Case: Automotive

Getting in-car voice commands right the first time, without sending them to the cloud

Situation

- Example: a car’s voice assistant controls 41 vehicle functions.
- A model turns each spoken command into the exact function and settings.
- Simple commands run on board; the rest go to a cloud model.

Solution

- About 4,000 commands generated from 96 hand-written examples.
- A small Qwen3 model (1.7B) trained to return the function or ask back.
- Tested on 76 unseen commands.
- Upload to scored model: 1 h 34 min.

Results (measured)

93%

of commands right first time, up from 50%

1.7s

per command, on CPU only (4 cores)

Challenge

- Off the shelf, the same model got only 38 of 76 commands right.
- It never asked back: unclear commands were executed on a guess.
- It invented settings, such as seat heating level 5 on a 0 to 3 scale.

Example

- Command: “I’m cold, seat heating to three and ambient light to warm white.”
- Model returns: seat heating, front right, level 3; ambient light, warm white.
- “Heating on please.” The model asks which heating is meant.

Where this applies

- Assistants that map speech or text to a fixed set of actions.
- Comfort, infotainment, charging and trip planning in the vehicle.
- Service, workshop and production-floor terminals.

Illustrative use case, measured on synthetic data: 76 held-out commands, 17 Nov 2026. No customer data.