Analysis and Detection of “Pink Slime” Websites in Social Media Posts

Abdullah Aljebreen
Temple University
Philadelphia, USA
aaljebreen@temple.edu

Weiyi Meng
Binghamton University
Binghamton, NY, USA
meng@binghamton.edu

Eduard C. Dragut
Temple University
Philadelphia, USA
edragut@temple.edu

ABSTRACT

Local news outlets play a vital role in providing trusted and relevant information to communities and addressing their specific needs and concerns. The emergence of news outlets posing as local sources and their spread on social media present a significant challenge in the digital information landscape. This paper presents a comprehensive study investigating posts featuring “pink slime” news, which is a term that has been used to refer to these news outlets due to its deceptive nature. By analyzing a large dataset of posts, we gain valuable insights into the patterns of these posts and the origin of these posts. We show in this work that extracting syntactical features proves valuable in developing a classification approach for detecting such posts and that the approach achieves 92.5% accuracy. We also show that our approach achieves near-perfect detection when grouping the posts by URL.

CCS CONCEPTS

• **Information systems → Social networks**; **Social networking sites**; *Clustering and classification.*

KEYWORDS

Social Media, News, Tweets, Misinformation, Information Integrity, Classification

ACM Reference Format:
Abdullah Aljebreen, Weiyi Meng, and Eduard C. Dragut. 2024. Analysis and Detection of “Pink Slime” Websites in Social Media Posts. In *Proceedings of the ACM Web Conference 2024 (WWW ’24), May 13–17, 2024, Singapore, Singapore.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3589334.3645588

1 INTRODUCTION

During the 2016 United States presidential election and the surge of the term *fake news*, Americans exhibited higher levels of trust in local news outlets compared to national ones [32]. The preference for local news demonstrated a belief that it offered a more trustworthy alternative to the perceived bias and sensationalism of national news [39]. This trust became a vulnerability that internal and external actors attempted to take advantage of to disseminate misinformation and shape public opinion. A report by the U.S. Senate Intelligence Committee in 2018 [33] found more than half a million posts by external-operated Twitter accounts impersonating local news outlets during the 2016 election [32]. Exploiting public trust in local news began with the appearance of many news sites that posed as local news which have been labeled as “pink slime” news [6].

“Pink slime,” officially known as “lean finely textured beef,” is an informal term used to describe a low-cost, processed beef product that is typically added to ground beef as a filler to reduce the overall fat [37]. However, the term has also been coined by journalists to refer to news outlets that appear to be local news, but in reality, they are not [55]. While pink slime is not a formally recognized term, it serves as a practical depiction of these outlets in our research. Our choice stems from the absence of an established, rigorously defined concept to name this particular phenomenon.

Although there have been sightings of these pink slime news outlets since 2012, they have started to be more noticeable in the year ahead of the 2020 United States presidential election [7]. These media platforms adopt the names of cities and towns across every U.S. state, with almost no local reporters or physical newsrooms [47]. While those outlets claim to promote local journalism, their operations divert readership away from traditional local newspapers, raising concerns about their true intentions [45]. Similar to conventional news sources, pink slime publications utilize social media for promotion. They often employ attention-grabbing headlines, commonly referred to as clickbait, to expand their presence on those platforms [4, 29]. Furthermore, some pink slime outlets resort to automated bots and fake accounts to artificially boost their engagement statistics, making it more difficult for users to discern the authenticity of their content [11].

We present a novel study on the presence of pink slime URLs in social media. Our study has four directions. First, we collect a large number of posts from X, formally known as Twitter, over 300K pink slime posts and 500K for national and local news, and curate them, e.g., we discard those with broken or expired URLs. Second, we study their textual organization by comparing the text of a post with that of the document (or news article) it references. This allows us to unearth interesting patterns of creation and dissemination of pink slime posts as compared to the rest of the news. For example, many of the pink slime posts copy the first sentence of the article it references. In addition, in most of those posts, the sentences are cut short, at arbitrary positions. Since length is not an issue such a finding suggests that those posts are generated by bots. Third, we aim to understand the factors– ranging from article post content to observed dynamics of user interaction with such posts [18, 19], like the number of likes– on detecting pink

WWW ’24, May 13–17, 2024, Singapore, Singapore

ACM ISBN 979-8-4007-0171-9/24/05
https://doi.org/10.1145/3589334.3645588

slime posts and explain which of those features contribute more toward detection. For instance, we observe that features related to the length or the time of the posts are more useful than the interaction measures, such as the replies or the reposts. We show that we achieve 92.55% accuracy in detecting pink slime posts. Finally, in all the studies mentioned above, we contrast pink slime posts with posts that reference national and local news websites. Overall, this study advances our understanding of pink slime news in the digital age and highlights the need for proactive interventions. By employing computational and analysis approaches, our paper makes the following contributions.

- We present a comprehensive analysis of posts sharing URLs from pink slime news websites, bridging the gap between journalism and data science.
- We utilize our analysis and insights into the sharing patterns within these posts and develop a set of features tailored for posts associated with pink slime news.
- We develop a classification approach that detects posts with URLs from pink slime among posts with news from other sources. Our approach achieves 92.55% accuracy without relying on the textual content of the news article.

The paper is organized as follows: In Section 2, we discuss related work and news reports about pink slime journalism. We define our problem in Section 3. In Section 4, we present the process of building and processing our dataset. We analyze posts from our dataset and apply processing techniques from previous work in Section 5. We present our detection model and its results in Section 6. We discuss future directions and our conclusion in Sections 7 and 8.

2 BACKGROUND

Pink slime journalism is a relatively new phenomenon. Hence, it remains an understudied topic in academia compared to similar concepts, such as satire, fake news, misinformation, and media bias [9, 18, 25, 27, 30, 31, 43, 51, 53, 54]. During our investigation of this subject, we discovered that there is minimal research available in the literature that defines and examines the effects of pink slime news. For example, some LLMs, like ChatGPT3, acknowledge the shortage of sufficient academic research (Figure 1) [34].

Numerous news articles and investigative reports have been published examining the notion of pink slime journalism. These attempts aim to uncover the origin of these outlets and trace their dissemination [6, 28, 32, 55]. Alongside these news reports, a few research papers investigate the topic of pink slime websites, examining two primary aspects: the website's attributes and its published content, as well as the consumption of their content [24, 29]. Furthermore, some studies investigate the characteristics of networks comprising multiple outlets rather than individual ones [38].

Bengani studies the history of the uprising of pink slime outlets and their networks [6, 7]. Several investigative reports talk about how pink slime outlets are more active during the times of elections [4, 7]. An article from The New York Times alleges that certain media outlets are being directed by political groups and corporate P.R. firms to continuously publish positive stories about a specific candidate or negative ones about their opponents [1]. Some local news outlets in swing states, such as Michigan or Pennsylvania, have reported the spread of pink slime outlets in their states [14, 41],

Figure 1: A screenshot showing the response of ChatGPT when asked about academic research on pink slime.

Pink slime news outlets aim to compete with existing local news outlets in the same localities. Therefore, an important research problem is studying how they behave compared with the corresponding local establishments. The 2021 paper, by Karell and Agrawal, studied the content of 122,054 news articles from pink slime websites and compared them with 90,689 news articles from the corresponding local news outlets. Their findings showed that the pink slime websites have three distinct behaviors: publishing relevant online text and data, such as gasoline prices, public reports, or press releases; capitalizing on national political controversies; and utilizing emotional appeals to attract moderate readers and encourage further engagement [24]. A more recent work, by Moore et al., claims to present the first empirical analysis of individuals' engagement with pink slime journalism. While the study reveals that a relatively small number of Americans visit pink slime websites, the findings emphasize the significance of further research on this content type [29]. Their results indicate that over 9 million Americans accessed a pink slime website during the 2020 election.

The research by Royal and Napoli studies one of the largest networks of pink slime outlets as a potential future modern model of local news reporting [38]. The finding of this study suggests that the network fails to adequately address the informational needs of the communities it claims to represent. These findings raise doubts about the effectiveness of its approach, which relies on generating automated content with minimal human-written stories to sustain a vast network of outlets nationwide. While local newspapers are facing financial challenges, this study highlights that automated, large-scale national operations are an inadequate substitute for the resource-intensive efforts of traditional local news. However, the model followed by such pink slime outlets may provide insights for those seeking to promote genuine local news in the future [38].

Overall, we observe a dearth of systematic, quantitative research on the phenomenon of pink slime news, particularly with empirical evidence. Hence, we believe that there are many opportunities for data-driven researchers on pink slime journalism, its consumption, and circulation on social media platforms. Our work is motivated mainly by the lack of similar work on this problem from a computational perspective. We seek to uncover unique features and behaviors associated with posts that share pink slime URLs and compare those to posts that share URLs of news articles from broadly recognized local and national news outlets. Understanding the distinctive attributes of those posts is an important first step since it may enable social media platforms to create a more responsible environment by flagging posts that do not seek to inform, but rather seek to promote propaganda and persuasion [13, 23, 43, 48, 50, 51].

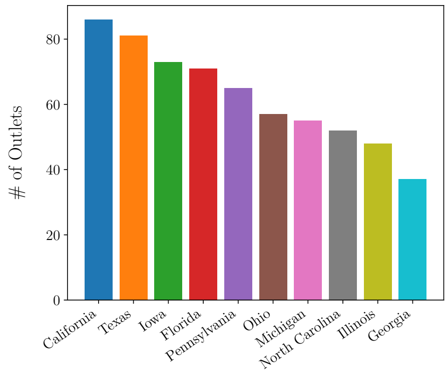

Figure 2: The distribution of pink slime news for the top 10 states.

Table 1: The distribution of pink slime outlets with respect to networks. The column *coverage* shows the number of outlets per type of local coverage (city-county-region-state)

Network	outlets	states	coverage
Star News	11	10	(4-0-0-7)
The Record	10	8	(8-0-1-1)
LGIS	36	1	(33-3-0-0)
Metric Media	989	50	(530-141-271-47)
Local Report	48	10	(20-9-19-0)
Franklin Archer	193	51	(187-2-1-3)
Locality Labs	16	2	(16-0-0-0)
American Independent	5	5	(0-0-0-5)
American Catholic	6	6	(0-0-0-6)

3 PROBLEM

In this work, we focus on a particular study of social media posts that share links from pink slime websites, which we call *pink slime posts* throughout the paper. We identify two tasks: **First**, we seek to identify the distinguishing characteristics of posts containing URLs from pink slime news websites with a comparison to posts with national and local news URLs. **Second**, by leveraging the findings and insights gained from our analysis we aim to show that we can find sufficient features for an efficient detection approach for these posts. While one can create lists of pink slime websites and use them to detect such posts, we aim toward a more comprehensive solution, which is able to tag such posts for new, not previously seen pink slime websites. Such technology can prove valuable as it can track the emergence and vanishing of pink slime news outlets, a pattern often influenced by significant political and societal occurrences over time, like election and vaccination campaigns.

4 DATASET

In our paper, we experiment with posts that contain news URLs from pink slime, local and national news outlets. Therefore, we built a dataset with posts in each category. We describe the process of collecting and curating the data in this section. The process has two steps. First, we construct a list of news outlets in each category. Second, we use the Twitter API to collect posts (tweets) that share URLs from each outlet in our list. We give the details below.

4.1 Collecting News Outlets

4.1.1 pink slime Outlets. For the pink slime category, we compiled a comprehensive list of pink slime outlets. The list is built using data from several sources, including the Tow Center and The New York Times [2, 7, 47]. The list includes a total of 1,313 outlets that are spread across all 50 states and DC. Each of these outlets is associated with a particular network with a total of 8 networks. Table 1 shows more details about these networks. We also give the distribution for the top 10 states with the largest presence of pink slime outlets in Figure 2.

4.1.2 Local and National Outlets. Regarding national category outlets, we aim to make sure that the news outlets in our list represent different leanings in the political spectrum. Therefore, we use the Media Bias Chart from AllSides.com [1] and select 25 outlets from left, left-leaning, centrist, right-leaning, and right stances. In the local news category, we first identify the localities where the pink slime outlet operates, then we manually search in Google News for outlets using those locality keywords. We collect a list of 50 local news outlets, such as azcentral, Chicago Tribune, and PennLive.

4.2 Collecting Posts with News URLs

4.2.1 Pink Slime Posts. After constructing the comprehensive list of outlets, the next step is gathering posts that share URLs for news articles from each website in our list. To achieve this, we use the Twitter API to query posts containing the prefix URL of the news website (e.g., www.newsoutlet.com). We collect posts between January 2016 to May 2023 [26]. We download the text content of the news article from the URL included in each post. We collect 348k posts with pink slime URLs, but after cleaning only contain accessible URLs. Figure 3 shows the distribution of these posts per network, top 10 outlets, and top 10 states.

Upon investigating the temporal distribution of the collected posts, we noticed a surge in the volume of pink slime posts between May 2022 and September 2022. This temporal trend is shown in Figure 4, which illustrates the overall distribution and the distribution of posts in the originating outlets, networks, and states. This surge is attributed primarily to three outlets: Arizona Business Daily, Ohio Business Daily, and Chicago City Wire. The majority of the posts from those outlets, precisely 95.82%, 99.90%, and 82.68%, were exclusively concentrated in this five-month window. Moreover, our investigation revealed that 87.17%, 99.87%, and 80.41% of these posts originated from an account that belonged to the respective outlets.

4.2.2 National and Local Posts. We collect posts from both national and local news outlets, employing the same time frame utilized in our pink slime posts collection. Specifically, we build a dataset comprising 500k posts from national news and an equivalent number from local news outlets. We also download the URLs included in those posts. This dataset is used in our subsequent experiments and in-depth analyses.

[1] https://www.allsides.com/media-bias/media-bias-chart

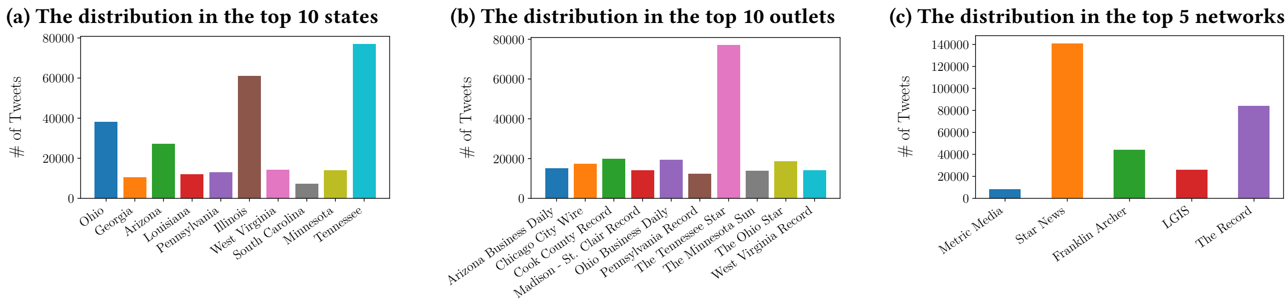

Figure 3: Distribution of pink slime posts in our datasets per states, outlets, and networks.

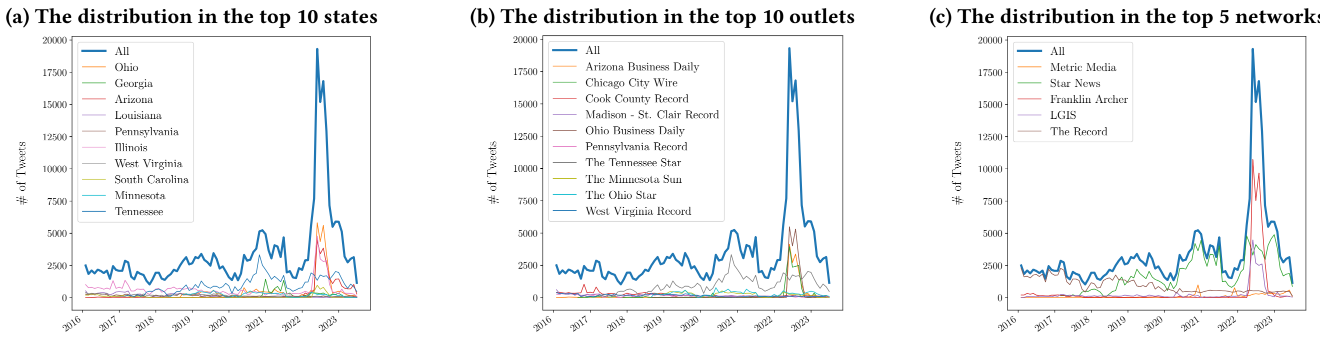

Figure 4: Temporal distribution of pink slime posts in our datasets per states, outlets, and networks.

5 ANALYZING WRITING PATTERNS

Following our data collection, we manually examined a sample of 1000 pink slime posts from our dataset to explore potential directions for analysis. Throughout our inspection, we found instances where the posts were posted by accounts that belonged to news outlets, others that appeared to be posted by bot accounts, and others by ordinary users. We believe that this observation applies to all pink slime posts. However, what interested us more was the common behavior of these accounts, specifically their composing and editing practices in pink slime posts. For example, we noticed that pink slime posts exhibit less discussion and contextual information compared to other news posts. For example, we came across many instances where posts featured incomplete sentences that were copied from the shared news article and subsequently truncated (Figure 5). In this example, the title of the shared news article is copied and cut short. This action was not deemed necessary due to the length restrictions of posts. In other examples, the truncated text is a quote from within the body of the article. Following these findings, we sought to explore the relationship between the text of the post and the text of the pink slime news article more systematically.

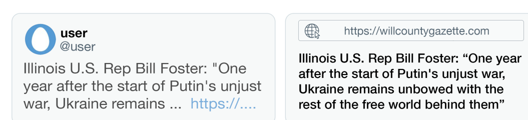

Figure 5: An example of an excerpt is copied from the news article and abruptly truncated.

5.1 Relationships Between Posts and News Articles

There are many works in the literature that analyze and study posts text, but we found few that focus on posts with news URLs or URLs, in general, [10, 16, 44, 49]. One work analyzes the changes to the news article headlines after being shared on Twitter and shows that these edits differ from one news outlet to another [17]. Another work studies general hyperlinked posts indicates that not all posts containing URLs are composed in the same way [3]. Their findings show that some posts simply copy the title of the URL, while others include quotes from the news article. Some may contain additional commentary or text written by the user, while others may include multiple elements such as hashtags, images, or mentions of other users [3]. These findings apply to all posts with URLs, including news URLs and more importantly pink slime URLs. The nature of such posts can have an impact on how convincing they appear to other users and how likely they are to be shared or spread further. By analyzing the text of the post, such as whether it contains only the title of the article or additional commentary, researchers can gain insights into the motivations, bias, or stance of the users who share these links [21, 42]. For instance, a potential study could explore whether posts featuring the news article's title are more likely to be posted by automated bots, whereas posts including excerpts from the article or user-generated content have a higher likelihood of human sharing. Next, we show the results and implications of applying the method of segmentation [3] on the posts in our dataset.

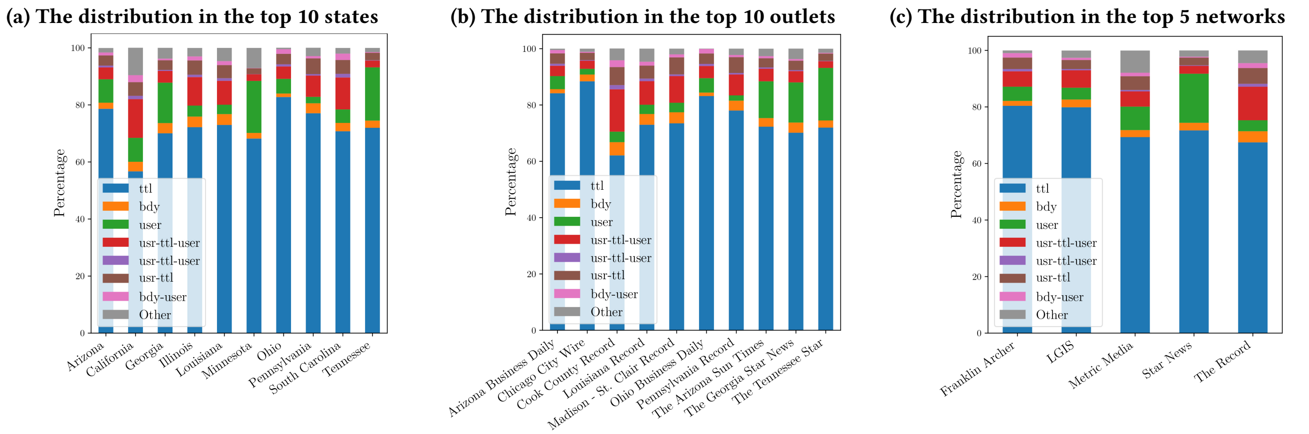

Figure 6: The observed segmentation patterns for pink slime posts. We organize them by state, outlet, and network.

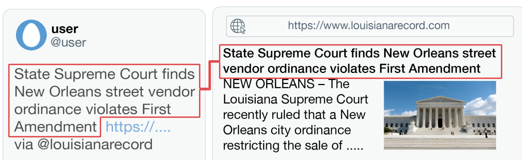

Figure 7: An example post where the title is followed by mention reference to the source of the URL.

5.2 Segmentation of the post text

To investigate the *editing behaviors* performed on posts with pink slime news URLs, we use the segmentation method for posts with URLs, in which we inspect the similarity between the text of the post and the title and text content of the shared URL [3]. To achieve that, we follow the method proposed by Aljebreen et al., which is summarized as follows: (1) we check if the title of the news article is copied into the post; if yes then we mark it as a title segment or *ttl*. (2) For the remaining, unsegmented, text of the post we check if it contains quotes copied from the main body of the article. If yes, we mark them as body segments, or *bdy*. (3) The remaining parts of the post are assumed to be added by the user and, hence, are marked as user segments, or *usr*. We followed the algorithm described in [3], which mostly follows known string matching strategies. We complement it with several heuristics to take into account minor edits to the title and body when copied to the post, such as adding hashtags or user handles or shortening the copied text from the news article.

5.2.1 Setup. To study the distinctive characteristics of pink slime posts compared to other news posts, we conducted segmentation across all three datasets: pink slime, national, and local news posts. We use all 305k collected pink slime posts. As for national and local news posts, we sampled the same number of posts from each dataset (Section 4.2.2) while adhering to the same time interval and post frequency as the pink slime dataset.

5.2.2 Results. In the segmentation process, each processed post is assigned a pattern based on its segments. For instance, a post with the pattern: (*ttl*) consists only of a title segment, while another with

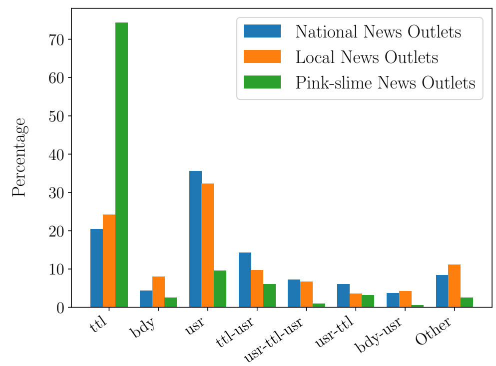

Figure 8: A comparison of each dataset's most common segmentation patterns.

(*bdyusr*) pattern has a quote from the body of the article followed by user-added text. The results reported by [3] indicate that the most common 3 patterns are *ttl*, **usr**, and *bdy*. In our experiments, the observed distribution is different from theirs and the patterns differs in each dataset. In Figure 8, we show the most common 7 segmentation patterns. The results reveal distinct patterns in pink slime posts as compared to national and local news posts. Pink slime posts predominantly exhibit *ttl* pattern, accounting for over 74.36%, whereas this pattern is observed in only 20.49% and 24.2% of national and local news posts, respectively. On the other hand, the most prevalent pattern in national and local news posts is *usr*, present in 35.64% and 32.29% respectively, whereas pink slime posts exhibit this pattern in only about 9.66% of cases.

The overall results give a clear indication of the posting behavior of pink slime posts as a population: they include little user text. Moreover, the segmentation distribution differs between some outlets. In Figure 6, we show the segmentation distribution of pink slime posts among different networks, outlets, and states. We will show later, in Section 6, that these patterns identified in the posts with news URLs, across the three categories, can also serve as useful features for a machine-learning classification approach designed to detect such content. By analyzing a large dataset of posts, an ML model can learn to identify common patterns that are associated

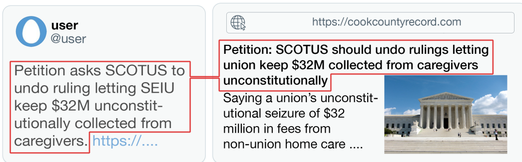

Figure 9: An example of a pink slime post where a paraphrased version of the URL's title appears.

with pink slime news URLs, and the dominant features are related to the presence or lack of user-generated text.

5.3 Special Cases of Segmentation

Once we have reached the results presented in this section, which reveal the predominant editing patterns of the posts in each category, our focus shifts towards identifying more specific markers within the range of these patterns. This is particularly crucial in the context of our primary subject of investigation, which is the pink slime posts. In the following sections, we delve into three noteworthy instances where we can discern additional features associated with these patterns.

5.3.1 Outlet reference. One of the most common patterns is when the title appears in the post followed by a segment that is assumed to be added by the user. We notice that the nature of *usr* segments differs in length and textual content. In many instances, the posts refer to the source of the shared news article by mentioning its user handle in Twitter, e.g. *via @BreakingNews*. An example of this posting behavior is shown in Figure 7. We further search the datasets for these patterns and we find that it is more common in the national and local posts than in the pink slime posts. Around 49% of the national news posts with *ttl-usr* pattern contain this case, while only 21% from the local posts and 14% of the pink slime posts have them.

5.3.2 First sentence quote. Another observed case is the location of the quoted text, the *bdy* segments, in the news article itself. This was brought to our attention by noticing the presence of many posts with incomplete quotes in pink slime posts, which appear to be copied from the beginning of the news article. We show two examples, in Figure 10, from both local (top posts) and pink slime (bottom posts) datasets. We study posts with *bdy* pattern in all datasets and attempt to determine the overall prevalence of this copying pattern. Our analysis finds this pattern in 57% of pink slime news posts with *bdy*, 27% of local posts, and less than 1% of national news posts.

5.3.3 Paraphrased title. We observed that certain posts with *usr* segments, often presented a paraphrased version of the article's title instead of an addition or discussion by the user. For example, the post shown in Figure 9 is marked to have *usr* segments. However, the title of the news article appears to be paraphrased and copied into the post. To confirm and test these observations on a larger level, we employed the GPT3 language model [34] to detect the paraphrasing of the title. For each post, we passed both the *usr* segment and the title and asked GPT3 to determine if the two pieces of text are paraphrased of each other. However, upon manually reviewing a sample of 100 posts that we processed with GPT3, we

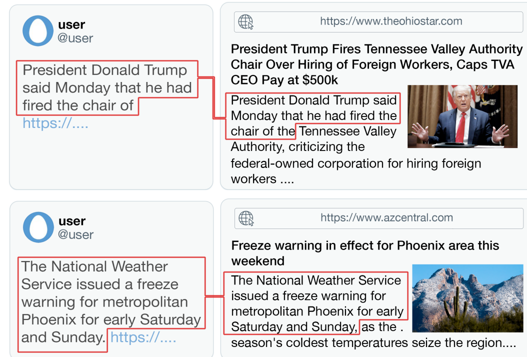

Figure 10: Two examples of posts, the top one is a pink slime post and the bottom is a local news post, where the first sentence of the article is copied into the post.

discovered numerous instances of inaccuracies and false positives. Among this sample, we found 21 posts to be false negative and 44 were false positive. These findings highlighted the limitations of relying solely on automated methods in such problems.

6 DETECTING PINK SLIME POSTS

The findings of our analysis in the previous section provide us with an opportunity to look for a set of features that can be used to develop a classification algorithm for pink slime posts. In addition, our goal is to identify the feature subset that is the most important for detecting posts with pink slime URLs.

6.1 Methodology

The use of the Random Forest machine learning model presents a compelling approach for classifying posts containing pink slime news. Random Forest is a powerful ensemble learning technique that combines multiple decision trees to make accurate predictions [12]. Its ability to handle high-dimensional data and capture complex interactions among features makes it well-suited for the task of classifying posts based on the presence of pink slime news [8, 35]. By leveraging the ensemble nature of Random Forest, we can effectively harness the collective wisdom of multiple decision trees to improve the overall classification performance. Additionally, the model's interpretability, compared to deep neural networks [5], allows us to gain insights into the key features driving the classification, enabling us to understand the distinguishing characteristics of posts with pink slime news URLs [46].

6.2 Features

In our experiments, we try different combinations of features and attempt to find the most useful ones for detecting pink slime posts among other news posts. In addition to the segmentation signals, we discussed in the previous section, we extracted additional features from the posts. We organize the features into six groups: Segmentation, Other patterns, Special elements, Interaction, Structural, Temporal, and Misc features. Table 2 summarizes these feature sets.

Segmentation features In the segmentation features, we use the patterns of the results when we process the posts using the segmentation patterns, such as *ttl-usr, bdy, usr* ... etc. We converted

Table 2: The 6 groups of features utilized in the pink slime posts detection experiments.

Feature	Description
A) Segmentation features	
post_pattern	(*T, B, BU* ...) the segmentation pattern
pattern_length	(0-...) number of segments in the pattern
ttl_presence	(0,1) whether *Title* segment is present
bdy_presence	(0,1) whether *Body* segment is present
usr_presence	(0,1) whether *User* segment is present
B) Other patterns features	
outlet_reference	(0,1) whether it has a reference to the outlet after the title, e.g. via @News
first_sentence_quote	(0,1) whether it has a direct quote the first sentence in the article
URL_location	(0,1,2) location of the URL in the post
C) Special elements features	
hashtags_count	(0-...) the number of hashtags
hashtags_length	(0-...) the length of hashtags in characters
mentions_count	(0-...) the number of user mentions
emojis_count	(0-...) the number of emojis
D) Interaction features	
repost_count	(0-...) the number of reposts
reply_count	(0-...) the number of replies to the post
quote_count	(0-...) the number of times the post was quoted
like_count	(0-...) the number of likes
bookmark_count	(0-...) the number of times the post was was bookmarked
E) Structural features	
post_length	(0-...) the length of the post
sentence_count	(0-...) the number of sentences
punctuation_count	(0-...) the number of punctuation in the post, e.g. *!?,.* ...etc
F) Temporal features	
hour_of_day	(0-23) the hour of the day the post was posted
day_of_week	(0-6) the day of the week the post was posted
G) Misc. features	
annotation_count	(0-...) the number of annotations
geo_info	(0-1) whether the geo info is available in the post

this nominal feature into multiple binary ones. For example, 1 if *ttl-usr* appears and 0 otherwise.

Other patterns features Here we use the observations we discussed in Section 5.3 and extract two features related to them and the post patterns. In addition, we add a within-posts location feature of a URL. A URL may appear at the beginning or end of a post, or somewhere in the middle of it.

Special elements features Hashtags, mentions, and emojis play significant roles in posts. Hence, we transform their presence into distinct features for analysis, such as the number of hashtags and their length.

Interaction features One of the important parameters that we collected during the gathering of the posts was *public metrics* which is a count of the engagement interactions that the post receives in real-time, which are Replies, Likes, Peposts, and Bookmarks. We use these interaction metrics as features for our posts.

Structural features Other textual properties were extracted during the processing of the posts. These properties include the length of the post, the number of sentences, and the number of characters in the post that are punctuations.

Temporal features Includes two features that are related to the posting time, we focus on the time of the day the posts were posted. Also, we use the day of the week as another temporal feature.

Misc features Lastly, we incorporate two additional metrics that we collected from the Twitter API: *annotation_count* is the number of mentioned entities recognized in the post and geo_info is the geographical location of the post’s origin.

We do not include www.newsoutlet.com as features. The reason is that we want the model to learn features that are independent of the name of an outlet. Such a model will be more robust to identifying unseen pink slime posts.

6.3 Experimental Setup

In our experimental study, we utilize a cross-validation approach to assess the RF model performance. The posts in our datasets are randomly divided into five folds based on the outlets they belong to. Four of these folds constitute the training set, while the fifth fold is reserved for testing. We create a model based on the training set’s features and then apply this model to the unseen testing set. This process is repeated five times, each time designating a different fold for testing. Our reported performance metrics are based on the average results obtained through this approach. Notably, this evaluation includes the application of a zero-shot detection technique, where no pink slime outlet is used during both the training and testing phases.

6.3.1 Evaluation metrics. In evaluating our technique, we rely on essential metrics like *F1* score and *accuracy*. The *F1* score offers a balanced assessment of precision and recall, crucial for evaluating the model’s ability to correctly identify pink slime posts while minimizing false positives. On the other hand, the *accuracy* measure provides an overall measure of correctness in post classification and shows the proportion of all correctly detected posts.

6.3.2 Hyperparameters. Tuning hyperparameters plays a critical role in enhancing the performance of Random Forest models. Through *random search cross-validation* which performs small-scale k-fold experiments we randomly experiment with different combinations of hyperparameters to optimize the model’s performance. The parameters used in our model and our choice are as follows: the number of decision trees in the forest ($num_t rees \in [300, 500, 700]$) and the maximum depth of these trees ($max_d epth \in [10, 20, 30]$).

6.4 Results

We trained our RF model to detect pink slime posts in three distinct settings, compared to national news, local news, and all news. Our model exhibits an accuracy of 0.926 and an F1 score of 0.8 when trained with all news. The model performs similarly when analyzing

Table 3: The results for identifying dominant features when only using local, national, and both settings with Random Forest.

Datasets		ALL	ALL-{A,B}	ALL-{C,D}	{A,B}	{C,D}	{E,F,G}	{A,B,E,G}
pink slime vs All news	Acc.	**0.93**	0.79	0.90	0.77	0.70	0.82	**0.91**
	F1	**0.80**	0.52	0.75	0.51	0.45	0.64	**0.77**
pink slime vs national	Acc.	**0.92**	0.78	**0.91**	0.77	0.68	0.85	0.89
	F1	**0.75**	0.58	**0.76**	0.50	0.44	0.68	0.74
pink slime vs local	Acc.	**0.92**	0.83	0.87	0.78	0.69	0.82	**0.91**
	F1	**0.76**	0.62	0.69	0.56	0.44	0.65	**0.77**

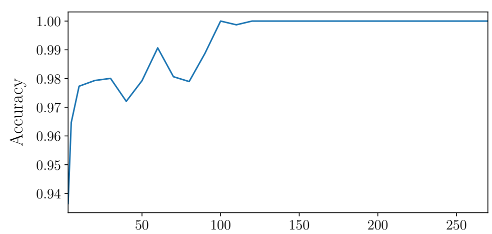

Figure 11: The increase in the accuracy when the number of common URLs increases. The collective detection label is always correct after 100 URLs.

solely local news. However, in comparison to national news, it displays a slightly lower accuracy (Table 3).

6.5 Identifying Dominant Features

In Table 3, we show the results of our ablation study, where we aim to discover the most valuable set of features. Our results demonstrate that a combination of segmentation, structural, and annotation features deliver performance results most closely to the performance of the entire set of features. On the other hand, features that related post elements, such as hashtags, mentions, or emojis, show to be less useful.

6.6 Collective Detection by URLs

To enhance the precision of detecting individual posts, we employ a grouping strategy. This strategy involves grouping posts that share the same URL, and subsequently, we designate the majority label within each group as the collective label for all posts associated with that URL. We resolve any tie by randomly assigning labels to the entire group. Our assumption is that posts in our dataset tend to have similar attributes as a group. Therefore, this grouping process should boost the accuracy of detecting pink slime URLs in individual posts. Also, it should simplify the classification task by providing a single, collective label for groups of posts that share identical URLs.

To assess the grouping concept, we examine subsets of the evaluation dataset, each consisting of post groups with equal sizes and shared URLs. In Figure 11, we show the increased accuracy as we have more URLs in a group. The accuracy reaches 100% around 100 posts per URL. We believe that this is an interesting finding as it shows that posts in isolation do not provide sufficient opportunity for learning, but collectively they do. Such observation is encountered in other works that study social media such as named entity recognition, entity linking, and misinformation [15, 22, 36, 40].

7 FUTURE WORK

The emergence of pink slime news websites should motivate researchers to study not just this type of news, but to monitor its evolution and look for any new kinds of news that might appear in the future, especially in social media. As for our current problem, there are several aspects of pink slime news we did not investigate. To gain a deeper understanding of the evolving nature of pink slime news and its impact over time, future research can focus on the user aspect of these posts. Studying the user accounts that post URLs from pink slime outlets or engaging with them can provide a further understanding of the extent of reach these outlets have in social media. Another promising avenue for future research is to study posts that are replies to pink slime posts and the content of the conversation around them [42]. Most pink slime outlets do not have comment sections on their website, and these replies can be valuable substitutes for comments. Further exploration of the content of pink slime posts and their accompanying replies may look into deeper text mining, such as sentiment [52] or stance detection [20]. Such analysis can be helpful, especially when looking at posts about controversial topics within public discourse. This extended analysis has the potential to enhance our understanding of the nuances and impact of pink slime posts in shaping online conversations.

8 CONCLUSION

In conclusion, our study unveils the distinctive nature of social media posts with URLs from pink slime outlets in contrast to posts with other national or local news. We identify unique patterns derived from the user behavior when sharing pink slime posts. We employ these valuable insights into extracting an extensive set of features and developing effective machine learning models to detect pink slime posts, ultimately aiding in the ongoing effort towards a more secure and safe social media ecosystem. Our approach achieves an accuracy of 92.55% and an F1 score of 0.80. Moreover, we show that taking advantage of multiple posts with the same URL can improve our results to reach a perfect prediction of Pink-slime posts. Our work contributes to the broader dialogue on evolving information dissemination and equips researchers with essential tools to study this issue.

ACKNOWLEDGMENTS

The authors thank the reviewers for their comments and suggestions. This work is supported in part by the U.S. NSF under grants #2236352 and #1838145.

REFERENCES

[1] Davey Alba and Jack Nicas. 2020. As Local News Dies, a Pay-for-Play Network Rises in Its Place. https://www.nytimes.com/2020/10/18/technology/timpone-local-news-metric-media.html

[2] Davey Alba and Jack Nicas. 2020. Here Are the Hundreds of Sites in a Pay-to-Play Local News Network. https://www.nytimes.com/2020/10/20/technology/timpone-network-pay-to-play-local-news.html

[3] Abdullah Aljebreen, Weiyi Meng, and Eduard Dragut. 2021. Segmentation of Tweets with URLs and its Applications to Sentiment Analysis. *Proceedings of the AAAI Conference on Artificial Intelligence* 35, 14 (May 2021), 12480–12488. https://doi.org/10.1609/aaai.v35i14.17480

[4] Lorenzo Arvanitis and McKenzie Sadeghi. 2022. Just in time for the midterms: New Entry Joins "Pink Slime" Pseudo-News Sweepstakes. https://www.newsguardtech.com/special-reports/american-independent-pink-slime-network/

[5] Clément Bénard, Gérard Biau, Sébastien da Veiga, and Erwan Scornet. 2021. Interpretable Random Forests via Rule Extraction. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 130)*, Arindam Banerjee and Kenji Fukumizu (Eds.). PMLR, 937–945. https://proceedings.mlr.press/v130/benard21a.html

[6] Priyanjana Bengani. 2019. Hundreds of 'pink slime' local news outlets are distributing algorithmic stories and conservative talking points. https://www.cjr.org/tow_center_reports/hundreds-of-pink-slime-local-news-outlets-are-distributing-algorithmic-stories-conservative-talking-points.php

[7] Priyanjana Bengani. 2020. As election looms, a network of mysterious 'pink slime' local news outlets nearly triples in size. https://www.cjr.org/analysis/as-election-looms-a-network-of-mysterious-pink-slime-local-news-outlets-nearly-triples-in-size.php

[8] Gérard Biau and Erwan Scornet. 2 015. A random forest guided tour. *TEST* 25 (2 015), 197–227. https://api.semanticscholar.org/CorpusID:14518730

[9] Ceren Budak. 2019. What Happened? The Spread of Fake News Publisher Content During the 2016 U.S. Presidential Election. In *The World Wide Web Conference* (San Francisco, CA, USA) *(WWW '19)*. Association for Computing Machinery, New York, NY, USA, 139–150. https://doi.org/10.1145/3308558.3313721

[10] Roshni Chakraborty, Maitry Bhavsar, Sourav Dandapat, and Joydeep Chandra. 2017. A Network Based Stratification Approach for Summarizing Relevant Comment Tweets of News Articles. In *Web Information Systems Engineering – WISE 2017*, Athman Bouguettaya, Yunjun Gao, Andrey Klimenko, Lu Chen, Xiangliang Zhang, Fedor Dzerzhinskiy, Weijia Jia, Stanislav V. Klimenko, and Qing Li (Eds.). Springer International Publishing, Cham, 33–48.

[11] Nicole Cohen. 2015. From Pink Slips to Pink Slime: Transforming Media Labor in a Digital Age. *The Communication Review* 18 (05 2015), 98–122. https://doi.org/10.1080/10714421.2015.1031996

[12] Antonio Criminisi, Jamie Shotton, and Ender Konukoglu. 2012. Decision Forests: A Unified Framework for Classification, Regression, Density Estimation, Manifold Learning and Semi-Supervised Learning. *Found. Trends Comput. Graph. Vis.* 7 (2012), 81–227. https://api.semanticscholar.org/CorpusID:62574244

[13] Emilio Ferrara, Onur Varol, Filippo Menczer, and Alessandro Flammini. 2021. Detection of Promoted Social Media Campaigns. *Proceedings of the International AAAI Conference on Web and Social Media* 10, 1 (Aug. 2021), 563–566. https://doi.org/10.1609/icwsm.v10i1.14772

[14] Jim Friedlich. 2021. A new report shows how partisan websites are eroding trust in legitimate local news sites in Pennsylvania. This is how we can restore faith in local journalism. https://www.lenfestinstitute.org/local-journalism/pennsylvania-media-trust-report-newsguard-fake-news/

[15] Kathleen Gregory, Paul Groth, Helena Cousijn, Andrea Scharnhorst, and Sally Wyatt. 2019. Searching Data: A Review of Observational Data Retrieval Practices in Selected Disciplines. *J. Assoc. Inf. Sci. Technol.* 70, 5 (2019), 419–432. https://doi.org/10.1002/ASI.24165

[16] Weiwei Guo, Hao Li, Heng Ji, and Mona Diab. 2013. Linking Tweets to News: A Framework to Enrich Short Text Data in Social Media. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Hinrich Schuetze, Pascale Fung, and Massimo Poesio (Eds.). Association for Computational Linguistics, Sofia, Bulgaria, 239–249. https://aclanthology.org/P13-1024

[17] Xingzhi Guo, Brian Kondracki, Nick Nikiforakis, and Steven Skiena. 2022. Verba Volant, Scripta Volant: Understanding Post-Publication Title Changes in News Outlets. In *Proceedings of the ACM Web Conference 2022* (Virtual Event, Lyon, France) *(WWW '22)*. Association for Computing Machinery, New York, NY, USA, 588–598. https://doi.org/10.1145/3485447.3512219

[18] Lihong He, Chao Han, Arjun Mukherjee, Zoran Obradovic, and Eduard C. Dragut. 2020. On the dynamics of user engagement in news comment media. *WIREs Data Mining Knowl. Discov.* 10, 1 (2020).

[19] Lihong He, Chen Shen, Arjun Mukherjee, Slobodan Vucetic, and Eduard Dragut. 2021. Cannot predict comment volume of a news article before (a few) users read it. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 15. 173–184.

[20] Marjan Hosseinia, Eduard C. Dragut, and Arjun Mukherjee. 2019. Pro/Con: Neural Detection of Stance in Argumentative Opinions. In *SBP-BRiMS*, Vol. 11549. Springer, 21–30.

[21] Marjan Hosseinia, Eduard C. Dragut, and Arjun Mukherjee. 2020. Stance Prediction for Contemporary Issues: Data and Experiments. In *SocialNLP@ACL*. 32–40.

[22] Wen Hua, Kai Zheng, and Xiaofang Zhou. 2015. Microblog Entity Linking with Social Temporal Context. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data* (Melbourne, Victoria, Australia) *(SIGMOD '15)*. Association for Computing Machinery, New York, NY, USA, 1761–1775. https://doi.org/10.1145/2723372.2751522

[23] William Williamson III and James Scrofani. 2019. Trends in Detection and Characterization of Propaganda Bots. In *52nd Hawaii International Conference on System Sciences, HICSS 2019, Grand Wailea, Maui, Hawaii, USA, January 8-11, 2019*, Tung Bui (Ed.). ScholarSpace, 1–6. https://hdl.handle.net/10125/60148

[24] Daniel Karell and Anjali Agrawal. 2021. Small town propaganda: The content and emotions of politicized digital local news in the United States.

[25] Dhruv Khattar, Jaipal Singh Goud, Manish Gupta, and Vasudeva Varma. 2019. MVAE: Multimodal Variational Autoencoder for Fake News Detection. In *The World Wide Web Conference* (San Francisco, CA, USA) *(WWW '19)*. Association for Computing Machinery, New York, NY, USA, 2915–2921. https://doi.org/10.1145/3308558.3313552

[26] Qingyuan Liu, Eduard C Dragut, Arjun Mukherjee, and Weiyi Meng. 2015. Florin: a system to support (near) real-time applications on user generated content on daily news. *VLDB Endowment* 8, 12 (2015), 1944–1947.

[27] Yang Liu and Yi-Fang Brook Wu. 2018. Early detection of fake news on social media through propagation path classification with recurrent and convolutional networks. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence* (New Orleans, Louisiana, USA) *(AAAI'18/IAAI'18/EAAI'18)*. AAAI Press, Article 44, 8 pages.

[28] Jessica Mahone and Philip Napoli. 2020. Hundreds of hyperpartisan sites are masquerading as local news. This map shows if there's one near you. https://www.niemanlab.org/2020/07/hundreds-of-hyperpartisan-sites-are-masquerading-as-local-news-this-map-shows-if-theres-one-near-you/

[29] Ryan C Moore, Ross Dahlke, Priyanjana Bengani, and Jeffrey Hancock. 2023. The Consumption of Pink Slime Journalism: Who, What, When, Where, and Why? https://doi.org/10.31219/osf.io/3bwz6

[30] Preslav Nakov, Giovanni Da San Martino, and Firoj Alam. 2022. Fact-Checking, Fake News, Propaganda, Media Bias, and the COVID-19 Infodemic. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining* (Virtual Event, AZ, USA) *(WSDM '22)*. Association for Computing Machinery, New York, NY, USA, 1632–1634. https://doi.org/10.1145/3488560.3501395

[31] Jeppe Nørregaard, Benjamin D. Horne, and Sibel Adali. 2021. NELA-GT-2020: A Large Multi-Labelled News Dataset for The Study of Misinformation in News Articles. https://api.semanticscholar.org/CorpusID:91184550

[32] Brendan Nyhan. 2019. Americans Trust Local News. That Belief Is Being Exploited. https://www.nytimes.com/2019/10/31/upshot/fake-local-news.html

[33] The Select Committee on Intelligence. 2018. OPEN HEARING ON FOREIGN INFLUENCE OPERATIONS' USE OF SOCIAL MEDIA PLATFORMS. https://www.intelligence.senate.gov/sites/default/files/hearings/CHRG-115shrg30959.pdf

[34] OpenAI. 2021. *GPT-3 Language Model*. . https://openai.com/gpt-3/

[35] Thais Mayumi Oshiro, Pedro Santoro Perez, and José Augusto Baranauskas. 2012. How Many Trees in a Random Forest?. In *Machine Learning and Data Mining in Pattern Recognition*, Petra Perner (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 154–168.

[36] Chenwei Ran, Wei Shen, and Jianyong Wang. 2018. An Attention Factor Graph Model for Tweet Entity Linking. In *Proceedings of the 2018 World Wide Web Conference* (Lyon, France) *(WWW '18)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 1135–1144. https://doi.org/10.1145/3178876.3186012

[37] Rita-Marie Cain Reid. 2014. You say "lean finely textured beef," I say "pink slime". *Food and drug law journal* 69 4 (2014), 625–47, iii.

[38] Asa Royal and Philip Napoli. 2021. Local Journalism's Possible Future: Metric Media and its Approach to Community Information Needs.

[39] John Sands. 2019. Local news is more trusted than national news — but that could change. https://knightfoundation.org/articles/local-news-is-more-trusted-than-national-news-but-that-could-change/

[40] Wei Shen, Jianyong Wang, and Jiawei Han. 2014. Entity linking with a knowledge base: Issues, techniques, and solutions. *IEEE Transactions on Knowledge and Data Engineering* 27, 2 (2014), 443–460.

[41] Dominick Sokotoff and Katherina Sourine. 2022. Pseudo local news sites in Michigan reveal nationally expanding network. https://www.michigandaily.com/news/community-affairs/pseudo-local-news-sites-michigan-reveal-nationally-expanding-network/

[42] Timo Spinde, Elisabeth Richter, Martin Wessel, Juhi Kulshrestha, and Karsten Donnay. 2023. What do Twitter comments tell about news article bias? Assessing

the impact of news article bias on its perception on Twitter. *Online Social Networks and Media* , (2023), . https://api.semanticscholar.org/CorpusID:260794379

[43] Marija Stanojevic, Jumanah Alshehri, Eduard C. Dragut, and Zoran Obradovic. 2019. Biased News Data Influence on Classifying Social Media Posts. In *NewsIR (CEUR Workshop Proceedings, Vol. 2411)*. 3–8.

[44] Axel Suarez, Dyaa Albakour, David Corney, Miguel Martinez, and José Esquivel. 2018. A Data Collection for Evaluating the Retrieval of Related Tweets to News Articles. In *Advances in Information Retrieval*, Gabriella Pasi, Benjamin Piwowarski, Leif Azzopardi, and Allan Hanbury (Eds.). Springer International Publishing, Cham, 780–786.

[45] Margaret Sullivan. 2022. Beware partisan 'pink slime' sites that pose as local news. https://www.washingtonpost.com/media/2022/06/05/pink-slime-west-cook-news-school-race-grading/

[46] Alfred Kar Yin Truong. 2009. Fast growing and interpretable oblique trees via logistic regression models. In . , , . https://api.semanticscholar.org/CorpusID:57884720

[47] Josef Verbanac and Claire Golding. 2022. Pink Slime: Fake-local News. https://iffy.news/pink-slime-fake-local-news/

[48] Prashanth Vijayaraghavan and Soroush Vosoughi. 2022. TWEETSPIN: Fine-grained Propaganda Detection in Social Media Using Multi-View Representations. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Marine Carpuat, Marie-Catherine de Marneffe, and Ivan Vladimir Meza Ruiz (Eds.). Association for Computational Linguistics, Seattle, United States, 3433–3448. https://doi.org/10.18653/v1/2022.naacl-main.251

[49] Zhongyu Wei and Wei Gao. 2015. Gibberish, Assistant, or Master? Using Tweets Linking to News for Extractive Single-Document Summarization. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval* (Santiago, Chile) *(SIGIR '15)*. Association for Computing Machinery, New York, NY, USA, 1003–1006. https://doi.org/10.1145/2766462.2767835

[50] Fan Yang, Eduard Dragut, and Arjun Mukherjee. 2021. Improving Evidence Retrieval with Claim-Evidence Entailment. In *RANLP 2021*. 1553–1558.

[51] Fan Yang, Eduard C. Dragut, and Arjun Mukherjee. 2020. Claim Verification under Positive Unlabeled Learning. In *ASONAM*. 143–150.

[52] Fan Yang, Eduard C. Dragut, and Arjun Mukherjee. 2020. Predicting Personal Opinion on Future Events with Fingerprints. In *COLING*. 1802–1807.

[53] Fan Yang, Arjun Mukherjee, and Eduard Dragut. 2017. Satirical News Detection and Analysis using Attention Mechanism and Linguistic Features. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Martha Palmer, Rebecca Hwa, and Sebastian Riedel (Eds.). Association for Computational Linguistics, Copenhagen, Denmark, 1979–1989. https://doi.org/10.18653/v1/D17-1211

[54] Yigeng Zhang, Fan Yang, Yifan Zhang, Eduard Dragut, and Arjun Mukherjee. 2020. Birds of a feather flock together: Satirical news detection via language model differentiation.

[55] Ryan Zickgraf. 2022. How 'pink slime' journalism exploits our faith in local news. https://www.washingtonpost.com/outlook/2022/08/12/pink-slime-jounrnalism-local-news/