Online-Journals.org

iJAC International Journal of Advanced Corporate Learning

iJAC | eISSN: 1867-5565 | Vol. 19 No. 2 (2026) | OPEN ACCESS

https://doi.org/10.3991/ijac.v19i2.58755

TLIC PAPER

AI-Powered Avatars in Medical Education: Advancing Virtual Coaching for Clinical Readiness

Fernando Salvetti[1] (✉), **Barbara Bertagni**[1,2], **Roxane Gardner**[3], **Jenny Rudolph**[4]

[1]e-REAL Labs at Logosnet, Turin, Italy – Lugano, Switzerland –Houston, TX, USA – New York, NY, USA

[2]Center for Medical Simulation, Boston, MA, USA

[3]Brigham and Women's Hospital/Children's Hospital/Massachusetts General Hospital/Harvard Medical School, Boston, MA, USA

[4]Center for Medical Simulation/Beth Israel Deaconess Medical Center, Boston, MA, USA

salvetti@logosnet.org

ABSTRACT

The increasing complexity of healthcare environments requires innovative educational approaches to better prepare students and professionals for the realities of clinical practice. Traditional lectures and limited role-play simulations often fall short in providing repeated, standardized, and unbiased opportunities to practice essential communication skills. At e-REAL Labs, we have developed Conversation Mastery™, a proprietary platform powered by AI-driven avatars—embodied conversational agents (ECAs)—specifically designed to enhance virtual coaching in medical education. These avatars recreate difficult and emotionally charged clinical conversations, offering real-time adaptive feedback and fostering critical skills such as conflict resolution, empathy, and inclusive communication. Unlike generic AI coaching tools, the system integrates generative AI, natural language processing, and behavioral analytics into a controlled and curriculum-aligned framework. Learners benefit from personalized, structured feedback that directly supports the acquisition of clinical competencies and professional readiness. In collaboration with the Harvard Center for Medical Simulation, we conducted a controlled study comparing lecture-based instruction, traditional role-play, and AI-driven avatar training among nursing students. Results demonstrate that students training with avatars achieved significantly higher success in conflict de-escalation, improved confidence, and stronger adherence to best practices. By embedding Conversation Mastery into simulation-based education, institutions can expand practice opportunities, ensure bias-aware evaluation, and support compliance with GDPR, HIPAA, and emerging AI governance frameworks. This paper extends previous work presented during The Learning Ideas Conference 2025 (held at Columbia University in New York City and online) by providing a deeper exploration of the pedagogical, ethical, and regulatory implications of integrating AI-driven avatars into healthcare training. We argue that these technologies represent a new paradigm for clinical readiness, complementing faculty expertise while scaling access to high-quality, adaptive, and equitable learning experiences.

KEYWORDS

medical education, AI-powered avatars, virtual coaching, clinical readiness, healthcare simulation, bias mitigation

Salvetti, F., Bertagni, B., Gardner, R., Rudolph, J. (2026). AI-Powered Avatars in Medical Education: Advancing Virtual Coaching for Clinical Readiness. *International Journal of Advanced Corporate Learning (iJAC)*, 19(2), pp. 80–102. https://doi.org/10.3991/ijac.v19i2.58755

This article is an expanded version of a paper presented at The Learning Ideas Conference, held in New York, NY, USA, June 11–13, 2025. Article submitted 2025-11-19. Revision uploaded 2026-03-19. Final acceptance 2026-03-23.

1 INTRODUCTION

The growing complexity of healthcare delivery requires advanced educational approaches to prepare students and professionals for the communicative and ethical challenges of clinical practice. Navigating difficult conversations, addressing conflicts, and fostering diversity, equity, and inclusion (DEI) are essential skills for ensuring safe, effective, and compassionate care. While traditional coaching and classroom-based methods have value, they often reveal limitations: variability in instructor feedback, restricted opportunities for practice, and subjective biases in assessment. The emergence of AI-powered avatars—also known as Embodied Conversational Agents (ECAs)—offers a transformative pathway, providing **adaptive, repeatable, and unbiased** training experiences that replicate real-world interactions in a controlled environment.

1.1 Background and rationale

Healthcare education is increasingly facing higher levels of complexity: patients often have multiple comorbidities, teams are becoming larger and more interdisciplinary, and expectations for patient-centered care and equity continue to rise. In this setting, the ability to have difficult conversations, manage conflicts, and communicate across cultural differences has become just as important as technical skills. Traditional teaching methods—such as lectures, workshops, or role-playing—offer valuable foundations; however, they often suffer from inconsistencies in facilitation, limited chances for repetition, and variations in the quality of feedback. Students may complete only a few simulated interactions before starting clinical practice, leaving important gaps in their preparedness.

The introduction of AI-powered avatars, or Embodied Conversational Agents (ECAs), addresses these limitations by enabling structured, adaptive, and repeatable practice that mirrors real-world complexity. These agents can reproduce a wide spectrum of emotional responses and communicative styles, offering learners the chance to experiment, reflect, and improve in a safe environment. Unlike static simulations or scripted actors, avatars powered by generative AI respond dynamically, adapting to tone, pacing, and word choice, thereby providing an experience that is both immersive and tailored to the learner's performance [1–3].

At e-REAL Labs in New York and Turin—the research and development hub owned by Logosnet (Houston, US), Logos Knowledge Network (Lugano, CH), and Logos Centro Studi (Turin, IT)—we have advanced this vision by designing **Conversation Master©, a proprietary solution for medical education and training** [4]. Building on simulation-based education principles, Conversation Mastery integrates generative AI, natural language processing (NLP), and behavioral analytics to create adaptive, context-sensitive learning environments [5–7]. In collaboration with the Harvard Center for Medical Simulation (Boston, US), the system is being tested and refined within clinical training pathways, allowing for rigorous evaluation of its usability, effectiveness, and curricular integration.

Fig. 1. Representative embodied conversational agents (ECA), performing as avatars or digital humans, boosted by artificial intelligence, designed and developed by the e-REAL Labs

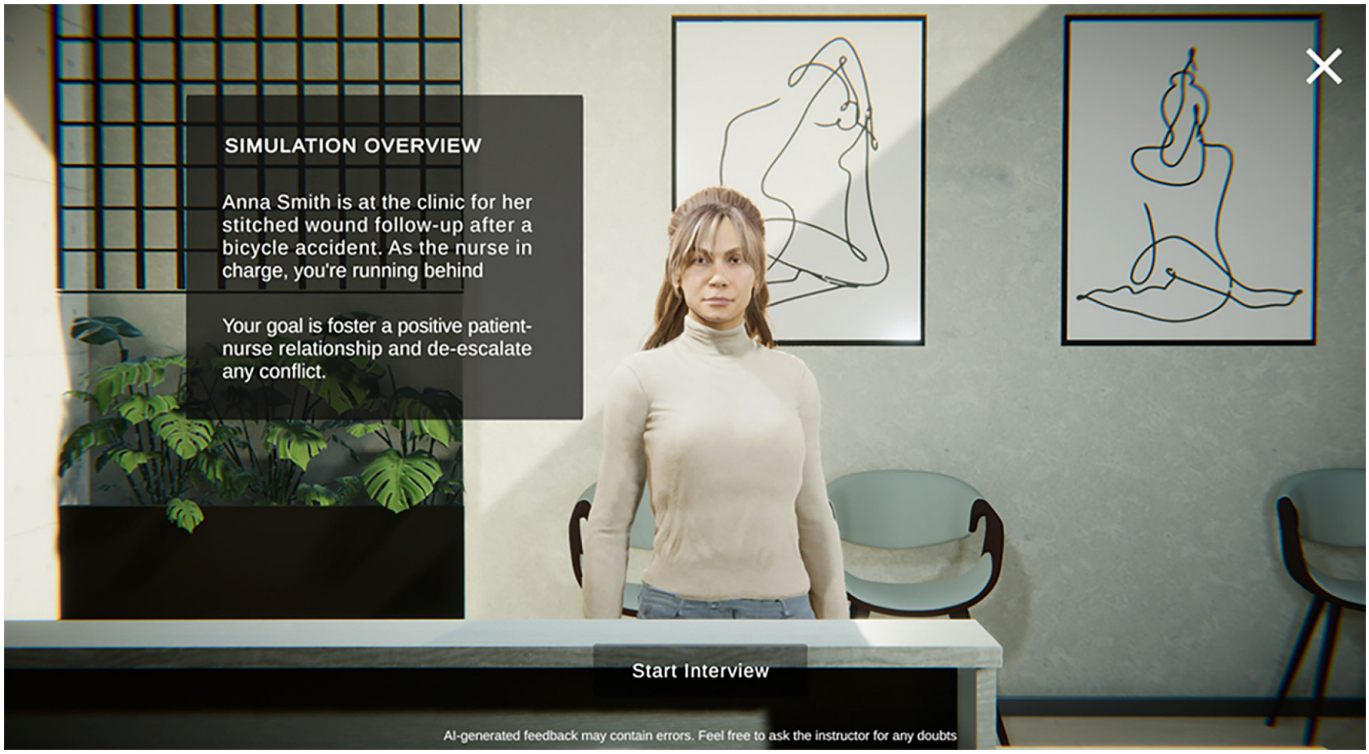

Fig. 2. An avatar within a simulation scenario (conflict de-escalation) implemented by the e-REAL Labs

1.2 Objectives of the study

The purpose of this study is to explore how AI-powered avatars can effectively enhance medical education by expanding opportunities for deliberate practice,

standardizing assessment, and boosting learners' readiness for high-stakes clinical encounters. While previous research has emphasized the value of simulation and standardized patients, there remains a critical need for scalable solutions that combine realism with repeatability. Conversation Mastery has been created to meet this need, providing healthcare students and professionals the chance to rehearse essential communication skills in environments that are both emotionally authentic and pedagogically rigorous.

The specific objectives of this paper are fourfold:

Effectiveness Analysis – To systematically compare AI-driven avatars with traditional human-led coaching and classroom teaching, assessing measurable outcomes such as conflict resolution success rates, the ability to handle emotionally charged encounters, and improvements in feedback delivery.

Skill Development Metrics – To evaluate key communication and professional competencies, including empathy, bias awareness, clarity of explanation, and procedural adherence, using validated rubrics that align with OSCE-based frameworks.

Ethical and Regulatory Considerations – To investigate the implications of deploying AI-driven avatars in clinical education, focusing on data protection, decision transparency, and compliance with regulatory frameworks such as GDPR, HIPAA, and the EU AI Act. This includes reflecting on the boundaries of emotional AI and ensuring that its role in education is ethically sound and culturally sensitive.

Implementation Model – To propose a structured and adaptable model for embedding AI-driven coaching into existing curricula for medical and nursing schools, residency programs, and continuing professional development. This includes guidance on integrating Conversation Mastery into diverse institutional contexts, ensuring that technology complements rather than replaces faculty expertise.

By addressing these objectives, the study seeks to demonstrate that AI-driven avatars are not merely novel digital tools but constitute a scalable and bias-sensitive educational paradigm. In doing so, it contributes evidence to support the integration of advanced conversational simulations into the broader ecosystem of simulation-based medical education.

1.3 Structure of the paper

This paper is organized into different sections, each addressing a distinct dimension of the study.

Section 2 – Background and Literature Context: provides an expanded review of existing approaches to clinical communication training, including lectures, standardized patients, and role-play, highlighting their strengths and limitations. It also situates AI-powered avatars within the broader evolution of simulation-based education.

Section 3 – The Conversation Mastery Platform: introduces the architecture, features, and pedagogical design of the solution, with particular emphasis on its adaptive generative AI engine, branching scenario structure, and real-time feedback system. This section also outlines the Responsible Technology Principles guiding development and deployment.

Section 4 – Scenario-Based Training in Medical Education: presents the portfolio of scenarios tested and validated within the platform, organized into seven clusters (e.g., early clinical exposure, conflict de-escalation, serious illness disclosure). Each cluster is analyzed for its pedagogical contribution and adaptability to cultural and linguistic contexts.

Section 5 – Methods and Results of the Pilot Study: details the research design and comparative evaluation of AI-powered avatars against traditional methods.

Quantitative and qualitative outcomes are presented, including learner performance metrics, confidence measures, and adherence to best practices.

Section 6 – Discussion and Implications: interprets the findings, considering implications for curriculum integration, ethical and regulatory compliance, and long-term scalability in medical and nursing education.

Section 7 – Conclusions and Future Directions: summarizes the main contributions of the study and outlines priorities for further research, including longitudinal evaluations, cross-institutional pilots, and the potential role of avatars in continuing professional development.

This structure ensures a coherent progression from context-setting, through platform presentation and empirical evidence, to reflection and projection, thereby situating Conversation Mastery as a robust and future-oriented contribution to medical education.

1.4 Background and literature context

Communication is widely recognized as a cornerstone of medical practice, influencing patient satisfaction, adherence to treatment, and even clinical outcomes. Over the past decades, medical education has progressively integrated structured approaches to communication skills training, ranging from lectures and workshops to more immersive modalities such as standardized patients (SPs), high-fidelity manikins, and role-play among peers. Each of these approaches contributes valuable learning opportunities, yet persistent gaps remain in terms of scalability, consistency, and sustainability.

Lectures and classroom-based learning provide foundational knowledge and theoretical frameworks for professional communication. They are cost-effective and easy to scale across large cohorts but offer limited opportunities for experiential practice or individualized feedback.

Role-play and peer-to-peer simulations introduce interactivity, allowing learners to rehearse clinical scenarios. However, outcomes are often inconsistent, as the quality of the exercise depends heavily on the skills of peers or instructors playing the patient role. Moreover, feedback is typically subjective, shaped by interpersonal dynamics rather than standardized criteria.

Standardized patients (SPs), widely used in Objective Structured Clinical Examinations (OSCEs), represent a more reliable method for simulating real-world encounters. Trained actors can replicate patient presentations and provide feedback to learners. Despite their value, SP-based simulations face practical constraints: they are costly, time-consuming, logistically demanding, and limited in availability. Learners may only have access to a few SP encounters throughout their education, reducing opportunities for deliberate practice and longitudinal skill development.

Simulation-based education with high-fidelity manikins has transformed technical skills training (e.g., resuscitation, anesthesia, emergency medicine). However, while these technologies excel at procedural and physiological replication, they fall short in capturing the subtleties of human communication, emotion, and bias management—competencies that are increasingly recognized as crucial to patient-centered care.

Against this backdrop, the emergence of AI-powered avatars or Embodied Conversational Agents (ECAs) represents a pivotal advancement. These avatars combine natural language processing, generative AI, and behavioral analytics to create dynamic and adaptive interactions. Unlike SPs, avatars do not fatigue, vary in performance, or require scheduling. Unlike scripted role-play, they respond flexibly

to tone, vocabulary, and conversational strategy, ensuring that no two encounters are identical. By embedding assessment rubrics within the system, AI avatars can also deliver standardized, objective, and real-time feedback, addressing one of the most critical limitations of human-based training.

Fig. 3. Mia Williams is an embodied conversational agent, or avatar, created at the e-REAL Labs and available online and on-site (multimodal), e.g., projected over a wall or on a totem. This avatar is also an agentic application because it is an AI-driven system that autonomously adapts, learns, and interacts with users, making context-aware decisions to optimize engagement and outcomes in real time, such as playing a video or displaying a chart

Fig. 4. Other representative multimodal avatars from the e-REAL Labs

Fig. 5. These images highlight AI-driven simulation environments designed by the e-REAL Labs, providing interactive and immersive learning experiences that enhance decision-making, communication, and crisis management skills. The same experiences are also available online as a desktop solution: reduced immersion, same interaction

Fig. 6. These images depict AI-driven avatars in a virtual simulation for conflict de-escalation and empathy training. The simulation allows learners to navigate emotionally sensitive conversations in healthcare settings, improving communication and patient-centered care

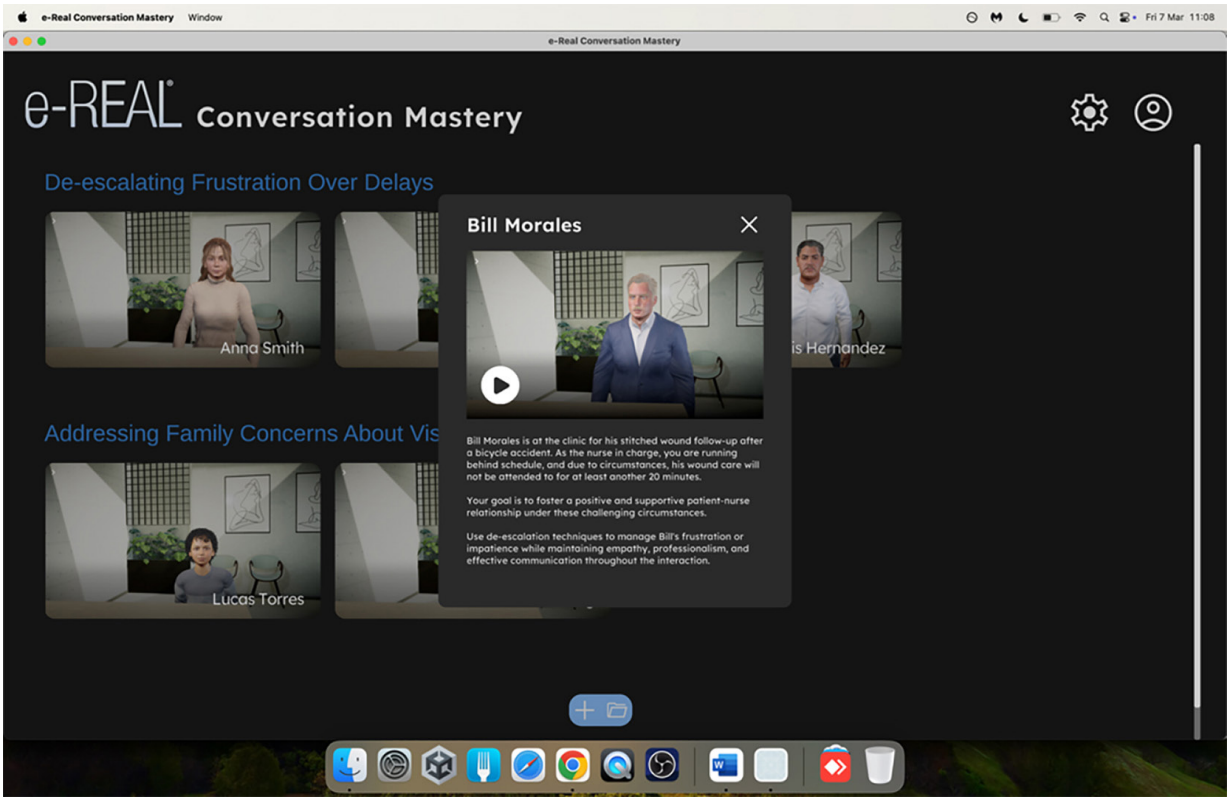

Fig. 7. This image illustrates the selection of initial avatars to begin a scenario that challenges learners to de-escalate frustration and effectively manage patient concerns in a delayed care situation. By utilizing AI-driven avatars, the platform allows healthcare professionals to practice empathetic communication, professionalism, and conflict resolution skills in realistic, high-pressure interactions

Fig. 8. This image highlights an immersive medical simulation environment created at the e-REAL Labs, where AI-driven avatars and a realistic hospital setting form a high-fidelity training space for healthcare professionals to practice critical decision-making and patient interactions

In the literature, the concept of deliberate practice—structured, repetitive engagement with feedback—has been strongly linked to expertise development in clinical and non-clinical domains. AI-driven conversational platforms align closely with this principle by allowing unlimited rehearsal of complex communication scenarios. They further support competency-based medical education (CBME) by mapping performance directly to curricular outcomes and assessment frameworks, including OSCE standards. Recent studies in digital health education underscore the potential of AI-enabled simulation to enhance learner confidence, reduce cognitive bias, and improve transfer of skills to clinical settings. Yet, despite growing enthusiasm, empirical evidence remains limited, particularly regarding integration into full curricula, long-term impact, and ethical considerations such as data privacy and transparency of AI decision-making. This study seeks to contribute to this emerging body of knowledge by presenting evidence from the Conversation Mastery platform and situating it within the broader trajectory of simulation-based medical education. By doing so, it highlights both the opportunities and the challenges of adopting AI-driven avatars as a complement to established methods, offering a bridge between traditional teaching, standardized assessment, and next-generation educational technologies.

Conversation Mastery distinguishes itself from traditional AI coaching tools through its **high adaptability, real-time contextual awareness, and the ability to personalize learning** based on individual user performance. Unlike static chatbots or scripted training modules, our avatars are designed to:

- **Provide immediate and context-specific feedback**, dynamically adjusting responses based on user behavior.
- **Incorporate emotional intelligence**, adapting feedback delivery based on **tone analysis and emotional cues**.
- **Support bias mitigation training**, helping users recognize and adjust unconscious biases in **healthcare decision-making and interactions**.
- **Facilitate multi-user interactions**, enabling team-based simulations that improve **collaborative problem-solving skills**.

The AI-powered feedback system is a core feature of the *Conversation Mastery* platform. It provides learners with structured, real-time insights aligned with institutional rubrics. After each interaction, participants receive feedback mapped to their institution's specific assessment framework, ensuring consistency and relevance. The system evaluates key communication skills—such as clarity in introductions, acknowledgment of emotions, problem-solving effectiveness, and explanation structure—while dynamically adapting feedback to individual performance. For example, in a conflict de-escalation scenario, if a participant demonstrates active listening and validates the other party's concerns, the system reinforces these behaviors with positive feedback. Conversely, if key de-escalation techniques—such as maintaining a calm tone or offering clear, step-by-step resolutions—are missing, the platform provides targeted suggestions for improvement. By aligning AI-driven feedback with institutional rubrics, *Conversation Mastery* ensures a structured, context-aware learning experience that directly enhances real-world communication skills.

Instructors and learners benefit from a structured recap of their interactions, with performance indicators visually represented to facilitate reflection and skill refinement. By integrating AI-generated insights with human coaching, the platform ensures a blended learning approach, where technology enhances but does not replace expert instruction. The system's multimodal feedback mechanism allows for continuous improvement, making training sessions more engaging, adaptive, and results-driven [8].

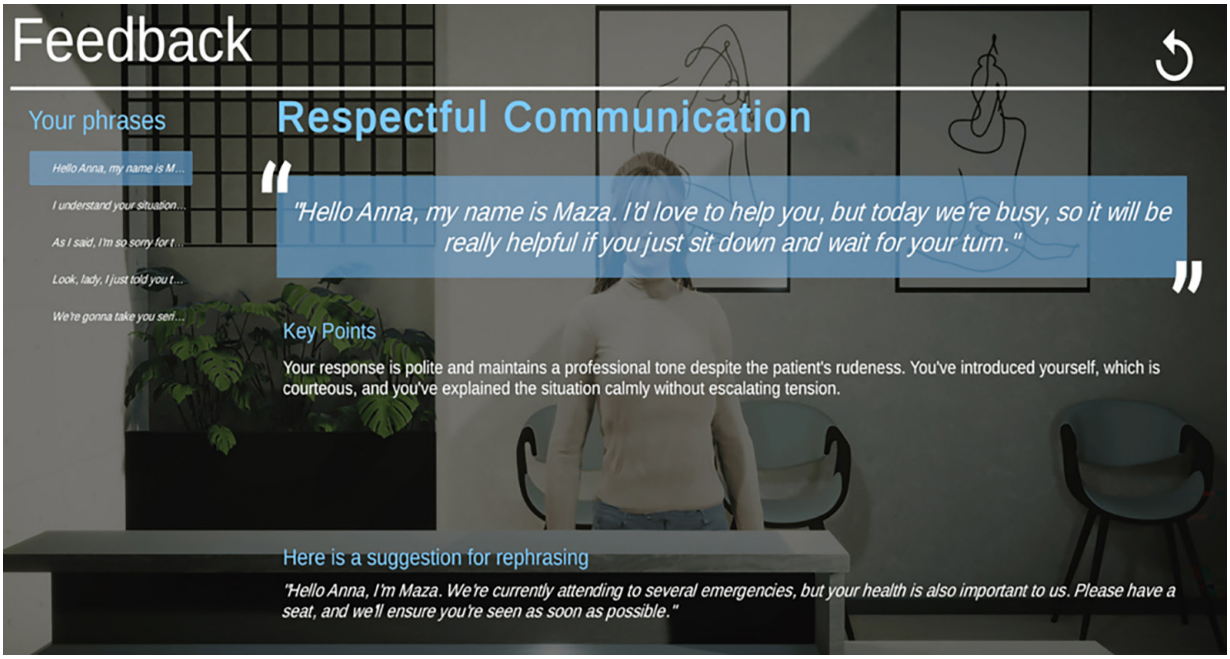

Fig. 9. It illustrates the AI-driven feedback system within the e-REAL Conversation Mastery platform, which focuses on respectful communication. The system evaluates the participant's response, highlighting strengths such as maintaining professionalism despite patient frustration. Additionally, it offers a refined rephrasing suggestion, ensuring a more empathetic and patient-centered approach. This feature helps learners enhance their de-escalation techniques and communication clarity in real-world interactions

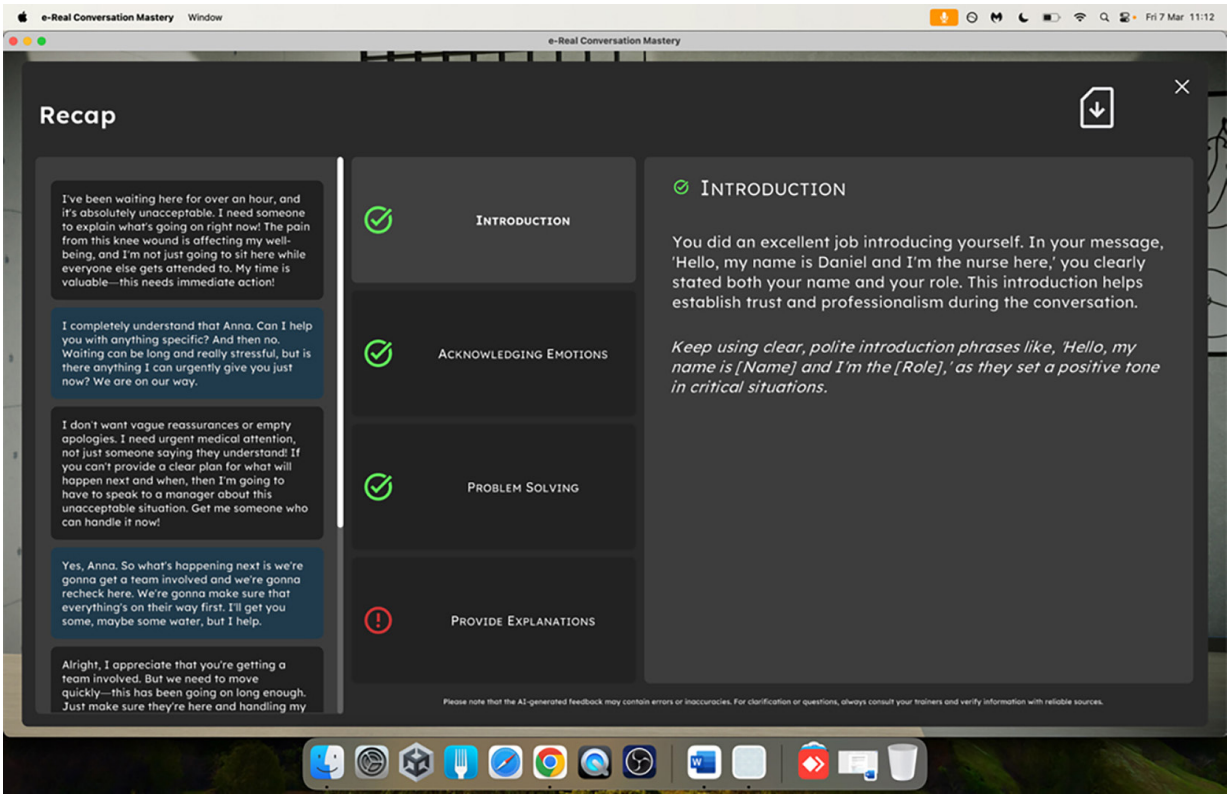

Fig. 10. This image presents a different feedback layout: in this case the recap categorizes responses into key competencies such as introduction, acknowledging emotions, problem-solving, and providing explanations, with visual indicators highlighting areas of strength and improvement. The system provides specific, actionable feedback, reinforcing effective communication techniques while guiding learners to refine their approach in real-world de-escalation scenarios

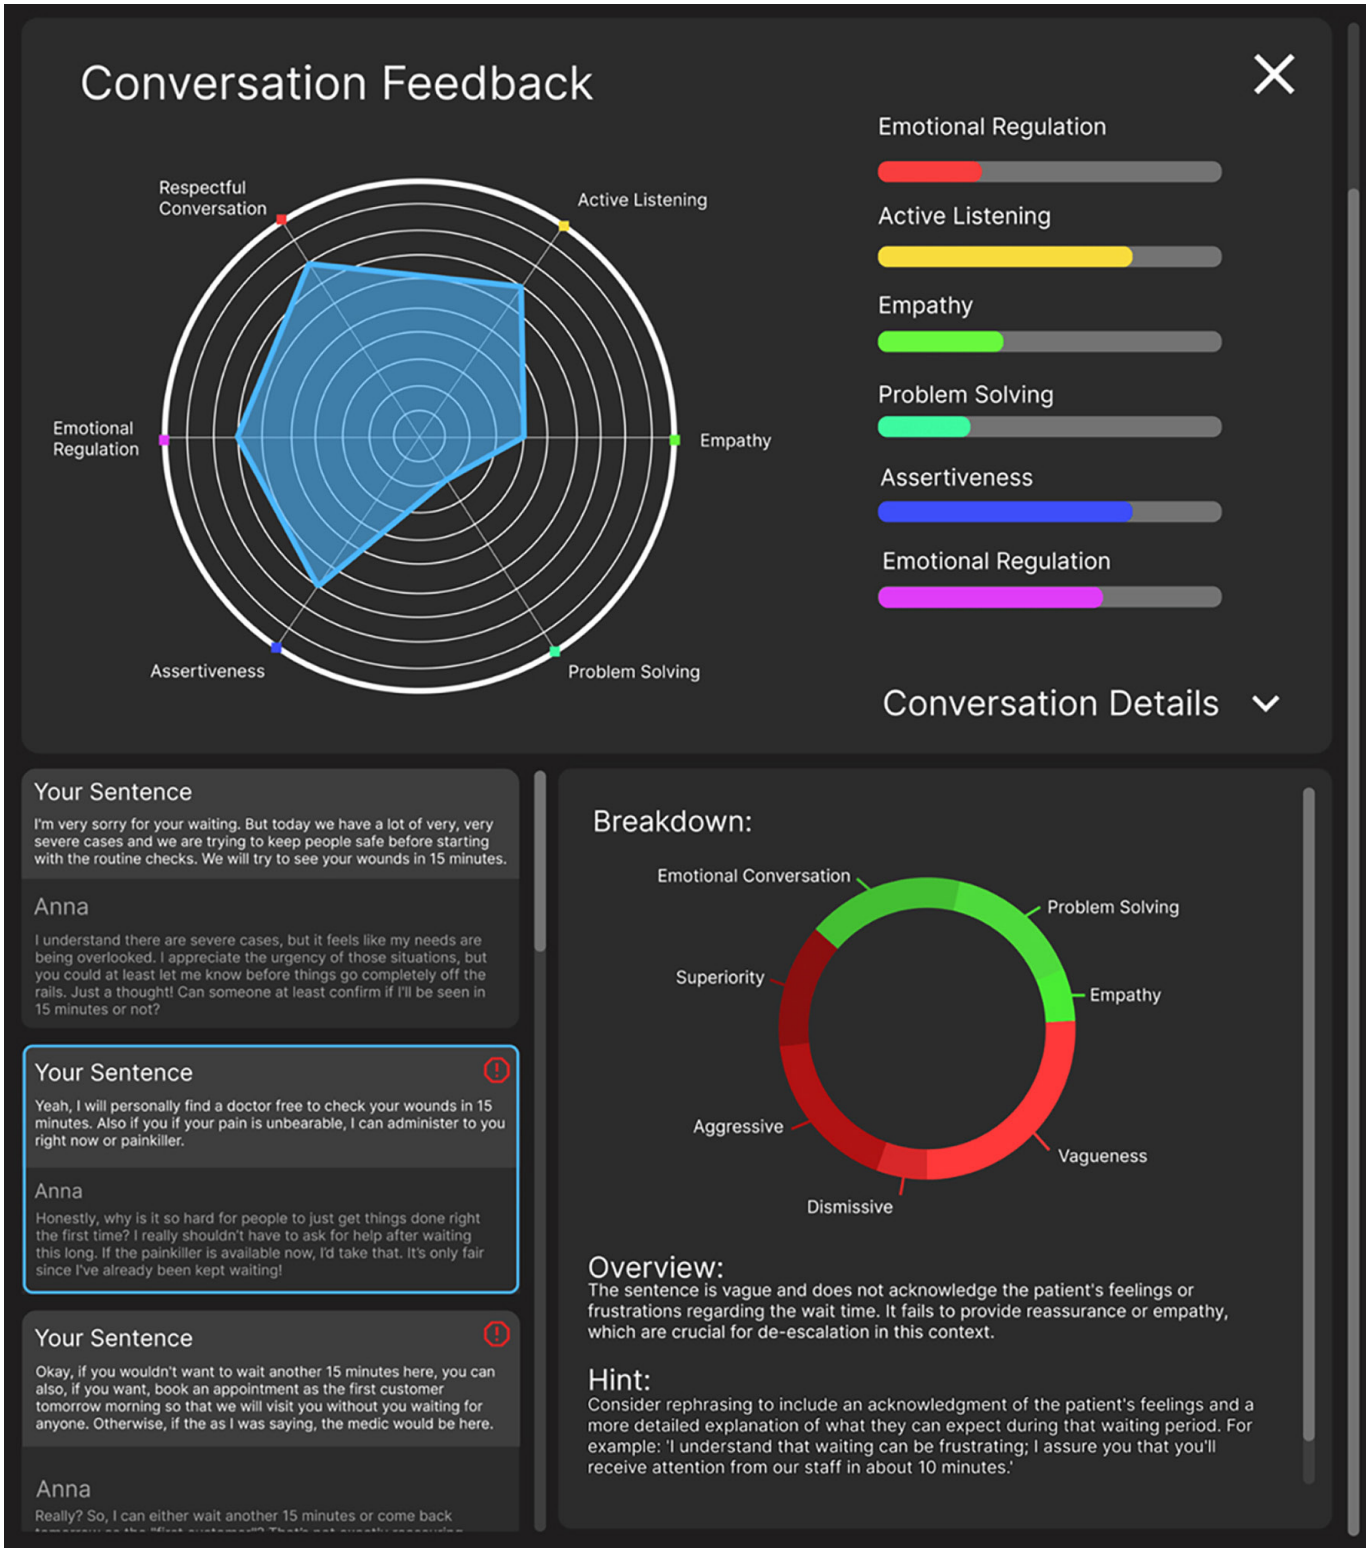

Fig. 11. This image presents detailed AI-driven conversation feedback, using a radar chart and sentiment breakdown to assess key communication skills such as empathy, active listening, problem-solving, and emotional regulation. The system highlights areas for improvement, particularly in addressing patient frustration and avoiding vague or dismissive language. The overview and hint section provide actionable guidance, encouraging learners to adopt more empathetic and reassuring communication strategies, making this an effective tool for refining de-escalation and patient interaction skills

By leveraging deep learning algorithms and multimodal AI, Conversation Mastery ensures that training sessions remain engaging, adaptive, and directly applicable to real-world challenges – improving learners' readiness.

1.5 The Conversation Mastery platform

Conversation Mastery is a proprietary platform designed to align with Objective Structured Clinical Examination (OSCE) standards, supporting both formative and summative assessment. Learners are immersed in a broad range of scenarios—covering early patient encounters, disclosure of serious diagnoses, end-of-life discussions, conflict de-escalation, interprofessional communication, and patient education—each adaptable by difficulty, cultural context, and linguistic variation.

The system combines multiple technological layers into a seamless user experience:

- ***Student Input Interface*** – Learners interact with avatars using text or speech, supported on laptops, mobile devices, immersive displays, or VR setups.
- ***Generative AI Conversational Engine*** – Provides context-aware, multilingual dialogue across 40+ languages.
- ***Deterministic Validation Layer*** – Ensures clinical accuracy and compliance with curricular guidelines, preventing unsafe or misleading outputs.
- ***3D Simulation Environment*** – High-fidelity avatars and adaptive backgrounds that adjust to scenario complexity.
- ***Feedback and Analytics Engine*** – Delivers real-time scoring, identifies missed opportunities, and generates structured reports for learners and faculty.

This architecture ensures that every interaction is both authentic and accountable, blending the adaptability of generative AI with the reliability of rule-based clinical logic.

A distinguishing feature of Conversation Mastery is its portfolio of validated scenarios. Organized into seven main categories, these scenarios span the full spectrum of clinical communication challenges:

- *Early Clinical Exposure* (e.g., patient introductions, informed consent, explaining procedures).
- *Serious Illness & Prognosis Disclosure* (e.g., delivering bad news, managing denial, cross-cultural adaptations).
- *End-of-Life Conversations* (e.g., palliative care, family disagreements, withdrawal of life-sustaining treatment).
- *Conflict De-Escalation* (e.g., patient frustration, family aggression, discriminatory comments).
- *Emotionally Charged Encounters* (e.g., anxiety, stigma, treatment refusal).
- *Interprofessional Communication* (e.g., handovers, addressing unprofessional conduct, resolving disagreements).
- *Patient Education and Health Promotion* (e.g., diabetes self-management, smoking cessation, safe medication use).

Each scenario is adjustable by emotional intensity and complexity, supporting progressive learning from structured to high-stakes encounters. All cases are adaptable linguistically and culturally, ensuring relevance across diverse clinical contexts.

The platform is guided by a set of Responsible Technology Principles, ensuring that innovation remains human-centered, ethical, and inclusive:

- ***Human-Centered Immersion*** – Simple interfaces, multimodal interaction (speech, gesture, visual cues), and adaptive feedback loops aligned with Rapid Cycle Deliberate Practice.
- ***Inclusive Accessibility*** – Avatars available in >40 languages, scenarios tailored to cultural contexts, and compliance with WCAG 2.1 AA and EN 301 549 standards for accessibility.
- ***Ethical AI*** – Transparent design, explicit disclosure of the avatars' digital nature, and rigorous consent procedures for any biometric or learning data.
- ***Evidence-Based Impact*** – Simulations built on cognitive psychology and constructivist pedagogy, benchmarked to deliver 3 hours of training impact equivalent to ~8 hours of traditional classroom learning.
- ***Privacy and Data Governance*** – Compliance with GDPR, HIPAA, and FERPA; anonymized learning data retained for ≤24 months.
- ***Sustainability and Scalability*** – Modular architecture (portable or fixed labs, VR, digital signage) designed for reuse and long-term adaptability.
- ***Partnership and Co-Design*** – Faculty and institutional partners actively involved in scenario customization and iterative improvement.

Taken together, these features position Conversation Mastery as a sole-source solution, unavailable from competing vendors and uniquely capable of combining adaptive generative AI with deterministic clinical safeguards. By offering standardized yet flexible training opportunities, the platform reduces reliance on scarce human resources, supports faculty with intuitive editing tools, and ensures learners can repeatedly practice complex skills without risk to patients.

2 SCENARIO-BASED TRAINING IN MEDICAL EDUCATION

2.1 Rationale for scenario-based learning

Scenario-based learning has long been a cornerstone of simulation in healthcare, offering opportunities to practice communication and decision-making in safe yet realistic environments. What distinguishes Conversation Mastery is the breadth, adaptability, and progressive calibration of its scenarios, each aligned with OSCE (Objective Structured Clinical Examination) frameworks. This ensures not only pedagogical rigor but also seamless integration into both formative and summative assessments.

Unlike standardized patients (SPs), who are constrained by cost, availability, and human variability, AI-powered avatars can deliver a consistently high-fidelity experience across a wide range of situations, while also being customizable in language, culture, and clinical focus. This combination allows institutions to replicate the diversity of real-world encounters and to ensure equity of learning opportunities across large cohorts.

2.2 Scenario categories

The scenario library is structured into seven main clusters, each addressing a distinct competency domain. Every scenario can be delivered at multiple levels

of complexity, from low-stress introductions to high-stakes emotionally charged interactions, thereby supporting longitudinal progression and deliberate practice.

2.3 Early clinical exposure

Designed for novice learners, these scenarios focus on rapport-building, obtaining consent, and explaining simple procedures. They help students develop confidence during their first patient interactions. Examples include:

- Introducing oneself to a patient and clarifying roles.
- Taking a structured patient history while fostering rapport.
- Explaining a routine procedure (e.g., blood draw) in jargon-free language.
- Communicating delays in care or the unavailability of a specialist.

2.4 Serious illness and prognosis disclosure

These scenarios guide learners through some of the most delicate aspects of medical practice—delivering bad news and discussing prognosis. The avatars reproduce a variety of patient and family reactions, from denial to intense grief. Examples include:

- Communicating a cancer diagnosis with clarity and empathy.
- Discussing poor prognosis with a family member who resists acceptance.
- Handling cross-cultural differences in openness about terminal illness.
- Explaining unexpected surgical complications with transparency.

2.5 End-of-life conversations

End-of-life care requires clinicians to balance compassion with medical and ethical clarity. Scenarios include:

- Introducing hospice or palliative care options to a supportive or resistant family member.
- Discussing withdrawal of life-sustaining treatment in the ICU.
- Negotiating care goals with families whose cultural or religious views complicate decision-making.

2.6 Conflict de-escalation

This cluster addresses a growing need in clinical practice: managing aggression, frustration, and mistrust. Learners practice acknowledging emotions without defensiveness, setting boundaries, and maintaining professionalism. Examples include:

- Calming a patient upset by long waiting times.
- Responding to anger over visitation restrictions.

- Addressing discriminatory remarks from patients or families.
- Restoring trust after accusations of neglect or poor organization.

2.7 Emotionally charged encounters

These scenarios emphasize empathy and emotional intelligence, helping clinicians navigate refusal, anxiety, and stigma. Examples include:

- Supporting a patient anxious before surgery.
- Consulting with a patient living with chronic pain and frustration.
- Managing refusal of treatment while respecting autonomy.
- Communicating with families minimizing psychiatric symptoms.

2.8 Interprofessional communication

Effective teamwork is a critical determinant of patient safety. These scenarios extend beyond doctor–patient exchanges to include intra-team dynamics. Examples include:

- Delivering structured handovers under time pressure (e.g., SBAR).
- Clarifying contradictory information during handover.
- Addressing unprofessional behavior or microaggressions from colleagues.
- Resolving open disagreements about treatment plans in front of patients.

2.9 Patient education and health promotion

Preventive care and patient empowerment are central to modern healthcare. Scenarios in this cluster support learners in tailoring explanations to different levels of health literacy. Examples include:

- Educating patients newly diagnosed with diabetes on self-management.
- Encouraging smoking cessation with motivational interviewing techniques.
- Explaining safe medication use to older adults with sensory limitations.
- Promoting vaccination during routine check-ups.

2.10 Pedagogical strengths

Several features make these scenarios particularly valuable in medical education:

- ***OSCE Alignment*** – Scenarios are structured to map directly to OSCE domains (communication, professionalism, empathy, cultural competence).
- ***Progressive Difficulty*** – Encounters can evolve from structured scripts to fully adaptive dialogues, matching learner growth.
- ***Cultural and Linguistic Adaptability*** – Each case can be localized to reflect the patient demographics of a given institution or region.

- ***Emotional Breadth*** – Learners are exposed to the full spectrum of emotions, from mild anxiety to high-intensity grief or aggression.
- ***Interprofessional Dimension*** – Training is not limited to patient communication but also supports team-based competencies.

2.11 Implications for clinical readiness

By spanning early exposure to high-stakes encounters, the Conversation Mastery scenario portfolio creates a spiral curriculum of deliberate practice. Students can revisit scenarios at increasing levels of difficulty, reinforcing competencies through repetition and reflection. Moreover, the platform enables institutions to move beyond sporadic simulation events toward **continuous and scalable integration of communication training**, ensuring that learners graduate with the readiness to navigate not only technical procedures but also the emotional and ethical complexities of patient care.

3 METHODS AND RESULTS OF THE PILOT STUDY

Our collaboration with the Harvard Center for Medical Simulation enables us to rigorously test and refine Conversation Mastery in real-world healthcare settings. This partnership allows us to:

- Validate the efficacy of AI-driven coaching and training using data from real medical simulation programs.
- Improve scenario realism and responsiveness based on insights from medical educators and healthcare professionals.
- Ensure compliance with medical education best practices, aligning AI-driven training with competency-based evaluation frameworks.

Key training use cases include:

- *Conflict De-Escalation*: Simulating high-stress scenarios to train professionals in managing difficult conversations with patients and colleagues.
- *Bias Awareness and Diversity, Equity and Inclusion Training*: Using AI avatars to challenge cognitive, cultural and intercultural biases in healthcare interactions and decision-making.
- *Performance Feedback and Leadership Coaching*: Helping medical educators and clinical supervisors refine their coaching techniques through real-time AI-driven debriefing.

By embedding Conversation Mastery into simulation-based medical education and training, we are promoting scalable, effective, and bias-aware AI-driven coaching that adapts to diverse learning environments.

3.1 Study design

The pilot study was conducted in collaboration with the Harvard Center for Medical Simulation (CMS), involving a cohort of nursing students enrolled in a

communication skills module. The primary aim was to compare the effectiveness of three instructional modalities in preparing learners for conflict de-escalation and emotionally charged conversations:

1. ***Lecture-based instruction*** – conventional classroom delivery using slides, exemplars, and guided discussion.
2. ***Traditional simulation*** – role-play with trained standardized patients (SPs), supported by instructor debriefing.
3. ***AI-driven avatar training*** – repeated practice with Conversation Mastery avatars, offering adaptive branching dialogues and structured feedback.

3.2 Participants and procedure

- **Sample size:** 90 students, randomized equally across the three groups (n = 30 per group).
- **Duration:** Each group received two hours of instruction, followed by a structured practice component (traditional or AI-based).
- **Assessment tool:** Performance was measured using a validated rubric covering empathy, communication clarity, emotional regulation, bias awareness, and adherence to institutional best practices.
- **Evaluator calibration:** Two independent faculty assessors scored each encounter, with inter-rater reliability calculated to ensure consistency.

3.3 Measures

Three categories of outcomes were examined:

- **Success rate** – ability to resolve or de-escalate the conflict presented in the scenario.
- **Confidence** – self-reported readiness to apply communication strategies in clinical practice.
- **Adherence to best practices** – alignment with guidelines on conflict resolution and patient-centered communication.

3.4 Results

The findings indicated a clear performance gradient across the three groups:

- **Lecture group:** 45% average success rate in conflict de-escalation; moderate theoretical understanding but low transfer to practice.
- **Traditional simulation group:** 67% success rate; higher realism than lectures, but limited by the availability of SPs and variability in performance.
- **AI-driven avatar group:** 89% success rate; significantly higher confidence scores and stronger adherence to best practices.

Students in the avatar group also engaged in substantially more practice sessions—an average of 12.3 encounters per learner compared with 2–3 in the SP group.

This reinforced the value of deliberate practice, where repetition and adaptive feedback drive skill acquisition.

3.5 Qualitative feedback

Open-ended reflections collected after the sessions revealed additional insights:

- **Lecture group:** students appreciated theoretical clarity but reported "not knowing how to apply strategies under pressure."
- **Traditional simulation group:** participants valued realism but noted variability in the feedback they received.
- **AI-driven avatar group:** learners highlighted the opportunity for "endless rehearsal without judgment," with several emphasizing how the structured feedback helped them identify blind spots in empathy or tone.

3.6 Limitations

As a pilot study, the sample size was modest, and results are limited to one institution and one competency area (conflict de-escalation). Further research with multi-institutional cohorts, diverse scenarios, and longitudinal follow-up is needed to confirm sustained impact on clinical performance.

3.7 Key takeaways

The pilot results suggest that AI-driven avatars complement and enhance existing training modalities, offering scalable, repeatable, and bias-aware practice that translates into measurable gains in learner performance. While lectures provide conceptual grounding and simulations with SPs foster realism, Conversation Mastery™ enables high-frequency, adaptive practice that addresses both the cognitive and emotional demands of clinical readiness.

Metric	Group A (Lecture)	Group B (Simulation)	Group C (AI Avatars)
Avg. Number of Practice Sessions per Student	0	1	12.3
Success Rate (%)	45%	67%	89%
Avg. Strategy Application Score (1-5)	2.1	3.4	4.6
Avg. Confidence Rating (1-5)	2.3	3.1	4.5
Avg. Comfort in Conflict Situations (1-5)	2.0	3.0	4.3
Adherence to Best Practices (%)	38%	62%	85%

Fig. 12. This table compares the effectiveness of lecture-based, simulation-based, and AI-driven training in conflict de-escalation for nursing students. Increased practice opportunities correlate with higher success rates, confidence, and adherence to best practices. AI-driven training (Group C) showed the highest success (89%), suggesting its value as a complementary tool alongside instructor-led methods.

4 DISCUSSION AND IMPLICATIONS

4.1 Curriculum integration and faculty workflow

Conversation Mastery fits naturally within competency-based medical education by mapping scenario design and feedback to OSCE domains and validated rubrics. The platform's Competency Practice Module (3 scenarios × 3 variations) structures repetition with variation, delivering ~9 distinct encounters per competence in compact 10–15 minute sessions and recommending ~20 dialogues (~3.5 hours) to reach an initial level of operational readiness—formatting that scales across lectures, small groups, and self-practice without altering learning efficacy.

Faculty retain pedagogical control through the Conversation Mastery Editor, which allows ongoing adaptation and creation of scenarios aligned to local protocols, languages, and cultural norms, without extra licensing fees; this supports true co-design and iterative curricular refinement. Integration with institutional systems can be staged: cohort-level analytics for classroom use; pseudonymous individual tracking for formative coaching; and full LMS/SSO interoperability (SCORM/xAPI/LTI, grade passback) for summative uses and program-level dashboards.

This layered approach enables schools to begin with a lightweight approach and progressively formalize assessment as evidence accumulates.

4.2 Scenario portfolios as engines of deliberate practice

The breadth and calibratability of the seven scenario clusters—ranging from Early Clinical Exposure to End-of-Life and Interprofessional Communication—allow for the creation of spiral curricula that revisit the same competence with increasing intensity while maintaining contextual realism (e.g., denial, cultural expectations, microaggressions, discriminatory remarks). Because each case is linguistically and culturally adaptable, institutions can mirror local patient populations and policy environments, strengthening transfer to practice.

4.3 Ethics, safety, and governance

From a governance perspective, Conversation Mastery implements privacy-by-design aligned with GDPR, HIPAA, and FERPA, limits processing to training delivery and feedback, and explicitly does not store biometric data; anonymized learning data follow defined retention windows.

Responsible Technology Principles further codify transparent agent behavior, explicit disclosure of the avatars' digital nature, consent management, and accessibility requirements (WCAG 2.1 AA/EN 301 549) to ensure inclusion and learner autonomy. These safeguards are pertinent to current debates about emotion-related AI: in this context, emotional cues are used to improve conversational realism and feedback rather than for profiling, consistent with ethical boundaries articulated in the platform documentation.

4.4 Faculty role and educational culture

AI-driven coaching is not a substitute for expert educators; it amplifies their reach. Faculty time shifts from repeating baseline drills to targeted debrief and feedback

synthesis, informed by standardized analytics and conversation transcripts. The embedded rubrics reduce evaluator variability, while the editor empowers faculty to encode institutional norms and escalate scenario complexity over time.

4.5 Scalability, access, and delivery modalities

Operational scalability is supported by multi-modal delivery: web/desktop for ubiquitous access, digital totems for classroom and skills labs, and VR for high-immersion deployments; licenses and minimal technical specs are documented to streamline adoption.

International roll-outs benefit from multilingual capability (>40 languages) and cultural localization, with the sole-source nature of the platform ensuring a unified feature set and support model across sites.

4.6 Implications for assessment and program evaluation

Because practice time is counted only during active avatar interaction and immediate feedback, institutions can budget exposure precisely and compare cohorts on comparable workloads. At higher integration levels, grade passback and KPI dashboards enable alignment with course objectives and program outcomes, supporting CQI cycles and accreditation evidence.

4.7 Limitations and future work

Pilot evidence indicates superior success rates and confidence with AI-driven practice relative to lectures and single-shot simulations; nonetheless, larger multi-institutional, longitudinal studies are needed to verify durability of gains, downstream impact on patient experience/safety, and fairness across learner subgroups. Future work should examine optimal blending (e.g., pre-brief → avatar practice → human debrief), adaptive spacing for retention, and cross-cultural validity of scenario variants at scale.

4.8 Practical pathways to adoption

A pragmatic path begins with one Competency Practice Module in a targeted course (e.g., conflict de-escalation), using cohort analytics only; iterate scenarios with the editor, then add individual tracking and LMS integration as the curriculum matures. For high-visibility teaching spaces or OSCE prep, pair desktop access with a digital totem installation; consider VR selectively for cohorts that benefit from embodied perspective-taking. This staged approach balances educational rigor, cost, and change management while preserving faculty ownership of pedagogy.

5 CONCLUSIONS AND FUTURE DIRECTIONS

This study contributes evidence that AI-powered avatars can play a transformative role in medical education, particularly in preparing learners for complex, emotionally

charged, and ethically sensitive encounters. By offering adaptive, repeatable, and unbiased practice opportunities, Conversation Mastery addresses limitations inherent in traditional lectures, role-play, and standardized patient encounters. The pilot results demonstrate that learners trained with avatars not only achieve higher success rates in conflict de-escalation but also develop greater confidence and adherence to best practices.

Beyond quantitative outcomes, the qualitative feedback highlights an equally important dimension: learners valued the ability to practice "without judgment," repeat encounters as needed, and receive structured, objective feedback. These findings resonate with the principles of deliberate practice, suggesting that avatars can provide the high-frequency, feedback-rich rehearsal that underpins skill mastery.

From a broader perspective, Conversation Mastery represents more than a technological innovation; it is a pedagogical shift. Its design reflects a commitment to responsible AI use—ensuring transparency, inclusivity, accessibility, and data privacy. The system complements rather than replaces human faculty, freeing educators to concentrate on higher-order reflection and debriefing. By enabling cultural and linguistic customization, it also supports global scalability, reducing inequities in access to high-quality simulation-based education.

Looking ahead, several priorities emerge for future research and development:

1. ***Longitudinal Impact*** – Studies should examine whether improvements achieved through avatar-based practice translate into sustained performance gains in clinical settings.
2. ***Multi-Center Validation*** – Expanding pilots across diverse institutions and cultural contexts will test generalizability and support global adoption.
3. ***Curricular Integration Models*** – Research should explore optimal ways to embed avatar-based training within medical and nursing curricula, from undergraduate education to residency and continuing professional development.
4. ***Advanced Analytics*** – Future iterations of the platform may incorporate predictive learning analytics, enabling early identification of learners at risk of underperformance and supporting personalized remediation.
5. ***Ethical Horizons*** – Continued dialogue is needed around the boundaries of emotional AI in education, ensuring that innovation remains aligned with professional values and regulatory frameworks.

In conclusion, AI-driven avatars have the potential to establish a new paradigm of clinical readiness: one where learners repeatedly practice difficult conversations in safe yet authentic conditions, supported by faculty and guided by ethical AI. By bridging the gap between theoretical instruction and real-world complexity, Conversation Mastery™ contributes to a more competent, compassionate, and inclusive healthcare workforce.

6 REFERENCES

[1] C. Justine, S. Joseph, P. Scott, and C. Elizabeth, *Embodied Conversational Agents*. Cambridge, MA: The MIT Press, 2000. https://doi.org/10.7551/mitpress/2697.001.0001

[2] B. Lugrin, C. Pelachaud, and D. Traum, Eds., *The Handbook on Socially Interactive Agents: 20 years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 2: Interactivity, Platforms, Application,* 1st ed., ACM Books, Vol. 48. Association for Computing Machinery, New York, NY, USA, 2022. https://doi.org/10.1145/3563659

[3] M. Wahde and M. Virgolin, "Conversational agents: Theory and applications," in *Handbook on Computer Learning and Intelligence: Volume 2: Deep Learning, Intelligent Control and Evolutionary Computation*, 2022, pp. 497–544.
[4] Https://www.e-real.net
[5] S. Fernando, B. Barbara, G. Roxane, and M. Rebecca, "Effective extended reality: A mixed-reality simulation demonstration with intelligent avatars, digitized and holographic Tools," in *Innovative Approaches to Technology-enhanced Learning for the Workplace and Higher Education. Proceedings of The Learning Ideas Conference 2022*, G. David, A. Michael, and P. Antonella, Eds., Springer, Cham, 2023. https://doi.org/10.1007/978-3-031-21569-8_35
[6] S. Fernando, B. Barbara, and C. Ianna, "Intelligent digital humans for bias-free recruitment interviews: A diversity & inclusion training program," in *Creative Approaches to Technology-Enhanced Learning for the Workplace and Higher Education. Proceedings of the Learning Ideas Conference 2023,* G. David, A. Michael, and P. Antonella, Eds., Springer, Cham, 2023. https://doi.org/10.1007/978-3-031-41637-8_36
[7] S. Fernando and B. Barbara, *Learning 4.0. Advanced Simulation, Immersive Experiences and Artificial Intelligence, Flipped Classrooms, Mentoring and Coaching*. Angeli, Milan, 2018.
[8] S. Fernando, B. Barbara, C. Ianna, G. Roxane, R. Jenny, and M. Rebecca, "Feedback reimagined: Generative AI and conversational avatars as your new training partners," in *Creative Approaches to Technology-Enhanced Learning for the Workplace and Higher Education. Proceedings of the Learning Ideas Conference 2024.* G. David, A. Michael, and P. Antonella, Eds., Springer, Cham, 2024. https://doi.org/10.1007/978-3-031-73427-4_13

7 AUTHORS

Fernando Salvetti is with the e-REAL Labs at Logosnet, Turin, Italy – Lugano, Switzerland –Houston, TX, USA – New York, NY, USA (E-mail: salvetti@logosnet.org).

Barbara Bertagni is with the e-REAL Labs at Logosnet, Turin, Italy – Lugano, Switzerland –Houston, TX, USA – New York, NY, USA and also with the Center for Medical Simulation, Boston, MA, USA.

Roxane Gardner is with the Brigham and Women's Hospital/Children's Hospital/Massachusetts General Hospital/Harvard Medical School, Boston, MA, USA.

Jenny Rudolph is with the Center for Medical Simulation/Beth Israel Deaconess Medical Center, Boston, MA, USA.