From Paper to Practice: Generative AI as a Bridge for Translating Research Findings into Design

DONGHOON SHIN, University of Washington, USA

Despite the wealth of knowledge generated by the human-computer interaction (HCI) community, a significant research-practice gap persists. Practitioners often struggle to apply academic findings to their work due to dense formatting, lack of specific contextualization, and the high cognitive friction of integrating external literature into active design workflows. My doctoral research addresses this rigor-relevance paradox by exploring the use of generative AI as a mediator for translational science in design. Through a series of system-building and empirical studies, I investigate how generative AI can be leveraged to (1) transform static academic papers into accessible, practitioner-friendly artifacts, (2) tailor scholarly insights to individual designers' unique design contexts, and (3) integrate evidence-informed guidance directly into in situ design workspaces. In this proposal, I outline my research carried out so far, including the transformation of academic insights into a prescriptive format, its customized translation pipelines, and AI-powered collaborative plugins, and discuss my future plans for evaluating the longitudinal impact of these systems on real-world design practice.

1 PROBLEM DOMAIN AND THE SPECIFIC PROBLEMS ADDRESSED

The core problem domain of my doctoral research is translational science in design. Although the HCI research community generates valuable design implications intended to influence practice, these insights rarely reach or meaningfully shape everyday design work [15, 25]. Designers consistently find scholarly papers difficult to consume due to dense jargon, lengthy prose, and a lack of immediate applicability to the fast-paced, constraint-driven realities of professional practice [7, 23]. This creates a persistent research-practice gap that limits the societal impact of HCI research and leaves designers without reliable, evidence-informed guidance when they need it most.

Scholarly papers typically present findings and design implications in abstract, jargon-heavy prose that lacks the visual, prescriptive, and inspirational qualities practitioners seek under tight deadlines [21, 25]. Reading academic literature is often described as cognitively taxing, and the style does not easily translate into the quick, actionable takeaways required during ideation and iteration [22, 23]. As a result, even high-value insights remain underutilized, with designers frequently relying on intuition, personal experience, or ad-hoc searches instead of the cumulative knowledge produced by the HCI community [23, 28].

Compounding this issue, creating high-quality practitioner-friendly translational artifacts, such as design cards or short-form videos, is a manual and labor-intensive process. Most HCI researchers lack the time, specialized skills, and resources required, and face limited incentives or clear guidance on what designers actually need [8, 26]. Consequently, translational materials remain scarce, unevenly distributed, and unable to keep pace with the rapid growth of HCI publications, leaving the majority of design implications inaccessible [25].

Moreover, even when translational artifacts are created, they are typically generic and fail to account for a designer's specific project context, including target users, interaction modalities, application domain, design stage, and organizational priorities [22]. This mismatch exacerbates the well-known rigor-relevance paradox [16], making academically rigorous knowledge feel disconnected from the messy realities of real-world design. As a result, designers often perceive these materials as too abstract or misaligned, reducing their willingness to engage with them [7, 22].

Lastly, while designers value evidence-informed support delivered in situ—directly within their workflow as they iterate on strategies, mockups, or prototypes—existing tools fail to retrieve, contextualize, or embed relevant research insights within familiar environments (*e.g.*, Figma, Miro). As a result, designers must bridge this gap themselves,

Author's address: Donghoon Shin, dhoon@uw.edu, University of Washington, Seattle, WA, USA.

introducing substantial cognitive friction [23, 28]. Likewise, gathering timely, diverse feedback from target audiences during early prototyping remains logistically challenging and costly, especially for niche user groups [24]. Consequently, designers risk overlooking critical user needs and misalignments among stakeholder expectations.

Together, these interconnected challenges substantially limit the practical impact of HCI scholarship and hinder designers' ability to leverage evidence-informed, user-centered insights in their daily practice.

2 ORIGINAL KEY IDEA, CORRESPONDING RESEARCH QUESTIONS, AND HYPOTHESIS

The core idea of this dissertation is that generative AI can function as a scalable, customizable, and integrative boundary spanner that transforms static academic papers into dynamic, practitioner-centric resources, while preserving scholarly validity and enabling meaningful human-AI collaboration. Rather than replacing human creativity or expertise, I posit that generative AI can act as an intelligent partner that handles the labor-intensive aspects of translation, contextualization, and integration, allowing designers to focus on higher-order strategic and creative decisions.

By leveraging generative AI models (*e.g.*, LLMs, text-to-image models), my research investigates if end-to-end generative AI pipelines can automatically extract design implications from papers, reformat them into accessible and visually engaging artifacts, tailor content precisely to an individual designer's project context, and embed evidence-informed guidance directly into existing design tools and workflows. Throughout my research, I aim to address the full translational pipeline—from retrieval and summarization to real-time application and audience simulation—while incorporating mechanisms for provenance tracking, hallucination mitigation, and iterative human refinement to maintain trust and accuracy.

More specifically, my dissertation is guided by three core research questions:

RQ1. How can generative AI effectively translate HCI design implications into prescriptive formats without losing validity? This question explores the use of generative AI to create practitioner-friendly outputs such as design cards and short-form videos. It examines how design elements (*e.g.*, actionable titles, concise descriptions, visuals, and evidence summaries) can be designed to best enhance perceived inspirability and generativity, while maintaining scholarly grounding. I hypothesize that AI-generated artifacts will significantly outperform raw paper text baselines on validated scales of inspirability, generativity, actionability, relevance, and validity.

RQ2. To what extent does context-aware customization of translational artifacts improve perceived relevance, validity, and utility for designers? This question investigates how generative AI can incorporate designer-provided context (*e.g.*, target users, interaction modalities, application domain, design stage) to produce tailored insights. It tests whether customized artifacts outperform generic ones even when source papers are not perfectly relevant to the designer's own design context. I hypothesize that such customization will significantly increase perceived relevance, validity, and utility, especially in ambiguous or cross-domain design scenarios.

RQ3. How does generative AI-enabled real-time integration of research insights into design workflows and audience simulation affect designer creativity, cognitive load, and decision quality? This question examines in situ tools (*e.g.*, Figma and Miro plugins, persona-simulated agents) that deliver contextualized insights or simulate audience feedback directly within the designer's canvas. It assesses their impact on iteration speed, evidence-informed decision-making, and overall design outcomes. I hypothesize that well-designed human-AI collaboration patterns—incorporating iterative refinement, provenance signals, and explicit accuracy checks—will preserve scholarly accuracy while substantially improving adoption, trust, and long-term integration of research into everyday design practice.

3 BRIEF OVERVIEW OF RELATED WORK

Prior HCI work has long advocated for translational resources such as design cards to bridge the research-practice gap by distilling scholarly insights into compact, visually engaging formats that facilitate ideation and stakeholder communication [1, 7, 14]. However, these efforts have largely relied on manual curation by researchers or volunteers, which limits scalability, leads to inconsistent coverage across topics, and fails to address the highly individualized contexts that each designer brings to their projects [8, 10]. Early design-card initiatives demonstrated value in communicating human insights and domain knowledge [6], yet they could not scale to match the growing volume of HCI publications or adapt dynamically to specific user needs [22].

Recent advances in large language models and multimodal generative AI have enabled automated summarization of scholarly articles [3, 5], generation of research highlights [20], creation of posters or slides [12, 27], and even production of visuals for public engagement [11]. While these applications show promise in making scientific literature more accessible, they have not yet systematically tackled the unique challenges of translational science in design—specifically the need to preserve academic validity alongside high levels of inspirability, contextual relevance, and direct actionability within professional workflows [21]. Commercial tools for converting PDFs to videos exist but lack researcher-in-the-loop control, hallucination safeguards, and tight integration with design software [2].

Recent works on generative agents has demonstrated the potential of LLM-driven personas to simulate believable human behaviors and social interactions [17, 18], opening new opportunities for audience testing and collaborative feedback [13]. In parallel, AI-supported creativity tools have begun to assist with UI iteration and inspirational search [4, 19]. However, relatively few systems incorporate verified scholarly evidence directly into the designer's canvas or support real-time synthesis of audience perspectives grounded in marketing data [23, 24]. This dissertation builds on and extends these strands by developing end-to-end generative AI systems that address gaps in scalability, contextual relevance, workflow integration, and audience-inclusive feedback for science communication in design.

4 CONTRIBUTION REALIZED/EXPECTED IN THE FIELD OF HCI

The technical contributions of my works include a suite of implemented and evaluated generative AI architectures that operationalize translational science at scale. This work provides open-source pipelines and interaction paradigms for creating practitioner-friendly artifacts, including an automated design-card generator using LLMs and text-to-image models [25], a human-AI collaborative short-form video creation pipeline [9], an LLM-powered context-aware design card customization system [22], a Miro plugin for theorizing design with AI [28], a Figma plugin for real-time research integration during UI prototyping [23], and an audience-driven persona-agent system for poster feedback [24]. By openly sharing the system architectures, prompt-chaining strategies, and integration mechanisms, this research offers reusable technical foundations for future translational AI systems in HCI.

Empirically, my work contributes a nuanced understanding of how designers perceive and engage with AI-translated scholarly knowledge. Through mixed-methods evaluations with professional designers and experts/evaluators, my work demonstrates that generative AI artifacts significantly outperform both raw academic text and non-customized baselines across key translational metrics. Specifically, AI-generated and context-customized materials are consistently rated higher in inspirability, generativity, relevance, actionability, and validity [22, 25]. The studies also provide quantitative evidence that in situ AI-assisted research translation reduces cognitive load, enhances creativity during ideation and prototyping, and improves design decision quality [23, 28], while maintaining scholarly accuracy through built-in validation mechanisms [23, 24].

In addition, my research contributes concrete, actionable design implications and heuristics for building trustworthy AI-driven translational tools. These include best practices for hallucination mitigation and provenance tracking (*e.g.*, interactive visual links that map generated suggestions back to original paper text), strategies for seamless integration into existing design software, and interaction techniques that effectively balance automation with designer agency [22–24]. These guidelines are intended to support both HCI researchers aiming to disseminate their work more effectively and industry practitioners developing evidence-informed design tools.

Collectively, my research advances a methodological shift in addressing the persistent research-practice gap in HCI [7, 15]. By showing that generative AI can contextualize and deliver academic insights directly into a designer's active workflow, this work reduces the translational burden on both researchers and practitioners. It is expected to increase the real-world adoption of HCI scholarship in industry, inform the design of ethical AI tools that augment human creativity, and provide scalable resources that benefit diverse practitioner communities across design domains.

5 RESEARCH CARRIED OUT SO FAR AND PLANNED AHEAD

All studies conducted in this doctoral research follow an iterative human-centered design methodology. This approach begins with preliminary or formative interviews to understand practitioner needs, moves into the implementation of the design support systems, and concludes with rigorous mixed-methods evaluations. These evaluations combine validated self-reported scales, qualitative interviews, technical benchmarks (*e.g.*, latency, accuracy), crowdsourced validation, and expert accuracy checks on generated artifacts to ensure both systemic reliability and user-centered efficacy while mitigating hallucinations.

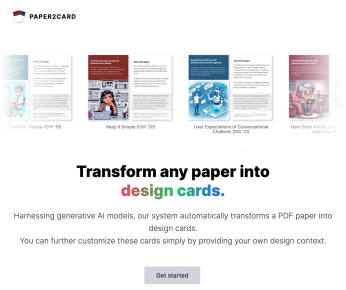

(a) Paper2Card (CHI '24) [25]

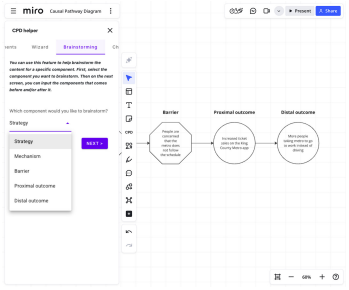

(b) AI-assisted CPD (CHI '24) [28]

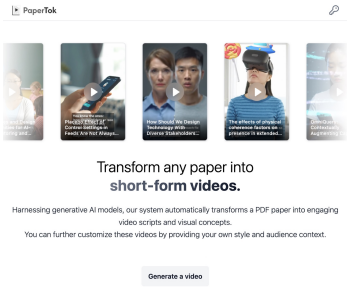

(c) PaperTok (CHI '26) [9]

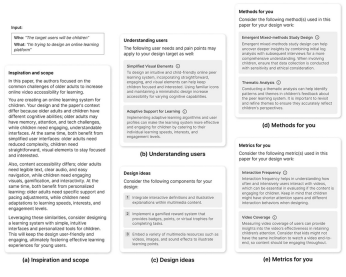

(d) What About My Design Context? (DIS '25) [22]

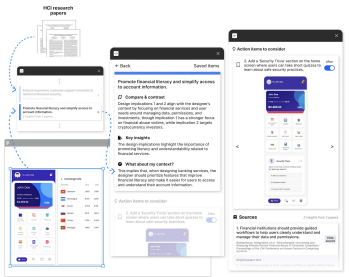

(e) ReFinE (DIS '26) [23]

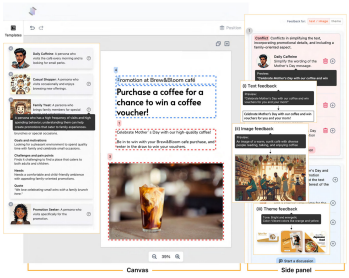

(f) PosterMate (UIST '25) [24]

Fig. 1. Research carried out so far

5.1 Transforming the Format of Scholarly Knowledge

My early research focused on bridging the format mismatch between dense academic prose and fast-paced practitioner needs. Paper2Card [25] introduced an end-to-end pipeline that leverages LLMs and text-to-image models to automatically convert design implications from HCI papers into visually rich, digestible design cards. Informed by preliminary interviews with twelve designers, an evaluation with twenty-one designers and twelve paper authors showed that the AI-generated cards were perceived as significantly more inspiring and generative than the original text, while maintaining comparable levels of core academic values (*i.e.*, validity, generalizability, originality).

Similarly, PaperTok [9] explored short-form video as a modern dissemination medium. It established a human-AI collaborative workflow that transforms academic papers into engaging research videos. Drawing on formative insights from eight science communicators, a mixed-methods study with eighteen researchers and a crowdsourced evaluation with one hundred audience members demonstrated that the generated videos were significantly more engaging and entertaining than existing automated tools, while preserving comparable informational value.

5.2 Contextual Adaptation and Early In Situ Tools

Building on format transformation, subsequent work investigated dynamic contextual adaptation. In What About My Design Context? [22], I developed an LLM-powered pipeline that generates design cards tailored to a designer's specific users, interaction modalities, application domain, and current design stage. Informed by interviews with fifteen designers, I built an end-to-end system to tailor translational artifacts to individual designers' design context. An evaluation with twenty designers found that the context-aware cards were rated significantly higher in relevance, actionability, validity, generativity, and inspiration compared to non-customized versions—even when source papers were only loosely related.

In the AI-assisted Causal Pathway Diagram project [28], we introduced a Miro plugin that supports evidence-informed brainstorming. The tool guides designers in creating causal pathway diagrams with real-time AI suggestions, helping them map strategies, mechanisms, barriers, and outcomes while integrating their own findings with published literature. A study with twenty designers showed substantial reductions in cognitive workload, increased creativity, and improved ability to communicate strategic rationales to stakeholders.

5.3 Advanced In Situ Integration in Prototyping Workflows

Most recently, my research has achieved deeper, in situ integration directly within designers' workspaces. ReFinE [23] is a Figma plugin that automatically infers design context from a selected UI mockup, retrieves relevant HCI papers, synthesizes clusters of design implications, and generates editable HTML action overlays on the canvas. Technical evaluations validated the accuracy of context extraction, retrieval, and visualization. A within-subjects user study with twelve professional designers and students confirmed that ReFinE led to significantly more evidence-informed design edits, drastically reduced cognitive load, and improved both the quality and speed of prototyping.

PosterMate [24] introduced an audience-driven collaborative system that generates multiple persona agents from a marketing brief. These agents provide component-level feedback on poster elements and engage in moderated multi-agent discussions to resolve conflicts and reach evidence-backed consensus. Studies with twelve designers and one hundred online participants showed the system effectively surfaces overlooked audience perspectives and produces coherent, integrated feedback.

5.4 Planned Ahead: Longitudinal Deployment and Dataset Construction

Research planned for the remainder of my doctoral timeline will shift from system building toward ecological validation and foundational infrastructure for the HCI community.

First, I will conduct extensive longitudinal field deployments to evaluate the sustained real-world impact of in situ translational systems such as ReFinE and PosterMate within professional enterprise design teams. Spanning six to twelve months, these studies will combine telemetry logging, ethnographic observation, and experience sampling to measure long-term adoption rates, the influence of AI-translated research on final shipped products, and organizational shifts toward evidence-informed decision-making.

Second, to address the critical bottleneck in evaluating and training future translational design supports, I will construct and open-source a large-scale, high-quality dataset that maps HCI design implications to concrete UI patterns and practitioner contexts. While current generative models rely on generic training data lacking the nuanced rigor of HCI scholarship, this structured, expert-annotated dataset will link specific academic guidelines to visual UI mockups and causal mechanisms. The dataset will provide the community with a robust benchmark for assessing the epistemic accuracy of LLMs in design tasks, reduce the risk of algorithmic hallucination, and establish a lasting infrastructural asset at the intersection of AI, design practice, and translational science.

6 QUESTIONS TO BE ADDRESSED DURING THE COLLOQUIUM

I would particularly value in-depth discussion on the following three questions during the colloquium to help shape the final phase of my dissertation, refine its methodological rigor, and strengthen its broader theoretical and applied impact:

- **Framing the narrative:** I have had clear objectives and a broader direction in mind whenever I ran each project, but since they ultimately result in separate AI tools and pipelines (*e.g.*, Paper2Card, PaperTok, ReFinE, PosterMate), these might come off as an unfocused portfolio of systems. How can I more effectively weave these different artifacts into a unified theoretical framework around "translational science" that feels cohesive from start to finish?
- **Balancing engineering and HCI:** For a body of work so heavily steeped in generative AI pipelines, prompt-chaining, and technical system-building, the engineering efforts might overshadow the human-centered findings. How can I better emphasize and elevate the empirical HCI contributions (*e.g.*, shifts in designer cognition, workflow integration) to an academic committee?
- **Scoping the final evaluation:** My planned next step is a 6- to 12-month longitudinal deployment within an enterprise design team. Given the notoriously complex nature of enterprise environments, confounding variables, and the timeline constraints of a PhD, is this too ambitious? How should I narrow or scope this field study so that it is realistic to complete, while still capturing meaningful metrics of real-world impact?

REFERENCES

[1] Tessa Aarts, Linas K Gabrielaitis, Lianne C De Jong, Renee Noortman, Emma M Van Zoelen, Sophia Kotea, Silvia Cazacu, Lesley L Lock, and Panos Markopoulos. 2020. Design Card Sets: Systematic Literature Survey and Card Sorting Study. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference*. 419–428. https://doi.org/10.1145/3357236.3395516

[2] Hussam Alkaissi and Samy I McFarlane. 2023. Artificial Hallucinations in ChatGPT: Implications in Scientific Writing. *Cureus* 15, 2 (2023). https://doi.org/10.7759/cureus.35179

[3] Tal August, Lucy Lu Wang, Jonathan Bragg, Marti A Hearst, Andrew Head, and Kyle Lo. 2023. Paper Plain: Making Medical Research Papers Approachable to Healthcare Consumers with Natural Language Processing. *ACM Transactions on Computer-Human Interaction* 30, 5 (2023), 1–38. https://doi.org/10.1145/3589955

[4] Sara Bunian, Kai Li, Chaima Jemmali, Casper Harteveld, Yun Fu, and Magy Seif Seif El-Nasr. 2021. VINS: Visual Search for Mobile User Interface Design. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–14. https://doi.org/10.1145/3411764.3445762

[5] Isabel Cachola, Kyle Lo, Arman Cohan, and Daniel S Weld. 2020. TLDR: Extreme Summarization of Scientific Documents. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. 4766–4777. https://doi.org/10.18653/v1/2020.findings-emnlp.428

[6] David Chung and Rung-Huei Liang. 2015. Understanding the Usefulness of Ideation Tools with the Grounding Lenses. In *Proceedings of the Third International Symposium of Chinese CHI*. 13–22. https://doi.org/10.1145/2739999.2740002

[7] Lucas Colusso, Cynthia L Bennett, Gary Hsieh, and Sean A Munson. 2017. Translational Resources: Reducing the Gap Between Academic Research and HCI Practice. In *Proceedings of the 2017 Conference on Designing Interactive Systems*. 957–968. https://doi.org/10.1145/3064663.3064667

[8] Lucas Colusso, Ridley Jones, Sean A Munson, and Gary Hsieh. 2019. A Translational Science Model for HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–13. https://doi.org/10.1145/3290605.3300231

[9] Meziah Ruby Cristobal*, Hyeonjeong Byeon*, Tze-Yu Chen*, Ruoxi Shang*, Donghoon Shin*, Ruican Zhong*, Tony Zhou*, and Gary Hsieh. 2026. PaperTok: Exploring the Use of Generative AI for Creating Short-form Videos for Research Communication. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems*. https://doi.org/10.1145/3772318.3790553

[10] Peter Denning, Jim Horning, David Parnas, and Lauren Weinstein. 2005. Wikipedia Risks. *Commun. ACM* 48, 12 (2005), 152–152. https://doi.org/10.1145/1101779.1101804

[11] Wiebke Finkler and Bienvenido León-Anguiano. 2019. The power of storytelling and video: a visual rhetoric for science communication. *Journal of Science Communication* 18, 5 (2019). https://doi.org/10.22323/2.18050202

[12] Tsu-Jui Fu, William Yang Wang, Daniel McDuff, and Yale Song. 2022. DOC2PPT: Automatic Presentation Slides Generation from Scientific Documents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 634–642. https://doi.org/10.1609/aaai.v36i1.19943

[13] Juhye Ha, Hyeon Jeon, Daeun Han, Jinwook Seo, and Changhoon Oh. 2024. CloChat: Understanding How People Customize, Interact, and Experience Personas in Large Language Models. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–24. https://doi.org/10.1145/3613904.3642472

[14] Gary Hsieh, Brett A Halperin, Evan Schmitz, Yen Nee Chew, and Yuan-Chi Tseng. 2023. What is in the Cards: Exploring Uses, Patterns, and Trends in Design Cards. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–18. https://doi.org/10.1145/3544548.3580712

[15] Donald A Norman. 2010. The Research-Practice Gap: The Need for Translational Developers. *Interactions* 17, 4 (2010), 9–12. https://doi.org/10.1145/1806491.1806494

[16] Jill W Paine and Fred Delmhorst. 2020. A Balance of Rigor and Relevance: Engaged Scholarship in Organizational Change. *The Journal of Applied Behavioral Science* 56, 4 (2020), 437–461. https://doi.org/10.1177/0021886320961218

[17] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–22. https://doi.org/10.1145/3586183.3606763

[18] Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2022. Social Simulacra: Creating Populated Prototypes for Social Computing Systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–18. 10.1145/3526113.3545616

[19] Seokhyeon Park, Yumin Song, Soohyun Lee, Jaeyoung Kim, and Jinwook Seo. 2025. Leveraging Multimodal LLM for Inspirational User Interface Search. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–22. https://doi.org/10.1145/3706598.3714213

[20] Tohida Rehman, Debarshi Kumar Sanyal, Prasenjit Majumder, and Samiran Chattopadhyay. 2022. Named Entity Recognition Based Automatic Generation of Research Highlights. In *Proceedings of the Third Workshop on Scholarly Document Processing*. https://aclanthology.org/2022.sdp-1.18/

[21] Corina Sas, Steve Whittaker, Steven Dow, Jodi Forlizzi, and John Zimmerman. 2014. Generating Implications for Design through Design Research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1971–1980. https://doi.org/10.1145/2556288.2557357

[22] Donghoon Shin, Tze-Yu Chen, Gary Hsieh, and Lucy Lu Wang. 2025. What About My Design Context?: Exploring the Use of Generative AI to Support Customization of Translational Research Artifacts. In *Proceedings of the 2025 ACM Designing Interactive Systems Conference*. 1210–1227. https://doi.org/10.1145/3715336.3735686

[23] Donghoon Shin, Bingcan Guo, Jaewook Lee, Lucy Lu Wang, and Gary Hsieh. 2026. ReFinE: Streamlining UI Mockup Iteration with Research Findings. In *Proceedings of the 2026 ACM Designing Interactive Systems Conference*.

[24] Donghoon Shin, Daniel Lee, Gary Hsieh, and Gromit Yeuk-Yin Chan. 2025. PosterMate: Audience-driven Collaborative Persona Agents for Poster Design. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology*. 1–20. https://doi.org/10.1145/3746059.3747769

[25] Donghoon Shin, Lucy Lu Wang, and Gary Hsieh. 2024. From Paper to Card: Transforming Design Implications with Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–15. https://doi.org/10.1145/3613904.3642266

[26] Spencer Williams, Ridley Jones, Katharina Reinecke, and Gary Hsieh. 2022. An HCI Research Agenda for Online Science Communication. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–22. https://doi.org/10.1145/3555591

[27] Sheng Xu and Xiaojun Wan. 2022. PosterBot: A System for Generating Posters of Scientific Papers with Neural Models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 13233–13235. https://doi.org/10.1609/aaai.v36i11.21738

[28] Ruican Zhong*, Donghoon Shin*, Rosemary Meza, Predrag Klasnja, Lucas Colusso, and Gary Hsieh. 2024. AI-Assisted Causal Pathway Diagram for Human-Centered Design. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–19. https://doi.org/10.1145/3613904.3642179