Star

Living Care:

Human Agentic Interaction Model for Digital Health

Foreword

For nearly two decades, Star has always been the bridge between technology and user experience for our clients, from ambitious startups to some of the world's most iconic brands. We fundamentally believe that technology is only an enabler for positive impact when the product or solution is something people truly want to use - and that's why the user experience is so critical.

Healthcare has always aimed to put people first. But too often, legacy systems and fragmented digital experiences have forced both patients and providers to work around the technology, rather than having it work for them.

With its intelligent systems' ability to understand context and adapt to each individual's unique needs, agentic AI offers a new possibility of how care can be delivered and experienced.

For the first time, technology can help people be more human by creating a digital health experience that is dynamic, adaptive and always built on a clinical foundation, all while combining care safety and equity for every individual.

We have been at the forefront of every emerging technology wave: from mobile, to cloud, to platforms and now artificial intelligence. Throughout, our mission has always been democratizing technology and putting human experience at the center.

We believe agentic AI is the next true lever to unlock personalized care and wellness at a scale, and with a level of empathy, previously unimaginable. We are thrilled to partner with the healthcare sector on this next phase and create a more engaging, human experience for everyone in the care journey.

Mahesh Naphade
Head of Healthcare & Life Sciences, Star

Agentification of telehealth and digital care

The rapid growth of telehealth, remote patient monitoring and digital health platforms have ushered in unprecedented opportunities and new pressures for healthcare providers, payers and technology innovators. With the global digital health market projected to reach $946 billion by 2030, the stakes for delivering effective digital experiences have never been higher.

Yet, most digital health solutions today still rely on static software and rigid interfaces. These platforms were designed primarily to execute tasks or transmit information, not to actively adapt to each patient's unique circumstances or deliver truly engaging experiences.

As adoption across telehealth and wearables and remote patient monitoring accelerates, those who can create personalized and engaging experiences at scale will gain quicker market share.

Understanding the AI evolution: generative vs agentic

To appreciate the transformation ahead, it's essential to distinguish between two related but fundamentally different AI paradigms now reshaping digital health: generative AI and agentic AI.

Generative AI excels at creating content based on prompts. It is reactive: users provide input, and the model produces a specific output, such as a patient education document, a clinical summary or a response to a question. Generative AI does not autonomously plan, decide or take action across multiple steps. Its main value lies in content generation speed.

Agentic AI is a proactive, autonomous system capable of independently pursuing multi-step goals with minimal human intervention. Rather than waiting for step-by-step prompts, agentic AI receives a high-level objective, breaks it into sub-tasks, makes decisions based on context, and executes a coordinated sequence of actions to achieve that goal.

Agentic AI combines the strengths of generative AI (content creation), predictive AI (forecasting), natural language processing (understanding), machine learning (continuous improvement) and robotic process automation (task execution) into an intelligent, goal-oriented framework

Critically, agentic AI can call upon generative AI as one of its tools. For example, an agentic system managing a post-discharge care journey might use generative AI to draft personalized medication reminders, create educational visuals tailored to a patient's literacy level or compose follow-up messages, while the agentic layer orchestrates when, how and through which channel these assets are delivered, continuously adapting based on patient response and clinical data.

From static app to generative solutions and interfaces

Traditionally, digital health interfaces are designed as static applications with limited configuration options: product teams define a set of screens, menus and workflows, then ship the same structure to every user. At best, these apps offer light configuration (language, theme, a few preferences), but the core interaction model - what you see, in what order, and how you interact - remains largely the same whether you are a 25-year-old new parent or an 80-year-old managing multiple chronic conditions. Even when chatbots or simple LLM features are added on top, they usually sit inside this fixed frame, returning text responses in a chat window without changing the surrounding interface or workflow.

Most of today's "agentic" AI deployments in healthcare still reflect this paradigm. They focus on efficient use cases such as virtual assistants handling FAQs or schedulers booking appointments. These solutions reduce administrative burden but do not fundamentally transform the patient experience, because they operate on pre-designed interaction paths and cannot restructure the experience around each individual in real time.

Recent research from Google shows that modern large language models can go further. They can generate not only content but also complete, high-quality custom user interfaces on the fly for a wide range of prompts. In controlled user studies, people strongly preferred these dynamically generated interfaces over traditional text-only responses, with preference rates exceeding 80% across multiple task types. This demonstrates that generative models can act as effective "UX creators," defining the interaction logic, assembling layouts and components, tailored to a given goal and context of each individual person.

This Generative UI capability is a fundamental shift from both static apps and simple chatbots. Instead of building a single portal for all patients and bolting a chatbot onto it, the system can create bespoke, task-specific micro-experiences for each individual and situation.

A human agentic interaction goes a step further: it uses agentic AI to decide what sequence of actions a patient needs, when to engage them, and through which combination of modalities (voice, text, touch, visual), based on their unique needs and constraints and clinical context, then uses Generative UI to render the right interface for that moment.

Personalization without limitation with HAI Model

We see agentic AI combined with Generative UI as the catalyst for true personalized orchestration, system where the digital health experience is provided by an individually created application for ea person and clinical scenario governed by a robust **Human AI Interaction (HAI) Model** that ensures patient safety and care integrity.

What is the Human Agentic Interaction (HAI) Model

A human centric model which defines how people interact with AI agents and how these agents should be built and behave to allow for trust, safety, transparency and control for each user.

That means we now have the technology to create experiences that are:

Dynamically tailored to each individual's cognitive, emotional, and contextual profile, including health literacy, impairments, language and environment.

Implemented on demand across modalities: within minutes, the system can assemble a personalized user experience across voice, text, visuals and haptics into a ready-to-use application.

Built on a clinical foundation, incorporating relevant medical history, current clinical data, and evidence-based guidelines with embedded safety guardrails.

Continuously updated in real time as user state and feedback shift, with interfaces and workflows adapting as symptoms, behavior, comprehension or context change.

Seamlessly orchestrated across devices and cloud: data flows securely between edge devices (wearables, phones, home sensors) and cloud systems, with privacy-preserving local processing at the edge

Regulatory-safe and audit-ready by design: every AI-driven adaptation, decision and data exchange is logged with full traceability, supporting different regulations such as HIPAA, GDPR and FDA.

HAI model is a key framework for deploying adaptive experiences at scale in regulated environments. It helps clarify how AI and humans collaborate safely, what can be adapted automatically, when human oversight is required and how accountability is maintained.

Human AI interaction • Clinical governance & safety rules • Decision authority & escalation protocols • Regulatory compliance & audit trails • Human-in-the-loop patterns • Explainability	**= Foundational layer** What should we adapt, when, under what circumstances, when must a clinician review or approve
Agentic AI orchestration • Clinical goal reasoning, intent preservation • Multi-step care pathway • Patient context interpretation • Workflow coordination across channels/ devices	**= Orchestration layer** What sequence of action will achieve this clinical goal based on the patient's specific data/ input
HAI executive orchestration • Dynamic interface generation • Modality selection	**= Execution layer** Where and how do I render this experience for this patient's abilities and context
Human - AI interaction • Clinical governance & safety rules • Decision authority & escalation protocols • Regulatory compliance & audit trails • Human-in-the-loop patterns • Explainability	**= Foundational layer** What should we adapt, when, under what circumstances, when must a clinician review or approve etc.
Agentic AI orchestration • Clinical goal reasoning, intent preservation • Multi-step care pathway • Patient context interpretation • Workflow coordination across channels/ devices • Decision-making with guardrails	**= Orchestration layer** What sequence of action will achieve this clinical goal based on the patient's specific data/ input

In addition, advances in edge AI - like Google AI Edge, MLC Chat and ONNX Runtime - are enabling systems to process data locally, bringing dynamic, sub-second responsiveness. As these capabilities mature, interfaces will not simply update once, but will continuously adapt in real time as users interact, sensing changes in context and adjusting modalities proactively.

Achieving this level of adaptation requires advanced AI under the hood. The system must integrate several capabilities:

Application goal: The AI keeps the intent of the application, such as medication treatment plan, as the core focus

User context: It factors in every user's unique context - this will require some input from the patient but once done, the AI learns and adapts. External conditions can also be context signals.

Real-time feedback loop: The AI monitors user actions as implicit and explicit feedback and alters the interface accordingly.

Multimodal interaction: The AI engages users through multiple interaction modalities, choosing the ones that best fit the user's needs at each moment.

This multimodal flexibility is a direct reflection of the system's dynamic ability in curating the experience. The system essentially decides how it "talks", how it "listens", how it "sees" and most importantly how it "understands" the user:

Adaptive output: The system modulates its communication style, for example, speaking out instructions versus displaying them, showing a diagram or playing a video demonstration or using vibration cues on a wearable device. For instance, if a patient has low literacy, the agent leans more on audio/visual guidance.

Adaptive Input: The system is not limited to button clicks or forms. It can process natural language voice commands such as "I'm not sure what to do next," interpret them as a signal of confusion, and adjust the guidance or escalate to a human.

The most important part to point out is that these modality choices are not pre-programmed branches but are decided agentically by the AI within the perimeters defined by HAI Model, making the experience both human and unique in every moment.

Implication to patience experience and care

It is critical to note that AI-led experience and solution curations does not remove clinicians from the loop; rather it's about preserving clinical guardrails while enhancing patient agency. In a HAI governed model, AI personalizes how care is delivered, while clinicians remain accountable for what care is delivered and when critical decisions are made.

While the HAI Model dynamically orchestrates and adapts every aspect of the user experience, the underlying clinical pathways and decision logic stay grounded in established protocols and human oversight. The system does not invent new treatment plans; rather, it ensures that validated plans are communicated in ways that are individually understandable, accessible and actionable for each patient.

The strategic challenge is achieving the right balance between automation and human judgment. Evidence from "centaur" models in healthcare shows that hybrid human-AI decision frameworks can outperform either clinicians or algorithms alone, especially when roles and decision rights are clearly defined. Successful HAI Model deployment follows this centaur pattern: AI handles personalization, monitoring and routine guidance, while automatically flagging risk or deviation for clinician review and intervention.

For example, an agent might draft a patient's personalized weekly care plan and let the patient tweak it, but also send a summary to a nurse for oversight. This ensures clinical standards are upheld – protocols can be baked into the AI's decision rules and workflow.

The impact is a care journey that feels more engaging and empowering for every patient, while safeguarding clinical quality every step of the way. This is not about "using AI to replace people"; it is about using AI to make that interaction in care more personal, and more human.

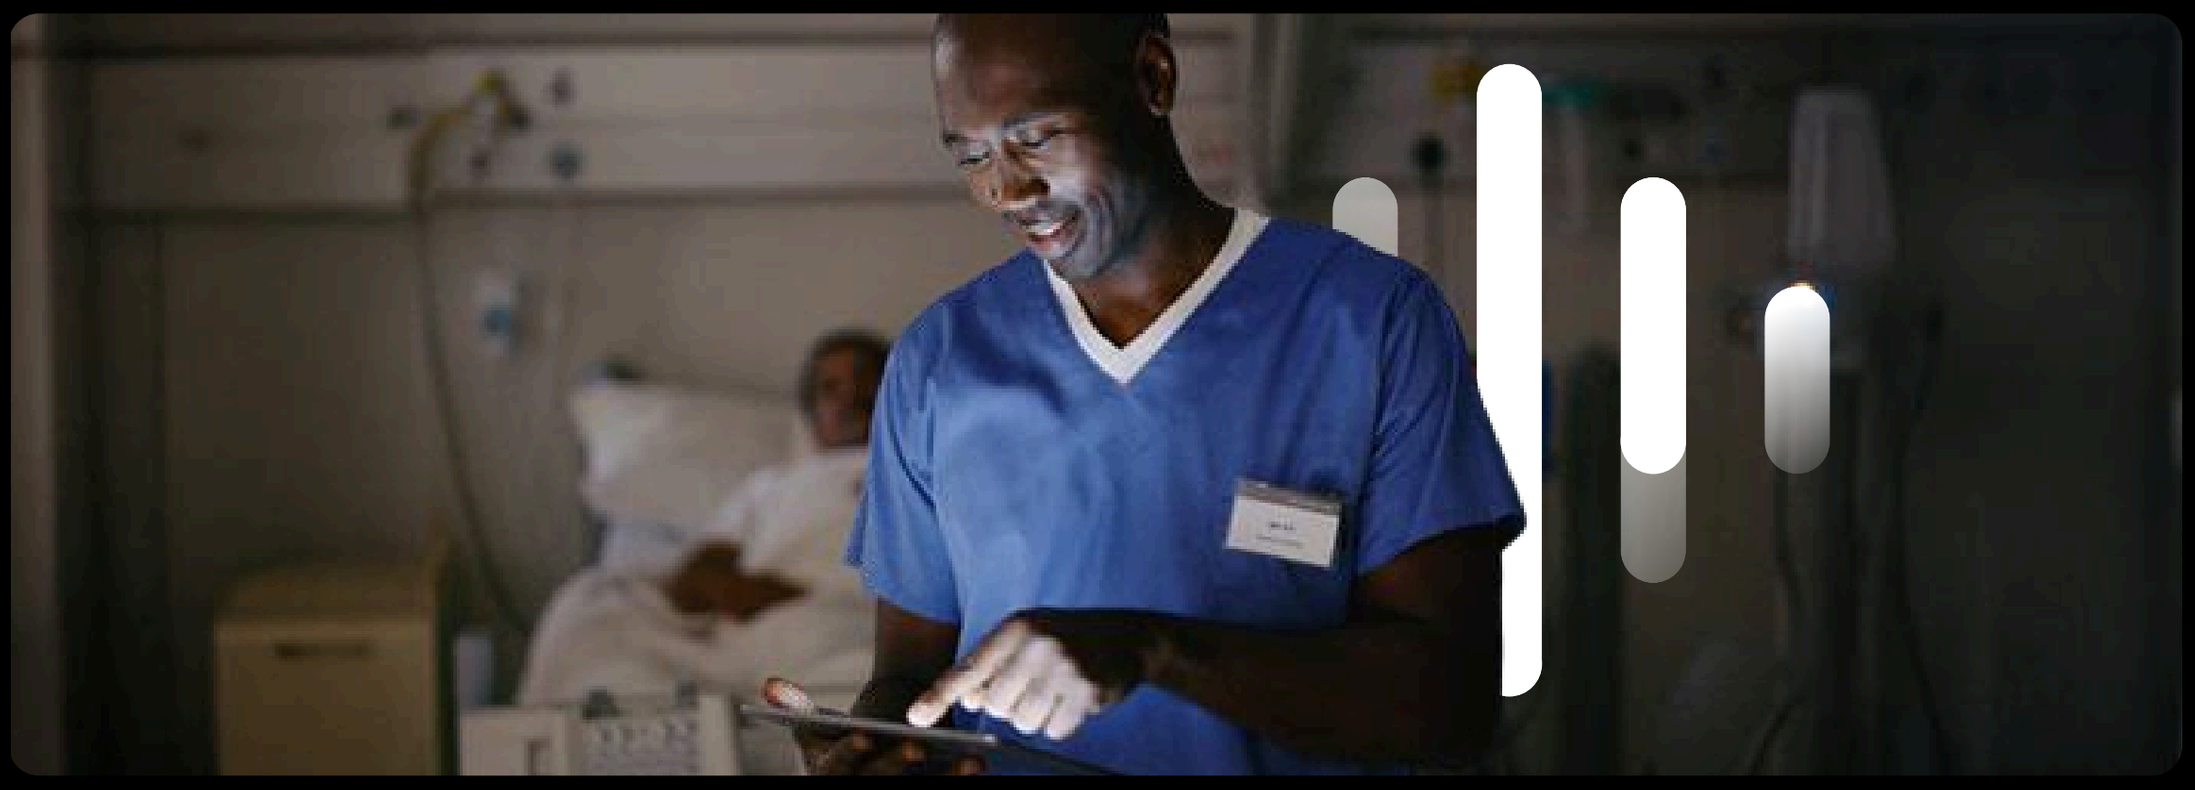

How individualized patient application is created: core principles

While the underlying technology is complex, the essence of the AI approach is elegantly simple from the user and clinician perspective:

Patient data and context capturing: The system begins by collecting data about the user, this includes core demographics, key medical history and even clinical signals from wearables or device integrations. All data collection must follow regulatory logging and audit protocols

Deliver across preferred devices, with continuous support: Once the individualized application is generated, it is delivered seamlessly to the patient's preferred device, whether that is a smartphone, tablet, wearable or another interface. The experience remains responsive across platforms. As the patient interacts, the system continually monitors feedback and evolving user context, adapting in real time and ensuring ongoing personalization, safety and support throughout the care journey.

Intent recognition, agentic experience curation: Advanced AI then analyzes this combined dataset to identify each patient’s needs, abilities, and impairments. It interprets intent, applies safety protocols and dynamically constructs an individualized application. Because the UI is generated and content rendered by AI agent for the user,it can switch seamlessly between modalities (voice ➜ screen ➜ haptic), depending on user context and requirement

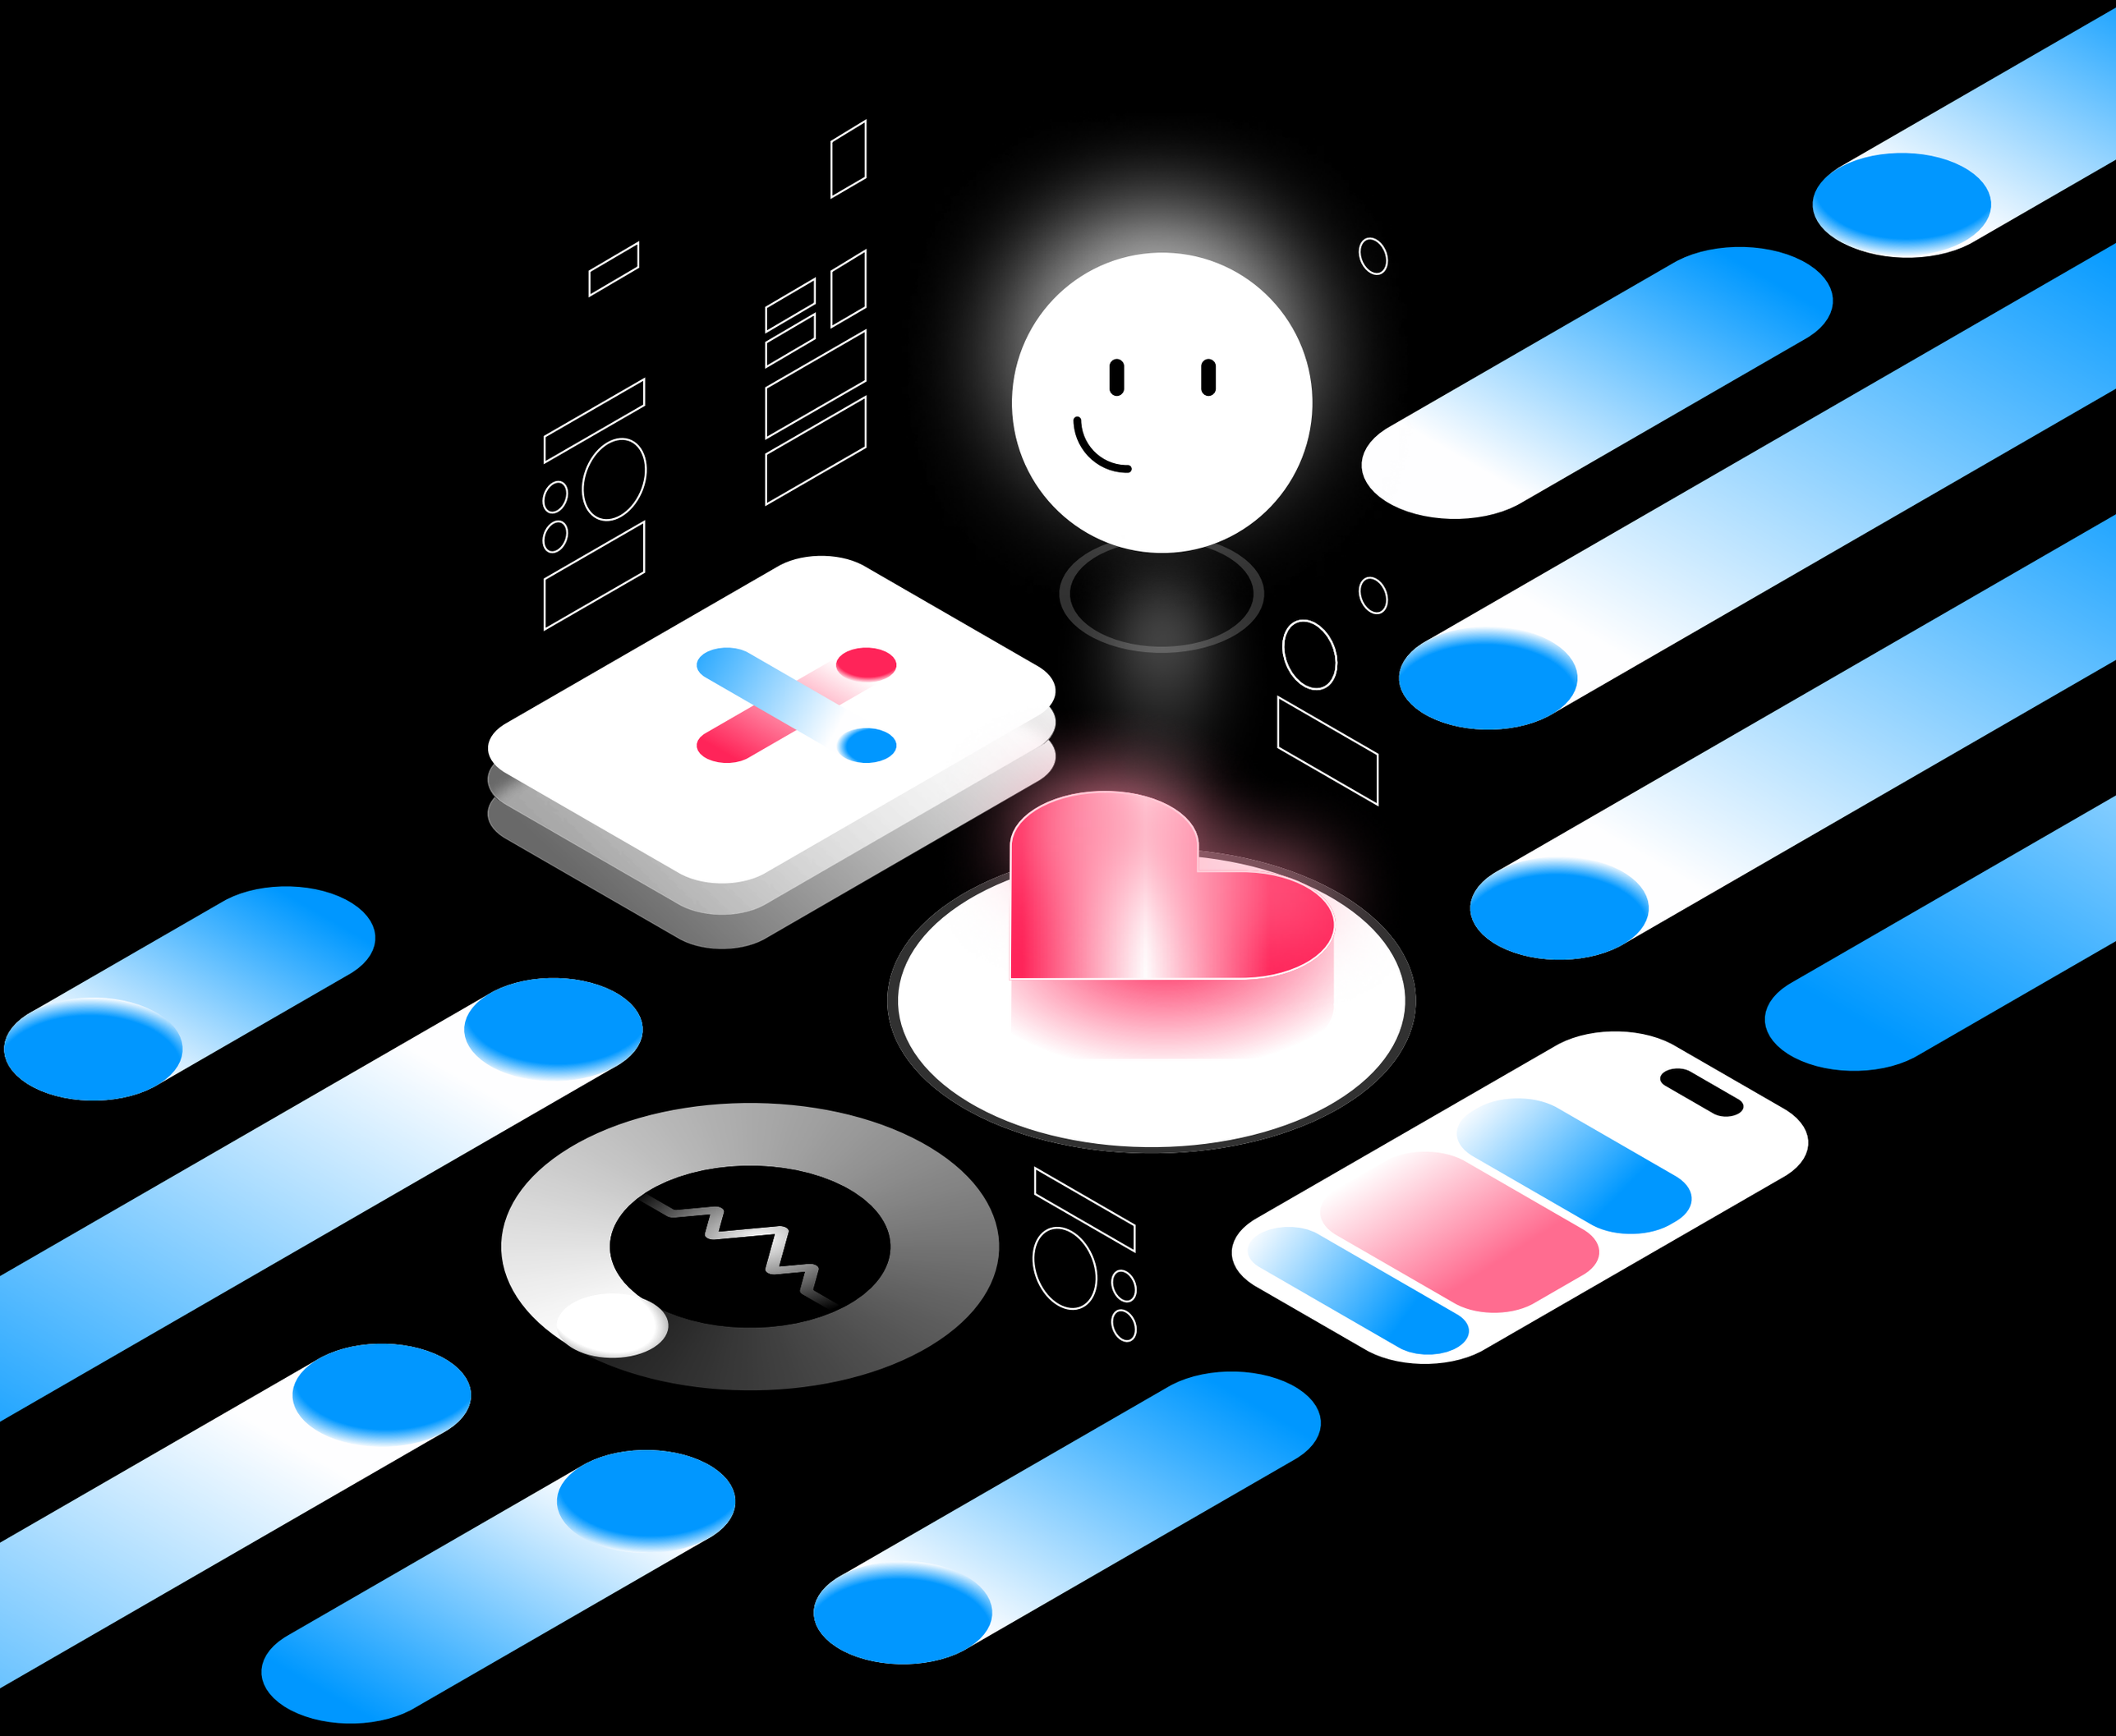

Experience, (re)imagined

The following use case scenarios illustrate how the HAI Model transforms the experience from one-size-fits-all tools into individualized patient applications safely, ethically and grounded in clinical safety

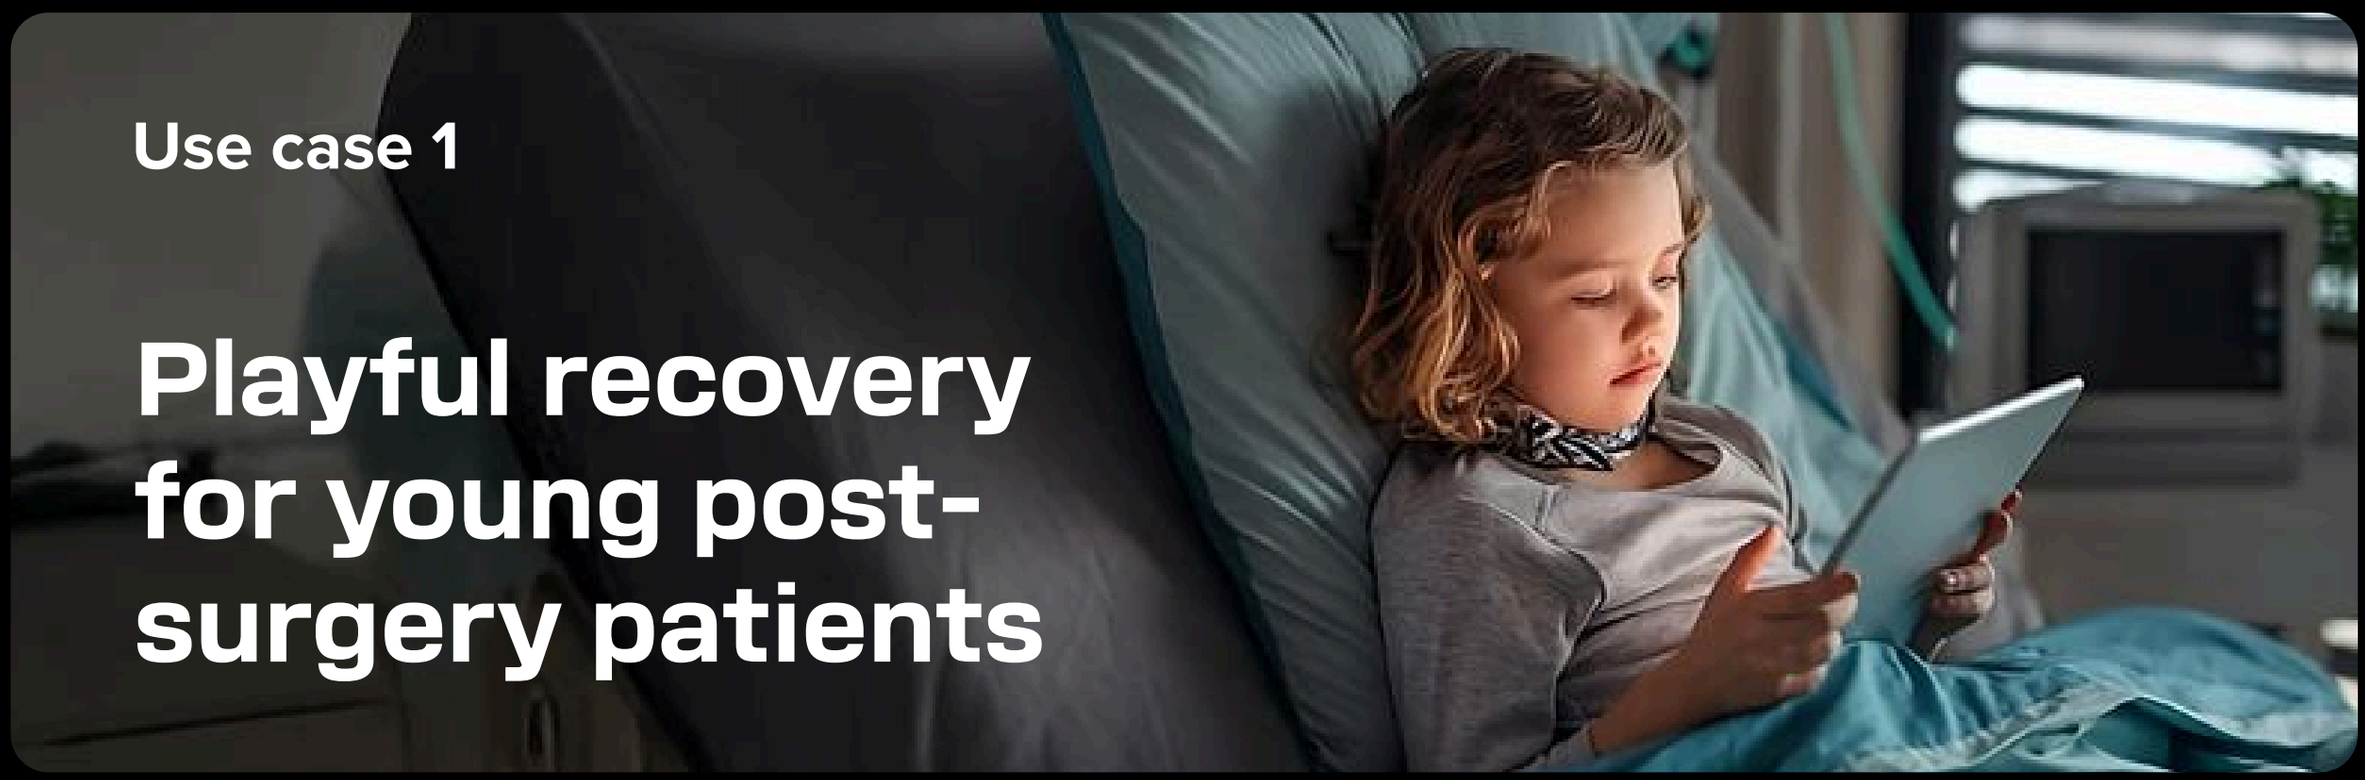

Use case 1

Playful recovery for young post-surgery patients

Profile:

Luca, 7 years old, recovering from appendectomy. Cannot read fluently, has limited attention span, needs to feel involved but not responsible. Parents guide the medical tasks.

Unique needs:

- **Developmental cognition:** Limited reading skills, short attention span, concrete rather than abstract thinking.
- **Shared agency:** Can express basic feelings but cannot interpret medical instructions or make care decisions.
- **Emotional containment:** High sensitivity to fear; needs reassurance and simple, non-threatening explanations.
- **Parent-mediated safety:** Adults must remain the final decision-makers for wound checks, medication, and escalation.

Possible HAI Model guided experience:

Interactive mode	Feature	Description
Visual and audio	Icon-led UI: age-appropriate interactions	Large, colorful icons and faces. Pictures, emoji and short audio promts over text; one micro-step at a time to match attention span and comprehension
Logic	Role-aware logic	Symptom-expression steps are routed to Luca (e.g., “how you feel”), while all clinically consequential actions (wound checks, escalation) are routed to parents with explicit caregiver framing.

Use case 2

An elderly patient after knee replacement

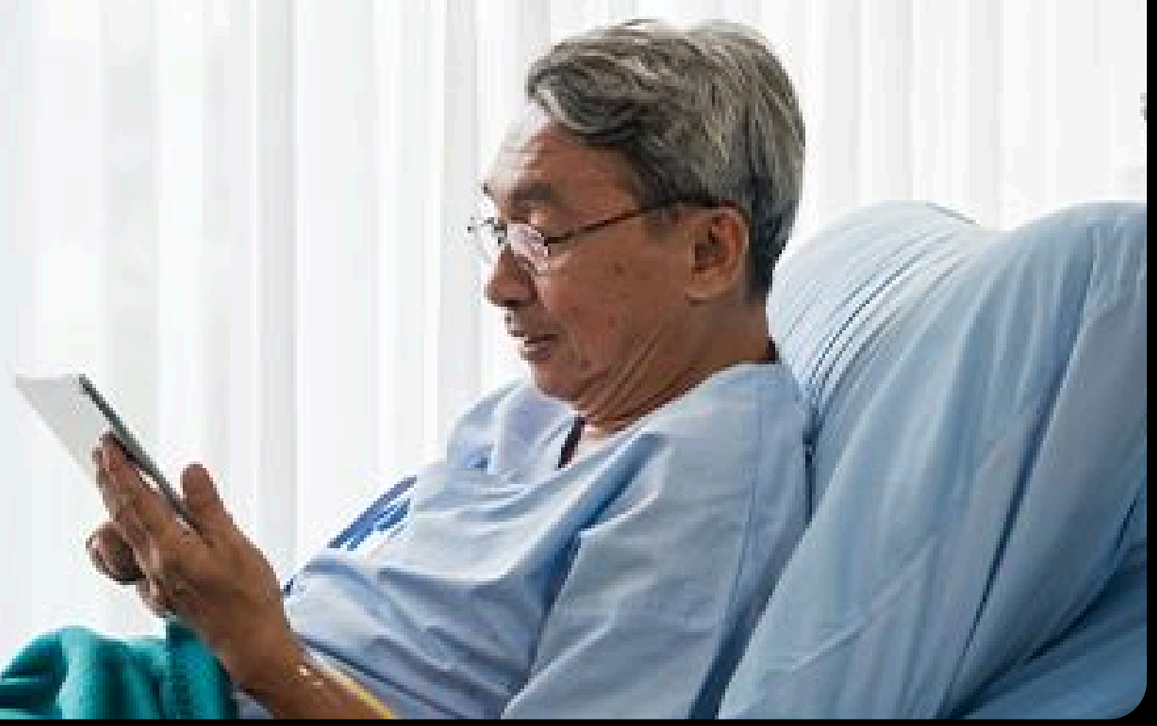

Profile:

A 79-year-old man recovering from knee replacement, with hand tremor, age-related vision limitations, low digital literacy, but intact cognition and strong desire for independence.

Unique needs:

- **Motor and vision constraints:** Difficulty with fine-touch precision and small text or low contrast; easily frustrated by complex digital interactions.
- **Low digital literacy:** Struggles with non-linear navigation, branching flows and gesture-heavy patterns.
- **Identity and dignity:** Rejects infantilizing tones or aesthetics; wants to be treated as competent and responsible.
- **Safety-critical rehab adherence:** Needs to follow a structured protocol while avoiding over- or under-exertion.

Possible HAI Model guided experience:

Interactive mode	Feature	Description
Visual	Accessibility-first interaction templates	Layouts with large, stable touch targets, high contrast, and linear “Next/Back” flows when motor and vision constraints are detected.
Tactile	Protocol-constrained personalization	Symptom-expression steps are routed to Luca (e.g., “how you feel”), while all clinically consequential actions (wound checks, escalation) are routed to parents with explicit caregiver framing.

Use case 3

Post-stroke rehabilitation patient

Profile:

A 45-year-old professional recovering from a mild stroke affecting speech articulation and fine-motor control, cognitively intact but easily fatigued, motivated to regain function and return to work.

Unique needs:

- **Input limitations:** Voice input is unreliable; precise tapping, tracing and writing are effortful.
- **Fatigability:** Cognitive and motor tasks quickly induce tiredness and frustration if sessions are too long or too intense.
- **Risk and prevention:** Requires ongoing monitoring for blood pressure and early neurological warning signs.

Possible HAI Model guided experience:

Interactive mode	Feature	Description
Visual/ Input	Impairment-aware modality selection	Prioritizes large-tap, low-precision touch interactions and simple gestures tuned to reduced fine-motor control.
Flow logic	Dynamic workload	Session length, difficulty and break intervals are adapted in real time based on performance and reported fatigue, within clinician-defined safe ranges.

Our experts

Mahesh Naphade,
Head of Healthcare & Life Sciences
mahesh.naphade@star.global

Martin Fix,
Technology Director
mfix@star.global

About Star

Star is a global technology consultancy that supports industry leaders on their digital journey. By connecting business strategy with technology execution, we deliver solutions that help enterprises innovate, optimize and scale.