Published: 02 January 2026

HAI '25: International Conference on Human-Agent Interaction
November 10 - 13, 2025
Yokohama, Japan

POSTER

FeedQUAC: Quick Unobtrusive AI-Generated Commentary

TAO LONG, Autodesk Inc., San Francisco, CA, United States

KENDRA WANNAMAKER, Autodesk Inc., San Francisco, CA, United States

JO VERMEULEN, Autodesk Inc., San Francisco, CA, United States

GEORGE FITZMAURICE, Autodesk Inc., San Francisco, CA, United States

JUSTIN MATEJKA, Autodesk Inc., San Francisco, CA, United States

Open Access Support provided by:

Autodesk Inc.

FeedQUAC: Quick Unobtrusive AI-Generated Commentary

Tao Long*
Autodesk Research
San Francisco, California, USA
long@cs.columbia.edu

Kendra Wannamaker
Autodesk Research
Toronto, Ontario, Canada
kendra.wannamaker@autodesk.com

Jo Vermeulen
Autodesk Research
Toronto, Ontario, Canada
jo.vermeulen@autodesk.com

George Fitzmaurice
Autodesk Research
Toronto, Ontario, Canada
george.fitzmaurice@autodesk.com

Justin Matejka
Autodesk Research
Toronto, Ontario, Canada
justin.matejka@autodesk.com

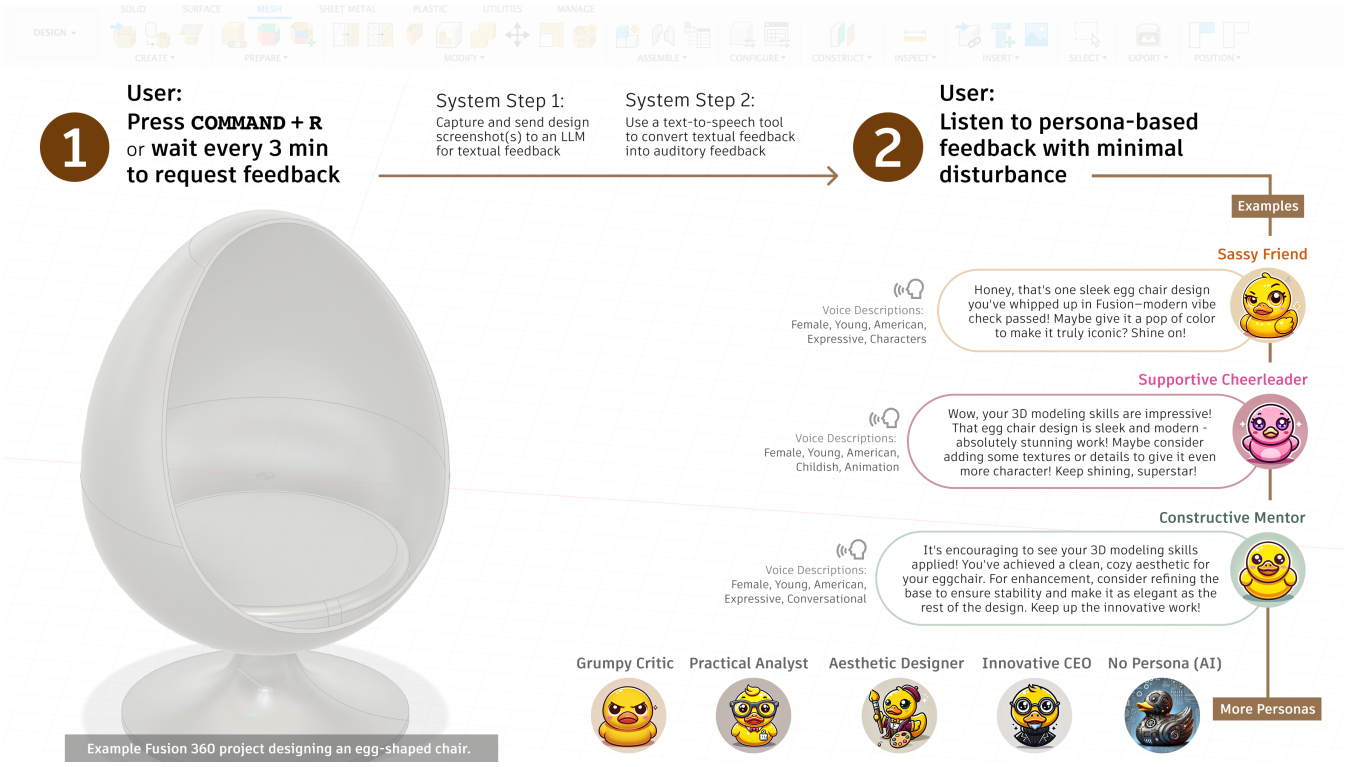

Figure 1: FeedQUAC is an ambient design companion that provides real-time, read-aloud AI feedback from diverse personas. Through shortcut keys, minimal screen space, and auditory cues, it offers quick, unobtrusive support that enhances design without disrupting workflow. See full paper here: https://arxiv.org/abs/2504.16416

Abstract

Design thrives on feedback. However, gathering constant feedback throughout the design process can be labor-intensive and disruptive. We explore how AI can bridge this gap by providing effortless, ambient feedback. We introduce FeedQUAC, a lightweight design companion that delivers real-time, read-aloud, AI-generated commentary from diverse personas based on live screenshots of the designer's workspace. FeedQUAC is always available, context-aware, ambient, playful, and iteration-aware. In a design probe with eight 3D CAD designers, participants highlighted convenience, playfulness, confidence boosts, and inspiration. Our findings suggest that ambient interaction is a valuable consideration for both designing and evaluating future creativity support systems.

*Also affiliated with Columbia University.

HAI '25, Yokohama, Japan

ACM ISBN 979-8-4007-2178-6/25/11
https://doi.org/10.1145/3765766.3765861

CCS Concepts

• **Human-centered computing → Interactive systems and tools**; *Auditory feedback*; *Empirical studies in interaction design.*

Keywords

design feedback, 3D design, creativity supports, evaluation, review, critique, CAD, persona, agents, ambient technology, AI companion

ACM Reference Format:
Tao Long, Kendra Wannamaker, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 2025. FeedQUAC: Quick Unobtrusive AI-Generated Commentary. In *13th International Conference on Human-Agent Interaction (HAI '25), November 10–13, 2025, Yokohama, Japan.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3765766.3765861

1 Introduction: Need for Constant AI Feedback

Feedback is crucial for understanding and improving designs [35], supporting iterative refinement across design stages and expertise levels [3, 5, 14, 15, 37]. Prior work highlights the value of short-form, real-time feedback throughout the process: quick, targeted check-ins enable timely adjustments and reduce late-stage issues [11, 15]. Yet frequent human feedback is hard to obtain in practice due to availability constraints, social dynamics, and limited diversity of perspectives [9, 13, 16, 23]. Our formative exploration of multiple feedback forums further suggested additional friction—slow response times, high-effort context posts, and uneven quality—often misaligned with a designer's current stage [9].

Recent AI tools built on language and vision models show strong context understanding, aiding tasks like brainstorming [4, 19, 20, 24] and drafting [22, 25, 33, 36]. However, in the reviewing stage of the creative process [12], most tools target prompt engineering [6, 38], academic writing and peer review [8, 26, 31], or language practice [27, 30]. By contrast, relatively little work examines visual design evaluation—whether open-ended artistic visions or mechanical requirements—leaving a gap in AI tools for continuous feedback across design phases.

We design FeedQUAC, an ambient AI companion that observes live screenshots of the designer's workspace, produces persona-driven, iteration-aware commentary, and reads it aloud on demand—akin to "rubber duck debugging" for design [32], and informed by HAI and ambient technology works [17, 18, 21, 29, 34].

2 FeedQUAC System

FeedQUAC is designed to be always-available to 3D CAD designers, as it sits as a small floating duck icon anchored at the bottom-right of the editor—always on top yet out of the way—and can be invoked instantly via hotkeys for emoji validation (`Command/Windows + E`) or a short persona-driven voice review (`Command/Windows + R`); optional automatic cycles (e.g., every 30 seconds, 3 minutes, or 5 minutes) keep feedback coming without manual effort.

Together with the optional interaction cycles, FeedQUAC's tiny overlay and audio-first delivery keep the interactions ambient, preserving focus and screen space in tool-dense editors. Besides, screenshot-based prompts and a brief running memory make it context-aware, while referencing recent edits and allowing partial views makes it iteration-aware.

FeedQUAC also generates persona-driven commentary, distinct voices, and the playful duck avatar make feedback feel playful, and switching personas on demand provides diverse perspectives and tones out-of-the-box. Specifically, adapt insights from prior work on conversational agent persona design and linguistic cues [7, 10, 28], FeedQUAC includes eight ready-made personas spanning tone, focus, and humanness: *Mentor* (constructive), *Cheerleader* (supportive), *Critic* (direct/grumpy), *Designer* (aesthetic focus), *Analyst* (practical constraints), *CEO* (vision/impact), *Friend* (sassy/conversational), and *Robot/No Persona* (neutral/robotic). Each persona is defined by a tailored prompt, a distinct text-to-speech (TTS) voice, and a corresponding duck avatar. Switching personas changes the style and lens of the feedback, helping designers quickly converge on what suits their current need in the moment.

FeedQUAC is an Electron-based, cross-platform desktop overlay that screenshots the workspace, pairs it with a selected persona and context memory, sends it to gpt-4-vision-preview [2] for a concise critique, speaks it via ElevenLabs' neural TTS model [1], and shows an inline transcript, with only 5–10 second end-to-end latency.

3 Evaluation and So-What

We ran a 70-minute remote design-probe with eight experienced Fusion 360 users who continued their own in-progress projects using FeedQUAC after a brief demo. During the active design period, the system generated 106 feedback items (13.25 per session), mixed between auto-cycles and hotkeys (61 timer-triggered, 45 manual). All eight personas were used; Critic, CEO, Mentor, and Designer were most popular, with most participants sampling multiple personas early and then settling on one that fit the task. For user experiences, participants reported that FeedQUAC was easy to use and always available (8/8), worthwhile (7/8), and enjoyable (8/8). They described it as lower-stakes compared to human feedback (perceived as low-stakes: AI 8/8 vs. human 3/8), conducive to making progress (8/8), and supportive in providing validation and confidence (6/8). Many rated it useful to very useful (5/8), and some noted that it exceeded what they could achieve alone (4/8).

Ambient Interaction: Supporting design focus via subtle cues. Participants consistently described FeedQUAC as quick, lightweight, and unobtrusive in an already dense CAD UI: "*[the FeedQUAC window] is minimal. I don't want my screen to be cluttered, [since] I already have a lot of tools in [it].*" (P2) All 8/8 rated hotkeys as ambient; 6/8 did not find the voice disruptive; 7/8 found floating emojis non-disruptive. Designers liked letting a 3-minute auto-cycle run in the background and using hotkeys for situational check-ins. The trade-off is that background feedback can be missed during heads-down work. For example, P4 was deeply immersed in their design work and even missed an instance of feedback.

Towards a future of ambient creativity support tools (CSTs). FeedQUAC demonstrates how quick, ambient, persona-driven commentary can fit naturally into professional CAD workflows. By balancing the attention it demands, FeedQUAC offers low-stakes, diverse, and useful feedback that boosts confidence and sparks ideas with minimal effort. We therefore argue for ambient CSTs as a distinct design value: keep the companion always available yet unobtrusive, respect attention, and offer help at the periphery. Future researchers consider ambient interaction and minimal tool presence

not just as UI/UX improvement, but as essential design values and evaluation metrics when building CSTs. Future evaluations of ambient CSTs should look beyond disturbance, assessing effort–reward trade-offs, mental demand, alignment with creative rhythms, and perceptions of control, transparency, and collaboration.

References

[1] 2025. AI Voice Generator & Text to Speech. https://elevenlabs.io

[2] 2025. OpenAI Platform. https://platform.openai.com

[3] Heidi L. Andrade, Joanna Hefferen, and Maria E. Palma. 2019. Formative Assessment in the Arts. In *Handbook of Formative Assessment in the Disciplines*. Routledge. Num Pages: 20.

[4] Tyler Angert, Miroslav Suzara, Jenny Han, Christopher Pondoc, and Hariharan Subramonyam. 2023. Spellburst: A Node-based Interface for Exploratory Creative Coding with Natural Language Prompts. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23)*. Association for Computing Machinery, New York, NY, USA, Article 100, 22 pages. https://doi.org/10.1145/3586183.3606719

[5] Frederik Anseel, Filip Lievens, and Eveline Schollaert. 2009. Reflection as a strategy to enhance task performance after feedback. *Organizational Behavior and Human Decision Processes* 110, 1 (2009), 23–35.

[6] Ian Arawjo, Chelse Swoopes, Priyan Vaithilingam, Martin Wattenberg, and Elena L. Glassman. 2024. ChainForge: A Visual Toolkit for Prompt Engineering and LLM Hypothesis Testing. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 304, 18 pages. https://doi.org/10.1145/3613904.3642016

[7] Ananya Bhattacharjee, Jina Suh, Mahsa Ershadi, Shamsi T. Iqbal, Andrew D. Wilson, and Javier Hernandez. 2024. Understanding Communication Preferences of Information Workers in Engagement with Text-Based Conversational Agents. https://doi.org/10.48550/arXiv.2410.20468 arXiv:2410.20468 [cs].

[8] Tuhin Chakrabarty, Vishakh Padmakumar, Faeze Brahman, and Smaranda Muresan. 2024. Creativity Support in the Age of Large Language Models: An Empirical Study Involving Professional Writers. In *Proceedings of the 16th Conference on Creativity & Cognition* (Chicago, IL, USA) *(C&C '24)*. Association for Computing Machinery, New York, NY, USA, 132–155. https://doi.org/10.1145/3635636.3656201

[9] Ruijia Cheng, Ziwen Zeng, Maysnow Liu, and Steven Dow. 2020. Critique me: exploring how creators publicly request feedback in an online critique community. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW2 (2020), 1–24.

[10] Yoonseo Choi, Eun Jeong Kang, Seulgi Choi, Min Kyung Lee, and Juho Kim. 2025. Proxona: Supporting Creators' Sensemaking and Ideation with LLM-Powered Audience Personas. https://doi.org/10.48550/arXiv.2408.10937 arXiv:2408.10937.

[11] Jane L. E*, Yu-Chun Grace Yen*, Isabelle Yan Pan, Grace Lin, Mingyi Li, Hyoungwook Jin, Mengyi Chen, Haijun Xia, and Steven P. Dow. 2024. When to Give Feedback: Exploring Tradeoffs in the Timing of Design Feedback. In *Proceedings of the 16th Annual ACM Conference on Creativity and Cognition*. https://doi.org/10.1145/3635636.3656183

[12] Linda Flower and John R. Hayes. 1981. A Cognitive Process Theory of Writing. *College Composition and Communication* 32, 4 (1981), 365–387. https://eric.ed.gov/?id=EJ256235

[13] Katy Ilonka Gero, Tao Long, and Lydia B Chilton. 2023. Social Dynamics of AI Support in Creative Writing. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 245, 15 pages. https://doi.org/10.1145/3544548.3580782

[14] Emily Jensen, Sriram Sankaranarayanan, and Bradley Hayes. 2024. Automated Assessment and Adaptive Multimodal Formative Feedback Improves Psychomotor Skills Training Outcomes in Quadrotor Teleoperation. In *Proceedings of the 12th International Conference on Human-Agent Interaction* (Swansea, United Kingdom) *(HAI '24)*. Association for Computing Machinery, New York, NY, USA, 185–194. https://doi.org/10.1145/3687272.3688322

[15] Sneha R Krishna Kumaran. 2022. *Setting, tracking and visualizing goals to seek feedback on in-progress creative work*. Ph. D. Dissertation. University of Illinois at Urbana-Champaign.

[16] Chinmay E Kulkarni, Michael S Bernstein, and Scott R Klemmer. 2015. PeerStudio: rapid peer feedback emphasizes revision and improves performance. In *Proceedings of the second (2015) ACM conference on learning@ scale*. 75–84.

[17] James C Lester, Sharolyn A Converse, Susan E Kahler, S Todd Barlow, Brian A Stone, and Ravinder S Bhogal. 1997. The persona effect: affective impact of animated pedagogical agents. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 359–366.

[18] Jiale Li, Jiayang Li, Jiahao Chen, Yifan Li, Shijie Wang, Hugo Zhou, Minjun Ye, and Yunsheng Su. 2024. Evolving agents: Interactive simulation of dynamic and diverse human personalities. *arXiv preprint arXiv:2404.02718* (2024).

[19] Vivian Liu, Tao Long, Nathan Raw, and Lydia Chilton. 2023. Generative Disco: Text-to-Video Generation for Music Visualization. arXiv:2304.08551 (Apr 2023). https://doi.org/10.48550/arXiv.2304.08551 arXiv:2304.08551 [cs].

[20] Vivian Liu, Jo Vermeulen, George Fitzmaurice, and Justin Matejka. 2023. 3DALL-E: Integrating Text-to-Image AI in 3D Design Workflows. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference* (Pittsburgh, PA, USA) *(DIS '23)*. Association for Computing Machinery, New York, NY, USA, 1955–1977. https://doi.org/10.1145/3563657.3596098

[21] Tao Long and Lydia B. Chilton. 2023. Challenges and Opportunities for the Design of Smart Speakers. In *AAAI Workshop on User-Centric Artificial Intelligence for Assistance in At-Home Tasks*. Washington, DC. https://doi.org/10.48550/arXiv.2306.05741 arXiv:2306.05741 [cs].

[22] Tao Long, Katy Ilonka Gero, and Lydia B Chilton. 2024. Not Just Novelty: A Longitudinal Study on Utility and Customization of an AI Workflow. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 782–803. https://doi.org/10.1145/3643834.3661587

[23] Tao Long and Weirui Peng. 2023. PortfolioMentor: Multimodal Generative AI Companion for Learning and Crafting Interactive Digital Art Portfolios. https://doi.org/10.48550/arXiv.2311.14091 arXiv:2311.14091 [cs].

[24] Tao Long, Dorothy Zhang, Grace Li, Batool Taraif, Samia Menon, Kynnedy Simone Smith, Sitong Wang, Katy Ilonka Gero, and Lydia B. Chilton. 2023. Tweetorial Hooks: Generative AI Tools to Motivate Science on Social Media. In *Proceedings of the 14th Conference on Computational Creativity (ICCC '23)*. Association for Computational Creativity, Ontario, Canada, 44–53.

[25] Piotr Mirowski, Kory W. Mathewson, Jaylen Pittman, and Richard Evans. 2023. Co-Writing Screenplays and Theatre Scripts with Language Models: Evaluation by Industry Professionals. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 355, 34 pages. https://doi.org/10.1145/3544548.3581225

[26] Seyed Parsa Neshaei, Roman Rietsche, Xiaotian Su, and Thiemo Wambsganss. 2024. Enhancing Peer Review with AI-Powered Suggestion Generation Assistance: Investigating the Design Dynamics. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (Greenville, SC, USA) *(IUI '24)*. Association for Computing Machinery, New York, NY, USA, 88–102. https://doi.org/10.1145/3640543.3645169

[27] Jeongeon Park and DaEun Choi. 2023. AudiLens: Configurable LLM-Generated Audiences for Public Speech Practice. In *Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) *(UIST '23 Adjunct)*. Association for Computing Machinery, New York, NY, USA, Article 122, 3 pages. https://doi.org/10.1145/3586182.3625114

[28] Jeongeon Park, Bryan Min, Kihoon Son, Jean Y. Song, Xiaojuan Ma, and Juho Kim. 2025. ChoiceMates: Supporting Unfamiliar Online Decision-Making with Multi-Agent Conversational Interactions. https://doi.org/10.48550/arXiv.2310.01331 arXiv:2310.01331 [cs].

[29] Alisha Pradhan and Amanda Lazar. 2021. Hey Google, do you have a personality? Designing personality and personas for conversational agents. In *Proceedings of the 3rd Conference on Conversational User Interfaces*. 1–4.

[30] Kathrin Seßler, Tao Xiang, Lukas Bogenrieder, and Enkelejda Kasneci. 2023. Peer: Empowering writing with large language models. In *European Conference on Technology Enhanced Learning*. Springer, 755–761.

[31] Lu Sun, Aaron Chan, Yun Seo Chang, and Steven P. Dow. 2024. ReviewFlow: Intelligent Scaffolding to Support Academic Peer Reviewing. In *Proceedings of the 29th International Conference on Intelligent User Interfaces* (Greenville, SC, USA) *(IUI '24)*. Association for Computing Machinery, New York, NY, USA, 120–137. https://doi.org/10.1145/3640543.3645159

[32] David Thomas and Andrew Hunt. 2019. *The Pragmatic Programmer: your journey to mastery*. Addison-Wesley Professional.

[33] Sitong Wang, Samia Menon, Tao Long, Keren Henderson, Dingzeyu Li, Kevin Crowston, Mark Hansen, Jeffrey V. Nickerson, and Lydia B. Chilton. 2024. ReelFramer: Human-AI Co-Creation for News-to-Video Translation. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–20.

[34] Mark Weiser and John Seely Brown. 1996. The Coming Age of Calm Technology. https://calmtech.com/papers/coming-age-calm-technology.html

[35] David C. Wynn and Anja M. Maier. 2022. Feedback systems in the design and development process. *Research in Engineering Design* 33, 3 (July 2022), 273–306. https://doi.org/10.1007/s00163-022-00386-z

[36] Kevin Yang, Yuandong Tian, Nanyun Peng, and Dan Klein. 2022. Re3: Generating Longer Stories With Recursive Reprompting and Revision. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing* (Abu Dhabi, United Arab Emirates, 2022-12), Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (Eds.). Association for Computational Linguistics, 4393–4479. https://doi.org/10.18653/v1/2022.emnlp-main.296

[37] Yu-Chun Grace Yen, Steven P Dow, Elizabeth Gerber, and Brian P Bailey. 2017. Listen to others, listen to yourself: Combining feedback review and reflection to improve iterative design. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 158–170.

[38] J.D. Zamfirescu-Pereira, Richmond Y. Wong, Bjoern Hartmann, and Qian Yang. 2023. Why Johnny Can't Prompt: How Non-AI Experts Try (and Fail) to Design LLM Prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. Association for Computing Machinery, New York, NY, USA, Article 437, 21 pages. https://doi.org/10.1145/3544548.3581388