



# AudienceView: AI-Assisted Interpretation of Audience Feedback in Journalism

William Brannon  
wbrannon@mit.edu  
MIT Center for Constructive  
Communication & MIT Media Lab  
Cambridge, MA, USA

Doug Beeferman  
doughb5@mit.edu  
MIT Center for Constructive  
Communication & MIT Media Lab  
Cambridge, MA, USA

Hang Jiang  
hjian42@mit.edu  
MIT Center for Constructive  
Communication & MIT Media Lab  
Cambridge, MA, USA

Andrew Heyward  
aheyward@media.mit.edu  
MIT Center for Constructive  
Communication & MIT Media Lab  
Cambridge, MA, USA

Deb Roy  
dkroy@mit.edu  
MIT Center for Constructive  
Communication & MIT Media Lab  
Cambridge, MA, USA

## Abstract

Understanding and making use of audience feedback is important but difficult for journalists, who now face an impractically large volume of audience comments online. We introduce AudienceView, an online tool to help journalists categorize and interpret this feedback by leveraging large language models (LLMs). AudienceView identifies themes and topics, connects them back to specific comments, provides ways to visualize the sentiment and distribution of the comments, and helps users develop ideas for subsequent reporting projects. We consider how such tools can be useful in a journalist's workflow, and emphasize the importance of contextual awareness and human judgment.

## CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**; **Web-based interaction**; **Visualization toolkits**; • **Computing methodologies** → **Natural language processing**.

## Keywords

natural language processing, text analysis, journalism, YouTube

### ACM Reference Format:

William Brannon, Doug Beeferman, Hang Jiang, Andrew Heyward, and Deb Roy. 2024. AudienceView: AI-Assisted Interpretation of Audience Feedback in Journalism. In *Companion of the 2024 Computer-Supported Cooperative Work and Social Computing (CSCW Companion '24), November 9–13, 2024, San Jose, Costa Rica*. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3678884.3681821>

## 1 Introduction

Journalists have a longstanding interest in the identity of the audience and how that audience responds to the journalists' work [19], though this impulse is balanced against the need to preserve

journalistic independence. Journalism's traditional means of obtaining and considering audience feedback, however, have been upended in recent years by the rise of the Internet, which provides far more input from readers and viewers than ever before. In place of letters to the editor and Nielsen audience figures, journalists now have access to detailed clickstream information, engagement statistics from social media, and above all a wealth of reader comments. Leveraging recent advances in natural language processing, themselves driven by the same increase in the scale of web data, we aim to simplify the task of engaging with these comments and deriving actionable insights from them.

This paper introduces AudienceView,<sup>1</sup> an AI-assisted tool for journalists to make sense of audience feedback on YouTube. We have developed AudienceView in partnership with PBS's Frontline team, and development is thus targeted toward YouTube-hosted video journalism. The tool is, however, modular and extensible to other comment sources. Our evaluation applies it to all Frontline documentaries during the more than 10 years between August 2013 and January 2024 (250 in total), comprising just over 599,000 comments.

AudienceView complements prior work on AI for qualitative analysis and sensemaking [4, 8, 16] in the particular context of journalism. Some of its components can also be viewed as a kind of summarization, building on a long literature in NLP on that question. Because of the diversity of the comments and use cases, we have designed AudienceView with two goals in mind: a) providing more than one way to visualize and understand the data, and b) keeping a close connection to the underlying comments, and surfacing them wherever possible. Both are intended to complement, rather than replace, human judgment in producing news.

## 2 Related Work

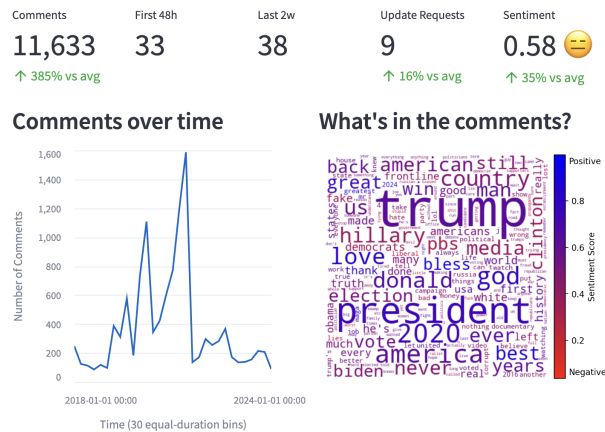
*NLP for sensemaking and qualitative analysis.* Both commercial [12, 22] and academic [16] tools use AI for qualitative analysis on textual data. These tools, however, are sometimes expensive and remain within a qualitative analysis paradigm that requires considerable effort from the analyst. A number of NLP tools and topic



This work is licensed under a Creative Commons Attribution International 4.0 License.

*CSCW Companion '24, November 9–13, 2024, San Jose, Costa Rica*  
© 2024 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-1114-5/24/11  
<https://doi.org/10.1145/3678884.3681821>

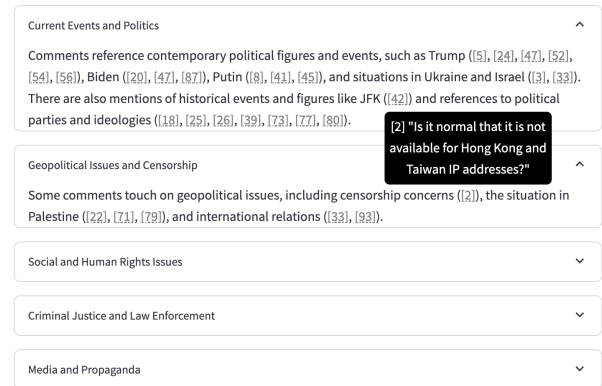
<sup>1</sup>A deployed version of the tool is available at <https://frontline.ccc-mit.org/>, and the underlying code is available at <https://github.com/mit-ccc/AudienceView-demo>.



(a) An example of the dashboard view shown for each video.

## Themes

The comments cover a wide range of themes, indicating that the video content might be complex or touch on various issues. Here are the major themes identified from the comments:



(b) GPT-generated themes at the channel level.

Figure 1: Selected views of the app interface: a video-level dashboard of comments, and channel-level generated themes.

detection algorithms have been developed to allow more automated sensemaking from textual data. In addition to classic methods like latent Dirichlet allocation [5], more recent neural network-based techniques include top2vec [1] and BERTopic [9], often based on the general idea of clustering dimension-reduced embeddings from a language model. While these methods are powerful, they require technical knowledge and are inaccessible to non-technical users. Generative AI tools like ChatGPT are more accessible, but have issues for journalistic sensemaking, especially a lack of integration with news and social sites and limited context-window lengths. We can help fill this gap, inspired by similar efforts in the domain of survey responses [4], with a journalism-focused tool that leverages these advanced NLP methods while also handling collection of audience comments and providing a user-friendly interface.

*Digital journalism.* The question of audience understanding has received much attention in journalism and journalism scholarship over the years [19]. Journalists’ ability to acquire audience feedback at all is partly a function of technology, and doing so has become much easier in the digital-media era. Social media in particular has become central to, and sped up, journalistic practice [11], and much work also argues that these changes amount to “ambient” journalism [10] or “context collapse” [13] that brings the audience more deeply into the process and inspires changes in journalists’ presentation of news. Reporters and editors now put much more effort into understanding the audience [23], especially via click-stream or traffic-tracking tools like Chartbeat [7]. Comments, while important, often receive less attention because of the unpleasant nature of abuse and trolls [19]. On the AI side, most work on AI in journalism has considered AI as a direct producer of news articles [20], or AI in journalism education [17]. We aim rather to use AI for sense-making, helping journalists understand audience views.

## 3 System Overview

AudienceView has both backend and frontend components, with the user-facing frontend implemented with the Streamlit framework [21]. In deployment, various aspects can be customized: the YouTube channel to report on, the OpenAI [15] (GPT-4 by default) and local language models [24] used to interpret the comments, the prompts given to downstream language models, and more. It is intended to require limited technical setup and provide an accessible interface for frontend users. At several points in the interface, we make sure to connect generated topics, suggestions, and themes back to the underlying comments (as shown in Figure 1b). This is a deliberate design choice intended to give the user more direct engagement with audience feedback, reducing the need to trust an AI system.

*Backend.* Before frontend use, the system has several backend processes which collect and process YouTube data. Running these steps before deployment is essential to reduce latency for the end user. In addition to collecting comments and videos, these processes calculate summary statistics for videos, create sentiment scores via a local transformer model<sup>2</sup> [2], and generate the topic clusters and video- and channel-level themes and suggestions discussed below.

### 3.1 Video tab

The video tab provides detailed information on each video, shown in Figure 1a. It features controls to select a video from the list of all videos in the channel, with options to sort the list chronologically, alphabetically and by measures of engagement. Once a video has been selected, the tab provides summary statistics about its comments, including their average sentiment, and a visualization of their distribution over time. To provide more detailed insight into the comments, there is also a word cloud of the most common terms they contain, color-coded by sentiment score.

<sup>2</sup>Hugging Face: cardiffnlp/twitter-roberta-base-sentiment-latest

*Comment themes.* The video tab also offers AI-generated themes detected in the video’s comments. These themes are generated by GPT-4 in response to a prompt that includes 100 randomly sampled comments and asks the model to “summarize the major themes of the comments, and please cite examples.” The frontend shows the comments the model cites for each theme in tooltips, allowing the user to examine both original content and high-level AI synthesis.

*Suggestions.* This tab also features GPT-4-generated suggestions for future content in light of the comments, intended as an aid to ideation grounded in audience feedback, rather than a to-do list or ready-made editorial agenda. These suggestions similarly use 100 randomly sampled comments, and ask “what kinds of new documentary content should [news org] work on?” As with themes, the prompt instructs the model to cite comments supporting its answers, and the frontend makes them easily accessible in tooltips. (Full prompts are available in the source code.)

### 3.2 Channel tab

The channel tab provides information on the overall YouTube channel, separate from any particular video. The top of this tab features summary statistics like those at the top of the video tab – among others, the most recent date YouTube data was collected, the number of views, the number of comments, and the average sentiment score of the comments.

*Comment themes and suggestions.* As on the video tab, the channel tab includes GPT-generated themes and suggestions, using the same prompt and random sampling approach to comments. The comments in this case, however, are sampled from the entire channel rather than only one video, providing a higher-level view of audience interests and feedback.

Topic	Share of Comments	Avg. Negativity	Sentiment Diversity
Commentary on documentary content and timestamps.	7.1%	0.47	✓✓✓
Concerns about aviation safety and industry practices	2.5%	0.62	✓
Child welfare and family struggles	11.4%	0.61	✓✓✓
Coronavirus pandemic	3.8%	0.68	✓
Concerns about meat consumption and production practices	2.4%	0.54	✓✓
Environmental issues and climate change	3.2%	0.62	✓
Drug addiction and substance abuse	2.4%	0.62	✓✓
End-of-life care and experiences with the healthcare system.	3.4%	0.62	✓✓
Criticism of police conduct and performance	4.1%	0.73	✓
Discussions on racism and racial inequality	2.3%	0.68	✓
Criticism of financial systems and corporate greed	10.8%	0.61	✓✓
Israeli-Palestinian conflict	2.2%	0.68	✓
Conflict involving Russia and Ukraine	7.1%	0.64	✓✓
Conflict and opinions on the U.S. involvement in the Middle East	8.2%	0.69	✓
Documentaries	8.5%	0.45	✓✓✓
Criticism of PBS and allegations of bias	5.1%	0.62	✓✓✓
PBS Frontline documentaries and reporting	6.0%	0.37	✓✓✓
Criticism of China and the Chinese Communist Party (CCP)	3.3%	0.63	✓✓
Criticism of Joe Biden and his presidency	2.5%	0.57	✓✓✓
Opinions on Trump and Obama	3.7%	0.58	✓✓✓

Figure 2: The topics detected in channel-wide comments.

*Comment topics.* In addition to themes, we provide an alternative way to explore the topics contained in the comments, using a pipeline that includes all comments rather than only a random sample. These topics are detected by applying HDBSCAN [6, 14] clustering to sentence embeddings from the all-mpnet-base-v2 sentence-transformers model [18], after dimensionality reduction with UMAP [3]. GPT-4 is used to label the detected clusters based on a random sample of the comments in each. The frontend displays them in a table (Figure 2), with information on the percentage of all comments each cluster accounts for, measures of the variance of sentiment in each cluster, and the ability to browse the constituent comments (not shown in the figure).

*Change alerts.* AudienceView also flags videos which are experiencing unusually high or low rates of commenting, unusually positive or negative sentiment, or high numbers of comments requesting an updated version of the video. These help journalists detect changes in audience interest quickly. As baselines to compare actual levels to, we use an exponential smoothing model for comments and a simple weighted average by month for sentiment. Requests for updated videos are rare (but informative) enough that a baseline of each video receiving 0 of them is useful.

*“Superfans”.* The “superfan” commenters section indicates the commenters with the highest (i.e., most positive) average sentiment score across all included comments. To avoid bias from commenters with few posts, only those commenters with at least 200 comments are included.

## 4 Discussion and Evaluation

The Frontline team at PBS has provided helpful feedback throughout the process of developing AudienceView. We have recently begun conducting structured user interviews; the one interview completed so far had a positive reaction and pointed particularly to the topic detection and change alerts as helpful features. We aim to develop the tool further, focusing especially on two priorities: first, continued refinement of the interface with existing users, under current IRB approval; second, with future approval, more systematic measurement of its efficacy in newsroom workflows. In the future, expansion to modalities or comment sources other than YouTube would also expand the set of potential users.

### Acknowledgments

We gratefully acknowledge both assistance from the PBS Frontline team and the helpful feedback of our colleagues at the MIT Center for Constructive Communication.

### References

- [1] Dimo Angelov. 2020. Top2Vec: Distributed Representations of Topics. (Aug. 2020). arXiv:2008.09470
- [2] Francesco Barbieri, Jose Camacho-Collados, Luis Espinosa Anke, and Leonardo Neves. 2020. TweetEval: Unified Benchmark and Comparative Evaluation for Tweet Classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics, Online, 1644–1650. <https://doi.org/10.18653/v1/2020.findings-emnlp.148>
- [3] Etienne Becht, Leland McInnes, John Healy, Charles-Antoine Dutertre, Immanuel W H Kwok, Lai Guan Ng, Florent Ginhoux, and Evan W Newell. 2019. Dimensionality Reduction for Visualizing Single-Cell Data Using UMAP. *Nature Biotechnology* 37, 1 (Jan. 2019), 38–44. <https://doi.org/10.1038/nbt.4314>
- [4] Doug Beeferman and Nabeel Gillani. 2023. FeedbackMap: A Tool for Making Sense of Open-ended Survey Responses. In *Computer Supported Cooperative Work*

- and *Social Computing*. ACM, Minneapolis MN USA, 395–397. <https://doi.org/10.1145/3584931.3607496>
- [5] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022.
- [6] Ricardo J. G. B. Campello, Davoud Moulavi, and Joerg Sander. 2013. Density-Based Clustering Based on Hierarchical Density Estimates. In *Advances in Knowledge Discovery and Data Mining*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, Jian Pei, Vincent S. Tseng, Longbing Cao, Hiroshi Motoda, and Guandong Xu (Eds.). Vol. 7819. Springer Berlin Heidelberg, Berlin, Heidelberg, 160–172. [https://doi.org/10.1007/978-3-642-37456-2\\_14](https://doi.org/10.1007/978-3-642-37456-2_14)
- [7] Chartbeat, Inc. 2024. Chartbeat. <https://chartbeat.com/>
- [8] Ariel Goldman, Cindy Espinosa, Shivani Patel, Francesca Cavuoti, Jade Chen, Alexandra Cheng, Sabrina Meng, Aditi Patil, Lydia B Chilton, and Sarah Morrison-Smith. 2022. QuAD: Deep-Learning Assisted Qualitative Data Analysis with Affinity Diagrams. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. ACM, New Orleans LA USA, 1–7. <https://doi.org/10.1145/3491101.3519863>
- [9] Maarten Grootendorst. 2022. BERTopic: Neural Topic Modeling with a Class-Based TF-IDF Procedure. (March 2022). [arXiv:2203.05794](https://arxiv.org/abs/2203.05794)
- [10] Alfred Hermida. 2012. Social Journalism: Exploring How Social Media Is Shaping Journalism. In *The Handbook of Global Online Journalism* (1 ed.), Eugenia Siapera and Andreas Veglis (Eds.). Wiley, Chichester, UK, 309–328. <https://doi.org/10.1002/9781118313978.ch17>
- [11] Regina Lawrence. 2020. Campaign News in the Time of Twitter. In *Controlling the Message*, V. Farrar-Myers and J. Vaughn (Eds.). NYU Press, New York, 93–112. <https://doi.org/10.18574/nyu/9781479886357.003.0005>
- [12] Lumivero. 2023. NVivo. <https://lumivero.com/products/nvivo/>
- [13] Alice E. Marwick and danah boyd. 2011. I Tweet Honestly, I Tweet Passionately: Twitter Users, Context Collapse, and the Imagined Audience. *New Media & Society* 13, 1 (Feb. 2011), 114–133. <https://doi.org/10.1177/1461444810365313>
- [14] Leland McInnes, John Healy, and Steve Astels. 2017. HdbSCAN: Hierarchical Density Based Clustering. *JOSS* 2, 11 (2017), 205. <https://doi.org/10.21105/joss.00205>
- [15] OpenAI. 2024. OpenAI API. <https://platform.openai.com/>
- [16] Cassandra Overney, Belén Saldías, Dimitra Dimitrakopoulou, and Deb Roy. 2024. SenseMate: An Accessible and Beginner-Friendly Human-AI Platform for Qualitative Data Analysis. In *Proceedings of IUI '24*. ACM, Greenville, 922–939. <https://doi.org/10.1145/3640543.3645194>
- [17] John V. Pavlik. 2023. Collaborating With ChatGPT: Considering the Implications of Generative Artificial Intelligence for Journalism and Media Education. *Journalism & Mass Communication Educator* 78, 1 (2023), 84–93. <https://doi.org/10.1177/10776958221149577>
- [18] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings Using Siamese BERT-Networks. 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)*. Association for Computational Linguistics, Hong Kong, China, 3980–3990. <https://doi.org/10.18653/v1/d19-1410>
- [19] James G. Robinson. 2019. The Audience in the Mind’s Eye: How Journalists Imagine Their Readers. *Columbia J. Rev.* (2019). <https://doi.org/10.7916/d8-drvj-wj06>
- [20] Jonathan Stray. 2019. Making Artificial Intelligence Work for Investigative Journalism. *Digital Journalism* 7, 8 (2019), 1076–1097. <https://doi.org/10.1080/21670811.2019.1630289>
- [21] Streamlit. 2021. Streamlit – A Faster Way to Build and Share Data Apps. <https://streamlit.io/>
- [22] VERBI Software. 2022. MAXQDA. <https://www.maxqda.com/>
- [23] Kelsey Whipple and Jeremy Shermak. 2018. Quality, Quantity and Policy: How Newspaper Journalists Use Digital Metrics to Evaluate Their Performance and Their Papers’ Strategies. *ISOJ Journal* 8, 1 (2018), 67–88.
- [24] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick Von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-Art Natural Language Processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics, Online, 38–45. <https://doi.org/10.18653/v1/2020.emnlp-demos.6>