

Evaluating Visual Prompting Modalities for Generative AI-Assisted UI Design

Original

Evaluating Visual Prompting Modalities for Generative AI-Assisted UI Design / Calò, T., De Russis, L.. - ELETTRONICO. - 15713:(2025), pp. 171-181. (IS-EUD 2025: the 10th International Symposium on End-User Development Munich (DEU) 16-18 June 2025) [10.1007/978-3-031-95452-8_11].

Availability:

This version is available at: 11583/3000511 since: 2025-08-04T15:56:55Z

Publisher:

Springer

Published

DOI:10.1007/978-3-031-95452-8_11

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Springer postprint/Author's Accepted Manuscript

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: http://dx.doi.org/10.1007/978-3-031-95452-8_11

(Article begins on next page)

Evaluating Visual Prompting Modalities for Generative AI-Assisted UI Design

Tommaso Calò^[0000–0002–3200–2348] and Luigi De Russis^[0000–0001–7647–6652]

Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy
{tommaso.calo, luigi.derussis}@polito.it

Abstract. The emergence of Generative AI systems has dramatically transformed design workflows across multiple domains. While text prompting remains the dominant interaction method with these systems, visual prompting—the practice of guiding AI generation through visual or semantic structure—offers designers potentially greater control and expressivity. A critical challenge for design practitioners lies in understanding how different visual prompting modalities integrate into established creative processes and impact design outcomes. This study examines this challenge in the context of UI mockup generation, a domain with well-defined semantic elements and established design methodologies. We evaluate two visual prompting approaches: *free-form visual prompting*, which creates outputs based on hand-drawn sketches, and *semantic-constrained prompting*, which uses predefined visual vocabularies to guide generation. Through experiments with 13 design practitioners, we explore how different prompting modalities impact both designer experience and output quality in the creation of UI artifacts. Results reveal that free-form visual prompting offers superior intuitiveness and expressiveness for ideation, while semantic-constrained prompting produces higher quality and fidelity outputs. Our findings suggest that effective visual prompting strategies should adapt to different stages of the design process, with implications for generative AI applications in design practice. We propose a hybrid approach that leverages the strengths of both modalities throughout the creative workflow, potentially offering design practitioners across domains a more balanced framework for designer-AI collaboration.

Keywords: Visual Guidance, Generative AI, UI Design, User Interfaces, Design Practitioners, Image Generation, Human-AI Collaboration

1 Introduction and Background

User interface designers face daily challenges in creating designs that are effective, usable, and innovative. They often draw inspiration from existing design samples to come up with new ideas [17]. There are two main types of resources that support this process of finding inspiration [19]. The first type are design gallery platforms, such as Dribbble [2] and Behance [1], which allow designers to

browse through a collection of designs and find examples that are interesting or useful for their work. The second type are design inspirational tools that suggest examples based on certain types of design input, such as a sketch or an existing design, using algorithms to determine image similarity [4, 26, 34, 37]. While these approaches can be helpful in finding inspiration, they have limitations. Browsing through design galleries can be overwhelming and lead to a shift in design ideas away from the original focus, while relying too much on examples with similar styles can lead to design fixation and hinder the originality of the work [21, 28]. Artificial Intelligence has emerged as a potential tool for supporting and enhancing the creative process [36]. AI systems can be trained to generate ideas and outputs based on a set of rules or guidelines, enabling the efficient production of a wide range of options, and providing a variety of applications and systems to support professionals in various visual art fields, such as graphic design [38], UI design [39], webtoon [23], digital art [41], and new media art [33]. Scholars are increasingly studying AI-powered, AI-enhanced, or AI-assisted human creativity [8, 10, 11, 14, 29], reporting the application of AI in the creative industry [3] and art industries [9]. Designers and design researchers have also discussed and practiced data-driven design [13, 18, 22]. Recent advancements in text-to-image (T2I) model capabilities have enabled the generation of realistic images from textual descriptions [5]. These models not only streamline the process of visualizing concepts but may also spark unexpected creative inspirations for users [35]. To find a balance between targeted and serendipitous inspiration, Mozaffari et al. [30] proposed a style-based generative adversarial network (StyleGAN) trained on a large dataset of existing interface designs. It generates a diverse yet focused set of examples based on a preliminary design input. While it can generate a diverse range of interface mockups, the user has no control over the specific layout or design elements of the mockup, which is crucial for a designer to express ideas that are inherently visuo-spatial by nature. Recently, Garg et al. [15] proposed a method to guide design exploration using diffusion models, [42] through different modalities. However, it remains unclear which modality, if any, can work better for specific design tasks and stages of the design process. A growing alternative to text-only prompting is **visual prompting**, which allows users to provide more explicit guidance through sketches, segmentation maps, or semantic labels. This technique is exemplified by ControlNet [43], a deep learning framework that enables fine-grained control over AI-generated outputs by conditioning the model on structured inputs such as edge maps, human poses, or depth maps. By incorporating structured constraints, ControlNet allows designers to generate AI-assisted content that remains consistent with their original composition while still leveraging the generative capabilities of diffusion models. There has been limited focus on comparing different visual prompting modalities for controlling the generation of design artifacts. Our work focuses on evaluating the effectiveness and user experience of two different visual prompting modalities—sketches and semantic-colored drawings—for generating design artifacts. We build upon previous knowledge in several ways: First, the use of sketches as a modality builds upon previous research on the use of visual representations in

the design process [25]. Studies have shown that sketches can be used to quickly generate and communicate design ideas, allowing designers to explore a wide range of possibilities in a relatively short amount of time. Additionally, sketches are often used as a way to capture informal, early-stage design ideas, which can then be refined and developed further [24]. Second, the use of semantic-colored drawings [20, 44], where each color represents a UI element category, allows for a clear representation of the design elements and their relationships. In this paper, we aim to answer the following research questions: **RQ1:** How do different modalities of UI mockup generation, such as sketches and semantic drawings, compare in terms of their unique strengths, perceived usefulness, and impact on the creative process? **RQ2:** What are the trade-offs between expressiveness, efficiency, and user satisfaction when considering the level of control and abstraction afforded by each modality in UI mockup generation? We answer these questions by exploring two distinct AI-assisted modalities: one based on sketches, and the other based on semantic-colored drawings. Sketch-based mockup generation involves the use of hands-drawn sketches to represent the desired design. Semantic-based mockup generation, instead, involves the use of colored drawings that convey specific design elements. The choice of these two modalities was based on a careful review of the literature, indicating that sketching and semantic drawing are commonly used in the design process due to their ability to balance precision and speed (e.g., [7, 27, 31]). Sketch-based mockup generation allows for quick exploration of ideas while semantic-colored drawing-based mockup generation provides more accuracy in representing the final product.

2 Study: Evaluating Visual Prompting Modalities

This study aimed to evaluate the sketch-to-mockup and semantic drawing-to-mockup visual prompting modalities in terms of time demand, ease, creative expressivity, and intuitiveness. The study was conducted with 13 human-computer interaction students from the first-year master’s program in Computer Engineering. The participants were asked to use both modalities to create a mockup of a mobile application, retrieve the AI-generated mockups, and then rate their experience with each modality and result on various dimensions.

Participants The study was conducted remotely with 13 participants who gave their explicit and informed consent to participate. All participants in the study had received training in UI design and programming as part of their curriculum (e.g., through a dedicated Human-Computer Interaction course they all passed), indicating that they had some level of experience in this area. The sample consisted of 2 participants who identify themselves as female and 11 people who identify as male, with an age range of 23 to 25 years old.

2.1 Procedure

Participants completed a two-phase task designing a mobile UI through sketching and semantic drawing. The task evaluated time demand, expressiveness,

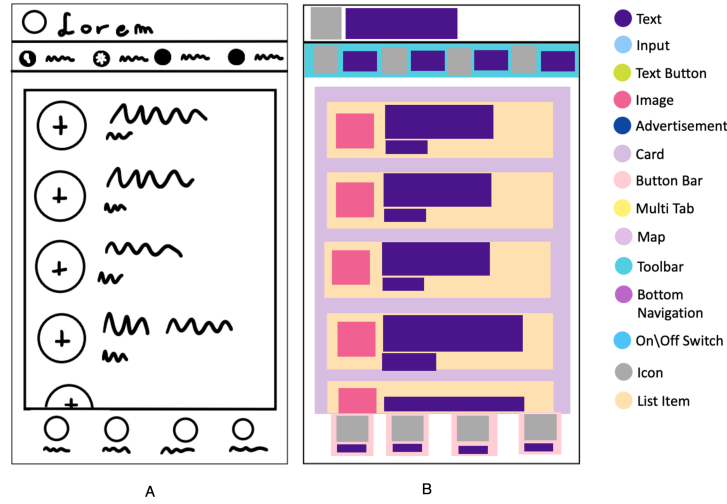


Fig. 1. An example of the same user interface represented in the sketch-based and semantic-based modalities. On the left, we can see a hand-drawn sketch of the user interface (A), with various elements such as buttons, text fields, and icons depicted in a freeform manner. On the right, we can see the same user interface represented in the semantic-based modality (B). These elements are filled with different colors, with each color representing a specific design element such as buttons, text fields, and icons.

intuitiveness, and ease of use in the first phase, followed by assessing the quality and fidelity of generated mockups in the second phase.

In phase one, participants received instructions to design a mobile interface using their preferred digital drawing tools. They first created a sketch, then semantically drew the same interface. This order followed the natural design progression from unstructured ideas to concrete iterations, allowing us to evaluate each modality’s strengths in a contextualized design process. Figure 1 shows an example interface in both modalities. After completing this phase, participants rated their experience on a 1-5 Likert scale (5 being most positive) and provided rationales for their ratings.

In phase two, participants’ drawings were processed by a network (see “Implementation” section), and the results were shown to them. Participants then evaluated the generated mockups’ quality and how well they respected their imagined interface (fidelity), using the same Likert scale. They also had the opportunity to provide additional comments and suggestions at the end of the study.

Implementation To translate sketches and semantic colored drawings to mockups we finetuned a Pix2Pix [40] model for sketch-based mockup generation and the SPADE [32] model for semantic-based mockup generation, following the same procedure used in [30]. We trained the model on the RICO dataset [12],

which uniquely provides both sketches and semantic annotations of the same UI designs. It is important to acknowledge that the quality of the generated mockups in our study may not be optimal, as the focus of this research is on evaluating the visual prompting modalities rather than the generation results themselves. While the generated mockups may exhibit some artifacts or inconsistencies, the rapid progression of AI technologies suggests that the quality of the generated mockups will likely improve in the future. Nevertheless, our findings and discussions regarding the strengths, weaknesses, and user preferences of the sketch-to-mockup and semantic drawing-to-mockup are expected to hold independently of the specific model performances.

3 Results and Discussion

Table 1. Results of sketch-to-mockup and semantic drawing-to-mockup on intuitiveness, ease of use, perceived time demand, and expressiveness (first phase), as well as the quality and fidelity of the AI-generated mockups (second phase), on a 1 to 5 Likert scale.

	Sketch-to-mockup	Semantic drawing-to-mockup
Intuitiveness	4.54 ± 0.78	3.15 ± 1.28
Ease of use	3.38 ± 1.04	3.69 ± 1.25
Time Demand	3.54 ± 1.13	3.77 ± 1.09
Expressiveness	4.15 ± 0.55	3.38 ± 1.12
Quality	1.50 ± 0.58	2.75 ± 0.50
Fidelity	3.38 ± 2.20	3.60 ± 0.55

The study analyzes the sketch-based and semantic-based visual prompting modalities by gathering participant feedback on their level of intuitiveness, ease of use, and time demand during the generation process, as well as the quality and fidelity of the generated mockups. The results of the questions provide insight into how participants perceive these modalities and their relative strengths and weaknesses. Numerical results are reported in Table 1, and the following sections describe the quantitative and qualitative aspects of those results.

3.1 Intuitiveness

The comparison between modalities revealed a clear preference for the intuitiveness of sketch-based visual prompting. Most participants (11/13) rated the sketch-based modality as fully intuitive (score 5), while only two gave semantic generation a score of 5, and most (8/13) rated it 3 or lower.

Participants favoring the sketch-based approach cited that sketching allowed them to "show others what you have in mind" and "represent at a higher level each part of the layout" (participants 1 and 6). The sketch-based modality made

it easier to "imagine the layout" (participant 9) and "pour the idea in mind into drawings" (participant 10).

Some participants found the semantic-based modality challenging, expressing difficulty understanding the "rules about the colors" (participant 4) and finding it "not so clear how it can be useful" (participant 1). However, a few found semantic generation more intuitive, appreciating the "map about element types" (participant 10).

Overall, the sketch-based modality emerged as superior for intuitiveness, facilitating effective communication and translation of mental concepts into visual form, while the semantic-based modality was hindered by its color-coding rules and perceived complexity.

Ease of use The comparison between the sketch-based and semantic-based visual prompting modalities for ease of use yielded mixed results, with no statistically significant difference in the mean scores between the two methods. This indicates that both approaches presented challenges and advantages to the participants.

For the sketch-based visual prompting modality, a majority of participants (7 out of 13) rated it as difficult to use, assigning scores of 3 or lower. The primary reasons cited for this difficulty were the need for "drawing skills and tools" (participant 9) and the requirement of "some drawing skills are needed to convey information in a reasonable and clear way" (participant 1). Furthermore, some participants found it "difficult to use the mouse to draw" (participants 3, 4, and 11), while others found it difficult to use because they "had to use different devices to complete the task" (participant 13).

On the other hand, the semantic modality also had its share of ease-of-use concerns. Participant 9 noted that the semantic approach "needs to remember the colors according to the legend," which can be cognitively demanding. However, some participants found the semantic modality easier to use because it "only requires squares and rectangles" (participant 4) and is "just a mechanical exercise" (participant 3). Participant 10 also found the semantic generation easy to use because it is "much more descriptive" and "easier regarding positioning the items" as compared to sketch generation.

The diverse opinions on the ease of use of both visual prompting modalities suggest that individual preferences, skills, and familiarity with the tools play a significant role in determining the perceived difficulty or ease of use. The sketch-based visual prompting modality may be more challenging for those lacking drawing skills or struggling with input devices, while the semantic modality may be seen as cognitively demanding in terms of color association but appreciated for its simplicity and descriptiveness in positioning elements.

Time demands The time demand comparison yielded mixed results. Participants had varying opinions on each approach's efficiency.

Some found sketch generation quicker due to its ability to represent layouts at a higher level. Participant 5 mentioned it "only took me some minutes to complete" the sketch, suggesting faster translation of mental concepts. However, others cited challenges affecting efficiency, including "difficulty drawing with a

mouse” (participant 10) and “difficulty in getting a decent result” (participant 3).

For semantic generation, some found it less time-demanding. Participant 8 considered it “the quickest one because it’s similar to sketch generation but with even fewer details to represent,” while participant 3 found it “fairly quick to draw what I was thinking in the semantic way.” Conversely, others found it more time-consuming. Participant 1 reported spending “very much time in finding a software to draw colored squares,” while participant 12 noted “it can take a long time because there is often a difference between ideas and reality.”

Time efficiency ultimately depends on the user’s tool familiarity, ability to translate ideas visually, and layout complexity.

Expressivity The comparison between the visual prompting modalities in terms of expressivity revealed a clear preference for the sketch-based visual prompting modality. Most participants (10 out of 13) rated the sketch generation as highly expressive, assigning it a score of 4 or 5, indicating that they found it to be fully capable of conveying their ideas. In contrast, the majority of participants (7 out of 13) gave the semantic generation a score of 3 or lower, suggesting that they perceived it as less expressive.

Participants who favored the sketch-based visual prompting modality cited several reasons for their preference. They found that sketching “makes it easiest to describe how you want the final result to look like” (participant 3) and “irons out some ambiguity” (participant 5). The ability to “draw almost everything” (participant 5) and express “the alignments and the position of the items” more easily (participant 6) were also highlighted as advantages of the sketch-based visual prompting modality.

On the other hand, some participants found the semantic generation to be less expressive. They felt that it “just conveys information with colored and filled squares that don’t reflect very well the actual elements” (participant 1) and is “too restrictive” (participant 4). Participants also encountered difficulties in defining “the role of each component” (participant 7) using the semantic approach.

These findings suggest that while the sketch-based visual prompting modality is generally preferred for its expressivity, the semantic-based modality may still have a role to play in certain contexts or for specific design tasks.

Quality and fidelity of the generated mockups The evaluation of the generated mockups revealed that the semantic-based modality outperformed the sketch-based modality in terms of both quality and fidelity to the original idea. The mean quality score for the semantic-based modality was 2.75, considerably higher than the 1.50 mean score for the sketch-based modality. This suggests that the mockups generated using the semantic-based approach were generally perceived as being of higher quality compared to those generated using the sketch-based approach.

Similarly, the semantic-based modality scored higher in terms of fidelity to the original idea, with a mean score of 3.60 compared to 3.38 for the sketch-

based modality. Although this difference is not as pronounced as the quality score difference, it still shows a significant difference. It indicates that the sketch-based approach was less effective in preserving the designer’s original intent and not accurately translating their ideas into the generated mockups.

Additional Comments and Design Implications The analysis of participant comments reveals diverse opinions regarding the best method for generating mockups. Several participants proposed hybrid solutions integrating elements of both modalities. Participant 7 suggested incorporating color conventions into sketching, combining expressive freedom with semantic clarity. Participant 11 acknowledged each approach’s limitations and proposed leveraging their respective strengths. Participant 12 took a context-dependent stance, arguing that method selection should depend on project circumstances.

Our findings indicate that sketch-based approaches are more intuitive and expressive, aligning with research highlighting sketching’s importance for design ideation [6, 16]. However, this approach requires drawing skills and suitable input devices.

The semantic-based approach, while less intuitive, produced higher quality mockups with better fidelity. Its color-coding and predefined elements contributed to more consistent translations of designers’ intentions [32].

These results suggest each approach has unique merits, with selection depending on specific designer needs and project requirements. Future research could focus on developing hybrid tools that integrate both approaches, allowing designers to leverage their respective strengths at different design stages.

4 Conclusions and Future Work

The paper investigated sketch-based and semantic-based visual prompting modalities for AI-assisted mockup generation, evaluating their intuitiveness, ease of use, time demand, expressiveness, quality, and fidelity.

The findings suggest the sketch-based modality is more intuitive and expressive, enabling designers to quickly convey ideas and explore concepts. However, its effectiveness depends on designers’ drawing skills and input devices. The semantic-based modality, while less intuitive and expressive, produced higher quality mockups with better fidelity to original designs.

Based on these findings, we conclude that a hybrid approach combining strengths of both modalities offers a promising direction for future research. By integrating sketching and semantic drawing capabilities, AI-assisted tools could provide designers with more flexible, expressive, and accurate means of creating high-quality mockups.

Future work should focus on: (1) larger-scale studies with diverse designer samples to validate findings; (2) developing and evaluating hybrid tools integrating both modalities; (3) conducting longitudinal studies to track how preferences evolve as designers become familiar with AI-assisted tools; (4) exploring integration into existing design workflows; and (5) addressing ethical implications of AI-assisted design tools.

References

1. Behance - search projects: Photos, videos, logos, illustrations and branding. (2022), <https://www.behance.net>, <https://www.behance.net>
2. Dribbble - discover the world's top designers and creative (2022), <https://dribbble.com/>, <https://dribbble.com/>
3. Anantrasirichai, N., Bull, D.: Artificial intelligence in the creative industries: a review. *Artificial Intelligence Review* **55** (01 2022). <https://doi.org/10.1007/s10462-021-10039-7>
4. Behrang, F., Reiss, S.P., Orso, A.: Guifetch: Supporting app design and development through gui search. In: *Proceedings of the 5th International Conference on Mobile Software Engineering and Systems*. p. 236–246. MOBILE-Soft '18, Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3197231.3197244>, <https://doi.org/10.1145/3197231.3197244>
5. Betker, J., Goh, G., Jing, L., Brooks, T., Wang, J., Li, L., Ouyang, L., Zhuang, J., Lee, J., Yufei Guo, e.a.: Improving image generation with better captions. *Computer Science* (2023), <https://cdn.openai.com/papers/dall-e-3.pdf>
6. Buxton, B.: *Sketching user experiences: getting the design right and the right design*. Morgan Kaufmann (2010)
7. Calò, T., De Russis, L.: Style-aware sketch-to-code conversion for the web. In: *Companion of the 2022 ACM SIGCHI Symposium on Engineering Interactive Computing Systems*. p. 44–47. EICS '22 Companion, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3531706.3536462>, <https://doi.org/10.1145/3531706.3536462>
8. Carter, S., Nielsen, M.: Using artificial intelligence to augment human intelligence. *Distill* (2017). <https://doi.org/10.23915/distill.00009>, <https://distill.pub/2017/aia>
9. Chen, W., Shidujaman, M., Xuelin, T.: Aiart: Towards artificial intelligence art (02 2020)
10. Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P.: In-fog: Interpretable representation learning by information maximizing generative adversarial nets. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 29. Curran Associates, Inc. (2016), <https://proceedings.neurips.cc/paper/2016/file/7c9d0b1f96aebd7b5eca8c3edaa19ebb-Paper.pdf>
11. Colton, S., Wiggins, G.: Computational creativity: The final frontier? *Frontiers in Artificial Intelligence and Applications* **242**, 21–26 (01 2012). <https://doi.org/10.3233/978-1-61499-098-7-21>
12. Deka, B., Huang, Z., Franzen, C., Hibschan, J., Afergan, D., Li, Y., Nichols, J., Kumar, R.: Rico: A mobile app dataset for building data-driven design applications. In: *Proceedings of the 30th Annual Symposium on User Interface Software and Technology*. UIST '17 (2017)
13. Dove, G.: *Codesign with data* (2015)
14. Fischer, G., Nakakoji, K.: Amplifying Designers' Creativity with Domain-Oriented Design Environments, pp. 343–364. Springer Netherlands, Dordrecht (1994). <https://doi.org/10.1007/978-94-017-0793-025>, <https://doi.org/10.1007/978-94-017-0793-025>
15. Garg, A., Jiang, Y., Oulasvirta, A.: Controllable gui exploration. *arXiv preprint arXiv:2502.03330* (2025)

16. Goldschmidt, G.: The dialectics of sketching. *Creativity research journal* **4**(2), 123–143 (1991)
17. Gonçalves, M., Cardoso, C., Badke-Schaub, P.: What inspires designers? preferences on inspirational approaches during idea generation. *Design Studies* **35** (01 2013). <https://doi.org/10.1016/j.destud.2013.09.001>
18. Ha, D., Eck, D.: A neural representation of sketch drawings. In: 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings. OpenReview.net (2018), <https://openreview.net/forum?id=Hy6GHpkCW>
19. Herring, S.R., Chang, C.C., Krantzler, J., Bailey, B.P.: Getting inspired! understanding how and why examples are used in creative design practice. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. p. 87–96. CHI '09, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1518701.1518717>, <https://doi.org/10.1145/1518701.1518717>
20. Hong, S., Yang, D., Choi, J., Lee, H.: Inferring semantic layout for hierarchical text-to-image synthesis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 7986–7994 (2018)
21. Jansson, D., Smith, S.: Design fixation. *Design Studies* **12** (01 1991). [https://doi.org/10.1016/0142-694X\(91\)90003-F](https://doi.org/10.1016/0142-694X(91)90003-F)
22. King, R., Churchill, E.F., Tan, C.: Designing with data: Improving the user experience with a/b testing (2017)
23. Ko, H.K., An, S., Park, G., Kim, S.K., Kim, D., Kim, B., Jo, J., Seo, J.: We-toon: A communication support system between writers and artists in collaborative webtoon sketch revision. In: Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology. UIST '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3526113.3545612>, <https://doi.org/10.1145/3526113.3545612>
24. Landay, J., Myers, B.: Sketching interfaces: toward more human interface design. *Computer* **34**(3), 56–64 (2001). <https://doi.org/10.1109/2.910894>
25. Landay, J.A., Myers, B.A.: Interactive sketching for the early stages of user interface design. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. p. 43–50. CHI '95, ACM Press/Addison-Wesley Publishing Co., USA (1995). <https://doi.org/10.1145/223904.223910>, <https://doi.org/10.1145/223904.223910>
26. Lee, B., Srivastava, S., Kumar, R., Brafman, R., Klemmer, S.R.: Designing with interactive example galleries. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. p. 2257–2266. Association for Computing Machinery, New York, NY, USA (2010). <https://doi.org/10.1145/1753326.1753667>, <https://doi.org/10.1145/1753326.1753667>
27. Li, Y., Liu, H., Wu, Q., Mu, F., Yang, J., Gao, J., Li, C., Lee, Y.J.: Gligen: Open-set grounded text-to-image generation. *CVPR* (2023)
28. Marsh, R., Landau, J., Hicks, J.: How examples may (and may not) constrain creativity. *Memory and Cognition* **24**, 669–680 (09 1996). <https://doi.org/10.3758/BF03201091>
29. Miller, A.: The Artist in the Machine: The World of AI-Powered Creativity (01 2019). <https://doi.org/10.7551/mitpress/11585.001.0001>
30. Mozaffari, M.A., Zhang, X., Cheng, J., Guo, J.L.: Ganspiration: Balancing targeted and serendipitous inspiration in user interface design with style-based generative adversarial network. In: Proceedings of the 2022 CHI Conference on Hu-

- man Factors in Computing Systems. CHI '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3491102.3517511>, <https://doi.org/10.1145/3491102.3517511>
31. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019)
 32. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 2332–2341 (2019)
 33. Qiao, H., Liu, V., Chilton, L.: Initial images: Using image prompts to improve subject representation in multimodal ai generated art. In: Creativity and Cognition. p. 15–28. Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3527927.3532792>, <https://doi.org/10.1145/3527927.3532792>
 34. Ritchie, D., Kejriwal, A.A., Klemmer, S.R.: D.tour: Style-based exploration of design example galleries. In: Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology. p. 165–174. Association for Computing Machinery, New York, NY, USA (2011). <https://doi.org/10.1145/2047196.2047216>, <https://doi.org/10.1145/2047196.2047216>
 35. de Rooij, A., Biskjaer, M.M.: Expecting the unexpected: A review of surprise in design processes. In: Gray, C., Chehade, E.C., Hekkert, P., Forlano, L., Ciuccarelli, P., Lloyd, P. (eds.) DRS2024: Boston, 23–28 June (2024), <https://doi.org/10.21606/drs.2024.333>
 36. Shneiderman, B.: Human-Centered AI. Oxford University Press (2022)
 37. Swearngin, A., Dontcheva, M., Li, W., Brandt, J., Dixon, M., Ko, A.J.: Rewire: Interface design assistance from examples. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. p. 1–12. Association for Computing Machinery, New York, NY, USA (2018). <https://doi.org/10.1145/3173574.3174078>, <https://doi.org/10.1145/3173574.3174078>
 38. Ueno, M., Satoh, S.: Continuous and gradual style changes of graphic designs with generative model. p. 280–289. Association for Computing Machinery, New York, NY, USA (2021). <https://doi.org/10.1145/3397481.3450666>, <https://doi.org/10.1145/3397481.3450666>
 39. Wang, B., Li, G., Zhou, X., Chen, Z., Grossman, T., Li, Y.: Screen2words: Automatic mobile ui summarization with multimodal learning. pp. 498–510 (10 2021). <https://doi.org/10.1145/3472749.3474765>
 40. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 8798–8807 (2017)
 41. Yurman, P., Reddy, A.V.: Drawing conversations mediated by ai. In: Creativity and Cognition. p. 56–70. C&C '22, Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/10.1145/3527927.3531448>, <https://doi.org/10.1145/3527927.3531448>
 42. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3836–3847 (2023)
 43. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3836–3847. IEEE (2023), <https://arxiv.org/abs/2302.05543>

44. Zhu, P., Abdal, R., Qin, Y., Wonka, P.: Sean: Image synthesis with semantic region-adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5104–5113 (2020)