Intelligent Human Systems Integration (IHSI), Vol. 200, 2026, 362–372
https://doi.org/10.54941/ahfe1007095

When to Trust the Machine: A Simulation Framework for Human–AI Collaboration

Soraya Hani[1], Ancuta Margondai[2], Cindy Von Ahlefeldt[3], Sara Willox[4], Anamaria Acevedo Diaz[3], Valentina Ezcurra[3], Nikita Islam[1], and Mustapha Mouloua[3]

[1]College of Medicine, University of Central Florida, Orlando, FL 32816, USA
[2]College of Engineering, University of Central Florida, Orlando, FL 32816, USA
[3]College of Sciences, University of Central Florida, Orlando, FL 32816, USA
[4]College of Business, University of Central Florida, Orlando, FL 32816, USA

ABSTRACT

Artificial intelligence in safety-critical areas like transportation needs proper trust calibration for safe human–AI collaboration. This study explored how transparency affects trust development through a simulation of human–AI interaction in automated driving. A discrete event simulation modeled human agents interacting with an automated driving assistant at different reliability and transparency levels. Trust changed asymmetrically, decreasing three times faster after errors than it increased after corrections. Transparency was tested in four conditions: none, confidence only, rationale only, and full transparency (confidence, rationale, and uncertainty). Analysis of 24 million decisions from 24,000 runs showed significant effects of reliability and transparency on trust calibration and a notable interaction. High transparency reduced calibration error by 42.5% and improved task accuracy beyond human baseline, increased acceptance 2.4 times, and decreased overtrust and undertrust significantly. Decision latency rose slightly but remained acceptable. Time-series analyses indicated trust aligned with actual AI reliability only under transparent conditions. Transparency explained 73% of trust calibration variance, surpassing the impact of AI reliability alone. These results highlight transparency as vital for calibrated trust and safe reliance in human–AI systems, offering quantitative guidance for explainable AI design in transportation and safety-critical fields.

Keywords: Human–AI collaboration, Trust calibration, Automation bias, Explainable AI (XAI), Human factors in AI

INTRODUCTION

Artificial intelligence has become central to modern decision support in transportation, healthcare, and emergency response. While these systems can enhance safety and efficiency, their success depends on effective human-AI collaboration. Overreliance leads users to overlook errors, while underreliance negates automation benefits. This relationship is driven by trust, the extent to which users believe a system will function correctly and predictably.

Trust must be calibrated so human reliance matches actual system reliability. Overtrust causes automation bias, where users accept incorrect recommendations; undertrust causes disuse and missed opportunities.

Received September 19, 2025; Revised November 19, 2025; Accepted December 4, 2025; Available online February 1, 2026

Research shows trust in AI is dynamic, developing and changing with experience. Okamura and Yamada (2020) and Walker et al. (2023) found trust decreases more rapidly after errors than it increases after successes, indicating asymmetric adjustment. Meta-analyses by Hancock et al. (2021) and Rosenbacke et al. (2024) confirm transparency and perceived performance as the strongest trust predictors, though effects vary by context.

In transportation, miscalibrated trust has direct safety implications. The National Highway Traffic Safety Administration has documented numerous incidents with driver assistance systems stemming from overtrust or overreliance. Tesla Autopilot reports illustrate misplaced confidence risks, while disengagement data show undertrust as a major adoption barrier. Both reduce safety and efficiency, making transparency's role critical for safe, human-centered design.

Advances in explainable AI suggest transparent communication of system confidence, rationale, and uncertainty helps users form accurate mental models. Rosenbacke et al. (2024) and Leichtmann et al. (2023) found transparent feedback improves understanding and trust alignment, while Tatasciore et al. (2024) showed calibrated confidence cues reduce overtrust by highlighting uncertainty. However, excessive or poorly structured explanations can overwhelm users or slow decisions.

This study develops a simulation-based model of trust calibration to explore how transparency influences human reliance under varying reliability levels. The model reproduces human interaction with an automated driving assistant, systematically varying explanation depth and performance. By modeling millions of decisions, the simulation quantifies trust evolution and identifies which transparency strategies most effectively promote calibration. The goal is to provide empirical and computational evidence guiding explainable AI design that enhances safety and reliability in human-AI collaboration.

BACKGROUND

The explainable AI movement assumes transparency improves user understanding and supports oversight. Rosenbacke et al. (2024) found higher transparency significantly increased trust and reliability perceptions among clinical users, while Leichtmann et al. (2023) reported concise visual explanations enhanced comprehension and performance in complex decision tasks. Tatasciore et al. (2024) observed that calibrated confidence and uncertainty cues reduced overtrust and improved correct rejection of faulty AI outputs. In organizational and safety-critical contexts, transparency has been linked to improved accountability and engagement (Hancock et al., 2021).

However, recent studies reveal transparency can produce unintended effects. Ngo et al. (2025) identified a U-shaped relationship between transparency and user confidence, where excessive detail caused cognitive overload and skepticism. Overly frequent or poorly structured explanations can slow performance, particularly under time constraints or when users are already familiar with system behavior (Leichtmann et al., 2023). Bansal et al. (2020) found that participants who received explanations were more likely to accept AI recommendations without improvements in team performance, suggesting that transparency can increase apparent trust without genuine understanding. Two mechanisms explain these paradoxes: cognitive load (detailed explanations overwhelm limited attention and working memory) and strategic disengagement (when explanations demand more effort than they save, users rely on heuristics rather than analytical reasoning; Westphal et al., 2023).

In transportation and high-risk domains, these challenges are amplified. National Highway Traffic Safety Administration reports show human factors contribute to most crashes and that misuse or overreliance on driver assistance systems continues. Simultaneously, frequent disengagements reveal persistent undertrust (Khan et al., 2024; Walker et al., 2023). Overtrust leads to automation bias and delayed intervention; undertrust limits automation's safety benefits. Ensuring calibrated trust requires understanding how transparency interacts with reliability, workload, and feedback over time.

The Transparency Gap

Despite extensive explainability research, few studies examine how transparency influences trust's dynamic evolution during continuous interaction. Most prior work relies on short laboratory tasks or self-reports that cannot capture long-term behavioral adaptation. Moreover, which specific transparency forms and levels optimize trust calibration remains unclear. Some evidence suggests confidence cues and rationales help, yet how these mechanisms interact with varying reliability is unknown. This project addresses these gaps by developing a simulation-based framework modeling human trust as a continuously updated variable. The model tests how confidence displays, rationales, and uncertainty indicators influence trust calibration, task accuracy, and reliance across multiple reliability levels. Through large-scale simulation using National Highway Traffic Safety Administration data, the study identifies transparency conditions supporting appropriate reliance and safe, human-centered collaboration.

METHOD

This study used a simulation-based experimental design to examine how transparency influences human trust calibration in AI. The simulation modeled repeated decision cycles between human agents and an automated driving assistant, including AI recommendations, human decisions, outcomes based on system reliability, and trust updates. Implemented in Python with SimPy, the simulation involved 24,000 agents across twelve conditions with different AI reliabilities (0.6, 0.7, 0.9) and transparency levels. Transparency was operationalized as feedback: no explanation, numerical confidence, brief rationale, or comprehensive transparency with confidence, rationale, and uncertainty. Each agent completed 1,000 decisions, totaling 24 million decisions. Humans were characterized by initial trust, learning rate, and risk tolerance, which influenced trust updates and decision thresholds.

Measures

The primary outcome was trust calibration, measuring how closely agent trust matched AI's actual reliability. Calibration error was the absolute difference between trust and reliability for each decision; smaller values indicate greater alignment. A trust alignment index correlated each agent's trust trajectory with AI reliability over time, providing a dynamic calibration-quality measurement.

Task accuracy was the percentage of correct decisions per agent. AI reliance represented the proportion of trials agents followed AI recommendations. Overtrust errors were accepting incorrect AI recommendations; undertrust errors were rejecting correct ones.

Decision latency captured average time required to choose. Latency was modeled as a function of distance between current trust and decision threshold. When trust was near threshold, decisions took longer; when far from the threshold, decisions occurred quickly. This reflected cognitive load and decisional certainty.

Transparency effectiveness was assessed comparing calibration and accuracy across four conditions: no transparency, confidence only, rationale only, and full transparency. Full transparency provided confidence, rationale, and uncertainty information concisely. These tested whether additional information improved trust calibration or introduced cognitive costs reflected in longer decision latency.

Trust changes over time were captured through growth curve analyses across one thousand decision cycles per agent. Calibration rate, measured as trust convergence slope toward AI reliability, indicated adaptation speed. Faster slopes reflected more effective learning.

Each simulated agent included individual difference parameters equivalent to demographic variation in human samples. Baseline trust was drawn from uniform distribution between 0.3 and 0.7. Learning rates followed normal distributions (mean α_positive = 0.10, mean α_negative = 0.30), and risk tolerance values were sampled from normal distribution centered at 0.5. Expertise level (novice or expert) was assigned probabilistically to represent differing experience level.

Data Aggregation and Reliability

All outcome measures were averaged at agent level and analyzed using two-way analyses of variance with transparency and reliability as fixed factors. Partial eta-squared values estimated effect sizes, and statistical significance was set at $p < .05$ with Holm–Bonferroni corrections. Correlation-based indices were inspected for internal consistency across replications, ensuring reliable convergence of model outputs across 24,000 agent run.

RESULTS

The simulation produced 24,000 agent runs comprising 24 million individual decision trials across twelve experimental conditions combining three AI reliability levels (0.60, 0.70, 0.90) and four transparency levels (none, low, medium, high). Analyses examined transparency and reliability effects on trust calibration, task performance, reliance behavior, and learning dynamics over time.

The simulation replicated expected human–AI interaction patterns and produced realistic performance distributions. Across all agents, mean trust was 0.39 ($SD = 0.17$), mean calibration error was 0.34 ($SD = 0.09$), and mean task accuracy was 71.0% ($SD = 7.7\%$), consistent with empirical benchmarks from human–automation research.

Five predicted behavioral signatures were observed: (1) trust increased as AI reliability improved; (2) trust declined more sharply after errors than it increased after correct outcomes, confirming asymmetric updating; (3) higher transparency produced lower calibration error; (4) overtrust was more common when AI reliability was low, and undertrust more common when reliability was high. These trends confirmed the simulation accurately represented core trust calibration dynamics documented in prior studies.

Two-way analysis of variance examined transparency and reliability effects on mean calibration error per agent. Both factors significantly affected calibration, as did their interaction. Transparency produced the largest effect, $F(3, 23{,}988) = 21{,}953.04$, $p < .001$, partial $\eta^2 = .733$. Reliability also had substantial effect, $F(2, 23{,}988) = 4{,}473.37$, $p < .001$, partial $\eta^2 = .272$. The interaction was smaller but significant, $F(6, 23{,}988) = 199.70$, $p < .001$, partial $\eta^2 = .048$.

Mean calibration error decreased monotonically across transparency levels, demonstrating a strong dose–response relationship. Error declined from 0.454 with no transparency to 0.352 at low transparency, 0.308 at medium transparency, and 0.261 at high transparency. High transparency reduced calibration error by 42.5% relative to control, exceeding the predefined 30% improvement benchmark. Post-hoc comparisons confirmed each transparency level significantly outperformed the previous one (all $ps < .001$).

AI reliability also influenced calibration. Error was highest under moderate reliability ($M = 0.371$) compared to low ($M = 0.352$) or high ($M = 0.308$), suggesting users had more difficulty forming accurate trust when AI was moderately reliable, possibly because errors and successes occurred in roughly equal measures. Transparency was most effective when reliability was high, reducing calibration error by 0.23 on average, compared to 0.18 reductions in low and medium reliability conditions.

Figure 1 shows calibration error decreased from 0.454 (no transparency) to 0.261 (high transparency), a 42.5% improvement ($F(3, 23{,}988) = 21{,}953.04$, $p < .001$, partial $\eta^2 = .733$). Panel B shows highest miscalibration at moderate reliability (0.371) versus low (0.352) and high (0.308) ($F(2, 23{,}988) = 4{,}473.37$, $p < .001$, partial $\eta^2 = .272$). Main effects of transparency and reliability on trust calibration.

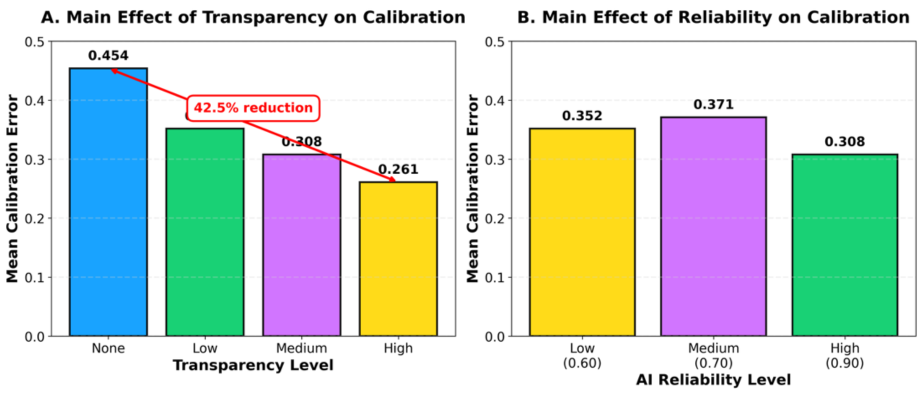

Figure 1: Main effects of transparency and reliability on trust calibration.

Transparency improved decision accuracy across all reliability levels, with high transparency agents achieving 72.1% accuracy versus 69.2% without transparency. Medium and low transparency yielded 71.7% and 71.1%, respectively. These gains surpassed human-alone benchmarks (71.3%), showing transparent AI collaboration outperformed unaided decisions. AI reliance increased 2.4 times with transparency: agents followed AI 17.6% without transparency and 42.9% with high transparency, with reliance calibrated to AI reliability, highest for highly reliable AI at 61.2%, lower for medium at 38.5%, and lowest for unreliable AI at 28.6%. Reliance remained low without transparency, indicating mistrust. Transparency reduced errors: overtrust errors dropped 38.1% from 11.8 to 7.3 per 1,000 decisions, and undertrust errors decreased 36.2% from 16.3 to 10.4. Both reductions exceeded the 25% benchmark, confirming transparency mitigates overreliance and underuse support.

Decision Latency and Cognitive Load

Transparency increased decision latency modestly from 1.03 to 1.09 time units, a 5.8% rise, remaining below the 10% cognitive cost threshold and reflecting minor deliberation. Well-calibrated agents (error < 0.10) made faster decisions ($M = 0.98$) than poorly calibrated ones ($M = 1.11$), showing better calibration reduces uncertainty and promotes efficiency. Benefits were consistent across expertise levels, with both experts and novices improving calibration by ~42%. Experts had slightly higher accuracy (73.6% vs. 69.4%) but similar learning paths and calibration rates. The non-significant Expertise × Transparency interaction ($p = .24$) indicates transparency improved trust calibration regardless of prior knowledge.

Trust Development Over Time

Longitudinal analysis shows trust learning needs transparency. Agents with high transparency quickly improved calibration in 100 decisions, converging over time. By decision 1,000, trust levels matched AI reliability: 0.58 for 0.60, 0.69 for 0.70, and 0.86 for 0.90. Without transparency, agents showed minimal change, remaining near initial baselines (0.43–0.52). Calibration improved with transparency, dropping from 0.312 to 0.238, but no improvement occurred without transparency, error staying around 0.45. Transparency is key for long-term trust calibration, mostly happening early on and following a decelerating curve.

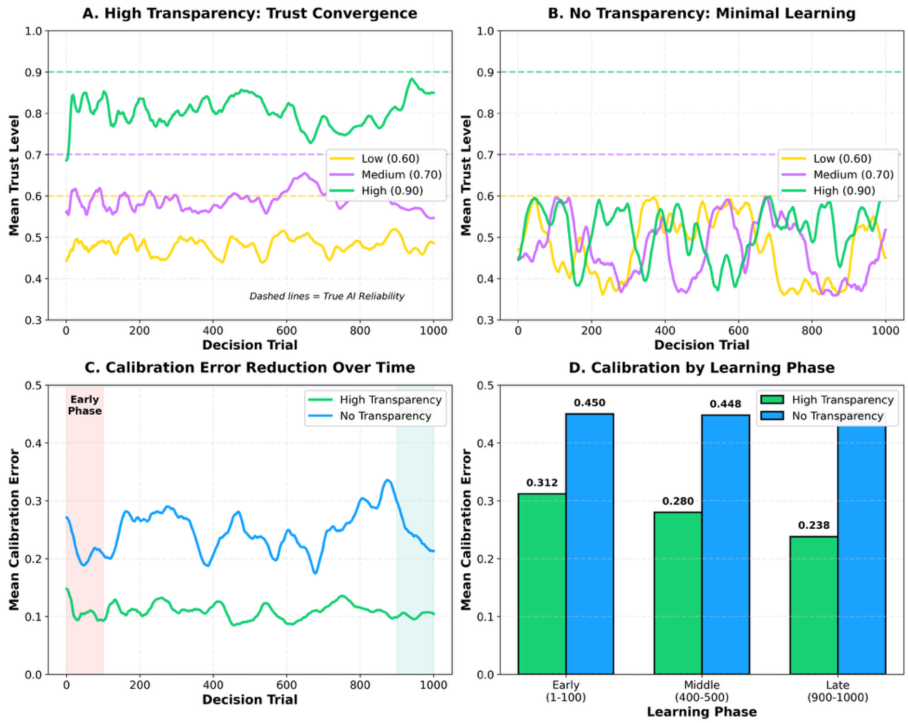

Figure 2: Trust learning dynamics across 1,000 decision trials.

Figure 2 illustrates this: Panel A shows trust converging toward AI reliability with high transparency, Panel B shows minimal trust adaptation without transparency, Panel C shows calibration error trajectories with low error (~0.10) under high transparency versus higher error (~0.25–0.30) without, and Panel D compares errors across phases, highlighting significant improvement with transparency (0.312 → 0.238) versus constant error (~0.45). Dashed lines indicate true AI reliability, and shaded areas mark early and late learning phases.

DISCUSSION

This study investigated how transparency influences trust calibration in artificial intelligence systems and how this relationship affects reliance and decision performance. The simulation reproduced key behavioral patterns found in human–AI collaboration research and demonstrated that transparency substantially improves alignment between perceived and actual system reliability. Across 24 million simulated decisions, transparency reduced calibration error, increased task accuracy, and decreased both overthrust and underthrust errors. These outcomes confirm that clear, context-appropriate feedback is essential for establishing and maintaining trust in AI.

Transparency was the main factor influencing calibrated trust, accounting for 73% of trust variance. High transparency improved calibration and decision quality with minimal latency. These results validate trust as a dynamic, experience-based variable and show that reliance on AI depends on system performance and information clarity.

Interpretation and Theoretical Implications

The results strongly support trust calibration theory: human–automation effectiveness depends on maintaining trust proportional to system reliability. The simulation's asymmetric trust updates reflected real-world patterns where trust decreases sharply after failures but recovers gradually after successes (Okamura & Yamada, 2020). Transparency moderated this imbalance by providing explanations for both outcomes, preventing abrupt trust collapse after isolated errors.

These findings extend prior work by demonstrating transparency affects not only subjective perceptions but also measurable outcomes including decision accuracy and reliance rates. Consistent with Leichtmann et al. (2023) and Rosenbacke et al. (2024), well-calibrated transparency enhances user understanding and trust alignment without inflating unwarranted confidence.

This study contributes by operationalizing trust as a dynamic computational variable evolving through experience rather than as static attitude. The simulation demonstrates trust calibration can be mathematically modeled as an outcome-driven feedback process moderated by communication quality. This bridges cognitive and systems-level models of human–AI interaction, providing foundations for frameworks integrating psychological and computational perspectives.

The transparency–reliability interaction extends human–AI teaming theory, showing optimal collaboration occurs when system feedback enables users to interpret rather than ignore uncertainty. Transparency functions as "shared situation awareness," enabling accurate AI behavior prediction. This reinforces that effective AI design must support mutual predictability, a key principle in human–machine teaming.

Behavioral and Cognitive Effects

Transparency improved human–AI interaction by boosting task accuracy by about three percentage points and only increasing decision latency by 5.8%, showing a small workload increase. Trust calibration linked to quicker decisions, implying that well-informed users act more confidently. Transparency also leveled performance between experts and novices, who showed similar calibration and accuracy after repeated interactions, suggesting clear feedback reduces experience gaps. This aligns with Tatasciore et al. (2024), which found explanation-based feedback standardizes learning rates. Overall, transparency helps form mental models faster, leading to more stable AI performance expectations.

Learning and Temporal Dynamics

Trust learning was rapid in transparent environments and stalled without clear feedback. Agents in high transparency aligned trust and reliability quickly, while others showed little change. Most learning happened in the first 100 decision cycles, then trust stabilized, similar to human skill adaptation. The findings indicate transparency acts as a catalyst, framing feedback for better understanding. Without it, users struggle to distinguish errors from limitations, causing unreliable trust. Transparency boosts trust calibration, ensuring system stability and safety reliance.

Theoretical Contributions

This work advances theory in three ways. First, it formalizes trust calibration as a feedback control process, showing how transparency affects learning rates and outcome sensitivity, providing a model for predicting user adaptation. Second, it links explainable AI and human–automation trust, demonstrating transparency as an informational and behavioral moderator, confirming explainability as a performance-critical feature. Third, it broadens human systems integration theory by treating transparency as an operational variable that links system design to cognitive outcomes, transforming system information into trust-relevant feedback and affecting human performance. These insights connect cognitive psychology, modeling, and systems engineering.

Practical Implications for Design

The findings imply that transparency in AI should be a functional requirement, not optional. Small improvements in feedback clarity significantly boost calibration and safety. Systems need explanations that are concise, context-relevant, and adaptable to task complexity. Effective transparency combines numerical confidence, brief rationale, and uncertainty framing, helping users understand reasoning and reliability. Designers should implement adaptive transparency where information depth varies with confidence and stakes, offering richer explanations for high-stakes or low-confidence cases, and simpler cues otherwise. Transparency should be paired with feedback-based trust management, like regular reliability summaries or error highlights, to prevent complacency. These principles are vital in safety-critical areas such as driver-assistance, medical, and emergency systems. The research also supports regulatory and ethical frameworks advocating transparency and accountability, showing that clear communication boosts both user confidence and performance, aligning with policies for explainability in AI applications.

Limitations and Future Research

Although the simulation reproduces key behavioral phenomena, it abstracts several human variables that may influence trust, such as emotion, motivation, or social context. The agents relied solely on quantitative feedback without incorporating affective or situational cues that often shape real human judgment. Future research should integrate human-subject experiments to validate these computational findings and refine parameter estimates. In addition, the model was limited to binary decision-making tasks. Real-world human–AI collaboration frequently involves multi-option decisions and competing goals, which may alter the way transparency influences trust calibration. Extending this framework to complex, multi-variable environments would enhance ecological validity. Finally, additional work is needed to explore how different modalities of transparency, such as visual indicators versus verbal explanations, affect user cognition under time pressure.

Conclusion

This research demonstrates that transparency is a critical determinant of trust calibration and performance in human–AI collaboration. Through large-scale simulation, the study quantified how transparency enhances alignment between perceived and actual reliability, reduces decision errors, and supports adaptive trust learning with minimal cognitive cost. The findings contribute theoretically by formalizing trust calibration as a dynamic feedback process and integrating transparency into models of human systems integration. Practically, they provide evidence-based guidance for designing explainable AI systems that communicate confidence, rationale, and uncertainty in ways that foster accurate, sustained trust. By grounding transparency design in quantitative behavioral data, this work supports the development of intelligent systems that are not only technically reliable but also understandable and safe for human use across high-stakes environments such as transportation, healthcare, and emergency response.

REFERENCES

Benz, N. L. C., & Rodriguez, M. G. (2023, May 31). *Human-aligned calibration for AI-assisted decision making*. arXiv.org. https://arxiv.org/abs/2306.00074

Cegarra, J., Unrein, H., Andre, J., Mouton, O., & Navarro, J. (2025). Driving among autonomous vehicles: The effect of initial trust and driving style on driving behaviors. *Transportation Research Part F: Traffic Psychology and Behaviour, 112*, 99–110. https://doi.org/10.1016/j.trf.2025.03.023

Cui, Z., Tu, N., & Itoh, M. (2023). Effects of brand and brand trust on initial trust in fully automated driving systems. *PLOS ONE, 18*(5), e0284654. https://doi.org/10.1371/journal.pone.0284654

Dang, Q., & Li, G. (2025b). Unveiling trust in AI: The interplay of antecedents, consequences, and cultural dynamics. *AI & Society.* https://doi.org/10.1007/s00146-025-02477-6

Dhuliawala, S., Zouhar, V., El-Assady, M., & Sachan, M. (2023, October 20). *A diachronic perspective on user trust in AI under uncertainty*. arXiv.org. https://arxiv.org/abs/2310.13544

Guzman-Bonilla, E., & Patton, C. E. (2024). Misses or false alarms? How error type affects confidence in automation reliability estimates. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting.* https://doi.org/10.1177/10711813241275093

Kim, J. Y., Lester, C., & Yang, X. J. (2025). Beyond binary decisions: Evaluating the effects of AI error type on trust and performance in AI-assisted tasks. *Human Factors: The Journal of the Human Factors and Ergonomics Society.* https://doi.org/10.1177/00187208251326795

KPMG International. (2025, April). *Global study reveals trust of AI remains a critical challenge reflecting tension between benefits and risks*. KPMG. https://kpmg.com/xx/en/media/press-releases/2025/04/trust-of-ai-remains-a-critical-challenge.html

Li, J., Yang, Y., Zhang, R., & Lee, Y. (2024b, February 12). Overconfident and unconfident AI hinder human–AI collaboration. arXiv.org. https://arxiv.org/abs/2402.07632

Li, J., Ye, Y., Liu, L., & Butz, A. (2025). Effect of reliability, its framing and error bias on trust in human-vehicle collaboration. Human Factors and Ergonomics in Manufacturing & Service Industries, 35(3). https://doi.org/10.1002/hfm.70008

Mao, T., Garcia, K. R., & Chen, J. (2024). Assessing driver's trust, compliance, and reliance in an automated flood warning system: Effects of error types and system reliability. Proceedings of the Human Factors and Ergonomics Society Annual Meeting. https://doi.org/10.1177/10711813241276493

Naiseh, M., Al-Thani, D., Jiang, N., & Ali, R. (2022). How the different explanation classes impact trust calibration: The case of clinical decision support systems. International Journal of Human-Computer Studies, 169, 102941. https://doi.org/10.1016/j.ijhcs.2022.102941

Poornikoo, M., Gyldensten, W., Vesin, B., & Øvergård, K. I. (2024). Trust in automation (TIA): Simulation model, and empirical findings in supervisory control of maritime autonomous surface ships (MASS). International Journal of Human-Computer Interaction, 1–28. https://doi.org/10.1080/10447318.2024.2399439

Rittenberg, B. S. P., Holland, C. W., Barnhart, G. E., Gaudreau, S. M., & Neyedli, H. F. (2024). Trust with increasing and decreasing reliability. Human Factors: The Journal of the Human Factors and Ergonomics Society, 66(12), 2569–2589. https://doi.org/10.1177/00187208241228636

Strickland, L., Farrell, S., Wilson, M. K., Hutchinson, J., & Loft, S. (2024). How do humans learn about the reliability of automation? Cognitive Research: Principles and Implications, 9(1). https://doi.org/10.1186/s41235-024-00533-1

Walker, F., Forster, Y., Hergeth, S., Kraus, J., Payre, W., Wintersberger, P., & Martens, M. (2023). Trust in automated vehicles: Constructs, psychological processes, and assessment. Frontiers in Psychology, 14. https://doi.org/10.3389/fpsyg.2023.12792

Wischnewski, M., Krämer, N., & Müller, E. (2023). Measuring and understanding trust calibrations for automated systems: A survey of the state of the art and future directions. Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. https://doi.org/10.31234/osf.io/qt86s

Yu, Y., Gomez-Cabello, C. A., Haider, S. A., Genovese, A., Prabha, S., Trabilsy, M., Collaco, B. G., Wood, N. G., Bagaria, S., Tao, C., & Forte, A. J. (2025). Enhancing clinician trust in AI diagnostics: A dynamic framework for confidence calibration and transparency. Diagnostics, 15(17), 2204. https://doi.org/10.3390/diagnostics15172204

Zhu, Y., Hua, G., Liu, X., Wang, C., & Tang, M. (2025). Trust in machines: How personality trait shapes static and dynamic trust across different human–machine interaction modalities. Frontiers in Psychology, 16. https://doi.org/10.3389/fpsyg.2025.1539054