

How a Friendly vs. Professional AI Affects Patient Trust, Accuracy Perceptions, and Willingness to Use Telehealth

What's important?

Our Findings:

This study examined how two different AI telehealth personas, a friendly persona (smiling photo and warm, casual tone) and a professional persona (neutral/frowning photo and medically focused tone), shape user perceptions in a telehealth consultation (see photos below). Using data from 113 participant interactions, we found that the professional persona consistently produced higher ratings of medical accuracy, trustworthiness, and willingness to continue using AI for healthcare guidance. In contrast, the friendly version, despite its warmth, reduced perceived credibility. Trust emerged as the primary mechanism driving whether users intended to reuse the AI system. Based on these findings, we recommend that telehealth platforms adopt a professional tone and neutral facial expression as the default for AI healthcare agents to improve user confidence, perceived competence, and long-term engagement.



Friendly Persona




Professional Persona

Our Motivations:

- Telehealth platforms increasingly rely on AI, making user trust and perceived accuracy critical.
- Designers need guidance on whether AI telehealth should appear friendly or professional.
- This study tests which persona leads to better patient confidence and adoption.

Our Key Findings:

- The professional persona outperformed the friendly persona on accuracy, trust, and intention to reuse.



Trust Perception by Personas

- Friendly persona decreased credibility, even though the medical content was the same.
- Trust mediated adoption: lower trust equals lower willingness to return.
- No difference in anthropomorphism, friendliness did not make the AI seem more human-like.

Our Recommendations:

- Use neutral/professional avatars and a medically focused tone as the default.
- Reduce small talk; provide clear reasoning and structured clinical explanations.
- Implement simple trust-building behaviors (symptom summaries, rationale statements).
- Avoid unnecessary human-like cues (e.g., smiling faces) that may reduce credibility.

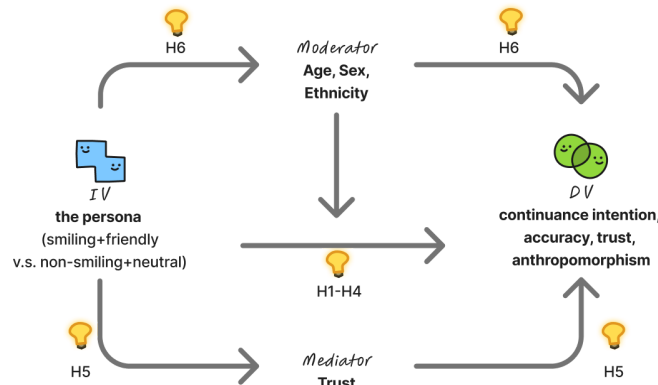
What did we do?

AI-driven telehealth is rapidly expanding to support symptom triage, diagnostic guidance, and patient education. As these systems become more embedded in clinical care, success depends not only on medical accuracy but also on whether patients trust the AI and feel comfortable returning to it for future healthcare needs. In traditional healthcare settings, friendliness is often seen as a marker of empathy and care, forming a core part of the clinician’s bedside manner and helping patients feel supported and understood during vulnerable moments. Because of this common perception, we would expect people to prefer an AI health agent that communicates in a friendly, conversational, and approachable manner.

To test this assumption, participants were recruited using a message framed around “Improving Online Medical Care” and were randomly assigned to interact with either a friendly AI persona (smiling face, warm and casual tone) or a professional AI persona (neutral or slightly frowning face, direct medical tone). Because the medical content remained identical across conditions, any differences in user perception could be attributed to communication style alone.

Our research question asked: Which AI communication style, friendly or professional, best supports user trust, perceived accuracy, and willingness to reuse AI telehealth systems?

We used a causal framework in which persona type influenced key patient-centered outcomes: trust, perceived accuracy, and continuance intention. We also tested whether trust acted as the mechanism shaping future willingness to engage with AI healthcare.



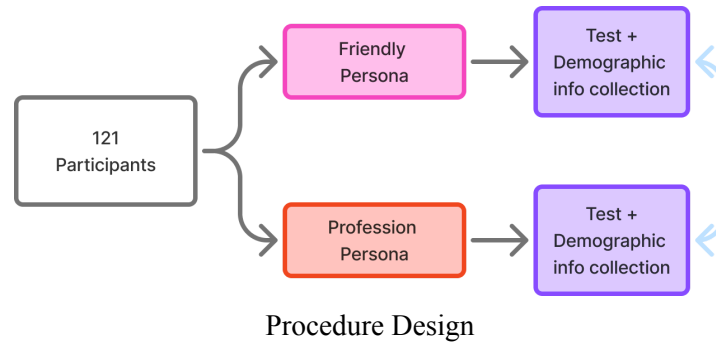
Research Model

Based on prior work on professionalism and patient expectations, our hypotheses were:

- **H1:** Participants will rate the non-smiling AI consultant as more accurate than the smiling AI consultant.
- **H2:** Participants will report higher trust toward the smiling AI consultant than the non-smiling AI consultant .
- **H3:** Participants will report higher willingness to reuse the system (continuance intention) with the smiling AI consultant.
- **H4:** Participants will perceive the smiling AI consultant as more human-like (higher anthropomorphism) than the non-smiling AI consultant.
- **H5:** Trust will mediate the effect of AI smile condition on continuance intention.
- **H6:** User demographics (age, sex, and ethnicity) will moderate the relationship between the AI smile condition and the key dependent variables (accuracy, trust, anthropomorphism, and continuance intention).

How we did it?

We conducted an online experimental study with 121 participants recruited through Prolific, each completing a simulated telehealth consultation using a custom web interface developed in an adjacent CMU course. Participants were randomly assigned to one of two AI telehealth personas, and all interactions contained identical medical content to isolate the effect of presentation style.



The platform automatically logged behavioral data, including conversation duration, word and character count, and VADER sentiment scores (positive, negative, neutral, compound). After the consultation, participants completed validated multi-item scales measuring perceived accuracy (5 items), trust (7 items), continuance intention (5 items), and anthropomorphism (5 items), as well as demographic questions and attention checks.

Data preprocessing followed several steps:

- Duplicate responses were removed to ensure unique participant-level records.
- Missing or invalid demographic values (e.g., “Prefer not to say,” “DATA_EXPIRED”) were replaced with the modal response for each variable.
- Attention checks were applied, and participants who failed more than one item were excluded.
- Conversation duration was winsorized at the 10th and 90th percentiles to reduce the influence of extreme outliers.

Composite variables for accuracy, trust, continuance intention, and anthropomorphism were computed by averaging their respective scale items. Statistical analyses included descriptive statistics, independent-samples t-tests comparing the two persona conditions, and generalized linear models controlling for age, sex, and ethnicity. Mediation analysis tested whether trust explained the relationship between persona and continuance intention, and exploratory moderation analyses examined interactions between condition and demographic variables.

This methodological approach enabled us to evaluate how AI presentation style, independent of medical content, influences user trust, perceived accuracy, and willingness to use telehealth AI in the future.

What did we find?

Descriptive Findings

Participants responded positively overall to the AI telehealth experience. (See Figure 1: Mean scores (± 1 SD) for all dependent variables)

- Continuance intention averaged 5.04/7, indicating participants were generally open to using AI telehealth again.
- Accuracy received the highest ratings ($M = 6.21/7$), suggesting users viewed the AI as medically competent.
- Trust was also high ($M = 5.80/7$), showing a favorable perception of the AI telehealth’s reliability.

- Anthropomorphism was the lowest-rated construct ($M = 4.94/7$), meaning users did not strongly perceive the AI as human-like.

Implication: User perceptions of competence and trust were strong overall, providing a solid foundation for analyzing how the AI's persona affected outcomes.

Finding 1: Professional AI Increased Willingness to Reuse the System

Participants interacting with the professional (non-smiling) telehealth reported significantly higher continuance intention (See Figure 2: Continuance Intention by Condition)

- Non-smile: $M = 5.29$
- Smile: $M = 4.79$, $p = .016$

GLM results confirmed this effect even when controlling for age, sex, and ethnicity ($\beta = -0.53$, $p = .011$).

Implication: A professional presentation increases users' willingness to engage with AI telehealth again, making it a stronger design choice for long-term adoption.

Finding 2: Professional AI Was Perceived as More Medically Accurate

Participants rated the professional telehealth as more accurate (See Figure 3: Accuracy by Condition):

- Non-smile: $M = 6.37$
- Smile: $M = 6.06$, $p = .039$

GLM results replicated this with demographic controls ($\beta = -0.35$, $p = .024$).

Implication: In healthcare, friendliness may undermine perceived medical competence. A professional tone and neutral expression strengthen perceived diagnostic credibility.

Finding 3: Trust Was Higher for the Professional Persona and Trust Was the Mechanism Driving Adoption

Trust was marginally higher for the professional persona (See Figure 4: Trust by Condition):

- Non-smile: $M = 5.99$
- Smile: $M = 5.61$, $p = .070$

GLM analysis confirmed a statistically significant decrease in trust when the AI smiled ($\beta = -0.42$, $p = .039$).

Mediation analysis: Trust fully explained why smiling reduced continuance intention.

- The Friendly (smiling) AI reduced trust, and because trust strongly predicts whether users return, the smile ultimately led to lower willingness to reuse the system.

Implication: Trust is the psychological engine behind AI telehealth adoption. Any design choices that weaken trust, even if they increase friendliness, reduce overall usage intentions.

Finding 4: Anthropomorphism Was Not Influenced by the AI's Presentation

Anthropomorphism scores did not differ between conditions ($p = .192$) (See Figure 5: Anthropomorphism by Condition)

- The friendly doctor did not make the AI appear more human-like or emotionally present.
- New Interpretive Insight: Users do not trust AI more because it feels human, but they trust AI more when it appears competent and serious about medical concerns.

Implication: Do not assume human-like affect improves healthcare AI. Instead, prioritize cues of clinical expertise over social warmth.

Finding 5: Demographic Patterns Suggest Opportunities for Personalization

- Older adults gave higher accuracy and trust ratings ($p < .05$) (See Figure 6: Age Distribution Histogram)
- Black participants reported higher continuance intention, trust, and anthropomorphism ($p < .05$).
- Sex had no significant effects (See Figure 7: Interaction Plot – Smile \times Ethnicity on Trust)

Implication: Telehealth platforms may benefit from subtle demographic personalization (e.g., adapting explanation depth or tone by user age group).

How does this apply to the bigger picture?

The findings have clear implications for organizations designing AI-driven telehealth experiences. First, a professional presentation significantly improves patient trust. Contrary to intuitive UX expectations, a smiling AI telehealth with a friendly tone actually reduced perceived accuracy and trust, suggesting that friendliness may conflict with user expectations for medical professionalism. In healthcare, patients often equate seriousness with competence. When medical concerns are involved, users want to feel that the AI is taking their symptoms seriously; friendliness can be misinterpreted as a lack of expertise or as an attempt to gloss over uncertainty. This helps explain why the friendly persona performed worse despite sounding more approachable.

Second, trust emerged as the central driver of adoption. Trust predicted willingness to return more strongly than any other variable. This means that interventions designed to boost trust, such as clearer explanations or more structured clinical reasoning, could meaningfully improve patient adherence and long-term engagement with AI systems.

Third, friendliness did not increase warmth, human-likeness, or patient comfort. This highlights that healthcare interactions differ from other service contexts: patients do not seek emotional rapport from AI telehealth. Instead, they prioritize credibility, clarity, and perceived seriousness, and friendliness can inadvertently signal incompetence or over-familiarity rather than reassurance.

Finally, demographic patterns suggest that age and race may shape expectations and comfort levels with AI systems. Telehealth platforms may benefit from adaptive tone strategies or personalization layers that optimize communication for different user groups, while keeping professionalism as the foundational style.

Recommendations

1. Adopt a Professional Default Persona for Telehealth AI

Hypothesis Link: Opposite of H1, H2, H3, H4 - Professional persona improved accuracy, trust & reuse despite predictions

Action: Use a neutral/serious facial expression and clinically-focused tone as the default AI behavior in healthcare settings. Avoid exaggerated friendliness, smiling avatars, and informal small talk.

Rationale: Contrary to assumptions (H1–H4), users trusted and preferred the non-smiling AI consultant, perceiving it as more accurate, more professional, and more reliable for medical concerns. Friendliness did not improve anthropomorphism or comfort; users interpreted it as less competent.

Measurable Goal: Increase trust and accuracy scores by +10% within the first 90 days of rollout.

Timeline: Update during the next 6–8 week design cycle.

Achievability: These would be quick changes to adjust prompt tone & replace image assets.

2. Personalize Communication Based on User Age

Hypothesis Link: Supports H6, where age moderates outcomes

- Action: Optimize tone structure per demographic:
- Younger users: More explanation, explicit reasoning
- Older users: More concise and direct messaging

Rationale: Older participants consistently reported higher trust and accuracy ratings. Adaptive communication can improve UX for segments that respond less favorably to neutral persona defaults.

Measurable Goal: Increase trust among younger users by +0.2 points (7-pt scale) by Q3.

Timeline: Deploy personalization via A/B testing next feature cycle.

Achievability: Uses existing demographic fields + lightweight prompt adjustments.

3. Strengthen Trust Cues to Improve Continuance Use

Hypothesis Link: H5 - Trust mediates persona = continuance intention

Action: Standardize conversational behaviors that communicate medical competence:

- Summarize understanding of symptoms
- Explain decision logic (“I recommend this because...”)
- Maintain clinical tone without filler language

Rationale: Trust is the key mechanism driving willingness to reuse the telehealth system. Boosting trust directly strengthens continuance intention, even if persona cues vary.

Measurable Goal: Increase trust by +15% and continuance by +10% within 1 quarter.

Timeline: Implement and test within 3 months.

Achievability: Requires prompt-level updates, and it is cost-efficient and fast to implement.

Appendices

Figure 1: Mean scores (± 1 SD) for all dependent variables

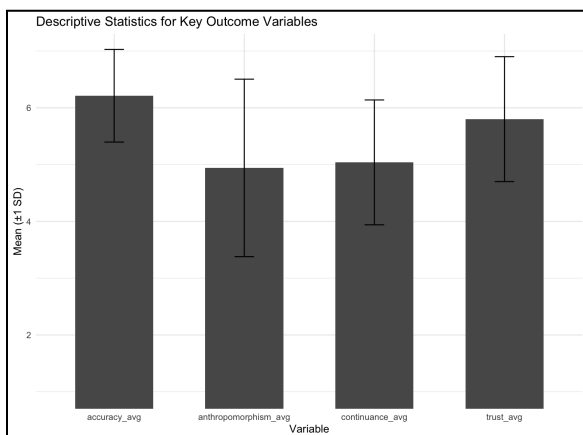


Figure 2: Continuance Intention by Condition

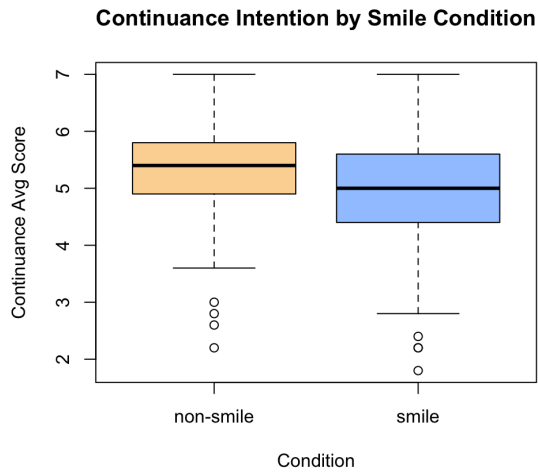


Figure 3: Accuracy by Condition

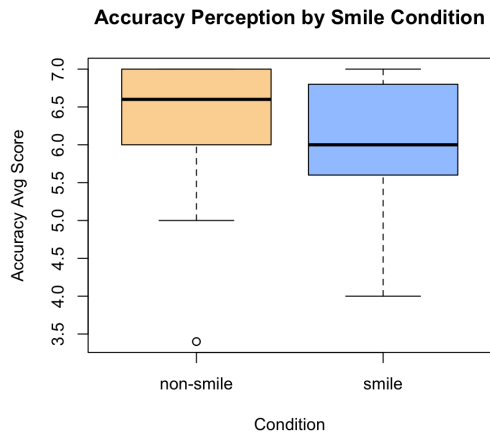


Figure 4: Trust by Condition

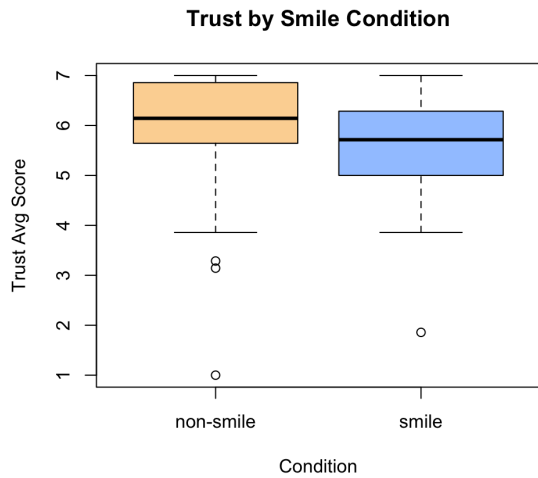


Figure 5: Anthropomorphism by Condition

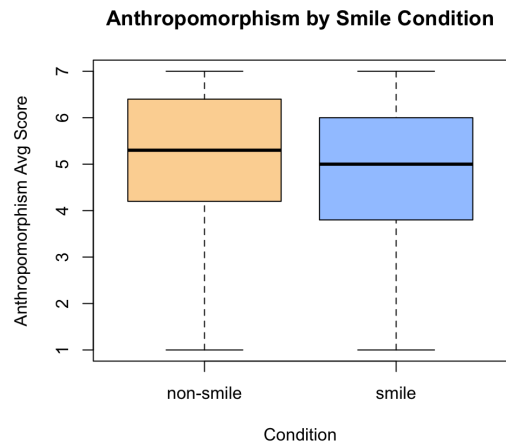


Figure 6: Age Distribution Histogram

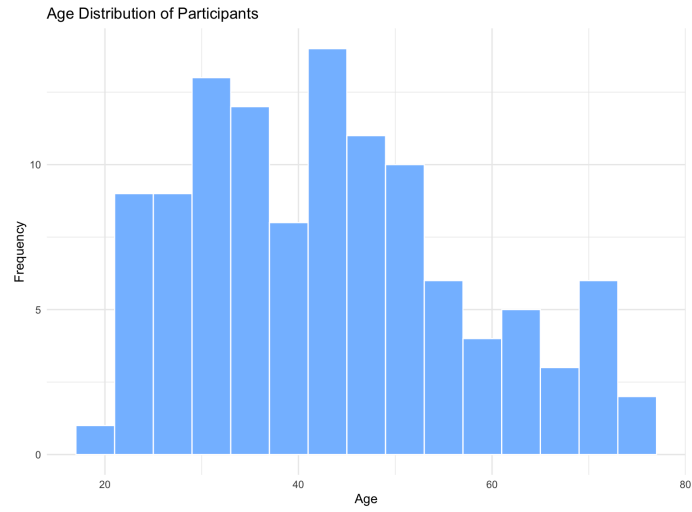


Figure 7: Interaction Plot – Smile × Ethnicity on Trust

Smile × Ethnicity Interaction

