entities along a one-dimensional spectrum, it cannot allow for situations where there is a victim. Therefore, the most obvious next step we can take is to move to a semantic model of classification that consists of three categories: "Hero," "Villain," and "Victim".

We improved our original results by incorporating disambiguation techniques and taking into consideration news story headlines. However, most of the time, authors used cultural cross-references to present the characters of news stories, and the meaning of these characters depend strongly on their previous actions and contexts Because of these difficulties, more work is required to detect implicit value judgments.

We also have plans to integrate this system with the Northwestern InfoLab News Tropes Project [3]. NTP seeks to identify the use of tropes as narrative frames for news articles. These tropes range from biblical, such as "David and Goliath", to fables, such as "Setting the fox to guard the henhouse". Many of them include heroes, villains, and victims. Thus, role detection could be used to confirm the presence of a trope–e.g. if the trope has a hero and the article has no hero, then the article cannot have that trope. It could also be used for elaboration–e.g. we know that this trope has a villain (Goliath) and a hero (David), and so by identifying the hero and villain in an article known to have this specific trope, we can say a lot about the entities cast in these roles.

ACKNOWLEDGMENTS

Our work at the Intelligent Information Lab (InfoLab) at Northwestern University has been supported by the National Science Foundation, under grant #IIS-0917261; and by Google Inc. Miriam Boon has also been supported by a Cognitive Science Fellowship from Northwestern University.

REFERENCES

[1] De Vreese, C. H. News framing: Theory and typology. *Information Design Journal & Document Design*, 13, 1 (2005).
[2] Green, M. C. Transportation into narrative worlds: The role of prior knowledge and perceived realism. *Discourse processes*, 38, 2 (2004), 247-266.
[3] Boon, M. L. *Augmenting Media Literacy with Automatic Characterization of News along Pragmatic Dimensions*. ACM, City, 2017.
[4] Pan, Z. and Kosicki, G. M. Framing analysis: An approach to news discourse. *Political communication*, 10, 1 (1993), 55-75.
[5] Scheufele, D. A. Framing as a theory of media effects. *Journal of communication*, 49, 1 (1999), 103-122.
[6] Anker, E. Villains, victims and heroes: Melodrama, media, and September 11. *Journal of Communication*, 55, 1 (2005), 22-37.
[7] Feldman, R. Techniques and applications for sentiment analysis. *Communications of the ACM*, 56, 4 (2013), 82-89.
[8] Uhl, M. W. Explaining U.S. consumer behavior with news sentiment. *ACM Trans. Manage. Inf. Syst.*, 2, 2 (2011), 1-18.
[9] Manaman, H. S., Jamali, S. and AleAhmad, A. Online reputation measurement of companies based on user-generated content in online social networks. *Computers in Human Behavior*, 54 (2016), 94-100.
[10] Kawai, Y., Fujita, Y., Kumamoto, T., Jianwei, J. and Tanaka, K. Using a sentiment map for visualizing credibility of news sites on the web. In *Proceedings of the Proceedings of the 2nd ACM workshop on Information credibility on the web* (Napa Valley, California, USA, 2008). ACM, [insert City of Publication],[insert 2008 of Publication].
[11] Bollen, J., Mao, H. and Zeng, X. Twitter mood predicts the stock market. *Journal of Computational Science*, 2 (2011 2011), 1-8.
[12] Vargo, C. J., Guo, L., McCombs, M. and Shaw, D. L. Network issue agendas on Twitter during the 2012 US presidential election. *Journal of Communication*, 64, 2 (2014), 296-316.
[13] Bird, S., Klein, E. and Loper, E. *Natural language processing with Python: analyzing text with the natural language toolkit.* " O'Reilly Media, Inc.", 2009.

[14] Lewis, J., Bartlett, A. and Atkinson, P. Hidden in the Middle: Culture, Value and Reward in Bioinformatics. *Minerva*, 54, 4 (Dec 2016), 471-490.
[15] Ettema, J. S. Crafting cultural resonance. *Journalism*, 6, 2 (2005), 131-152.
[16] Berger, A. A. *Narratives in popular culture, media, and everyday life*. Sage, 1997.
[17] Finlayson, M. A. *The Story Workbench: An Extensible Semi-Automatic Text Annotation Tool*. City, 2011.
[18] Garten, J., Hoover, J., Johnson, K. M., Boghrati, R., Iskiwitch, C. and Dehghani, M. Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior Research Methods* (2017), 1-18.
[19] Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D. and Manning, C. *A multi-pass sieve for coreference resolution*. Association for Computational Linguistics, City, 2010.