Human Factors in Design, Engineering, and Computing, Vol. 238, 2026, 718–727
https://doi.org/10.54941/ahfe1008303

When Helpful AI Feels Too Trustworthy: Engagement and Reliance in Design Ideation

Ting Liao and Jiacheng Sun

Department of Systems Engineering, Stevens Institute of Technology, Hoboken, NJ 07030, USA

ABSTRACT

Artificial intelligence (AI) assistants are increasingly used in engineering ideation, but their effectiveness depends on both the AI model's capabilities and how users interpret and act on AI feedback. This paper synthesizes human-factors findings from the development and evaluation of the AI Design Assistant (AIDA), a modular testbed for iterative engineering ideation, focusing on trust calibration, feedback presentation, user agency, and design exploration. Across studies, adaptive and socially responsive feedback generally increased trust and perceived usefulness, yet trust did not reliably correspond to measured design improvement. More feedback yielded mixed results: longer or highly structured responses were not consistently associated with stronger outcomes, while concise feedback and longer user revisions were associated with larger score gains in turn-level analyses. Users also applied AI guidance selectively; only 38.6% of actionable suggestions were implemented, although most revisions used at least one AI-generated suggestion. These findings suggest that the next generation of AI design assistants needs to support calibrated trust, reduce unnecessary information-processing demands, preserve explicit user control over suggestion use, and distinguish user perceptions from measured design outcomes. The results provide practical design guidance for human-centered AI systems that augment, rather than replace, human judgment during creative engineering work.

Keywords: Human-AI collaboration, Decision support, Trustworthy AI, Adaptive system

INTRODUCTION

Artificial intelligence (AI) has become an active participant in engineering ideation, supporting the generation, critique, evaluation, and refinement of design concepts (Chen et al., 2025; Wang et al., 2025). However, its effectiveness depends not only on model capability but also on how users interpret and respond to AI-generated feedback (Yang et al., 2020; Tankelevitch et al., 2024). Designers must decide whether recommendations are relevant, how much attention to allocate to them, and whether following them is likely to improve the design rather than make the system appear helpful or competent (Lee and See, 2004; Gmeiner et al., 2023; Schemmer et al., 2023).

Received September 26, 2026; Revised October 30, 2026; Accepted November 7, 2026; Available online December 1, 2026

© 2026 The Authors. This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 License.
For more information, see https://creativecommons.org/licenses/by-nc-nd/4.0/

These decisions are particularly difficult in open-ended design, where the value of AI feedback is often uncertain and may involve trade-offs among competing design objectives. The challenge is to support appropriately calibrated reliance (Lee and See, 2004; Hoff and Bashir, 2015). To investigate these relationships, we developed the AI Design Assistant (AIDA), a modular testbed that supports iterative engineering ideation and enables systematic variation of AI characteristics and interaction behaviors (Mathews Renji et al., 2024; Mathews Renji and Liao, 2025).

This paper synthesizes findings across the AIDA studies around a recurring tension: features that make AI support more responsive and engaging can increase trust and perceived usefulness without necessarily producing corresponding gains in design quality. We focus on four related issues, including trust and reliance, feedback presentation, selective use of AI suggestions, and the balance between exploration and refinement, to identify implications for the design of AI-supported engineering ideation.

BACKGROUND

Trust and Reliance in Human-AI Design Collaboration

Trust plays a central role in determining whether users rely on AI-generated recommendations, especially when they cannot directly verify the quality of the system output. In human automation research, trust is not simply a positive attitude toward a system. It reflects a willingness to depend on the system under uncertainty and to allow its recommendations to influence task decisions (Lee and See, 2004; Hoff and Bashir, 2015). Both excessive and insufficient trust can create problems. Excessive trust can lead users to accept recommendations with limited scrutiny, while insufficient trust can cause them to ignore useful assistance. Recent research emphasizes appropriate reliance rather than maximizing trust. The objective is to align users' reliance with the actual capabilities and limitations of the AI system (Jacovi et al., 2021; Schemmer et al., 2023; Mehrotra et al., 2024).

This calibration is more difficult in engineering ideation than in tasks where system performance can be evaluated against a correct answer. Design concepts are typically assessed across multiple criteria, such as novelty, feasibility, sustainability, and usefulness, and improvement in one dimension may come at the expense of another (Shah et al., 2003; Han et al., 2021). Users often need to decide whether AI feedback is worth following before its contribution to the design is fully known. Design also unfolds over time through repeated cycles of refinement and evaluation. A recommendation may appear reasonable and specific, but its actual value may only become clear after further revision or evaluation. In this context, trust is partly a judgment about the credibility of guidance rather than a direct response to observable performance.

AI further complicates this judgment because the presentation of feedback can itself influence perceived trustworthiness. Fluent language, personalization, supportive tone, and emotional acknowledgment can make an AI assistant appear more attentive and competent (Rheu et al., 2021; Liao and Yan, 2023). These characteristics can improve engagement and make

critique easier to receive. However, they may also increase trust without providing stronger evidence that the recommendation will improve the design. Prior research on conversational agents similarly shows that social and emotional cues can influence trust and perceived helpfulness even when the underlying task capability remains unchanged (Rheu et al., 2021; Liao and Yan, 2023).

AI Feedback and User Agency

AI feedback also creates cognitive demands when users must identify relevant information, compare suggestions with their own goals, and decide what to revise. This shifts some effort from producing ideas to monitoring, evaluating, and integrating system output (Tankelevitch et al., 2024; Xu et al., 2025). More information is not necessarily more useful. Feedback should instead make the revision target clear and provide enough information for users to judge whether a recommendation is relevant and actionable (Gmeiner et al., 2023; Zhang et al., 2025).

These decisions are closely related to user agency, which concerns users' ability to direct the collaborative process and determine how AI contributions influence their work. In human-AI collaboration, users may adopt a suggestion directly, modify it, combine it with their own ideas, defer it, or reject it in favor of another direction (Moruzzi and Margarido, 2024; Rezwana and Ford, 2025). Agency in this setting does not depend on minimizing AI use or maximizing acceptance. Rather, it depends on whether users can make informed choices about when and how AI should influence the developing concept. Research increasingly describes such control as something that evolves throughout the interaction and depends partly on the mechanisms the interface provides for inspecting, modifying, and redirecting AI contributions (Tankelevitch et al., 2024; Issak et al., 2026).

Human-AI collaboration should not be evaluated simply by how much AI-generated content appears in the final result. Studies of AI-assisted writing show that users can rewrite suggested text and develop AI-proposed ideas through subsequent work (Zhou and Sterman, 2024; Bhat et al., 2026). Preserving this selectivity is especially important in ideation because generative AI can expand the solution space while also encouraging fixation on AI-proposed directions (Wadinambiarachchi et al., 2024). Effective AI support must help users understand and evaluate suggestions while preserving their ability to choose which directions to develop.

AIDA as a Modular Testbed for Design Ideation

To test how these system and interaction characteristics shape user trust, reliance, and design outcomes, we developed AIDA as a modular experimental platform for early-stage engineering design. AIDA is a web-based conversational assistant in which users submit a concept, receive evaluations and feedback, and may revise the idea over multiple turns. This iterative structure reflects design as a process of generating, evaluating, and refining ideas rather than producing a final solution in one step (Boudier et al., 2023).

AIDA evolved through several phases of development. Early phases used structured stage-gate logic with fine-tuned RoBERTa classifiers for design evaluation (Figure 1a). Later phases incorporated GPT-based evaluation and increasingly adaptive responses. Additional experiments varied human-like versus abstract appearance, detailed versus concise feedback, and empathetic versus non-empathetic behavior (Mathews Renji et al., 2024; Mathews Renji and Liao, 2025).

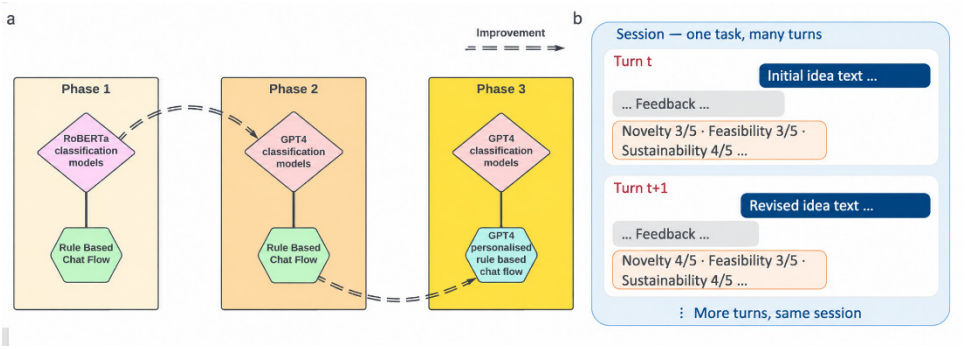


Figure 1: Illustrations of (a) AIDA development phases; (b) an interaction session.

Later studies expanded AIDA to four sustainable-design tasks: washing machines, household lighting, heating, ventilation, and air-conditioning (HVAC) systems, and toilet systems (Figure 1b). Users submit a concept, receive criterion-specific scores and feedback, and may revise and resubmit it. In these analyses, each criterion—novelty, feasibility, and sustainability—was scored from 1 to 5, and the composite score was their unweighted mean. Session-level change was the final minus initial composite score; turn-level change compared consecutive scored submissions. The logs preserve user submissions, AI responses, scores, and their order, enabling complementary analyses of appraisal and revision behavior.

EMPIRICAL RESULTS AND CROSS-STUDY OBSERVATIONS

Divergence Between Ideation Change and User Appraisal

Across 183 iterative ideation sessions from 66 users, the mean AIDA-assigned composite score increased from 2.823 to 3.434, corresponding to a mean within-session increase of 0.611 points. In the linked survey sample of 167 sessions from 53 users, mean trust was 3.410 and mean perceived usefulness was 3.629 on five-point scales, both significantly above the midpoint. Trust averaged 11 items after reverse-coding three risk/caution items; usefulness used one item. These analyses draw on an archive of 219 sessions from 73 users, with sample sizes determined by score and survey availability.

However, trust did not reliably correspond to AI-assessed ideation change (i.e., session-level change in the composite score). In the joint sample of 141 sessions from 48 users, AI-assessed ideation change was not significantly correlated with trust ($r = -.051$, $p = .581$) or perceived usefulness

($r = -.074$, $p = .419$). The same pattern remained after controlling for initial AIDA-assigned composite score, AIDA version, design task, and participant fixed effects. AI-assessed ideation change also remained nonsignificant in the adjusted models for trust and perceived usefulness. Correlation tests used participant-cluster bootstrap resampling, and adjusted models used participant-clustered standard errors.

These findings suggest that users respond to both what the AI contributes and how the interaction feels. Socially supportive feedback can improve engagement, but it may also make the system appear more reliable before the value of its recommendations is established. This creates a calibration challenge: interaction cues that make AI pleasant to use should not become substitutes for evidence of design value.

Feedback Quantity and Effectiveness

The AIDA studies also show that additional feedback may not necessarily be more useful. Longer or highly structured responses were not consistently associated with stronger outcomes. In the turn-level analysis of 426 consecutive submission pairs, shorter AIDA responses and longer user revisions were associated with larger increases in the AIDA-assigned composite score. Response and revision length were measured by word count; mixed models included design task and random intercepts for participants and sessions nested within participants.

This does not imply that feedback should always be short. Rather, it suggests that information volume may introduce additional processing demands. Users must identify the relevant issue, interpret the recommendation, connect it to their own idea, and decide what to change. AI can increase cognitive and metacognitive work even when it reduces effort elsewhere (Tankelevitch et al., 2024; Xu et al., 2025).

Feedback should consequently prioritize actionable information. A concise response that identifies the problem, explains why it matters, and proposes a possible next step may be more useful than an exhaustive critique. Additional detail can be provided when needed.

Selective Use of AI Suggestions

The analysis showed that users did not simply accept or reject AIDA feedback as a whole. Among 125 feedback–revision pairs from 78 sessions and 40 users, 80.0% of revisions implemented at least one AI suggestion, but only 38.6% of 435 actionable suggestions were implemented. One engineering-design expert compared the pre-feedback concept, AI response, and next revision. Implementation required substantive new content addressing a suggested action; pre-existing content and topical overlap alone did not count. Thus, users generally incorporated suggestions selectively. Semantic similarity, computed as cosine similarity between Sentence-BERT representations of the AI response and next revision, did not reliably distinguish observable implementation. Absence of implementation in the next revision does not establish deliberate rejection.

AI Guidance in Design Revision

AIDA's effects also differed across design criteria. Improvements were often stronger in feasibility and sustainability than in novelty. In one comparison, adaptive feedback was associated with greater sustainability score gains but lower novelty score gains than the generic iterative baseline. These results suggest that criterion-based feedback is associated with different patterns of revision across design dimensions. The pattern raises the possibility that feedback focused on particular criteria makes some improvement paths more salient than others. Prior research has similarly shown that generative AI can support ideation while sometimes narrowing exploration or contributing to design fixation (Wadinambiarachchi et al., 2024).

It is essential to understand what type of improvement an AI system encourages. Systems should help users recognize when they are exploring new directions versus refining an existing concept.

Evaluation Measures

A final cross-study finding concerns evaluation measures. Trust, perceived usefulness, AI-assessed ideation change, semantic similarity, and observed suggestion implementation capture different aspects of the interaction, and AIDA analyses show that these measures can diverge.

Design support systems can create value through changes in both the design artifact and the user experience. These outcomes may be related, but one does not establish the other (Remy et al., 2020; Rhys Cox et al., 2025). Similarly, trust and usefulness are distinct from behavioral reliance, which concerns whether users actually act on AI suggestions (Davis, 1989; Lee and See, 2004; Schemmer et al., 2023).

Automated evaluations require the same caution. An AI-generated score reflects alignment with the system's own criteria, but it should not automatically be treated as an independent measure of design quality. Stronger claims about design improvement require independent evaluation (Remy et al., 2020; Park et al., 2026). Table 1 summarizes these cross-study lessons.

Table 1: Cross-study lessons from AIDA.

Interaction Issue	Evidence Across AIDA Studies	Implication
Trust and engagement	Adaptive, empathetic, and socially responsive feedback can increase favorable perceptions even when AI-assessed ideation change is limited or does not correspond to user appraisal.	Support trust through understandable evidence rather than social cues alone.
Feedback quantity	Longer or more structured responses were not consistently associated with stronger outcomes.	Prioritize relevant, actionable information and allow additional detail on demand.
Selective reliance	Users typically implemented some, but not all, AI suggestions, and semantic similarity did not reliably indicate implementation.	Preserve explicit options to accept, adapt, defer, or reject AI suggestions.

(Continued)

Table 1: Continued.

Interaction Issue	Evidence Across AIDA Studies	Implication
Exploration vs. refinement	Criterion-based feedback tended to strengthen some dimensions more than others.	Make the direction and trade-offs of AI guidance visible.
Evaluation	Trust, usefulness, AI-assessed ideation change, semantic similarity, and observed implementation did not behave as interchangeable measures.	Evaluate perception, behavior, and design outcomes separately.

DESIGN IMPLICATIONS

The cross-study findings suggest several priorities for designing and evaluating AI-assisted ideation systems (Table 1).

First, AI systems should support **engagement without treating increased trust as the objective**. Empathy and supportive language can remain part of the interaction, but recommendations should also indicate the criterion being addressed and the reasoning behind the suggestion. This helps users distinguish conversational confidence from design evidence (Schemmer et al., 2023).

Second, feedback should be designed for **cognitive economy**. Rather than maximizing detail, the assistant can prioritize a small number of relevant revision targets and make additional explanation available when requested. This can reduce unnecessary processing while preserving access to richer information.

Third, systems should preserve **selective user control**. Users should be able to accept, adapt, defer, or reject suggestions without implying that more adoption is necessarily better. Future interfaces could make these decisions visible so that both users and researchers can distinguish different forms of reliance.

Fourth, AI systems should make **design trade-offs visible**. Suggestions could be grouped according to whether they support novelty, feasibility, sustainability, or another objective. Users could also choose whether they want to explore alternatives, refine the current concept, or evaluate trade-offs. Such features could help prevent the AI's evaluation criteria from silently defining what counts as progress.

DISCUSSION

The AIDA studies show that engagement, trust, reliance, and AI-assessed ideation change are related but distinct aspects of human-AI collaboration. Adaptive and socially responsive feedback can make an assistant feel more useful and trustworthy, even when changes in AIDA-assigned composite scores are limited or concentrated in specific criteria. This is especially important in open-ended design, where users often cannot immediately verify whether a recommendation will improve the concept.

At the same time, users demonstrated selective use of AI guidance. They typically incorporated some suggestions while leaving others unused, indicating that human judgment continued to shape how AIDA feedback

entered subsequent revisions. The design objective of future AI support should be to support better decisions about when and how AI feedback is used, rather than to maximize AI influence. Effective systems should help users understand the purpose of a recommendation, recognize relevant trade-offs, compare alternatives, and decide whether the feedback aligns with their design goals.

We acknowledge several limitations in this paper. The studies synthesized here used partially overlapping datasets and should be interpreted as complementary analyses rather than independent replications. Some design measures were generated by AIDA during interaction and reflected changes within its evaluation framework rather than independent ground truth. In addition, observed relationships involving feedback length, empathy, sentiment, and revision behavior are primarily associative and should not be interpreted as causal. The studies also focused largely on short text-based ideation tasks, which may not fully represent professional design practice.

Future work can use AIDA's modular architecture to examine these issues more directly. Promising directions include measuring trust and reliance across individual turns, independently manipulating feedback characteristics, using external design evaluation, and testing features such as progressive feedback disclosure, explicit acceptance or rejection of suggestions, and separate modes for exploration and refinement.

CONCLUSION

Across AIDA studies, favorable perceptions of AI support did not necessarily correspond to larger gains in AIDA-assigned composite scores. Trust did not reliably track AI-assessed ideation change, additional feedback was not always beneficial, users selectively incorporated AI suggestions, and different forms of guidance were associated with different patterns of change across design criteria.

These findings suggest that effective AI design support should help users evaluate and selectively act on feedback rather than simply increase engagement or reliance. Concise feedback, visible trade-offs, and meaningful user control may allow AI to support creative engineering work while preserving the human judgment that remains essential to the design process.

ACKNOWLEDGMENT

This work was supported by the National Science Foundation (NSF) under Award No. 2301846 and 2543591.

REFERENCES

Bhat, A., Le Quéré, M.A., Naaman, M. and Jakesch, M. (2026) 'Reactive writers: How co-writing with AI changes how we engage with ideas', in *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, pp. 1–21. https://doi.org/10.1145/3772318.3791529

Boudier, J., Sukhov, A., Netz, J., Le Masson, P. and Weil, B. (2023) 'Idea evaluation as a design process: Understanding how experts develop ideas and manage fixations', *Design Science*, 9, e9. https://doi.org/10.1017/dsj.2023.7

Chen, L., Song, Y., Guo, J., Sun, L., Childs, P. and Yin, Y. (2025) 'How generative AI supports human in conceptual design', *Design Science*, 11, e9. https://doi.org/10.1017/dsj.2025.2

Davis, F.D. (1989) 'Perceived usefulness, perceived ease of use, and user acceptance of information technology', *MIS Quarterly*, 13(3), pp. 319–340. https://doi.org/10.2307/249008

Gmeiner, F., Yang, H., Yao, L., Holstein, K. and Martelaro, N. (2023) 'Exploring challenges and opportunities to support designers in learning to co-create with AI-based manufacturing design tools', in *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, pp. 1–20. https://doi.org/10.1145/3544548.3580999

Han, J., Forbes, H. and Schaefer, D. (2021) 'An exploration of how creativity, functionality, and aesthetics are related in design', *Research in Engineering Design*, 32(3), pp. 289–307. https://doi.org/10.1007/s00163-021-00366-9

Hoff, K.A. and Bashir, M. (2015) 'Trust in automation: Integrating empirical evidence on factors that influence trust', *Human Factors*, 57(3), pp. 407–434. https://doi.org/10.1177/0018720814547570

Issak, A., Rezwana, J. and Harteveld, C. (2026) 'Control is a trajectory, not a point: Conceptualizing control in human-AI co-creativity', in *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*, pp. 1–17. https://doi.org/10.1145/3772318.3790861

Jacovi, A., Marasović, A., Miller, T. and Goldberg, Y. (2021) 'Formalizing trust in artificial intelligence: Prerequisites, causes and goals of human trust in AI', in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 624–635. https://doi.org/10.1145/3442188.3445923

Lee, J.D. and See, K.A. (2004) 'Trust in automation: Designing for appropriate reliance', *Human Factors*, 46(1), pp. 50–80. https://doi.org/10.1518/hfes.46.1.50_30392

Liao, T. and Yan, B. (2023) 'Let's chat if you are unhappy: The effect of emotions on interaction experience and trust toward empathetic chatbots', in *Volume 3B: 49th Design Automation Conference (DAC)*, V03BT03A057. American Society of Mechanical Engineers. https://doi.org/10.1115/DETC2023-115318

Mathews Renji, N. and Liao, T. (2025) 'Designing trust: How system-level design attributes shape AI-assisted design ideation', in *Volume 4: 22nd International Conference on Design Education, 30th Design for Manufacturing and the Life Cycle Conference, and 37th International Conference on Design Theory and Methodology*, V004T06A042. https://doi.org/10.1115/DETC2025-168783

Mathews Renji, N., Veera, S.C.M., Yan, B. and Liao, T. (2024) 'Augmenting engineering design with AI: Introducing the AI Design Assistant (AIDA)', in Volume 6: 36th International Conference on Design Theory and Methodology, V006T06A016. https://doi.org/10.1115/DETC2024-143765

Mehrotra, S., Degachi, C., Vereschak, O., Jonker, C.M. and Tielman, M.L. (2024) 'A systematic review on fostering appropriate trust in human-AI interaction: Trends, opportunities and challenges', *ACM Journal on Responsible Computing*, 1(4), pp. 26:1–26:45. https://doi.org/10.1145/3696449

Moruzzi, C. and Margarido, S. (2024) 'A user-centered framework for human-AI co-creativity', in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, pp. 1–9. https://doi.org/10.1145/3613905.3650929

Park, H., Huynh, K., Eiband, M., Dillmann, J., Mayer, S. and Sedlmair, M. (2026) 'Evaluating generative AI in the lab: Methodological challenges and guidelines', in *Proceedings of the 31st International Conference on Intelligent User Interfaces*, pp. 599–616. https://doi.org/10.1145/3742413.3789065

Remy, C., MacDonald Vermeulen, L., Frich, J., Biskjaer, M.M. and Dalsgaard, P. (2020) 'Evaluating creativity support tools in HCI research', in *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, pp. 457–476. https://doi.org/10.1145/3357236.3395474

Rezwana, J. and Ford, C. (2025) 'Human-centered AI communication in co-creativity: An initial framework and insights', in *Proceedings of the 2025 Conference on Creativity and Cognition*, pp. 651–665. https://doi.org/10.1145/3698061.3726932

Rheu, M., Shin, J.Y., Peng, W. and Huh-Yoo, J. (2021) 'Systematic review: Trust-building factors and implications for conversational agent design', *International Journal of Human–Computer Interaction*, 37(1), pp. 81–96. https://doi.org/10.1080/10447318.2020.1807710

Rhys Cox, S., Bøjer Djernæs, H. and van Berkel, N. (2025) 'Beyond productivity: Rethinking the impact of creativity support tools', in *Proceedings of the 2025 Conference on Creativity and Cognition*, pp. 735–749. https://doi.org/10.1145/3698061.3726924

Schemmer, M., Kuehl, N., Benz, C., Bartos, A. and Satzger, G. (2023) 'Appropriate reliance on AI advice: Conceptualization and the effect of explanations', in *Proceedings of the 28th International Conference on Intelligent User Interfaces*, pp. 410–422. https://doi.org/10.1145/3581641.3584066

Shah, J.J., Smith, S.M. and Vargas-Hernandez, N. (2003) 'Metrics for measuring ideation effectiveness', *Design Studies*, 24(2), pp. 111–134. https://doi.org/10.1016/S0142-694X(02)00034-0

Tankelevitch, L., Kewenig, V., Simkute, A., Scott, A.E., Sarkar, A., Sellen, A. and Rintel, S. (2024) 'The metacognitive demands and opportunities of generative AI', in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, pp. 1–24. https://doi.org/10.1145/3613904.3642902

Wadinambiarachchi, S., Kelly, R.M., Pareek, S., Zhou, Q. and Velloso, E. (2024) 'The effects of generative AI on design fixation and divergent thinking', in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pp. 1–18. https://doi.org/10.1145/3613904.3642919

Wang, W.F., Lu, C.T., Ponsa i Campanyà, N., Chen, B.Y. and Chen, M.Y. (2025) 'AIdeation: Designing a human-AI collaborative ideation system for concept designers', in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–28. https://doi.org/10.1145/3706598.3714148

Xu, X., Konnova, A., Gao, B., Peng, C., Vo, D. and Dow, S.P. (2025) 'Productive vs. reflective: How different ways of integrating AI into design workflows affect cognition and motivation', in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–15. https://doi.org/10.1145/3706598.3713649

Yang, Q., Steinfeld, A., Rosé, C. and Zimmerman, J. (2020) 'Re-examining whether, why, and how human-AI interaction is uniquely difficult to design', in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13. https://doi.org/10.1145/3313831.3376301

Zhang, C., Ju, K., Bidoshi, P., Yen, Y.C.G. and Rzeszotarski, J.M. (2025) 'Friction: Deciphering writing feedback into writing revisions through LLM-assisted reflection', in *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, pp. 1–27. https://doi.org/10.1145/3706598.3714316

Zhou, D. and Sterman, S. (2024) 'Ai.llude: Investigating rewriting AI-generated text to support creative expression', in *Proceedings of Creativity and Cognition*, pp. 241–254. https://doi.org/10.1145/3635636.3656187