

Sujet traité : Il y a un nouveau podcast qui va vous épater / There's a New Hit Podcast That Will Blow Your Mind


Source : Wall Street Journal Date : 5 octobre 2024

## There's a New Hit Podcast That Will Blow Your Mind

K

WSJ [wsj.com/tech/ai/google-notebooklm-ai-podcast-deep-dive-audio-c30a06b3](https://www.wsj.com/tech/ai/google-notebooklm-ai-podcast-deep-dive-audio-c30a06b3)

October 5, 2024



By Ben Cohen

Have you heard about the latest hit podcast? It's called Deep Dive—and you have to check it out.

Each show is a chatty, 10-minute conversation about, well, any topic you could possibly imagine. The hosts are just geniuses. It's like they know everything about everything. Their voices are soothing. Their banter is charming. They sound like the kind of people you want to hang out with.

But you can't. As it turns out, these podcast hosts aren't real people. Their voices are entirely AI-generated—and so is everything they say.

And I can't stop listening to them.

Welcome to a Deep Dive into the thrilling world of Formula One. And let me tell you—it's way more than just fast cars and Champagne.

This experimental audio feature released last month by Google is not just some toy or another tantalizing piece of technology with approximately zero practical value.

It's one of the most compelling and completely flabbergasting demonstrations of AI's potential yet.

"A lot of the feedback we get from users and businesses for AI products is basically: That's cool, but is it useful, and is it easy to use?" said Kelly Schaefer, a product director in Google Labs.

This one is definitely cool, but it's also useful and easy to use. All you need to do is drag a file, drop a link or dump text into a free tool called NotebookLM, which can take any chunk of information and make it an entertaining, accessible conversation.

Google calls it an "audio overview." You would just call it a podcast.

One of the coolest, most useful parts is that it makes podcasts out of stuff that nobody would ever confuse for scintillating podcast material.

Wikipedia pages. YouTube clips. Random PDFs. Your college thesis. Your notes from that business meeting last month. Your grandmother's lasagna recipe. Your resume. Your credit-card bill! This week, I listened to an entire podcast about my 401(k).

And then I found myself listening to an oddly captivating Deep Dive into that day's edition of the Federal Register.

"Oh, the Federal Register—people hear that and think: *boring*," one AI host says.

“Right! Dusty old government documents,” the other robo-host replies. “But I actually think they’re kind of interesting. It’s like a little peek behind the curtain to see how things really work.”

So, for anyone listening who might have thought: The Federal Register—that’s not for me. Hopefully, we’ve changed your mind a little bit today.

Each conversation is between the same male and female voices—Google says they don’t have names—and it takes only a few minutes to create one. If you didn’t know anything about it, you wouldn’t guess it was automatically generated. It doesn’t sound like other AI slop. It just sounds like any other podcast. And once you start playing with it, you’ll probably become obsessed with it.

It has become increasingly popular across Silicon Valley and among the people who are smartest about artificial intelligence, like Andrej Karpathy, a deep-learning expert who co-founded OpenAI and led Tesla’s team working on computer vision for Autopilot. Lately, he’s had another AI project on his mind—and in his ears.

“Deep Dive is now my favorite podcast,” Karpathy wrote on X. “The more I listen, the more I feel like I’m becoming friends with the hosts, and I think this is the first time I’ve actually viscerally liked an AI. Two AIs! They are fun, engaging, thoughtful, open-minded, curious.”

He likes them so much that he’s listened to Deep Dives on the philosophy of Ludwig Wittgenstein, the scientific process of oxidative phosphorylation, Mars, gold, Arnold Schwarzenegger and pomegranates. He even generated a podcast series called Histories of Mysteries—10 episodes based on Wikipedia pages ranging from Atlantis to the Bronze Age.

The first time I heard a Deep Dive, it reminded me of an AI version of Acquired, the hit podcast about business history. When I told the hosts, they tried it out for themselves, opening NotebookLM and uploading a Google Doc with links to some of their sources for a recent show on Microsoft. Not their notes—just a bunch of links. They found the results to be mind-blowing.

“I don’t know whether to be amazed or terrified,” Acquired co-host David Rosenthal said.

Exactly. There have only been two times when I had my mind blown by AI. The first was my introduction to ChatGPT. This was the second.

It’s especially notable that it was made by Google. The search titan has been scrambling to catch up in the AI race since the release of ChatGPT, and this bonkers audio feature is the company’s first truly original and viral product in that time.

I had lots of questions about it—and for answers, I turned to the real people behind the AI voices.

They work in Google Labs, a division within the tech giant that incubates and builds AI-powered products like NotebookLM—LM for “language model.” When it rolled out last year, Google called it a “virtual research assistant,” a personalized AI to summarize documents, simplify complex ideas and brainstorm connections in the source material. They began thinking about how to expand those capabilities and kept coming back to one question of their own.

“How do we bring the utility and magic of AI to people with as little work as possible?” said Raiza Martin, the Google Labs product manager for NotebookLM.

For most people, the utility and magic of AI is feeding a blob of text to a chatbot with basic instructions: Make this into something else.

*Here are the bullet points from a strategy meeting. Now write me a catchy marketing slogan.*

What the Google Labs team realized was that a blob of text could be transformed into anything—even a podcast.

There are many reasons you might want to turn information into an audio conversation. Productivity. Creativity. Procrastination. Inspiration. Education. Maybe you’re an auditory learner and prefer listening to reading. Maybe you find podcasts more engaging than pitch decks. Maybe you want the same information in a different way to spark fresh ideas.

Here’s how it works. When you upload your source material, NotebookLM instantly digests it and spits out a basic written summary. Under the hood, a model is writing and constantly editing a script for the conversation “based on the goal of being entertaining—and, crucially, bringing out the insights,” says Steven Johnson, the editorial director of Google Labs.

That’s the real breakthrough. It doesn’t merely summarize the source material. The product is specifically designed to focus on the most interesting and surprising parts—and it does.

But what the AI hosts say is just as important as how they say it: like humans.

It’s like they’re finding the sweet spot on the human-robot spectrum for audio. Right!  
And speaking of voices...

This audio feature wouldn’t exist without Google’s recently developed voice technology, but it wouldn’t be enjoyable if not for an important design choice: To make the conversations sound better, Google had to make them rougher.

“It’s a very counterintuitive thing,” Johnson said. “But if you had two perfect scripts talking to each other in complete sentences, nobody would listen to it. It would just sound too robotic.”

Instead, the robotic voices speak like actual people. They break up their speech with “like,” “um” and “you know.” They stammer. They pause. They introduce points with “OK, so the thing is.” They reinforce each other’s points with “totally” and “oh, 100%.” They sound so much like real podcasters that I kept waiting for one of them to read a FanDuel ad.

Before he joined Google Labs two years ago, Johnson was a bestselling author, and he was impressed by the quality of the Deep Dive conversations when he uploaded his books about innovation and technology into NotebookLM. I felt the same way when I tried it with my own work.

So I asked someone else to participate in a similar experiment: Nobel Prize-winning economist Claudia Goldin.

She was game. “That would be fun,” she told me. “I often think of my papers as podcasts.” I sent over Deep Dives about her groundbreaking research, her latest book and even her Nobel lecture. She quickly sent back a review: “Wow!” She loved the AI voices, their summaries and one pithy line, which the hosts riffed as if they were real people.

“Goldin would say, ‘Don’t try to boil the ocean,’ ” the female AI said.

“Where did *that* come from?” the real Goldin asked.

There’s a moment like that in all of the podcasts I’ve generated. And every time I hear one, I want to make another one.

When I saw a Googler tweet about a Deep Dive on the history of potatoes, I decided to listen to a Deep Dive on the history of Mr. Potato Head.

“You might think you know this spud—but trust me, there’s more to him than meets the eye.”

“Let’s peel back the layers.”

You might think you know this spud—but trust me, there’s more to him than meets the eye. Oh, absolutely. We’re going way beyond plastic parts and goofy grins to uncover how Mr. Potato Head reflects bigger trends.

Anyone who has used AI knows that it’s not perfect. These voices sometimes mispronounce names or misunderstand facts. They occasionally make strange noises. You won’t hear them hallucinate and start yammering about nonsense. But you also won’t hear novel ideas or anything genuinely hilarious.

“I suspect it’s because humor is in many ways the opposite of translation and summarization,” Johnson explained on his blog. “Humor is all about surprise, about defying expectations, going off script in just the right way.”

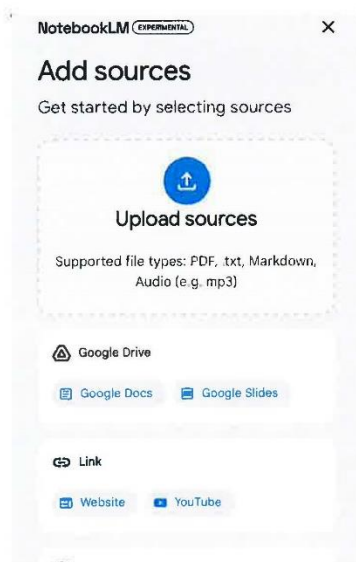
Or at least that's what he thought. Since then, he's come across a few conversations that made him crack up, and he now thinks the hosts can be funny with the right steering.

### How to make your own NotebookLM podcast



*Nous avons expérimenté cet outil. Vraiment fascinant.*

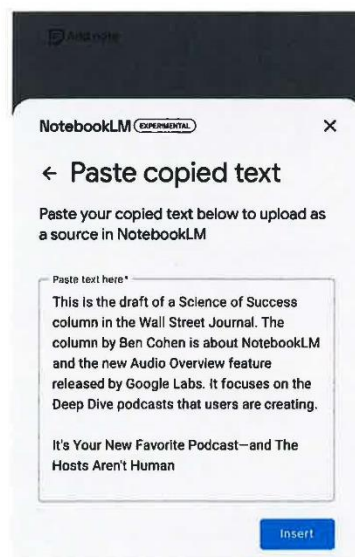
Click on the plus sign to start a new notebook.



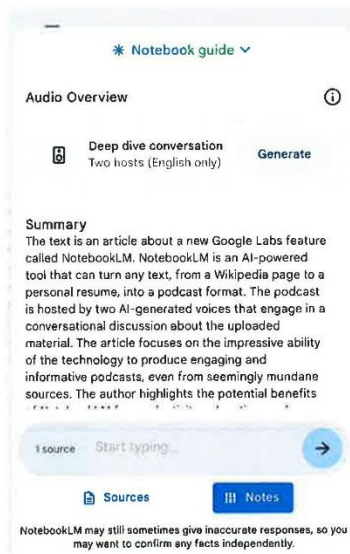
Upload your source material. You can drag or drop a file, attach a Google Doc, link to websites and YouTube videos or paste copied text. You can upload up to 50 sources.



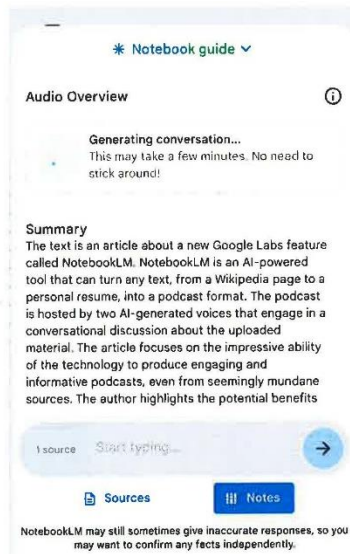
I chose the text option.



I uploaded a draft of this column with a brief introduction.



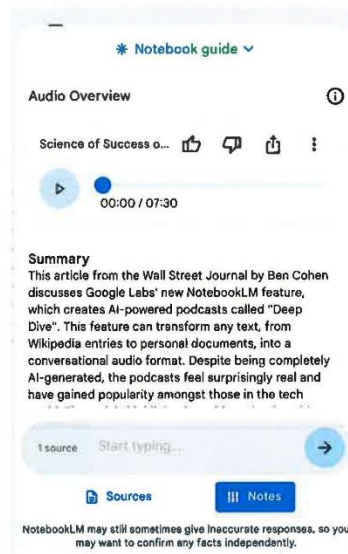
After a few seconds, NotebookLM generated a written summary. To listen to a Deep Dive conversation, click the "Generate" button.



Wait for the model to generate the conversation. This usually takes a few minutes.

2-3 min.





That's it. Now you can listen to the Deep Dive.

As it happens, the product evolving in exciting and unexpected directions is one of the Google Labs team's two metrics for success.

"Did people use it for what we thought they would use it for?" Schaefer said. "And maybe even more important: Did they identify, like, a thousand other ways to use it that we never would have imagined?"

I wanted to hear more about that idea—so I opened NotebookLM, uploaded a transcript of the interview and listened to my new favorite AIs offer their takes on another episode of Deep Dive.