

New in Hedy: Speaker Names, Meeting Detection, and Document Exports

Hedy now puts real names to the voices in your transcript and carries them into your recap and notes, spots meetings starting on Mac and Windows, and saves sessions as Word or PDF.

Published by Julian Pscheid · August 14, 2026

[Read this article online: https://www.hedy.ai/post/speaker-names-in-transcripts/](https://www.hedy.ai/post/speaker-names-in-transcripts/)



Overhead view of four colleagues mid-conversation around a round wooden table with coffee mugs, an open notebook, and a face-down iPhone

Quick answer Hedy's recent releases focus on one thing: knowing who said what, and doing something useful with it. Transcripts now carry real names instead of "Speaker 1" wherever the conversation gives Hedy something to go on, and those names feed the summary, detailed notes, and to-dos it writes. On Mac and Windows, Hedy notices a call starting in Zoom, Teams, or Webex and offers to start a session. And a session can leave Hedy as a Word document or PDF. Free for 5 hours/month, \$12.99/month for unlimited .

A transcript that says "Speaker 2 committed to the deadline" is only half useful. You still have to work out who Speaker 2 was, and if your summary and your notes disagree about it, you end up not trusting either one.

That's the problem the last few Hedy releases went after.

Hedy works out who is talking

Hedy already separated voices in a conversation. What's new is that it now works out the names behind them.

It reads the conversation for actual evidence: someone introducing themselves, or another participant addressing them by name. Your calendar invite and any Topic context (/help/understanding-hedy-topics-organize-your-sessions-unlock-deeper-insights/) you've set help it match the right person to the right voice. When the evidence isn't there, Hedy leaves that speaker numbered instead of guessing, so a session where two people are named and one isn't is normal and expected.

The part that matters most is what happens next. Whatever Hedy works out is applied once and then used everywhere it writes: the transcript, the recap, the detailed notes, and your to-dos all draw on the same set of names, rather than each making its own call about who was who.

You stay in charge of it. Sessions with speaker turns show a row of speaker chips above the transcript (/help/understanding-hedy-transcripts/) , and tapping one lets you type the right name. A name you set is never overwritten by Hedy's inference, and it follows the session into copies, shares, downloads, and your other devices. If you rename someone after your recap and notes are already written, those keep the old name until you refresh them, and Hedy prompts you to do exactly that. Renaming works on your own finished transcripts; a session shared with you, one still running, or the cleaned view are read-only.

Speaker names come from Hedy's on-device Nemotron engine, which is what separates the voices in the first place. If you're on Whisper or a cloud speech provider, switch under Settings !' Speech & AI and your next session will carry names. It runs on Apple Silicon Macs, iPhone 12 and newer, iPad Air 4 and newer, Windows, and 64-bit Android.

Hedy notices the meeting starting

On Mac and Windows, Hedy watches for a call beginning in Zoom, Microsoft Teams, or Webex, plus FaceTime on the Mac, and offers to start a session for it.

It's a small thing that removes the most common way people lose a session: joining the call, getting straight into it, and remembering Hedy twenty minutes later. You get the offer, you accept, and Hedy starts capturing.

Hedy stays bounded either way. It's watching for one of those apps to pick up your microphone, not listening to anything, and nothing is captured until you say yes.

On iPhone and Android there's a calendar version of the same idea: connect Google Calendar, turn on Auto-detect meetings, and Hedy offers to start a session as a scheduled meeting begins.

Take a session anywhere

You can now save a session out of Hedy as a real document. Open one and choose Download Summary , Download Detailed Notes , or Download All Data , then pick your format: plain text, Markdown, Word, or PDF.

Word and PDF keep your headings, lists, quotes, and tables rather than flattening everything into one long block of text, so a recap you send to a client looks like a document instead of a text dump. PDF covers the characters in Hedy's bundled font, and for anything outside it Hedy points you to Word instead.

Two things worth knowing about what lands in each file. Download All Data gathers the recap, detailed notes, your session notes, highlights, chat history, and to-dos, but not the transcript. The transcript has its own Download Transcript action and saves as a text file.

If you've ever hesitated about putting your meetings in a tool because you weren't sure you could get them back out, this is the answer.

Every action item from a project, together

Sessions that belong to a Topic (/help/understanding-hedy-topics-organize-your-sessions-unlock-deeper-insights/) now share a To-Dos tab. Open the Topic and you see every action item from every session in it, sorted by due date or by the session it came from.

You can edit and complete them right there, and hide the completed ones once a project gets long. Completing an item there also completes it in its original session, because there's only one copy being shown two ways. A Topic that someone shared with you is read-only.

Before this you had to open sessions one at a time to reconstruct what you owed people. For anything running longer than a couple of meetings, that was the tedious part.

Also worth knowing

A few smaller things landed alongside those:

- Decide about audio before you start. On Mac, Windows, iPhone, iPad, and Android, a Save Session Audio switch now sits in the start sheet, so you choose per session whether audio is kept instead of digging through Settings. As always, get consent from the people you're with before saving audio.
- Session merges keep everything. Combining two sessions now carries the chat history and session notes across, not just the transcript.
- More languages and newer models. Deepgram adds Swiss German, Armenian, Gujarati, Nepali, and Punjabi, plus more regional variants. New OpenAI sessions run on its current speech models, and any model you already had selected keeps working. Both providers use your own API key.
- Adjustable suggestions on Mac. Suggestion popups now carry their own menu for hiding alerts, changing how often they arrive, and choosing whether they follow your cursor.
- Background imports on iPhone and iPad. On iOS and iPadOS 26, importing audio or video keeps running when you switch apps or lock your screen.

Getting the update

Hedy updates itself on Mac and Windows. On iPhone, iPad, and Android, grab it from the App Store or Google Play. Everything above is in the current release.

If you've been using Whisper, the one thing worth doing today is switching to Nemotron under Settings !' Speech & AI . It's what unlocks the speaker names, and we're working toward making it the default so nobody has to think about it.

Hedy AI · Live AI Coaching for Important Conversations

Try Hedy free: <https://www.hedy.ai/downloads/>

<https://www.hedy.ai/post/speaker-names-in-transcripts/>