

Q4 2025

October – December 2025 · 4 changes

Every episode of My Weird Prompts is written, voiced and assembled by an automated pipeline. This changelog records its major updates — new models, new voices, new stages — and the first episode each one produced. Small prompt tweaks happen weekly and are not listed.

24 November 2025 · Infrastructure · Pipeline V3

Two AI hosts, one voice memo

The earliest episodes in the archive. Daniel records a voice memo, Gemini 2.5 Flash transcribes it and writes a two-host dialogue between Corn and Herman, and cloned voices read it out.

- Script and transcription: Gemini 2.5 Flash.
- Voices: Chatterbox voice clones, run on fal.ai.
- Roughly \$2 of compute per episode.

FIRST EPISODE: [#2 Local STT For AMD GPU Owners](#)

11 December 2025 · Voices · Pipeline V4

Pipeline V4: auditioning new voice engines

Text-to-speech moved to Fish Audio. This began a month of trying different voice engines (Fish Audio, then Inworld, then Chatterbox hosted on Replicate) before the show settled on running its own.

FIRST EPISODE: [#52 System Prompts vs. Fine-Tuning: Are We Building Solutions for Problems That Don't Exist?](#)

21 December 2025 · Script model · Pipeline V4

Scripts move to Gemini 3 Flash

The script writer was upgraded from Gemini 2.5 Flash to Gemini 3 Flash. From two days later it ran in preview, and wrote nearly 2,000 episodes before April 2026 brought a run of model changes.

FIRST EPISODE: [#67 Scaling or Pivoting AI for Code](#)

23 December 2025 · Infrastructure · Pipeline V5

Pipeline V5: production moves to Modal

Generation moved onto Modal, a serverless GPU platform, which has run the pipeline ever since. The V5 label then stayed in place for nine months; for a while V4 and V5 were both still running.

FIRST EPISODE: [#75 Community Depth vs. Raw Power in Local AI](#)