



OPEN

DATA DESCRIPTOR

# Using AI to Summarize US Presidential Campaign TV Advertisement Videos, 1952–2012

Adam Breuer<sup>1</sup>✉, Bryce J. Dietrich<sup>2</sup>✉, Michael H. Crespin<sup>3,4</sup>, Matthew Butler<sup>5</sup>, JA Pryse<sup>3</sup> & Kosuke Imai<sup>6</sup>

This paper introduces the largest and most comprehensive dataset of US presidential campaign television advertisements, available in digital format. The dataset also includes machine-searchable transcripts and high-quality summaries designed to facilitate a variety of academic research. To date, there has been great interest in collecting and analyzing US presidential campaign advertisements, but the need for manual procurement and annotation has led many to rely on smaller subsets. We design a large-scale, parallelized, AI-based analysis pipeline that automates the laborious process of preparing, transcribing, storyboarding, and summarizing videos. We then apply this methodology to the 9,707 presidential ads from the Julian P. Kanter Political Commercial Archive. We conduct extensive human evaluations to show that these transcripts and summaries match the quality of manually generated alternatives. We illustrate the value of this data by including an application that tracks the genesis and evolution of current focal issue areas over seven decades of presidential elections. Our analysis pipeline and codebase also show how to use LLM-based tools to obtain high-quality summaries for other video datasets.

## Background & Summary

Studies of political campaign advertisements help us understand what messages political elites send to voters and how opposing campaigns strategically interact. Campaign ads also play a critical role in identifying salient issues that matter to the electorate.

In the United States, television ads are the primary means through which campaigns reach voters. In the 2024 US election alone, political campaigns aired ads on broadcast television more than 1.47 million times at a cost of over 5.3 billion dollars between Super Tuesday and November 1st<sup>1,2</sup>. Despite the recent trend towards online advertising, broadcast television advertising continues to increase in real amounts year-over-year. Television also continues to dwarf alternatives, accounting for 48% of ad spend, more than triple the value of online and social media advertising campaigns<sup>2,3</sup>. Therefore, it is no surprise that TV advertisements for the US presidential campaign have been a focal point of considerable research in political communication<sup>4,5</sup>, political psychology<sup>6–8</sup> and political behavior<sup>9–12</sup>.

Despite this scholarly interest, two challenges remain that hinder further development of studies on US presidential campaign TV ads. First, the literature on ads has repeatedly emphasized the importance of analyzing how the content and style of ads have evolved over time [e.g.,<sup>13,14</sup>]. However, such analysis is largely precluded by the fact that ad data are limited or unavailable for elections prior to 1996. Second, even when ads from recent elections are available, extracting relevant information from the videos for downstream analyses has required costly human hand-coding, limiting the scope of analyses to a specific set of variables (see the Wesleyan Media Project<sup>15</sup> and its predecessor Wisconsin Advertising Project<sup>16,17</sup>). In part due to these limitations, scholars have focused primarily on analyzing small subsets of US presidential campaign TV ads [e.g.,<sup>18</sup>].

<sup>1</sup>Dartmouth College, Department of Computer Science and Department of Government, Hanover, NH, USA.

<sup>2</sup>Purdue University, Department of Political Science, West Lafayette, IN, USA. <sup>3</sup>University of Oklahoma, Carl Albert Congressional Research and Studies Center, Norman, OK, USA. <sup>4</sup>University of Oklahoma, Department of Political Science, Norman, OK, USA. <sup>5</sup>University of Iowa, Digital Scholarship & Publishing Studio, Iowa City, IA, USA. <sup>6</sup>Harvard University, Department of Government and Department of Statistics, Cambridge, MA, USA. ✉e-mail: [adam.breuer@dartmouth.edu](mailto:adam.breuer@dartmouth.edu); [bjdietri@purdue.edu](mailto:bjdietri@purdue.edu)

Ads	9,707
Elections	1952–2012
Democratic Ads	5,859
Republican Ads	3,848
Total Hours of Video	102.95
Format	mp4: 720 × 480

**Table 1.** Summary of the Presidential Ad dataset.

**Overview of the dataset: Ad videos, transcripts, and summaries, 1952–2012.** We address these problems by publicly releasing the largest collection of US presidential campaign TV advertisements to date from the Julian P. Kanter Political Commercial Archive<sup>19</sup>. There are 9,707 ads for presidential candidates starting with the 1952 election where Dwight D. Eisenhower faced Adlai Stevenson and ending with the 2012 Obama v. Romney election. Although the ads primarily feature major candidates, the data also include commercials from lesser-known candidates who did not garner the nomination, permitting analysis of campaign messaging on either side of primaries and major party nominations.

To facilitate research on this dataset, we develop a large-scale, parallelized AI-based analysis pipeline that automates the laborious process of preparing, transcribing, and summarizing videos. The proposed methodology is designed to facilitate the neutrality, comprehensiveness, and coherence that are desirable for downstream academic research. We then apply this analysis pipeline to the ad dataset and release each ad with an accompanying high-quality transcript and an efficient 50-word summary. In addition to these transcripts and summaries, we also provide extensive metadata and an online search engine so that researchers can easily access the ads, determine the subsets of ads that are relevant to their research topic or time period, and use them in their downstream analyses.

At the time of writing, LLMs that are capable of summarizing video files are not yet available. Therefore, our analysis pipeline itself may be of independent interest, as it permits researchers to efficiently generate neutral and comprehensive summaries for each video in a large video dataset. The proposed methodology improves upon a recent study, which shows that automated coding using machine learning algorithms performs as well as human coding<sup>20</sup>. Compared to that study, our transcripts are of higher quality. Moreover, instead of just coded variables, we are also able to provide a natural language summary of each ad’s audio *and* visual content that is designed to facilitate efficient data exploration.

We validate each step of our pipeline used to generate quality AI-generated summaries. Specifically, we validate our video preprocessing step, AI video transcription step, and AI video summarization step by comparing stratified samples of these outputs with human-generated outputs. Our main validation result is that human specialists find the quality of our LLM-generated ad transcripts and summaries to be at least equal to, and often superior to, the quality of transcripts and summaries composed by human research assistants across multiple key dimensions. Importantly, we also show that the high quality of our transcripts and summaries remains consistent across key sources of variation in the literature, including political party and time. This result and several others from the validation exercises outlined below will give applied researchers a great deal of confidence when they use both the raw ads and the associated summaries and transcripts for their own research.

Finally, we also substantively validate and illustrate our dataset by considering an application that measures the historical evolution of current key electoral issue areas, as viewed through a keyword-assisted topic model estimated on the transcripts of all ads in the dataset. We show that the seven decades of data exhibit trends in campaign issues that correspond to historically recognized developments in campaign issue agendas.

**History of the Kanter campaign advertisement video collection.** A former staffer for Adlai Stevenson, Julian P. Kanter started the Kanter political TV ad collection in 1956 when he realized that political commercials were being discarded, or later taped over, after airing. He obtained the early film ads from radio and television studios, ad agencies, and the campaigns themselves. He later traded blank tapes in exchange for the already aired commercials. The collection moved to the University of Oklahoma in 1986 when it outbid the Smithsonian to purchase approximately 25,000 radio and television ads from Julian Kanter for \$250,000<sup>21</sup>, including approximately 21,000 television ads ranging from 1950 to 1984<sup>22</sup>. The archive has increased in size as Kanter and others continued collecting ads and political advertising firms donated their film and video. At the University of Oklahoma, the archive was originally housed in the Political Communications Center and is now part of the Carl Albert Congressional Research and Studies Center’s archival collections<sup>23</sup>. Unfortunately, there are no presidential ads available in the dataset beyond the 2012 election. This is largely due to a resource constraint within the archive itself that limits active collecting and the reluctance of political consultants to donate commercials for current or potential future active candidates. We hope to add these recent ads to the dataset in the future and process them by applying the same techniques described in this paper. In the meantime, the missing ads limit what scholars are able to say about the last three elections in 2016, 2020, and 2024.

**The US presidential advertisement video dataset.** Although the collection has grown to more than 100,000 commercials from campaigns up and down the ticket, we focus on the 9,707 ads for the US presidency from 1952–2012. Table 1 reports an overall summary of the ad collection. We do not make any claim that the collection contains the *complete* universe of ads (and it is not known which ads are missing), or that it is a

representative sample of all ads. However, it is the most complete collection available in such a long time frame. Indeed, for a substantial number of items in the collection, the archive has the only known copy anywhere. To emphasize this point, we asked two undergraduate research assistants to check whether the ads in our dataset also appear in *The Living Room Candidate*<sup>24</sup>, which is the largest known collection of presidential ads (365 ads matching our 1952–2012 time period) available online. Using a random sample of size  $N = 500$  from our collection, they found only 10.6 percent of our ads were available on that website.

Over the years, the collection has been used or cited not only in political science research<sup>10,25–29</sup>, but also in *amicus curiae* briefs<sup>30</sup>, documentaries<sup>31,32</sup>, podcasts<sup>33</sup>, and in many academic disciplines, including advertising<sup>34</sup>, area studies<sup>35</sup>, cinematography<sup>36</sup>, communication<sup>7,37–40</sup>, history<sup>41,42</sup>, and library and information science<sup>22,43</sup>. Nevertheless, access to the data has been significantly limited by the fact that until recently, anyone wishing to examine a substantial number of Kanter ads had to visit the archives in person. Our goal is to remove this barrier and make the dataset both widely available to scholars and also amenable to efficient research curation via the release of accompanying transcripts, summaries, and an online interface that permits keyword searches that return the corresponding relevant subset of ads.

**Transcripts.** In addition to the ads themselves, we also release a high-quality transcript for each ad. These transcripts serve three purposes. First, they permit practitioners to search the ad dataset for keywords of interest in order to identify ads that are relevant to their specific research topic. We provide online text search functionality on our project website <https://campaignTVads.org/> (see *Usage Notes* below). Second, transcripts allow practitioners to study the dataset using novel methods from the recent renaissance in text-based research methods for social science [see e.g.,<sup>44–47</sup>]. Finally, transcripts are also a key ingredient of our ad summarization workflow, as we describe below.

**50-word summaries and visually representative ‘storyboard’ still frames.** Historically, most research on political ads has focused on small subsets of ads, in no small part due to the high manual labor costs associated with viewing ads. Specifically, while each individual ad is relatively brief, our dataset contains more than 100 hours of ads in total, which can prevent practitioners from conducting a thorough data exploration. Transcripts are helpful in this regard, but reading transcripts requires nearly as much time as watching videos and also ignores the rich visual information in the ads.

To address this problem, we take inspiration from a rich scholarly tradition of ‘campaign ad storyboards’—time-synced transcript segments associated with a visually representative selection of still frames extracted from each campaign ad (see, e.g.,<sup>48,49</sup>). Specifically, to supplement the ad transcripts, we develop a novel method to automatically extract a set of still frames that comprehensively represent the visual contents of each ad. We include these ‘storyboard’ still frames alongside each ad and transcript in our dataset. We then use the transcripts and these visually representative still frames to generate a high-quality 50-word summary of each ad designed to significantly accelerate the process of manually parsing the dataset. These summaries offer an efficient synopsis of *both* audio and visual content, and each summary can be parsed by a human in approximately 4 seconds.

**Metadata.** Finally, we also provide additional metadata about each ad. Our files include the candidate name, political party, election, ad title (if known), original video technology format (e.g., 16mm film), and ad duration.

## Methods


In this section, we discuss the proposed automated procedure that takes a historical campaign TV ad and produces a complete transcript, a visually representative selection of still frames, and an efficient, high-quality 50-word AI-generated summary designed to facilitate downstream research. We begin by describing how we preprocessed video files, as the original raw files contained extra segments at the beginning and end that were not part of the ad. We then show how to generate ad summaries using LLMs.

Importantly, our summarization pipeline serves as readily-adaptable model. By providing the pipeline and data for all intermediate outputs, such as transcripts and key frames, we enable other researchers to generate their own custom-directed summaries or ad labels merely by adjusting the provided prompts and recomputing the final step of the pipeline, bypassing the significant cost and complexity of preprocessing, transcription, and frame extraction steps and their associated human validation workflows.

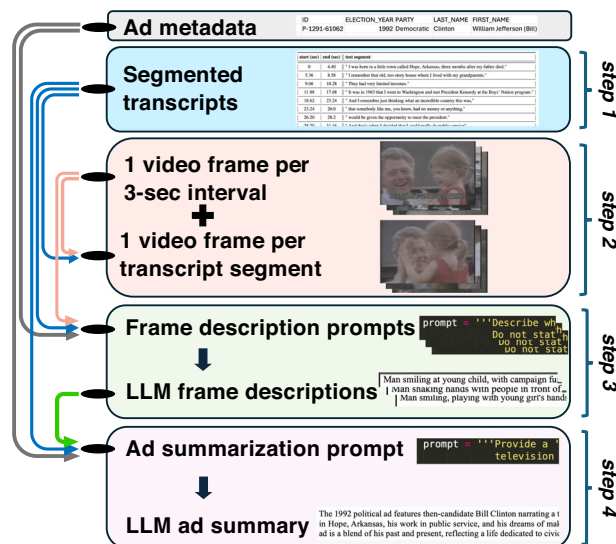
We illustrate our approach throughout this section by considering the running example of Bill Clinton’s famous 1992 ‘Hope, Arkansas’ ad (ID P-1291-61062). This ad (<https://vimeo.com/992368244>) is a challenging example as it contains a diverse set of themes and images.

**Preprocessing videos to extract political ads.** Recordings of historical television ads typically contain ‘pre-roll’ and/or ‘post-roll’ clips that include, for example, the ad metadata, the name of those who recorded the ad, and the television station. It is important to trim these segments so that they do not hinder efficient viewing, and so their content is not included in the downstream ad summary we generate. However, automatic detection of transition points between the pre/post-roll clips and an ad itself is a challenging task because these segments vary widely in terms of length, format, and content.

We develop an automated trimming algorithm based on the following key observation: despite their diversity, nearly all political ads begin and end with a scene that contains video footage of people and/or audio of speech, such as voiceover or a fading in (or out) scene depicting a candidate or interviewer speaking. We leverage this observation by performing, on each video, (1) automated video detection of the timestamps of spoken words, (2) automated video person detection, and (3) automated video scene detection, using the Google Video Intelligence API<sup>50</sup>. As an aside, we note that researchers who wish to apply our methodology to their own



**Fig. 1** Preprocessing example: Time stamps  $w$  and  $w'$  are the start of the first spoken word and the end of the last word spoken, respectively. Points  $p$  and  $p'$  are the start of the first person and end of the last person depicted. Lastly,  $s$  and  $s'$  are the start of the first detected scene and the end of the last detected scene that contain a word or person, respectively. We keep the segment from  $s$  to  $s'$  as the advertisement, and we trim on either side.



**Fig. 2** Ad video summarization workflow. Dependencies between intermediate outputs are indicated by arrows on the left.

datasets may skip using Google Video Intelligence API for spoken word timestamp detection and simply use the Whisper beginning/ending timestamps instead, as they are computed in **Step 1** below anyway. We then trim the beginning and end segments of the video file, yielding the video clip that spans from the beginning of the earliest scene that contains a person or spoken word to the end of the last scene that contains a person or spoken word.

Figure 1 illustrates this process via an example. Here,  $w$  represents the start of the first spoken word, and  $p$  indicates the moment the first detected person appears onscreen. We seek  $s$ , which is the beginning of the first scene that contains  $w$  or  $p$ . Similarly,  $w'$  and  $p'$  represent the end of the last spoken word and the time the last detected person leaves the screen. Finally,  $s'$  is the endpoint of the last scene that contains a word or person, i.e., the last scene that contains  $w'$  or  $p'$ . In this case, we take  $[s, s']$  as the ad segment, which contains all of these six points.

**Generating ad summaries that limit mismeasurement-by-omission.** Our video summarization workflow is designed to overcome two primary challenges. First, at the time of writing, mainstream LLM tools accept audio and image inputs, but they *lack the capability to accept video inputs*. Second, even if video-processing LLM tools were available, our use case does not directly conform to popular video description tasks' objectives. Specifically, because our video summaries may later be used for downstream statistical analysis, our primary objective is not to generate summaries that are articulate or interesting. Instead, our main goal is to ensure descriptive completeness and minimize potential sources of bias and mismeasurement. In particular, we do not want to omit important descriptive information.

We design a summarization workflow that both (i) accepts video inputs and (ii) limits bias and omission errors that might affect downstream analyses. As shown in Fig. 2, we address these two challenges by dividing the summarization of each video into three independent tasks designed to favor descriptive completeness at the expense of redundancy: (**Step 1**) transcription, (**Step 2**) video segmentation, and (**Step 3**) frame description. We then merge these many (possibly redundant) summaries into a single comprehensive ad summary (**Step 4**). We now describe each step in detail.

start (sec)	end (sec)	text segment:
0	4.40	"I was born in a little town called Hope, Arkansas, three months after my father died."
5.36	8.58	"I remember that old, too-story ( <i>sic</i> ) house where I lived with my grandparents."
9.06	10.28	"They had very limited incomes."
11.88	17.68	"It was in 1963 that I went to Washington and met President Kennedy at the Boys' Nation program."
18.62	23.24	"And I remember just thinking what an incredible country this was,"
23.24	26.0	"that somebody like me, you know, had no money or anything,"
26.20	28.2	"would be given the opportunity to meet the president."
28.20	31.16	"And that's when I decided that I could really do public service"
31.16	32.64	"because I cared so much about people."
33.18	36.56	"I worked my way through law school with part-time jobs, anything I could find."
37.4	40.46	"And after I graduated, I really didn't care about making a lot of money."
40.54	42.68	"I just wanted to go home and see if I could make a difference."
43.72	48.22	"We've worked hard in education and health care to create jobs,"
48.36	50.16	"and we've made real progress."
50.44	53.06	"Now it's exhilarating to me to think that as president,"
53.2	55.76	"I could help to change all our people's lives for the better"
55.76	58.16	"and bring hope back to the American dream."

**Table 2.** Example transcript and automatically detected text segments (and their corresponding start and end timestamps) extracted from Clinton’s ca. 1992 “Hope, Arkansas” ad (ad ID P-1291-61062).

*Step 1: Transcription and transcript segmentation.* After video preprocessing, we perform automated transcription and transcript segmentation using the largest and most performant version of OpenAI’s LLM-based Whisper transcription model: WHISPER-LARGE-V3<sup>51</sup>. This model improves upon a previous version (V2) that already obtained transcription accuracy within a fraction of a percent of professional human transcribers across a variety of text benchmarks in terms of word error rates<sup>51</sup>.

The Whisper model is open source, which permits us to parallel-process many videos simultaneously on the local resources of a cloud-based server computer. Specifically, a standard Google Cloud Compute virtual machine with 48 cores transcribes our approximately 10,000 videos in a couple of hours by running parallel batches of 48 videos at a time.

Importantly, transcription with Whisper returns two outputs: the video transcripts and the segmentation of the transcript into short phrases that are separated by a natural pause in the speaker’s voice or a change in the person speaking. Table 2 presents an example of the transcription output. Each text segment comes with the start and end points recorded in seconds, which we later combine with other information extracted from the video to determine the trimming start and end times.

*Step 2: Video key frame extraction.* Our second task is to extract an ordered set of key frames (i.e., images of video frames) from each video that are *comprehensive* in terms of reflecting each video’s visual contents and progression of imagery. More specifically, we seek a time-ordered set of frames that contains all of the important visual content of the ad.


Recent work has accomplished this by randomly sampling one frame for every 15 seconds of video, as well as one frame from the first and last two seconds of each ad<sup>52</sup>. However, our objective differs: rather than merely collecting a representative subset of frames, we aim to extract a frame sequence that comprehensively covers all visually relevant content in the ad.

We accomplish this through two complementary strategies. First, we determine when important images are shown according to the ad’s narrator or speaker via the transcription segments detected in **Step 1** above. Specifically, we observe that the visual contents of a particular ad are almost always closely time-synced to its transcript. For example, the vast majority of scene changes correspond to transitions between speech segments, which we detected in **Step 1** above. Moreover, most visually important imagery is onscreen around the middle of a narrator’s (or speaker’s) speech segment. Therefore, for each video, we extract a frame corresponding to the central moment of each speech segment. In the example shown in Table 2, we extract the frame that is depicted at the timestamp that falls in the middle of the phrase “I was born in ...” (i.e. the frame at  $(4.4 - 0)/2 = 2.2$  sec), another frame from the middle of the phrase “I remember that ...” (i.e. the frame at 6.97 sec), and so on.

However, it is possible that relying only on speech segments to detect important images may miss some important frames. This is because a small number of ads also avoid speech during key segments for dramatic effect or avoid spoken words entirely. For example, Clinton ran several ads that depict only text set to an upbeat jazz backing track. We account for these rare but important stylistic choices by extracting a ‘transcription-agnostic’ set of video frames at regular 3-second intervals. We then merge this set with the transcription-directed video frames described above.

Figure 3 shows the distribution of counts of key frames extracted from the ads. The ads have a mean of 9.78 text segment-centered frames plus 12.24 regular-intervaled frames (~22 frames total). Note that in this figure, many ads have approximately 10 or 20 regular interval frames, as this corresponds to standard ad spot lengths of 30 or 60 seconds.





**Fig. 3** Distributions of counts of frames extracted from ads in **Step 2**.

**Step 3: Obtaining video key frame descriptions.** We generate a brief summary of each video key frame by submitting it to a multimodal LLM: GPT-4-VISION, accompanied by a custom prompt that provides context using the ad metadata. Specifically, we submit the following prompt for each key frame:

*Describe what is depicted in this video frame in no more than 15 words. Do not state that the frame depicts a vintage advertisement, and do not comment on the image quality. If the image includes text, then state that it includes text and also include a summary of the text that is shown. For context, this video frame is a still taken from an advertisement for the {ELECTION YEAR} presidential campaign of {PARTY} {CANDIDATE}. The transcript of the entire ad is: {TRANSCRIPT}*

In some cases, attack ads avoid advocating for a specific candidate, and instead focus entirely on criticizing an opponent. For these cases, we replace the phrase “...an advertisement for the {ELECTION YEAR} presidential campaign of {PARTY} {CANDIDATE}...” with the phrase “an advertisement for the ...{ELECTION YEAR} presidential election. This ad is anti-{CANDIDATE} and pro-{PARTY}...”

**Step 4: Video summarization.** We now describe how we obtain LLM-generated summaries of the ads. Specifically, we generate each ad summary by querying GPT using a customized summarization prompt that provides (1) the ad transcript generated in **Step 1**; (2) the time-ordered set of approximately 20 descriptions of video imagery generated in **Step 3** above including the merged speech-segment-based key frames and regular interval frames (merged by ascending timestamps); and (3) a contextual description of the ad that contains relevant ad metadata. We consider the following prompt, which we submit to the GPT-4 API:

*Provide a 50 word summary of a political television ad for the academic community. Your summary should not exceed 50 words. For context, this ad was for the {ELECTION YEAR} presidential campaign of {PARTY} candidate {CANDIDATE}. The transcript of the entire ad is:*

*{TRANSCRIPT}*

*The ad video depicts a set of scenes that can be described as follows:*

*{First ad frame description}*

*{Second ad frame description}*

*...*

We note that our prompt repeats its emphasis on the 50-word length limit. This is by design: during testing, we noted that including only a single reference to the 50-word limit in the prompt resulted in summaries that exceeded 50 words over 10% of the time, whereas the doubly-repeated emphasis results in 50-word summaries in nearly all cases. We also note that when ads criticize a candidate rather than advocate for one, we apply the same prompt modification (context phrase modification) as we did for its frame description prompts above.

**Example summary output.** Our method returns the following ad summary for the Hope, Arkansas example ad: “The 1992 political ad features then-candidate Bill Clinton narrating a touching dialogue of his humble beginnings in Hope, Arkansas, his work in public service, and his dreams of making a difference as President. Visually, the ad is a blend of his past and present, reflecting a life dedicated to civic engagement.”

**Scaling up to large video datasets.** We design our codebase to accommodate even very large-scale applications (including those that are orders of magnitude larger than the project described here) by using the MPI<sup>53</sup> parallelization standard via OPEN MPI<sup>54</sup> and the MPI4PY<sup>55</sup> Python module. Specifically, we note that the workflow described above exhibits what is known as ‘embarrassingly parallel’ at the level of videos (i.e. each video can be computed fully independently in parallel). Key frame description generation (**Step 3**) can also be parallelized at the level of individual ad frames. We conduct all processing steps described above by computing

on approximately 50 workloads simultaneously, which is sufficient for our purposes. However, we note that our codebase and techniques may easily be applied to much larger datasets (e.g., millions of videos) *without* requiring additional compute time simply by scaling up the processor count.

## Data Records

All presidential ads and associated transcripts, summaries, key frames, and metadata are available for *bulk download* on Harvard's Dataverse at: <https://dataverse.harvard.edu/dataverse/ctap><sup>56</sup>. We have organized the bulk data into subdirectories with one directory for each election year so that scholars interested in a particular subset of elections may bulk-download just that subset of interest. The filenames for each transcript, summary, and ad key frame begin with the filename of the video to which they correspond. That video filename also matches a corresponding row in the dataset metadata file that includes the ad's metadata. In the *Usage Notes* section below, we also describe our online interface designed to allow researchers to efficiently explore the data.

## Technical Validation I: Data Accuracy, Consistency, & Balance versus Human Coders

Our goal in this validation section is to show that (1) our workflow produces high-quality preprocessed videos, transcripts, and summaries *overall*, and (2) the quality is *consistent across features of interest*. We accomplish this by showing that across these three tasks, our methodology performs at least as well as human coders using standard measures of quality and this performance is independent of features such as partisanship and year.

**Stratified Random Sampling.** Our dataset is highly diverse, spanning not only competing political parties, but also seven decades of changes to speech styles and issue areas under discussion. The dataset also spans the revolution in audiovisual technology from black-and-white video cameras with mono vinyl-cut audio to high-definition color video with digitally processed audio. To ensure the integrity of downstream research, it is advantageous to explicitly test whether any of these sources of variance (including aspects that are imperceptible to human senses) cause our AI-based workflow to exhibit errors that affect one party or a certain time range of videos more than others.

To accomplish this, we employ a stratified sampling validation strategy. Specifically, instead of validating uniformly randomly sampled ads, we instead validate by sampling batches, where each batch contains one sampled ad for each party-year combination. For tasks that required additional statistical power, we asked our human research assistants to label additional batches of stratified samples. As a result of this sampling strategy, the high performance on the validation tests below indicates high quality that is *balanced* across the primary sources of variation in the dataset. We also test this explicitly by showing that quality has no statistically significant relationship with features of interest.

**Preprocessing Validation.** We empirically validate our preprocessing procedure that trims pre-roll and post-roll clips out of each ad's video file. Three undergraduate research assistants were asked to manually identify the beginning and end of political ads for 30 stratified sample batches of ads (i.e., 900 video files—see the “Stratified Random Sampling” section above). We then compare these human-coded ad beginning and ending timestamps with those inferred by our automated procedure.

**Quality Evaluation Metrics.** Our pipeline already ensures that no dialogue or onscreen people are ‘over-trimmed’ from the video, and that any pre-dialogue or post-dialogue segments of the ad will also be sampled and included in our AI summarization prompt via our regular-spaced 3-second interval key frame samples. Therefore, our primary goal in this section is to show that our automated procedure removes ‘pre-roll’ and ‘post-roll’ without overtrimming more than 3 seconds of pre-dialogue or post-dialogue imagery from the ads.

**Preprocessing Evaluation Results.** Compared to our human coders, our automated preprocessing procedure ‘overtrims’ the start (‘pre-roll’) of the videos by more than 3 seconds in approximately 2.44% of the stratified-sampled ads. Similarly, compared to our human coders, our automated preprocessing procedure ‘overtrims’ the end (‘post-roll’) of the videos by more than 3 seconds in approximately 3.33% of the stratified-sampled ads. This ensures that our automated preprocessing step obtains high accuracy for this task.

We also measure cases where our automated procedure ‘undertrims’ the beginning or end of the ads (i.e. leaves additional video frames) compared to humans, though such ‘errors’ are generally benign. Specifically, our automated procedure ‘undertrims’ the videos ‘pre-roll’ by more than 3 seconds in approximately 3.44% of the sampled ads, and it ‘undertrims’ the ‘post-roll’ in approximately 12.67% of the sampled ads. In cases where there is ‘undertrimming’ of the end, the mean undertrim leaves 2.38 seconds of additional video compared to human coders.

We also explicitly test whether errors affect one party or a certain time range of videos more than others. Specifically, we regress the difference between human-coded times and automated times for each ad on party, election year, and a party-election-year interaction term by estimating models of the form:

$$\text{Error}_{\text{type}} = \beta_0 + \beta_1 \text{Republican} + \beta_2 \text{ElectionYear} + \beta_3 \text{Republican} \times \text{ElectionYear} \quad (1)$$

We refit this model for different error types including (1) the videos’ start time errors, (2) the videos’ end time errors, and (3) the videos’ sum of absolute errors, i.e.,  $\text{Error}_j = |\text{Error}_{j,\text{start}}| + |\text{Error}_{j,\text{end}}|$ . In all cases, non-intercept coefficients  $\beta_1, \beta_2, \beta_3$  are statistically insignificant ( $p\text{-value} > 0.11$ ). This indicates that errors compared to human coders have no statistically significant relationship to these features of interest. We report details in the Appendix.

We also test the intercoder reliability of our three undergraduate research assistants on this task. Specifically, we uniformly subsampled a subset of ~200 of these sampled ads and asked all three research assistants to identify

start and end timestamps for these ads. We then compute Intraclass Correlation Coefficient (ICC) for the raters using the appropriate two-way random effects, absolute agreement, single rater/measurement model following the standard approach<sup>57</sup>. Our three raters have an ICC of 0.97 for their start times and 0.98 for their end times, indicating excellent inter-coder agreement. More concretely, across the stratified samples, the mean absolute difference between any two human coders is 0.62 seconds for ‘pre-roll’ timestamps and 0.67 seconds for ‘post-roll’ timestamps. This ensures that our human-generated timestamps provide a high-quality standard of comparison.

**Automated Transcription Validation.** Our ad transcripts from **Step 1** are generated by OpenAI’s WHISPER-LARGE-V3 model, which has already been extensively validated in other contexts<sup>51</sup>. Therefore, our goal is just to ensure that its performance on our ads is consistent with its known performance on benchmark video datasets. To accomplish this, we ask our human coders to manually transcribe a smaller stratified sample of 4 stratified batches of sampled ads (i.e., 120 ads—see “Stratified Random Sampling” Section above).

**Quality Evaluation Metrics.** The standard evaluation metric used to compare the automated transcripts to manually generated transcripts is Word Error Rate (WER). WER represents the count of errors in the AI transcript divided by the word count  $n$  of the reference (correct) transcript<sup>58</sup>. Here, an error is defined as either a substituted word (denoted by  $s$ ), a deleted/missing word (denoted by  $d$ ), or a superfluously inserted word (denoted by  $i$ ). Thus, we compute WER as:

$$\text{WER} := \frac{s + d + i}{n} \quad (2)$$

A WER of 0 indicates that the candidate text (i.e., the automated transcript) and the reference text (i.e., human transcript) are identical, while a WER of 0.10, for example, indicates that there are 10 errors for every 100 words in the reference transcript. For reference, the original Whisper paper<sup>51</sup> showed state-of-the-art WERs across 14 standard benchmark datasets ranging from a WER of 0.027 on ‘maximally clean’ high-definition speech data recorded in e.g. sound isolating voiceover booths (*LibriSpeech Clean* data) to a WER of 0.364 on speech audio recorded in meeting rooms (*AMI SDM1* data). In general, state-of-the-art transcription obtains WER values of 0.16–0.26 on standard ‘in-situ’ speech benchmark datasets<sup>51</sup>. For comparison, a single human transcriber has a WER of approximately 0.14 on these same benchmark datasets.

**Transcription Validation Results.** We compute two WER metrics for our validation. First, we compute the WER for each transcript and then take the mean over the transcripts. The transcripts have a mean WER of 0.0353. Second, we also consider the overall WER, which we compute by concatenating the 120 transcripts and calculating a single overall WER. This yields an overall WER of 0.0329. In both cases, we conclude that our transcripts exhibit roughly 3 errors per 100 words. This means that our transcripts exhibit higher accuracy than standard benchmark transcriptions of virtually all standard transcription datasets except for ‘maximally clean’ high-definition speech data recorded in voiceover booths (e.g. transcriptions of *LibriSpeech Clean* data obtain WER that differs from ours by a margin of less than a single error in every 100 words [see<sup>51</sup>]).

We also find that these rare transcription errors are qualitatively negligible in terms of their effect on the meaning of the text. Table 3 shows 5 transcripts randomly selected from the subset of validation transcripts that exhibited any errors. Note that errors are generally limited to minor spelling differences or pronouns, and do not change the substantive meaning of the text.

Finally, we test whether transcription errors affect one party or time range of ads more than others. Specifically, we model transcription accuracy as a function of party, election year, and a party-election-year interaction term. Here, we are careful to adopt modeling choices that reflect the word-error distributions of the transcripts. In particular, it is undesirable to model word errors directly, as errors within the same transcript are not independent. It is also undesirable to adopt WER as the response, as WER has a lower bound (0) but no upper bound in the presence of insertions, and the data-generating process for WER is unclear with respect to insertions.

Therefore, we consider a transcript-level response variable that is bounded and linear in the count of per-transcript errors. Specifically, we define a response variable,  $R_j$ , that captures transcription accuracy as the fraction  $0 \leq R_j \leq 1$  of words in a human transcript  $j$  that are correctly recovered by the AI transcript:

$$R_j := \frac{n_i - s_j - d_j}{n_i}.$$

We then use Beta regression to model this fractional response. Because the logit link used in the Beta regression is undefined when the response takes endpoint values, we adopt the standard approach of<sup>59,60</sup> and squeeze response values by a negligible amount:  $R_j \in [0, 1] \rightarrow R'_j \in (0, 1)$ , by defining  $R'_j := (R_j(N - 1) + 0.5)/N$ , where  $N$  is the sample size (i.e. 120). Now, we can proceed with the standard Beta regression where the transcription accuracy (response variable)  $R'_j$  follows a beta distribution with mean  $\mu_j$  and precision  $\phi$ , where the mean is linked to covariates via the standard logit link function  $g(x)$ :

$$R'_j \sim \text{Beta}(\mu_j \phi, (1 - \mu_j) \phi), \quad g(\mu_j) = \beta_0 + \beta_1 \text{Republican}_j + \beta_2 \text{ElectionYear}_j + \beta_3 \text{Republican}_j \times \text{ElectionYear}_j \quad (3)$$



ID: P-1232-57589 (WER: 0.0494): “For Hispanics like me, <i>this present</i> <del>[discretion in]</del> elections may be the most important in our history because Hispanics everywhere have a real chance to be heard, to elect a president who cares about our needs, our proud traditions, our strong family values. But don’t take my word for it. Ask my father-in-law. As president, I will have a lot of reasons to help Hispanics everywhere because I’ll not only be answering to my grandchildren, I will be answering to history.”
ID: P-1905-108595 (WER: 0.0133): “The people of Florida face an important choice do <del>you</del> <del>[we]</del> elect a president who will raise your taxes, or fight for working families? A president who will spread your income, or let you keep what is yours? A president who will spend us deeper into debt, or control spending and eliminate waste? Tough times take experienced bipartisan leadership. Join me and vote for senator John McCain. I am John McCain and I approve this message.”
ID: P-334-15180 (WER: 0.109): <del>whose</del> <del>[who is]</del> bill did Congress pass to keep the lid on oil prices? Udall Bayh <del>Sh</del> <del>[c]</del> hriver Wallace Sanford <del>Shapp</del> <del>[Schaap]</del> Carter Harris? Nope, Senator Scoop Jackson. Who’s the only candidate to win the Sierra Club’s highest environmental award? Wallace Sanford <del>Shapp</del> <del>[Schaap]</del> Harris Carter <del>Sh</del> <del>c</del> river Bayh Udall? Nope, Senator Scoop Jackson. Which candidate insists <del>that</del> we stop exporting our jobs to Russia? Sanford <del>Shapp</del> <del>[Schaap]</del> Harris Wallace Udall Carter <del>Sh</del> <del>[c]</del> hriver Bayh? Nope, Senator Scoop Jackson. Let’s have a president who does something. Elect Senator Scoop Jackson while there is still time to save the 70’s.
ID: P-1631-83508 (WER: 0.0118): “As Governor, George W. Bush gave big oil a tax break, while opposing health care for 220,000 kids. Texas now ranks 50th in family health care. He’s left the minimum wage at \$3.35 an hour, let polluters police themselves. Today, Texas ranks last in air quality. Now Bush promises the same one trillion <del>dollars</del> from Social Security to two different groups. He squanders the surplus on a tax cut for those making over \$300,000. Is he ready to lead America?”
ID: P-1840-104205 (WER: 0.0108): “There’s nothing wrong with making a decision and then changing your mind. But if you never commit to what you believe in, who will ever commit to you? John Kerry has changed his mind on all these important issues. Now, there’s nothing wrong with a little indecision as long as your job doesn’t involve any responsibility. John Kerry has changed his mind time and time again. If you thought you could trust him, you might want to change your mind too. Clubfo <u>[u]</u> rGrowth.net is responsible for the content of this advertising.”

**Table 3.** Transcription error examples. For five transcripts randomly selected from validation transcripts with nonzero errors. Words in the reference transcript that the AI transcript missed are denoted by strikethrough text. Words (or single letters) that the AI transcript erroneously added are denoted by bold text.

Once again, we find that all non-intercept coefficients,  $\beta_1, \beta_2, \beta_3$ , are statistically insignificant (all  $p > 0.5$ ), indicating that the high transcription accuracy we observe is consistent regardless of party and year. We defer details to the Appendix.

**Video Summarization Validation.** The goal of our AI-generated ad summaries is to facilitate academic research on political campaign ads. As such, our workflow should provide researchers with a coherent, neutral, accurate summary that efficiently conveys the main themes and relevant or interesting aspects in each ad. Therefore, we seek summaries that exhibit high quality across *multiple* key dimensions. We note that this goal differs in some important respects from the objective functions used to train standard LLMs, which emphasize only fluency (via token prediction accuracy).

**Quality Evaluation Metrics.** To measure these dimensions, we ask human subject matter experts to evaluate the summaries based on the 4 quality evaluation criteria of<sup>61</sup> and<sup>62</sup>: *coherence, consistency, fluency*, and *relevance*. Table 4 gives the definitions of each of these criteria.

**Comparison Benchmark: Human-Generated Summaries.** We compare a validation set of our 50-word AI-generated summaries from **Step 4** to 50-word summaries composed by undergraduate research assistants (we note that in recent work on political ads, most ad metadata is produced by undergraduate research assistants). Specifically, for each ad in 4 stratified sample batches (i.e. 120 sampled ads), we randomly selected one of our undergraduate RAs and asked them to compose a summary via the following prompt:

*We are releasing a large collection of political ads to the academic community. Summaries of these ads will also be released. Please help this project by writing a 50-word summary of the ad you have been assigned. The summary should be written in a way that is useful for academic research. Also, since we know the names of the candidates, please do not write their names in your summary. Rather, just refer to them as ‘the candidate’ or ‘the opposing candidate.’*

**Expert Human Annotators and Online Evaluation Setup.** To ensure fairness in the evaluation, we created an online evaluation tool using a JavaScript embedded within a Qualtrics survey. This tool randomly displayed an ad and two summaries, labeled “Summary A” and “Summary B.” The human summary was randomly assigned to either “Summary A” or “Summary B,” and the LLM-generated summary assigned the remaining label. We then asked political science graduate students to answer four questions about each summary, corresponding to the four summary evaluation metrics identified in the literature: coherence, consistency, fluency, and relevance<sup>[62, 396]</sup>. For example, to evaluate a summary’s consistency, which is defined as the “factual alignment between the summary and the summarized source,” we asked to what extent someone agreed or disagreed with the following statement: “The summary is factually consistent with the advertisement.” Annotators could select from 5 responses ranging from “strongly disagree” to “strongly agree”, which we coded on a 5-point integer Likert scale as  $-2$  and  $2$ , respectively. We ran the evaluation until each ad in the stratified sample had been rated by three graduate students.

Dimension	Definition
Coherence	Coherent and easy to understand.
Consistency	Factually consistent with the advertisement.
Fluency	Contains the main themes of the advertisement.
Relevance	Useful for research on presidential advertisements.

**Table 4.** Dimensions of summary quality rated by human coders.

Dimension	AI mean	Human mean	diff.	t-test 95% C.I.	t-statistic	p-value	V-statistic	p-value
Coherence	1.52	1.37	0.15	(−0.018, 0.313)	1.753	0.081*	1493.5	0.104
Consistency	1.54	1.61	−0.07	(−0.231, 0.087)	0.895	0.371	1856.0	0.491
Fluency	1.64	1.49	0.15	(0.005, 0.289)	2.041	0.042**	1261.5	0.028**
Relevance	1.51	1.43	0.08	(−0.075, 0.242)	1.037	0.301	1609.5	0.146
Aggregate (N = 480 pairs)	1.55	1.48	0.08	(−0.001, 0.153)	1.956	0.052*	2258.5	0.009***

**Table 5.** Video summarization validation results. AI-generated summaries perform at least as well as human summaries. *t*-statistics are from 2-sided *t*-tests. *V*-statistics are from paired samples Wilcoxon tests (Note: \**p* < 0.1; \*\**p* < 0.05; \*\*\**p* < 0.01). The sample size of *N* = 120 paired ratings per quality dimension corresponds to 120 stratified-sampled movies, each with one AI-generated summary and one human-written summary. Each summary was rated by three independent raters, and their ratings were averaged to produce a single score per summary per quality dimension.

**Summary Validation Results.** We employ a range of statistical methods to analyze the comparative quality of human- and AI-generated ad summaries. Table 5 reports mean responses of three graduate research assistants across the four quality dimensions. We include standard paired *t*-tests as well as Wilcoxon signed-rank tests; the latter tests are more robust given that our response variables are Likert-scale ratings. According to both types of tests, the mean score of our AI summaries taken across all quality dimensions outperforms the human summaries by a statistically significant margin (Table 5, bottom row). Our AI summaries also exhibit significantly better fluency—that is, they are more likely to convey the main themes of each ad. Other dimensions exhibit either a slight advantage for our AI summaries, or no statistically significant difference between AI and human summaries.

While paired *t*-tests and Wilcoxon signed-rank tests provide valuable baseline comparisons, they do not account for within-rater dependencies or control for key features such as year and party. To address these aspects, we also estimate a Mixed-Effects Ordinal Logit (MEOL) model, a form of Cumulative Link Mixed Model that appropriately models the ordinal (Likert) dependent variable while incorporating random effects to account for dependencies in the data. Our MEOL model includes:

- *Treatment effect of AI vs. human authorship* ( $\beta_1$ ) measures the causal effect of AI-generated summaries on ratings;
- *Covariates for year, party, and their interaction* ( $\beta_2, \beta_3, \beta_4$ ) control for systematic differences in ads;
- *Random intercepts for raters* ( $u_{\text{Rater}}$ ) capture latent differences in individuals' scoring tendencies (within-rater correlation);
- *Random intercepts for ads* ( $u_{\text{Ad}}$ ) account for different ads' latent variation in summary difficulty.

Specifically, the MEOL model estimates the probability that, for one of the four quality dimensions *d*, a given summary rating *Y* falls within a particular Likert category. The MEOL model takes the form:

$$\text{logit}(P(Y_d \leq k)) = \tau_k - \beta_1 \text{AI} + \beta_2 \text{Republican} + \beta_3 \text{ElectionYear} + \beta_4 \text{ElectionYear} \times \text{Republican} + u_{\text{Rater}} + u_{\text{Ad}} \quad (4)$$

Here, *Y* represents a Likert rating of a summary in one of the quality dimensions *d*, and  $P(Y \leq k)$  denotes the cumulative probability of receiving a rating in category *k* or lower. Note that the negative sign before the predictor effects means that larger values of  $X\beta$  decrease the cumulative probability  $P(Y_d \leq k)$  of a lower rating, shifting the probability mass toward higher ordinal categories. This matches our intuition that larger  $X\beta$  corresponds to a better rating. The thresholds  $\tau_k$  then define the latent decision boundaries between ordinal categories. The terms in parentheses capture the main predictors of interest that estimate the treatment effect of AI summaries on ratings and control for covariates such as party and year. Finally,  $u_{\text{Rater}}$  and  $u_{\text{Ad}}$  terms are random intercepts that account for individual-level variation in rater tendencies and latent characteristics of the ad summaries, respectively. By including both random effects (*u*-terms) and covariates, the MEOL model allows us to properly account for the pairwise structure of the experiment, where each rater evaluates both the human- and AI-generated summary for the same ad, while also controlling for individual rater biases and latent ad characteristics that could influence ratings.

	Coherence			Consistency			Fluency			Relevance		
	MEOL	OLS	OLS-FE	MEOL	OLS	OLS-FE	MEOL	OLS	OLS-FE	MEOL	OLS	OLS-FE
AI	0.335**	0.147*	0.147***	-0.171	-0.072	-0.072	0.620**	0.147**	0.147***	0.332**	0.083	0.083
	(0.164)	(0.078)	(0.055)	(0.186)	(0.072)	(0.056)	(0.177)	(0.069)	(0.052)	(0.165)	(0.073)	(0.055)
	[0.179]			[0.194]			[0.193]			[0.179]		
Republican	0.300	-2.655		0.341	4.280		0.317	-5.917		0.236	8.417	
	(0.278)	(11.035)		(0.283)	(10.921)		(0.265)	(8.901)		(.265)	(10.407)	
	[0.226]			[0.233]			[0.216]			[0.218]		
Election Year	0.013	0.002		0.020*	0.007*		0.008	0.0004		0.019*	0.006	
	(0.011)	(0.004)		(0.011)	(0.004)		(0.010)	(0.003)		(0.010)	(0.004)	
	[0.009]			[0.010]			[0.009]			[0.009]		
Republican × Election Year	-0.008	0.001		-0.002	-0.002		0.006	0.003		-0.021	-0.004	
	(0.016)	(0.006)		(0.016)	(0.005)		(0.015)	(0.004)		(0.015)	(0.005)	
	[0.014]			[0.014]			[0.013]			[0.013]		
Observations	720	720	720	720	720	720	720	720	720	720	720	720

**Table 6.** Human vs. AI Summary Results: Mixed-Effects Ordinal Logit (MEOL), OLS, & OLS with ad fixed effects (OLS-FE). SE's in parentheses; bootstrapped SE's in brackets; OLS and OLS-FE SE's are robust & ad (OLS) or rater (OLS-FE) clustered. Note: \*p < 0.1; \*\*p < 0.05; \*\*\*p < 0.01.

Table 6 reports the results of the MEOL regression computed for each of the four quality dimensions, Coherence, Consistency, Fluency, and Relevance. To assess statistical significance, we report both standard errors (in parentheses) and ad-clustered bootstrapped standard errors (in square brackets). While MEOL does not assume normality of residuals in the way that linear regression does, it nonetheless relies on the assumption that random effects are normally distributed. Therefore, bootstrapped standard errors allow us to test for significant effects even under potential violations of these assumptions.

Finally, we also estimate two variants of OLS linear regression for each quality dimension  $d$  and report it alongside the corresponding MEOL estimates. First, we estimate an OLS model with party, election year, and interaction controls as above. We report standard cluster-robust errors clustered by ad to account for potential within-ad rating correlations:

$$\text{OLS: } Y_d = \alpha + \beta_1 \text{AI} + \beta_2 \text{Republican} + \beta_3 \text{ElectionYear} + \beta_4 (\text{ElectionYear} \times \text{Republican}) + \varepsilon, \quad (5)$$

Second, we estimate an OLS model with ad fixed effects (OLS-FE), which fully absorb party and year controls, and we report the usual Huber-White standard errors. This Neymanian approach allows us to explicitly restrict measurement of the AI treatment effect to pairwise AI-vs-human summary comparisons:

$$\text{OLS-FE: } Y_d = \beta_1 \text{AI} + \gamma_{ad} + \varepsilon, \quad (6)$$

While OLS models do not account for the ordinal nature of the Likert ratings or the hierarchical structure of our data, they nonetheless provide useful and familiar baselines.

**MEOL and OLS Results of Human vs AI Summaries.** We find that AI summaries exhibit significantly better coherence, fluency, and relevance compared to their human-authored counterparts. Specifically, MEOL coefficients represent the log-odds of moving into a higher category of the ordinal response (i.e. rating). Therefore, we exponentiate these coefficients to obtain interpretable results. We find that our AI summaries have:

- 39.8% greater odds of obtaining a higher coherence score;
- 85.9% greater odds of obtaining a higher fluency score;
- 39.4% greater odds of obtaining a higher relevance score.

The remaining dimension, consistency, is the only dimension where the AI-generated summary was not rated significantly higher than the human-generated summary, suggesting that both are equally adept at capturing the factual information in the ad. We also find that the remaining covariates are statistically insignificant, with the exception of election year, which is moderately correlated with two of the four quality dimensions at the  $p < 0.1$  level. We note that the magnitude of these potential time effects remains small.

## Technical Validation II: Tracing the Historical Emergence of Key Issues

We briefly validate and illustrate our dataset by considering an application that measures the historical evolution of current key electoral issue areas. A central focus of political communication research is understanding how such issues emerge and shift over time<sup>13,14,52,63</sup>. This task not only demonstrates the dataset's substantive relevance, but also serves as a validation check: do the trends it reveals correspond to historically recognized developments in campaign issue agendas? We find that our primary topics are generally consistent—though not perfectly aligned—with a well-known coding of subjects and themes that analyzes the transcripts from a smaller set of 895 campaign advertisement from 1952–2000<sup>63</sup>. Such validation is especially valuable given the difficulty

Topic	Election	Keywords
Economy	2024, 2020, 2016	economy, employment, unemployment, inflation, growth, price, recession, markets, trade, wage
Immigration	2024	immigration, immigrant, border, naturalization, alien, asylum
Abortion	2024 & advertiser	abortion, pro-life, pro-choice, roe, roe v wade, termination, maternal,
Terrorism	2020	terrorism, isis, isil, homeland, radical, security, osama, laden, attack, 9-11, al-qaeda
Foreign Policy	2020	nato, china, russia, benghazi, alliance, foreign, war, peace
Supreme Court	2016	supreme, court, appointment, justices, ginsburg, barrett
Healthcare	2016	healthcare, obamacare, medicare, prescription, insurance
Taxation	advertiser	taxation, taxes, bracket, income
Crime	advertiser	crime, police, prison, sentence, criminal, gang, murder, rape, rapist, murderer, superpredator

**Table 7.** Keyword topics and their associated keywords. Keywords for topics 1 to 7 represent the top 3 issue areas by voter interest from the most recent 3 elections according to Pew Research (see the *Election* column). Keywords 8 and 9 on Taxation and Crime reflect two remaining issue areas that are top-3 issues (along with Abortion) per current advertiser attention.

of analyzing electoral issue dynamics prior to the 1996 election due to limited data availability and high labor costs involved in coding more than a few hundred ads<sup>63</sup>.

To accomplish this, we estimate a keyword-assisted dynamic topic model (dynamic keyATM) on the corpus of ad transcripts<sup>45</sup>. This model conveys two advantages for this application. First, the model allows topic proportions to evolve across election years. Second, it permits us to seed keyword topics corresponding to the top issue areas from the three most recent presidential elections (2024, 2020, 2016).

We select keywords for our dynamic keyATM model by choosing issue areas that reflect the three issues most important to voters in these three elections according to Pew Research<sup>64–66</sup>. Due to some overlapping issues across these elections, this yields a total of 7 keyword-seeded topics, which we show in Table 7. However, according to recent research, campaign messaging often focuses on a set of topics that diverge from voters’ priorities. For example, taxation and crime—two issues that do not appear in voters’ top three priorities—are among the top three issue areas in the 2024 election in terms of advertiser attention<sup>1,2</sup>. We therefore include two additional keyword-seeded topics to capture them. Finally, we complement these 9 keyword topics with 25 additional unsupervised topics to accommodate the range of (potentially unanticipated) textual themes.

Figure 4 plots the estimated proportions of topics for the 9 issue-area-focused topics over the seven decades of ads in the dataset. While an in-depth analysis is beyond the scope of this paper, we find that the estimated proportions of foreign policy issues are inversely correlated with domestic issues related to the economy and taxation.

For example, consistent with the theoretical conjecture that incumbents should emphasize economic issues during periods of prosperity<sup>63</sup>, ads in the 1950s often focus on the booming postwar economy and tax policy. These domestic topics gave way to foreign policy issues during periods when insurgent candidates attempted to draw attention to Cold War developments or when incumbents highlighted their opponents’ foreign policy positions. The 1964 campaign is illustrative: a spike in the foreign policy topic reflects Lyndon B. Johnson’s relentless critique—exemplified by the iconic “Daisy” ad—of Barry Goldwater’s escalatory support for deploying limited tactical nuclear weapons against communists in the immediate wake of the Cuban Missile Crisis and Johnson’s recent endorsement of the Nuclear Test Ban Treaty.


Economic and tax issues resurface during the stagflation era and again during Bill Clinton’s heresthetical<sup>67</sup> pivot to the “it’s the economy, stupid” strategy. These topics recede once more as foreign policy dominated the 2000s, particularly as George W. Bush focused his 2004 reelection campaign overwhelmingly on the War on Terror, favoring bellicose imagery and interventionist military rhetoric against Kerry’s more topically varied messaging that included a significant focus on healthcare reform (reflected in the spike in the healthcare topic in Fig. 4, top plot). In 2008, Barack Obama and John McCain maintained the central focus on terrorism and foreign affairs, though by then the discourse had largely shifted toward troop withdrawal and accountability. Obama’s campaign concurrently increased focus on economic topics as he refocused the contest through his “Wall Street versus Main Street” refrain and attempted to catch challenger McCain flat-footed on issues related to the contemporaneous onset of the financial crisis.

We also observe the nascence, ebb, and flow of the now-pervasive ‘crime’ issue area, beginning with the infamous ‘Horton weekend pass’ dog whistle that aired during Bush’s 1988 campaign, and resurging in Bill Clinton’s infamous ‘superpredator’ and ‘three strikes’ marketing. We note that political campaign supply-side measures of content such as these offer an interesting point of comparison to standard opinion polls that only capture the ‘demand-side’ of political campaigns.

Usage Notes

In addition to bulk dataset access described in the *Data Records* section above, we also provide a project website (<https://campaignTVads.org/>) with a convenient online interface designed to facilitate research by allowing practitioners to search and view videos that are relevant to their specific research questions.

**Videos by Candidate, Election, and/or Exact & Fuzzy Transcript Keyword Search.** The “Data” tab of our project website provides drop-down menus that subset the videos by year and candidate’s last name. It also provides a search bar that queries the transcripts of all the ads that were generated in Step 1 above. The



**Fig. 4** Evolution of key 2016–2024 election issues over time. We plot topic prevalence (measured via mean topic probability) over presidential elections from 1952–2012. Topics focus on 9 key issue areas, which represent the union of the top 3 issues from each of the most recent 3 presidential elections, 2016–2024, plus two additional topics (taxation and crime) that reflect remaining issues among top-3 issue areas in terms of advertiser attention. For presentation, we partition the issues into three plots corresponding to economic policy and social welfare, security and international affairs, and law and the courts. Error ribbons represent 90% credible intervals computed via the standard highest density interval method.

search engine then returns videos whose transcripts contain either exact or fuzzy matches, with the former being returned first. The videos returned by the search engine are hosted on Vimeo, where they can be *viewed online* and also individually *downloaded* for further analysis.

### Code availability

Our replication codes as well as codes for applying our workflow to other video datasets can be found on GitHub at: <https://github.com/BreuerLabs/AI-SummarizeVid>.

Received: 30 March 2025; Accepted: 7 July 2025;

Published online: 24 September 2025

### References

- Adimpack: The play for the white house (2024).
- Adimpack: The 2024 cycle in review (2024).
- Borrell Associates. The final analysis: What happened to political advertising in 2016 (and forever) (2017).
- Jamieson, K. H. *Packaging the presidency: A history and criticism of presidential campaign advertising* (Oxford University Press, 1996).
- Kaid, L. L. Political advertising. *Handbook of political communication research* **15**202 (2004).
- Brader, T. *Campaigning for hearts and minds: How emotional appeals in political ads work* (University of Chicago Press, 2006).
- Biocca, F. *Television and political advertising: Volume I: Psychological processes* (Routledge, 2013).
- Valentino, N. A., Hutchings, V. L. & Williams, D. The impact of political advertising on knowledge, internet information seeking, and candidate preference. *Journal of communication* **54**, 337–354 (2004).
- Goldstein, K. & Freedman, P. Campaign advertising and voter turnout: New evidence for a stimulation effect. *Journal of Politics* **64**, 721–740 (2002).
- Geer, J. G. *In defense of negativity: Attack ads in presidential campaigns* (University of Chicago Press, 2008).
- Gerber, A. S., Gimpel, J. G., Green, D. P. & Shaw, D. R. How large and long-lasting are the persuasive effects of televised campaign ads? results from a randomized field experiment. *American Political Science Review* **105**, 135–150 (2011).
- Fowler, E. F., Franz, M. M. & Ridout, T. N. *Political advertising in the United States* (Routledge, 2018).
- Petrocik, J. R., Benoit, W. L. & Hansen, G. J. Issue ownership and presidential campaigning, 1952–2000. *Political Science Quarterly* **118**, 599–626 (2003).
- Hayes, D. Candidate qualities through a partisan lens: A theory of trait ownership. *American Journal of Political Science* **49**, 908–923 (2005).
- Fowler, E. F., Franz, M. M. & Ridout, T. N. *Political Advertising in the United States* (Westview Press, Boulder, Colo., 2016).
- Goldstein, K., Franz, M. & Ridout, T. Political advertising in 2000. *Combined File [dataset]*. Final release. Madison, WI: The Department of Political Science at The University of Wisconsin-Madison and the Brennan Center for Justice at New York University (2002).



17. Goldstein, K. & Rivlin, J. Congressional and gubernatorial advertising, 2003–2004. *Combined File [dataset]. Final release. Madison, WI: The University of Wisconsin Advertising Project, The Department of Political Science at The University of Wisconsin-Madison* (2007).
18. Prior, M. Weighted content analysis of political advertisements. *Political Communication* **18**, 335–345 (2001).
19. Carl albert center julian p. kanter political commercial collection. <https://arc.ou.edu/>. Accessed on September 26, 2024.
20. Tarr, A., Hwang, J. & Imai, K. Automated coding of political campaign advertisement videos: An empirical validation study. *Political Analysis* **31**, 554–574, <https://doi.org/10.1017/pan.2022.26> (2023).
21. Harp, A. B. The 60 second candidate. *Sooner Magazine* **36**, 8–11 (2016).
22. Haynes, K., Kaid, L. & Rand, C. The political commercial archive: Management of moving image and sound recordings. *The American Archivist* **59**, 48–61 (1996).
23. Pryse, J. Practical remote workflow solutions for complex digital projects: Opportunities in a pandemic. *Collections* **18**, 258–279 (2022).
24. Clinton, W. J. The living room candidate: Presidential campaign commercials: 1952–2004. *American Museum of the Moving Image* (1992).
25. Kaid, L. L. & Johnston, A. Negative versus positive television advertising in us presidential campaigns, 1960–1988. *Journal of communication* **41**, 53–64 (1991).
26. Finkel, S. E. & Geer, J. G. A spot check: Casting doubt on the demobilizing effect of attack advertising. *American journal of political science* 573–595 (1998).
27. Damore, D. F. Using campaign advertisements to assess campaign processes. *Journal of Political Marketing* **3**, 39–59 (2003).
28. Prior, M. *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections* (Cambridge University Press, 2007).
29. Brader, T. *Campaigning for hearts and minds: How emotional appeals in political ads work* (University of Chicago Press, 2020).
30. Schofield, B. L. & Bailey, L. Amicus brief in support of defendant-appellant-cross-appellee in fox news v. tv eyes (2016).
31. Colt, S. (Director) Dobrowksi, H. (Producer). Prayer. politics. power. A Sarah Colt Productions Film for American Experience (2015).
32. Treitz, Z. (Director) Ambrose, M. (Producer). American conspiracy: The octopus murders (2024).
33. MacNicol, G. Wilder: Daughter dearest pt. 2: Politics and rose. I Heart Radio (June 29, 2023).
34. Beard, F. & Timke, E. Advertising in the archives: A guide to advertising and marketing archives around the world. *Advertising & Society Quarterly* **21** (2020).
35. Dover, E. Videostyle in presidential campaigns: Style and content of televised political advertising. *Great Plains Quarterly* **22**, 60 (2002).
36. Bernard, S. C. & Rabin, K. *Archival storytelling: a filmmaker's guide to finding, using, and licensing third-party visuals and music* (Routledge, 2020).
37. Biocca, F. *et al.* Television and political advertising. *Taylor & Francis eBooks DRM Free Collection* (1991).
38. Dalton, P. & McIlwain, C. Third-party “hatchet” ads: An exploratory content study comparing third-party and candidate spots from the 2004 presidential election. *Atlantic Journal of Communication* **19**, 129–151 (2011).
39. Biocca, F. *Television and Political Advertising: Volume II: Signs, Codes, and Images* (Routledge, 2014).
40. Dunaway, J., Searles, K., Fowler, E., Ridout, T. & Napoli, P. The effects of political advertising: assessing the impact of changing technologies, strategies and tactics. *Mediated communication: Handbooks of communication science* **7**, 1–22 (2019).
41. Webb, C. Freedom for all? blacks, jews, and the political censorship of white racists in the civil rights era. *American Jewish History* **94**, 267–297 (2008).
42. Nelson, R. A. Sources for archival research on film and television propaganda in the united states. *Film History* 333–340 (1989).
43. Albitz, R. S. Locating moving image materials for multimedia development: A reference strategy. *The Reference Librarian* **34**, 99–110 (2001).
44. Grimmer, J., Roberts, M. E. & Stewart, B. M. *Text as data: A new framework for machine learning and the social sciences* (Princeton University Press, 2022).
45. Eshima, S., Imai, K. & Sasaki, T. Keyword-assisted topic models. *American Journal of Political Science* **68**, 730–750 (2024).
46. Egami, N., Fong, C. J., Grimmer, J., Roberts, M. E. & Stewart, B. M. How to make causal inferences using texts. *Science Advances* **8**, eabg2652 (2022).
47. Fong, C. & Grimmer, J. Causal inference with latent treatments. *American Journal of Political Science* **67**, 374–389 (2023).
48. Freedman, P. & Goldstein, K. Measuring media exposure and the effects of negative campaign ads. *American journal of political science* 1189–1208 (1999).
49. Goldstein, K. M., Freedman, P. B. & Franz, M. M. *Campaign Advertising and American Democracy* (Temple University Press, Philadelphia, PA, 2009).
50. Li, F-F. Announcing Google Cloud Video Intelligence API, and more Cloud Machine Learning updates. *Google Cloud Blog*. Available at: <https://cloud.google.com/blog/products/gcp/announcing-google-cloud-video-intelligence-api-and-more-cloud-machine-learning-updates> (2017).
51. Radford, A. *et al.* Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, 28492–28518 (PMLR, 2023).
52. Fowler, E. F., Franz, M. M., Martin, G. J., Peskowitz, Z. & Ridout, T. N. Political advertising online and offline. *American Political Science Review* **115**, 130–149 (2021).
53. Walker, D. W. & Dongarra, J. J. Mpi: a standard message passing interface. *Supercomputer* **12**, 56–68 (1996).
54. Gabriel, E. *et al.* Open mpi: Goals, concept, and design of a next generation mpi implementation. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface: 11th European PVM/MPI Users' Group Meeting Budapest, Hungary, September 19–22, 2004. Proceedings* **11**, 97–104 (Springer, 2004).
55. Dalcin, L. & Fang, Y.-L. L. mpi4py: Status update after 12 years of development. *Computing in Science & Engineering* **23**, 47–54 (2021).
56. Breuer, A., Butler, M., Crespin, M. H., Dietrich, B. J. & Imai, K. Campaign television advertisement project (CTAP) dataset. DOI is associated with the first year of data available. Full archive & all years available at <https://doi.org/10.7910/DVN/7CJSLD> (2025).
57. McGraw, K. O. & Wong, S. P. Forming inferences about some intraclass correlation coefficients. *Psychological methods* **1**, 30 (1996).
58. Hamed, I. *et al.* Benchmarking evaluation metrics for code-switching automatic speech recognition. In *2022 IEEE Spoken Language Technology Workshop (SLT)*, 999–1005 (IEEE, 2023).
59. Smithson, M. & Verkuilen, J. A better lemon squeezer? maximum-likelihood regression with beta-distributed dependent variables. *Psychological methods* **11**, 54 (2006).
60. Smithson, M. & Verkuilen, J. Supplementary materials for: A better lemon squeezer? maximum-likelihood regression with beta-distributed dependent variables. *Psychological methods* **11**, 54 (2006).
61. Kryściński, W., Keskar, N. S., McCann, B., Xiong, C. & Socher, R. Neural text summarization: A critical evaluation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 540–551, <https://doi.org/10.18653/v1/D19-1051> (Association for Computational Linguistics, Hong Kong, China, 2019).

62. Fabbri, A. R. *et al.* Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics* **9**, 391–409 (2021).
63. Vavreck, L. *The Message Matters: The Economy and Presidential Campaigns*. Political science : American politics : Economics (Princeton University Press, 2009).
64. Doherty, C., Kiley, J. & Asheer, N. In tied presidential race, harris and trump have contrasting strengths, weaknesses. *Pew Research Center* **September** (2024).
65. Doherty, C., Kiley, J., Asheer, N. & Jordan, C. Election 2020: Voters are highly engaged, but nearly half expect to have difficulties voting. *Pew Research Center* **August** (2020).
66. Doherty, C., Kiley, J. & Johnson, B. 2016 campaign: Strong interest, widespread dissatisfaction. *Pew Research Center* **July** (2016).
67. Riker, W. *The Art of Political Manipulation*. Political science (Yale University Press, 1986).

## Acknowledgements

This project was funded by a grant from the National Science Foundation's Accountable Institutions and Behavior (AIB) program and the Established Program to Stimulate Competitive Research (EPSCoR). Award Numbers SES-2147635 and SES-2148928. Adam Breuer's research lab is funded by OpenAI. Financial support was provided by the University of Oklahoma Libraries' Open Access Fund.

## Author contributions

J.P. led the archival team, and J.P. and M.C. directed the student coders at the University of Oklahoma. B.D. led the validation effort using student coders at the University of Iowa. A.B. created the preprocessing protocol, video summarization workflow, and the automated transcripts and summaries. M.B. created various online tools that helped in the coding effort, as well as the project website. A.B., J.P., and K.I. designed the statistical validation. A.B. conducted the empirical illustration with topic modeling. K.I. helped structure and develop all aspects of the project and manuscript. All authors reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to A.B. or B.J.D.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025