Automated Pink Slime Detection from Social Network Features*

Christine Sowa Lepird[1[0000−0002−9350−5627]] and Dr. Kathleen M. Carley[1[0000−0002−6356−0238]]

Carnegie Mellon University, Pittsburgh PA 15213, USA csowa@andrew.cmu.com
http://www.casos.cs.cmu.edu

Abstract. A network of over a thousand seemingly local news sites was found to be controlled by a handful of partisan national organizations who have plans for expansion to new regions. These polarized sites (referred to as "pink slime") can undermine trust in the local news in susceptible communities. While over a thousand of these sites have been labeled and identified, discovering new sites remains challenging; the current process of discovering new sites involves a tedious IP address lookup process. This research proposes a methodology for detecting emerging pink slime sites through analyzing the network of Facebook pages that share content linking to these sites (as well as others of known credibility) and assigning credibility scores to the news domains in order to quickly flag any suspicious activity. It allows researchers to efficiently surveil the social media landscape to find new sources of pink slime as they emerge. This paper demonstrates the importance of the new features through machine learning validation and then applies the methods to a recent dataset for the discovery of new pink slime sites.

Keywords: pink slime · social network analysis · misinformation.

1 Introduction

Since 2019, there has been a rapid spread of regional news sources that casual viewers have a hard time differentiating from reliable local news sources [5]. These news sources, filled with mostly automated, low-quality, partisan reporting were dubbed "Pink Slime" by journalist Ryan Smith in 2012 [22]. Smith used the term to draw a comparison to the meat producers adding cheap additives to beef, in this case adding cheap reporting to a self-reported news outlet. Cohen described the many technologies and outsourcing practices available to publish

* This work was supported in part by the Office of Naval Research (ONR) Award N00014182106, the Knight Foundation, the Center for Computational Analysis of Social and Organizational Systems (CASOS), and the Center for Informed Democracy and Social-cybersecurity (IDeaS). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the ONR or the U.S. government.

automated reporting back in 2015 (before the invention of text generating platforms like ChatGPT) [10].

There is reason to be fearful of false reporting that purports to serve a particular region: while one in five local newspapers have closed in the past fifteen years [21], American trust in local news organizations has remained higher than that of national news organizations [11]. Half of all U.S. counties have only one weekly local newspaper [19]. To exploit the trust in and dearth of local reporting, organizations like Metric Media LLC have created almost 1,000 local news sites [6]. While there is a lack of authentic local reporting by local reports, it remains highly trusted, and creators of these networks are taking advantage of this trust.

Research looking into Metric Media found that "455 [of their sites] never featured a human-bylined article on their front page during the observation period" [20] and "many authors wrote stories for publications in several states" [20].

The organizations like Media Metric that control vast swaths of pink slime sites do not appear to have foreign ties [6], but they are currently financed by political candidates and political action committees with the hope of swaying election results. When speaking of threats to election integrity, Alex Stamos, director of the Stanford Internet Observatory, remarked "The issue ... is not going to be foreign interference. It's much more likely that legitimate domestic actors possibly operating under their own name — with LLCs or corporations with very shady funding that are not required to disclose what that funding is — are going to dominate the online conversation about the outcome of the election. [18]"

Pink slime news sources do not exist in silos. Many of the known sources of pink slime have their own associated social media accounts on platforms like Facebook to amplify the spread of the messaging to the community (as the names of these sites frequently have the targeted community in the domain name). Almost 70% of Americans get their news from the social media platform Facebook [12], but not all of this news is coming from quality news sources; 15% of referrals to fake news sites are coming from Facebook [13]. While research shows that platforms like Reddit and 4chan are not as susceptible to pink slime [8], the methods only looked into more national news subreddits instead of smaller communities and local subreddits.

This paper takes an algorithmic approach to determine the credibility of news domains based on the network of the public Facebook pages that are sharing the links with their followers. Additional features, including whether the name of the site contains a regional locale (extracted through natural language processing), reactions to the posts that share the news domain, and how frequently the domains appear in the dataset are included in a machine learning model to predict the labels of the news sites shared. The network features defined in this paper can be used as a ranking statistic for a fast human in the loop analysis of tens of thousands of domains. It uses a dataset surrounding the 2020 U.S. Presidential election and the "reopen" movement of the COVID-19 pandemic to validate the importance of the network features of shared domains. Finally, it applies these findings and features to the 2022 U.S. Midterm Election dataset to allow for a quick analysis of thousands of news domains to uncover new sources of pink slime.

2 Literature Review

The leading voice in this research has been that of Priyanjana Bengani, a senior research fellow at Columbia Journalism School's Tow Center for Digital Journalism. In order to find pink slime sites, she looked to see which websites shared the same tracking identifiers, IP addresses, and servers [5].

Bengani mapped out the IP addresses and servers these news domains utilize to find the relations between the organizations. While she found it was less common for pink slime sites to share Google Analytics IDs, a handful of *Media Metric*'s sites shared three of these IDs while Franklin Archer, Local Government Information Services (LGIS), and LocalityLabs shared three other unique identifiers. Other identifiers - like Quantcast and NewRelic IDs - were shared across various pink slime networks [5].

Bengani followed the RSS feeds of 189 *Media Metric* sites and found over 50,000 stories published over a two week period. Of these, only 15,000 (around 30%) of the stories were unique. Furthermore, only 100 of these stories were attributed to human reporters while the rest utilized APIs or press releases [5].

The process of finding these sites involved manual entry of domains into platforms that found sites that shared identifiers, but it would frequently lead to dead ends. "Websites with the same IP address are not necessarily related though. Websites hosted by companies like Squarespace and Wix... are likely to share IP addresses ... lookups on two websites - one hosted by Squarespace and one by Wix - showed more than five million domains for each of the two IP addresses. Similarly, "parked" domains - domains that have been registered but not linked to any web-hosting - are likely to share IP addresses. [5]" However, finding new sources of pink slime not controlled by the companies Bengani identified in her publications is an unanswered question. This paper aims to find new sources of pink slime using network measures.

3 Feature Engineering

By pulling the 1,000 most recent instances of known pink slime sources [15] being shared to Facebook pages via the CrowdTangle API some basic inferences can be made about the spread of this information. Furthermore, when looking at a network of Facebook pages to the parent organization of the pink slime sites they shared, there are many Facebook pages that are posting links to pink slime sites owned by multiple parent organizations. One such Facebook page, "Democrats of the Alachua County Area", linked to 5 different parent organizations of pink slime sites. Most of the Facebook pages sharing these sites were smaller (under 1,000 followers) and targeted a hyper-local area.

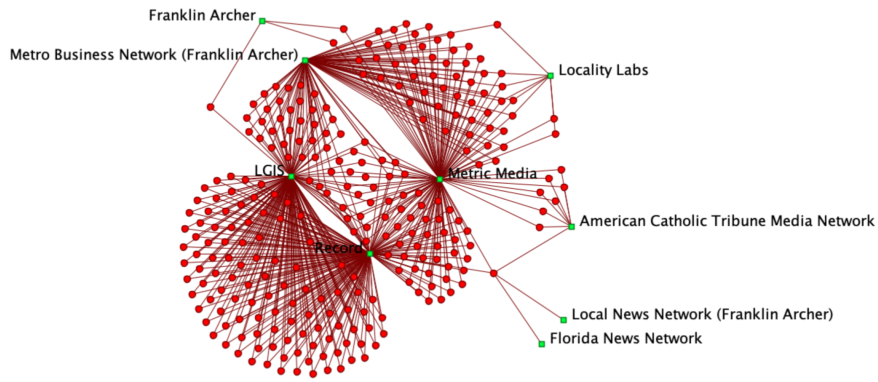

Fig. 1. Network visual of Facebook Pages linking to Parent Organizations of Pink Slime Sites

This background information informs the methodology in this paper that there are Facebook pages geared for local communities that share information going to multiple pink slime parent organizations. By searching Facebook pages for polarized topics pertaining to local communities, some of these pink slime websites can be found.

4 Methodology

Other academic publications can serve as an inspiration to understanding methods to find relevant websites of interest. In a precursor to PageRank, Kleinberg proposed a model that found authoritative webpages for search topics by analyzing the relationship between "authorities for a topic and those pages that link to many related authorities ... [called] hubs" [16]. Using the reasoning that "hubs and authorities exhibit what could be called a mutually reinforcing relationship: a good hub is a page that points to many good authorities; a good authority is a page that is pointed to by many good hubs" [16], they created a hub score that incorporated the authority scores of the nodes pointing to it. Likewise, they gave nodes a high hub score if it linked to nodes that are authorities on the subject. While others have taken these methods and applied them to citation networks and ranking academic journals, this research sets about viewing certain Facebook pages as hubs of disinformation and building out the network of authorities that they share.

The methodology for this research involves creating network features of the news domains to signal their credibility based on which hubs are sharing the authorities and what other authorities those hubs were sharing. In this research, Facebook pages act as the hubs, and the domains they share are the authorities. The labeled news domains are utilized to test the validity of the features. From there, the larger, unlabeled domain dataset was analyzed using the set of features to find new source of pink slime.

In order to determine the credibility of domains, the following equation was implemented. First, each Facebook page present in the dataset was given a score to indicate the proportion of content it shared that originated from a known source of fake news or pink slime (Noncredible Sharer Score). The higher the Noncredible Sharer Score, the lower the credibility of the Facebook Page. Then, the domains could be scored by averaging the Noncredible Sharer Score of the pages that shared the domain.

A common use case for this scoring algorithm would be to upload a Facebook dataset for topic of interest, run the algorithm and find the Noncredibility Scores, and then sort by descending Noncredibility Scores. This ranks sites from lowest credibility and is a place where the analyst can then use visual inspection and visit some of the highest scoring domains.

$$\text{Noncredible Sharer Score}_k = \frac{1}{n}\sum_{j=1}^{n} I_j \tag{1}$$

$$\mathrm{I}_j = \begin{cases} 1 & if \text{domain j is a labeled pink slime or fake news site} \\ 0 & otherwise \end{cases}$$

k = Facebook Page linking to domains
n = number of domains shared by Facebook Page$_k$

$$\text{Noncredibility Score}_z = \frac{1}{|K|}\sum_{k=1} \text{Noncredible Sharer Score}_k \tag{2}$$

K = set of Facebook Pages sharing a given unlabeled domain
In meta-network terms, in a network that linked Domains to Domains by the Facebook Pages that shared them (Domain x Facebook Page x Domain), the Noncredibility Score for any given domain would be the percentage of first degree neighbors that were known sources of pink slime or fake news.

For example, a page that shared 5 news articles, none of which link to domains that were labeled as fake news was assigned a score of zero. However, if a page shared 10 news articles and 5 of the articles were to domains labeled as fake news, that Facebook Page would be assigned a Noncredible Sharer Score of 0.50. Each Facebook Page that shared a domain had its score averaged to assign the domain a Noncredibility Score.

These network features, as well as those described in the Data Description, were utilized as inputs in a machine learning model to predict whether or not a domain in the dataset was pink slime. 70% of the known sources of fake news, pink slime, and real news are utilized in training the model to see if they can accurately predict the legitimacy of the 30% of withheld domains. The 30% of withheld domains are treated as unlabeled news sources during the calculation of the network measures. The XGBoost model, an open-source implementation of the gradient boosted trees algorithm, is used to perform the training and validation due to its high efficiency and accuracy [4].

5 Data Description

The social monitoring tool CrowdTangle [1] is used to collect posts containing the phrase "reopen" or "election" linking to outside URLs on public Facebook pages in the year of 2020. As this was the start of the global pandemic COVID-19, the "reopen" phrase was used to protest establishments being closed. The particular phrase was selected for this study due to the nature of pages insisting that their local region reopen. Due to the number of regional elections that coincided with the 2020 U.S. Presidential election, the election phrase is also utilized to acquire a dataset of the magnitude to build out the necessary network. This dataset includes 3,223,269 posts to public Facebook pages in 2020 that link to 61,692 unique domains.

Once the dataset of Facebook posts was acquired from Crowdtangle, the links to outside domains were extracted and were flagged if they went to known sources of real news, fake news, or pink slime.

Throughout this paper, "known sources of real news, fake news, or pink slime" will be defined as news articles originating from domains listed as "real, fake, or pink slime" in the media thesaurus compiled by the CASOS University Center at Carnegie Mellon University. The media thesaurus has been compiled from multiple publicly available lists of news media URLs and media organizations' Twitter accounts: Media Bias/Fact Check [2] lists many news sites and rates how factual and credible the reporting is for many; the George Washington University Dataverse [17] has a list of over 9600 Twitter accounts for media organizations, derived from over 160 million tweets between 2016 and 2020; the Columbia Journalism Review site has been a source for hundreds of "pink slime" news outlet domains [15] that often publish biased, algorithmically produced stories; there is also a Github repository [3] of unreliable, misleading, and/or "fake" news sources that includes lists from Snopes Field Guide, Melissa Zimdars' OpenSources, Wikipedia, and others. There is often overlap between these sources, particularly for the less factual news outlets; to resolve any conflicts that emerge between the sources, the thesaurus errs on the side of not labeling a news source in question as fake news. Additionally, the websites owned by Courier Newsroom were included and labeled as pink slime; while there has been some debate over whether these sites remain as pink slime after changing some funding structures, they exhibited many properties of pink slime sites in the time period for which the data for this research was collected.

Upon network observation in the dataset, certain known pink slime domains were pendant nodes in the network, only shared by a single news source (typically of the same name as the news domain). Since these isolated instances of pink slime would result in minimal network values (such as the Noncredibility Score), additional features were included in the feature selection phase. These included: the number of distinct Facebook pages that shared each domain, and the average reaction features to the Facebook page posts for each domain.

Finally, in order to determine whether a domain was targeting a local region, the domain name was compared to a set of United States-based Gazetteers [7].

The names of each county, city, region, and town in the country were analyzed to see if that specific string was included in the domain name.

Table 1 breaks down some of the statistical features of the larger dataset. The vast majority (96%) of the domains in this dataset are to currently unknown or unlabeled news sites per the media thesaurus. This reiterates the problem of finding vast swaths of unlabeled domains and not knowing where to start when manually reviewing tens of thousands of URLs to find a handful of pink slime (or even fake news) sites.

Of the labeled domains, only 1.3% are known sources of pink slime and 9.3% are to known sources of fake news. Real news domain occurred most frequently (on average each domain appeared 492 times in the dataset) with fake news domains seeing approximately half as many occurrences, and the average pink slime domain was embedded within only 72 posts to Facebook pages, most likely due to the more targeted, smaller audiences this type of news was designed to attract. The 63% of known pink slime domains contain a location name in the text of the domain. This is substantially higher than that of known fake and real news (17% and 22%, respectively) and speaks to the desire of the pink slime owners' goals of targeting specific locales.

Table 1. Feature analysis for the input dataset.

	Dataset Size — Average Occurrences — Average Pages — Location in Domain Name			
Pink Slime	30	71.6	3.1	63%
Fake News	201	251.3	9.17	17%
Real News	1,941	492.3	47.20	22%
Not labeled	59,751	35.6	3.08	31%

6 Results

To see if the labels of "pink slime", "real news", and "pink slime" could be predicted using the features extracted in the Data Description and Methods sections, a Naive estimator that assigned classifications to the dominant class (real news) was used as a baseline to compare to another estimator, the XGBoost model. In comparison, the XGBoost model demonstrated an increase in differentiating classification between real news and fake news. Fake news and real news saw an uptick in precision and recall when the actual estimator was used. However, class imbalance played a role in the classification of pink slime. Due to the extremely low support for pink slime instances in the dataset, the model struggled to find a threshold for pink slime that increased the model's accuracy.

While the accuracy value of the pink slime classification was low, the overall intent of the features are to use them as a starting point of analysis to find higher likelihood domains faster. Since the goal of this research isn't to do human-out-of-the-loop, the accuracy scores don't paint the entire picture of the model's performance. Overall its area under the curve (AUC) of 0.77 provides an acceptable value to use the network features of the model as a ranking heuristic in large scale data analysis [14]. When looking at how well the model could predict the presence of fake news, the ROC curve is slightly stronger with an AUC of 0.78.The presence of real news was similarly well predicted to that of pink slime. The ROC curve has an AUC of 0.77 shows an ability for this approach to be generalizable to these three difference news types.

When analyzing the feature importance of the variables included in the model, only two had an effect greater than there. The strongest one was the domain's Noncredibility Score. The other attribute, with 25 times less importance, was the average number of likes a post from the domain received. These results validate that using a network-based approach to assessing credibility can be used to predict news labels.

7 Analysis

The primary function of the attributes discussed in the Methodology was to apply these measures to new datasets to quickly rank and assess probable new pink slime news agencies. As seen in the Data Description, labeled pink slime sites are a small minority of the labeled news sources in a given dataset. Due to the issue of class imbalance, the general results for predicting pink slime show low accuracy but promising AUC curves. Given these results, using the model itself is not to output a list of two possible pink slime sites in a dataset of tens of thousands of domains. Rather, the intent is to apply the new attributes to new datasets and sort by the Noncredibility Score for a shorter, more targeted list of potential new pink slime sites. While it did not show measurable effect on the model, the attribute describing whether a domain contains a location can still be used as a filter for the quickest analysis.

Applying these methods to a more recent dataset to prove longevity of the features, a dataset of Facebook Pages posting about the 2022 United States Midterm Elections was analyzed to find new, emerging sources of pink slime. This dataset included 10,223 domains which would be challenging to analyze using visual inspection alone. Through analysis of the domains by Noncredibility Score, the first to jump out as a new source of pink slime is georgiastarnews.com should be included as a new source of pink slime. Further investigating on the site shows that it is owned by the same media group (Star News Digital Media) that owns *The Tennessee Star*, *The Ohio Star*, *The Michigan Star*, *The Minnesota Sun*, *The Virginia Star*, *The Georgia Star News*, *The Florida Capital Star*, *The Arizona Sun Times*, *The Wisconsin Daily Star*, *The Pennsylvania Daily Star*, and *The Connecticut Star*. *The Georgia Star* appeared in the dataset in 27 Facebook posts, most of which received two or fewer interactions (likes, shares, or comments). The most interacted with post (5 interactions) was sharing an article with the title, "Georgia Moves Forward with Plan to Implement Work Requirements for Medicaid Coverage" [9].

In addition to finding the pink slime sites that are targeting the local regions in the USA, this model can also find previously-uncategoritzed domains that have the same characteristics as other low credibility news sites.

8 Limitations

The primary limitation of this research is that not all unlabeled news sites are of unknown credibility. In order to combat this, it is recommended to perform exploratory data analysis on the unlabeled shared sites in the dataset and find the most shared sites, research their news label on Media Bias/Fact Check and other fact checking sites, and add the new label into the labeled news sources thesaurus.

9 Conclusion

This paper utilizes a novel method of network analysis to score and rank news domains by Facebook pages as misinformation versus authentic. It validates the importance of the metric through a classification algorithm that categorizes the domains as real news, fake news, or pink slime. It also applies the algorithms to a recent dataset to help the analyst uncover previously unlabeled pink slime domains.

While the individual utilizing this information would need to determine a scoring threshold below which he felt comfortable trusting the news domains, it provides a reliable metric to compare the trustworthiness of sites - in particular, sites for which this information does not exist by current formal means.

References

1. CrowdTangle, https://www.crowdtangle.com/
2. Media bias/fact check news, https://mediabiasfactcheck.com/
3. unreliable-news/data at master · hearvox/unreliable-news, https://github.com/hearvox/unreliable-news
4. XGBoost documentation — xgboost 1.7.4 documentation, https://xgboost.readthedocs.io/en/stable/
5. Bengani, P.: 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/, publication Title: Columbia Journalism Review

6. BENGANI, P.: The metric media network runs more than 1,200 local news sites. here are some of the non-profits funding them., https://www.cjr.org/tow_center_reports/metric-media-lobbyists-funding.php/
7. Bureau, U.C.: Gazetteer files, https://www.census.gov/geographies/reference-files/time-series/geo/gazetteer-files.html, section: Government
8. Burton, A.G., Koehorst, D.: Research note: The spread of political misinformation on online subcultural platforms . https://doi.org/10.37016/mr-2020-40, https://misinforeview.hks.harvard.edu/article/research-note-the-spread-of-political-misinformation-on-online-subcultural-platforms/
9. Byrne, M.: Georgia moves forward with plan to implement work requirements for medicaid coverage, https://georgiastarnews.com/news/georgia-moves-forward-with-plan-to-implement-work-requirements-for-medicaid-coverage/tcsquare/2022/12/01/
10. Cohen, N.S.: From pink slips to pink slime: Transforming media labor in a digital age **18**(2), 98–122. https://doi.org/10.1080/10714421.2015.1031996, https://doi.org/10.1080/10714421.2015.1031996, publisher: Routledge _eprint: https://doi.org/10.1080/10714421.2015.1031996
11. GOTTFRIED, J., LIEDKE, J.: Partisan divides in media trust widen, driven by a decline among republicans https://www.pewresearch.org/fact-tank/2021/08/30/partisan-divides-in-media-trust-widen-driven-by-a-decline-among-republicans/
12. Gramlich, J.: 10 facts about americans and facebook, https://www.pewresearch.org/fact-tank/2021/06/01/facts-about-americans-and-facebook/
13. Guess, A.M., Nyhan, B., Reifler, J.: Exposure to untrustworthy websites in the 2016 US election **4**(5), 472–480. https://doi.org/10.1038/s41562-020-0833-x, https://www.nature.com/articles/s41562-020-0833-x, number: 5 Publisher: Nature Publishing Group
14. Hosmer, D.W., Lemeshow, S.: Area under the ROC curve **2**, 160–164
15. JOURNALISM, T.C.F.D.: Domains as of august 3, 2020, https://datawrapper.dwcdn.net/TqILa/2/
16. Kleinberg, J.M.: Authoritative sources in a hyperlinked environment
17. Littman, J., Wrubel, L., Kerchner, D., Bromberg Gaber, Y.: News outlet tweet ids. https://doi.org/10.7910/DVN/2FIFLH, https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/2FIFLH
18. Murphy, H., Venkataramakrishnan, S.: Local news is drowning in 'pink slime' ahead of US election
19. Muse Abernathy, P.: THE EXPANDING NEWS DESERT, https://www.cislm.org/wp-content/uploads/2018/10/The-Expanding-News-Desert-10_14-Web.pdf
20. Royal, A., Napoli, P.M.: Local journalism without journalists? metric media and the future of local news **8**, 119–147. https://doi.org/10.56140/JOCIS-v8-2, https://jocis.org/2022/12/27/local-journalism-without-journalists-metric-media-and-the-future-of-local-news/
21. Takenaga, L.: More than 1 in 5 u.s. papers has closed. this is the result. https://www.nytimes.com/2019/12/21/reader-center/local-news-deserts.html
22. Tarkov, A.: Journatic worker takes 'this american life' inside outsourced journalism https://www.poynter.org/reporting-editing/2012/journatic-staffer-takes-this-american-life-inside-outsourced-journalism/