Proceedings of the Nineteenth International AAAI Conference on Web and Social Media (ICWSM 2025)

Opposites Attract? Ambivalence in Distinguishing Real and Fake News and Predicting their Spread

Mostofa Najmus Sakib, Francesca Spezzano, Anne Hamby

Boise State University - Boise (Idaho), USA

mostofanajmussak@u.boisestate.edu, {francescaspezzano, annehamby}@boisestate.edu

Abstract

This study expands the understanding of linguistic cues in fake news by exploring ambivalent language's role in distinguishing real from fake news and its impact on news spread. Unlike prior research focusing on positive or negative language separately, this work hypothesizes that fake news may exhibit higher ambivalence, aligning with its association with high-arousal emotions and writers' efforts to attract attention. Ambivalence, traditionally viewed as the co-occurrence of conflicting positive and negative elements, is extended here to include diverse dimensions such as textual, temporal, psychological, and content-related ambivalence.

Introduction

Fake news, defined as intentionally false articles (Ireton and Posetti 2018), threatens modern democracies, prompting efforts to understand its origins, spread, and impact (Lazer et al. 2018). Many people consume news online but struggle to assess its truth (Spezzano et al. 2021). Hence, research has focused on distinguishing real from fake news, using linguistic features and automated text analysis (Asubiaro and Rubin 2018)(Singh, Ghosh, and Sonagara 2021).

One feature that has reliably surfaced as a predictor of fake news is the presence of negative language (Gupta et al. 2022) (Damstra et al. 2021). However, research on attitudinal ambivalence – the degree to which people hold positive and negative evaluations at the same time – suggests that a focus on the valence of information alone may be too narrow (Priester and Petty 1996). According to this literature, the joint presence of positive and negative information may cause a reader to feel mixed or conflicted that is not prompted by the presence of negative (or positive) information alone. However, focusing on the mere presence (absence) of negative language overlooks the potential existence of ambivalent language.

In this paper, we build on prior work by exploring how ambivalent language – beyond just positive or negative tones – can help distinguish real from fake news, contributing to the growing literature on linguistic cues and fake news. There is reason to expect that fake news may feature higher levels of ambivalence (rather than merely higher levels of negative language): past research illustrates fake news' broader association with higher arousal emotions (Zhang et al. 2021) (Damstra et al. 2021) and writers' motives to attract and sustain attention (Zhang and Ghorbani 2020), all of which are potential outcomes of ambivalence (Rothman et al. 2017). We also examine whether the presence of ambivalence predicts the spread of news, across multiple datasets covering different subject areas. Furthermore, we make a contribution to the literature on ambivalence. Classic research conceptualizes ambivalence as the joint presence of positive and negative unipolar elements whose simultaneous presence creates conflict. Other unipolar and potentially conflicting elements exist in language as well (e.g., work-leisure ambivalence, emotional-cognitive ambivalence, past-present ambivalence, etc.). We thus conceptually expand the concept of valence-related ambivalence to introduce and test the effects of new operationalizations of ambivalence in the context of distinguishing between real and fake news and predicting news spread on social media.

This research connects fake news detection and ambivalence theory, showing how conflicting language features can signal fake news and predict its viral spread. It also situates findings within social science literature, exploring how ambivalence interacts with human psychology and social media dynamics to shape information dissemination.

Related Work

Language in Fake News. Research examining the linguistic characteristics of fake news has shown that it differs from real news in several ways: fake news articles tend to be shorter in length, with fewer and shorter words and sentences (Asubiaro and Rubin 2018), and use excess punctuation and capitalization (Bradshaw et al. 2020). Fake news is also more likely to feature first- and second-person pronouns (Singh, Ghosh, and Sonagara 2021) and more likely to feature more novel language than regular news (Vosoughi, Roy, and Aral 2018). Fake news is also reliably more likely to use affective and particularly negative language (Damstra et al. 2021), featuring more fear-, anger-, and sadness-related words than real news (Gupta et al. 2022). Thus, past research shows that valence, emotion, and content-related features have been shown to distinguish between real and fake news. One consistent feature of past research is that it examines the mere presence (absence) of linguistic features in their ability to predict fake news. In this work, we investigate whether the joint presence of some of these elements may be diagnostic.

News Sharing. Past work has repeatedly shown that emotional, arousing content is more likely to be shared (Berger 2011) (Vosoughi, Roy, and Aral 2018). While some past research on news article sharing has shown that the level of belief that the articles were truthful predictors likelihood of sharing (Di Domenico et al. 2021), other research suggests that social media sharing behavior can actually be unrelated to judgments about the article's accuracy (Pennycook and Rand 2021). Social media content may distract people from prioritizing truth when they decide what to share, favoring aspects that they perceive as more likely to engage others. For instance, Chen, Pennycook, and Rand (2021) found that message evocativeness predicted message-sharing intentions, and ambivalent messages may be more vivid and imagery-inducing than their non-ambivalent counterparts.

Attitudinal Ambivalence. Early research saw attitudinal ambivalence as purely positive or negative, but later work revealed that combined positive and negative evaluations create a distinct state that cannot be captured by the independent components alone. Attitudinal ambivalence is the degree to which people explicitly hold both positive and negative evaluations at the same time (Kaplan 1972). This situation often gives rise to a state of internal conflict, such that people feel torn about the attitude object (Priester and Petty 1996). Attitudinal ambivalence is aversive (Nordgren, Van Harreveld, and Van Der Pligt 2006), causing discomfort, especially when action is required. To ease this, individuals seek consistency, as cognitive dissonance theory shows inconsistency is unpleasant (Abelson et al. 1968). Consequently, prior research shows ambivalence can lead to biased thinking and impulsive actions (Hamby and Russell 2022). While traditionally viewed as an individual psychological state, we extend this concept to language, as discussed next.

Definition of Ambivalence and its Extension

In this paper, we examine whether fake news articles feature a higher level of **textual ambivalence** – that is, "the state of having two opposing feelings at the same time, or being uncertain of how [one feels]" (Cambridge Dictionary, 2018). Ambivalence does not have to be confused with neutrality, which is the absence of any strong feelings, whereas ambivalence focuses on strong opposing sensations (Schneider and Schwarz 2017). Past text analysis work routinely captures the valence of the text, and as aforementioned, one feature that has reliably surfaced as a predictor of fake news is the presence of negative language (Gupta et al. 2022) (Damstra et al. 2021). We thus move beyond the textual features examined by past research on fake news to consider how textual ambivalence may distinguish between real and fake news.

We computed textual ambivalence using a formula widely utilized in prior ambivalence research, as introduced by Priester and Petty (1996): $\frac{P+N}{2} - |P - N|$, where P and N refer to positive and negative sentiment, respectively, and are computed on the news text by using the VADER tool (Hutto and Gilbert 2014). By subtracting the absolute value of the difference between P and N, we help distinguish ambivalence from neutrality—if P and N are both large, then the value of ambivalence is high, whereas if the only P is large, then the value of ambivalence is low. There is reason to predict that textual ambivalence may be higher in fake (vs. real) news, due to fake news writers' motives to create articles. A primary motivation is financial: fake news may be created to encourage clicks and revenues from advertising. Fake news writers aim to attract and sustain readers' attention to enable its spread through social media (Zhang and Ghorbani 2020), and ambivalence has the potential to increase arousal and enhance information spread (Berger 2011).

In addition to extending the concept of attitudinal ambivalence to the context of news text (i.e., textual ambivalence), we leverage the fundamental concept to introduce new bases of ambivalence, based on fundamental aspects of news. First, all news is temporally embedded (Neiger and Tenenboim-Weinblatt 2016) and may adopt different tenses (past, present, and future). We refer to the joint presence of multiple tenses in an article as **temporal ambivalence** (*past-present* ambivalence, *past-future* ambivalence, *present-future* ambivalence). Second, news often includes a narrator who reveals mental status, such as thoughts and feelings (Van Krieken and Sanders 2021). We refer to the joint presence of conflicting mental states in an article as **psychological states ambivalence** (*emotion-cognitive* ambivalence, *happiness-sadness* ambivalence). Finally, news articles can discuss topics that may vary in terms of their substantive content (Zhao et al. 2011). We refer to the joint presence of conflicting content in an article as **content-related ambivalence** (e.g., *male-female* ambivalence; *work-leisure* ambivalence). In each case, the aforementioned formula used to calculate traditional ambivalence was used to calculate the values of the new ambivalence definitions. For example, the ambivalence values for Emotional vs. Cognitive were calculated in the following way using the scores provided by the LIWC tool (Tausczik and Pennebaker 2010): $\frac{T+C}{2} - |T - C|$, where T and C are the tone and cognitive dimensions, respectively. Table 5 in the Appendix lists the new ambivalence measures we defined, along with the LIWC dimensions used to compute them.

Research Questions. To study how classical ambivalence and its extension influence fake news prediction and news sharing, we pose the following research questions:

RQ1: Is fake (vs. real) news higher in (a) textual ambivalence, (b) temporal ambivalence, (c) psychological states ambivalence, and (d) content-related ambivalence?

RQ2: Is (a) textual ambivalence, (b) temporal ambivalence, (c) psychological states ambivalence, and (d) content-related ambivalence associated with news sharing behavior?

Experimental Framework

In order to study the impact of textual ambivalence and its extensions on fake news detection and news sharing prediction we proceeded as follows.[1] We used classical machine-learning approaches to (i) classify fake (vs. real) news and (ii) predict real and fake news sharing. For the latter, we addressed the problem as a four-class classification problem

[1]To enhance **reproducibility**, our code is available at https://anonymous.4open.science/r/Ambivalence-4104 Reproducing all the experiments on Google Colab with CPU takes about 5 hours.

Dataset	# Articles	# Fake Articles
PolitiFact (Shu et al. 2020)	935	417
GossipCop (Shu et al. 2020)	19,759	4,734
ReCOVery (Zhou et al. 2020)	2,029	665
Fake Health (Dai, Sun, and Wang 2020)	2,296	763

Table 1: Datasets considered in our study.

where the number of shares is divided into four bins of percentiles (25, 50, 75, 100), and the article labels are assigned according to the bin the number of shares of an article belongs to. Moreover, when predicting news sharing, we addressed the cases of fake and real news separately, so that we can compare and contrast the two cases.

We used LIWC plus textual, temporal, psychological states, and content-related ambivalence features in input to a Random Forest (RF) classifier to perform classification.[2] Our goal is not to outperform state-of-the-art models prediction models, but to examine how ambivalence influences fake news and sharing outcomes. Hence, we used LIWC features in input to our classification models as this is a well-established method in content-based fake news related research (Zhou and Zafarani 2020). Next, we used the mean decrease in impurity method to study whether textual ambivalence and its extensions were important features for the considered classification problems. If so, we performed Shapley value-based explanations (SHAP) to analyze how they contribute to fake news and news sharing prediction.

We considered the four popular datasets related to politics, gossip, and health reported in Table 1. Each dataset provides news text, label (fake or real), and the number of shares as tweets/retweets on the Twitter (now X) platform.

Results and Discussion

Table 2 reports the accuracy and F1 score of the Random Forest classifier (with the LIWC and 8 ambivalence features in input) we used to predict fake news and the number of shares of real and fake news on the considered datasets. Results are compared with a Dummy classifier generating predictions uniformly at random from the list of observed classes. As we can see, the accuracy achieved for fake news detection is always above 82% in three out of four datasets. Similarly, the considered classifier achieves an accuracy ranging between 30% and 58% when predicting news shares (four-class classification problem). In comparison, the Dummy classifier ranges between 18% and 26%.

Tables 3 and 4 summarize the impact of our considered ambivalence features on the detection of fake news and prediction of news shares. An ambivalence feature is reported in the tables if it appears among the top-30 most important features (sorted by importance value from top to bottom). The tables also report the results of the explainability analysis conducted with SHAP to find out whether a higher or lower value of a feature predicts fake or real news or is associated with high or low volumes of shares. Due to page limitations, the corresponding SHAP plots are reported in the Appendix.

[2] We also tested Logistic Regression and SVM but report only RF (200 estimators, random state = 100) as it performed best.

Dataset	Prediction Problem	Accuracy		F1	
		RF	Dummy	RF	Dummy
Politifact	Fake News	86.95	45.98	86.85	46.08
	# of Share (Fake)	38.12	17.85	31.96	17.78
	# of Share (Real)	37.25	20.19	28.71	20.98
GossipCop	Fake News	82.72	48.10	66.29	51.71
	# of Share (Fake)	57.52	23.33	32.35	26.02
	# of Share (Real)	30.96	24.22	27.48	24.55
ReCOVery	Fake News	82.55	47.78	76.19	48.81
	# of Share (Fake)	46.01	24.81	36.71	27.85
	# of Share (Real)	38.05	25.64	36.13	25.65
Fake Health	Fake News	69.99	48.69	61.65	49.69
	# of Share (Fake)	29.88	23.52	29.09	23.45
	# of Share (Real)	34.17	21.82	33.85	21.72

Table 2: Accuracy and F1 score for predicting Fake News and # of Shares.

The symbol "↑" indicates high values and the symbol "↓" indicates low values. We have left the direction blank when no strong relationship was found with an ambivalence feature.

RQ1: Ambivalence and Fake News Detection

We refer to Table 3 to answer RQ1. Several ambivalence features appear important to predict fake news in three of the considered datasets. These include features in the temporal, content-related, and psychological states (emotional-cognitive) ambivalence categories. Also, for the majority of these features, higher values are associated with fake news. Surprisingly, no ambivalence-related features are important to predict fake news in the ReCOVery dataset, while textual ambivalence appears as an important feature only in GossipCop. Looking more closely at the relationship between ambivalence and fake news prediction, we observe that: (i) a high presence of textual ambivalence in the news text predicts gossip-related fake news, and (ii) a high presence of emotion-cognitive ambivalence in the news text predicts people-related fake news (GossipCop and Politifact).

RQ2: Ambivalence and News Sharing

We refer to Table 4 to answer RQ2. We observe ambivalence features from all the defined categories appearing as important features to predict news shares, with textual ambivalence being important to predict news sharing. More specifically, our results show that: (i) a high presence of textual ambivalence increases gossip-related fake news sharing and decreases gossip-related real news sharing; (ii) a high presence of temporal past-present ambivalence generally increases news sharing (real and fake news in FakeHealth, and real news in GossipCop); (iii) high presence of temporal past-feature ambivalence increases real news sharing (GossipCop and FakeHealth); (iv) sad-happy and work-leisure ambivalence features are important to predict health-related real news sharing (ReCOVery and Fake Health), even if they show opposite behavior in the two datasets.

Discussion and Limitations

In addressing RQ1, ambivalence types differ in their association with real and fake news depending on context. This variation stems from how ambivalence functions in different genres and aligns with audience expectations and con-

Dataset	Ambivalence Features (Top 30)	Feature Value	News Prediction
Politifact	Past-Future	↑	Real
	Emotion-Cognitive	↑	Fake
	Work-Leisure	↑	Fake
	Past-Present	↑	Fake
	Female-Male	↑	Real
GossipCop	Emotion-Cognitive	↑	Fake
	Present-Future	↑	Real
	Female-Male	↑	Fake
	Textual	↑	Fake
ReCOVery	-		
Fake Health	Present-Future	↑	Fake
	Past-Present	↑	Real
	Work-Leisure		

Table 3: Impact of Ambivalence on Fake News Prediction

tent strategies in those contexts. For instance, "classic" (textual) ambivalence is associated with fake news only in the gossip-related dataset. Gossip content often thrives on emotional and sensational elements to capture attention and provoke curiosity. The presence of both positive and negative words may make the story appear more dramatic or intriguing, which aligns with the characteristics of fake gossip news designed to attract clicks. In political or health news, ambivalence might not function as a signal of falsity because ambivalence is often seen as an indicator of complexity or nuance, which is expected in real reporting. For example, a health article might balance the pros and cons of a medical treatment, and a political article might present arguments from both sides of a debate. The findings also indicated that emotional-cognitive ambivalence was associated with fake news in gossip and political contexts. Gossip often thrives on emotional engagement, but adding cognitive elements (e.g., reasoning or justification) can make fake stories feel more credible. For instance, a fake gossip story might combine emotional hooks (e.g., "shocking," "outrageous") with cognitive elements (e.g., "as reported by insiders") to both intrigue and rationalize. Fake political news often aims to provoke strong emotional responses (e.g., outrage, fear) while providing cognitive justifications to align with readers' preexisting beliefs or biases. For example, a story might say, "This policy will devastate families (emotion), and experts predict economic collapse if it continues (cognition)."

In addressing RQ2, textual ambivalence boosts sharing in fake gossip news but reduces it in real gossip. Fake gossip thrives on sensationalism, with mixed emotions making stories more provocative and shareable. In contrast, real gossip news demands clarity; ambivalence may appear as uncertainty or hesitation, undermining credibility and engagement. Additionally, past-present ambivalence tends to boost the sharing of both real and fake news. It may link past events to current situations, making the content feel relevant and rich in context. Readers may be more likely to share articles that connect historical context with current events because it provides a sense of continuity and depth (regardless of whether it is real or fake). Fake news may use past-present ambivalence to heighten credibility and drama, while real news may leverage it to provide depth and context. Finally, we found that sad-happy and work-leisure ambivalence types reduce news sharing in one health dataset

Dataset	News Type	Ambivalence Features (Top 30)	Feature Value	# of Share
Politifact	Fake	Textual		
		Female-Male		
		Past-Present		
	Real	Work-Leisure		
		Female-Male		
		Past-Future		
		Textual		
GossipCop	Fake	Past-Future	↓	↑
		Female-Male	↓	↑
		Present-Future	↑	↑
		Textual	↑	↑
		Work-Leisure		
	Real	Present-Future	↑	↑
		Textual	↑	↓
		Past-Present	↑	↑
		Emotional-Cognitive		
ReCOVery	Fake	Past-Present		
		Female-Male	↑	↓
		Textual		
	Real	Emotional-Cognitive	↑	↑
		Work-Leisure	↑	↓
		Sad-Happy	↑	↓
Fake Health	Fake	Past-Present	↑	↑
		Present-Future		
		Emotion-Cognitive		
		Work-Leisure		
	Real	Work-Leisure	↑	↑
		Past-Present	↑	↑
		Textual		
		Sad-Happy	↑	↑
		Present-Future		

Table 4: Impact of Ambivalence on News Share Prediction.

but increase it in another. These differing effects likely arise from audience perceptions, content framing, and platform dynamics. In one case, such ambivalence may seem confusing or inappropriate, lowering trust and engagement, while in the other, it may convey nuance or relatability, boosting shareability. These outcomes depend on how ambivalence aligns with audience expectations, emotional tone, and the health topic's context.

The present work is not without limitations. Although past work also does so, the operationalization of ambivalence focusing on linguistic measures may not be a sufficiently nuanced approach to capture these constructs. Additionally, while the used datasets cover politics, gossip, and health, they may not fully represent real and fake news in domains like science, finance, or global events.

Conclusions

Our study highlighted the role of ambivalence in distinguishing real from fake news and predicting news sharing behavior across different domains. Textual ambivalence, particularly in gossip-related fake news, and emotional-cognitive ambivalence in gossip and political fake news, emerged as key indicators of fakery. These findings highlight ambivalence's role in attracting attention and boosting engagement, impacting fake news detection and virality.

Acknowledgments

Part of this work has been supported by the US National Science Foundation award #1950599. We thank Rebecca Cataldo, Amir Marashifar, Jonathan Zhang, and Hyunji Jeong for their early contributions to this research.

References

Abelson, R. P.; Aronson, E. E.; McGuire, W. J.; Newcomb, T. M.; Rosenberg, M. J.; and Tannenbaum, P. H. 1968. Theories of cognitive consistency: A sourcebook.

Asubiaro, T. V.; and Rubin, V. L. 2018. Comparing features of fabricated and legitimate political news in digital environments (2016-2017). *Proceedings of the Association for Information Science and Technology*, 55(1): 747–750.

Berger, J. 2011. Arousal increases social transmission of information. *Psychological science*, 22(7): 891–893.

Bradshaw, S.; Howard, P. N.; Kollanyi, B.; and Neudert, L.-M. 2020. Sourcing and automation of political news and information over social media in the United States, 2016-2018. *Political Communication*, 37(2): 173–193.

Chen, X. C.; Pennycook, G.; and Rand, D. G. 2021. What makes news sharable on social media? Pre-Print.

Dai, E.; Sun, Y.; and Wang, S. 2020. Ginger Cannot Cure Cancer: Battling Fake Health News with a Comprehensive Data Repository. *ICWSM*, 14(1): 853–862.

Damstra, A.; Boomgaarden, H. G.; Broda, E.; Lindgren, E.; Strömbäck, J.; Tsfati, Y.; and Vliegenthart, R. 2021. What does fake look like? A review of the literature on intentional deception in the news and on social media. *Journalism Studies*, 22(14): 1947–1963.

Di Domenico, G.; Sit, J.; Ishizaka, A.; and Nunan, D. 2021. Fake news, social media and marketing: A systematic review. *Journal of Business Research*, 124: 329–341.

FORCE11. 2020. The FAIR Data principles. https://force11.org/info/the-fair-data-principles/.

Gebru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.

Gupta, A.; Li, H.; Farnoush, A.; and Jiang, W. 2022. Understanding patterns of COVID infodemic: A systematic and pragmatic approach to curb fake news. *Journal of business research*, 140: 670–683.

Hamby, A.; and Russell, C. 2022. How does ambivalence affect young consumers' response to risky products? *Journal of the Academy of Marketing Science*, 50(4): 841–863.

Hutto, C.; and Gilbert, E. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In *ICWSM*, volume 8, 216–225.

Ireton, C.; and Posetti, J. 2018. *Journalism, fake news & disinformation: handbook for journalism education and training*. Unesco Publishing.

Kaplan, K. J. 1972. On the ambivalence-indifference problem in attitude theory and measurement: A suggested modification of the semantic differential technique. *Psychological bulletin*, 77(5): 361.

Lazer, D. M.; Baum, M. A.; Benkler, Y.; Berinsky, A. J.; Greenhill, K. M.; Menczer, F.; Metzger, M. J.; Nyhan, B.; Pennycook, G.; Rothschild, D.; et al. 2018. The science of fake news. *Science*, 359(6380): 1094–1096.

Neiger, M.; and Tenenboim-Weinblatt, K. 2016. Understanding journalism through a nuanced deconstruction of temporal layers in news narratives. *Journal of communication*, 66(1): 139–160.

Nordgren, L. F.; Van Harreveld, F.; and Van Der Pligt, J. 2006. Ambivalence, discomfort, and motivated information processing. *Journal of experimental social psychology*, 42(2): 252–258.

Pennycook, G.; and Rand, D. G. 2021. The psychology of fake news. *Trends in cognitive sciences*, 25(5): 388–402.

Priester, J. R.; and Petty, R. E. 1996. The gradual threshold model of ambivalence: relating the positive and negative bases of attitudes to subjective ambivalence. *Journal of personality and social psychology*, 71(3): 431.

Rothman, N. B.; Pratt, M. G.; Rees, L.; and Vogus, T. J. 2017. Understanding the dual nature of ambivalence: Why and when ambivalence leads to good and bad outcomes. *Academy of Management Annals*, 11(1): 33–72.

Schneider, I. K.; and Schwarz, N. 2017. Mixed feelings: the case of ambivalence. *Current Opinion in Behavioral Sciences*, 15: 39–45. Mixed emotions.

Shu, K.; Mahudeswaran, D.; Wang, S.; Lee, D.; and Liu, H. 2020. Fakenewsnet: A data repository with news content, social context, and spatiotemporal information for studying fake news on social media. *Big data*, 8(3): 171–188.

Singh, V. K.; Ghosh, I.; and Sonagara, D. 2021. Detecting fake news stories via multimodal analysis. *Journal of the Association for Information Science and Technology*, 72(1): 3–17.

Spezzano, F.; Shrestha, A.; Fails, J. A.; and Stone, B. W. 2021. That's Fake News! Reliability of News When Provided Title, Image, Source Bias & Full Article. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1): 1–19.

Tausczik, Y. R.; and Pennebaker, J. W. 2010. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of language and social psychology*, 29(1): 24–54.

Van Krieken, K.; and Sanders, J. 2021. What is narrative journalism? A systematic review and an empirical agenda. *Journalism*, 22(6): 1393–1412.

Vosoughi, S.; Roy, D.; and Aral, S. 2018. The spread of true and false news online. *science*, 359(6380): 1146–1151.

Zhang, X.; Cao, J.; Li, X.; Sheng, Q.; Zhong, L.; and Shu, K. 2021. Mining dual emotion for fake news detection. In *The Web Conference*, 3465–3476.

Zhang, X.; and Ghorbani, A. A. 2020. An overview of online fake news: Characterization, detection, and discussion. *Information Processing & Management*, 57(2): 102025.

Zhao, W. X.; Jiang, J.; Weng, J.; He, J.; Lim, E.-P.; Yan, H.; and Li, X. 2011. Comparing twitter and traditional media using topic models. In *ECIR*, 338–349. Springer.

Zhou, X.; Mulay, A.; Ferrara, E.; and Zafarani, R. 2020. Recovery: A multimodal repository for covid-19 news credibility research. In *CIKM*, 3205–3212.

Zhou, X.; and Zafarani, R. 2020. A Survey of Fake News: Fundamental Theories, Detection Methods, and Opportunities. *ACM Comput. Surv.*, 53(5).

Paper Checklist

1. For most authors...
   (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? *Yes, the research advances science without violating any social norms.*
   (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? *Yes, abstract and introduction accurately reflect the paper's contributions and scope.*
   (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? *Yes, we applied classical machine learning techniques (classification, important feature extraction, and explanation) to answer our research questions and clarify that we used LIWC features in input to our classification models as this is a well-established method in content-based fake news related research (Zhou and Zafarani 2020; Spezzano et al. 2021).*
   (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? *NA, as we didn't collect any data. However, the analyzed datasets are widely used and studied by the research community.*
   (e) Did you describe the limitations of your work? *Yes, we have discussed the limitations of our work at the end of the "Discussion and Limitations" section.*
   (f) Did you discuss any potential negative societal impacts of your work? *No, we did not discuss any potential negative societal impacts of our work; if anything, the main application of our work is for societal benefit.*
   (g) Did you discuss any potential misuse of your work? *No, we did not discuss any potential misuse of our work; as mentioned above, the main application of our work is for societal benefit, and we do not foresee misuse of our work for ill intent.*
   (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? *Yes, data and models are described in the "Experimental Framework" section. To enhance reproducibility, our code is available at* https://anonymous.4open.science/r/Ambivalence-4104
   (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? *Yes, we have read the review guidelines and our paper conforms to them.*
2. Additionally, if your study involves hypotheses testing...
   (a) Did you clearly state the assumptions underlying all theoretical results? *NA, we do not have any hypotheses testing in our work.*
   (b) Have you provided justifications for all theoretical results? *NA*
   (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? *NA*
   (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? *NA*
   (e) Did you address potential biases or limitations in your theoretical framework? *NA*
   (f) Have you related your theoretical results to the existing literature in social science? *NA*
   (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? *NA*
3. Additionally, if you are including theoretical proofs...
   (a) Did you state the full set of assumptions of all theoretical results? *NA*
   (b) Did you include complete proofs of all theoretical results? *NA*
4. Additionally, if you ran machine learning experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? *Yes, to enhance reproducibility, our code is available at* https://anonymous.4open.science/r/Ambivalence-4104 *The datasets are publicly available.*
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? *Yes, as we mention in the "Experimental Framework" section, we have done 5-fold stratified cross-validation for all our considered datasets. Hyperparameters for the Random Forest model were (number of estimators = 200, random state = 100). All the hyperparameters were chosen with experiments.*
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? *No, we performed 5-fold cross-validation and averaged the results to mitigate the errors. So the variation in results was negligible and not reported.*
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? *Yes, as we mention in the "Experimental Framework" section, we ran our experiments on Google Colab with CPU. Reproducing all the experiments reported in the paper takes approximately 5 hours.*
   (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? *Yes, see our response to Question 1 (c).*
   (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? *No. Even if we use classification in our paper, our goal is not to develop a new fake news detector or a model predicting the number of news shares. Rather, we study the relationship between ambivalence and fake news or news sharing.*

5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
   (a) If your work uses existing assets, did you cite the creators? *Yes, we have properly cited the authors of the datasets we have used for our work.*
   (b) Did you mention the license of the assets? *NA. If the used datasets have any licenses, this is reported in their repository mentioned in the cited papers.*
   (c) Did you include any new assets in the supplemental material or as a URL? *Yes, to enhance reproducibility, our code is available at* https://anonymous.4open.science/r/Ambivalence-4104
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? *NA*
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? *NA*
   (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? *NA*
   (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? *NA*
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
   (a) Did you include the full text of instructions given to participants and screenshots? *NA*
   (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? *NA*
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? *NA*
   (d) Did you discuss how data is stored, shared, and deidentified? *NA*

Appendix

Table 5 lists the new ambivalence measures we defined in this paper, along with the LIWC dimensions used to compute them.

Table 5: New types ambivalence and associated LIWC dimensions.

Ambivalence Type	LIWC Dimension	Associated Words
Past vs. Present	Past	was, had, were, been
	Present	is, are, I'm, can
Past vs. Future	Past	was, had, were, been
	Future	will, going, have to, many
Present vs. Future	Present	is, are, I'm, can
	Future	will, going, have to, many
Emotional vs. Cognitive	Tone	good, well, new, love (pos); bad, wrong, too much (neg)
	Cognitive	but, not, if or know
Sad vs. Happy	Sad	sad, disappointed, cry
	Happy	good, well, happy, love
Work vs. Leisure	Work	work, school, working, class
	Leisure	fun, play, party
Female vs. Male	Female	she, her, girl, woman
	Male	he, his him,man

In the following *three* pages, we have included all the SHAP plots to describe the relationship between the impact of ambivalence features on fake news detection and the prediction of news shares. *Please note that* ***Textual ambivalence*** *and* ***Emotional-cognitive ambivalence*** *are referred to in the plots as* ***Ambivalence*** *and* ***toneCognitiveAmbiv****, respectively, for simplicity.* For all plots:

- the values on the **X-axis** represent the impact of a feature on the prediction. For continuous features, it typically shows the magnitude of the SHAP value, which is how much the feature contributes to the prediction.
- the **Y-axis** lists the ambivalence features, ordered by their importance or the impact they have on the model's output.
- the color indicates the feature's actual value (e.g., red for high values and blue for low values). This helps to understand how the feature's value influences the model's prediction.
- Each point represents a single prediction, and the position along the x-axis shows the SHAP value for that feature. A positive SHAP value indicates the feature pushed the model's prediction higher (e.g., fake news for fake news prediction as this label is coded with 1), while a negative SHAP value means it pushed the prediction lower (e.g., real news for fake news prediction as this label is coded with 0).

For instance, looking at past-future ambivalence in Figure 1, we see that higher values of this feature predict real news in the Politifact dataset, while, in the same plot, higher values of emotional-cognitive ambivalence predict fake news.

Relationship between Ambivalence and Fake News Prediction

Please note that no ambivalence feature results among the top-30 most important features for predicting fake news in the ReCOVery dataset, hence there is no corresponding plot for this dataset.

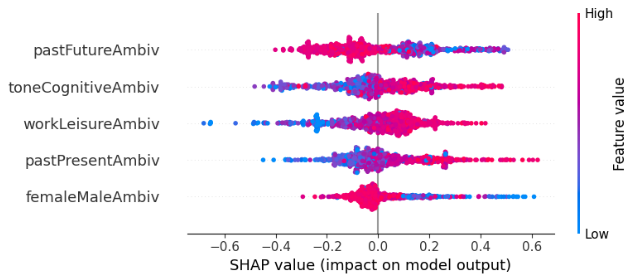

Figure 1: Politifact - Impact on News Prediction

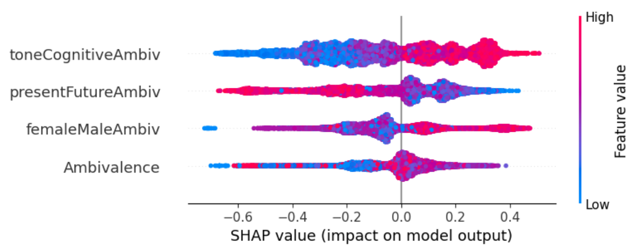

Figure 2: GossipCop - Impact on News Prediction

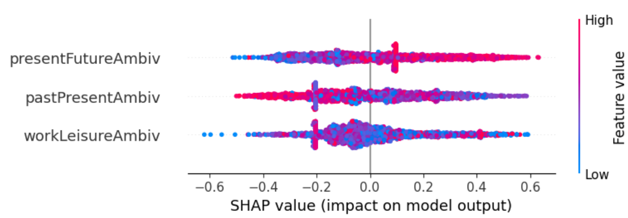

Figure 3: Fake Health - Impact on News Prediction

Relationship between Ambivalence and Fake News Share Prediction

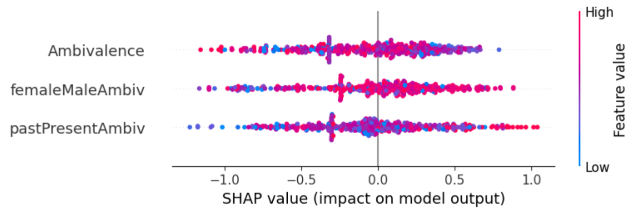

Figure 4: Politifact - Impact on # of Share (Fake)

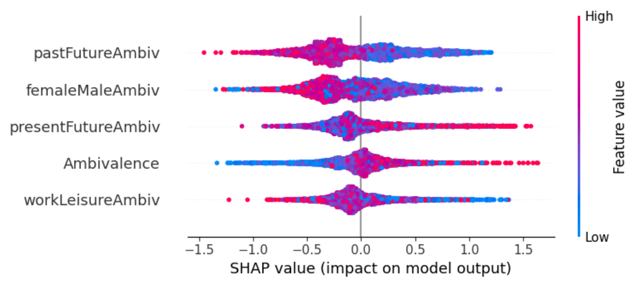

Figure 5: GossipCop - Impact on # of Share (Fake)

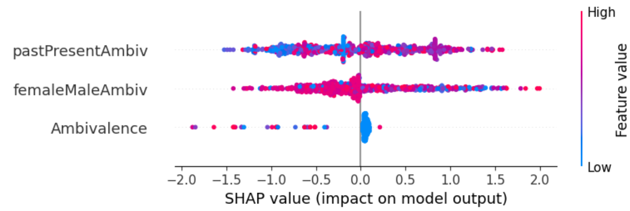

Figure 6: ReCOVery - Impact on # of Share (Fake)

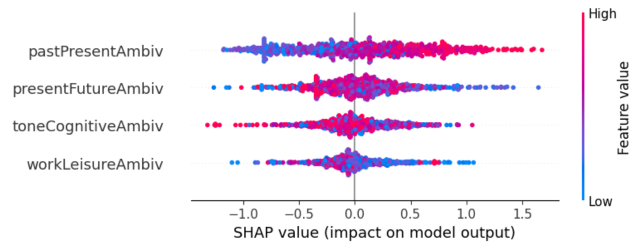

Figure 7: Fake Health - Impact on # of Share (Fake)

Relationship between Ambivalence and Real News Share Prediction

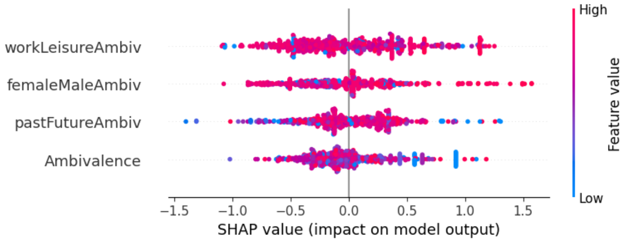

Figure 8: Politifact - Impact on # of Share (Real)

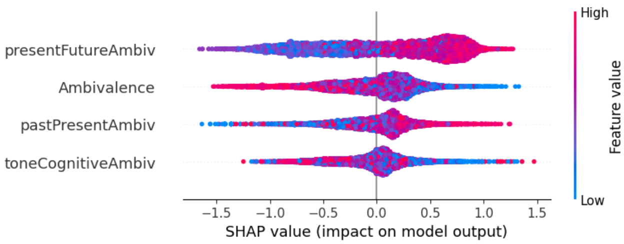

Figure 9: GossipCop - Impact on # of Share (Real)

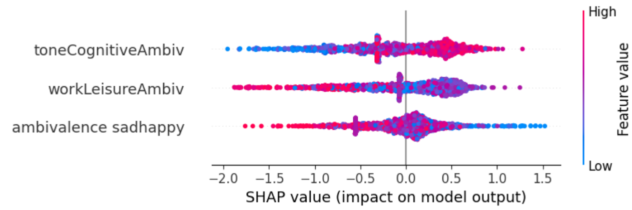

Figure 10: ReCOVery - Impact on # of Share (Real)

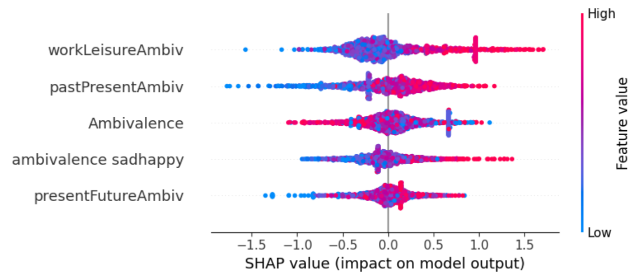

Figure 11: Fake Health - Impact on # of Share (Real)