

# An AI coach that learns your voice and speaks back to you, **as you.**


Inner Voice is a multimodal AI companion I prototyped: it interviews you about your goals, clones your voice while you talk, and then reflects your own aspirations and peak moments back to you, in your own voice.

Role: **Solo, full-stack**

Claude + Whisper + voice-clone TTS

**62** commits in ~3 weeks

Working prototype



The core idea: hear your own voice, refined and reflected back, as a coach.

---

## The idea

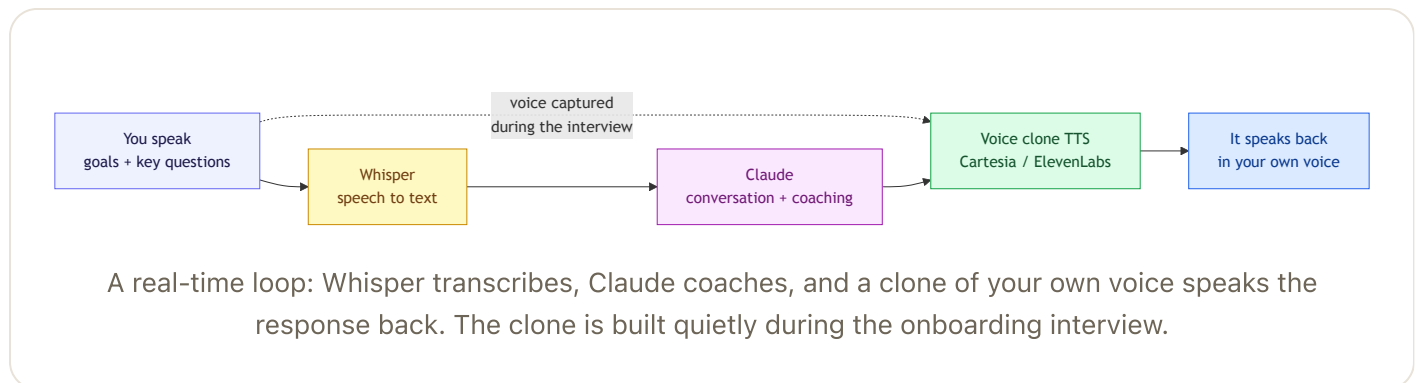
There is real power in hearing your *own* voice. I had been recording affirmations and playing them back to myself, and it works in a way another narrator never quite does. Inner Voice turns that into a product: an AI coach that speaks to you in your own voice, reminding you of the person you are trying to be and the wins you have already had.

It grows out of earlier wellness work of mine on countering negative self-talk. The wager is simple: a coach that sounds like you, recalling your peak moments back to you, is more believable and more motivating than any generic voice.

---

## How it works

The non-obvious part is the onboarding. Instead of asking you to read scripted lines into a microphone, Inner Voice **interviews you** about your goals and aspirations, and captures and refines your voice *while you answer*. The voice clone is a byproduct of a conversation you would want to have anyway. From there it runs a real-time multimodal loop:



End to end it is a React Native and Expo client over a Python FastAPI backend, wiring **Claude** for the conversation, **Whisper** for transcription, and **voice-cloning TTS** (Cartesia and ElevenLabs) for the reply, with real audio processing (noise reduction, segmentation) in between.

---

## See it in action

Demo coming soon

A short walkthrough, onboarding interview, voice capture, and the coach replying in your own voice, is being recorded. Check back, or ask me for a live look.

## Status, honestly

This is a **working prototype**, not a shipped product. It spins up and runs the full loop, and I paused it when work got busy. I am including it because the interesting part is not polish, it is that the hard, multimodal plumbing works: a real-time pipeline across three AI systems, plus an original onboarding idea that turns voice capture into something a user actually wants to do. It is proof I can take a non-obvious product concept and stand up the full multimodal stack behind it.

## By the numbers

**3**

AI systems in one loop:  
Claude, Whisper, voice-clone  
TTS

**62**

Commits, solo, in ~3 weeks

**2**

Tiers: RN/Expo client +  
FastAPI backend

**Real-time**

Speak, transcribe, coach,  
reply in your voice

**1**

Original idea: clone your  
voice during the interview

**Proto**

Spin-up-able prototype,  
demo in progress

*A personal prototype exploring voice, AI, and self-coaching. The takeaway is the capability: an original product idea, stood up end to end across a real multimodal AI stack. Happy to give a live walkthrough.*

