

Coaching agents vs. coaching people

Ed Schaefer · Version 1.0 · August 2026 · edward-schaefer.com/cards/coaching-agents · Free to share

The same craft on different fuel.

- With a person, growth lands on the person and stays there. An agent forgets everything when the session ends, so anything I'd call coaching has to land somewhere that survives it: the setup — instructions, role files, project memory, the checks I've built in. What actually improves is the setup and the person running it. The agent itself doesn't.
- Two ways I've watched this go wrong:
- **Coaching the model instead of the setup:** correcting an agent mid-chat feels like it worked, then reverts the next session because the lesson only ever lived in memory that just got wiped.
- **Projected growth:** crediting the agent for maturing when what actually got better was how well I'd edited my own instructions.
- Improvement gets built in three places, not one:
- the setup itself, which is the part I'm best at
- the human running it, the harder client, because it carries ego and history a setup file doesn't
- the working relationship between the two — I run retro-style passes over my setups (what did we learn, what broke, what needs improving, what did I use a lot of), some on a weekly pass, some as a monthly audit, holding the frequency loosely because agent pace outruns any fixed rule
- Some of the human-coaching canon transfers whole. Intent-based leadership does. So do blameless postmortems and DevOps/TDD's pipeline thinking. Scrum's ceremonies don't: they assume a team that persists session to session, and agents don't.
- This is also why the AI Operating Model lists "corrections land in the setup, not the chat" as one of its standing practices for teams: the same line, applied at team scale.
- Check yourself the way you'd check anyone you were coaching: what have you actually changed in your setup lately — which instruction did you rewrite, which check did you add? If the honest answer is nothing, and the only evidence you have that an agent's improving is that it "seems smarter," you're crediting the tool for your own editing — and that improvement stops the moment you stop making it.

What's behind this card

Read the full case for how one model upgrade improved every other agent I run at [Coaching Agents Isn't Coaching People](#).

This is one of six cards today. The field guide maps seventeen practice areas, and I'm growing this set as each one earns its own depth.

This is a fit question, not a recipe. Where your team actually sits against this list, and what that's costing, is what the [Engineering Ways of Working Diagnostic](#) is built to find.