

Frontend Diffusion: Empowering Self-Representation of Researchers and Designers with Multi-agent System

Zijian Ding¹, Qinshi Zhang², Mohan Chi³, Ziyi Wang¹

¹University of Maryland, College Park, ²University of California, San Diego, ³Purdue University
¹{ding, zoewang}@umd.edu, ²qiz065@ucsd.edu, ³chi70@purdue.edu

Abstract—With the continuous development of generative AI’s logical reasoning abilities, AI’s growing code-generation potential poses challenges for both technical and creative professionals. But how can these advances be directed toward empowering junior researchers and designers who often require additional help to build and express their professional and personal identities? We introduce Frontend Diffusion, a multi-agent coding system transforming user-drawn layouts and textual prompts into refined website code, thereby supporting self-representation goals. A user study with 13 junior researchers and designers shows AI as a human capability enhancer rather than a replacement, and highlights the importance of bidirectional human-AI alignment. We then discuss future work such as leveraging AI for career development and fostering bidirectional human-AI alignment of multi-agent systems.

Index Terms—Code Generation, Generative AI

I. Introduction

With the advancements in generative models’ logical reasoning capabilities [1], [2], technical productivity has been enhanced across various domains. In programming, generative models have increased the resolution rate of pull requests on GitHub from under 2% in 2023 to over 75% by mid-2025¹. The performance of AI in coding has not only spurred companies’ interest in adopting agentic workflows within software development but have also raised concerns among human professionals about the potential for AI to replace their roles. This anxiety has been particularly pronounced in creative industries, where screenwriters face competition from text-generation models [3] and illustrators contend with text-to-image models [4]–[6]. Such realities motivate us to explore a future where humans and AI coexist synergistically. Instead of automating humans out of the creative process, AI can—and should—act as a catalyst for enhancing self-expression, facilitating more effective personal and professional presentation, and freeing up time for higher-level intellectual pursuits.

This vision holds particular relevance for emerging scholars, such as early-career PhD and master’s students, who stand at a critical juncture in their academic journeys. For them, self-presentation is intricately tied to personal growth, skill development, and the formation of academic

identities. Despite their wealth of new ideas and scholarly potential, junior researchers often face challenges in building their professional brands and achieving online visibility. Creating a professional website—a platform for presenting one’s research and career aspirations—can be a daunting endeavor, often hindered by technical challenges and time constraints. Multi-agent workflows from design to development and reflection have the potential to address these challenges, not by replacing the researcher’s creative authority, but by enabling them to focus on more meaningful pursuits, such as crafting their scholarly narratives and refining their research agendas, ultimately emphasizing their unique academic identities and depth.

To support self-representation through agentic workflows, we developed Frontend Diffusion², an open-source, end-to-end multi-agent system transforming user-drawn layouts and thematic prompts into iteratively refined website code. We evaluated this system through a user study involving 13 participants with diverse technical backgrounds. The findings revealed that the AI tool functions not merely as a code generator but as a collaborative partner. Participants emphasized two major themes: AI as a Human Capability Enhancer, Not a Replacement, and Bidirectional Human-AI Alignment.

II. Related Work

The integration of code data in language model training [7] has made code generation a core capability of generative systems, with front-end code generation demonstrating particularly strong performance [8], [9]. Researchers have developed coding agents like OpenHands [10] and specialized techniques for UI code generation from screenshots [11], [12]. Different from traditional template-based approaches, AI-powered UI generation systems now enable personalized and adaptive interfaces. Tools like FrameKit [13] and PromptInfuser [14] support dynamic UI generation, while DesignAID [15] and Misty [16] provide conceptual inspiration and remixing capabilities for designers. AI can also offer real-time design feedback [17], [18] and automated heuristic evaluations [19], [20].

¹<https://www.swebench.com> (verified)

²Available at: <https://github.com/Carolzhangzz/frontenddiffusion>

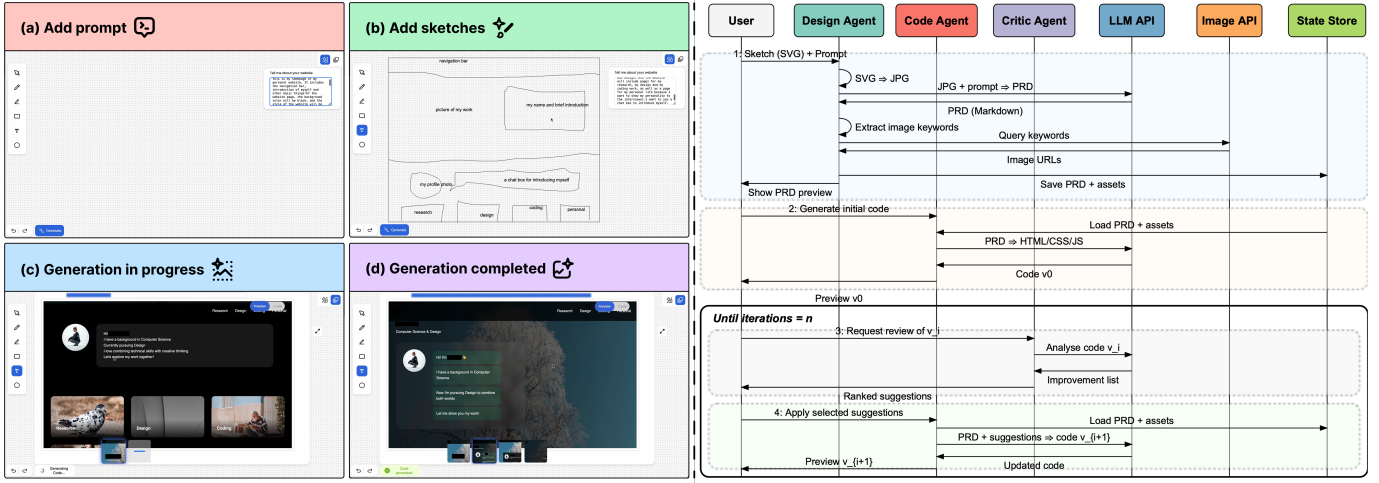


Fig. 1. Left: Website generation workflow: (a) user inputs prompt; (b) user draws layout of the website in sketch; (c) the system generates the first website; (d) the system completes generations four website iterations. Right: End-to-end multi-agent workflow of Frontend Diffusion. The user supplies a sketch and a textual prompt. (1) Design Agent converts the sketch to a Product-Requirements-Document (PRD) and retrieves illustrative images; (2) Code Agent translates the PRD and assets into runnable HTML/CSS/JavaScript; (3) Critic Agent evaluates the code, proposes improvements, and triggers regeneration until the maximum number of iterations n is reached.

Frontend development is beyond code generation and more like creative workflows. In this creative process, AI reduces the time needed for high-fidelity prototyping [21], lowers experimentation barriers [22], provides inspiration when users have broad but underspecified concepts [23], and enables parallel prototyping of diverse design solutions [24]. Yet, much of the existing literature emphasizes technical prowess over the ways AI can empower particular user groups to shape and communicate their own identities. Building on previous work, our work examines how a multi-agent system can be customized to help junior researchers and designers articulate and curate their professional personas.

III. System Design

We developed Frontend Diffusion, an end-to-end multi-agent workflow (Figure 1) that delegates distinct responsibilities to three cooperating agents:

- **Design Agent (Sketch-to-PRD):** The Design Agent translates the user’s visual sketch (SVG) and prompt into a structured PRD. After a format conversion of the sketch ($\text{SVG} \rightarrow \text{JPG}$) that improves model perception, the agent prompts an LLM to produce a markdown Product Requirements Document (PRD) containing layout, component semantics, and image searching keywords. Keywords in PRD such as [hero(landscape)] or [profile(large)] are extracted and issued to an image searching API³, whose URLs are injected back into the PRD. The PRD containing image links are stored in state for downstream agents.
- **Code Agent (PRD-to-Code):** Using the stored PRD, image links, and the user original prompt, the Code

Agent prompts LLM to generate the initial version of the website (v_0).

- **Critic Agent (Reflection & Refinement):** The Critic Agent performs automated code review and reflection. It analyses the current version (v_i), produces improvement suggestions (layout, accessibility, performance, etc.), and iterate on the site ($v_i \rightarrow v_{i+1}$) based on the suggestions. This agentic loop continues for up to $n=4$ iterations by default, and users can branch between versions via thumbnail previews.

All agents communicate through a shared memory layer and expose JSON-RPC APIs, making the multi-agent system model-agnostic and future-proof against generative models advance. The multi-agent system is based on Claude-3.5-Sonnet (2024-10-22), the state-of-the-art language model for coding as of Jan 2025

IV. Study Design

For evaluating the quality of generated personal websites, besides general objective requirements for websites, subjective factors—such as the content and message the user intends to convey—are more important. Therefore, we used qualitative user interviews to assess both the generated websites and the human-AI co-creation process. The experiment lasted approximately 45 minutes and consisted of three parts: a 5-minute participant onboarding, a 20-30 minute user study, and a 15-20 minute post-study interview. First, participants read and signed the consent form, granting permission for screen recording. Next, the researcher demonstrated how to use it by creating one website page. Participants then created personal website showcasing either professional content, such as research or design portfolios, or personal interests, such as an image

³<https://www.pexels.com>

gallery or reading list. After completing the user study, the researcher conducted a 15-20 minute post-study interview.

A. Demographic Information

In the pre-study survey, we collected demographic information, including participants' age, gender, educational level, experience in web development, design, and research, as well as whether they had a personal website (Y/N). Table I summarizes this information.

TABLE I

Participants' demographic information, including age, gender, educational level, and experience in web development (Dev), design, and research (in years).

PID	Demographic			Experience (yrs)		
	Age	Gender	Edu	Dev	Design	Research
1	25	F	Yr-1 PhD	0	0	3
2	25	F	Yr-1 PhD	0.1	0.2	4
3	23	F	Yr-1 Master	1	3	0.5
4	23	F	Yr-2 Master	0.5	3	3
5	28	F	Yr-3 PhD	1.5	4	5.5
6	31	M	Yr-3 PhD	0.5	1	10
7	29	M	Yr-1 PhD	0	0	4
8	23	F	Yr-1 Master	1	0	1
9	31	F	Yr-4 PhD	0	1	8
10	24	F	Yr-1 PhD	0.5	0	3
11	23	F	Yr-1 Master	3	1	0
12	29	M	Yr-3 PhD	1	0	4
13	27	M	Yr-1 PhD	2	3	3.5

B. Post-Study Interview

Following the completion of the website creation task, we conducted semi-structured interviews with each participant. The interview outline was as follows:

- User Experience
 - Process: Did you find the process of using the tool to generate a personal website smooth? Were there any steps that were particularly time-consuming or challenging?
 - Result: Did the final website meet your expectations? Do you think the tool accurately understood your intent between sketch input and the generated website? Why or why not?
- Improvement and Recommendation
 - Are there any features you would like the tool to add or improve? Did you feel the need for more interaction or guidance from the tool?
 - Would you recommend this tool to your peers? Why or why not?
- Broader Impacts and Perceptions
 - Does this tool increase or decrease your interest in frontend development, or even research?
 - Do you think this tool enhances human abilities or replaces them? Why?

C. Data Analysis

Interview transcripts were analyzed using Braun and Clarke's [25] six-phase thematic analysis framework. First, the data were read repeatedly to gain familiarity and generate initial codes. Codes were then collated into provisional themes, which were reviewed and refined for coherence and distinctiveness. Finally, each theme was defined and illustrated with representative quotations, and analytic memos were kept throughout to ensure transparency and rigor.

V. Results

We identified two major themes in our results: Theme 1 - AI as A Human Capability Enhancer and Bidirectional Human-AI Alignment.

Theme 1: When supporting the self-representation of junior researchers and designers, AI has the potential to enhance human capabilities by alleviating the cognitive and technical burden of repetitive tasks. For example, P3 addressed: "Without this tool, maybe very few people would have their own website; but once it's available, everyone might want one as a digital business card... With such tools, everyone can become a designer, frontend developer, or product manager, thereby driving the creation of more interesting products." This enables users to dedicate more time to higher-level ideation and reflection: "it's actually picking up intention... it's not going to try to second guess me" (P9), while P10 emphasized the irreplaceable role of humans as requirement providers, noting that "even the user doesn't have a standard perspective. The users are like the client, and this tool is like the contractor." Detailed supporting quotes for each subtheme can be found in Table II.

Theme 2 (Table III): We decompose this into subthemes based on the initiator of the alignment process as below:

- 1) Ai-initiated alignment
 - Onboarding support for new users.
 - Prompt guidance and refinement.
- 2) User-initiated alignment
 - Fine-grained user control over details.
 - Harnessing AI's unexpected creative sparks.
 - Communicating dynamic behaviors.

Notably, participants (P2, P6, P8, P10, P13) emphasized the need for fine-grained user control over details. Although the multi-agent system is developed to interpret high-level sketches and prompts for end-to-end website generation, users expressed a need for more granular control over intent expression. For example, P2 wanted to provide more precise parameters or specific regions. Similarly, P6 advocated for hierarchical prompting for different page sections: "I wonder if you can have a general prompt and then prompt also in the specific section with a little tag." More detailed participant quotes for each subtheme are provided in Table III.

TABLE II
Subthemes and representative user quotes for Theme 1 – AI as a Human Capability Enhancer.

Theme 1 – AI as a Human Capability Enhancer	
Subtheme (Description)	Participant Quotes
Website Generation Beyond Templates (AI generates personalized, targeted code)	<p>“With templates, you spend time reviewing different options, and even after selection, you must understand their code structure. This tool generates clean, targeted output without unnecessary elements.” (P2)</p> <p>“For those without UI design experience who struggle with platforms like Wix - which has numerous options requiring drag-and-drop interactions and color decisions - this tool streamlines the process by handling color schemes automatically. When I needed a floating music player, instead of searching through component libraries, the tool generated one directly with multiple variations.” (P7)</p>
Supporting Ideation, Team Brainstorming & Self-Reflection (AI facilitates collaborative prototyping and self-examination via rapid sketching and text input)	<p>“It provides accessibility and facilitates communication. In group projects, team members can individually contribute through text or sketches, collectively evaluate outcomes, and iteratively build upon selected concepts - making it particularly suitable for brainstorming sessions.” (P8)</p> <p>“From my perspective, a personal blog’s essential function lies in its capacity to aggregate and synthesize existing knowledge. This utility isn’t necessarily directed toward public consumption - it may serve primarily as a tool for self-examination, illuminating the processes of my internal thought world... The act of presenting to oneself necessarily precedes external presentation to the broader audience.” (P5)</p>
Empowering User Interface Development (Accelerates work for experts and reduces frustration for novices by bridging UI design gaps)	<p>“To achieve more specific outcomes, some UI/UX knowledge is essential. For instance, when implementing complex animations like a rotating solar system, rather than using general descriptions, I would specify technical requirements like ‘I want one sphere rotating around another sphere.’ The AI tool can then better interpret these technical terms like ‘rotate’ to generate the appropriate code.” (P8)</p> <p>“Frontend development can be very frustrating for some people. It can be very painful. People might literally pull their hands up. So this can certainly be a great stress elimination... I prefer back-end programming. I like writing algorithms... I’m not a big fan of frontend development. And when I was learning front, when I was trying to write frontend code for the first few times, I found it very frustrating.” (P13)</p>
Potential Educational Applications (Acts as a mentor with tutorials/examples, guiding learners toward professional standards)	<p>“By studying the code it generates, one can understand how the page is constructed... When I cannot envision certain animations, it provides exemplary solutions that I can learn from regarding how such effects are created.” (P11)</p> <p>“Because I want to see what recruiters and consultants want... if there’s a way that the AI can advise me, not just with the visual, but just with the overall look and feel and content of it. Let’s say you’re a business looking at my UX research portfolio or description of my career, then is this an OK page for that kind of thing?” (P6)</p>
Potential Research Applications (Expedites prototyping and supports accessible code for specialized groups)	<p>“My users are blind users, so I need the frontend to be sufficiently simple for screen readers like Voiceover to work effectively. If it can generate HTML that is both simple and well-structured, it would enhance my productivity.” (P2)</p> <p>“It proved immensely helpful when we needed to create an audio recording experiment website. Previously, we spent extensive time trying to build it with Wix... In 2020, it took me an exceptionally long time to set it up. This tool makes it much simpler - just designing the interface saved me substantial time, especially for components like the music bar and audio elements.” (P7)</p>

VI. Discussion

A. AI as a Human Capability Enhancer: AI Career Advising and Planning

The results demonstrate that generative AI systems can actively enhance users’ self-presentation and career planning by streamlining technical burdens and simulating industry perspectives. P6 explicitly noted that the AI tool could role-play recruiter and consultant viewpoints—offering concrete guidance on overall look, feel, and content that aligns with professional expectations. By providing a low-cost “try-on” environment for career scenarios, AI reduces both cognitive load and technical barriers, thereby

empowering users to experiment with diverse pathways.

B. Bidirectional Human–AI Alignment: Alignment in Hierarchical Multi-agent Systems

As AI reasoning and planning capabilities advance and multi-agent systems scale, they can decompose complex human objectives into extended task sequences and achieve more significant outcomes. In our interviews, participants underscored the importance of onboarding support for novice users, guidance and refinement of prompts, and fine-grained control over agent behavior. Enabling precise user intervention—such as hierarchical prompting [26], where a general instruction is augmented

TABLE III
Subthemes corresponding to user quotes for Theme 2 – Bidirectional Human–AI Alignment.

Theme 2 – Bidirectional Human–AI Alignment	
Subtheme (Description)	Participant Quotes
AI-initiated: Onboarding Support for New Users (Example sites, demo videos, and tutorials to build proper mental models)	<p>“I do think another thing that was challenging was coming up with a certain format for the website, just because I don’t have one yet and I don’t really know what a website, or what a grad student’s website should entail at this point.” (P1)</p> <p>“Even a tutorial or an example of, like, hey, this is what one person put and this is what - this is all the different images that came out from it. This is an example of what AI can do in this context. I think having that as a reference point for me to look at and say, okay, so the AI can develop - can be this creative even when given this.” (P12)</p>
AI-initiated: Prompt Guidance & Refinement (Interactive follow-up questions and iterative dialogue instead of one-off outputs)	<p>“The little bit of challenge comes with figuring out how to prompt it, because sometimes you need to have some prior experience to know what to expect to minimize the number of iterations. So I didn’t know whether I should use certain key terminology, like cards or models or something like that.” (P6)</p> <p>“I definitely think that it should ask maybe follow-up questions when generating the response...Follow up questions to maybe prompt the user to specify what they want.” (P1)</p>
User-initiated: Fine-Grained Control Over Details (Detailed element specs, hierarchical prompts, and previews e.g. colors, images.)	<p>“I wonder if you can have a general prompt and then prompt also in the specific section with a little tag.” (P6)</p> <p>“Some users may be hesitant to share their materials directly with AI systems. They might prefer generating just the structure, then populating it with their information locally... You could have a large model generate an overall template that’s sophisticated and well-structured, clearly indicating where user information should go, then use a local small model to input the personal content without uploading sensitive materials to servers.” (P2)</p>
User-initiated: Harnessing AI’s Unexpected Creative Sparks (Captures serendipitous ideas and enables remix/editing across versions)	<p>“The AI’s knowledge base may exceed humans’, potentially offering unexpected ideas...It introduces dynamic effects that I hadn’t considered when creating static sketches.” (P11)</p> <p>“With three or four versions, I might appreciate certain design elements from each. How could we combine them? Perhaps through editing capabilities, comments, or drag-and-drop functionality?” (P4)</p>
User-initiated: Communicating Dynamic Behaviors (Challenges expressing interactions via static sketches/text; needs flowcharts or state diagrams)	<p>“While my color instructions were minimal, the system comprehended that I wanted vibrant colors and a clean design. Although my sketch was static, it understood which elements should be dynamic, such as the floating music player that follows page scrolling - an option I hadn’t explicitly requested but found valuable.” (P7)</p> <p>“The underlines here aren’t actual content but rather elements meant to be toggled - expressing such toggle requirements in sketches poses a challenge.” (P12)</p>

by section-specific tags—requires a transparent multi-agent architecture: in our system, each agent’s role (Design, Code, or Critic) should be explicitly defined and visible to users. Moreover, the interface should allow users to address individual agents directly—for example, requesting additional design variations from the Design Agent or further review iterations from the Critic Agent—thereby facilitating targeted interventions and reinforcing bidirectional alignment by interacting with each agent. A promising avenue is to explore coordination both between humans and agents and among agents themselves [27].

C. Limitation and Future Work

There were several limitations in the current study that warrant further investigation. First, the user sample was relatively narrow with 13 junior researchers and designers. This sample did not encompass a broader demographic, such as participants from non-academic backgrounds. Second, this study leveraged publicly available online image repositories when generating website content to

avoid direct use of participants’ personal images. This approach resulted in the selection of images that participants perceived as irrelevant or misaligned with their personal or professional identity, potentially impacting the overall user experience. Future work could not only expand the sample size and enhance population diversity but also pursue the following three research directions:

- Studying how AI systems could augment—rather than replace—human capabilities over extended periods, for example in supporting career planning and skill development;
- Increasing the transparency of multi-agent systems to enable users to employ hierarchical prompting strategies to each individual agent;
- Developing localized solutions—such as open-weight models—to allow sensitive personal data to be processed and stored on users’ local devices.

In conclusion, we eagerly look forward to further research exploring the real-world impact of multi-agent systems.

References

- [1] D. Guo, D. Yang, H. Zhang, J. Song, R. Zhang, R. Xu, Q. Zhu, S. Ma, P. Wang, X. Bi et al., “Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning,” arXiv preprint arXiv:2501.12948, 2025.
- [2] OpenAI, “OpenAI o3-mini System Card,” Dec. 2025. [Online]. Available: <https://cdn.openai.com/o3-mini-system-card.pdf>
- [3] P. Mirowski, K. W. Mathewson, J. Pittman, and R. Evans, “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, ser. CHI ’23. New York, NY, USA: Association for Computing Machinery, Apr. 2023, pp. 1–34. [Online]. Available: <https://dl.acm.org/doi/10.1145/3544548.3581225>
- [4] H.-K. Ko, G. Park, H. Jeon, J. Jo, J. Kim, and J. Seo, “Large-scale Text-to-Image Generation Models for Visual Artists’ Creative Works,” in Proceedings of the 28th International Conference on Intelligent User Interfaces, Mar. 2023, pp. 919–933, arXiv:2210.08477 [cs]. [Online]. Available: <http://arxiv.org/abs/2210.08477>
- [5] H. H. Jiang, L. Brown, J. Cheng, M. Khan, A. Gupta, D. Workman, A. Hanna, J. Flowers, and T. Geburu, “AI Art and its Impact on Artists,” in Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society. Montréal QC Canada: ACM, Aug. 2023, pp. 363–374. [Online]. Available: <https://dl.acm.org/doi/10.1145/3600211.3604681>
- [6] R. Kawakami and S. Venkatagiri, “The Impact of Generative AI on Artists,” in Creativity and Cognition. Chicago IL USA: ACM, Jun. 2024, pp. 79–82. [Online]. Available: <https://dl.acm.org/doi/10.1145/3635636.3664263>
- [7] M. Chen, J. Tworek, H. Jun, Q. Yuan, H. P. d. O. Pinto, J. Kaplan, H. Edwards, Y. Burda, N. Joseph, G. Brockman, A. Ray, R. Puri, G. Krueger, M. Petrov, H. Khlaaf, G. Sastry, P. Mishkin, B. Chan, S. Gray, N. Ryder, M. Pavlov, A. Power, L. Kaiser, M. Bavarian, C. Winter, P. Tillet, F. P. Such, D. Cummings, M. Plappert, F. Chantzis, E. Barnes, A. Herbert-Voss, W. H. Guss, A. Nichol, A. Paino, N. Tezak, J. Tang, I. Babuschkin, S. Balaji, S. Jain, W. Saunders, C. Hesse, A. N. Carr, J. Leike, J. Achiam, V. Misra, E. Morikawa, A. Radford, M. Knight, M. Brundage, M. Murati, K. Mayer, P. Welinder, B. McGrew, D. Amodei, S. McCandlish, I. Sutskever, and W. Zaremba, “Evaluating Large Language Models Trained on Code,” Jul. 2021, arXiv:2107.03374 [cs]. [Online]. Available: <http://arxiv.org/abs/2107.03374>
- [8] C. E. Jimenez, J. Yang, A. Wettig, S. Yao, K. Pei, O. Press, and K. R. Narasimhan, “SWE-bench: Can language models resolve real-world github issues?” in The Twelfth International Conference on Learning Representations, 2024. [Online]. Available: <https://openreview.net/forum?id=VTF8yNQm66>
- [9] C. Si, Y. Zhang, Z. Yang, R. Liu, and D. Yang, “Design2Code: How Far Are We From Automating Front-End Engineering?” Mar. 2024, arXiv:2403.03163 [cs]. [Online]. Available: <http://arxiv.org/abs/2403.03163>
- [10] X. Wang, B. Li, Y. Song, F. F. Xu, X. Tang, M. Zhuge, J. Pan, Y. Song, B. Li, J. Singh, H. H. Tran, F. Li, R. Ma, M. Zheng, B. Qian, Y. Shao, N. Muennighoff, Y. Zhang, B. Hui, J. Lin, R. Brennan, H. Peng, H. Ji, and G. Neubig, “OpenHands: An Open Platform for AI Software Developers as Generalist Agents,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.16741>
- [11] Y. Wan, C. Wang, Y. Dong, W. Wang, S. Li, Y. Huo, and M. R. Lyu, “Automatically Generating UI Code from Screenshot: A Divide-and-Conquer-Based Approach,” Oct. 2024, arXiv:2406.16386 [cs]. [Online]. Available: <http://arxiv.org/abs/2406.16386>
- [12] J. Wu, E. Schoop, A. Leung, T. Barik, J. P. Bigham, and J. Nichols, “UICoder: Finetuning Large Language Models to Generate User Interface Code through Automated Feedback,” Jun. 2024, arXiv:2406.07739 [cs]. [Online]. Available: <http://arxiv.org/abs/2406.07739>
- [13] J. Wu, K. Todi, J. Chan, B. A. Myers, and B. Lafreniere, “FrameKit: A Tool for Authoring Adaptive UIs Using Keyframes,” in Proceedings of the 29th International Conference on Intelligent User Interfaces. Greenville SC USA: ACM, Mar. 2024, pp. 660–674.
- [14] S. Petridis, M. Terry, and C. J. Cai, “PromptInfuser: How Tightly Coupling AI and UI Design Impacts Designers’ Workflows,” in Designing Interactive Systems Conference. IT University of Copenhagen Denmark: ACM, Jul. 2024, pp. 743–756.
- [15] A. Cai, S. R. Rick, J. L. Heyman, Y. Zhang, A. Filipowicz, M. Hong, M. Klenk, and T. Malone, “DesignAID: Using Generative AI and Semantic Diversity for Design Inspiration,” in Proceedings of The ACM Collective Intelligence Conference. Delft Netherlands: ACM, Nov. 2023, pp. 1–11. [Online]. Available: <https://dl.acm.org/doi/10.1145/3582269.3615596>
- [16] Y. Lu, A. Leung, A. Swearngin, J. Nichols, and T. Barik, “Misty: UI Prototyping Through Interactive Conceptual Blending,” Sep. 2024.
- [17] P. Duan, J. Warner, and B. Hartmann, “Towards Generating UI Design Feedback with LLMs,” in Adjunct Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology. San Francisco CA USA: ACM, Oct. 2023, pp. 1–3. [Online]. Available: <https://dl.acm.org/doi/10.1145/3586182.3615810>
- [18] P. Duan, J. Warner, Y. Li, and B. Hartmann, “Generating Automatic Feedback on UI Mockups with Large Language Models,” in Proceedings of the CHI Conference on Human Factors in Computing Systems. Honolulu HI USA: ACM, May 2024, pp. 1–20.
- [19] J. Wu, Y.-H. Peng, X. Y. A. Li, A. Swearngin, J. P. Bigham, and J. Nichols, “UIClip: A Data-driven Model for Assessing User Interface Design,” in Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology. Pittsburgh PA USA: ACM, Oct. 2024, pp. 1–16.
- [20] Y. Lu, T. Kneare, S. Dutta, J. Blass, C. Kliman-Silver, and F. Bentley, “AI Is Not Enough: A Hybrid Technical Approach to AI Adoption in UI Linting With Heuristics,” in Extended Abstracts of the CHI Conference on Human Factors in Computing Systems. Honolulu HI USA: ACM, May 2024, pp. 1–7. [Online]. Available: <https://dl.acm.org/doi/10.1145/3613905.3637135>
- [21] K. M. Edwards, B. Man, and F. Ahmed, “Sketch2Prototype: rapid conceptual design exploration and prototyping with generative AI,” Proceedings of the Design Society, vol. 4, pp. 1989–1998, May 2024.
- [22] S. Palani and G. Ramos, “Evolving Roles and Workflows of Creative Practitioners in the Age of Generative AI,” in Creativity and Cognition. Chicago IL USA: ACM, Jun. 2024, pp. 170–184. [Online]. Available: <https://dl.acm.org/doi/10.1145/3635636.3656190>
- [23] S. R. Rick, G. Giacomelli, H. Wen, R. J. Laubacher, N. Taubenslag, J. L. Heyman, M. S. Knicker, Y. Jeddi, H. Maier, S. Dwyer, P. Ragupathy, and T. W. Malone, “Supermind Ideator: Exploring generative AI to support creative problem-solving,” Nov. 2023.
- [24] S. P. Dow, A. Glassco, J. Kass, M. Schwarz, D. L. Schwartz, and S. R. Klemmer, “Parallel prototyping leads to better design results, more divergence, and increased self-efficacy,” ACM Transactions on Computer-Human Interaction, vol. 17, no. 4, pp. 1–24, Dec. 2010.
- [25] V. Braun and V. Clarke, “Using thematic analysis in psychology,” Qualitative Research in Psychology, vol. 3, no. 2, pp. 77–101, Jan. 2006.
- [26] F. Gmeiner, N. Marquardt, M. Bentley, H. Romat, M. Pahud, D. Brown, A. Roseway, N. Martelaro, K. Holstein, K. Hinckley, and N. Riche, “Intent Tagging: Exploring Micro-Prompting Interactions for Supporting Granular Human-GenAI Co-Creation Workflows,” Feb. 2025, arXiv:2502.18737 [cs]. [Online]. Available: <http://arxiv.org/abs/2502.18737>
- [27] Y. Bai, Z. Ding, S. Wen, X. Chang, and A. Taylor, “From MAS to MARS: Coordination Failures and Reasoning Trade-offs in Hierarchical Multi-Agent Robotic Systems within a Healthcare Scenario,” Aug. 2025, arXiv:2508.04691 [cs]. [Online]. Available: <http://arxiv.org/abs/2508.04691>