

# Bridging Psychological Safety and Skill Guidance: An Adaptive Robotic Interview Coach

Wanqi Zhang  
University of Tennessee  
Knoxville TN USA  
wzhang79@vols.utk.edu

Jianguen He  
University of Tennessee  
Knoxville TN USA  
jianguen@utk.edu

Marielle Santos  
University of Tennessee  
Knoxville TN USA  
msanto10@vols.utk.edu

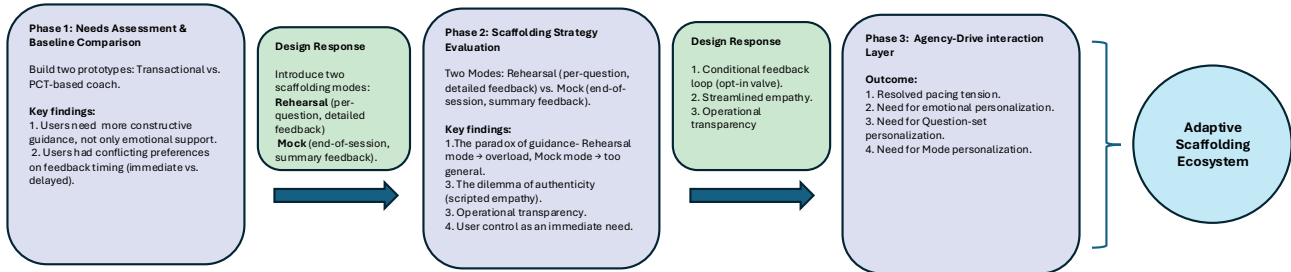


Figure 1: Three-phase developmental trajectory of the robotic coach, showing key findings and design changes culminating in the Adaptive Scaffolding Ecosystem.

## Abstract

Social robots hold promise for reducing job interview anxiety, yet designing agents that provide both psychological safety and instructional guidance remains challenging. Through a three-phase exploratory iterative design study ( $N = 8$ ), we empirically mapped this tension. Phase I revealed a “Safety–Guidance Gap”: while a Person-Centered Therapy (PCT) robot established safety, users felt insufficiently coached. Phase II identified a “Scaffolding Paradox”: rigid feedback caused cognitive overload, while delayed feedback lacked specificity. In Phase III, we resolved these tensions by developing an *Agency-Driven Interaction Layer*. Synthesizing our empirical findings, we propose the **Adaptive Scaffolding Ecosystem**—a conceptual framework that redefines robotic coaching not as a static script, but as a dynamic balance between affective support and instructional challenge, mediated by user agency.

## CCS Concepts

• **Human-centered computing** → *Empirical studies in HCI*; **User centered design**; *Empirical studies in interaction design*; • **Computer systems organization** → **Robotics**.

## Keywords

Human-Robot Interaction, interview training, social robots

## ACM Reference Format:

Wanqi Zhang, Jianguen He, and Marielle Santos. 2026. Bridging Psychological Safety and Skill Guidance: An Adaptive Robotic Interview Coach. In



This work is licensed under a Creative Commons Attribution 4.0 International License. *HRI Companion '26, Edinburgh, Scotland, UK*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2321-6/2026/03  
<https://doi.org/10.1145/3776734.3794502>

*Companion Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI Companion '26), March 16–19, 2026, Edinburgh, Scotland, UK.* ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3776734.3794502>

## 1 Introduction

Reducing social anxiety and improving performance-related skills are central promises of social robots. While screen-based AI systems expand access to interview training [3, 13], physically embodied agents elicit a stronger sense of social presence [2, 6] and reduce perceived social pressure compared to humans [5, 7]. This combination positions social robots as ideal partners for high-stakes evaluative contexts such as job interviews, where anxiety can significantly undermine performance [9].

Existing interventions typically bifurcate into two approaches: skill-focused training (e.g., AI platforms analyzing non-verbal behavior [4]) or exposure-based therapy (e.g., VR simulations for desensitization [3]). However, few systems successfully integrate both. While cognitive-behavioral frameworks dominate robot-based interventions [10], Person-Centered Therapy (PCT)—emphasizing genuineness, unconditional positive regard, and empathic understanding—offers a powerful mechanism for establishing immediate psychological safety and mitigating acute anxiety [11]. Yet, safety alone is insufficient for skill acquisition. Effective skill development requires instructional scaffolding: structured support that guides learners through their Zone of Proximal Development (ZPD) [14, 15].

A central challenge in Human–Robot Interaction (HRI) is addressing the imbalance between cognitive tutoring and affective support [16]. Prior educational research suggests that *adaptive scaffolding*—which dynamically adjusts support according to the learner’s evolving understanding—outperforms fixed strategies in fostering self-regulated learning [1]. However, excessive external

regulation can undermine learner agency [8], creating a tension between providing guidance and preserving autonomy.

To explore this tension, we conducted a three-phase iterative design study ( $N = 8$ ). We assessed the “Safety–Guidance Gap” in a purely empathetic robot (Phase I), identified a “Scaffolding Paradox” in conflicting feedback strategies (Phase II), and resolved these tensions through an Agency-Driven Interaction Layer (Phase III). This study addresses the following research questions:

- **RQ1:** How can we design a PCT-based robot for establishing psychological safety in interview coaching?
- **RQ2:** How do different scaffolding strategies (immediate vs. delayed) affect cognitive load and perceived utility in robot-supported learning situations?

We make the following contributions:

- **Preliminary evidence on the “scaffolding paradox.”** We find that while Rogerian empathy establishes safety, it lacks instructional utility. Conversely, rigid feedback introduces cognitive costs, illustrating a practical tension in HRI design.
- **A preliminary framework for adaptive scaffolding.** We outline an ecosystem in which robotic coaches dynamically calibrate affective tone and intervention intensity, mediated by user agency, allowing systems to move between “nurturer” and “instructor” roles as user needs evolve.

## 2 System Implementation and Methodology

### 2.1 System Architecture

To ensure ecological validity, we utilized the Misty II social robot to simulate physical presence. The robot was controlled via a custom web-based application (React.js/Node.js) integrating the OpenAI GPT-4 Realtime API. This architecture enabled low-latency, multimodal interactions, where system prompts guided verbal and non-verbal behaviors (e.g., nodding) over WebRTC. Researchers monitored system stability via a dashboard without intervening in the autonomous dialogue.

### 2.2 Iterative Design Study Design

We adopted a three-phase user-centred design process over Three weeks. The study employed a “within-subjects design”, engaging the same cohort across all phases: **Phase I (Baseline):** Comparing PCT principles against a transactional control. **Phase II (Scaffolding Evaluation):** Testing conflicting feedback strategies (Rehearsal vs. Mock). **Phase III (Validation):** Evaluating the final Agency-Driven system.

**Experimental Setup.** Sessions were conducted in a soundproof usability lab. Participants sat face-to-face with the robot (1m distance) to replicate professional interview dynamics. Researchers monitored conversation flow from an adjacent room via a one-way mirror to ensure system stability without disrupting the participant’s psychological “safe space.”

### 2.3 Participants

We recruited eight university students ( $N = 8$ ; 6 female, 2 male) aged 18–27 ( $M = 21.25$ ,  $SD = 3.24$ ) with self-reported moderate-to-high interview anxiety. Participants came from diverse backgrounds (e.g., Psychology, Nursing) and were active job seekers.

## 2.4 Data Collection and Analysis

We employed a mixed-methods approach targeting social perception, relational quality, and intervention efficacy.

**Quantitative Measures.** We utilized three instruments: (1) *Robotic Social Attributes Scale (RoSAS-SF)*: A concise measure capturing Warmth and Competence. (2) *Barrett–Lennard Relationship Inventory (B–L RI :mini)*: Operationalizes core PCT conditions to quantify the “therapeutic alliance”, specifically empathic understanding, unconditional positive regard, and congruence. (3) *Measure of Anxiety in Selection Interviews (MASI)*: Measured pre- and post-interaction levels using three subscales: Communication, Social, and Behavioral Anxiety.

**Qualitative Measures.** Semi-structured interviews followed each session to probe specific pain points (e.g., feedback timing, cognitive load). Transcripts were analyzed using reflexive thematic analysis.

## 3 Iterative Design and Findings

Our iterative process revealed a progression from establishing safety to navigating cognitive trade-offs, culminating in an agency-driven solution.

### 3.1 Phase I: The Safety–Guidance Gap

To address RQ1, we developed two contrasting system architectures using GPT-4 on the Misty II robot.

- **Baseline (Transactional):** Simulating standard tools, the robot functioned as a neutral data collector, using static behaviors and brief acknowledgments (e.g., “Understood”) to move through questions without emotional engagement.
- **PCT Condition (Empathetic):** Designed to replicate a counselor, the prompt enforced Rogerian principles (unconditional positive regard). Crucially, it included actuation instructions, transforming the LLM into a multimodal engine that matched emotional tone with gestures (e.g., warm nods, welcoming arm movements).

**3.1.1 Findings. Quantitative Validation.** Results confirmed that PCT successfully established psychological safety. On the RoSAS-SF scale, the PCT condition achieved a Warmth score ( $M = 6.06$ ,  $SD = 0.73$ ) significantly higher than the Baseline ( $M = 2.06$ ,  $SD = 0.98$ ;  $t(7) = -9.24$ ,  $p < .001$ ). Similarly, participants reported a significantly higher therapeutic alliance on the B-L RI:mini ( $d = 2.73$ ,  $p < .001$ ). Competence ratings showed no significant difference ( $p = .34$ ), suggesting warmth alone does not drive perceived capability.

**Qualitative: The Safety–Guidance Gap.** Despite the high safety, feedback revealed two critical design flaws:

- **Need for Constructive Guidance:** Participants found the exclusive focus on positivity pedagogically limiting.
- **Tension of Intervention Timing:** Participants were polarized regarding feedback timing. One group prioritized immediate analysis to capitalize on fresh memory, while others preferred deferring feedback to the end to preserve conversational flow.

This divergence highlighted that a static timing strategy could not meet conflicting cognitive needs, motivating the exploration of scaffolding strategies in Phase II.

### 3.2 Phase II: The Scaffolding Paradox

Addressing the need for guidance (RQ2), Phase II tested two conflicting scaffolding strategies with the same cohort:

- **Rehearsal Mode (Strong Scaffolding):** Prioritized immediate instruction. The robot provided detailed professional feedback after *every* response, creating a high-density “perform–listen–learn” structure.
- **Mock Mode (Weak Scaffolding):** Prioritized realism. The robot withheld feedback until a summative report at the end, preserving conversational flow.

3.2.1 *Findings: Critical Design Tensions.* The evaluation revealed a **Scaffolding Paradox**, where neither static extreme met user needs:

- **Cognitive Overload (Rehearsal):** While the feedback was useful, the rigid pacing hindered processing.
- **Dilemma of Authenticity:** The mandatory PCT reflection, when repeated mechanically alongside critique, damaged trust.
- **The Specificity Gap and Need for Agency:** Conversely, Mock Mode felt too generic. Consequently, users rejected the binary choice, explicitly requesting the ability to regulate the flow.

These findings indicated that a static system—whether “Strong” or “Weak”—was insufficient. Users required **Agency** (to control pacing), **Modulated Empathy** (to ensure genuineness), and **Operational Transparency** (to manage expectations).

## 4 Phase III: The Interaction Layer: Agency-Driven Mode

### 4.1 The Final Version: The Agency-Driven Interaction Layer

Addressing the Phase 2 trade-off between cognitive overload (strong scaffolding) and the specificity gap (weak scaffolding), we developed the **Agency-Driven Mode**. The system prompt was refined to operationalize control, authenticity, and transparency:

- **Conditional Feedback Loop (Opt-in Valve):** To balance guidance with flow, the robot explicitly asked, “Would you like professional feedback?” after each response. “Yes” triggered a brief PCT reflection and concise critique; “No” skipped the evaluation to maintain momentum.
- **Streamlined Empathy:** We modulated PCT intensity by capping emotional reflections at one sentence. This kept the robot’s empathy organic and professional rather than performative.
- **Operational Transparency:** The introduction established a clear mental model (e.g., “8–10 questions, ~20–30 minutes”) to mitigate the “black box” anxiety identified in Phase 2.

### 4.2 Phase 3 Testing and Analysis

4.2.1 *Findings: Agency as an Anxiety Buffer.* The Agency-Driven Mode achieved high usability validation. Qualitative analysis revealed that control over the interaction flow acted as a psychological buffer against anxiety. Notably, participants reported that the **ability to decline feedback** was as valuable as the feedback

itself, successfully reframing the interaction from evaluation to collaboration.

4.2.2 *Emerging Design Tension: The Need for Deep Personalization.* While the Agency-Driven Mode resolved pacing issues, extended usage revealed a hierarchy of needs requiring deeper personalization across three levels:

- **Emotional:** Users desired adaptable robot personalities (e.g., distinct interviewer profiles) rather than a fixed style.
- **Content:** Participants requested proficiency-adaptive content, distinguishing between general behavioral questions for novices and domain-specific inquiries for experts.
- **Mode:** Users envisioned structural flexibility beyond manual opt-ins, such as uninterrupted simulation modes for advanced practice versus supported modes for learning.

**Conclusion for Part III.** The Agency-Driven Mode resolved the *Scaffolding Paradox* by restoring user control. However, the demand for nuanced customization suggests that manual agency is merely a transitional step. These findings point toward an **Adaptive Scaffolding Ecosystem** that dynamically tailors personality and difficulty to the user’s evolving state—a direction elaborated in the General Discussion.

## 5 Evaluation of the Final System

We evaluated the success of the final Agency-Driven Mode through comparative quantitative analysis and holistic user satisfaction ratings.

### 5.1 Quantitative Validation: Efficacy & Safety (Phase I vs. III)

To validate the Agency-Driven ecosystem, we analyzed its impact on anxiety reduction (Efficacy) and compared social attributes against the purely empathetic baseline (Safety).

5.1.1 *Intervention Efficacy (MASI Results).* A series of paired-samples *t*-tests on pre- vs. post-interaction MASI scores revealed that the system successfully reduced acute anxiety.

- **Social Anxiety:** Showed the strongest reduction, dropping significantly from  $M = 3.35$  ( $SD = 0.34$ ) to  $M = 2.65$  ( $SD = 0.45$ ;  $t(7) = 6.30$ ,  $p < .001$ ). This reduction suggests an association between the Agency-Driven interaction and decreased fear of negative evaluation.
- **Communication Anxiety:** Significantly decreased from  $M = 3.02$  ( $SD = 0.72$ ) to  $M = 2.56$  ( $SD = 0.61$ ;  $t(7) = 2.43$ ,  $p = .045$ ), suggesting potential improvement in verbal confidence, though this result was not statistically significant after Bonferroni correction ( $\alpha_{corrected} \approx .017$ ).
- **Behavioral Anxiety:** Did not change significantly ( $M = 3.04$ ,  $SD = 1.01$  vs.  $M = 2.94$ ,  $SD = 0.68$ ;  $t(7) = 0.42$ ,  $p = .685$ ). This suggests that while the robot effectively alleviates cognitive and social distress (e.g., fear of judgment), altering deep-seated physiological responses (e.g., trembling, heart rate) may require longer-term exposure beyond three sessions.

5.1.2 *Robot Perception (RoSAS-SF)*. To assess social perception, we analyzed the three core dimensions of the RoSAS-SF scale using Kruskal-Wallis tests:

- **Warmth:** A significant difference was found across conditions (Statistic = 18.02,  $p < .001$ ). All PCT conditions (S1–S3) significantly outperformed the control ( $p < .001$ ), proving that the Agency-Driven Mode successfully maintained the robot’s emotional appeal.
- **Discomfort:** The Agency-Driven Mode (S3) elicited significantly lower discomfort than the control ( $p = .031$ ). This indicates that providing user agency effectively mitigated the “uncanny” or evaluative stress often associated with robotic feedback.
- **Competence:** No significant differences were found across conditions ( $p = .755$ ). However, S3 achieved the highest numerical mean ( $M = 6.25$ ), suggesting that the addition of scaffolding began to trend toward higher perceived task-expertise.

5.1.3 *Perceived Therapeutic Alliance (B–L RI:mini)*. The core relational bond remained robust across all phases (Statistic = 12.33,  $p = .006$ ). Post-hoc Dunnett’s tests confirmed that all PCT-based interactions achieved significantly higher alliance scores than the control: S1 ( $p < .001$ ), S2 ( $p < .001$ ), and S3 ( $p < .001$ ).

## 5.2 System Effectiveness and Satisfaction

Participants rated the final system’s holistic impact on ten dimensions (0–100 scale). The results highlight three key strengths:

- **High Utility:** Participants rated “Information usefulness” ( $M = 91.25$ ,  $SD = 11.26$ ) and “Support for stress management” ( $M = 92.50$ ,  $SD = 11.65$ ) as the highest attributes. Consequently, retention intent was high (“Will continue practicing,”  $M = 92.5$ ).
- **Confidence Building:** Despite the high-stakes context, users reported substantial “Increased confidence” ( $M = 83.75$ ,  $SD = 19.23$ ), verifying that agency-driven feedback supported self-efficacy.
- **Acceptable Artificiality:** Lower scores for “Natural behavior” ( $M = 65.00$ ) reflect the inherent artificiality of robotic role-play. Importantly, participants nonetheless felt “Ready for a human interview” ( $M = 76.25$ ), validating the system’s role as a preparatory tool rather than a human replacement.

## 6 Discussion

Our findings offer empirical evidence for bridging the gap between emotional support and instructional guidance in HRI. We discuss two key implications: the role of agency as a scaffolding strategy and the theoretical shift toward adaptive ecosystems.

### 6.1 Agency as a Scaffolding Strategy

Current educational robots often treat “user control” merely as a UI feature. However, our Phase III results suggest that in high-stakes, anxiety-inducing contexts, **agency itself is a scaffolding strategy**. By allowing users to toggle between “Learning Mode” (Feedback On) and “Flow Mode” (Feedback Off), the Opt-in mechanism enabled users to self-regulate their exposure to critique. This aligns with

Self-Determination Theory: preserving autonomy buffers the social threat usually associated with evaluation [12]. Unlike Phase II’s rigid pacing, which forced users into a Zone of Proximal Development (ZPD) they were not emotionally ready for, the Agency-Driven layer allowed users to enter the ZPD on their own terms.

### 6.2 Toward an Adaptive Scaffolding Ecosystem

While the *Manual Agency* (Opt-in) proved effective, qualitative feedback (e.g., decision fatigue) points to the necessity of *Intelligent Adaptation*. Participants eventually requested a system “smart enough to know when to step in” (P8).

We propose a theoretical shift from manual choices to an **Adaptive Scaffolding Ecosystem**.

- **Theoretical Basis:** As Azevedo et al. [1] demonstrated, adaptive scaffolding—which dynamically adjusts to the learning emergent, understanding—is significantly more effective in facilitating self-regulated learning than fixed scaffolding (like our Phase II). Furthermore, Woolf et al. [16] emphasize that intelligent tutors must “redress the cognitive versus affective imbalance.”
- **Future Architecture:** Based on this, future robot coaches should not rely on a static persona. Instead, they should utilize **implicit sensing** (e.g., analyzing speech latency, prosody, or hesitation) to infer anxiety states in real-time. Based on this data, the system should autonomously slide along two axes—*Affective Tone* and *Intervention Intensity*—shifting between a “Nurturer” role to reduce anxiety and an “Instructor” role to build skills, without requiring explicit user commands.

### 6.3 Limitations

Our study is limited by a small sample size ( $N = 8$ ) and a short duration, which may introduce novelty effects. Additionally, our evaluation relied on self-reported metrics (RoSAS-SF, B–L RI:mini); future work should incorporate objective performance benchmarks (e.g., mock interview scores graded by HR experts) and physiological data (e.g., heart rate variability) to verify skill acquisition and stress reduction objectively.

## 7 Conclusion

This study addressed the fundamental tension between psychological safety and instructional guidance in robotic coaching. Through a three-phase iterative design, we identified a **Scaffolding Paradox**: rigid guidance overloads the user, while pure empathy leaves them unskilled.

We contribute an **Agency-Driven Interaction Layer** that resolves this paradox. By integrating a conditional feedback loop with streamlined empathy, our system achieved high user satisfaction (91.3/100 utility) without sacrificing the therapeutic bond established by PCT ( $p > .05$  for Warmth retention). Ultimately, this work suggests that the next generation of AI coaches must evolve to **adaptive ecosystems** that attune to the learner’s shifting cognitive and emotional needs, bridging the divide between the comfort of a counselor and the rigor of a coach.

## References

- [1] Roger Azevedo, Jennifer G Cromley, and Diane Seibert. 2004. Does adaptive scaffolding facilitate students' ability to regulate their learning with hypermedia? *Contemporary educational psychology* 29, 3 (2004), 344–370.
- [2] Daniel David, Silviu-Andrei Matu, and Oana Alexandra David. 2014. Robot-based psychotherapy: Concepts development, state of the art, and new directions. *International Journal of Cognitive Therapy* 7, 2 (2014), 192–210.
- [3] Wenxiu Geng, Chao Zhou, and Yulong Bian. 2024. Change gently: an agent-based virtual interview training for college students with great shyness. *Virtual Reality* 29, 1 (2024), 12.
- [4] Balasaheb Jadhav, Avadhut Sawant, Arnav Shah, Pranamya Vemula, Abhijeet Waikar, and Srushti Yadav. 2024. A comprehensive study and implementation of the mock interview simulator with ai and pose-based interaction. In *2024 1st International Conference on Cognitive, Green and Ubiquitous Computing (IC-CGU)*. IEEE, 01–05.
- [5] Katarzyna Kłęczek, Andra Rice, and Maryam Alimardani. 2024. Robots as mental health coaches: a study of emotional responses to technology-assisted stress management tasks using physiological signals. *Sensors* 24, 13 (2024), 4032.
- [6] Jamy Li. 2015. The benefit of being physically present: A survey of experimental works comparing copresent robots, telepresent robots and virtual agents. *International Journal of Human-Computer Studies* 77 (2015), 23–37.
- [7] Tatsuya Nomura, Takayuki Kanda, Tomohiro Suzuki, and Sachie Yamada. 2020. Do people with social anxiety feel anxious about interacting with a robot? *Ai & Society* 35, 2 (2020), 381–390.
- [8] Reinhard Pekrun. 2006. The control-value theory of achievement emotions: Assumptions, corollaries, and implications for educational research and practice. *Educational psychology review* 18, 4 (2006), 315–341.
- [9] Deborah M Powell, David J Stanley, and Kayla N Brown. 2018. Meta-analysis of the relation between interview anxiety and interview performance. *Canadian Journal of Behavioural Science/Revue canadienne des sciences du comportement* 50, 4 (2018), 195.
- [10] Samira Rasouli, Moojan Ghafurian, and Kerstin Dautenhahn. 2025. Co-Design and User Evaluation of a Robotic Mental Well-Being Coach to Support University Students' Public Speaking Anxiety. *ACM Transactions on Computer-Human Interaction* 32, 3 (2025), 1–70.
- [11] Carl R Rogers. 1957. The necessary and sufficient conditions of therapeutic personality change. *Journal of consulting psychology* 21, 2 (1957), 95.
- [12] Richard M Ryan and Edward L Deci. 2000. Self-determination theory and the facilitation of intrinsic motivation, social development, and well-being. *American psychologist* 55, 1 (2000), 68.
- [13] Matthew J Smith, Emily J Ginger, Michael Wright, Katherine Wright, Laura Boteler Humm, Dale Olsen, Morris D Bell, and Michael F Fleming. 2014. Virtual reality job interview training for individuals with psychiatric disabilities. *The Journal of nervous and mental disease* 202, 9 (2014), 659–667.
- [14] Lev S Vygotsky. 1978. *Mind in society: The development of higher psychological processes*. Vol. 86. Harvard university press.
- [15] David Wood, Jerome S Bruner, and Gail Ross. 1976. The role of tutoring in problem solving. *Journal of child psychology and psychiatry* 17, 2 (1976), 89–100.
- [16] Beverly Woolf, Winslow Burleson, Ivon Arroyo, Toby Dragon, David Cooper, and Rosalind Picard. 2009. Affect-aware tutors: recognising and responding to student affect. *International Journal of Learning Technology* 4, 3-4 (2009), 129–164.

Received 2025-12-08; accepted 2026-01-12