Podcast Outcasts: Understanding Rumble's Podcast Dynamics

Utkucan Balci[1], **Jay Patel** [1], **Berkan Balci** [2], **Jeremy Blackburn** [1]
[1] Binghamton University
[2] Middle East Technical University
ubalci1@binghamton.edu, jpatel67@binghamton.edu, e252615@metu.edu.tr, jblackbu@binghamton.edu

Abstract

The rising popularity of podcasts as an emerging medium opens new avenues for digital humanities research, particularly when examining video-based media on alternative platforms. We present a novel data analysis pipeline for analyzing over 13K podcast videos (526 days of video content) from Rumble and YouTube that integrates advanced speech-to-text transcription, transformer-based topic modeling, and contrastive visual learning. We uncover the interplay between spoken rhetoric and visual elements in shaping political bias. Our findings reveal a distinct right-wing orientation in Rumble's podcasts, contrasting with YouTube's more diverse and apolitical content. By merging computational techniques with comparative analysis, our study advances digital humanities by demonstrating how large-scale multimodal analysis can decode ideological narratives in emerging media format.

1 Introduction

In today's world, visual elements play an important role in communication and engagement (Ling et al., 2021). The rise of social media, video-sharing platforms, and visual-centric content has transformed how we perceive information. This shift has transformed various media formats, including podcasts. The integration of visual elements into podcasts has given rise to video podcasts, making it increasingly popular (Grunfeld, 2023).

From Joe Rogan's $200M exclusivity deal with Spotify (Rosman et al., 2022), to Andrew Tate's misogynistic rhetoric that led him to get banned from mainstream platforms (Wilson, 2022), and Donald Trump's podcast appearances during the 2024 US election period (DeLetter, 2024; Mahdawi, 2024; Lewis, 2024), the cultural impact of this medium is increasingly visible in mainstream and alternative platforms alike. For scholars in the digital humanities, these developments open new avenues for exploring how language, imagery, and ideology intersect to shape collective understandings of politics and culture.

In particular, Rumble, a self-described neutral video-sharing platform (Brown, 2022), hosts a range of high-profile, often deplatformed, controversial figures, e.g., Donald Trump, Alex Jones, Andrew Tate, who have cultivated large followings despite bans or restrictions on sites like YouTube (Mak, 2021; Farah, 2023; Klee, 2023). Notably, in August 2024, following the arrest of Telegram's CEO in France under allegations of failing to adequately moderate content on the platform, Rumble's CEO announced his departure from Europe, due to concern of encountering comparable challenges to his platform (Cebi, 2024). To date, no research has explored whether right-wing podcasters are merely a segment of Rumble's podcasting selection or if the platform serves as a bastion for right-wing propaganda.

To address this gap, our study integrates computational text analysis and visual embedding techniques with a novel data analysis pipeline to answer the following research question: How do political biases manifest in the narratives and imagery of video podcasts on Rumble?

Drawing on 13K podcast videos, equivalent to 526 days (more than 750K minutes) of video content, from both YouTube and Rumble, our approach integrates speech-to-text transcription with transformer-based topic modeling and contrastive learning for image analysis. Our research reveals a clear pattern: Rumble exhibits a noticeable right-wing bias in its audio and visual content, whereas YouTube primarily remains apolitical, concentrating on mainstream subjects.

Contributions. We make several contributions. First, we conduct the first large-scale data-driven study on video podcasts, where we provide a layered analysis of platform bias on video podcast

Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities, pages 48–62
May 3-4, 2025

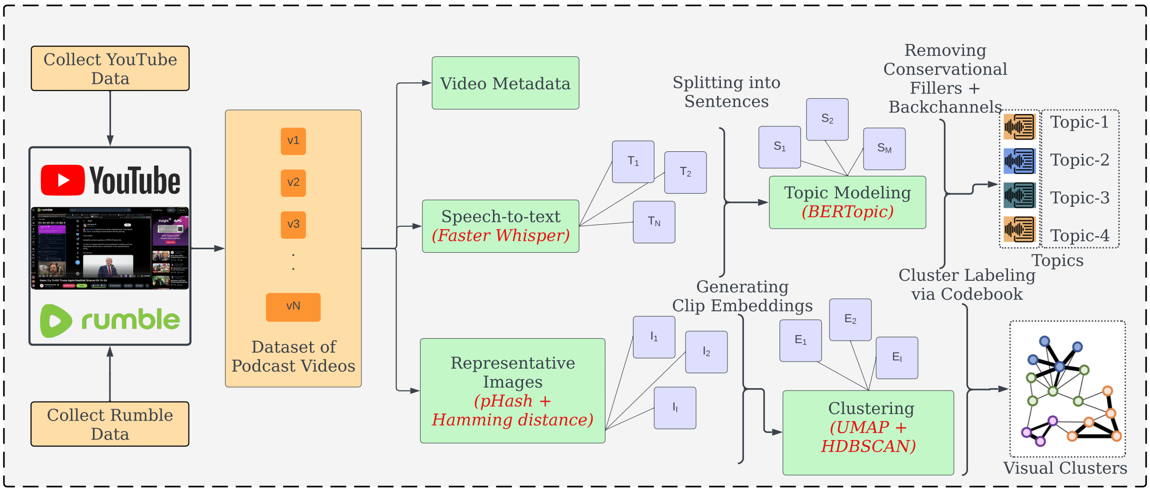

Figure 1: Our Podcast video processing and analysis pipeline: 1) Collect data from Rumble and YouTube, 2) Extract video metadata 3) Use 'faster-whisper' to generate transcriptions, 4) Use BERTopic with MPNet vectors to find topics, 5) Sample representative images of podcast videos, 6) Generate CLIP Embeddings of the representative images to cluster them using UMAP and HDBSCAN.

content, moving beyond simply confirming that Rumble is predominantly right-wing. Second, we make a methodological contribution (see Figure 1) by demonstrating how the rise of video podcasts necessitates new analytical techniques. As the social media landscape shifts, alongside API restrictions on platforms like Twitter and Reddit, in addition to video podcasts gaining prominence, our paper showcases how state-of-the-art methods on speech-to-text transcription, transformer-based topic modeling, and visual content analysis can be combined to offer a holistic analysis of multimedia content. Our methodology provides a blueprint for future studies on political multimodal content engagement in the context of evolving social platforms, particularly in video podcast content. Overall, our contributions enrich digital humanities by providing new avenues for interpreting multimodal political communication and understanding the cultural dynamics of digital media.

2 Background & Related Work

The term "podcast" was mentioned for the first time in 2004 (Robertson, 2019). In 2006, the PEW Research Center defined podcasting as a method of distributing audio and video content online, distinguishing it from earlier formats by enabling automatic transfers to users' devices for on-demand consumption, often on portable digital music players such as MP3 players or iPods (Pew, 2006). Although the core of this definition remains relevant, the influence of podcasts has evolved remarkably over time. Particularly with the widespread use of social media, podcasts are now viewed by millions of users on video streaming platforms (Escandon, 2024).

Over a 12-month period spanning parts of both 2022 and 2023, nearly half of the US adult population reported having listened to a podcast, with one-fifth frequently doing so multiple times a week (Shearer et al., 2023). This proportion increases to one-third among young adults under 30. Of the U.S. adults who listened to a podcast during this interval, 46% were Republicans and 54% Democrats, with 65% of Republican and 69% of Democratic listeners tuning into news-related podcasts.

Podcasts as vectors of political discourse. A range of studies have explored the impact of political podcasts on individuals' political engagement and attitudes (Cho et al., 2023; Euritt, 2019; Lee, 2021; Kim et al., 2016a; MacDougall, 2011; Rae, 2023; Sterne et al., 2008). Notably, consuming podcasts is linked to heightened levels of personalized politics, a process where individuals integrate new information into their existing ideological frameworks to develop more personalized political understandings (Bratcher, 2022). (Kim et al., 2016b) further explored the relationship between partisan podcast consumption, emotional responses, and political participation, finding that selective exposure to partisan podcasts can shape emotional reactions to political candidates, thereby affecting political engagement. (Chadha et al., 2012) also observed a positive correlation between using podcasts for news and increased political participation, suggesting that podcasts might boost political involvement among individuals. Rizwan et al. (Rizwan et al., 2025) analyzed over 9,000 episodes from 31 U.S. political podcast channels, finding that many popular shows had a majority of episodes containing at least one toxic segment. While many of these studies have focused on YouTube podcasts and estimating the ideology of YouTube channels (Dinkov et al., 2019; Lai et al., 2022), there has yet to be a large-scale, data-driven analysis of the political bias in popular YouTube podcast channels.

What is Rumble? Launched in 2013 as a YouTube alternative, Rumble gained notable attention during the COVID-19 pandemic (McCluskey, 2022). The number of monthly users on the platform increased from 1.6 million in the Fall 2020 to 31.9 million by the beginning of 2021 (Pramod, 2021) and eventually hit a peak of 80 million active users monthly by the end of 2022 (Brown, 2022). While the platform's founder asserts its neutrality (Brown, 2022), Rumble has become particularly known for being a haven for right-leaning public figures, including Andrew Tate, Rudy Giuliani, and Alex Jones (Farah, 2023). Despite its popularity, research on this platform is limited. Previous work (Stocking et al., 2022) estimated that over 75% of US adults who regularly use Rumble for news are Republicans or lean towards the Republican Party. This survey also notes that Rumble is a regular news source for 2% of the American population. While Rumble has been mentioned in research related to the alt-right and the Russian invasion of Ukraine (Chen and Ferrara, 2023; Aliapoulios et al., 2021), and some of Andrew Tate's Rumble channel podcast episodes have undergone analysis (Sayogie et al., 2023), similar to YouTube, a large-scale, data-driven research analyzing political bias in popular Rumble podcast channels has yet to be conducted.

3 Dataset

To collect podcast videos from Rumble, we develop a custom crawler that extracts video information from the "Podcasts" section on the home page of rumble.com (Rumble, 2023). This crawler systematically navigates through the URLs, scanning pages in this section until no new pages are found. We initially ran our crawler in October 2022 and conducted a follow-up in early 2023 to ensure coverage of the entire year. In the first week of July 2023, we revisited the video pages in our collection to update their metadata and remove any podcast videos that were no longer accessible. The rationale for this approach is to allow at least six full months for the metadata of each video (e.g., views) to stabilize and reflect their actual values. As a result, we compile a dataset of 6,761 videos from 246 channels, posted between August 27, 2020, to January 1, 2023. To remove non-English content from our dataset we perform language verification (see Appendix A for details) and transcribe the podcast videos using a reimplementation of OpenAI's Whisper (OpenAI, 2022). A more comprehensive look at this dataset can be found in the corresponding dataset paper (Balci et al., 2024). As we aim to analyze popular Rumble podcast channels, we limit our dataset to include the top 100 channels with the highest cumulative podcast video views. This subset comprises a total of 6,272 videos, accounting for 99% of all podcast views on Rumble. Table 3 presents the top 20 channels by cumulative views, along with their total number of videos and average view counts in our dataset. We refer to this dataset as D_{rumble} throughout the remainder of this paper.

YouTube. Using the YouTube API, we extract video metadata categorized as podcasts from YouTube's list of top 100 popular podcast creators (YouTube, 2023). Our manual inspection of these channels revealed non-English content and videos unrelated to podcasts (e.g., music and gospel). To refine our dataset, we used the following criteria: 1) videos must be categorized under the Podcast tab within the channel's playlists, 2) the content must be in English, and 3) genres unrelated to podcasts (e.g., gospel and music) are excluded. In the refinement process, we randomly select and manually inspect 5 videos from each playlist, subsequently eliminating playlists that failed to meet our criteria. This process yields a dataset of more than 20K videos from 69 channels, with all videos available and their metadata collected during the first week of July 2023. For a comparative analysis with the Rumble dataset, we adjust the YouTube dataset to match the monthly video distribution and the total number of podcast videos in the Rumble dataset. This way, by aligning the dataset with the specific months, we account for the potential in-

	YouTube				Rumble			
no.	**Top 5 Topic Words**	**Left**	**Center**	**Right**	**Top 5 Topic Words**	**Left**	**Center**	**Right**
1	saying, im, know, mis, straight	–	–	–	vaccine, vaccinated, vaccines, vaccination, unvaccinated	✓	✓	✓
2	bengals, nfc, raiders, eagles, afc	–	–	–	ballots, mailin, ballot, absentee, harvesting	✓	✓	✓
3	billionaire, richest, multimillionaire, mil, paid	–	–	–	ukrainians, crimea, putin, ukraine, ukrainian	✓	✓	✓
4	niggas, nigga, ns, dappin, doin	–	–	–	roe, abortion, abortions, wade, prolife	✓	✓	✓
5	book, books, chapter, bestseller, chapters	✓	✓	✓	mask, masks, masking, n95, masked	–	✓	✓
6	ukrainians, crimea, putin, ukraine, ukrainian	✓	✓	✓	rumble, rumbles, rumblecom, rants, rumblers	–	–	–
7	interview, interviews, interviewer, interviewing, interviewed	–	✓	–	biden, bidens, joe, administration, antibiden	✓	–	✓
8	feel, antioch75, wesh, pico, recant	–	–	–	alito, clarence, justices, roberts, gorsuch	✓	✓	✓
9	lakers, clippers, nets, knicks, celtics	–	–	–	desantis, ron, desantiss, crist, trumpdesantis	–	–	–
10	vaccine, vaccinated, vaccines, vaccination, unvaccinated	✓	✓	✓	democrats, dems, republicans, gop, twoparty	✓	–	✓
11	entrepreneur, entrepreneurship, entrepreneurs entrepreneurial, business	–	–	–	inflation, inflationary, reduction, hyperinflation, inflations	✓	✓	✓
12	roe, abortion, abortions, wade, prolife	✓	✓	✓	book, books, chapter, bestseller, chapters	✓	✓	✓
13	rudn, stk, know, presses, ironically	–	–	–	mainstream, media, medias, lamestream, trusts	–	–	–
14	masks, mask, masking, n95, masked	–	✓	✓	lefties, leftism, lefts, left, leftists	–	–	✓
15	rapping, rap, hiphop, hip, hop	–	–	–	tweet, retweeted, retweet, tweeted, tweets	–	–	–
16	sober, beers, drink, beer, drunk	–	–	–	youtubes, youtube, youtubers, youtuber, demonetized	–	–	–
17	corvette, lamborghini, bentley, honda, mercedes	–	–	–	denier, congressperson, reelection, hillary, caucusing	✓	–	✓
18	numbers, numerals, staggering, number, digits	–	–	–	fbi, fbis, disband, disbanded, informants	–	–	–
19	dangs, bagot, shrugs, sarcastically, becca	–	–	–	border, borders, crossings, immigration, apprehensions	–	–	✓
20	podcasting, podcasts, podcaster, podcast, podcasters	–	–	–	science, scientific, scientists, antiscience, scientist	–	–	✓

Table 1: Comparison of the top 20 topics on YouTube and Rumble. The presence of a checkmark signifies that the topic appears in the top 20 topics of baseline political podcasts.

fluence of simultaneous events on the focus and content of the discussions. Next, we eliminated non-English content following the methodologies outlined in Appendix A. Overall, we collect 6,272 podcast videos using youtube-dl (ytdl, 2006). Table 3 in the Appendix displays the top 20 channels by cumulative views for both $D_{youtube}$ and D_{rumble}, including their total number of videos and average view counts in our YouTube dataset. We refer to this dataset as $D_{youtube}$ throughout the remainder of this paper.

Political podcast channels. To compare D_{rumble} and $D_{youtube}$ from a political perspective, we draw on a pre-established classification of YouTube channels into left, center, and right (Dinkov et al., 2019; Boesinger et al., 2024). After applying the same selection and refinement process used in our $D_{youtube}$ extraction, we obtain 7,755 videos across these three ideological categories. Next, we exclude channels that appear in either D_{rumble} or $D_{youtube}$ to prevent the influence of duplicate podcasts in our analyses. This process removes Steven Crowder's channel from our political podcast sample, as it already exists in D_{rumble}. This step is crucial to prevent the influence of identical podcasts from skewing our analyses. Finally, to ensure balanced and comparable analyses, we sample 500 videos per category, matching the monthly distribution patterns in D_{rumble} and $D_{youtube}$. We refer to this dataset as $D_{political}$, with D_{left}, D_{right}, and D_{center} denoting its left, right, and center subsets, respectively.

Speech-to-Text transcription. For the transcription of podcast videos, we use faster-whisper (Klein, 2023), a reimplementation of OpenAI's Whisper (OpenAI, 2022) via CTranslate2 (OpenNMT, 2019), in conjunction with Silero's Voice Activity Detection (Silero, 2021). This combination is particularly effective in handling challenges (e.g., long pauses and background music) present in many videos in our dataset. We use the large-v2 model of Whisper in our analysis and use English as the language parameter. In total, we spend 658 hours (27 days) with NVIDIA A100 GPU with 80GB of Memory to generate their speech-to-text transcriptions.

4 Is there a political bias in the videos of podcast channels on Rumble?

To explore political bias in Rumble podcasts, we perform a quantitative analysis using speech-to-text transcriptions. Initially, we examine political orientations by comparing the popular topics on D_{rumble} and $D_{youtube}$ with those in $D_{political}$. We aim to determine if the discussions align with those typically found on channels known for their political activism or ideological bias, establishing a foundational understanding of the political characteristics inherent in the analyzed content.

Subsequently, we examine centroid cosine similarities across topics using transformer-based sentence embeddings, which allow us to facilitate a deeper inference of potential political alignments or biases present within the discourse. Our analysis extends to channel-based political stances, where we evaluate the political leanings of the podcast videos from channels on D_{rumble} and $D_{youtube}$. This broader perspective helps us understand the diversity of political views on these platforms and whether there is a tendency towards certain political ideologies.

Topic model. We use BERTopic (Grootendorst, 2022), a transformers-based topic modeling technique, in conjunction with MPNetv2 embeddings to extract meaningful topics used by D_{rumble} and $D_{youtube}$, and $D_{political}$. We use this combination because of its ability to discern semantic similarities and differences among documents (Hanley et al., 2023; Yang et al., 2023). In line with prior research (Hanley et al., 2023), we split transcripts into sentences and extract their embeddings using MPNEt-base-v2 model. Our manual inspection of transcripts finds that 2% of all podcast videos in our dataset are missing punctuation. For these specific transcriptions, we split the speech-to-text outputs into sentences using a model (Guhr et al., 2021), which achieves an F1 score of 0.94 for predicting sentence endings in English text.

Postprocessing. We implement three postprocessing steps. To refine our analysis, we remove English stop words from the topic keywords using Scikit-learn's CountVectorizer function (Pedregosa et al., 2011). Next, we exclude topics that comprise fewer than 5 keywords. This decision is based on our observation from manually inspecting the top 100 most popular topics, which indicates that topics with few keywords predominantly consist of generic sentences that are mostly identical (e.g., "Ok."). Subsequently, we filter out topics characterized by conversational fillers and backchannels, e.g., "hmm," "yeah," "oh," "uh," "so," and "well," if these appear among a topic's top five keywords. For this purpose, we use a keyword list derived from previous work (Kim, 2004), which is constructed based on annotated conversational speech data from the Linguistic Data Consortium and standard scoring tools (NIST, 2003). This step is crucial as our primary goal is to enhance the interpretability of our results. Nonetheless, we perform no additional postprocessing due to the intrinsic characteristics of podcast content, which may include casual or mundane discussions. We treat the remaining generic topics as indicative of everyday conversation, providing a richer, more nuanced understanding of our findings.

Examining the political alignment of topics. To identify political bias in D_{rumble} and $D_{youtube}$, we initially assess the extent to which the topics they focus on align with those in $D_{political}$. To

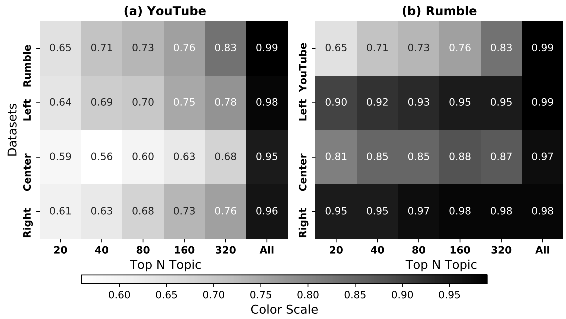

Figure 2: Heatmaps illustrating cosine similarity patterns among the top N topic centroids, comparing platforms (YouTube vs. Rumble) and ideological categories (left-wing, center, right-wing). Darker shades indicate higher centroid cosine similarity.

achieve this, we compare the most popular topics of D_{rumble} and $D_{youtube}$ with those of $D_{political}$. Table 1 presents the top 20 topics of D_{rumble} and $D_{youtube}$. A checkmark indicates if a topic also appears among the top 20 topics in a political podcast sample, where a topic can appear in more than one political leaning.

Among the popular topics of D_{rumble}, 70% align with D_{right}, 50% with D_{left}, and 40% with D_{center}. We find that D_{rumble} focuses primarily on topics heavily discussed in politics or those that can be attributed to political discussions, with a few exceptions (topics #6, #12, #15, #16, and #20), which are related to social media, books, science, and mundane conversations.

In contrast to D_{rumble}, our analysis shows that $D_{youtube}$ has less alignment with political spectrums, aligning 25%, 20%, and 30% with D_{right}, D_{left}, and D_{center}, respectively. This indicates a reduced focus on political subjects overall. Instead, $D_{youtube}$ tends to feature content centered around more apolitical life interests, e.g., sports (topics #2 and #9), sport cars (#17), or music (#15). We also note that, while $D_{youtube}$'s most popular topics are generally more mainstream than political, the presence of topics related to the Russian invasion of Ukraine (#6) and Roe v Wade overturn decision (#12), masks (#14), and vaccines (#10) suggest that popular podcast channels of YouTube can also facilitate discussions around political and social issues.

Centroid cosine similarities with political podcasts. To understand the overall similarity between documents from different groups, previous research (Balci et al., 2023) examined the cosine similarities of the embedding vector centroids.

Building on this, we explore the centroid cosine similarities between D_{rumble} and $D_{youtube}$, as well as their relationship to $D_{political}$. Using MPNet sentence embeddings, our analysis involves performing a layered examination of centroid cosine similarities across varying levels of topic prevalence. Our rationale for this approach is based on an observation made during our earlier analysis, where we noted that a holistic comparison results in high similarity scores, possibly due to occurrences of mundane conversations common in many podcast videos. So, we perform our analysis beginning with the top 20 topics and expanding exponentially in base 2 across five tiers, from 20 to 320 topics. This approach allows us to examine the overall similarity across different tiers of topic frequency, covering nearly 20% of the sentences in D_{rumble} and $D_{youtube}$ after postprocessing (See Figure 4 in the Appendix). We also present the centroid cosine similarities that cover all topics.

To determine centroid cosine similarities between the datasets, we first calculate the centroids of the top N topics for each dataset. The similarity is then assessed using the cosine similarity between these centroids for the top N topics of each dataset. This method provides a nuanced view of the semantic connections between D_{rumble} and $D_{youtube}$ in comparison to $D_{political}$, across multiple strata of topic concentration.

As seen in Figure 2, D_{rumble} exhibits similarity scores of ≥ 0.95 with D_{right} across all ranks. In comparison, the similarity scores are ≥ 0.90 with D_{left}, ≥ 0.81 with D_{center}, and ≥ 0.65 with $D_{youtube}$. These results indicate high centroid cosine similarities with Rumble's podcast videos. However, this high similarity causes D_{rumble}'s relationships with $D_{political}$ to appear more closely aligned than they might actually be. To address this, we normalized Rumble's centroid cosine similarities with the $D_{political}$ datasets. This adjustment helps eliminate the influence of non-political content in computed similarities, providing a more detailed understanding of D_{rumble}'s overall similarity with political content. As a result, we find that D_{rumble} has centroid cosine similarity scores of ≥ 0.75 with D_{right} across all ranks, compared to ≥ 0.47 with D_{left}. Further details are provided in Figure 5 in the Appendix.

When we look at Figure 2, we see considerably lower centroid cosine similarities between $D_{youtube}$ and $D_{political}$. Furthermore, we find that D_{rumble} shows less similarity with $D_{youtube}$ compared to D_{rumble}'s centroid cosine similarities with $D_{political}$. Although these similarity scores increase with the topic size, it is evident that D_{rumble} shows more pronounced centroid cosine similarities with $D_{political}$, particularly with D_{right}, in contrast to $D_{youtube}$.

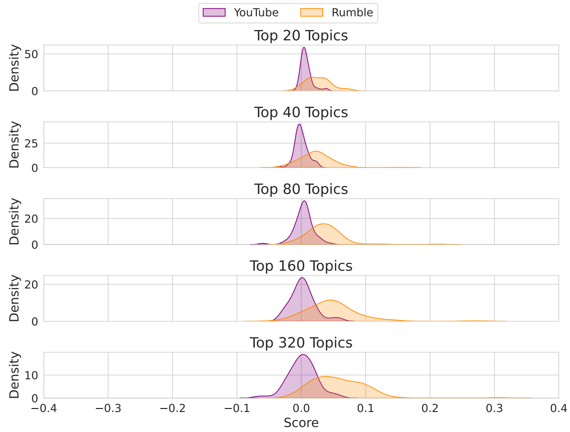

Figure 3: Density plots of political alignment scores for Rumble and YouTube channels. Scores represent ideological orientation and range from -1 to 1, where negative values denote a left-leaning bias and positive values suggest a right-leaning inclination.

Channel-based ideological alignments. Finally, we compare the channel-level similarities between D_{rumble} and $D_{youtube}$ with D_{left} and D_{right}. To measure their similarities, we calculate the percentage of intersecting topics. Specifically, for the podcast videos of each channel on D_{rumble} and $D_{youtube}$, we identify the top 320 topics and evaluate their intersection with the top N topics of D_{left} and D_{right}, where N increments exponentially in base 2 from 20 to 320. This method allows us to assess the breadth of topics covered by the podcast videos of each channel on D_{rumble} and $D_{youtube}$ and how they intersect with the political spectrum at various levels. By exponentially increasing N for the D_{left} and D_{right}, we can measure how their content aligns or diverges from the broader topic set of D_{rumble} and $D_{youtube}$.

To quantify this similarity, we compute the difference in intersection percentages with D_{left} and D_{right} topics:

$$SimScore_{T_i} = \frac{|C \cap R_{T_i}| - |C \cap L_{T_i}|}{|C|}$$

where C represents the set of topics of a given channel, and R_{T_i} and L_{T_i} correspond to the top T_i topics from D_{left} and D_{right}, respectively. For our purposes, C is the set of the top 320 topics, and $T_i = \{20, 40, 80, 160, 320\}$.

Figure 3 plots density distributions of political similarity scores for D_{rumble} and $D_{youtube}$. Complementing this, Figure 6 in the Appendix displays the distribution of left-wing and right-wing similarity scores for the top 320 topics in D_{rumble} and $D_{youtube}$ channels. This scatter plot also includes R-squared and slope values derived from a linear regression analysis, providing further insights into the patterns observed in Figure 3. It is evident that $D_{youtube}$ predominantly clusters around the neutral score (0) across all top N topics, whereas D_{rumble} exhibits a distribution skewed towards the right-wing, indicated by predominantly positive (right-wing leaning) maximum densities. This is another indication of Rumble podcasts' overall right-wing political leaning.

Takeaways. Rumble's popular podcasts lean predominantly towards political topics. Our analysis shows that this political focus is reflected not only in the platform's overall content but also in the individual leanings of specific channels. Their topical focus aligns most with right-wing podcasts, where we also find Rumble has over 0.95 centroid cosine similarity with right-wing podcast content. Moreover, there is a clear inclination towards right-wing content at the channel level. This contrasts with YouTube, where podcasts have a broader focus, covering a wide array of mainstream topics and interests beyond the political sphere. Our results are further supported when we compare the word usages between $D_{youtube}$ and D_{rumble} (detailed in Appendix B), where we find that D_{rumble} aligns with general right-wing narratives on topics related to abortion, elections, and the January 6 Capitol attack.

5 What are the most widely used visual elements? Do they share commonalities with politically motivated podcasts?

Similar to Rumble, the literature on the usage of visual elements in podcasts is also relatively scant. Recognizing this gap, we now focus on the visual topics covered in podcast videos. By examining these visual topics, we aim to have a foundational understanding on how podcasts on Rumble use visual strategies beyond mere auditory content. Based on our previous results, we hypothesize that podcasts on Rumble also use politically motivated visual elements that align with those found in right-wing podcasts. To investigate this, we first extract representative frames from the podcast videos. Subsequently, we apply a clustering technique to the representative images (i.e., visual elements) we identify and analyze the visual clusters that are most frequently used in D_{rumble} and $D_{youtube}$ channels.

Extracting representative video frames. To effectively analyze the visual clusters, our first step is to extract representative video frames. This approach helps us avoid clusters of sequential and almost identical images from the same video. We begin by extracting frames from each podcast video at a rate of one frame per second. Adopting a technique used in previous research (Zannettou et al., 2018), we first apply perceptual hashing (pHash) to each sampled frame. This method extracts representative feature vectors from the images, capturing their visual characteristics. We then measure the similarity between frames by calculating the Hamming distance and set a threshold to identify frames with meaningful visual differences. To establish this threshold, we tested 20 sample videos from both D_{rumble} and $D_{youtube}$. Starting with the second frame, we eliminate frames that fell below a varying threshold θ compared to any of the previous video frames, ranging from $\theta = 5$ to $\theta = 50$ in increments of 5. This evaluation is conducted by three authors of this paper who individually analyze the extracted frames for each sampled video at each θ level, focusing on two metrics: 1) minimizing the number of duplicate images, and 2) maximizing the number of visually distinct images. In the end, the annotators reached a unanimous agreement (Fleiss' Kappa 1.0) on setting the threshold at $\theta = 20$. Figure 7 in the Appendix shows the distribution of representative frames per video for each dataset.

Clustering. We leverage OpenAI's CLIP (Radford et al., 2021) to generate embeddings, using its top performing model, *ViT-L/14@336px*. Our clustering approach is inspired by techniques used in BERTopic (Grootendorst, 2022) and Top2Vec (Angelov, 2020). This methodology first reduces the dimensionality of these embeddings with UMAP (McInnes et al., 2018). Subsequently, we input these reduced-dimension embeddings into HDBSCAN (McInnes et al., 2017), an algorithm that excels in generating dense clusters without the need for predefining cluster sizes. This flexibility allows us to explore thematic topics organically, without the constraint of limiting the visual clusters

YouTube					Rumble				
no.	**Label | (% Channels)**	**Left**	**Center**	**Right**	**no.**	**Label | (% Channels)**	**Left**	**Center**	**Right**
1	Captioned images – (46)	–	–	✓	1	Joe Biden – (34)	–	–	–
2	Guests (Video conference) – (23)	✓	–	✓	2	Jen Psaki – (31)	✓	–	✓
3	Smart Phones – (21)	–	–	–	3	Covid-19 News – (31)	–	–	✓
4	Cartoons – (19)	–	–	–	4	Hillary Clinton – (31)	–	–	–
5	Nostalgic Photos – (19)	–	–	–	5	Ron Desantis – (30)	–	–	✓
6	Basketball Court – (19)	–	–	–	6	Kamala Harris – (29)	–	–	✓
7	Google Image Queries – (18)	–	–	–	7	Guests (Video conference) – (28)	✓	–	✓
8	Typing (keyboard) – (18)	–	–	–	8	Canadian Politics – (27)	–	–	✓
9	Space – (18)	–	–	–	9	Captioned images – (24)	–	–	✓
10	Podcast Studio – (16)	–	–	–	10	Tucker Carlson – (23)	–	–	–
11	Joe Rogan – (16)	–	–	–	11	Joe Biden (w/ mask) – (23)	–	–	–
12	Money – (14)	–	–	–	12	Rand Paul – (22)	–	–	–
13	Typing (smart phone) – (14)	–	–	–	13	Anthony Fauci – (21)	–	–	✓
14	Science – (14)	–	–	–	14	Whoopi Goldberg – (21)	–	–	–
15	Instagram – (14)	–	–	–	15	Karine Jean-Pierre – (20)	–	–	–
16	Fire Images – (14)	–	–	–	16	Joe Biden (News) – (19)	–	–	–
17	Kardashians – (13)	–	–	–	17	Gavin Newsom – (19)	–	–	–
18	Animals – (13)	–	–	–	18	Press conference – (19)	–	–	–
19	Photographers – (13)	–	–	–	19	Joe Rogan– (19)	–	–	–
20	Clocks – (13)	–	–	–	20	Bill gates – (19)	–	–	–

Table 2: Comparison of the Top 20 visual clusters detected through image clustering (manually labelled) on YouTube and Rumble. The presence of a checkmark signifies that the topic appears in the top 20 visual themes of left-wing, center, or right-wing podcasts.

to a specific number.

Finding clusters of widely used visual elements. To determine the most commonly used visual clusters across various channels, we start by identifying the clusters that appear in the highest number of channels for each dataset. Starting from the highest ranked clusters for each dataset, three authors of this paper examine 20 randomly sampled images (or the entire set if a visual cluster comprised $\leq$ 20 images) and labeled the clusters based on the codebook provided in Appendix C. This process is repeated until we have a definitive list of the top 20 visual clusters for each dataset, where we do not include clusters that are primarily composed of frames without meaningful visual content (e.g., black screens or solid colors, including those showing only a channel logo).

Top visual clusters of Rumble and YouTube. Table 2 displays the most frequently used visual clusters across D_{rumble} and $D_{youtube}$, and their alignments with those in $D_{political}$. Figure 8 shows top-10 clusters for each platform. For D_{rumble}, we observe that the most prevalent visual clusters align with our earlier findings, focusing predominantly on political figures. Notably, while the majority of politicians are associated with the left-wing (e.g., Joe Biden, Kamala Harris, and Hillary Clinton), we also see politicians and political commentators that are recognized for their right-wing perspectives (i.e., Tucker Carlson, Ron DeSantis, and Rand Paul). We also observe that the majority of the visual elements of Canadian politics topic are related to Justin Trudeau. We also observe numerous anti-vaccine-related news items in the Covid-19 News topic. Additionally, we encounter a visual topic related to Anthony Fauci, the former Chief Medical Advisor to the President during the COVID-19 pandemic, who has been a target of criticism from right-wing figures, including former President Trump himself (Collins and Liptak, 2020). Interestingly, Bill Gates also appeared among the top 20 visual clusters of Rumble, who has been at the center of COVID-19 related conspiracy theories deployed by the right-wing (McNeil-Willson, 2022). Comparing these findings with the top 20 visual clusters from D_{left}, D_{right}, and D_{center}, we find alignments of 10%, 40%, and 0% respectively. This suggests that Rumble's podcasts exhibit meaningfully more visual commonalities with right-wing podcasts.

Our results from $D_{youtube}$'s most widely used visual clusters also align with our previous findings, as these visuals consist of mostly apolitical and more mainstream themes (e.g., cartoons, basketball court, and Kardashians). When comparing these results to the top 20 most widely used visual clusters in $D_{political}$, we find 5% alignment with D_{left}, 10% with D_{right}, and no alignment (0%) with D_{center}.

Takeaways. Rumble podcasts' visual content is primarily political, with popular visual clusters aligning closely with right-wing podcasts. We observe that these clusters predominantly feature political figures. While these clusters largely showcase left-wing politicians, the political commentators within them are typically associated with right-wing viewpoints. One possible explanation for this could be the dominance of the Democratic Party in the US government during the majority of our dataset's timeline. This may suggest that Rumble's podcasts use visuals of these politicians while critiquing them, stimulating their viewers beyond merely using audio. On YouTube, we consistently find a dominance of apolitical visual clusters, aligning with our prior observations. This contrast further underscores Rumble's non-neutral political stance.

6 Discussion & Conclusion

In this paper, we present the first large-scale data-driven study on podcast videos, where we analyzed the audio-visual content of popular Rumble and YouTube podcast channels, focusing on their political leanings. We present a methodology that can use multimodalilties for understanding video podcast content. Our analysis of over 13K podcast videos demonstrates a right-wing bias in Rumble's content, which sharply contrasts with YouTube's more apolitical content. This dichotomy highlights the role of platforms in either reinforcing or challenging existing political narratives. Our findings suggest that Rumble's video podcast content is predominantly right-wing content, potentially creating a distinct echo chamber effect (Efstratiou et al., 2023). This phenomenon is critical to understand, as it potentially exacerbates societal polarization in a yet underexplored area, e.g., podcasts.

Our findings also emphasize the need to consider both audio and visual elements in media studies. While textual content has been extensively analyzed in social media research, through this work, we emphasize the need to consider both audio and visual content when studying podcast videos, as cues from both modalities can be useful for understanding political leanings. Furthermore, our study makes a valuable contribution to digital humanities by demonstrating how a multimodal, computational data analysis pipeline can deepen our understanding of cultural and political narratives in digital media. By integrating advanced speech-to-text transcription, transformer-based topic modeling, and visual content analysis, our approach bridges computational methods with humanistic inquiry. This methodological innovation not only expands the digital humanities toolkit but also provides a blueprint for exploring how audio, visual, and textual cues collectively shape public discourse and societal ideologies in emerging digital platforms.

6.1 Limitations

This work is subject to certain limitations. First, the data collection was not conducted live, which means some content may have been missed. Furthermore, as we rely on content creators' labeling to create our initial set of podcast videos, we might miss some podcast videos that are not labeled by their creators. Our reliance on tools like faster-whisper, BERTopic, and CLIP, could introduce errors due to their inherent limitations, e.g., Whisper is known for hallucinating content (Mittal et al., 2024; Koenecke et al., 2024) and BERTopic can generate higher number of outliers than expected (Egger and Yu, 2022). These factors should be considered when interpreting our findings.

Our analysis has other limitations. For instance, our labeling of the visual clusters in Rumble and YouTube podcasts was mainly guided by our domain knowledge, yet some channel owners might challenge our categorizations. Another limitation of our study involves assessing how the content of Rumble and YouTube podcasts aligns with political orientations without analyzing the sentiment of this content. While this methodology was in line with our research objectives, it is important to recognize that including sentiment analysis might have offered additional insights into the emotional tone and impact of the podcast content. Finally, our results are based on popular podcast videos from Rumble and YouTube and should not be generalized to video podcasts as a whole.

Ethics statement. Our project, which exclusively uses publicly accessible data and does not involve human subjects, is not classified as human subjects research according to the guidelines of our institution's Institutional Review Board (IRB). We adhere to established ethical standards in social media research and the application of shared measurement data. Additionally, we only use third-party models with publicly available licenses. We do not anonymize people if they are public figures (i.e., podcast channel owners on YouTube or Rumble).

7 Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-2046590.

References

Max Aliapoulios, Emmi Bevensee, Jeremy Blackburn, Barry Bradlyn, Emiliano De Cristofaro, Gianluca Stringhini, and Savvas Zannettou. 2021. A large open dataset from the parler social network. In *ICWSM*, volume 15, pages 943–951.

Waleed Ammar, George Mulcaire, Yulia Tsvetkov, Guillaume Lample, Chris Dyer, and Noah A Smith. 2016. Massively multilingual word embeddings. *arXiv:1602.01925*.

Dimo Angelov. 2020. Top2vec: Distributed representations of topics. *arXiv:2008.09470*.

Utkucan Balci, Chen Ling, Emiliano De Cristofaro, Megan Squire, Gianluca Stringhini, and Jeremy Blackburn. 2023. Beyond fish and bicycles: Exploring the varieties of online women's ideological spaces. In *WebSci*, pages 43–54.

Utkucan Balci, Jay Patel, Berkan Balci, and Jeremy Blackburn. 2024. idrama-rumble-2024: A dataset of podcasts from rumble spanning 2020 to 2022. *Workshop Proceedings of the 18th International AAAI Conference on Web and Social Media.*

Léopaul Boesinger, Manoel Horta Ribeiro, Veniamin Veselovsky, and Robert West. 2024. Tube2vec: Social and semantic embeddings of youtube channels. In *ICWSM*, volume 18, pages 2084–2090.

Tegan R. Bratcher. 2022. Toward a deeper discussion: A survey analysis of podcasts and personalized politics. *AJC*, 30(2):188–199.

Abram Brown. 2022. Is rumble, a right-wing social media company, already the next meme stock? `https://www.forbes.com/sites/abrambrown/2021/12/02/rumble-spac-ipo-social-media-conservative`.

Gizem Nisa Cebi. 2024. Arrest of telegram ceo in france sparks global concerns over free speech.

Monica Chadha, Alex Avila, and Homero Gil de Zúñiga. 2012. Listening In: Building a Profile of Podcast Users and Analyzing Their Political Participation. *J. Inf. Technol. Politics*, 9(4):388–401.

Emily Chen and Emilio Ferrara. 2023. Tweets in time of conflict: A public dataset tracking the twitter discourse on the war between ukraine and russia. In *ICWSM*, volume 17, pages 1006–1013.

Yoon Y. Cho, Ahran Park, and Jinho Choi. 2023. Motives for using news podcasts and political participation intention in South Korea: The mediating effect of political discussion. *Media International Australia*, 187(1):39–56.

Ann Clifton, Sravana Reddy, Yongze Yu, Aasish Pappu, Rezvaneh Rezapour, Hamed Bonab, Maria Eskevich, Gareth Jones, Jussi Karlgren, Ben Carterette, and Rosie Jones. 2020. 100,000 Podcasts: A Spoken English Document Corpus. In *COLING*, pages 5903–5917.

Kaitlan Collins and Kevin Liptak. 2020. Trump trashes fauci and makes baseless coronavirus claims in campaign call. `https://www.cnn.com/2020/10/19/politics/donald-trump-anthony-fauci-coronavirus/index.html`.

Michal Mimino Danilak. 2021. Language detection library ported from Google's language-detection. `https://pypi.org/project/langdetect/`.

Emily DeLetter. 2024. "that's down and dirty": Donald trump asks comedian theo von about cocaine, alcohol use.

Yoan Dinkov, Ahmed Ali, Ivan Koychev, and Preslav Nakov. 2019. Predicting the Leading Political Ideology of YouTube Channels Using Acoustic, Textual, and Metadata Information. *Preprint*, arxiv:1910.08948.

Alexandros Efstratiou, Jeremy Blackburn, Tristan Caulfield, Gianluca Stringhini, Savvas Zannettou, and Emiliano De Cristofaro. 2023. Non-polar opposites: analyzing the relationship between echo chambers and hostile intergroup interactions on reddit. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 17, pages 197–208.

Roman Egger and Joanne Yu. 2022. A topic modeling comparison between lda, nmf, top2vec, and bertopic to demystify twitter posts. *Frontiers in Sociology*, 7.

Rosa Escandon. 2024. "video podcasting" growing in popularity. `https://www.forbes.com/sites/rosaescandon/2024/04/29/video-podcasting-growing-in-popularity`.

Alyn Euritt. 2019. Public circulation in the NPR Politics Podcast. *Popular Communication*, 17(4):348–359.

Hibaq Farah. 2023. What is rumble, the video-sharing platform "immune to cancel culture"? The Guardian.

gop.gov. 2024. It is now a fact that pelosi's sham january 6th committee was designed to be a political witch-hunt. https://www.gop.gov/news/aspx?DocumentID=758.

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv:2203.05794*.

Abel Grunfeld. 2023. Why has video podcasting become increasingly popular? `https://riverside.fm/blog/why-video-podcasting-is-popular`.

Oliver Guhr, Anne-Kathrin Schumann, Frank Bahrmann, and Hans Joachim Böhme. 2021. Fullstop: Multilingual deep models for punctuation prediction.

Hans WA Hanley, Deepak Kumar, and Zakir Durumeric. 2023. Happenstance: Utilizing Semantic Search to Track Russian State Media Narratives about the Russo-Ukrainian War On Reddit. In *ICWSM*, volume 17, pages 327–338.

Jiyoung Kim, Yeon-Ok Lee, and Han Woo Park. 2016a. Delineating the complex use of a political podcast in South Korea by hybrid web indicators: The case of the Nakkomsu Twitter network. *TFSC*, 110:42–50.

Joungbum Kim. 2004. *Automatic detection of sentence boundaries, disfluencies, and conversational fillers in spontaneous speech*. Ph.D. thesis, Citeseer.

Youngju Kim, Yonghwan Kim, and Yuan Wang. 2016b. Selective exposure to podcast and political participation: The mediating role of emotions. *IJMC*, 14(2):133–148.

Miles Klee. 2023. Rumble is down. will russell brand's allegations knock it out? https://www.rollingstone.com/culture/culture-features/russell-brand-allegations-rumble-1234851624/.

Guillaume Klein. 2023. faster-whisper. https://github.com/guillaumekln/faster-whisper.

Allison Koenecke, Anna Seo Gyeong Choi, Katelyn Mei, Hilke Schellmann, and Mona Sloane. 2024. Careless whisper: Speech-to-text hallucination harms. *ArXiv:2402.08021*.

Angela Lai, Megan A Brown, James Bisbee, Joshua A Tucker, Jonathan Nagler, and Richard Bonneau. 2022. Estimating the ideology of political youtube videos. *Political Analysis*, pages 1–16.

Changho Lee. 2021. News Podcast Usage in Promoting Political Participation in Korea. *AJMMC*, 7(2):107–120.

Helen Lewis. 2024. Trump's red-pill podcast tour.

Chen Ling, Ihab AbuHilal, Jeremy Blackburn, Emiliano De Cristofaro, Savvas Zannettou, and Gianluca Stringhini. 2021. Dissecting the meme magic: Understanding indicators of virality in image memes. *CSCW*, 5(CSCW1):1–24.

Robert C. MacDougall. 2011. Podcasting and Political Life. *American Behavioral Scientist*, 55(6):714–732.

Arwa Mahdawi. 2024. "edgelords" and "butt-sniffers": Will trump's tour of hyper-masculine podcasts win over young men? | arwa mahdawi.

Aaron Mak. 2021. Gab is furious that donald trump signed up for another right-wing social network. https://slate.com/technology/2021/06/donald-trump-rally-rumble-gab-parler.html.

John D. McCarthy. 2022. Are antifa and black lives matter related? - the washington post. https://www.washingtonpost.com/politics/2022/02/08/antifa-blm-extremism-violence/.

Megan McCluskey. 2022. Rumble offers joe rogan $100 million to switch platforms. https://time.com/6145835/joe-rogan-rumble-podcast-offer/.

Leland McInnes, John Healy, and Steve Astels. 2017. hdbscan: Hierarchical density based clustering. *J. Open Source Softw.*, 2(11):205.

Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. 2018. Umap: Uniform manifold approximation and projection. *Journal of Open Source Software*, 3(29):861.

Richard McNeil-Willson. 2022. Understanding the# plandemic: Core framings on twitter and what this tells us about countering online far right covid-19 conspiracies. *First Monday*.

Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *NIPS*, 26.

Jeremiah Milbauer, Adarsh Mathew, and James Evans. 2021. Aligning multidimensional worldviews and discovering ideological differences. In *EMNLP*.

Ashish Mittal, Rudra Murthy, Vishwajeet Kumar, and Riyaz Bhat. 2024. Towards understanding and mitigating the hallucinations in nlp and speech. In *11th ACM IKDD CODS and 29th COMAD*, pages 489–492.

NIST. 2003. The rich transcription fall 2003 (rt-03f) evaluation plan. http://www.nist.gov/speech/tests/rt/rt2003/fall/docs/rt03-fall-eval-plan-v9.pdf.

OpenAI. 2022. Whisper. https://github.com/openai/whisper.

OpenNMT. 2019. Ctranslate2. https://github.com/OpenNMT/CTranslate2.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in python. *JMLR*, 12:2825–2830.

Pew. 2006. What is podcasting? https://www.pewresearch.org/journalism/2006/07/19/what-is-podcasting/.

Naga Pramod. 2021. Rumble is experiencing massive growth as people ditch big tech. https://reclaimthenet.org/rumble-is-experiencing-massive-growth.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR.

Maria Rae. 2023. Podcasts and political listening: Sound, voice and intimacy in the Joe Rogan Experience. *Continuum*, 37(2):182–193.

Naquee Rizwan, Nayandeep Deb, Sarthak Roy, Vishwajeet Singh Solanki, Kiran Garimella, and Animesh Mukherjee. 2025. Dynamics of toxicity in political podcasts. *arXiv preprint arXiv:2501.12640*.

Jamie Robertson. 2019. How podcasts went from unlistenable to unmissable. https://www.bbc.com/news/business-49279177.

Katherine Rosman, Ben Sisario, Mike Isaac, and Adam Satariano. 2022. Spotify bet big on joe rogan. it got more than it counted on. https://www.nytimes.com/2022/02/17/arts/music/spotify-joe-rogan-misinformation.html.

Rumble. 2023. Podcasts. https://web.archive.org/web/20230131060622/https://rumble.com/category/podcasts.

Frans Sayogie, Muhammad Farkhan, Hendrio Putra Julian, Hilman Syauqiy Fauza Al Hakim, Muhammad Guntur Wiralaksana, et al. 2023. Patriarchal ideology, andrew tate, and rumble's podcasts. *3L: Southeast Asian Journal of English Language Studies*, 29(2).

Elisa Shearer, Jacob Liedke, Katerina Eva Matsa, Michael Lipka, and Mark Jurkowitz. 2023. Podcasts as a source of news and information. *Pew Research Center*, 16.

Carter Sherman. 2024. Many republicans support abortion. are they switching parties because of it?

Silero. 2021. Silero vad: pre-trained enterprise-grade voice activity detector (vad), number detector and language classifier. https://github.com/snakers4/silero-vad.

Jonathan Sterne, Jeremy Morris, Michael Brendan Baker, and Ariana Moscote Freire. 2008. The Politics of Podcasting. *Fibreculture*.

Galen Stocking, Amy Mitchell, Katerina Eva Matsa, Regina Widjaya, Mark Jurkowitz, Shreenita Ghosh, Aaron Smith, Sarah Naseer, and Christopher St Aubin. 2022. The role of alternative social media in the news and information environment. *Pew Research Center*.

Josh Wilson. 2022. The downfall of andrew tate and its implications.

Jonghyeon Yang, Hanme Jang, and Kiyun Yu. 2023. Analyzing geographic questions using embedding-based topic modeling. *ISPRS International Journal of Geo-Information*, 12(2):52.

YouTube. 2023. youtube-popularcreators. https://www.youtube.com/podcasts/popularcreators.

ytdl. 2006. youtube-dl. https://github.com/ytdl-org/youtube-dl.

Savvas Zannettou, Tristan Caulfield, Jeremy Blackburn, Emiliano De Cristofaro, Michael Sirivianos, Gianluca Stringhini, and Guillermo Suarez-Tangil. 2018. On the origins of memes by means of fringe web communities. In *IMC*, pages 188–202.

A

Language verification for podcasts. In addition to our initial step of excluding non-English channels and playlists, following previous work (Clifton et al., 2020), we run language detection on podcast video descriptions. For this purpose, we use langdetect library (Danilak, 2021), which is a Python implementation of Google's languagedetection library in Java. We also remove URLs from video descriptions before running language detection. During a manual inspection of videos flagged as non-English, we observe that these videos have short descriptions (e.g., social media platforms and their URLs) that could cause mislabeling their languages. Consequently, we conduct a manual inspection of these videos and videos with no description, and exclude "Monarky" channel from Rumble, due to its content being in a language other than English.

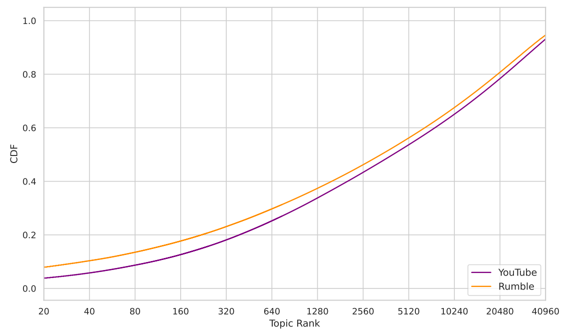

Figure 4: CDF of the proportion of sentences covered cumulatively at each topic rank in YouTube and Rumble podcast videos. Topic ranks start at 20 and increase exponentially.

B

Misalignment analysis. To further solidify our findings for RQ1, we analyze the differences in word usage between D_{rumble} and $D_{youtube}$. To do this, we leverage the methodology proposed by Milbauer et al. (Milbauer et al., 2021), which trains word2vec models for each community, and aligns their words using a linear translation function MultiCCA (Ammar et al., 2016). If a community's word projection does not match the same word in

YouTube				Rumble			
Channel	# Views	# Podcasts	Avg. Views	Channel	# Views	# Podcasts	Avg. Views
H3 Podcast	183M	108	1.7M	The Dan Bongino Show	133M	576	231K
Philip DeFranco	143M	156	918K	Steven Crowder	42M	212	198K
rSlash	111M	223	500K	The Post Millennial	13M	10	1.3M
No Jumper	107M	465	231K	RepMattGaetz	9.7M	45	216K
Bailey Sarian	10M	34	3M	TateSpeech by Andrew Tate	7.8M	3	2.6M
IMPAULSIVE	89M	39	2M	The JD Rucker Show	7.3M	38	194K
REVOLT	88M	61	1.4M	The Charlie Kirk Show	5.7M	215	26K
YMH Studios	77M	130	598K	The Rubin Report	5.0M	174	28K
Gecko's Garage - Trucks For Children	70M	42	1.6M	Glenn Greenwald	4.8M	24	201K
FLAGRANT	67M	51	1.3M	HodgeTwins	4.6M	152	30K
Dr. Sten Ekberg	64M	50	1.3M	Senator Ron Johnson	4.5M	1	4.5M
Lex Fridman	64M	59	1M	Devin Nunes	4.2M	64	66K
The 85 South Comedy Show	63M	45	1.4M	vivafrei	4.2M	178	23K
NBC News	61M	313	196K	Dinesh D'Souza	4.1M	208	20K
The Pat McAfee Show	58M	161	365K	Russell Brand	4.0M	48	83K
FreshandFit	55M	226	246K	TheSaltyCracker	3.8M	62	62K
Critical Role	51M	26	1.9M	Ben Shapiro	3.2M	297	10K
CinnamonToastKen	47M	41	1.1M	TimcastIRL	3.1M	326	9K
Jordan B Peterson	47M	43	1M	The Trish Regan Show	2.9M	190	15K
48 Hours	46M	10	4.6M	Joe Pags	2.4M	134	18K

Table 3: Top 20 podcast video channels of YouTube and Rumble, by their cumulative views, total number of videos, and average views.

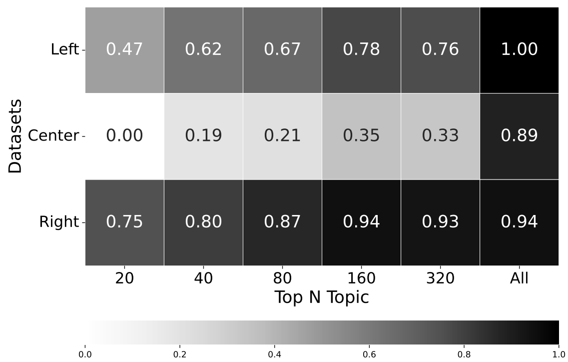

Figure 5: Heatmap illustrating the normalized cosine similarities among the top N topic centroids of Rumble versus left-wing, center, and right-wing podcasts. Darker shades denote greater centroid cosine similarity.

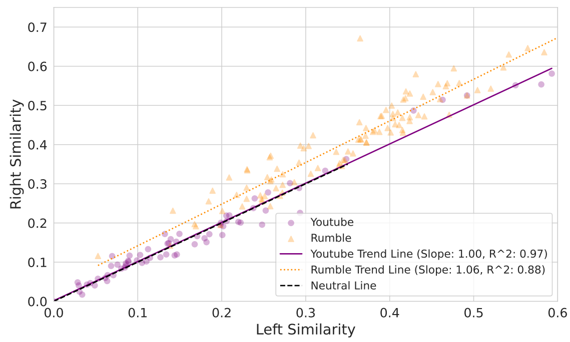

Figure 6: Scatter plot showing right-wing and left-wing similarity distributions for the top 320 topics in podcast videos of popular YouTube and Rumble podcast channels, with R-squared and slope values from linear regression.

Rumble	YouTube	Alignment
Republicans	Democrats	0.8787
Democrat	Republican	0.7717
Dems	Democrats	0.6986
Leftists	Right-wingers	0.6231
Hillary Clintons	Trumps	0.5761
Pro-choice	Pro-life	0.5560
Progressive	Conservative	0.5190
Pro-Trump	Anti-Trump	0.4732
Witch Hunt	January 6th	0.4571

Table 4: Identified misaligning word pairs between popular podcast channels of YouTube and Rumble.

another community, we consider these words are *misaligned*. This way, by identifying misaligned word pairs with political meanings, e.g., Democrat's usage of "Republican" and Republican's usage of "Democrat, " we can have an understanding of a community's political positioning.

Training. We follow the preprocessing steps proposed by Milbauer et al., where we tokenize each sentence, remove hyperlinks, and lowercase all characters. Next, we train Word2Vec skip-gram models (Mikolov et al., 2013) for $D_{youtube}$ and D_{rumble} using 100 dimensions and a maximum vocabulary of 30,000 words. We anchor the top 5K common words of these datasets and translate them using MultiCCA.

Results. Table 4 presents identified misaligning word pairs between $D_{youtube}$ and D_{rumble}, along with their cosine similarities. Similar to our previous example, we find many misaligning word pairs in the context of "Democrats vs Republicans." This

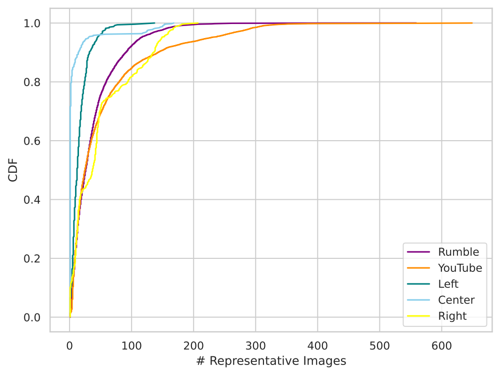

Figure 7: CDF of the representative frames for each podcast video for each dataset. We see center leaning podcasts use less variety of visual elements compared to other datasets.

is evident from Republicans & Democrats, Democrat & Republican, Dems & Democrats, Leftists & Right-wingers, Hillary Clintons & Trumps, Progressive & Conservative, and Pro-Trump & Anti-Trump word pairs.

Additionally, we identify Pro-choice & Pro-life and Witch Hunt & January 6th pairs, which further indicate that D_{rumble} aligns with general right-wing narratives on these topics (McCarthy, 2022; gop.gov, 2024; Sherman, 2024). Overall, these results further solidify our findings from RQ1, demonstrating that Rumble's podcast content exhibits a pronounced right-wing bias, a trend that remains evident even when compared to YouTube's predominantly apolitical content.

C

C.1 Codebook for Visual Element Clustering and Labeling

The goal of this codebook is to provide a systematic approach for labeling the top 20 clusters of visual elements identified in this study. The process involves collaboration among three researchers and, when necessary, external validation through online resources.

Cluster Elimination Criteria. Clusters were excluded from labeling if they lacked meaningful content. A cluster was considered meaningful if it contained distinguishable and recognizable visual elements. This was determined by the consensus of the three researchers.

Labeling Process. Our codebook involves three different cases for the labeling process.

C.1.1 Consensus Labeling.

A cluster is labeled when all three researchers reach an agreement on the appropriate label.

1. Each researcher independently analyzes the cluster and proposes a label.
2. The label is finalized if all researchers agree.

C.1.2 Partial Agreement.

If at least one researcher is unable to label the cluster, but the remaining researchers agree on a label, further validation is sought through online resources.

1. Perform a Google search query based on the proposed label.
2. Check for a corresponding Wikipedia page or other reputable sources.
3. If validation is confirmed, the proposed label is accepted.

C.1.3 No Initial Agreement.

If none of the researchers can label a cluster, external validation is sought through investigating the source videos of the visual elements.

1. Investigate source podcast videos to gather more information about the visual elements in the cluster.
2. Based on the findings from the investigation, conduct a Google search query.
3. Validate the information with a Wikipedia page or another reputable source, if applicable.
4. Assign a label based on the validated information.

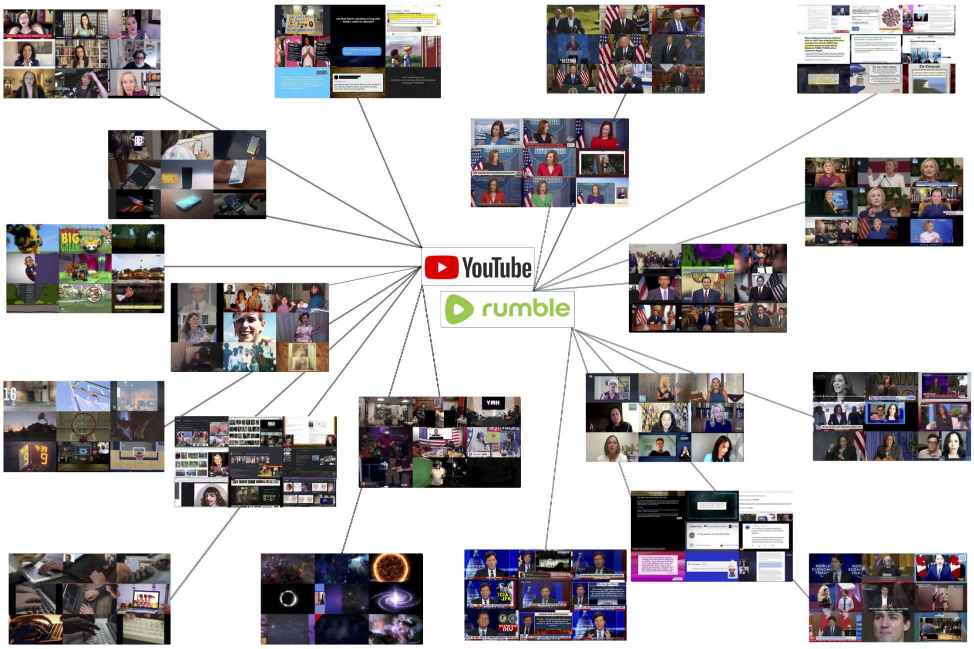

Figure 8: Comparison of visual topics between Youtube and Rumble, extracted through clustering, showing top-10 clusters for each platform (Refer to Table 2).