

CASE STUDY

Newsroom: Michigan Radio

About the newsroom

Michigan Radio WUOM-FM is the NPR affiliate at the University of Michigan with a news team covering the entire state. Prior to AP's involvement, the station developed a newsgathering application called Minutes that scrapes city council meetings posted on popular video sharing services and transcribes them, allowing Michigan Radio reporters more visibility into public meetings they are unable to attend in person. The cloud-based application was developed in house with funding from the Google News Initiative.

Project goals

The goals were to build improvements to the existing Minutes application, specifically by adding new features to summarize city council meeting transcripts, and to generate alerts from those transcripts. In doing so, they hoped to make city council meetings more accessible and decrease the amount of time the newsroom spent on attending meetings in person.

Why is this project important to the newsroom?

The proposed new features, summarization and alerting, would help improve awareness of what transpired in a city council meeting where a reporter couldn't attend in person. Summarization would eliminate the need to read through the entire transcript of hourslong meetings. Alerting would give proactive notice to journalists of a keyword or topic that matches a reporter's interests.

Engineering process

As this was an existing product, the Northwestern University Medill Knight Lab developers began with a comprehensive examination of the code already deployed to production. The developers noted performance issues with the existing Google Cloud speech-to-text transcription service. Using a tool that calculates "word error rates," a statistical measurement of how different the words in a sample text are to a reference text, Michigan Radio's stakeholders agreed

that an overhaul was needed. (See the various tools compared in the Appendix.) The developers settled on OpenAI's Whisper's small.en model for its balance of transcription speed and accuracy.

Instead of pursuing summarization as originally planned, most of the development period was spent on rebuilding the core transcription service with OpenAI's Whisper. Concurrently, as part of a class, students at the University of Michigan built out the alerting feature.

Were the goals met?

One of the goals changed along the way. After the developers noticed the transcript quality issues, most of the development time was allocated to replacing the transcription system. The updated Minutes application sees significant improvement in the quality of transcripts, lower costs of transcription, along with the ability to pull more frequent updates from city councils.

The second goal was met with the development of an alerting system. Reporters were able to choose keywords and get alerts whenever those words showed up in the transcript of a meeting.

Major challenges

The original developer of the Minutes app at Michigan Radio moved to a new role at the University of Michigan during the project period. The Northwestern team was able to complete development and deploy the updates on both transcription and alerting for Michigan Radio. However, this reflects a difficulty of the project— it requires a full-time developer on staff in order to manage the web application.

Because city council meetings come from a variety of websites which each have their own systems for uploading and labeling videos, it was difficult to consolidate all of the meetings onto one application smoothly. The source material for the city council meetings was not standardized among different city websites, and so it was difficult to maintain and debug all of the issues for each website that uploaded city council meetings.

Future work

The program successfully alerts reporters to keywords said in meetings, but more work needs to be done to fully automate the process, debug the system, and make the transcription more consistent for the scale of meetings being inputted.

STAKEHOLDER REACTION

“Minutes is a great tool to keep up with the local government meetings I can’t go to in person. I use it to skim transcripts to get a general idea of the meeting, without having to spend hours sitting through the whole thing. And the new keyword alert system allows me to track issues and find sources in communities that I normally don’t cover.”

DUSTIN DWYER
REPORTER AND PRODUCER

And because so much time was spent reworking the transcription system, the project was unable to work on the summarization of city council meetings. In the future, the newsroom could work on being able to not only get alerts from keywords found in transcripts, but to also extract information from transcripts such as summaries, automating a process that reporters originally did.

Development team

The original project was developed in-house by Michigan Radio. For this project, work was distributed between Michigan Radio and Northwestern University Medill Knight Lab under the guidance of Professor Jeremy Gilbert. Additional development work was performed by University of Michigan students of SI-699 “User-Centered Agile Development” under the guidance of Professor Mark Newman.

For Michigan Radio: Brad Gowland, Dustin Dwyer

For Knight Lab: Jeremy Gilbert, Scott Bradley, Joe Germuska

For University of Michigan: Mark Newman

For The Associated Press: Aimee Rinehart, Ernest Kung

Core components of the system

Input: Government meeting videos posted online

Code: Python, Javascript

Database: Redis

Integrations: YouTube via API, Vimeo via API, OpenAI Whisper (open-source package)

Output: Web portal, Email

Hosting: Google Cloud Platform

As the code from this project is not being open sourced due to commercial sensitivities, the developers provided these additional notes:

The Minutes project is driven by a suite of configurations for monitoring civic media feeds. Newly discovered media are injected into a pipeline of queued processing stages that register the media with the system, then download and transcribe the media. Transcripts are subsequently written to a search index that is used by the SearchMinutes website.

The registry and queues in the Minutes project are handled by the popular Redis database. The system is written in Python, and OpenAI’s Whisper transcription model provides the transcripts.

Users can search the transcripts on searchminutes.com and sign up to receive email alerts whenever certain terms appear in the transcripts. The system's alert task runs every hour on new transcripts added to the database. Users can choose what hour they want the email alert to be sent and whether they want to receive alerts every day or only weekly. The email shows the section of the transcript where the key word appeared and links to the timestamp in the video when it was mentioned.

Appendix

Figure 1. Word error rates as calculated by the **JiWER** Python package. Recordings of complete city council meetings from four Michigan municipalities were used for testing. A reference transcript was generated by humans for each meeting using the Happy Scribe service. The **JiWER** tool then calculated the statistical difference in the words between the reference transcript and the transcript generated by the various AI-powered services. A lower percentage means a more accurate transcript.

Service	Meeting 1	Meeting 2	Meeting 3	Meeting 4
YouTube	22.35%	296.05%	17.45%	12.33%
Google Cloud	67.90%	51.51%	60.13%	42.56%
Whisper base	19.17%	14.31%	14.99%	12.71%
Whisper small	19.55%	12.49%	11.90%	11.38%
Whisper medium	15.48%	12.29%	10.03%	8.20%
Trint	25.72%	17.67%	20.36%	12.38%

Figure 2. Process map

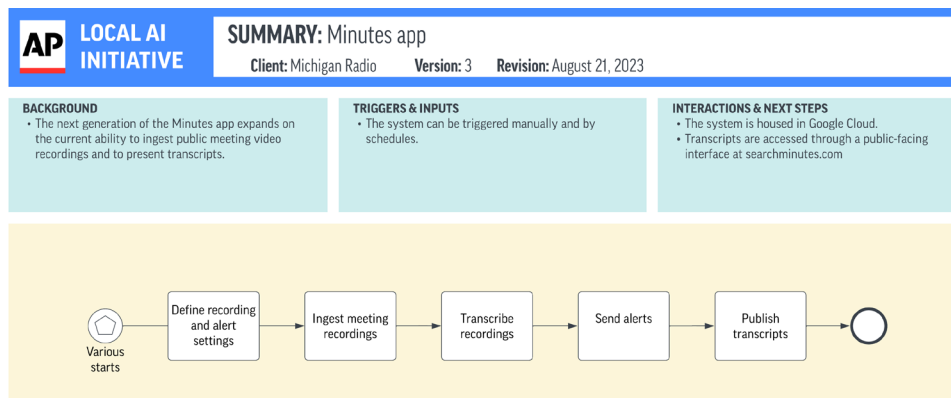


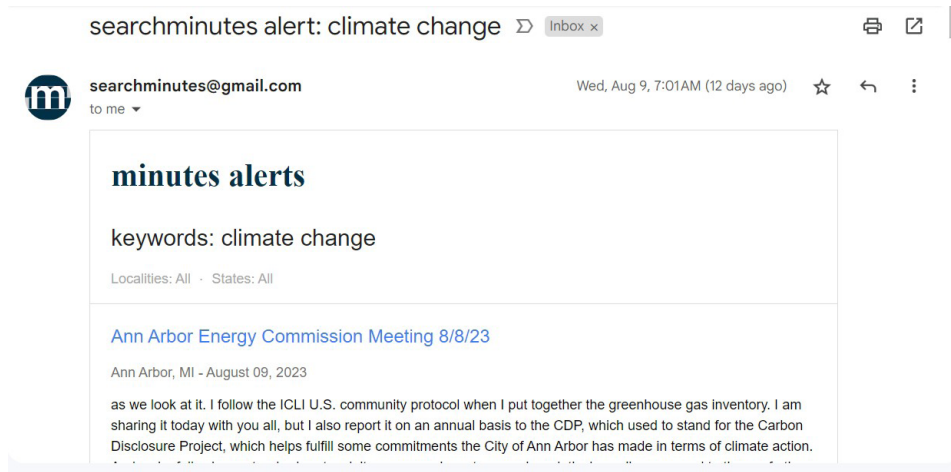
Figure 3. Transcript view

TIMESTAMP	TRANSCRIPT
00:00:00	you you you
00:01:30	you you So I go home and shave okay, we're ready gentleman. He was a super I Call to order the regular meeting of the Mount Pleasant City Commission Monday, August 14th 2023 at 7 p.m. Please rise for the pledge
00:02:50	The pledge of allegiance to the flag of the United States of America and to the Republic for which it stands, one nation, under God, indivisible, with liberty and justice for all. The Mount Pleasant City Commission recognizes the Anishinaabe whose ancestral lands the city stands upon today. We acknowledge the Saginaw Chippewa Indian tribe and recognize that the three bands of Ojibwe, Saginaw, Black River, and Swan Creek have resided on this land for over 200 years. Our acknowledgement of the full history of the tribe and our commitment to be good stewards to these ancestral lands. Our expression of gratitude and appreciation to those ancestors who came

Figure 4. Keyword searches

Figure 5. Keyword-based alerts

Figure 6. Alerts by email





THE ASSOCIATED PRESS

For more information, go to ap.org/ai