

OPERATOR FIELD GUIDE NO. 01

FREE — NO SIGN-UP WALL

The Meeting Coach

Build an AI that grades your meetings and tells you how to fix them

A coach that reads the transcript of any meeting you were in and tells you how much you talked, how many questions you actually asked, how long your longest monologue ran, and the three things to do differently next time. With the exact line you said, and the better version of it.

20 minutes

TIME TO BUILD

Free tier

COST

No code

SKILL LEVEL

A transcript

ALL YOU NEED

WHY THIS ONE FIRST

Nobody tells you that you talk too much

Most professional feedback stops arriving somewhere around your first management job. People stop telling you that you talk too much, that you open with a disclaimer instead of a claim, that you asked one question in fifty minutes. Not because you fixed it. Because it got awkward to say.

A transcript doesn't find it awkward.

This came out of our bonus episode, where Ceci described the meeting coach she built for herself. Her framing was blunt:

"It's probably like a side effect of my insecurity."

CECILIA MANDUCA — EP 11

Then she described what it does, and Bea's reaction on tape was *"This is brilliant. Another one that I will copy."*

THE DETAIL THAT MAKES IT WORK

It reads the transcript, not the summary. Summaries are written to be pleasant. Transcripts contain the sentence where you cut someone off.

WHAT YOU NEED

Three things, and you almost certainly have all of them

YOU NEED	OPTIONS THAT WORK
A transcript source	Any notetaker that saves transcripts — Granola, Otter, Fathom, Fireflies, tl;dv, Zoom's own cloud transcript, Google Meet, Teams. Or paste a transcript in by hand.
An AI you can paste into	Claude or ChatGPT. Free tier is fine to start.
A calendar (optional)	Only needed for the automated version in Step 4. Google Calendar, Outlook, whatever.

You do not need: a paid notetaker, a specific company, connectors, an ops team, or permission from anyone. If you can get a transcript of a meeting you were in, you can run this today.

ONE THING TO CHECK FIRST

Recording and transcribing meetings has legal and cultural rules that vary by country and company — in some places every participant must consent. Check your own policy before you point anything at a call. If recording is off the table, this still works on a meeting you transcribe yourself, or on your half of the conversation from your own notes.

PICK YOUR MODEL

Do not run this on the cheap tier

The single most common mistake with this workflow is running it on a cheap, fast model. Transcripts are long. A one-hour meeting is 8,000 to 12,000 words. Small models lose the middle of long documents, and the middle is where the useful evidence lives.

WHAT YOU'RE DOING	ON CLAUDE	ON CHATGPT
Building the skill One-off, needs to get the structure right	Fable 5	GPT-5.6 Sol
Running it on each meeting Routine, repeated	Sonnet 5	GPT-5.6 Terra
Reviewing your setup for security See Step 6	Fable 5, high effort	GPT-5.6 Sol
Do not use for this	Haiku	Luna

Why Luna and Haiku are the wrong call here: both are built for cheap, high-volume, short tasks. On long-context recall — which is exactly what reading a full transcript is — Luna scores around 41% against Sol's 91%. It will confidently quote things nobody said. For a workflow whose entire value is accurate evidence, that's disqualifying.

THE RULE, SO THIS SURVIVES THE MODEL NAMES CHANGING

Build with the strongest model you have access to, run with the balanced mid-tier, and never run long-document analysis on the cheap tier. Model names turn over every few months. That rule doesn't.

STEP 1

Get one transcript

Pick a meeting from the last week where you did most of the talking. Not your best meeting. That's the point.

Export the transcript as plain text. Almost every notetaker has a copy-transcript or download-transcript option. You want the raw dialogue with speaker labels, like this:

TRANSCRIPT FORMAT

Alex (00:04)
So I wanted to walk through the Q3 numbers before we get into the roadmap.

Priya (00:11)
Sure. Do you want me to pull the dashboard up?

Speaker labels matter — without them the coach can't tell your talking from theirs. If your tool doesn't label speakers, most will let you tag them once and remember.

STEP 2

Run the coach



This is the whole workflow in one prompt. Paste it into Claude or ChatGPT, fill in the bracketed parts, paste your transcript underneath.

It works standalone, in a plain chat window, with no setup at all. Step 4 turns it into something reusable.

THE PROMPT

You are an executive meeting coach. You are honest, specific and unsentimental.
You are not encouraging. Your job is to make me better, not to make me feel good.

Below is the transcript of a meeting I was in.

MY NAME IN THE TRANSCRIPT: [your name exactly as it appears]

MY ROLE: [e.g. Chief of Staff / Head of Finance / Founder]

WHAT THIS MEETING WAS FOR: [one line – e.g. "get the team to commit to a launch date" or "understand why churn moved"]

WHAT A GOOD OUTCOME WOULD HAVE BEEN: [one line]

Analyse ONLY what is in the transcript. Follow these rules absolutely:

1. Every single claim you make must be supported by a VERBATIM QUOTE from the transcript. If you cannot quote it, do not claim it.
2. Never invent, paraphrase or reconstruct a quote. Quote exactly or say "no evidence in transcript".
3. If the meeting was genuinely good, say so. Do not manufacture problems to seem useful. An honest "three things went well and one thing to watch" is a valid result.
4. Do not comment on my personality. Comment on behaviour I could change on Monday.

Produce exactly this, and nothing else:

THE NUMBERS

- My share of the words spoken: X% (and the rough split for everyone else)
- Questions I asked: X total – of those, X were genuinely open (invited thinking) and X were closed or rhetorical. List the open ones verbatim.
- My longest uninterrupted turn: approximately X words, on the topic of [topic]. Quote the first line of it.
- Times I interrupted someone mid-sentence: X. Quote each one.
- Times someone tried to speak and I continued: X. Quote each one.
- Decisions actually reached: list them. For each, say whether it has a named owner and a date, or does not.
- Open loops left at the end: anything raised and never resolved.

WHAT WORKED

Two or three things I did well, each with the quote that proves it.

THE VERDICT

Three sentences maximum. The single most important pattern in how I showed up. Lead with the claim. No preamble, no cushioning.

THREE FIXES

Exactly three, ordered by how much they'd change the outcome. For each:

- **The behaviour:** what I did, with the verbatim quote
- **Why it cost me:** tied to the stated purpose of this meeting
- **Say this instead:** an actual rewrite of my actual words. Not general advice. The specific sentence I should have said.

ONE THING TO TRY IN MY NEXT MEETING

A single, concrete, testable action. Something I can either do or not do, and know afterwards which it was.

TRANSCRIPT:

[paste transcript here]

What comes back

You'll get numbers you can't argue with and a verdict you probably can. Both are useful. The numbers are the part that changes behaviour — most people are genuinely surprised by their own talk share, and nobody guesses their longest monologue correctly.

RULE 1 IS THE LOAD-BEARING INSTRUCTION

Without *"every claim must have a verbatim quote,"* you get fluent, generic, plausible coaching that would apply to any meeting ever held. With it, you get evidence. If your output has no quotes in it, the model ignored the rule — say *"you made claims without quoting the transcript, redo it"* and it will comply.

Rule 3 matters more than it looks. An AI asked for criticism will always find criticism. Explicitly permitting "this was fine" is what makes the bad reports trustworthy.

STEP 3

Make it mean, or don't

Ceci's version is deliberately harsh:

"It's really mean sometimes, but I want it mean."

CECILIA MANDUCA — EP 11

That's a preference, not a requirement. Add one of these lines to the prompt:

DIAL IT UP

Be blunt to the point of uncomfortable. Do not soften anything. Assume I am experienced and can take it. Skip all praise unless it is genuinely notable.

DIAL IT DOWN

Be direct but constructive. Lead with what worked. Frame fixes as experiments rather than failures.

If you make it too harsh you will stop reading it, and a coach you avoid is worth nothing. If you make it too kind you'll read it forever and never change. Start neutral, run it four or five times, then adjust based on whether you're actually opening the reports.

STEP 4

Turn it into something reusable

Pasting a prompt works, but you won't keep doing it. Here's how to install it properly on each platform.

On Claude

OPTION A — A PROJECT (EASIEST, 5 MINUTES)

1. In Claude, create a new Project called `Meeting Coach`.
2. Open the Project's custom instructions and paste in the prompt from Step 2 — everything from *"You are an executive meeting coach"* down to but not including `TRANSCRIPT:`.
3. Replace the bracketed fields with your real name and role once, permanently.
4. From then on, every new chat in that Project only needs the transcript pasted in. Everything else is remembered.

OPTION B — A SKILL (BETTER, REUSABLE, INVOKABLE BY NAME)

1. Create a folder called `meeting-coach`.
2. Inside it create a file called `SKILL.md`.
3. Put this at the top of the file:

SKILL.MD FRONTMATTER

```
---
name: meeting-coach
description: Analyse a meeting transcript and score how the user showed up —
  talk share, questions asked, longest monologue, interruptions, decisions
  reached. Returns a verdict and three specific fixes with rewrites. Use when
  the user shares a meeting transcript, asks how a meeting went, asks for
  feedback on their own performance in a call, or says "coach this meeting".
---
```

4. Underneath the `---`, paste the full prompt from Step 2.
5. Zip the folder so you get `meeting-coach.zip`, then rename it `meeting-coach.skill`.
6. Upload it in Claude's skill settings.

Now you can say *"run meeting-coach on this"* and paste a transcript, in any conversation.

OPTION C – AUTOMATE IT (THE VERSION CECI ACTUALLY RUNS)

If you have Claude connected to your calendar and your notetaker, you can go further and have it run without you asking. Ask Claude, in a Project that already has the skill:

SCHEDULED TASK

Set this up as a scheduled task. Every weekday at 18:00, look at my calendar for meetings that happened today where I was a participant. For each one that has a transcript available in [your notetaker], run the meeting-coach analysis. Then give me ONE combined summary for the day: the numbers per meeting, and a single pattern you noticed across all of them. Do not send me one report per meeting – I will not read five reports.

That last sentence is the difference between a habit and an ignored notification.

THE INTERACTIVE VERSION

Ceci's lives as an artifact in Cowork rather than as an email: a list of her meetings that she clicks into when she wants feedback, rather than a report that arrives whether she wants it or not. If you're building this, prefer pull over push. Feedback you requested lands very differently from feedback that turned up.

On ChatGPT

OPTION A – A PROJECT

1. Create a new Project called `Meeting Coach`.
2. Put the Step 2 prompt into the Project instructions, with your name and role filled in.
3. Paste transcripts into chats inside the Project.

1. Go to *Explore GPTs* → *Create*.
2. Name it `Meeting Coach`.
3. In *Instructions*, paste the Step 2 prompt.
4. Under *Conversation starters*, add: `Here's a transcript, coach me.`
5. Turn **off** web browsing and image generation. It doesn't need them, and switching them off stops it wandering off to look up meeting-facilitation blogs instead of reading your transcript.
6. Save as private.

OPTION C - TASKS

ChatGPT's scheduled Tasks can prompt you rather than fetch for you. Set one for Friday afternoon: *"Ask me to paste this week's toughest meeting transcript, then run the Meeting Coach analysis."* Less automatic, but it survives having no connectors at all.

STEP 5

Add the tracking layer

This is the part missing from most versions of this workflow, including the one on the episode. A verdict tells you about one meeting. A trend tells you whether you're actually getting better.

After each run, append one line to a running file or sheet:

DATE	MEETING	MY TALK SHARE	OPEN QUESTIONS	LONGEST TURN	FIX I WAS GIVEN
11 Aug	Roadmap review	68%	2	340 words	Stop answering my own question

Then once a month:

MONTHLY REVIEW

Here is my meeting log for the last month. Do not coach individual meetings. Tell me: what is measurably improving, what is not moving at all, and which fix I have now been given three times and clearly haven't implemented. Quote the dates.

The third one is the useful one. Being told the same thing repeatedly, with dates attached, is harder to ignore than being told it once.

STEP 6

Make it fail loudly, and audit it

Two rules from the episode that apply to everything you build, not just this.

Make it fail loudly

If you automate this and the transcript isn't available, the run must tell you. A workflow that silently does nothing is worse than no workflow, because you'll believe it's running. Add to any scheduled version:

FAILURE INSTRUCTION

If you cannot find a transcript for a meeting, do not skip it silently. Say explicitly: "No transcript found for [meeting name] – coach did not run."

"If they can't reach the information, make it say: Hi, I'm the fairy, I couldn't reach the information."

BEATRICE ALIPRANDI – EP 11

Silence is not success.

Audit your own setup

Once you've built two or three of these, point your strongest model at all of them and ask it to find what's broken. On the episode, Bea did exactly this and it found an authentication flaw in one of her own scheduled tasks, left over from reconnecting some connectors.

SECURITY AUDIT

Go through every skill, scheduled task and automation I have set up. For each one, check: what data can it read, what can it write or send, whose credentials is it using, and what happens if a connection breaks or a permission changes. Flag anything where the access is broader than the job requires. Be specific about the risk, not general.

Run it on Claude Fable 5 or ChatGPT Sol. This is the one job where you want the expensive model — a reviewer that rubber-stamps everything is worse than no reviewer.

Every three prompts, check for bugs. Every ten, check for security.

HONEST LIMITS

When not to do this

Honest limits, because a guide that only sells you the upside isn't a guide.

- **Sensitive meetings.** Some conversations should not be transcribed at all, and some shouldn't leave your machine. Ceci retired a different automation partly for this reason — the output would have sat on a calendar she'd then have to make private. If a meeting is confidential, it's outside this workflow. Full stop.
- **Interviews and performance conversations.** Don't point this at meetings where you're assessing another person. The transcript then contains their material, not yours, and coaching yourself on it slides quickly into analysing them.
- **Meetings where talking a lot was the job.** A talk share of 80% in a presentation is correct. Tell the coach the purpose in the prompt or it will scold you for doing your job properly. This is why the **WHAT THIS MEETING WAS FOR** field exists.
- **The first week will be unpleasant, and that's fine.** Everyone's first three reports are worse than they expected. Judge the workflow on week four, not day one.

TROUBLESHOOTING

When it misbehaves

"It's giving me generic advice."

It's ignoring the verbatim-quote rule. Reply: *"Every claim needs a verbatim quote from the transcript. Redo it and cut anything you can't quote."*

"The percentages look wrong."

They're estimates from word counts, not a measurement instrument. Treat them as directional. If you want them exact, ask it to count words per speaker and show the arithmetic.

"It's quoting things I never said."

You're on too small a model, or the transcript is too long for the context window. Move up a tier, or split the transcript in half and run it twice.

"It's too nice."

Add the harsh calibration from Step 3, and add: *"You are being too generous. Assume your first draft was too kind and redo it."*

"It missed the actual problem in the meeting."

Tell it what the problem was and ask it to find the evidence. If there's no evidence in the transcript, the problem may have been in the room rather than in the words — tone, faces, who wasn't speaking. That's the limit of the tool, and it's a real one.

THE POINT OF THIS

A meeting coach isn't really about meetings

It's the smallest possible version of a bigger idea: the parts of your job you can't see are the parts you can't improve, and a transcript is just a mirror you can't argue with.

The busywork was never the job. But neither was flying blind.

This guide came out of Top of the Ops, Season 1 Bonus — ["Ride or Die AI Workflows"](#), where Bea and Ceci went through every automation they actually run. The rest of the Operator Field Guides live on the [Free Guides](#) page.

Built something better? Tell us. We'll publish the improved version and credit you.