Less Redraw, More Explore: Suggestion and Completion for Sketch-to-Image

Zeyu Zhao
Computer Science
Swansea University
Swansea, United Kingdom
mattjonez@gmail.com

Connor Rees
Swansea University
Swansea, United Kingdom
connor.a.rees@swansea.ac.uk

Gavin Bailey
Swansea University
Swansea, United Kingdom
g.bailey@swansea.ac.uk

Matt Jones
Swansea University
Swansea, United Kingdom
matt.jones@swansea.ac.uk

Simon Robinson
Swansea University
Swansea, United Kingdom
s.n.w.robinson@swansea.ac.uk

Jennifer Pearson
Swansea University
Swansea, Wales, United Kingdom
bagelobagel@gmail.com

Abstract

Sketch-to-image systems let users transform simple line drawings into realistic images, but current workflows force users into tedious redraw-regenerate cycles that slow creative exploration. We introduce two complementary interaction techniques that reduce iteration friction: *AutoSketch*, which extends partial sketches through AI-driven completions (*pre-generation* support), and *BackSketch*, which transforms generated images back into editable sketches at multiple abstraction levels (*post-generation* support). In a study with 30 participants, the results indicate that both techniques can improve exploration and expressiveness compared to a baseline sketch-to-image system, while *AutoSketch* also can increase users' sense of agency and co-creation with the AI. We contribute new evidence that shifting support before or after generation opens distinct pathways for balancing user control and system initiative. Together, our results establish pre- and post-generation assistance as a design space for co-creative sketch-to-image systems.

CCS Concepts

• **Human-centered computing** → *Interaction techniques*; • **Applied computing** → *Image composition.*

Keywords

Generative AI, sketch-to-image, creativity, user interface, image generation, user study

ACM Reference Format:
Zeyu Zhao, Connor Rees, Gavin Bailey, Matt Jones, Simon Robinson, and Jennifer Pearson. 2026. Less Redraw, More Explore: Suggestion and Completion for Sketch-to-Image. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26), April 13–17, 2026, Barcelona, Spain.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3772318.3791026

CHI '26, Barcelona, Spain

ACM ISBN 979-8-4007-2278-3/26/04
https://doi.org/10.1145/3772318.3791026

1 Introduction

Sketching is a well-established medium for visual thinking and creative expression. In fields such as design and art, sketches serve as fast, flexible externalisations of visual ideas. With the rise of generative AI, sketch-to-image systems now allow users to transform rough line drawings into realistic images, broadening access to digital design tools [10, 29]. By grounding generation in visual marks rather than text prompts, these systems support more intuitive ideation for visually oriented users.

However, current interactive workflows remain limited. A typical system accepts a sketch, produces an image, and leaves refinement up to the user. If the result is unsatisfactory, the user must redraw or adjust their sketch and regenerate, a process that is often tedious and unpredictable [20, 31]. Prior research has made progress through algorithmic advances, such as ControlNet [38], T2I-Adapter [27], DeepFaceDrawing [9] and SketchFlex [20], as well as interactive interfaces like GANPaint Studio [3], Dream Lens [26] and AdaptiveSliders [16]. Yet most approaches still treat sketches as static inputs, offering limited support for the iterative, back-and-forth refinement that is central to creative practice [32].

We address this gap with two sketch-to-image systems that preserve a sketch-only interface while enabling more dynamic, iterative workflows:

- **AutoSketch:** Users can draw partial sketches and invoke AI completions that add meaningful content. Multiple completions can be added to a single sketch, and each completion can be accepted, edited or undone, supporting progressive refinement through human-AI collaboration. This helps users decide what to do next after an initial idea by offering concrete extensions to their sketch.
- **BackSketch:** After the user generates an image from their initial sketch, the generated image is converted into a set of simplified sketches at different levels of abstraction. Users can edit their original sketch or switch to one of these suggestions, reducing the need to redraw when exploring new directions. This helps users quickly branch out into alternative ideas without starting over, making iteration more efficient and less effortful.

Together, these systems move beyond existing static input-output pipelines to support more fluid, iterative sketch-to-image interaction. This paper makes **three contributions** to research on creativity support and human-AI co-creation: (i) **Two novel interaction techniques for sketch-to-image workflows**—pre-generation completion (AutoSketch) and post-generation suggestion (BackSketch)—that reduce iteration cost and expand creative possibilities while maintaining a sketch-only interface; (ii) **Design insights and implementation strategies for embedding iterative support into sketch-to-image pipelines**, highlighting how reversible AI contributions (e.g., undoable completions, selectable alternative sketches) help balance user agency and system initiative; and, (iii) **Empirical evidence from a 30-participant study** indicating that these techniques foster exploration, expressiveness, and a stronger sense of co-creative partnership compared to a baseline system, establishing pre- and post-generation support as a generalisable design space for co-creative AI tools.

In the remainder of this paper we first review related work, then describe the architecture of our two novel designs. We evaluate both designs through a user study against a baseline system, focusing on creativity, usability and satisfaction, and conclude by highlighting key insights from this work.

2 Related work

Sketch-to-image synthesis enables intuitive, human-centred creation from sparse line drawings. Early GAN-based systems such as SketchyGAN [10] and Scribbler [30] translated sketches into realistic images by learning colour, texture and detail. Subsequent work expanded to domain-specific generation (e.g., faces in DeepFaceDrawing [9]) and richer representations that leverage stroke order and vector encodings (Sketch-R2CNN [18]). With diffusion models, ControlNet [38] and T2I-Adapter [27] provide fine-grained conditioning of large models such as Stable Diffusion [28] using edges, scribbles or poses. SketchFlex [20] further improves spatial-semantic alignment with region-based prompts and refinement pipelines. However, these approaches largely treat the sketch as a *static* input, supporting one-way translation from sketch to image with limited support for iteration.

2.1 Interactive image generation interfaces

Various interfaces have been proposed to make generative workflows more controllable and interactive. Promptify [4] supports prompt engineering with language guidance and clustering; GANPaint Studio [3] and Paint by Word [5] enable semantic edits via brushes or text; Opal [22] and 3DALL-E [23] extend editing to multimodal and 3D contexts; and, AdaptiveSliders [16] offers dynamically adjusted, semantically aligned sliders for diffusion-based editing. While effective, these systems primarily operate at the prompt or latent level. In contrast, our work centres sketch-level input and enables post-generation *sketch suggestion* (BackSketch) and pre-generation *sketch completion* (AutoSketch).

2.2 Interactive sketch editing and iteration

Reversible workflows project generated results back into interpretable forms to support iteration. Prompt-to-Prompt [15], CLIP Interrogator[1], and Image2StyleGAN [1] refine outputs via captions, attention or latent embeddings. For sketch contexts, sketch simplification [36] and contour extraction [7] produce editable outlines from images. Building on these ideas, BackSketch closes the loop by converting generated images into multiple simplified sketches that users can select and refine, while AutoSketch supports progressive development by extending partial sketches with AI completions. Related domain-specific tools (e.g., SketchAI [12]) demonstrate the value of assistive sketching but do not generalise to a sketch-only pipeline that integrates either pre- or post-generation support.

2.3 Creativity and iterative design support

Creativity support tools emphasise exploration across alternatives. Design Galleries [24] and Dream Lens [26] help navigate large spaces of outputs; GANravel [14] and GEM-NI [37] provide interpretable manipulation via latent controls or graph workflows; and, CompSketch [34] treats sketches as interactive design objects for parallel prototyping. Prompt-oriented tools (e.g., MagicPrompt[2] and Automatic1111[3]) further support iteration at the text level. Similarly, recent work such as Block-to-Detail Scaffolding [31] demonstrates how simple sketch inputs can be progressively refined into more detailed iterations, highlighting the role of incremental sketch refinement in supporting creativity. Our contribution complements this literature by offering two *sketch-only* mechanisms that embed iterative exploration directly in the sketching workflow, reducing redraw effort and enabling fluid, co-creative refinement.

2.4 Recent advances in co-creative systems

Recent work in co-creative generative systems has expanded how users steer, remix, and iterate on AI-augmented visual ideas. CreativeConnect [11] enables reference recombination for early-stage ideation, generating blended variations from curated visual materials to support divergent exploration. L.Ink [8] introduces controllable procedural growth for sketching, showing how well-scoped unpredictability can provoke reflection-in-action and stimulate experimental workflows. ImaginationVellum [25] frames the entire canvas as a spatial prompt, using proximity-dependent intent tags, generative strokes, and ideation histories to fluidly traverse design spaces.

Compared to these systems—which combine sketches with text prompts, reference materials, or spatial layout—BackSketch and AutoSketch remain sketch-only and place iterative exploration directly within the sketching interaction itself. BackSketch offers contrastive, post-generation entry points by surfacing alternative sketch interpretations, while AutoSketch provides pre-generation completions that expand ideas as they are being drawn. Together, these two techniques introduce complementary access points into the design space of creativity support, allowing users to explore alternatives without switching modalities or relying on external scaffolds.

3 Sketch-to-image system designs

In order to explore alternative interaction techniques, we first constructed a baseline sketch-to-image system built on the `gpt-image-1` model provided by OpenAI. The interface offers a

[1] https://github.com/pharmapsychotic/clip-interrogator

[2] https://huggingface.co/spaces/Gustavosta/MagicPrompt-Stable-Diffusion

[3] https://github.com/AUTOMATIC1111/stable-diffusion-webui

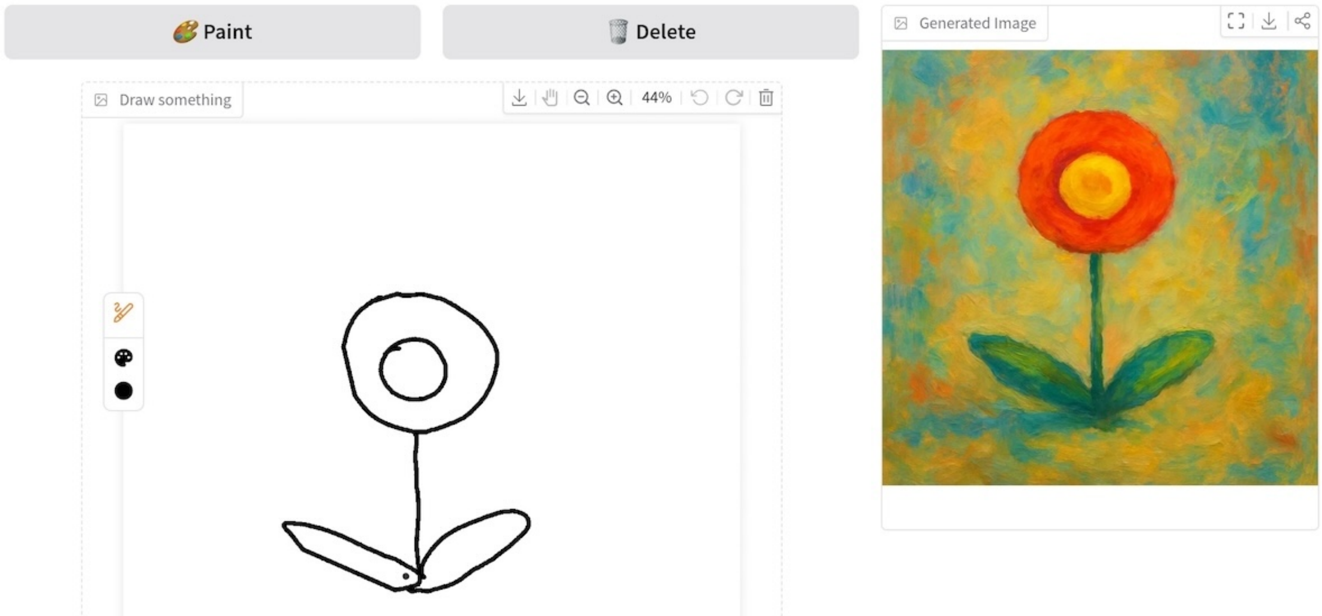

Figure 1: The interface of the baseline sketch-to-image system we created to use as a control. Users begin by drawing a sketch on the canvas (left). When the user presses *Paint*, the system generates a corresponding image in a painting-like style (right).

minimal image generation workflow and serves as a reference point for comparison to our novel sketch-to-image interaction techniques. BackSketch and AutoSketch build on this baseline to add their pre- and post-generation support.

A core motivation across our two enhanced systems is to reduce the friction inherent in iterative sketch-to-image workflows – where users frequently redraw-regenerate between sketch adjustments and model outputs [4, 21]. Inspired by long-standing ideas that highlight fluid, reversible iteration as central to creative work—such as Schön's "reflection-in- action" [33] and Buxton's call for rapid, low-cost sketch transformations [6]—we designed techniques that support iteration at different stages of the creative loop. AutoSketch reduces pre-generation friction by providing in-place sketch completions, helping users elaborate ideas without restarting their sketch. Conversely, BackSketch reduces post-generation friction by surfacing alternative sketch interpretations derived from the model's output, enabling users to branch, revise, or re-interpret without redrawing from scratch. Together, these two techniques support our goal of improving *iteration fluidity*, helping users move easily between idea, sketch and image without heavy commitment or effort. Our approach aligns with earlier co-creative systems that promote fluid iteration through suggestion or progressive refinement [17, 19, 24], but extends them by offering a sketch-only workflow that stays reversible, lightweight, and continuous.

3.1 Baseline system

The baseline system, as shown in Fig. 1, provides a minimalist interface – essentially a plain sketch-to-image mapping using the underlying AI image generation model directly with no auxiliary suggestion or completion tools. The interface consists of a sketchpad with standard drawing controls, two buttons (*Paint* and *Delete*), and an output image container. Users draw freely on the sketchpad and press *Paint* to generate a painting-style artistic image. If unsatisfied, users can refine their sketch and *Paint* again, or *Delete* to reset the interface. This workflow reflects the one-shot, linear workflow common to sketch-to-image systems [10], where iteration often requires redrawing or repeated regeneration [4, 21].

The prompt used for AI image generation is:

> *Create a painting-style image based on the input sketch. The image should follow the structure of the sketch closely, and be rendered in a colourful, vibrant painting style. Use soft brush strokes, rich textures, and a vivid palette with natural lighting. The final result should look like a traditional painting, full of life and colour.*

Our approach of constructing a baseline for comparison aligns with prior work in generative co-creative systems, where comparison interfaces serve to isolate the effects of additional scaffolding features (e.g., Inkspire was compared to ControlNet [19]). By comparing our enhanced systems (BackSketch and AutoSketch) directly against this streamlined baseline, we can more clearly attribute improvements in creativity support to our added interaction mechanisms.

3.2 AutoSketch: sketch-to-image + completion

The AutoSketch system extends the baseline interface to introduce two additional buttons: *Sketch* and *Undo*. Instead of completing a full sketch, users can choose to draw only part of the object and press *Sketch*. The system then automatically recognises the drawing and adds one semantically meaningful element to extend it. If the completion aligns with the user's intent, it can be kept or refined; if not, the *Undo* button restores the sketch to its prior state. Figure 2 shows an example of how a user might use this process to iteratively refine their sketch to achieve the desired final image.

The help from AI appears *before* the image generation, shifting system support earlier in the process, encouraging users to build on partial ideas rather than starting from scratch. The prompt used for sketch completion is:

> *Preserve the original drawing style and composition. Add one semantically valid and salient object to the sketch. Only make minimal, sensible additions that enhance the sketch without changing its overall layout. Do not fill in large areas or apply shading. The result should remain a clean, line-based black-and-white sketch.*

3.3 BackSketch: sketch-to-image + suggestions

BackSketch, as shown in Fig. 3 extends the baseline interface with a *Generated Sketches* panel. After an image is painted, the system takes this image as input and decomposes it into a sequence of four sketches that progressively increase in detail (from rough outlines to complete line drawings), as shown in Fig. 3. These four sketches, together with the user's original input sketch, are displayed as alternatives in the gallery. Users can select any sketch as the new basis for continued editing in the sketchpad or further generation, enabling post-generation inspiration and iterative exploration. For example, in Fig. 3, the user chose the fourth sketch, as they liked the detail of the mouse's paws and eye and edited it to emphasise the animals expression, remove the butterfly and added some cheese.

This workflow lowers iteration cost by providing multiple candidate sketches [36], and supports exploratory creativity [9, 15, 24]. Note that the sketch suggestions are generated *after* the image, allowing users to explore new directions without redrawing from scratch. The prompt used for generating sketch completions is:

> *You are helping to illustrate the process of drawing a black-and-white sketch from a complete image. You will produce four sequential sketches, arranged in a 2×2 grid within a single 1024×1024 image. Each sketch should occupy one quadrant (512×512) and represent a step in the drawing process:*
>
> *- Top-left: Step 1 – very rough sketch, only outer shapes*
> *- Top-right: Step 2 – clearer outlines with some internal lines*
> *- Bottom-left: Step 3 – facial features, folds, fine outlines*
> *- Bottom-right: Step 4 – complete outline drawing with all visible structures*
>
> *Use only clean, thin black lines on a white background. Do not use any shading, hatching, textures, cross-hatching, or filled areas. Each sketch must be purely line-based and clearly show its respective level of completeness.*

4 Evaluation

We designed a controlled study in order to compare AutoSketch and BackSketch against the baseline system, focusing on how each approach shapes creative processes and outcomes, as well as which system(s) participants preferred and why. Thirty participants completed the one-hour IRB-approved study, with two additional pilots used to refine procedure and timing. We recruited non-expert participants because early trials with focus groups showed that the systems primarily supported amateur and casual sketchers, while professionals tended to prefer more direct control over their sketches. Prior work on creativity support tools also shows that interaction burdens, iteration friction and system learnability are experienced most strongly by non-experts, who often lack formal drawing training yet still engage in open-ended creative exploration [6, 35]. Evaluating with this population therefore offered clearer insight into whether our sketch-only techniques genuinely lower barriers to iterative creativity. We employed a within-subjects design with full counterbalancing: each participant used all three systems, and the six possible orders were evenly distributed (five participants per order). A within-subjects design was essential because creative tasks vary widely across individuals. Having each participant use all three systems allows direct comparison while controlling for personal drawing style, confidence and ideation habits. Full counterbalancing mitigates order and learning effects, ensuring that differences in creativity, usability or preference can be attributed to the interaction techniques themselves rather than presentation sequence. To minimise bias, all the systems were anonymised throughout the study.

Each system was used for one session, defined as four minutes of *active time* – time spent sketching and making sketching decisions on the sketchpad canvas. Computation time was excluded, as our focus was on interaction techniques rather than back-end latency. Each session involved at least two distinct sketches. If participants had not already switched, they were prompted after two minutes to start a second sketch. Session duration and task structure were calibrated through pilot testing to balance realism with exploratory breadth. A four-minute window ensures participants experience a complete interaction loop—sketch, generate, revise—while keeping cognitive load manageable and avoiding fatigue in a three-system study. Requiring at least two distinct sketches encourages breadth and prevents overspecialisation on a single idea, promoting exploratory behaviour that aligns with our focus on iterative, divergent creativity [24].

Participants viewed the interfaces on a laptop and used its trackpad as the sole input device. They were free to sketch their own ideas, but optional inspiration was available via 20 prompt cards with suggested scenarios (e.g., "an electric jellyfish"). Participants could adapt prompts to simpler forms (e.g., "jellyfish") if desired, ensuring engagement without requiring advanced drawing skills. They were also asked not to reuse the same idea across systems to avoid bias.

4.1 Procedure

Upon arrival, participants were welcomed by two researchers and seated at a desk. The first researcher explained the study purpose, described the data to be collected (questionnaires and interaction logs), and then left the room until the end to minimise pressure. The second researcher stayed to observe and manually record behaviours, including numbers of image generations, sketch completions and undo actions (AutoSketch), and suggestions selected (BackSketch). To help maintain a calm atmosphere, quiet background music was played throughout.

Figure 2: An example workflow when using AutoSketch: sketch-to-image with pre-generation suggestions. Left: a user draws a partial sketch. Centre left: After pressing *Sketch*, the system adds a watering can above the flower. Centre right: After pressing *Sketch* again, the system adds another small flower and a plant pot. Finally, the user presses *Paint* to generate an image (right).

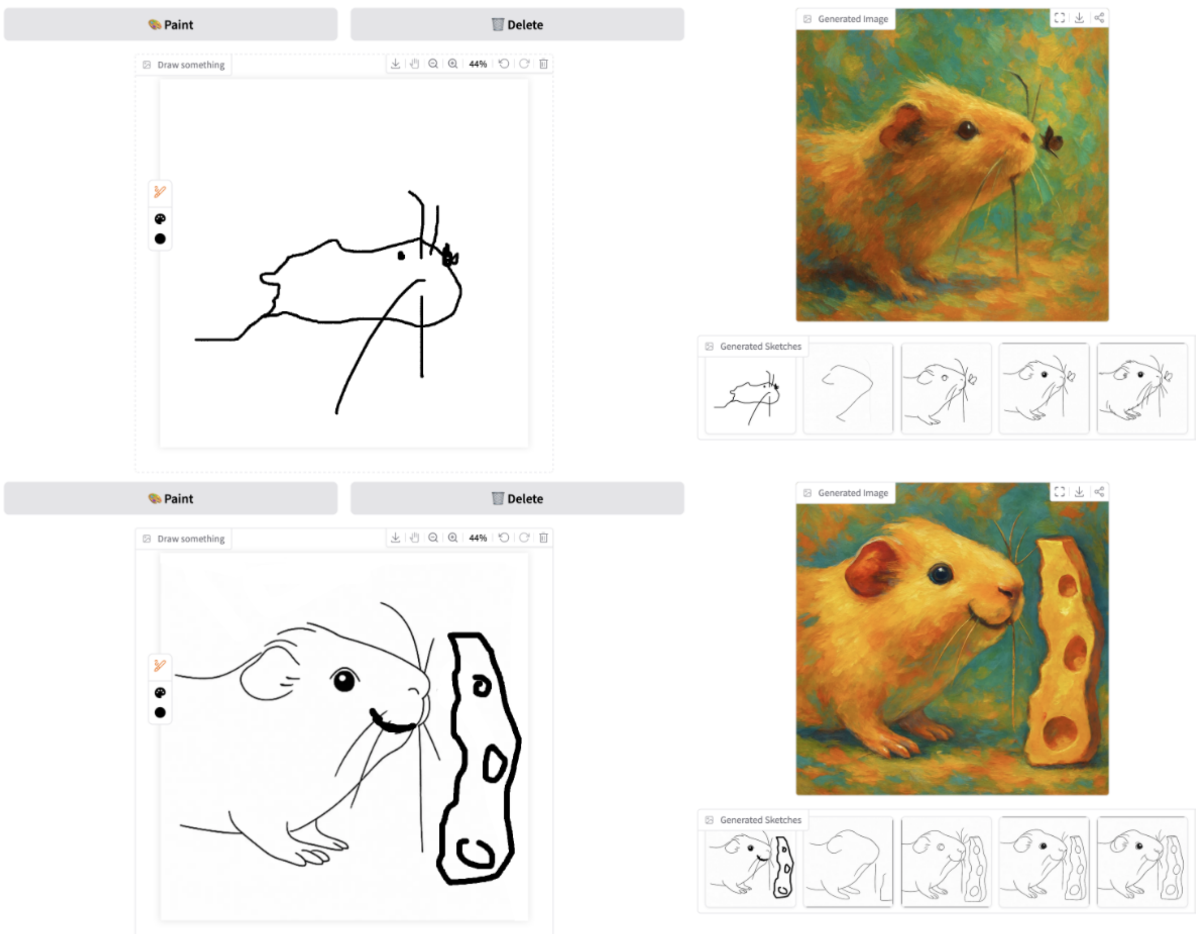

Figure 3: BackSketch: sketch-to-image with post-generation suggestions. Top: users sketch the image they would like to see (left pane), but after viewing the output (right pane), they can choose to replace their own sketch with an alternative sketch generated by the system (bottom right area). Bottom: the user has selected the fourth suggested sketch from the set of options in the original, and has made new edits (thicker lines) before generating a new image. Additional suggestions are generated each time to help the user keep refining their image.

For each of the three systems, the researcher gave a short tutorial with a simple example to ensure consistent understanding. Participants then worked independently for four minutes of active sketching time, while the second researcher remained nearby to track timings and provide technical support without interfering. After using each system, participants completed a short questionnaire about their experience. This cycle of demonstration, interaction, and evaluation was repeated for all three systems. At the end, the participant answered an open-ended question about their overall preference and received a £25 voucher as compensation. Each session lasted about one hour in total.

4.2 Measures

To evaluate the potential for our designs as creativity support tools, we developed a questionnaire addressing both process- and product-oriented dimensions of creativity. The design was guided by prior work on creativity evaluation frameworks [2] and the Mixed Initiative Creativity Support Index (MICSI) [17].

In adapting MICSI to our sketch-based setting, we removed several items that were either redundant or conceptually misaligned with short-form, rapid sketch-to-image interaction. Specifically, we omitted the paired questions on immersion, worth, alignment, and communication (e.g., *"I became so absorbed [. . .]"*, *"The output aligned with my goals"*) as these were strongly overlapping with the core constructs we retained (enjoyment, exploration, expressiveness, communication). We also removed the items on contribution, surprise and control alignment because they assume a stronger human-system separation than our sketch-only workflow supports: in our systems, agency emerges through iterative manipulation rather than through discrete system moves. This simplification follows the approach taken in the original MICSI work [17], where the authors found that questions within each category showed highly similar response patterns, allowing them to merge overlapping items without loss of analytical fidelity. Before finalising our questionnaire, we also conducted small-scale pre-study checks to ensure that the retained items covered the experiential distinctions that participants were most likely to comment on during rapid sketching tasks. Our final questionnaire therefore preserves the dimensions most relevant to lightweight, rapid sketch-based co-creation while avoiding redundancy and participant fatigue.

Participants rated how well the systems supported their creative process on seven dimensions as shown in Table 1. These items captured subjective experiences of creativity support, from enjoyment and expressiveness to whether the system was perceived as a collaborator rather than a passive tool. Participants also evaluated the quality (the creative product) of their outputs via three bipolar questions (1–7).

In terms of system usage, we focused on four behavioural metrics. For all systems we counted image generations – i.e., the number of times participants pressed the *Paint* button to create a painted output. For BackSketch we counted suggestion selections: how many times participants selected one of the suggested sketches to replace their current sketch. With AutoSketch we counted how many times participants pressed the button to invoke AI completion; and, how many times participants pressed the *Undo* button to revert a completion.

In addition to these and the questionnaire items, we collected qualitative feedback through open-ended questions after each session, asking participants to explain their ratings and share aspects they liked or disliked, or suggestions. Finally, after using all three systems, participants indicated their overall preference and explained their choice.

4.3 Participants

Thirty participants took part in the study. Demographic data were collected through a pre-study questionnaire, including participants' prior experience with sketching and AI-assisted art tools.

Participants' age distribution was primarily younger adults, with 57% in the 18–30 range (17 participants), 40% in the 31–40 range (12 participants), and a single participant over 50 (3%). Gender representation was broadly balanced, with 43% male (13 participants), 53% female (16 participants), and one participant identifying as non-binary (3%). Most participants (73%) reported that they did not regularly use AI-assisted art or design tools in their daily or professional lives, with only 17% (5 participants) indicating consistent use and 10% (3 participants) unsure. Similarly, regular sketching or painting for pleasure was uncommon: 70% (21 participants) responded "No," 20% (6 participants) "Yes," and 10% (3 participants) "Maybe." Prior exposure to AI image generation tools was more evenly split, with 57% (17 participants) having tried such tools at least once and 43% (13 participants) reporting no prior experience.

Overall, this suggests that the participant pool largely comprised non-experts, reflecting the intended focus of the study on casual, entertainment-oriented creativity support rather than professional art or design practice.

5 Results

The majority of questionnaire responses across all systems fell into the positive region of the scale (values 5–7). This suggests that participants found all three systems reasonably enjoyable and effective as creativity support tools. Importantly, none of the experimental systems (*BackSketch* and *AutoSketch*) were judged less favourably than the baseline, indicating that our added features did not introduce usability costs.

We conducted paired-sample t-tests of questionnaire responses. Because our study used a within-subjects design, we also tested for order effects by comparing responses according to whether a system was encountered first, second, or third. No significant order effects were found for any questionnaire item (all $p > 0.05$). Following the same approach as the original MICSI work [17], all statistical tests were conducted at an exploratory threshold of $\alpha = 0.05$, and no corrections for multiple comparisons were applied. Consistent with findings reported in the original MICSI study [17], only a small subset of questionnaire items yielded statistically significant differences (at $p < 0.05$).

Full details of all results are given in Table 2 and Fig. 4 in Appendix A; here we summarise specific question clusters. We conducted paired-sample t-tests of questionnaire responses. In addition to p-values, we report Cohen's d_z as an effect size for within-subject comparisons (with 0.2, 0.5, and 0.8 interpreted as small, medium, and large effects). Here we summarise specific question clusters.

Table 1: Creative *process* (Q1–Q7) and creative *product* (Q8–Q10) questions. For process questions, participants responded on 1–7 Likert-like scales (1 = strongly disagree, 7 = strongly agree). Product questions were again on a seven point scale, with participants indicating their strength of feeling in response to each of the bipolar options.

Q	Item	Subscale
1	I enjoyed using the system.	Enjoyment
2	I would be happy to use this system again.	Enjoyment
3	It was easy to explore different ideas or directions using this system.	Exploration
4	I was able to be creative while using the system.	Expressiveness
5	My attention was fully tuned to the activity (I forgot about the tool).	Immersiveness
6	I was able to effectively communicate what I wanted to the system.	Communication
7	At times, it felt like the system and I were collaborating as equals.	Agency
8	The outputs are very typical ↔ very novel.	Originality
9	The outputs are very simple ↔ very detailed.	Elaboration
10	I am very unsatisfied ↔ very satisfied with the outputs.	Overall satisfaction

Table 2: Question-wise descriptive statistics (mean, median) for each system and paired-sample *t*-test *p*-values for system comparisons. Green shading indicates statistically significant differences ($p < 0.05$).

	Baseline		BackSketch		AutoSketch		*t*-test *p*-value		
Q	**Mean**	**Median**	**Mean**	**Median**	**Mean**	**Median**	**Baseline vs. BackSketch**	**Baseline vs. AutoSketch**	**BackSketch vs. AutoSketch**
Q1	5.77	6.0	6.03	6.0	6.07	6.5	0.161	0.213	0.884
Q2	5.90	6.0	6.17	6.5	5.90	6.0	0.310	1.000	0.318
Q3	4.77	5.0	5.50	6.0	5.67	6.0	0.014	0.010	0.634
Q4	5.20	5.5	6.10	6.0	5.93	6.0	0.013	0.056	0.605
Q5	5.03	5.0	5.33	5.0	5.07	5.0	0.343	0.903	0.373
Q6	4.83	5.0	5.13	5.0	5.30	5.5	0.365	0.114	0.670
Q7	4.30	4.5	4.73	5.0	4.97	6.0	0.267	0.035	0.594
Q8	4.87	5.0	5.20	5.0	4.87	5.0	0.326	1.000	0.282
Q9	4.80	5.0	5.50	6.0	5.30	5.0	0.072	0.083	0.586
Q10	5.40	5.5	5.73	6.0	5.63	6.0	0.193	0.482	0.764

Enjoyment (Q1, Q2): Participants reported high levels of enjoyment across all three systems, with median ratings consistently at or above 6. No significant differences were found between systems, suggesting that each provided an engaging and enjoyable experience. Pairwise effect sizes for enjoyment were small (all $|d_z| \leq 0.26$), reinforcing the lack of meaningful differences in this dimension.

Exploration and Expressiveness (Q3, Q4): Significant improvements were observed for both exploration (Q3) and expressiveness (Q4) when comparing the proposed systems to the baseline. BackSketch and AutoSketch both scored higher than the baseline on exploration ($p = 0.014$ and $p = 0.010$, respectively), while BackSketch also outperformed the baseline on expressiveness ($p = 0.013$). AutoSketch showed a marginally higher expressiveness score than the baseline, though this was not a significant difference ($p = 0.056$). These differences correspond to medium-sized effects for exploration (BackSketch vs baseline: $d_z = 0.48$; AutoSketch vs. baseline: $d_z = 0.50$) and a medium-sized effect for expressiveness for BackSketch ($d_z = 0.48$), with a smaller but still non-trivial effect for AutoSketch ($d_z = 0.36$). These results indicate that both proposed systems better supported users in trying out new ideas and feeling creative while sketching.

Immersiveness and Communication (Q5, Q6): Median ratings were stable across systems, with no significant differences. While AutoSketch showed a small trend toward improved communication (Q6), this effect did not reach statistical significance. Effect sizes for immersiveness and communication were generally trivial (all $|d_z| \leq 0.30$), suggesting that these aspects were less influenced by the introduction of suggestions or sketch completions.

Agency (Q7): Participants reported higher perceived agency when using AutoSketch compared to the baseline system ($p = 0.035$). This difference corresponds to a small-to-medium effect size (AutoSketch vs. baseline: $d_z = 0.40$; BackSketch vs. baseline: $d_z = 0.21$), suggesting that the completion feature, which actively adds elements to a user's sketch, can make the interaction feel more like a collaborative partnership.

Creative product aspects (Q8–Q10): Scores on Originality (Q8), Elaboration (Q9), and overall satisfaction (Q10) were generally positive across all systems, with Elaboration showing the most benefit (BackSketch ($p = 0.072$); AutoSketch ($p = 0.083$)). The corresponding effect sizes for elaboration were in the small-to-medium range (BackSketch vs. baseline: $d_z = 0.34$; AutoSketch vs. baseline: $d_z = 0.33$), whereas originality and satisfaction showed only small effects (all $|d_z| \leq 0.24$). However, ultimately no significant differences were detected for product-focused questions.

5.1 Participant behaviour analysis

The action logs provide further insight into how participants engaged with each system. Behavioural metrics indicate that participants generated images at a similar rate across the three systems, averaging 1.55, 1.89, and 1.4 generations per drawing idea for the baseline, BackSketch, and AutoSketch, respectively. The metrics also show that the new interaction mechanisms in BackSketch and AutoSketch were actively adopted. On average, participants selected 1.57 suggested sketches per idea when using BackSketch. Similarly, participants invoked the completion feature in AutoSketch 1.9 times per idea, often more than once within a single drawing session. The undo action was also used regularly, averaging 0.7 per idea.

5.2 System preferences and qualitative feedback

At the end of the study we asked participants to select their overall favourite system from the three options and provide an optional explanation for their choice. The majority of participants chose AutoSketch (17 people; 57%), while ten chose BackSketch (33%) and three chose the baseline system (10%).

Qualitative feedback offers deeper insight into how participants experienced each system and what led to their preferences. For AutoSketch, many participants emphasised the value of the sketch completion feature, describing it as *"very clever," "interesting to see what it adds,"* and *"useful to help add details before painting."* Several noted that the AI's additions often inspired them to continue exploring new directions they *"would not have thought of"* on their own, supporting a more expressive and collaborative process. One participant reflected that the AI *"added detail that would have been fairly time consuming to add,"* while another highlighted how the completion undo function gave them *"the freedom to try unexpected ideas without risk."* In their preference justifications, participants frequently praised AutoSketch for providing *"the best balance of user creativity mixed with the AI,"* allowing them to refine sketches while still seeing surprising suggestions. Others valued that they could *"change my ideas while making them"* and *"never felt like I drew myself into a corner I couldn't get out of."* At the same time, some found the system's contributions overly assertive – such as *"adding random things I ended up deleting"* or producing outputs that felt more constrained by the AI's decisions than by their own.

Comments about BackSketch highlighted the suggestion gallery as an important source of inspiration. Participants described it as *"a good range of detail," "different sketch options,"* and *"a good starting point when making additions,"* which reduced the need to redraw from scratch. Several less confident sketchers recounted having *"options to work from which enabled me to be more creative."* Others noted that being able to select and refine AI-generated alternatives extended their expressive range: *"I felt I was able to make the output more detailed even when I had reached the end of my own creativity."* In their preferences, participants choosing BackSketch highlighted that they felt it was *"the most fun to see what it came up with,"* and that the gallery *"helped me achieve the best results, as it gave me a range of options to work from which enabled me to be more creative."* However, participants also pointed out limitations, such as the system *"putting its own ideas on what I was trying to communicate,"* or making small details hard to adjust once a suggestion was chosen.

By contrast, the baseline system was seen as both simpler and more restrictive, as might be expected. Supporters valued its straightforwardness, describing it as *"simple, less moving parts"* that let them focus on their own edits without interference. Some appreciated that the system *"captured what I was trying to draw quite well"* when their sketches were clear. A small number even preferred it overall, noting that it gave them *"more say in the product"* and allowed them to "focus on edits I want rather than predicting what AI would do." Yet many others criticised it for lacking opportunities for inspiration: *"there were less prompts as a result of no sketch option, meaning the results sometimes lacked inspiration."* Participants described struggling to depict specific concepts (e.g., *"it kept generating as an apple with a stalk even though I drew an ice cream with a flake"*). As one noted, *"capturing more detailed things would be better... like smile faces, sad faces."*

6 Discussion

While the study was exploratory, it provided pointers towards the timing, reversibility, and scope of AI interventions in sketch-based co-creation, to enable *iteration fluidity*. Such fluidity allows users to rapidly cycle between drawing, receiving system feedback and revising their intent without friction. This characteristic responds to design theory that emphasises re-framing, reflective conversation, and the co-evolution of problem and solution [13, 33].

6.1 Locating AI support in the creative loop

The results indicate that the timing of an AI intervention shapes creative exploration. AutoSketch and BackSketch do not merely offer different features; they reposition system initiative at distinct moments in the creative loop, with different consequences. AutoSketch intervenes before image generation, when intent is still forming. In this phase, participants used completions to overcome moments of uncertainty (*"what should I draw next?"*), treating AI additions as provisional scaffolding rather than commitments. This suggests that pre-generation support is particularly effective for idea elaboration: extending partial concepts, adding missing structure, and helping users move forward without fully specifying an outcome. Crucially, this benefit depended on strictly limited scope (one addition at a time) and easy reversibility, which kept completions from being interpreted as authoritative. BackSketch, by contrast, intervenes after generation, with users reacting to concrete outcomes. Qualitative feedback suggests this supported branching and redirection by reducing the need to redraw from scratch.

6.2 Integrating agency and constraint

Our results clarify that agency in mixed-initiative sketching is not simply about offering the user control, but about constraining system behaviour. Participants' responses suggest that both systems were perceived as collaborative when AI contributions remained negotiable. In AutoSketch, agency increased not because the system acted more, but because its actions were incremental, interpretable, and undoable. Participants tried completions speculatively, then accepted, edited, or rejected them. This pattern indicates that agency can emerge from trial without commitment, where users can probe system suggestions without handing-over authorship.

In BackSketch, agency took a different form. Selecting among alternative sketches shifted agency from fine-grained control to curatorial choice. While some participants noted loss of precision after adopting a suggestion, they still valued the ability to choose which interpretation to pursue. This suggests that a sense of agency can be preserved even when control is coarse-grained, provided users retain decisive selection power.

Importantly, tensions arose when system contributions appeared overly assertive or semantically misaligned. These moments underline a critical lesson for design: mixed-initiative systems must limit not only when they act, but how much they act. Richer or more autonomous AI behaviour might not necessarily enhance collaboration; in early-stage ideation, it may undermine it. These tensions are consistent with prior work showing that emergent AI behaviours can enable and complicate co-creative interaction [17].

6.3 Implications for research and design

From these observations, we suggest several implications that extend beyond the specific systems studied:

- Design AI support that is appropriate to the creative phase. Tools should explicitly target ideation stages (e.g., extension vs. reframing) rather than offering uniform assistance throughout the workflow.
- Make reversibility a first-class interaction primitive. Undo, replaceable alternatives, and low-cost rejection are not auxiliary usability features; they are essential for sustaining agency in co-creative sketching systems.
- Constrain system initiative to preserve interpretability. Incremental, legible contributions are more compatible with sketch-based ideation than holistic or high-confidence outputs, particularly, perhaps, for non-expert users.
- Evaluate collaboration through a process lens, not by outcomes alone. Our behavioural logs revealed meaningful engagement patterns even where product-level differences were modest, reminding us that process-oriented metrics are important for assessing co-creative tools.

6.4 Limitations and future research directions

While we have highlighted the few statistically significant differences, the majority of measures showed no reliable differences between systems, and these null results should be interpreted as equally important in understanding the boundaries of our findings. Our study focused on short, exploratory sketching sessions with non-expert users, which allowed us to isolate early-stage creative dynamics but limits generalisability. Longer-term studies are needed to examine how preferences and strategies evolve over time, particularly as users develop trust—or fatigue—with AI interventions. Similarly, professional illustrators may require different balances of initiative and constraint, especially in later-stage refinement tasks. Future work could therefore incorporate multi-session or longitudinal evaluations, as well as studies involving expert illustrators or designers, to understand how these mechanisms integrate into more sustained or professional workflows.

Additionally, while our sketch-only interface captured a focused interaction style, it likely does not reflect the increasingly multimodal nature of real-world creative practice, where sketching often interacts with text prompts, references or style specifications. Exploring hybrid sketch-text pipelines, or systems that adapt suggestion or completion assertiveness based on user behaviour, is a clear next step.

7 Conclusions

In this paper, we presented two novel sketch-to-image interaction techniques, AutoSketch and BackSketch, and compared them with a baseline system that represents the current norm in sketch-to-image system designs. Across quantitative ratings, behavioural logs and qualitative feedback, both approaches enhanced creativity support without reducing usability or enjoyment. Participants rated them higher for *exploration* and *expressiveness*, with AutoSketch also boosting perceived *agency*. They actively engaged with the new features, regularly selecting suggestions in BackSketch and invoking completions with undo in AutoSketch. Qualitative feedback highlighted that suggestions provided inspiration and variety, while completions fostered collaboration and surprise, tempered by the ability to undo.

Beyond these specific systems, as detailed in Section 6.2, the findings highlight a broader design space for co-creative sketch-to-image tools. These principles matter because they lower the cost of creative iteration. Reversible contributions let users test ideas without risk, preserving a sense of ownership over the evolving sketch. Shifting system support across the workflow opens different entry points for exploration, allowing users to extend ideas before generation or reframe them afterward. Balancing user and AI initiative ensures that assistance feels supportive rather than intrusive, enabling a collaborative dynamic in which the system contributes possibilities while users retain control.

Taken together, these insights position pre- and post-generation support as complementary strategies for advancing sketch-to-image interaction. Looking ahead, these ideas could be applied beyond sketch-to-image systems—for example in mixed sketch-and-text workflows or adaptive tools that adjust their support over time—and could benefit creative domains such as layout design, animation, and conceptual illustration.

Acknowledgments

We thank study participants for their contribution to this research. This work was supported by Engineering and Physical Sciences Research Council grant EP/Y010477/1.

References

[1] Rameen Abdal, Yipeng Qin, and Peter Wonka. 2019. Image2StyleGAN: How to Embed Images Into the StyleGAN Latent Space?. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, 4432–4441.

[2] Barrett R Anderson, Jash Hemant Shah, and Max Kreminski. 2024. Evaluating Creativity Support Tools via Homogenization Analysis. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '24)*. Association for Computing Machinery, New York, NY, USA, 1–7. doi:10.1145/3613905.3651088

[3] David Bau, Hendrik Strobelt, William Peebles, Jonas Wulff, Bolei Zhou, Jun-Yan Zhu, and Antonio Torralba. 2019. Semantic Photo Manipulation with a Generative Image Prior. *ACM Trans. Graph.* 38, 4 (July 2019), 59:1–59:11. doi:10.1145/3306346.3323023

[4] Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023. Promptify: Text-to-Image Generation through Interactive Prompt Exploration with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–14. doi:10.1145/3586183.3606725

[5] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. 2023. InstructPix2Pix: Learning to Follow Image Editing Instructions. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, Vancouver, BC, Canada, 18392–18402. doi:10.1109/CVPR52729.2023.01764

[6] Bill Buxton. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

[7] John Canny. 1986. A Computational Approach to Edge Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* PAMI-8, 6 (Nov. 1986), 679–698. doi:10.1109/TPAMI.1986.4767851

[8] Eric Nai-Li Chen, Joshua Kong Yang, Jeff Huang, and Tongyu Zhou. 2025. L.Ink: Procedural Ink Growth for Controllable Surprise. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3746059.3747702

[9] Shu-Yu Chen, Wanchao Su, Lin Gao, Shihong Xia, and Hongbo Fu. 2020. DeepFaceDrawing: Deep Generation of Face Images from Sketches. *ACM Trans. Graph.* 39, 4 (Aug. 2020), 72:72:1–72:72:16. doi:10.1145/3386569.3392386

[10] Wengling Chen and James Hays. 2018. SketchyGAN: Towards Diverse and Realistic Sketch to Image Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 9416–9425.

[11] DaEun Choi, Sumin Hong, Jeongeon Park, John Joon Young Chung, and Juho Kim. 2024. CreativeConnect: Supporting Reference Recombination for Graphic Design Ideation with Generative AI. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, 1–25. doi:10.1145/3613904.3642794

[12] Richard Lee Davis, Kevin Fred Mwaita, Livia Müller, Daniel C. Tozadore, Aleksandra Novikova, Tanja Käser, and Thiemo Wambsganss. 2025. SketchAI: A "Sketch-First" Approach to Incorporating Generative AI into Fashion Design. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–7. doi:10.1145/3706599.3719782

[13] Kees Dorst. 2015. *Frame Innovation: Create New Thinking by Design.* The MIT Press, Cambridge, Massachusetts.

[14] Noyan Evirgen and Xiang 'Anthony Chen. 2023. GANravel: User-Driven Direction Disentanglement in Generative Adversarial Networks. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA, 1–15. doi:10.1145/3544548.3581226

[15] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-or. 2022. Prompt-to-Prompt Image Editing with Cross-Attention Control. In *The Eleventh International Conference on Learning Representations*. ICLR, 19 pages.

[16] Rahul Jain, Amit Goel, Koichiro Niinuma, and Aakar Gupta. 2025. AdaptiveSliders: User-aligned Semantic Slider-based Editing of Text-to-Image Model Output. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–27. doi:10.1145/3706598.3714292

[17] Tomas Lawton, Francisco J Ibarrola, Dan Ventura, and Kazjon Grace. 2023. Drawing with Reframer: Emergence and Control in Co-Creative AI. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. ACM, Sydney NSW Australia, 264–277. doi:10.1145/3581641.3584095

[18] Lei Li, Changqing Zou, Youyi Zheng, Qingkun Su, Hongbo Fu, and Chiew-Lan Tai. 2021. Sketch-R2CNN: An RNN-Rasterization-CNN Architecture for Vector Sketch Recognition. *IEEE Transactions on Visualization and Computer Graphics* 27, 9 (Sept. 2021), 3745–3754. doi:10.1109/TVCG.2020.2987626

[19] David Chuan-En Lin, Hyeonsu B Kang, Nikolas Martelaro, Aniket Kittur, Yan-Ying Chen, and Matthew K. Hong. 2024. Inkspire: Sketching Product Designs with AI. In *The 37th Annual ACM Symposium on User Interface Software and Technology*. ACM, Pittsburgh PA USA, 1–6. doi:10.1145/3672539.3686339

[20] Haichuan Lin, Yilin Ye, Jiazhi Xia, and Wei Zeng. 2025. SketchFlex: Facilitating Spatial-Semantic Coherence in Text-to-Image Generation with Region-Based Sketches. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–19. doi:10.1145/3706598.3713801

[21] Vivian Liu and Lydia B Chilton. 2022. Design Guidelines for Prompt Engineering Text-to-Image Generative Models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI '22)*. Association for Computing Machinery, New York, NY, USA, 1–23. doi:10.1145/3491102.3501825

[22] Vivian Liu, Han Qiao, and Lydia Chilton. 2022. Opal: Multimodal Image Generation for News Illustration. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology (UIST '22)*. Association for Computing Machinery, New York, NY, USA, 1–17. doi:10.1145/3526113.3545621

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

[24] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. 1997. Design Galleries: A General Approach to Setting Parameters for Computer Graphics and Animation. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., USA, 389–400. doi:10.1145/258734.258887

[25] Nicolai Marquardt, Asta Roseway, Hugo Romat, Payod Panda, Michel Pahud, Gonzalo Ramos, Steven M. Drucker, Andrew D. Wilson, Ken Hinckley, and Nathalie Riche. 2025. ImaginationVellum: Generative-AI Ideation Canvas with Spatial Prompts, Generative Strokes, and Ideation History. In *Proceedings of the 38th Annual ACM Symposium on User Interface Software and Technology (UIST '25)*. Association for Computing Machinery, New York, NY, USA, 1–19. doi:10.1145/3746059.3747631

[26] Justin Matejka, Michael Glueck, Erin Bradner, Ali Hashemi, Tovi Grossman, and George Fitzmaurice. 2018. Dream Lens: Exploration and Visualization of Large-Scale Generative Design Datasets. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. Association for Computing Machinery, New York, NY, USA, 1–12. doi:10.1145/3173574.3173943

[27] Chong Mou, Xintao Wang, Liangbin Xie, Yanze Wu, Jian Zhang, Zhongang Qi, and Ying Shan. 2024. T2I-Adapter: Learning Adapters to Dig out More Controllable Ability for Text-to-Image Diffusion Models. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence (AAAI'24/IAAI'24/EAAI'24, Vol. 38)*. AAAI Press, 4296–4304. doi:10.1609/aaai.v38i5.28226

[28] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, New Orleans, LA, USA, 10674–10685. doi:10.1109/CVPR52688.2022.01042

[29] Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. 2016. The Sketchy Database: Learning to Retrieve Badly Drawn Bunnies. *ACM Trans. Graph.* 35, 4 (July 2016), 119:1–119:12. doi:10.1145/2897824.2925954

[30] Patsorn Sangkloy, Jingwan Lu, Chen Fang, Fisher Yu, and James Hays. 2017. Scribbler: Controlling Deep Image Synthesis with Sketch and Color. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 6836–6845. doi:10.1109/CVPR.2017.723

[31] Vishnu Sarukkai, Lu Yuan, Mia Tang, Maneesh Agrawala, and Kayvon Fatahalian. 2024. Block and Detail: Scaffolding Sketch-to-Image Generation. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology*. ACM, Pittsburgh PA USA, 1–13. doi:10.1145/3654777.3676444

[32] R. Keith Sawyer. 2021. The iterative and improvisational nature of the creative process. *Journal of Creativity* 31 (2021), 100002. doi:10.1016/j.yjoc.2021.100002

[33] Donald A. Schön. 1983. *The Reflective Practitioner: How Professionals Think in Action.* Basic Books, New York.

[34] Xinyu Shi, Shunan Guo, Jane Hoffswell, Gromit Yeuk-Yin Chan, Victor S. Bursztyn, Jian Zhao, and Eunyee Koh. 2025. Comprehensive Sketching: Exploring Infographic Design Alternatives in Parallel. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. ACM, Yokohama Japan, 1–8. doi:10.1145/3706599.3720182

[35] Ben Shneiderman. 2007. Creativity Support Tools: Accelerating Discovery and Innovation. *Commun. ACM* 50, 12 (Dec. 2007), 20–32. doi:10.1145/1323688.1323689

[36] Edgar Simo-Serra, Satoshi Iizuka, Kazuma Sasaki, and Hiroshi Ishikawa. 2016. Learning to Simplify: Fully Convolutional Networks for Rough Sketch Cleanup. *ACM Trans. Graph.* 35, 4 (July 2016), 121:1–121:11. doi:10.1145/2897824.2925972

[37] Loutfouz Zaman, Wolfgang Stuerzlinger, Christian Neugebauer, Rob Woodbury, Maher Elkhaldi, Naghmi Shireen, and Michael Terry. 2015. GEM-NI: A System for Creating and Managing Alternatives In Generative Design. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. Association for Computing Machinery, New York, NY, USA, 1201–1210. doi:10.1145/2702123.2702398

[38] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. 2023. Adding Conditional Control to Text-to-Image Diffusion Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. IEEE, 3836–3847.

A Likert response distributions

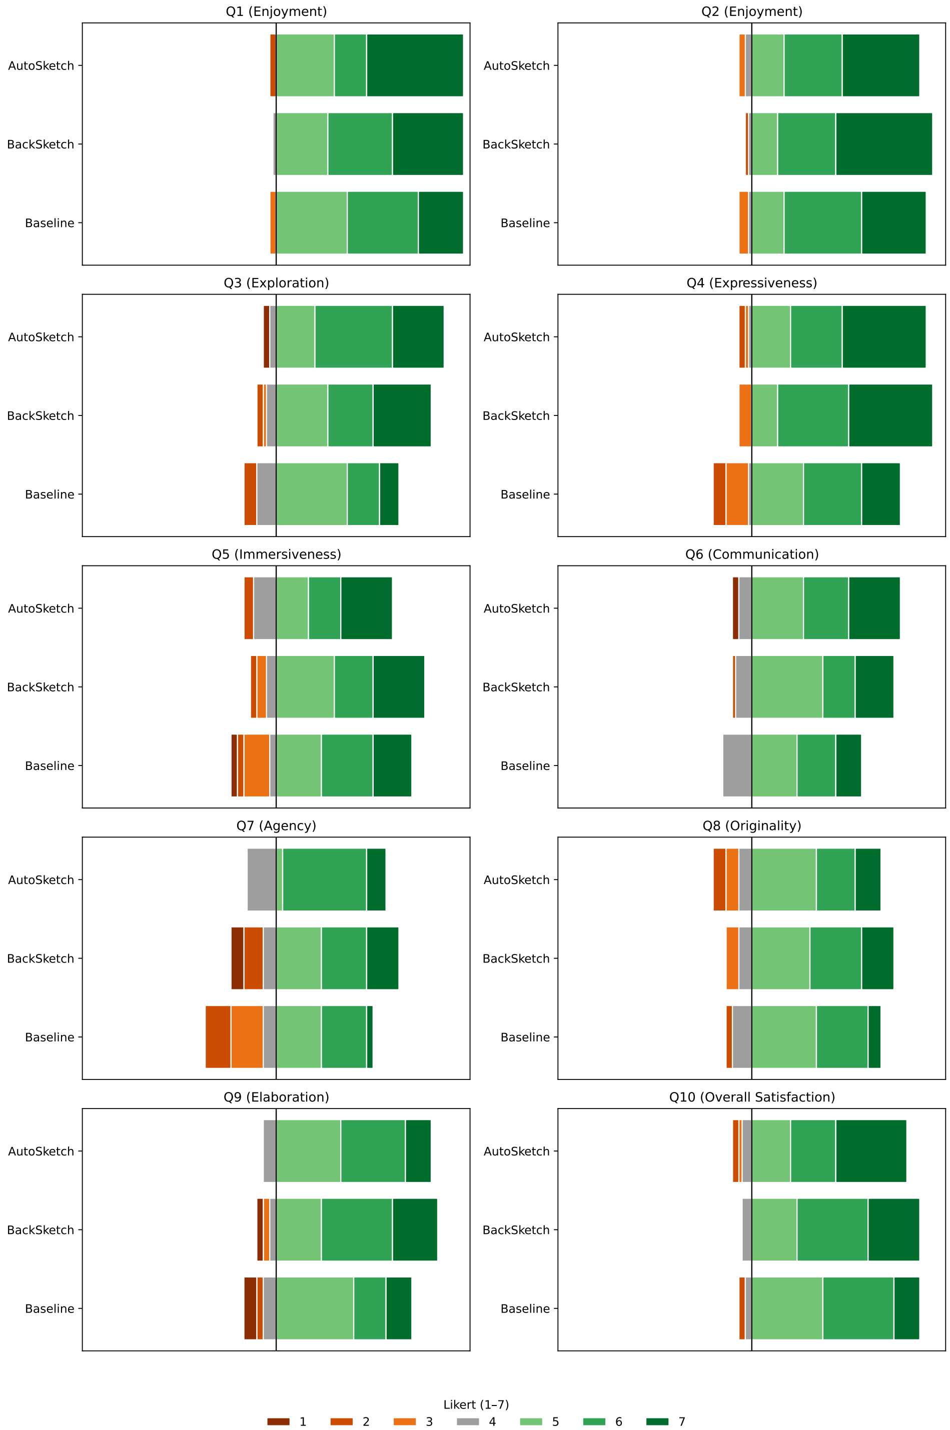

Figure 4: Diverging stacked bar charts showing the distribution of Likert responses (1–7) for all questionnaire items (Q1–Q10) across the three systems (Baseline, BackSketch, AutoSketch). Bars to the left of the centre line indicate more negative ratings, and bars to the right indicate more positive ratings.