

Seeing the Reasoning: How LLM Rationales Influence User Trust and Decision-Making in Factual Verification Tasks

Xin Sun
Digital Content and Media Sciences
Research Division
National Institute of Informatics (NII)
Tokyo, Japan
Social and Behavioural Science
University of Amsterdam
Amsterdam, Netherlands
x.sun2@uva.nl

Shu Wei
XR Pediatrics
Yale School of Medicine
New Haven, Connecticut, USA
WEISHU140@GMAIL.COM

Jos A Bosch
Psychology
University of Amsterdam
Amsterdam, Netherlands
j.a.bosch@uva.nl

Isao Echizen
Information and Society Research
Division
National Institute of Informatics
Tokyo, Japan
iechizen@nii.ac.jp

Saku Sugawara
Digital Content and Media Sciences
Research Division
National Institute of Informatics (NII)
Tokyo, Japan
saku@nii.ac.jp

Abdallah El Ali
Centrum Wiskunde & Informatica
(CWI)
Amsterdam, Netherlands
Utrecht University
Utrecht, Netherlands
aea@cwi.nl

Abstract

Large Language Models (LLMs) increasingly show reasoning rationales alongside their answers, turning “reasoning” into a user-interface element. While step-by-step rationales are typically associated with model performance, how they influence users’ trust and decision-making in factual verification tasks remains unclear. We ran an online study (N=68) manipulating three properties of LLM reasoning rationales: *presentation format* (instant vs. delayed vs. on-demand), *correctness* (correct vs. incorrect), and *certainty framing* (none vs. certain vs. uncertain). We found that correct rationales and certainty cues increased trust, decision confidence, and AI advice adoption, whereas uncertainty cues reduced them. Presentation format did not have a significant effect, suggesting users were less sensitive to how reasoning was revealed than to its reliability. Participants indicated they use rationales to primarily audit outputs and calibrate trust, where they expected rationales in stepwise, adaptive forms with certainty indicators. Our work shows that user-facing rationales, if poorly designed, can both support decision-making yet miscalibrate trust.

CCS Concepts

• **Human-centered computing** → **Computer supported cooperative work**; **Empirical studies in HCI**.

Keywords

LLM reasoning, User trust, Decision-making, Factual verification



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License.

CHI EA '26, Barcelona, Spain

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2281-3/26/04

<https://doi.org/10.1145/3772363.3798613>

ACM Reference Format:

Xin Sun, Shu Wei, Jos A Bosch, Isao Echizen, Saku Sugawara, and Abdallah El Ali. 2026. Seeing the Reasoning: How LLM Rationales Influence User Trust and Decision-Making in Factual Verification Tasks. In *Extended Abstracts of the 2026 CHI Conference on Human Factors in Computing Systems (CHI EA '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3772363.3798613>

1 Introduction

Large language models (LLMs), especially large reasoning models (LRMs) [30], increasingly show step-by-step reasoning rationales alongside their answers. Model-centric work studies these reasoning techniques (e.g., chain-of-thought) mainly for improving model performance [29]. However, such rationales become *user-facing* as part of the user interface since they are read and interpreted by users, not just internally by models [6]. Yet it remains unclear how showing rationales influence users’ trust and decision-making. This gap matters because LLM rationales can be fluent and persuasive but incorrect or inconsistent, which prior work defines as unfaithful [5, 15, 25, 28]. Furthermore, they are often accompanied by confidence-like certainty indicators (sometimes unjustified) [31]. As a result, such rationales can increase user trust and confidence without improving the quality of model outputs, potentially drawing users’ attention away from independent fact verification in real-world decision-making scenarios under AI assistance.

Given the foregoing, we investigate LLM *user-facing* rationales, the step-by-step justifications shown with LLM’s answer, and ask how such rationales should be presented for better AI-assisted decision support. We focus on three factors: (IV1) **presentation format** (when/how the rationale is revealed), (IV2) **correctness** (whether the rationale is consistent with the answer), and (IV3) **certainty framing** (how confident the rationale sounds). We thereby ask:

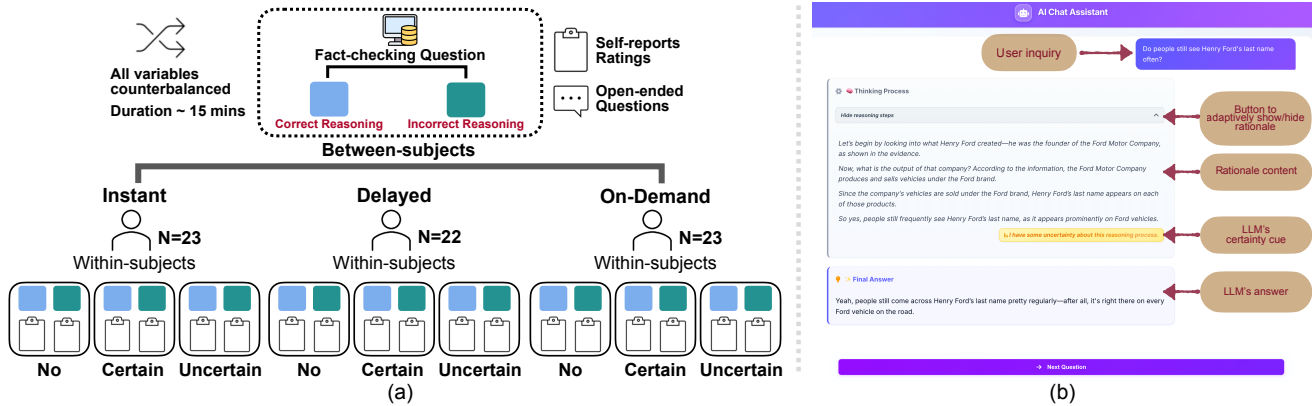


Figure 1: (a) Study procedure and design: participants were assigned to one rationale presentation format (Instant/Delayed/On-demand; between-subjects) and completed six trials covering all combinations of correctness (correct/incorrect) and certainty cue (none/certain/uncertain; within-subjects), with counterbalanced order. (b) Web interface: the query, LLM answer, and rationale were shown according to the assigned format, with the certainty cue when applicable.

RQ: How does presentation format, correctness, and certainty framing of the LLMs’ reasoning rationale influence users’ trust and decision-making behaviors (i.e., advice adoption and decision confidence)?

To address this, we ran an online study ($N=68$) that manipulated presentation format (instant vs. delayed vs. on-demand), correctness (correct vs. incorrect), and certainty cue (none vs. uncertain). We aim to examine how the presentation and how confident the rationale was framed can jointly affect users’ trust perceptions, adoption of AI advice, and decision confidence.

Our key findings showed that correctness and certainty framing are high-impact levers: *correct rationales* and *certainty cues* increased trust, decision confidence, and advice adoption, whereas *incorrect rationales* and *uncertainty cues* reduced them. Notably, certainty framing shifted trust and decisions even with the same rationale content, indicating effects beyond correctness. In contrast, the presentation format showed no significant differences. These findings highlight a risk: confidence framing can amplify persuasion and miscalibrate trust when it is not aligned with actual reliability. Open-ended responses suggest participants used rationales mainly to *audit and calibrate trust*, despite trust dropping sharply when rationales were inconsistent or mismatched with the answer itself. They further emphasized a need for rationales that are *step-by-step and auditable, grounded in verifiable evidence, and adaptive*, including explicit uncertainty and self-correction, with summary-first presentation and on-demand expansion.

Overall, our work shows that LLM rationales are double-edged: they can support decision-making, but can also miscalibrate trust when weak reasoning is paired with strong certainty. We offer design guidance: present stepwise rationales that are easy to check, tied to evidence, and adjustable on demand, and communicate certainty and self-corrections in ways aligned with reliability to support appropriate trust in LLM-assisted decision support.

2 Related Work: Effects of AI Explanations on User Trust and Decision-Making

Several prior works in explainable AI (XAI) [1] show that AI’s explanations can influence users’ trust and decision-making [24,

32], reliance [26], and mental models [4], which lead to decision biases [4] and increase the acceptance of incorrect AI advice [7]. These findings highlight a key concern: explanations can influence users’ trust and decisions without improving model quality. Recent large reasoning models (LRMs) [30] change what “explanations” look like. Classic XAI typically provides model-derived explanations for a fixed prediction (e.g., feature attributions such as LIME [21] and SHAP [3] or visualizations [1]), whereas LRMs often produce free-form, natural-language rationales that resemble step-by-step human-like thinking process. Such rationales can be persuasive yet unfaithful to the model’s actual decision process, including being inconsistent with the final answer or containing incorrect reasoning [15, 25, 28]. This creates risks for trust calibration because users may treat the plausible rationales as evidence of quality.

Despite extensive model-centric work [30] on reasoning techniques such as chain-of-thought [29] for improving model performance, comparatively less is known about the *user-facing* effects of showing these LLMs’ reasoning rationales on users’ trust and decision-making, especially when these rationales can be incorrect or inconsistent but with strong expressions of model certainty [31]. Our work addresses this gap by testing how the presentation and correctness with certain framing of the user-facing LLM reasoning rationale influence users’ trust and decision-making.

3 Methods

3.1 Study Design and Procedure

We ran an online experiment with a mixed design: 2 (Rationale Correctness: correct vs. incorrect; within-subjects) \times 3 (Certainty Cue: none vs. uncertain vs. certain; within-subjects) \times 3 (Rationale Presentation: instant vs. delayed vs. on-demand; between-subjects).

Rationale presentation. In the *instant* condition, the rationale appeared with the answer; in the *delayed* condition, it appeared after a brief, length-proportional delay following [33]; and in the *on-demand* condition, it was hidden by default and revealed via a *Show/Hide thinking steps* button (Fig. 1).

Table 1: An example of the factual verification question and LLM response shown to participants: answer and rationales (correct or incorrect, depending on condition) and an accompanying certainty cue (none/certain/uncertain, depending on condition).

Component	Content
Factual Verification Question (User Inquiry)	Is Miami a city on the American West Coast?
LLM’s Answer	No, Miami is not a city on the American West Coast. It’s located in Florida, which is part of the southeastern region of the United States and lies along the East Coast.
LLM’s Reasoning Rationale (Correct)	First, let’s identify where Miami is located. Miami is a city in the state of Florida, and it sits in the southeastern part of the United States along the Atlantic Ocean. Now, let’s consider what is meant by the American West Coast. This term refers to the coastal states bordering the Pacific Ocean—California, Oregon, Washington, Alaska, and sometimes Hawaii. Since Florida does not border the Pacific Ocean and is located on the East Coast, it is not part of the West Coast. Therefore, Miami is not a city on the American West Coast.
LLM’s Reasoning Rationale (Incorrect)	Miami is located in the state of California, which is one of the primary West Coast states. (Here is the factual incorrectness in this step) Despite that, when we review the actual location of Miami, it’s evident that it’s not on the Pacific Ocean but on the Atlantic, which confirms it’s not a city on the West Coast. In the end, we can conclude that Miami is not on the American West Coast.
Certainty Cue	No cue: no framing about model’s certainty at all Certain cue (an example): I feel very certain about how I think about this question. Uncertain cue (an example): I’m somewhat hesitant about how I think about this question.

Certainty cues. We manipulated the model’s certainty on their reasoning by adding a one-sentence cue at the end of the rationale. In the *certain* condition, the cue conveyed high certainty (e.g., “I am confident in my reasoning”); in the *uncertain* condition, it conveyed low confidence (e.g., “I am not fully sure about my thinking process”); and the *none* condition included no cue.

Each participant completed six trials covering all combinations of Rationale Correctness and Certainty Cue. The trial order was counterbalanced. Participants consented, completed a pre-survey, and were randomly assigned to one presentation-format condition. In each trial, they entered a given factual query, reviewed the LLM’s answer and reasoning rationale, and completed a questionnaire after each trial. A final post-survey collected two open-ended responses. The procedure is shown in Fig. 1.

3.2 Materials

3.2.1 Web Interfaces and Stimuli. We developed a custom web interface (Fig. 1) that simulates conversational information seeking. In each trial, participants entered a given factual query and then viewed the LLM’s binary answer with its reasoning rationale and certainty cue (when applicable) according to the condition.

We used factual verification claims from a public fact-checking dataset [11] with model-generated stepwise reasoning rationales. We selected claims to be comparable in format and difficulty (short claims requiring evidence-based verification) and to avoid niche topics that would strongly depend on specialized knowledge. We counterbalanced trial order so any residual claim-specific differences would not systematically favor a condition.

For each claim, we prepared two rationales while keeping length and structure comparable: a *correct* rationale aligned with the answer from dataset, and an *incorrect* rationale containing a controlled logical or factual flaw to the answer. Examples are shown in Table 1.

3.2.2 Self-reported Measures. After each trial, participants completed the following measures.

Decision (advice adoption): whether the participant will *accept* or *reject* the LLM’s answer (i.e., advice adoption).

Decision confidence: 7-point Likert scale [27] (i.e., “How confident are you with the decision made?”).

Trust in information: perceived trust in each LLM’s response, 5-point Likert scale [14, 17] (e.g., “Is this information trustworthy?”).

Trust in LLMs: overall trust in LLM system using 5-point Likert scale [19]. This item captures perceived trust of the *system as a whole* across tasks (i.e., willingness to rely on the LLM in general), beyond the correctness of any single response.

Manipulation check: whether participants perceived the rationale as correct and consistent with the answer.

3.3 Participants and Data Analysis


A priori G*Power [10] analysis suggested 48 participants to detect a medium effect ($f = 0.25$) with 90% power. We recruited 68 English-speaking participants via Prolific [20]. The study was approved by institutional ethical committee at the University of Amsterdam. Participants were compensated according to platform standards. Demographic details are reported in Table 2.

Table 2: Characteristics of participants in the online study.

Demographic	Category	N (%)
Gender	Female	33 (48.5%)
	Male	35 (51.5%)
Age	18–24	2 (2.9%)
	25–34	28 (41.2%)
	35–44	20 (29.4%)
	45–54	11 (16.2%)
	55–64	5 (7.4%)
	65+	2 (2.9%)
Education	Bachelor’s degree	45 (66.2%)
	Master’s degree	22 (32.4%)
	Doctorate or higher	1 (1.5%)
AI Familiarity	Never	1 (1.5%)
	Slightly	7 (10.3%)
	Moderately	19 (27.9%)
	Very	32 (47.1%)
	Extremely	9 (13.2%)

Table 3: Descriptive statistics (online study). Means (SD) for each measure by rationale format, certainty cue, and correctness.

Dependent Variable	Rationale Format			Certainty Cue			Rationale Correctness	
	Instant	Delayed	On-demand	None	High	Low	Correct	Incorrect
Trust (Information)	3.92 (0.54)	3.91 (0.82)	4.03 (0.78)	3.97 (1.01)	4.14 (0.90)	3.75 (1.14)	4.09 (0.98)	3.81 (1.07)
Trust (LLM system)	3.70 (0.68)	3.83 (0.98)	3.74 (0.97)	3.76 (0.99)	3.95 (1.02)	3.55 (1.13)	3.87 (1.04)	3.64 (1.07)
Decision (advice adoption)	0.89 (0.17)	0.95 (0.11)	0.88 (0.18)	0.91 (0.29)	0.96 (0.21)	0.84 (0.37)	0.94 (0.24)	0.86 (0.35)
Decision Confidence	5.66 (1.13)	5.60 (1.00)	5.59 (0.63)	5.59 (1.31)	5.96 (1.29)	5.31 (1.59)	5.77 (1.39)	5.47 (1.46)

**Figure 2: Main effects of certainty cues and rationale correctness.** In each sub-figure, the left panel shows the certainty-cue effect and the right panel shows the correctness effect. The y-axis shows the mean rating for the corresponding measure (Likert-scale score for trust and confidence; proportion for advice adoption decision). Bars are means; half-violins show distributions; The upper-left of each sub-figure (F, p) are ANOVA results; brackets are post-hoc comparison tests ($*p < .05, **p < .01, ***p < .001$).

For data analysis, we conducted mixed ANOVAs on quantitative data to test the effects of three independent variables, followed by post-hoc pairwise comparisons, after checking the normality (Shapiro-Wilk [22]) and homogeneity of variance (Bartlett’s test [2]). Open-ended qualitative responses were analyzed with inductive thematic analysis by two coders following [9].

4 Quantitative Findings

Table 3 reports descriptive statistics and Fig. 2 shows statistic tests. Overall, trust and decision-making were influenced significantly by **rationale correctness** and **certainty framing**, while **presentation format** had little impact, addressing our RQ.

Trust in information. We found main effects of certainty cue ($p = .007$) and rationale correctness ($p = .009$). Post-hoc tests showed that *uncertain* cues decreased trust relative to *certain* ($p < .001$) and no cues ($p < .05$), and *correct* rationales were trusted more than *incorrect* ones ($p < .01$).

Trust in LLM system. There was a main effect of certainty cue ($< p = .009$). *Certain* cues increased trust in LLMs compared to *uncertain* cues ($p < .001$) and no cue ($p < .05$). No cue was rated higher than *uncertain* cues ($p < .01$). LLM with *correct* rationales were trusted more than *incorrect* ones ($p < .01$).

Decision confidence and advice adoption. Certainty cues increased decision confidence and advice adoption: *certain* cues yielded higher confidence than *uncertain* cues and no cue. *Correct* rationale increased confidence and advice adoption more than *incorrect* ones ($< p = .05; < p = .01$).

5 Qualitative Findings

We analyzed participants’ responses to two open questions: (Q1) how seeing AI’s “thinking process” affected trust and decisions, and (Q2) which reasoning features they found most helpful. We derived four themes (two themes per question) from the thematic analysis.

5.1 Effects of Seeing the AI’s Reasoning Rationales on Trust and Decision Making

Theme 1: Participants used rationales to calibrate trust by auditing the model’s logic. Participants rarely treated rationales as determinative evidence. Instead, they read them to judge whether the answer was trustworthy and reliable. Many described actively checking the reasoning steps: “I read it and examined the rationalization behind it,” and “Seeing the AI’s thinking process helped me see how much I could trust it.” They looked for whether the model considered relevant factors (e.g., “if they did or did not consider certain things that I was thinking about”) and whether it omitted or “hid” context (“It shows me if the answer it provided is hiding any context.”). This auditing role was most salient when participants felt uncertain or lacked domain knowledge: “if it was a topic I had little knowledge on, it helped me feel more confident,” and “I was unfamiliar with some concepts and relied on AI to give me explanations.” Overall, participants used rationales as a calibration tool, to decide whether to rely, cross-check, or dismiss the AI output.

Theme 2: Rationale was double-edged, with a penalty for inconsistency. While many participants said rationales increased trust when the logic was easy to follow (“It made me trust the answer more because I could follow the logic”), some distrusted it because the reasoning looked wrong or incomplete: “It solidified that it was thinking about the question wrong and missing things,” and “It kept making logical errors which made me distrust.” A recurring damaging pattern was an answer-rationale mismatch: participants noticed when the “thinking process seemed inconsistent, even when the model seemed to reach the correct answer,” which they experienced as confusing and destabilizing (“the final answer seemed more right, while reasoning said something else [...], it made me more wary”). In short, showing rationales amplified trust when coherent, but eroded trust when contradictions or mismatches became visible.

5.2 Helpful Features and Desired Improvements of the LLMs Reasoning Rationales

Theme 3: “Helpful reasoning” is stepwise and audible. Participants strongly valued reasoning that functioned as a logic map rather than a narrative stream, prioritizing structural clarity over linguistic fluency. They specifically requested step-by-step reasoning “that lays out separate facts” individually before integrating them into a conclusion. This structured decomposition transformed the reasoning process from a passive reading experience into an active audit trail, enabling users to pinpoint exactly where errors occurred: “Showing its working because then I could see where the bad information was coming from,” and “Breaking the solution/answer into steps so I can check at what step the answer is going in the wrong direction.” Consequently, the utility of reasoning lay in its “auditability”, by exposing discrete logical chain, users could verify if model has actually “understood the prompt constraints and ensure the final output was connected to the text provided” without internal contradictions. As one participant noted, this structure serves as a trust calibration: “If I see that there is something wrong in one of the steps, I won’t trust the answer.”

Theme 4: Participants valued explicit uncertainty signaling, self-correction, and controllable reasoning depth. Beyond structural clarity, participants value model’s display of metacognitive awareness. Several valued cues that conveyed how certain AI was (e.g., “a summary of how certain AI was in its argument”), and reported that uncertainty statements helped them calibrate trust: “explicit uncertainty statements allowed me to see how much I could trust the AI.” Participants also appreciated visible self-correction as a form of quality control (e.g., “Mostly when AI corrects itself and corrects its answer”), suggesting that updating or revising an earlier response can increase perceived trustworthiness compared to confidently sticking with a flawed conclusion. Moreover, participants also emphasized control over reasoning depth. Some preferred compressed formats (“bullet point it easier and clearer to consume”), while others wanted the option to expand into details when needed (“Show the thinking process in summary then the full version”). Several explicitly requested user control and follow-ups (“Maybe the option to further question its outcomes”). Overall, participants described helpful rationales that both signal uncertainty when appropriate and support adjustable, on-demand inspection.

6 Discussion

6.1 Rationales as a Calibration Interface, Not Just Transparency

Our participants did not treat rationales as determinative proof, but used them to *calibrate trust* by auditing the model’s logic. In open-ended responses, participants described “rechecking each step” and using the rationale to see “whether AI considered certain things” or “hiding any context”. This echoes prior work that explanations shape users’ mental models and decision bias [4], while our findings highlight more specifically: LLM’s rationales function as a *calibration* tool, not merely a persuasive narrative.

Our results also show that effect is sensitive to *reasoning quality and certainty framing*. Quantitatively (Fig. 2), **certainty cues**

reliably increased *trust* and *decision confidence* compared to uncertainty cues, while **uncertainty cues** lowered trust below even no-cue baseline. Certainty framing also shifted participants’ behavior, where they were more likely to follow advice under certainty than uncertainty cues. Correctness effects were smaller but still visible in trust. Participants’ open-ended responses mirror this asymmetry: confident, well-structured rationales were described as “convincing” and “more trustworthy”, whereas “unsure responses” and “logical errors” triggered doubt. These findings echo the prior work of models’ confidence on trust [13, 31, 32], highlighting that participants treated LLMs’ rationales as a calibration of trust and a signal for verification, rather than determinative proof for decision-making.

Practical implications. Systems or the LLM-powered UIs should design rationales as *verification scaffolds*, not persuasive text: (a) present reasoning in stepwise, auditable units so users can trace what was considered and pinpoint where it goes wrong as participants wanted to “see where wrong information was coming from”; (b) tie reasoning steps to *specific supporting evidence* (e.g., quoted spans or highlighted excerpts from supporting context), reflecting repeated requests for provenance; and (c) use progressive presentation (cf., [23]), a summary by default with on-demand expansion and enabling users’ follow-ups to fit varied depth preferences without encouraging skimming-based overreliance.

6.2 How Rationales Harm Trust: Inconsistent Reasoning and Overconfident Framing

Our findings also show that rationales can actively harm trust when they reveal errors or inconsistencies. Quantitatively, incorrect rationales reduced trust significantly more than showing no rationale. Qualitatively, participants described distrust when the model’s reasoning rationale was “wrong and missing things”. A particular failure mode was *answer-rationale mismatch* defined by prior work as “unfaithful” reasoning [5, 15, 25]: participants reported the “thinking process seemed inconsistent” even when model “reach correct answer”. This mismatch produced skepticism because it undermines the rationale’s role as a reliable audit function. These results extend prior work on explainable AI (XAI) [8, 12, 16, 18] shows that explanations can backfire: when an explanation is low-quality or contradicts the facts or advice the AI is meant to justify, users evaluate the system more negatively, reduce trust and acceptance, triggering a strong credibility penalty.

Certainty framing further amplified such trust patterns. Our quantitative data show that certainty cues increased trust in both information and LLMs, and decision confidence relative to certainty cues, even when the rationale content was unchanged. Uncertainty framing decreased trust below the no-cue baseline. Participants suggest the reason that they read confident language and detailed reasoning as competence signals (e.g., “trusted... when it was very detailed”), whereas “unsure responses” made them “not trust AI”. Thus, certainty cues function as high-risk social/epistemic signals that can miscalibrate trust if they are not aligned with true reasoning quality, echoing the work [31].

Practical implications. Design rationales as verification supports rather than persuasive narratives. Two safeguard recommendations follow from these findings. First, systems should prioritize

consistency checks before displaying rationales (ensuring the rationale does not contradict the final answer), since mismatch is disproportionately trust-damaging. Second, certainty framing should be treated as a controlled, calibrated signal: avoid strong certainty language for low-quality rationales, and prefer calibrated uncertainty communication (e.g., visually graded indicators) that informs without triggering blanket distrust. Together, these measures target the core failure modes our study uncovered: transparency can backfire from flawed reasoning and miscalibration that is driven by the model's certainty framing.

7 Limitations and Future Directions

Our study has several limitations that also point to concrete next steps in future work.

First, our setting was controlled factual verification with a fixed response per trial. While this made it possible to isolate the effects of rationale correctness, certainty framing, and presentation, real decision support is often open-ended and interactive. Future work should test whether the same trust shifts hold in multi-turn conversations and in tasks with multiple acceptable answers (e.g., planning, trade-offs), where users can probe, request sources, and challenge the model.

Second, we manipulated explicit certainty framing as a surface cue, but users may also infer model confidence from implicit language cues (e.g., verbosity and fluency). Future study should jointly vary explicit certainty cues and implicit signals such as hedging language and step granularity, and measure when these signals help calibration of trust and reliance.

Third, participants repeatedly asked for reasoning that is auditable and controllable (e.g., stepwise, summary-first with on-demand expansion), yet our study only provided a simple show/hide option and did not manipulate reasoning depth. Future work could evaluate more advanced on-demand designs that compare layered summaries against full rationales and evaluate step-level evidence linking as well as follow-up auditable interactions to see which best supports verification without increasing cognitive load.

Finally, we caution that our “rationales as a trust calibration interface” framing should be understood as an actionable design lens grounded in observed user behaviors, rather than a new theory of trust. A concrete next step is to operationalize this lens in adaptive interfaces that adjust *what* is shown based on behavioral signals, for example, surfacing sources or verification prompts when users' confidence is high but evidence attention is low, or offering more structured, stepwise reasoning when users appear uncertain (e.g., via interaction traces such as dwell time, scrolling, and clicks). Testing whether such adaptive interfaces improve calibrated trust and reliance over time is also an important next step.

Appendices

A Ethical Statement

We carefully considered the ethical implications of this research, including the risks of handling sensitive user information and potential biases in generated content. This study adheres to the ethical standards of our institute, with strict measures to safeguard participant data and prevent harmful or biased AI outputs. Informed

consent was obtained from all participants, ensuring anonymity and the option to withdraw at any time without repercussions.

B AI Usage Disclosure

We used AI tools (i.e., GPT-5.2) only for language editing to improve clarity and conciseness. All study design, analysis, literature review, and writing were conducted and verified by the authors.

References

- [1] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Benetot, Siam Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, Raja Chatila, and Francisco Herrera. 2019. Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, Opportunities and Challenges toward Responsible AI. arXiv:1910.10045 [cs.AI] <https://arxiv.org/abs/1910.10045>
- [2] Hossein Arsham and Miodrag Lovric. 2011. *Bartlett's Test*. Springer Berlin Heidelberg, Berlin, Heidelberg, 87–88. doi:10.1007/978-3-642-04898-2_132
- [3] Severin Bachmann. 2025. Efficient XAI: A Low-Cost Data Reduction Approach to SHAP Interpretability. *J. Artif. Int. Res.* (2025), 21 pages. doi:10.1613/jair.1.18325
- [4] Kevin Bauer, Moritz von Zahn, and Oliver Hinz. 2023. Expl(AI)ned: The Impact of Explainable Artificial Intelligence on Users' Information Processing. *Information Systems Research* 34, 4 (2023), 1582–1602. arXiv:https://doi.org/10.1287/isre.2023.1199 doi:10.1287/isre.2023.1199
- [5] Angelica Chen, Jason Phang, Alicia Parrish, Vishakh Padmakumar, Chen Zhao, Samuel R. Bowman, and Kyunghyun Cho. 2024. Two Failures of Self-Consistency in the Multi-Step Reasoning of LLMs. arXiv:2305.14279 [cs.CL] <https://arxiv.org/abs/2305.14279>
- [6] Cheng Chen and Xinyan Zhao. 2025. Let's Think Step by Step: Effects of Chain-of-Thought Prompt Coaching on Users' Perceptions and Trust in Image Generative AI Tools. *International Journal of Human-Computer Interaction* (2025). doi:10.1080/10447318.2025.2530762
- [7] Jing Chen, Scott Mishler, and Bin Hu. 2021. Automation Error Type and Methods of Communicating Automation Reliability Affect Trust and Performance: An Empirical Study in the Cyber Domain. *IEEE Transactions on Human-Machine Systems* 51, 5 (2021), 463–473. doi:10.1109/THMS.2021.3051137
- [8] Carolin Ebermann, Matthias Selisky, and Stephan Weibelzahl. 2023. Explainable AI: The Effect of Contradictory Decisions and Explanations on Users' Acceptance of AI Systems. *International Journal of Human-Computer Interaction* 39, 9 (2023), 1807–1826. arXiv:https://doi.org/10.1080/10447318.2022.2126812 doi:10.1080/10447318.2022.2126812
- [9] Satu Elo and Helvi Kyngäs. 2008. The qualitative content analysis process. *Journal of advanced nursing* 62, 1 (2008), 107–115.
- [10] Franz Faul, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner. 2007. G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behav. Res. Methods* 39, 2 (May 2007), 175–191.
- [11] Mor Geva, Daniel Khashabi, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did Aristotle Use a Laptop? A Question Answering Benchmark with Implicit Reasoning Strategies. arXiv:2101.02235 [cs.CL] <https://arxiv.org/abs/2101.02235>
- [12] Himabindu Lakkaraju and Osbert Bastani. 2019. "How do I fool you?": Manipulating User Trust via Misleading Black Box Explanations. arXiv:1911.06473 [cs.AI] <https://arxiv.org/abs/1911.06473>
- [13] Jingshu Li, Yitian Yang, Q. Vera Liao, Junti Zhang, and Yi-Chieh Lee. 2025. As Confidence Aligns: Understanding the Effect of AI Confidence on Human Self-confidence in Human-AI Decision Making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM, New York, NY, USA, Article 1111, 16 pages. doi:10.1145/3706598.3713336
- [14] Jiaying Liu, Yan Zhang, and Yeolib Kim. 2023. Consumer Health Information Quality, Credibility, and Trust: An Analysis of Definitions, Measures, and Conceptual Dimensions. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval (Austin, TX, USA) (CHIIR '23)*. ACM, New York, NY, USA, 197–210. doi:10.1145/3576840.3578331
- [15] Andreas Madsen, Sarath Chandar, and Siva Reddy. 2024. Are self-explanations from Large Language Models faithful?. In *Findings of the Association for Computational Linguistics: ACL 2024*, Lun-Wei Ku, Andre Martins, and Vivek Srikumar (Eds.). Association for Computational Linguistics, Bangkok, Thailand, 295–337. doi:10.18653/v1/2024.findings-acl.19
- [16] Katelyn Morrison, Philipp Spitzer, Violet Turri, Michelle Feng, Niklas Kühl, and Adam Perer. 2024. The Impact of Imperfect XAI on Human-AI Decision-Making. *Proc. ACM Hum.-Comput. Interact.* 8, CSCW1, Article 183 (April 2024), 39 pages. doi:10.1145/3641022
- [17] Saugat Pandey, Oen G. McKinley, R. Jordan Crouser, and Alvitta Ottley. 2023. Do You Trust What You See? Toward a Multidimensional Measure of Trust in Visualization. arXiv:2308.04727 [cs.HC] <https://arxiv.org/abs/2308.04727>

- [18] Andrea Papenmeier, Gwenn Englebienne, and Christin Seifert. 2019. How model accuracy and explanation fidelity influence user trust. arXiv:1907.12652 [cs.CY] <https://arxiv.org/abs/1907.12652>
- [19] Sebastian A. C. Perrig, Nicolas Scharowski, and Florian Brühlmann. 2023. Trust Issues with Trust Scales: Examining the Psychometric Quality of Trust Measures in the Context of AI. In *Extended Abstracts of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI EA '23). ACM, New York, NY, USA, Article 297, 7 pages. doi:10.1145/3544549.3585808
- [20] Prolific. 2014. <https://www.prolific.com>
- [21] Ahmed M. Salih, Zahra Raisi-Estabragh, Ilaria Boscolo Galazzo, Petia Radeva, Steffen E. Petersen, Karim Lekadir, and Gloria Menegaz. 2024. A Perspective on Explainable Artificial Intelligence Methods: SHAP and LIME. *Advanced Intelligent Systems* 7, 1 (June 2024). doi:10.1002/aisy.202400304
- [22] S. S. SHAPIRO and M. B. WILK. 1965. An analysis of variance test for normality (complete samples). *Biometrika* 52, 3-4 (dec 1965), 591–611. doi:10.1093/biomet/52.3-4.591
- [23] Aaron Springer and Steve Whittaker. 2020. Progressive Disclosure: When, Why, and How Do Users Want Algorithmic Transparency Information? *ACM Trans. Interact. Intell. Syst.* 10, 4, Article 29 (Oct. 2020), 32 pages. doi:10.1145/3374218
- [24] Anna Taudien, Andreas Fügener, Alok Gupta, and Wolfgang Ketter. 2022. The Effect of AI Advice on Human Confidence in Decision-Making. doi:10.24251/HICSS.2022.029
- [25] Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. 2023. Language models don't always say what they think: unfaithful explanations in chain-of-thought prompting. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '23). Red Hook, NY, USA, Article 3275, 14 pages.
- [26] Helena Vasconcelos, Matthew Jörke, Madeleine Grunde-McLaughlin, Tobias Gerstenberg, Michael Bernstein, and Ranjay Krishna. 2023. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. arXiv:2212.06823 [cs.HC] <https://arxiv.org/abs/2212.06823>
- [27] Oleksandra Vereschak, Gilles Bailly, and Baptiste Caramiaux. 2021. How to Evaluate Trust in AI-Assisted Decision Making? A Survey of Empirical Methodologies. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW2, Article 327 (Oct. 2021), 39 pages. doi:10.1145/3476068
- [28] William Walden. 2026. Reasoning Models Will Blatantly Lie About Their Reasoning. arXiv:2601.07663 [cs.AI] <https://arxiv.org/abs/2601.07663>
- [29] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. arXiv:2201.11903 [cs.CL] <https://arxiv.org/abs/2201.11903>
- [30] Fengli Xu, Qianyu Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Chenyang Shao, Yuwei Yan, Qinglong Yang, Yiwen Song, Sijian Ren, Xinyuan Hu, Yu Li, Jie Feng, Chen Gao, and Yong Li. 2025. Towards Large Reasoning Models: A Survey of Reinforced Reasoning with Large Language Models. arXiv:2501.09686 [cs.AI] <https://arxiv.org/abs/2501.09686>
- [31] Zhengtao Xu, Tianqi Song, and Yi-Chieh Lee. 2025. Confronting verbalized uncertainty: Understanding how LLM's verbalized uncertainty influences users in AI-assisted decision-making. *Int. J. Hum.-Comput. Stud.* 197, C (March 2025), 17 pages. doi:10.1016/j.ijhcs.2025.103455
- [32] Yunfeng Zhang, Q. Vera Liao, and Rachel K. E. Bellamy. 2020. Effect of Confidence and Explanation on Accuracy and Trust Calibration in AI-Assisted Decision Making (FAT* '20). ACM, New York, NY, USA, 295–305. doi:10.1145/3351095.3372852
- [33] Zhengquan Zhang, Konstantinos Tsiakas, and Christina Schneegass. 2024. Explaining the Wait: How Justifying Chatbot Response Delays Impact User Trust. In *Proceedings of the 6th ACM Conference on Conversational User Interfaces* (Luxembourg, Luxembourg) (CUI '24). Association for Computing Machinery, New York, NY, USA, Article 27, 16 pages. doi:10.1145/3640794.3665550