

CARE AI

Care AI: our own AI for practising the hardest conversations.



Care AI is the intelligence inside Duty Line Training Systems. The callers and how they behave, the voices they speak with, the way a call is heard, read and judged, and the way it gets better are designed and built by us, for call-handling training. Underneath, it draws on a leading large language model for language, and on the voice engine you choose.

What's ours, and what it's built on

Ours: designed and built by us	Built on
– Every caller: who they are, what they know, what they'll give away and when – The rules that make them behave like the public: patience, distress, answering only what's asked – How each voice is cast and acted, and the places heard behind it – The algorithms that listen: trigger cards, the call sheet, pace, talking over, safety checks and their timing – How a call is judged: each caller's must-dos, the line's checks, and who is on the line – The model listener: how an experienced listener would take the call, and why – The sign reader: hands, body and face, read in the browser – How it learns from its own mistakes and from trainers' corrections	– A large language model from Anthropic (Claude), for the words each caller says and for reading a call – Voices: ElevenLabs or OpenAI in the cloud, or Chatterbox on your own server – Transcription: ElevenLabs in the cloud, or Whisper on your own server – Hand, body and face tracking: MediaPipe, running in the trainee's browser – Each can be swapped without changing the callers or the training

The callers

- **Written, not generated on the fly.** Each caller has a brief written by us: their life, what's happened tonight, what they're holding back, and what a good listener does for them.
- **Patience.** Rushed, dismissive or lecturing replies wear it down; kindness and good questions build it back. Too low, and they hang up.
- **Distress, line by line.** The caller rates how they feel as they speak, from calm to the worst; the rundown shows it as a graph, with what was said when it moved.
- **Behaving like the public.** They answer only what's asked, often only part of it, hesitate, go off track, and come from the customer's own area.
- **More than one person.** A caller can hand the phone to someone, or be asked to; the new person speaks as themselves, in their own voice. Others in the room can be heard.
- **At the end,** the caller says, in their own voice, how they felt and what made the difference.

The voices

- Cast for each caller by age, sex and where they're from, so a Leeds pensioner and a Belfast teenager don't sound alike
- Acted from the moment: whispering, crying, shouting, slurring or flat, following the caller's mood
- The place behind them: a busy street, a house fire and its alarm, a quiet kitchen at three in the morning
- Your own server can make every voice, so no caller's words leave your building

How it listens

- **Trigger cards** are raised by what the caller says, from phrase lists built by us and yours to change, with negation understood and "not when" phrases ("for my back") that stop a card
- **The call sheet** fills itself in from what's said: the place, what they've taken, who's with them
- **The trainee's delivery:** words a minute, talking over the caller, and long silences, fed back gently
- **Safety, timed:** when they asked about safety, and whether they asked directly about suicide

How it judges a call

After every practice, Care AI reads the whole call against that caller's must-dos and the line's own checks, and judges by meaning: something done in other words still counts, with the moment quoted. It knows who was on the line, so a parent ringing about a child isn't marked down for a check meant for the child. Checks that didn't apply are said so. Without Care AI, phrase matching stands in, so feedback never stops.

Show me how it's done

The model listener is an experienced listener for that line, in its own calm voice. It deals with every card the caller raises and gives the reason for each move. Its next line is prepared while the caller is still talking, then it takes a natural pause before answering, so it sounds like a person.

Sign language and the face

On video, the sign reader notices signing and prompts the call-taker to use a relay. It reads signs a service has taught it: each hand by which hand it is, the leading hand first (a left-handed signer is mirrored), handshape from the bend of each finger joint, and where the sign is made, measured from the shoulders. The face carries BSL's grammar, and it reads that too: eyebrows for questions, a head shake for "not", a nod for "yes", puffed cheeks for "very". Care AI turns the signs into English, marked for the call-taker to check. All of it runs in the browser: no video leaves the device.

How it learns

- **The caller reviews itself.** After each call, Care AI notes its own mistakes playing that caller (repeating itself, giving too much away, breaking character) and reads them next time.
- **The model listener learns from what it missed** on a demonstrated call, for next time on that line.
- **Trainers correct the rundown.** "Not right?" on any point keeps the correction, and it's used the next time a call with that caller is judged.
- Only the newest few lessons are kept for each caller, so they stay short and current.



Safeguards

- Every caller is fictional, and every page says so; trainees are asked not to bring real people's details into practice
- A monthly Care AI budget, with a warning at 80% and an option to stop when it's spent; each organisation can have its own budget too
- Every Care AI request is logged with its cost; organisations see their own use
- Callers requested by a service are drafted by Care AI but go live only after a person has tried and approved them
- Experienced call handlers review scenarios for realism and safety

Hear Care AI's callers for yourself in a demonstration: dutylinesystems.co.uk/demo