

Humanizing AI Through Avatars and Patient Engagement

Rayhan Beebeejaun

BEYOND ACCURACY

Performance is no longer the problem

Diagnostic imaging

Sensitivity and specificity at or above the best specialists.

Liu et al. (2019), The Lancet Digital Health

Risk prediction

Models routinely match or exceed expert clinical judgement.

Systematic review, J. Clinical Epidemiology (2024)

Everyday adoption

Still slow, still limited. The barrier is not technical.

Wu et al. (2024), NEJM AI

The gap is one of acceptance, not performance. Patients and clinicians must trust the tool before entrusting it with a decision about health.

WHY PATIENTS RESIST

**“Patients don’t reject intelligence.
They reject foreignness.”**

Uniqueness neglect

Patients refuse medical AI even when told it is more accurate — they believe their condition is unique, and that only a human can grasp its nuances.

Longoni, Bonezzi and Morewedge (2019), Journal of Consumer Research

Resistance softens

Acceptance rises sharply when AI is presented as a tool in the physician’s service — not as a replacement for the physician.

Gaczek et al. (2023), general health care settings

WHY CLINICIANS HESITATE

Trust requires the ability to verify

Opacity blocks responsibility

Facing critical judgements, clinicians often choose not to engage with AI at all: when a model is opaque, its output cannot be verified — and an unverifiable answer is one no physician can take professional responsibility for.

Lebovitz, Lifshitz-Assaf and Levina (2022), Organization Science

More than a technical choice

Whether a physician uses AI in practice depends on reputational, legal and relational factors — not only on how well the model performs.

Dai and Singh (2025), Journal of Marketing Research

Generic AI asks both patient and doctor to trust a stranger.

THE TWIN DOCTOR FRAMEWORK

Not a generic AI — a digital twin of the doctor

The system learns from the clinician's own practice and meets the patient with the recognisable identity of their treating physician.

1

Coherent identity

Faithfully reflects the doctor's appearance, voice and communication style.

2

Shared memory

Accesses the patient's clinical and relational history, guaranteeing continuity.

3

Transparent reasoning

Every recommendation carries its underlying logic, verifiable by the doctor.

4

Clinical supervision

The doctor remains responsible; the AI extends their reach to more patients.

FROM FRAMEWORK TO PRESENCE

What it takes for an avatar to feel human

- 1 **A real face** — photorealistic, and the same face for the whole conversation — identity never drifts.
- 2 **The doctor's own voice** — cloned from a few minutes of recordings, speaking the patient's language.
- 3 **Natural timing** — it answers in under half a second, the rhythm of a real exchange.
- 4 **It can be interrupted** — the patient can speak over it at any moment, exactly as in real conversation.
- 5 **It is never frozen** — it breathes and blinks between turns, so presence never breaks.

THE INTAKE JOURNEY

A doctor becomes a twin in about 25 minutes

Seven steps, consent first: a short guided session of photos, natural on-camera conversation, voice in each language, and explicit guardrails.

1 Consent

2 About the avatar

3 Headshots

4 Conversation videos

5 Voice samples

6 Persona and guardrails

7 Review and submit

Privacy by design

Explicit likeness consent · encrypted at rest, stored in the EU · raw uploads purged after 90 days · never used to train shared models

LIVE DEMO

Three twins, three roles

AI Shai

General-practice consultation

English · takes a history one question at a time, exactly like a real visit.

AI Barbara

Patient education — back pain

Italian · explains, never prescribes; red flags go straight to the treating doctor or emergency care.

AI Simone

Healthcare innovation and the SSN

Italian · expert conversation on bringing innovation into the national health system.

About ninety seconds each — live, unscripted.

SAFE BY DESIGN

Guardrails are part of the design, not an afterthought

One question per turn

No interrogation, no overload — the pace of a careful clinical interview.

Plain spoken language

No jargon unless the patient brings it; numbers and doses read out naturally.

Red-flag escalation

Chest pain, stroke signs, crisis: immediate handoff to emergency services.

Strict scope

Each twin declines what it is not qualified to discuss — by explicit instruction.

The doctor remains responsible. The twin extends care — it never replaces accountability.

CONCLUSION

The challenge is relational, not technological.

Measure trust, adherence and continuity of care — not only accuracy. Combine the scale of AI with the authenticity of the doctor–patient relationship.

Technology should strengthen — not weaken — the human experience of medicine.

Longoni, Bonezzi and Morewedge 2019 · Gaczek et al. 2023 · Lebovitz, Lifshitz-Assaf and Levina 2022 · Dai and Singh 2025