

This document was translated from English by machine and may contain mistakes. The English original is at <https://www.novaracg.com/2026/05/03/when-access-becomes-approximation-ai-avatars-clinical-risk-and-the-cost-of-getting-it-wrong/>.

When Access Becomes Approximation | AI Avatars, Clinical Risk, and the Cost of Getting It Wrong

Heather Grizzle · May 3, 2026

AI sign language avatars are entering healthcare before they meet clinical standards. This NCG position examines communication failure risk, governance gaps, and why premature deployment shifts liability onto patients and providers.

Conteúdo

1. Presence Is Not Proof
2. Where Failure Lands
3. The Missing Evaluation Framework
4. Iteration Has No Buffer Here
5. Recalibrating Readiness
6. A Signal, Not an Incident

The moment itself is ordinary on the surface. A patient walks into urgent care, checks in, sits through triage, and prepares to explain what brought them in. What should be routine quickly becomes something else. Instead of a stable communication pathway, the patient is handed a technological substitute that has not yet earned its place in that room. A sign language avatar is introduced as the bridge. Within seconds, it collapses under the weight of real conversation. The system cannot keep up with the pace, the nuance, or the structure of the interaction. The patient stops it and pivots, not because alternatives are better, but because they are at least predictable. The visit continues, but the standard has already shifted downward.



Presence Is Not Proof

This is where the current trajectory of AI in accessibility reveals its weakest point. There is a growing tendency to treat presence as proof. If a tool exists, if it can demonstrate output in controlled settings, it is assumed to be ready for real environments. Healthcare exposes that assumption immediately. Clinical communication is not linear. It is iterative, layered, and often fragmented. A provider moves quickly, revises questions mid-sentence, reacts to incomplete information, and expects the response to carry context forward. Any communication system operating in that environment must do more than translate words. It must sustain meaning across time, across interruptions, and across changing conditions. Without that continuity, the interaction begins to degrade, and once it degrades, it does not recover cleanly.

The introduction of AI avatars into this space, particularly by organizations positioned as leaders in accessibility infrastructure such as Sorenson, is not just a technological evolution. It is a shift in how access is being defined and delivered. That shift carries consequences that are not being fully acknowledged in current deployment practices. When a healthcare provider relies on a communication system, they are making an implicit commitment that the patient will be understood and that the exchange will support informed decision-making. If the system cannot meet that threshold, the failure does not remain contained within the tool. It extends into the clinical process itself.

Where Failure Lands

What makes this particularly concerning is not that the technology failed in a single instance. Failure is expected in development cycles. The concern is that the system was positioned in a context where failure carries immediate impact. The patient is forced to compensate. The provider proceeds without a validated communication channel. The interaction becomes a patchwork of partial understanding and workaround methods. This is not a neutral outcome. It introduces ambiguity into environments that require precision. It alters the conditions under which care is delivered.

The Missing Evaluation Framework

There is currently no consistent framework governing how these systems are evaluated before entering clinical use. There are no universally adopted benchmarks for sign



language accuracy in real-time medical scenarios. There are no standardized methods to test whether an avatar can maintain context across a full patient interaction, from intake to discharge. There is no clear taxonomy that distinguishes between different types of avatar systems, which leads to procurement decisions based on surface-level similarities rather than functional capabilities. In the absence of these structures, deployment decisions are being made on incomplete information, and the risks associated with those decisions are being absorbed at the point of care.

From a governance perspective, this is where the failure becomes structural. The issue is not that innovation is happening. The issue is that the sequence is misaligned. Systems are being introduced into high-stakes environments before the conditions for safe use are established. The burden of that misalignment does not fall on the developers or the vendors. It falls on the patient, who must navigate the gap between what was promised and what actually works in the moment.

Iteration Has No Buffer Here

The argument that this is part of an inevitable transition does not hold when examined through a clinical lens. Iteration belongs in controlled environments where variables can be isolated and outcomes can be measured without consequence. Healthcare does not offer that buffer. Every interaction carries weight. Every exchange contributes to a chain of decisions that affect diagnosis, treatment, and patient autonomy. Introducing instability into that chain is not a step forward. It is a redistribution of risk.

The concept of accessibility itself begins to erode under these conditions. Accessibility is not defined by the presence of a tool, nor by the intention behind its use. It is defined by whether the individual can participate fully and accurately in the interaction. If the system cannot support that participation consistently, it does not meet the standard, regardless of how advanced it appears in isolation. What emerges instead is a form of approximation that mimics access without delivering it.

What emerges instead
is a form of approximation
that mimics access
without delivering it.



Recalibrating Readiness

The path forward requires a recalibration of how readiness is determined. AI avatar systems must be subjected to validation processes that reflect the environments in which they will be used. This includes sustained interaction testing, context retention analysis, and scenario-based evaluation within clinical workflows. There must be transparency in how these systems operate, including clear disclosure when AI is being used in place of human interpretation. Most critically, there must be enforceable thresholds that define when a system is suitable for deployment and when it must remain in development.

UNTIL THOSE CONDITIONS ARE IN PLACE

The use of AI avatars in healthcare should be limited to controlled, low-risk applications where failure does not compromise the integrity of the interaction. Anything beyond that introduces uncertainty into a system that depends on clarity.

A Signal, Not an Incident

The urgent care visit is not an isolated incident. It is a signal. It reflects a broader pattern in which innovation is outpacing the structures designed to manage it. Closing that gap is not a matter of slowing progress. It is a matter of aligning it with the realities of the environments it seeks to serve.

Communication in healthcare is not a feature that can be optimized over time. It is a foundational component of care delivery. When that foundation is unstable, the effects are immediate and far-reaching. The decision to deploy a system that cannot yet sustain that foundation is not simply premature. It is avoidable.

That is the point at which governance must intervene.

Como citar

Grizzle, H. M. (2026, May 3). When Access Becomes Approximation | AI Avatars, Clinical Risk, and the Cost of Getting It Wrong. Novara Consulting Group. <https://www.novaracg.com/pt/2026/05/03/when-access-becomes-approximation-ai-avatars-clinical-risk-and-the-cost-of-getting-it-wrong/>

