MoMix: Sketches as Mutable Anchors for Navigating and Experimenting with Design Materials on Infinite Canvases

ANONYMOUS AUTHOR(S)

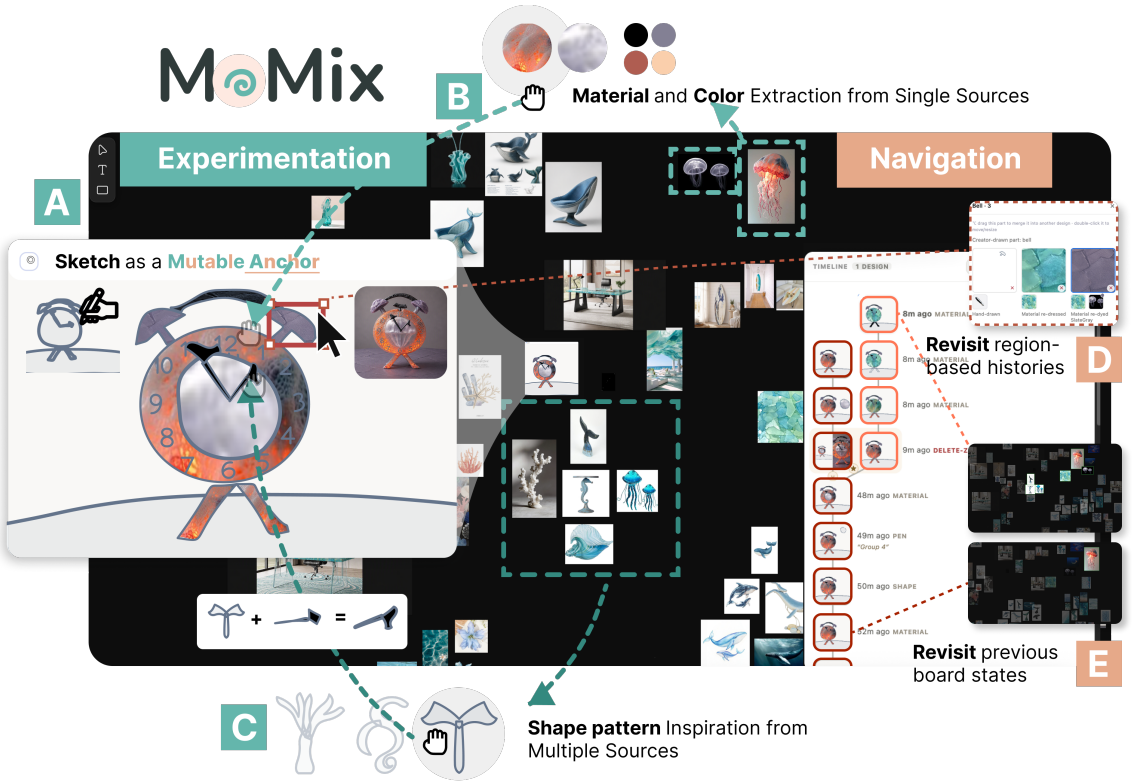


Fig. 1. *MoMix* treats a freehand sketch as a mutable anchor on an infinite canvas. **Experimentation**: (A) properties attach to sketch parts and stay editable; (B) material and color and (C) shape patterns are extracted from one or several references. **Navigation**: (D) a part's region-based history, and (E) a branching timeline of board states.

Infinite canvas systems let designers organize ideas before committing resources. Designers often test those ideas by sketching, prototyping concrete possibilities from visual references. To make experimentation with design materials tangible, we built a system that treats the designer's own sketches as mutable anchors. *MoMix* turns freehand sketches into editable parts for applying color, material, and shape from references, while region-based histories and a branching timeline trace provenance across design states. In a within-subjects study with 12 designers, participants rated *MoMix* outcomes as more novel, described greater cognitive flexibility in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© 2027 Copyright held by the owner/author(s). Publication rights licensed to ACM.
Manuscript submitted to ACM

developing an idea, reported a stronger sense of ownership over the outcome, and were more confident both in communicating their own design and in continuing a design left off at an earlier stage. Our work shows how mixing visual references with sketches may capture the way people work flexibly with design ideas.

CCS Concepts: • **Human-centered computing → Human computer interaction (HCI)**.

Additional Key Words and Phrases: sketching, design references, moodboards, infinite canvas, generative AI, creativity support, provenance, design history

ACM Reference Format:
Anonymous Author(s). 2027. MoMix: Sketches as Mutable Anchors for Navigating and Experimenting with Design Materials on Infinite Canvases. In *CHI Conference on Human Factors in Computing Systems (CHI '27), May 10–14, 2027, Pittsburgh, PA, USA.* ACM, New York, NY, USA, 31 pages. https://doi.org/10.1145/XXXXXXX.XXXXXXX

1 Introduction

Visual design is often akin to mosaic making. Instead of starting from *tabula rasa*, the designer usually draws from multiple visual references, either physically from their repertoire of moodboards or mentally from their past experiences, to recombine properties across them into new possibilities [10, 26, 33]. Infinite canvases extend this practice into a spatial workspace where design materials such as visual references, sketches, and emerging alternatives can be arranged together during visual ideation [39]. Within this space, sketching does more than record an idea already formed: designers additionally use it to assemble partial imagery, re-interpret forms, and restructure them as new possibilities emerge [20, 48, 53]. The incorporation of generative AI (GenAI) in recent creative workflows now makes these visual references and alternatives much cheaper to produce, causing infinite canvases to grow rapidly [1, 36, 51, 52]. As visual references accumulate, two challenges become progressively important to address: how can designers *navigate* large collections of references and trace how they shape an evolving design over time, and how can they *experiment* with different properties of those references to update and refine their designs?

By *navigation* we mean not only finding and revisiting useful references on a canvas, but also tracing how those references shape an evolving design. During authoring, designers revisit earlier states and alternatives to evaluate current choices, recover promising directions, and continue exploration without starting over [23, 39, 41]. Communication makes this harder still, because a moodboarding canvas encodes meaning implicitly: designers and their collaborators who return to the board later can inspect the visual references but not the role each one played, such as which idea it informed or the rationale behind a given direction [23, 31, 38, 40]. Existing provenance systems can record version history and design alternatives [4, 21, 22, 39, 41], or organize large canvases for browsing [16, 51, 52]. Both record that an asset came from a reference, but neither records what the reference contributed or which part of the design it changed.

Experimentation refers to how designers draw on different properties of references, for example color, material, and style, to update or enhance their designs. Recent systems have made these properties portable by abstracting them as visual tokens [42, 49] and pen- or region-based transfer mechanisms [43, 46]. However, these interactions still largely resolve into pixel-based generation: the properties of a reference become manipulable, but the design they produce does not itself stay *mutable*. Its parts do not persist as editable geometry that designers can continuously reshape, resize, regroup, or recombine. This is particularly limiting for shape, whose transformation is itself central to exploratory design [43, 44]. More broadly, existing systems provide limited functionality for continuously prototyping how reference properties transform an evolving design while keeping the design itself editable.

To address both of these gaps, we introduce *Mosaic Mixing (MoMix)*, an infinite-canvas system that keeps the designer's own sketch as a *mutable anchor* for *both* experimentation and navigation. We use the sketch as the design object itself because its parts naturally remain open to restructuring and different readings, supporting movement between ideas without requiring early commitment to a fixed representation [19, 55–57]. Rather than consuming the sketch as an input to produce a new image, *MoMix* keeps it persistent throughout the process: the anchor is *assembled* from the designer's own editable parts, and supports *polymorphic* transformation, in which different properties can reshape different parts while they remain editable. It also carries a *shared design imaginary*: not only the current design, but the references, alternatives, and contributions through which it developed, which the designer can revisit and collaborators can understand. *MoMix* realizes the mutable anchor by planarizing freehand strokes into persistent, addressable Scalable Vector Graphics (SVG) parts. For *experimentation*, designers can extract and apply color, material, and shape properties from individual or grouped references to selected sketch parts, while keeping the resulting parts editable. For *navigation*, region-based histories preserve the reference-driven transformations of individual parts, while a branching timeline and playback show how the overall design evolved across stages and which references informed each state.

We conducted a mixed-methods, within-subjects study with 12 designers comparing *MoMix* with a GenAI-integrated infinite-canvas baseline. We found that designer-initiated sketches introduced variation beyond model priors, contributing to more novel outcomes. As a mutable anchor, the sketch's polymorphic structure supported cognitive flexibility through part-level reshaping and reference recombination, encouraging deeper exploration before commitment. Participants also developed ownership not only over the final output but also over the process of constructing it, likening the experience to building with Lego. Finally, provenance histories turned the evolving sketch into a shared design imaginary for understanding and reflecting on how ideas develop over time and across collaborators.

In summary, this work contributes: (1) a formative study with eight experienced designers on how they navigate, experiment with, and communicate visual references on infinite canvases when GenAI enters their workflows; (2) *MoMix*, an infinite-canvas system that uses the designer's own sketches as mutable anchors to support polymorphic, part-level reference experimentation and provenance navigation; and (3) a within-subjects study with 12 designers showing what mutable anchors afford for experimentation and navigation. Together, these contributions show how a mosaic-like sketch-based anchor can help designers navigate references during exploration of design ideas. In doing so, we hope to capture how people work flexibly with the design ideas they accumulate.

2 Related Work

2.1 Sketching as a Cognitive Process for Early-Stage Ideation

Early-stage design differs from later refinement because the designer reasons about an artifact that does not exist. Because it has never been perceived, the designer cannot picture the design and must instead build it up from partial images of tentative aspects [20]. These image references are rarely invented but are borrowed from existing artifacts or carried over from analogous cases [10, 20]. Sketching is where this building up happens. The dialectics of sketching frames a sketch not as the representation of an image already held in mind but as the search for one. Designers draw and redraw until something meaningful can be read off the page, a process Goldschmidt terms *interactive imagery*, in which producing a visual display and reading a new one off it happen almost at the same time [20]. This rereading leads to serendipitous discovery through a recursive cycle of seeing, moving, and seeing again [48, 53]. In action, designers

impose imaginary lines that carve an undifferentiated shape into recognizable regions, then shift their attention to the regions they notice [20, 54].

Experiments on creative mental synthesis showed that a sketch plays two roles beyond holding the image [15, 56]—*combining* and *restructuring*. Combining moves parts around while each part stays whole, and people do it in their minds as well as on paper. Restructuring reads a different set of parts out of the same marks, and it needs a sketch [56, 57]. Goel's protocol studies describe sketches as supporting *lateral transformations*, in which one idea becomes a different idea because freehand marks are dense and open to more than one reading, leaving most attributes unspecified. Representations that already look like the final result allow only vertical refinement, in which the designer refines an idea that is already fixed [19]. Similarly, viewers who attended to different parts of an ambiguous drawing, or mentally rearranged them, produced nearly twice as many interpretations [55], and ambiguity in early representations has long been argued to protect against premature commitment to a first idea [28].

Taken together, these works describe the sketch as the medium of a distinctly material process: partial images borrowed from references are assembled on it, reread part by part, and restructured into ideas none of them contained. Its flexibility depends on the sketch remaining ambiguous, partitioned, and open to change.

2.2 Experimenting with Visual References

Designers rarely begin with a blank page. Early ideas develop through working back and forth with sketches, examples, and other reference materials [26, 33]. This process is less the execution of a fixed plan than what Schön calls a reflective conversation with the materials of a design situation, where each move makes the material talk back [47]. References, on this view, are not passive stimuli but epistemic objects: incomplete and question-generating materials whose reworking opens further possibilities [14].

Tools have accordingly invested in making a reference's properties portable—extractable from where they were found and applicable somewhere else. One line makes them portable in thought: SemanticCollage labels images so that vague impressions become searchable terms [31], *May AI?* recommends and explains references drawn from a designer's selections [29], GANCollage samples a GAN's latent space to reach beyond collected images [59], and CreativeConnect decomposes references into subject matter, action and pose, theme and mood, and arrangement, then recombines them into new candidates [10]. Another makes them portable in the hand: Brickify reifies reference images as tokens a designer can resize, group, and link [49]; DesignPrompt composes multimodal prompts from example images, colors, and words, steered by adjusting their mix rather than by precise wording [42]; FusAIn loads objects, colors, and textures into pens whose strokes place those properties onto a canvas [43]; and AI-Instruments goes furthest, allowing a brush to pick up content or style from one image and apply it to a segmented region of another [46]. Yet these systems ultimately produce pixel-based designs, where part structure is flattened into the image and local changes rely on further generation through prompting rather than direct manipulation of persistent, addressable geometry. This is especially limiting for shape, which cannot be transferred to a part and then directly reshaped, resized, or recombined as geometry [43].

The sketch itself could also work as a working design object. It has low fidelity, and the designer builds it up one stroke at a time, so ideas stay cheap to change and the drawing still keeps the designer's own hand. Guided-sketching tools support these strokes rather than replace them. The Drawing Assistant gives the user construction-line scaffolds and feedback on how closely their strokes follow those scaffolds, so the user can refine their own drawing [27]. PortraitSketch adjusts the shape and look of each stroke while the user draws it, but every final stroke is still one the user drew [61]. SketchSoup morphs whole drawings into in-between alternatives, and the designer can draw over those alternatives

and feed them back into the design [2]. Inkspire takes a generated design and turns it back into a rough low-fidelity sketch, so each new stroke can steer another round of generation [34]. In all of these tools, however, the designer can change only individual strokes or the sketch as a whole. They cannot point directly to a meaningful part of the drawing, nor to a single attribute taken from a reference. Our work instead makes the designer's sketch a *mutable anchor*. We assemble the anchor from the designer's strokes, and we planarize it into addressable parts, meaning we turn the strokes into separate regions that the designer can point to. This divides the drawing into flat, closed regions along its lines. The designer then transforms the sketch by binding reference attributes directly to any part, at different levels of detail and abstraction, without replacing the sketch itself.

2.3 Navigating Provenance on Infinite Canvases

Creativity tools tend to preserve what a designer makes but not how they made it. Recent work refers to this as *artifact bias*: tools emphasize finished outputs while the creation process stays invisible, and a growing line of research argues for treating the trace of that process as a first-class design material instead [23]. Systems have recorded the trace at two scales. Within a single creation, Chronicle captures a video history of an image-editing session and links finished content back to the operations that produced it [21], and HistoryPalette organizes generations into palettes by position, time, and prompt theme as reusable alternatives [4]. Across a canvas, the record instead spans many pieces and traces how each one derives from or branches off what came before: ImaginationVellum time-indexes every canvas action and offers temporal replay alongside a provenance visualization over generated results [39]; DesignTrace records how alternatives branch under shared semantic attributes, each node keeping its own edit history [41]; and Vistoria preserves the source of each multimodal card [16]. Outside image generation, Spatial Balancing shows the same idea in writing: it lays revision strategies out on a two-dimensional canvas so that writers can reason about revision as a trajectory rather than as a sequence of isolated edits [17]. Atelier traces every generated asset to the references, parameters, and outputs it descends from, weaving a navigable lineage into the canvas [22]. All of these tools record a link between an asset and the references it came from. Such a link identifies that two things are related, but not which property entered the design or where it landed.

The volume of references that GenAI makes it cheap to produce turns this limitation into a practical barrier. For example, tools such as Adobe Firefly Boards[1] and Lovart[2] now fold collection, generation, and remixing into a single canvas [1, 36]. Each new item brings its own links, so the record grows denser without growing more informative. And a link only helps someone who already knows what to look for—which is what a crowded board takes away. For the research community, the answer has been to impose structure after the fact, through semantic zoom [52], interpretable dimensions [51], or grouping by shared elements [16]. These approaches organize outputs as browsable items, compressing what a board holds rather than recording what it did. Accumulation also creates a second problem. AI support during ideation fixates designers on early examples and narrows what they produce [58], while AI outputs drift toward a shared center, reducing the collective diversity of what is produced [12].

Our work considers history at two complementary levels: region-based histories that connect references and prior states to the specific design parts they shaped, and a branching timeline that preserves the broader relationships among references, alternatives, and versions across the exploration process. Together, these views turn design history into a shared interpretive space that a teammate, or a designer's own later self, can re-enter and build on.

[1] Adobe Firefly Boards, https://www.adobe.com/products/firefly/features/boards.html, accessed 2026-08-27.
[2] Lovart, https://www.lovart.ai/, accessed 2026-08-27.

Study	ID	Experience	Design Expertise
Interview	F1	11 yrs	Architecture, Product, UI/UX
	F2	8 yrs	Graphic, Visual, 3D/Motion
	F3	9 yrs	Industrial, Product, UI/UX
	F4	8 yrs	Product, Graphic, 3D/Motion
Observation Group 1	F5	5 yrs	UI/UX, Game, Interaction
	F6	4 yrs	UI/UX, Media, Game
Observation Group 2	F7	5 yrs	UI/UX, Product, Architecture
	F8	7 yrs	Illustration, UI/UX

Table 1. Background of participants in the formative studies.

3 Formative Study

3.1 Participants and Procedure

To understand how designers organically work on infinite canvases, particularly how they navigate and experiment with large collections of visual references with the aid of GenAI, we conducted a formative investigation via four semi-structured interviews (F1–F4) and two collaborative observational studies (F5–F8) with experienced designers. Participant demographics are summarized in **Table 1**. The interviews focused on participants' accounts of how they resolve challenges, communicate evolving ideas, and feel about the integration of GenAI within existing practices. The two observational studies, conducted with designer pairs (G1: F5–F6 and G2: F7–F8), examined how the presence of a collaborator introduced additional challenges to reference navigation and experimentation during early-stage ideation. Participants were first given 20 minutes to collaboratively explore an initial design concept for a vase based on a researcher-provided brief by collecting 20 references. They then brought these materials into Figma to sketch, annotate, and organize references freely in a shared infinite-canvas workspace. After exchanging emerging ideas, they jointly developed a design direction for 30 minutes.

We used inductive qualitative analysis on the interview and observational data, informed by grounded-theory techniques through open coding and constant comparison [8]. We iteratively compared codes across participants and data sources, grouped them into higher-level themes, and interpreted these themes alongside prior work on moodboarding and reference-based design [23, 30, 31, 38, 40, 42, 43, 49] to derive design insights and goals for *MoMix*.

3.2 Design Insights

3.2.1 Using Sketches to Externalize Emerging Ideas. During early ideation, instead of finding examples of objects they intend to design, participants perused various materials, forms, and visual categories to use as references for conceptualization (F3, F4). F3 described this process as progressively identifying meaningful relationships while comparing the references' visual properties and recurring characteristics that could be recombined into an emerging concept. These connections were not always semantically obvious, as F4 mentioned, *"references from different categories, or even seemingly random and absurd combinations, could reveal unexpected possibilities that would not emerge from searching within a single object type."* This is consistent with existing literature on moodboards, describing them as deliberately ambiguous spaces where meaning is constructed through comparison, juxtaposition, abstraction, and reinterpretation across references [38]. As possibilities began to emerge, participants then used *low-fidelity sketches* to give partial ideas an initial visual form without committing to a complete solution (F1, F2, F3, F4). These findings

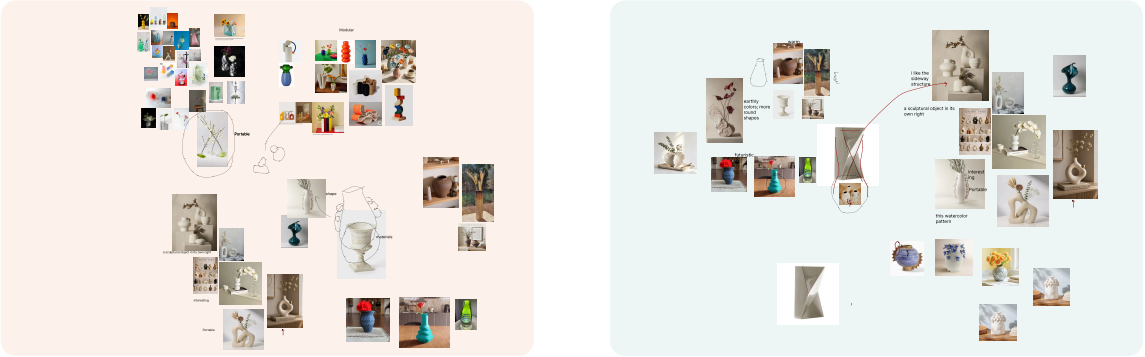


Fig. 2. Boards from the formative observation. Both groups (G1, G2) arranged references around emerging ideas and used sketches, lines, and arrows to show which reference informed what, yet the relationships stayed ambiguous: neither authorship nor the order of events is recoverable.

suggest that concept formulation is characterized by both reference exploration and representation. Sketches bridge the two by externalizing interpretations while preserving room for further exploration.

3.2.2 Balancing Local Control with Design Coherence. Once an initial concept emerges, participants' relationships with references shifted from *"something to look at"* towards *"something to try with."* Rather than treating each image as an indivisible whole, designers picked out visual dimensions to transfer into evolving designs (F1, F2, F3, F4). F2 described annotating references to indicate color, materials, or shapes she wanted to reuse. F3 similarly noted that *"Shape was especially prominent during early exploration, while color, material, and style were considered progressively as the design became more concrete."* Our group observations showed the same pattern: participants sketched beside references and used arrows or annotations to indicate how specific reference properties could inform sketch parts **(Figure 2)**, consistent with prior work on selectively extracting and recombining reference properties [42, 43, 49]. Such experimentation was highly iterative, as participants tried different combinations and local variations before committing (F1, F2, F3, F4). In contrast, existing GenAI workflows often modified the entire image rather than local components. Even when designers specified local changes through prompts, regenerating one component could unexpectedly alter previously accepted regions (F1, F3); as F4 put it, *"After modifying the image ten times, it becomes deformed. It has a really high communication cost."* This made local trial and error difficult without repeatedly re-negotiating the rest of the design to preserve their original intent (F2, F3, F4, F8). On the other hand, while sketches kept the evolving design grounded in the designers' own decisions, they were difficult to reuse or recombine once drawn. F3 noted that combining two sketch ideas, or integrating her sketch with another designer's, often required redrawing the design, making rapid experimentation with alternative combinations cumbersome. These findings suggest that flexible reference experimentation requires local control while keeping the evolving design coherent as a whole, motivating a mutable intermediate representation that accommodates direct, part-level experimentation.

3.2.3 Tracing Holistic Design Histories for Experimentation. Designers often organized references and sketches spatially around different ideas while leaving their relationships implicit. However, after a while, F1 noted *"I sometimes no longer remember what each version represents."* Prior work similarly shows that designers may struggle to recover why earlier

images, labels, or clusters mattered [31]. This ambiguity arises because spatial arrangements express relationships only implicitly [40], so the same arrangement admits multiple interpretations [38]. The problem compounds as the board grows, since designers must recover why a reference mattered and how it related to a particular design direction or state. Thus, participants valued preserving *holistic* histories. As F2 explained, *"I want to save every step of the exploration, so I can see what each step looked like and return to any previous step if an earlier version worked better."* These findings suggest that spatial exploration should keep its productive ambiguity while its history stays traceable.

3.2.4 Communicating Exploration Beyond the Moment It Happens. In G1, participants used keywords and spatial clusters to mark promising directions. When revisiting these materials, F5 and F6 could reconstruct *their own* intentions but struggled to infer what *the other* designer intended to pursue, even when those ideas had been discussed before. This observation aligns with prior work on the ambiguity of visual labels [31, 38], suggesting a need to make connections between references and design intentions more explicit to trace. Moreover, communication did not capture the full picture (F3, F4) of why references were selected and the directions encountered along the way. Even unselected references could inspire another designer, yet this potential disappeared when only the resulting board was shared. These gaps are particularly consequential in asynchronous workflows that lack the discussion [38] to convey the creative process and evolving intent that drove the final design [23]. As a result, the provenance of emerging ideas could become unclear, weakening participants' sense of authorship and ownership over their contributions. These findings suggest that collaborative reference work needs traces legible enough to carry design intent when the board is revisited without the discussion that produced it.

3.3 Design Goals

Drawing on the design insights from our formative study, we identified three design goals for an interactive interface that supports reference-driven experimentation and navigation on infinite canvases with large collections of references.

DG1: Assembling Emerging Design Ideas while Preserving Designer Intent and Control. Grounded in Sections 3.2.1 and 3.2.2, our system should support designers in progressively assembling and refining emerging ideas across large and diverse collections of explorable references, without prematurely committing to a design concept or disrupting decisions already made. Sketches offer particular potential as provisional, designer-authored intermediate representations that support this assembly: unlike generated images, they originate from designers' own actions and can give partial ideas form while remaining open to restructuring and reinterpretation.

DG2: Enabling Polymorphic Experimentation with Reference Properties. Informed by Section 3.2.2, our system should enable designers to rapidly try, combine, and compare visual properties such as shape, color, and material across different parts and granularities of an evolving design. Such polymorphic experimentation should remain visually represented and directly inspectable, providing immediate feedback for evaluating alternatives and making design decisions.

DG3: Sharing a Design Imaginary through Legible Visual Traces. Grounded in Sections 3.2.3 and 3.2.4, our system should preserve, in a visible intermediate form, how references inspire and shape evolving design directions, helping designers better externalize and record their ideas while reducing ambiguity in communicating design intent. These traces should form a shared design imaginary that a designer can re-enter later, whether it is their own board or one handed over, by keeping individual exploration and authorship distinguishable within a shared workspace and making reference-to-design-idea relationships understandable to someone who was not there when they were made.

4 *MoMix* System

In this section, we present *MoMix*, an infinite-canvas system that uses the designer's sketch as a mutable anchor for reference-driven experimentation and navigation. Rather than consuming the sketch as a one-off input to image generation, *MoMix* keeps it on the canvas as the design object itself. Designers place a sketch card among visual references, which turns freehand strokes into named, fillable parts. They can extract color, material, and shape from references and apply them to individual sketch parts. Clicking a part opens its region-based history, while a timeline records how the overall design evolved.

4.1 Key Features

4.1.1 Sketch as a Mutable Anchor for Assembling Emerging Design Ideas. During early-stage exploration, designers work across broad, ambiguous collections of references while their ideas remain partial and uncertain (**DG1**). *MoMix* therefore treats the sketch as a *mutable anchor*: a persistent, designer-authored object whose editable parts represent the evolving design itself. Unlike a generated image, the anchor is progressively refined, remaining open to restructuring as new possibilities emerge.

Designers create a sketch card on the infinite canvas among the references being explored. Designers can optionally pin the card so that it follows the currently selected reference as they move around the canvas. The sketch card provides drawing, layer-management, and high-fidelity generation controls **(Figure 3)**. As designers draw, each enclosed region of the drawing becomes a named, fillable *part*, such as the body, rim, or handle of a vase. Parts are what references are applied to and what history is attached to, so a part's record is also called its region-based history. These parts remain addressable and can be moved, resized, rotated, mirrored, copied, and further edited. The layer panel records how parts are assembled. Designers can group, rename, regroup, reorder, or subdivide parts, and continue adding new parts. Stacking determines overlap, while grouping allows several parts to be transformed together. Thus, the sketch remains assembled from parts that can be reorganized and recombined as the design develops.

Because the sketch remains low fidelity to preserve early-stage ambiguity, designers can request a high-fidelity rendering at any point without committing to it. The rendering reflects the sketch's silhouette, part structure, spatial arrangement, and reference properties, while the sketch remains the editable design object. Each render is stored in history and linked to the sketch state that produced it.

Persistent parts therefore let local changes happen without disrupting prior decisions, while keeping their links to source references and earlier states.

4.1.2 Sketch as a Mutable Anchor for Polymorphic Reference Experimentation. As an initial concept emerges, designers shift from viewing references to trying their visual properties directly (**DG2**). *MoMix* lets designers apply properties from individual or grouped references to specific sketch parts. Because transformations act on persistent parts, local variations can be explored without disrupting prior decisions.

Reference experimentation in *MoMix* is *polymorphic* across visual property and granularity. Designers can apply color, material, and shape to individual parts or groups, while references themselves can likewise be used individually or in groups.

Color. Selecting a reference reveals extracted color swatches that can be dragged onto a sketch part or group, with an immediate cancelable preview. Selecting multiple references instead produces a shared palette derived from recurring colors across the group.

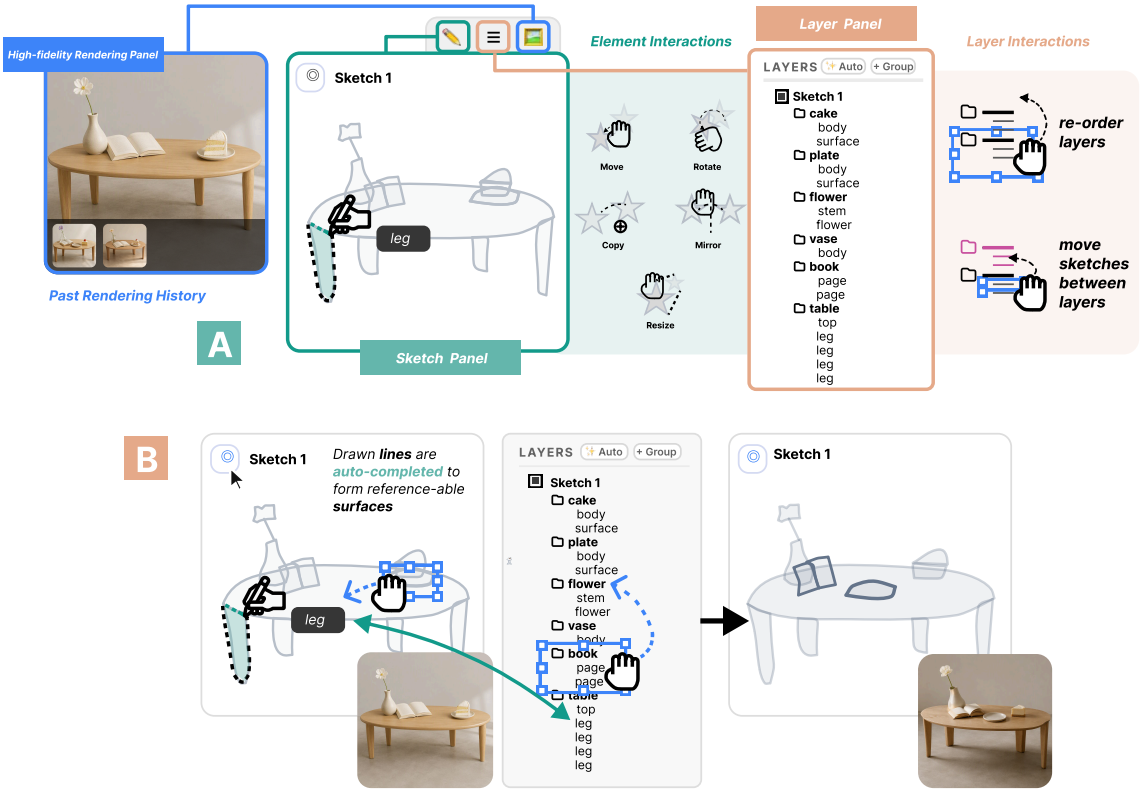


Fig. 3. The sketch card (A) is made up of the sketch panel, the layer panel, and a high-fidelity rendering panel. Any auto-completed drawn lines will form a part on the sketch panel, and direct manipulation on the groups and specific layers is reflected in the rendered image outputs (B).

Material. Designers lasso a reference region to extract a material and drag it onto one or more sketch parts. If a part already contains a color or material, the new property is fused with the existing one, supporting color–color, color–material, and material–material combinations **(Figure 4)**. Sketch parts can also be reset, while grouped references provide representative material samples for comparison.

Shape. Designers can lasso a form from one reference or select multiple references to derive their shared shape patterns. Applying a shape produces three alternatives ranging from closer to the reference to closer to the original sketch **(Figure 5)**. Designers preview these alternatives before confirming or canceling. Shape can be applied to either a single part or a group, and confirmed geometry remains editable for subsequent resizing, regrouping, recombination, or further transformation.

Together, these interactions support part-level experimentation while preserving provenance for later navigation.

4.1.3 Sketch as a Mutable Anchor for Navigating Reference Provenance and Design History. As references, alternatives, and versions accumulate, designers may lose track of how references relate to particular design directions, while collaborators may find these relationships even harder to recover (**DG3**). *MoMix* records each confirmed transformation

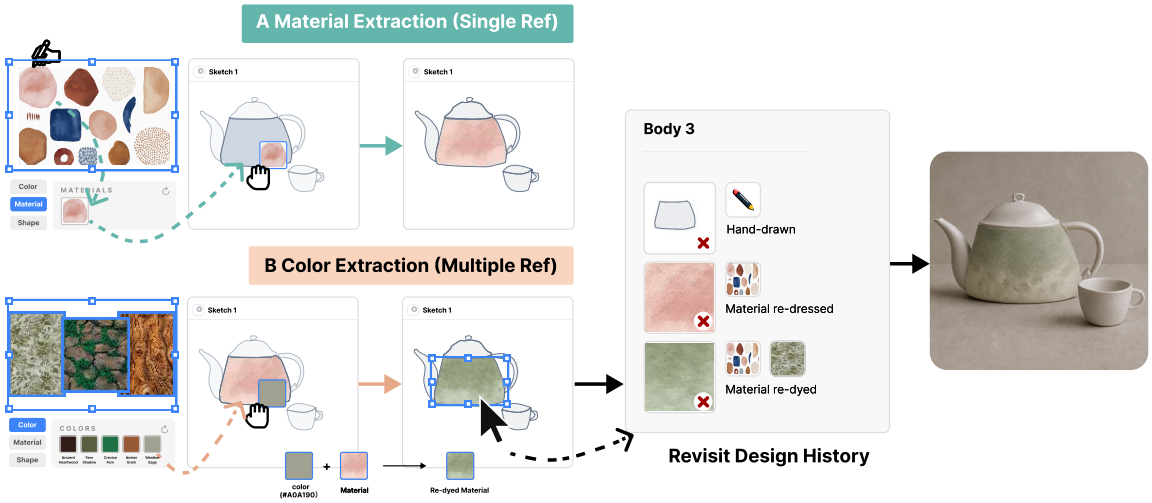


Fig. 4. Material and color can be extracted from single or multiple reference images, then applied to the fillable parts of the sketch. The history of these applications is preserved per part and reflected in the high-fidelity rendering.

with its source reference, visual property, affected sketch part, and design state. The sketch thus becomes an entry point for navigating how references shaped the evolving design, through region-based history and a branching timeline.

Region-Based History. Each sketch part retains its previous reference-driven transformations. Selecting a part reveals the colors, materials, and shapes tried on it, together with their source references. Designers can reapply a previous property to revisit an earlier direction. Grounding provenance in individual parts makes prior experimentation recoverable and reusable **(Figure 6A)**.

Timeline, Branching, and Playback. The timeline captures broader relationships among references and versions across exploration **(Figure 6B)**. Each confirmed transformation adds a version node. Selecting one restores the sketch, high-fidelity rendering, and surrounding board at that state, while highlighting the references that informed it. Confirmed high-fidelity states are automatically starred, users can star additional versions, and the timeline can be filtered to these states. Designers can restore and branch from an earlier state without overwriting the existing direction. Each branch opens as a separate sketch card while retaining its origin; duplicated cards likewise form new branches. Combining sketches merges their branches into a new node. Contributor colors distinguish individual transformations and authorship. For reflection and communication, designers can select two states and play back the process between them **(Figure 7)**. Playback shows the board and reconstructed sketch at each step, animating only the changed parts while highlighting the references that informed them. Each frame therefore exposes the design state, its influencing references, and its contributor.

Together, region-based history and playback make relationships among design ideas, references, and contributors legible over time. They support recovering earlier reasoning and continuing exploration, while helping collaborators understand how a direction developed. The mutable sketch thereby carries a shared design *imaginary*: not only the current design, but the references, alternatives, and contributions through which it evolved.

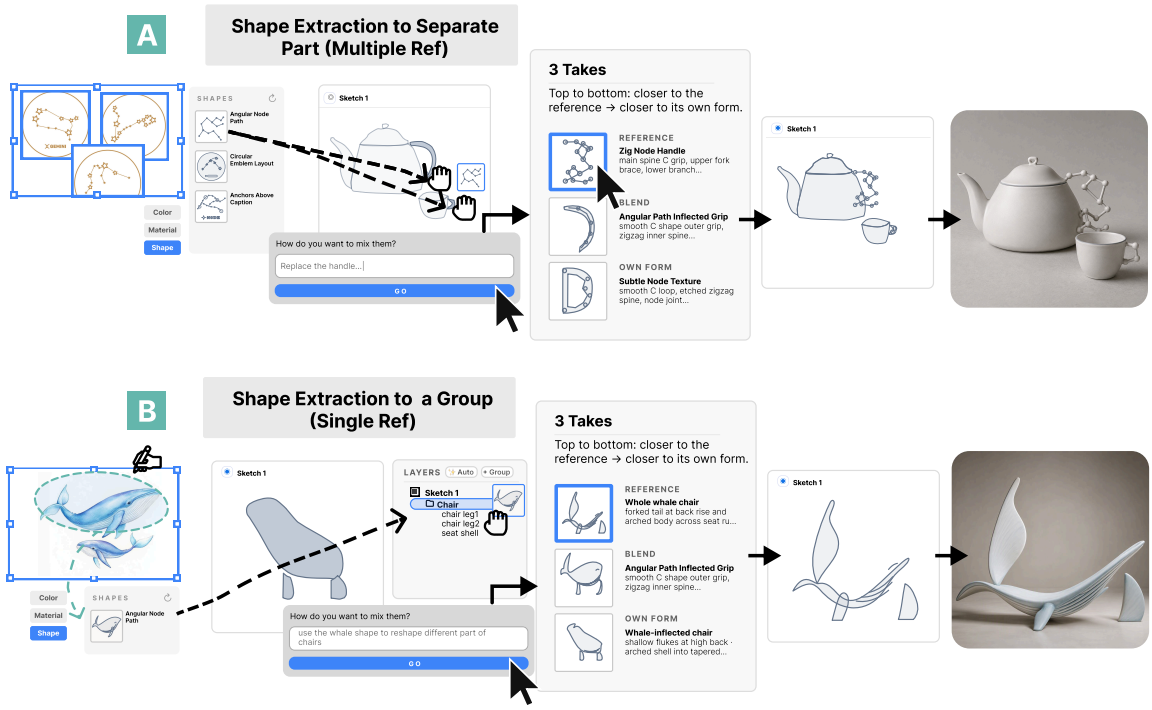


Fig. 5. Shapes can be extracted from multiple (A) or a single (B) visual references and applied to individual sketch parts or groups. Three potential merge styles are presented to the user, ranging from one that is closest to the reference to one that is closest to the sketch's original form.

4.2 Implementation

MoMix is a client-only web application built with React, TypeScript, and Vite. All state and logic run in the browser; a model-service layer handles reference extraction, material fusion, shape transformation, and high-fidelity rendering. Model outputs update the persistent vector sketch or its version state rather than replacing it with a flattened image.

4.2.1 Automatic Part Detection for Sketches. After each completed stroke, the sketch panel reprocesses the drawing. Strokes are simplified using the Ramer–Douglas–Peucker algorithm [13, 45]. Stroke ends are extended and snapped at intersections, while nearby endpoints are welded using a tolerance of 5% of the drawing diagonal, capped at one third of the shortest stroke. All constants were tuned empirically against observed failures. Crossings and T-junctions divide strokes into a segment network from which enclosed regions are detected. Unused extensions are then retracted so the visible ink remains faithful to the designer's original strokes.

4.2.2 Reference Property Extraction and Transformation. *MoMix* supports color, material, and shape extraction from individual or grouped references **(Figure 8A)**.

Color Extraction. For each reference, candidate colors are extracted locally using median-cut quantization [25] and near-duplicates are merged. GPT-5.6 then chooses the final swatches, restricted to colors from this pool. For grouped

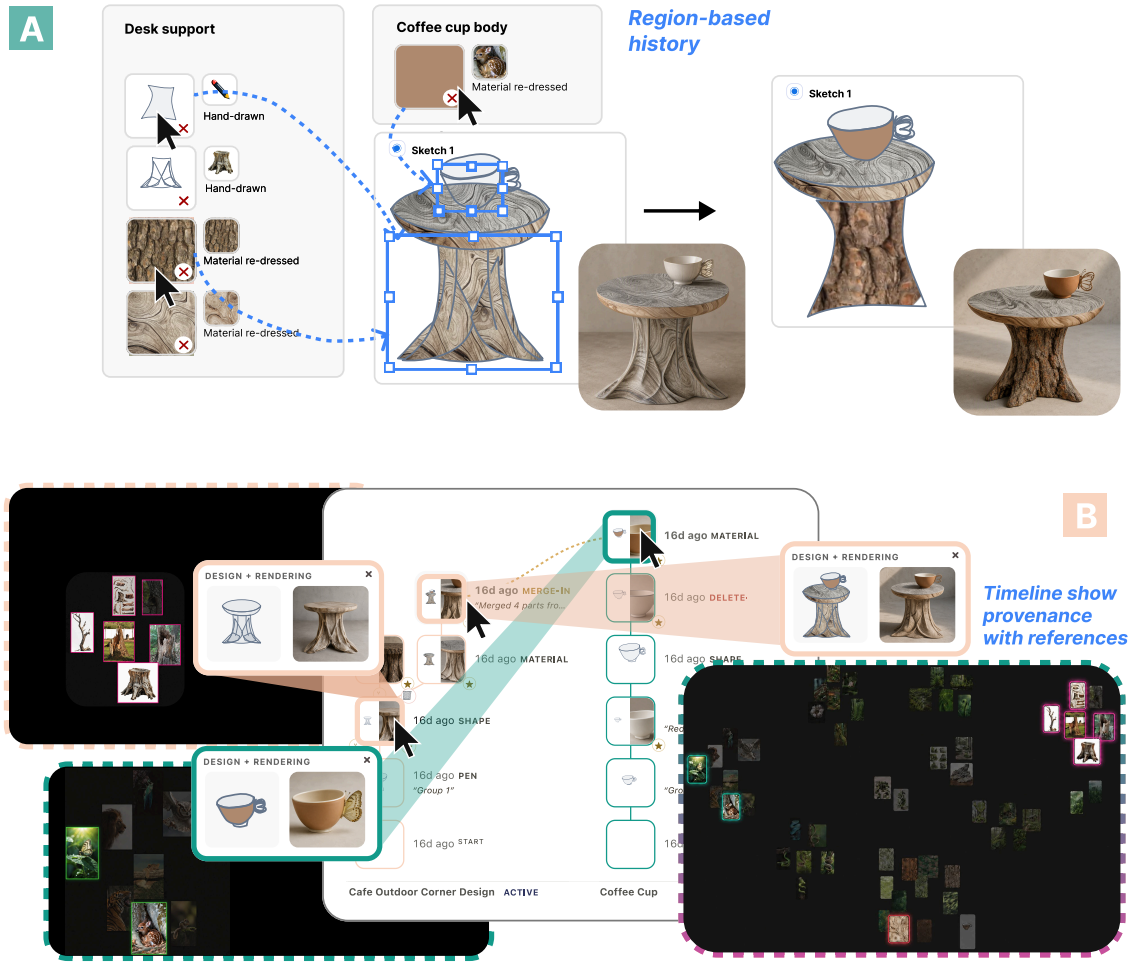


Fig. 6. Two levels at which the sketch anchors navigation of reference provenance and design history: (A) region-based history, in which each sketch part preserves the references and properties previously tried on it; and (B) the branching timeline embedded in the sketch card, which records how the overall design evolved across states. Colors distinguish one contributor from another.

references, five colors per image are pooled while duplicates are retained to emphasize recurring hues; GPT-5.6 selects a 10-color palette from this set, keeping every swatch traceable to source pixels.

Material Extraction and Fusion. A material sample is extracted from the designer's lasso and stored with its source-reference identifier. It is uniformly scaled and cropped to the target part rather than stretched or tiled. If the part already contains a material or color, the appearances are fused using FLUX.2 [klein]. For grouped references, GPT-5.6 selects one representative bounding box per image, which the client then crops deterministically.

Shape Extraction and Transformation. Shape operations run as code generation rather than image generation: the model writes SVG path data, so what comes back is editable geometry instead of pixels. For a single reference, GPT-5.6 reads a lassoed region together with an optional note and emits a colorless, part-labeled line drawing as SVG. For

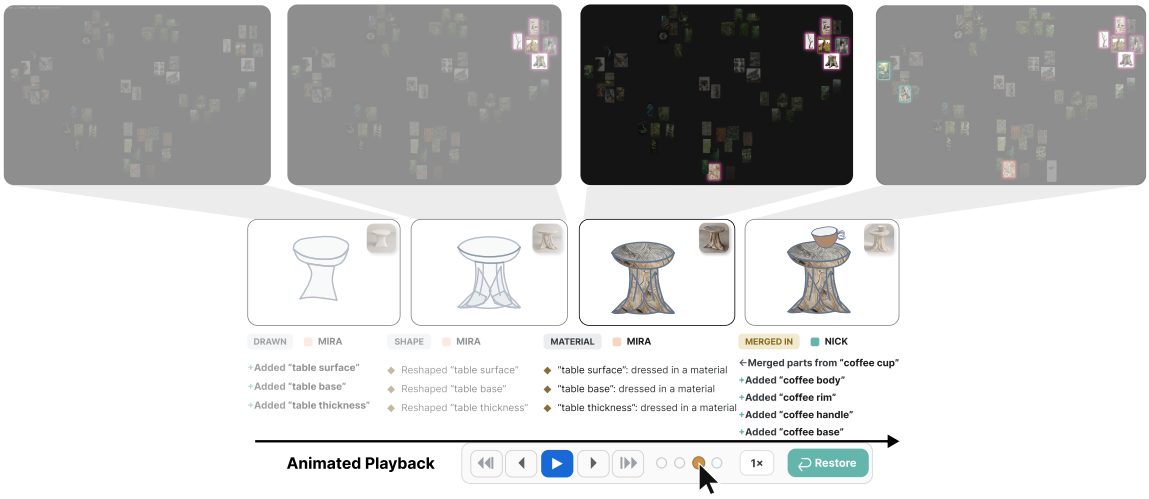


Fig. 7. Playback can be initiated by selecting two designs on the timeline. An animation then plays, transitioning the sketch, the state and structure of the infinite canvas, and the high-fidelity rendering from one design to the other.

grouped references, it identifies shared shape patterns and emits each as a separate line swatch. To apply a swatch, *MoMix* produces three variants ranging from close to the reference to close to the original sketch: Gemini 2.5 Flash first proposes a transformation direction for each variant, and GPT-5.6 then writes all three paths in a single call with differing reference-adherence strengths. Because a variant arrives as path data rather than as an image, the designer can reshape, resize, and recombine it as soon as it is confirmed, without a further round of generation.

4.2.3 High-Fidelity Rendering. High-fidelity rendering turns the current sketch into a realistic image, while the sketch remains the source of the design. Its composition, part boundaries, spatial arrangement, and per-part appearance are determined by the designer's sketch rather than composed anew by the model. *MoMix* exports the sketch together with its layer hierarchy, design intent, and part-level reference bindings. These inputs are assembled into a structured prompt with seven sections **(Figure 8B)**. Two state the situation: the *design intent* and the *attached references*. Two fix the structure: the *object structure* names every part the render may contain, and the *layout scaffold* ties each silhouette to the sketch. Three fix the appearance: *material and color references* bind each reference to the parts it may dress, *shape language* states the form details that must survive, and *integration* describes how the parts sit together as one scene. References remain bound to the parts they inform, and stacking order provides approximate depth. The sketch, references, and structured prompt are then passed to GPT-Image-2 to generate the high-fidelity render. Each render is stored with its version node while the underlying SVG remains editable. Thus, the render depicts a particular design state rather than replacing the mutable anchor.

4.2.4 Provenance and Version State. *MoMix* stores the sketch and provenance as SVG and JSON records in IndexedDB, allowing history to persist with the project. Each confirmed application records its type (color, material, or shape), resulting appearance, source reference, timestamp, contributor, and the part it affected. The first change also records the part's original state. Indexing records by part enables each part to expose and restore its own history independently. The same records form a version graph. Each confirmed change adds a node containing the full sketch, different sketch part appearances, board context, timestamp, contributor, and associated render. Parent pointers preserve branches.

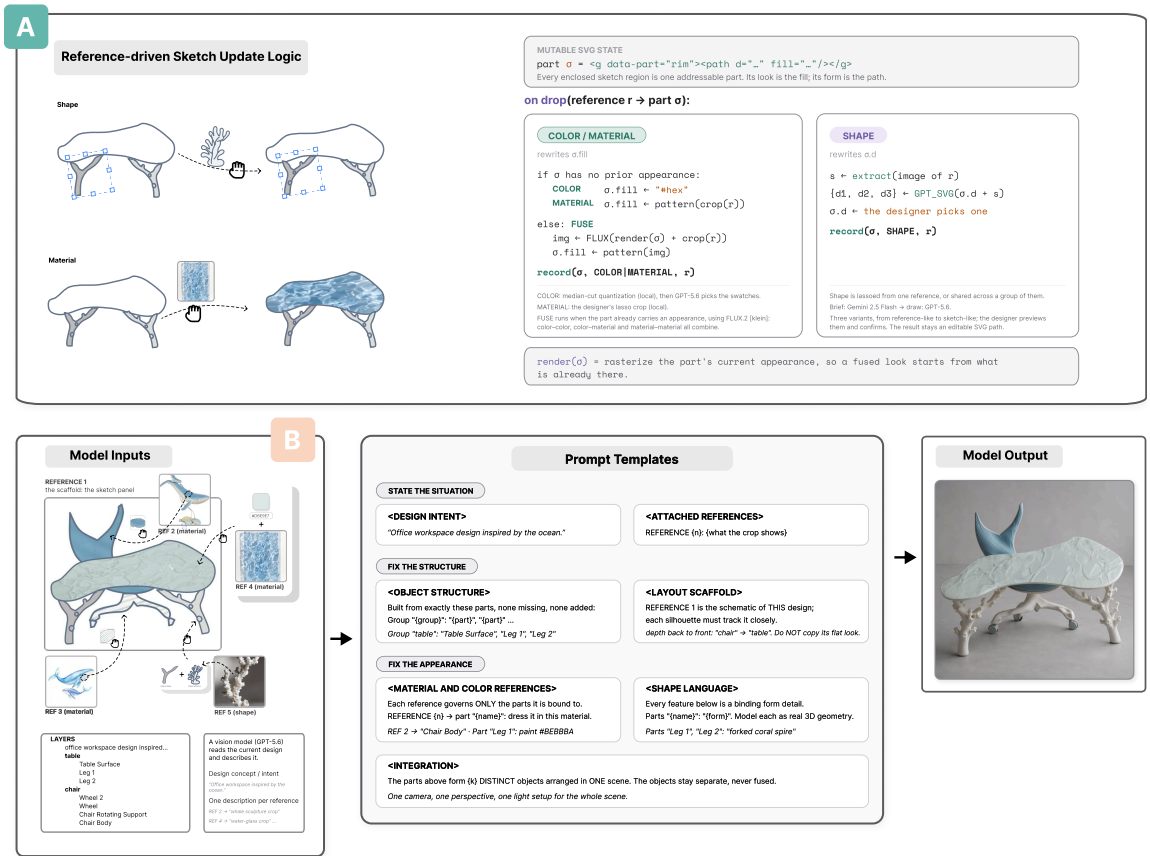


Fig. 8. *MoMix* supports reference-driven sketch transformation (A) by applying colors, materials, and shapes extracted from visual references to editable parts of an SVG sketch. The resulting sketch state, together with its layer structure, reference bindings, and design intent, is then organized into structured prompts (B) to generate a high-fidelity rendering.

Playback traverses these nodes, restores successive states, identifies changed parts, and highlights their informing references. Because both region history and the timeline use the same records, provenance can be accessed either from a design part or from a point in the overall process. A contributor is just a locally set name stamped on each commit, its color derived from the name; there are no accounts or live sync. The study seeded a board under one name and continued it under another.

5 User Study

To evaluate the creativity support *MoMix* offers and understand whether a mutable anchor can sustain reference-driven experimentation and navigation on infinite canvases, we conducted a within-subjects lab study with 12 participants. We compared *MoMix* with a baseline infinite-canvas workflow and structured our study around the following research questions:

RQ1: *How does the sketch as a mutable anchor shape the way designers experiment with design ideas using visual references?*

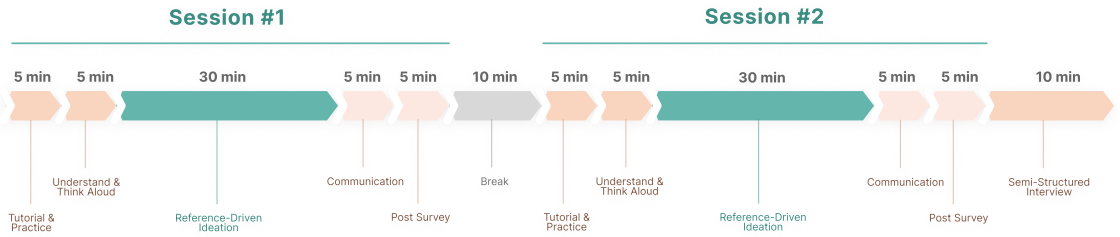


Fig. 9. User study procedure. The two-hour study consisted of two counterbalanced conditions, each including five minutes of understanding the existing board and thinking aloud, a 30-minute reference-driven experimentation task, a five-minute communication task, and a post-task survey. The study concluded with a 10-minute semi-structured interview.

RQ2: *How does the sketch as a mutable anchor affect the way designers navigate a shared design imaginary, both across their own design history and across a design inherited from an earlier stage?*

5.1 Participants

As *MoMix* targets designers who work with visual references during early-stage ideation on an infinite canvas, we recruited 12 participants (3 male, 9 female, aged 24–33, $M = 27.9$) with prior design experience. Their practices spanned UI/UX, product, graphic, illustration, architecture, and motion design, with an average of 4.6 years of experience. All had prior experience with GenAI for creative or design activities and some familiarity with moodboarding. Participant demographics are summarized in Appendix **Table 4**. Each participant received a 50 USD Amazon gift card.

5.2 Procedure

5.2.1 Apparatus and Tasks. Sessions were conducted on a computer using an infinite-canvas workspace with keyboard and mouse input. In the *MoMix* condition, participants used the complete interface described in Section 4.

The baseline reflected common interactions in commercial GenAI-integrated infinite-canvas workflows, such as Adobe Firefly Boards[3], Flowith[4], and Lovart[5]. It provided freehand sketching and prompt-based image generation. Participants could select references or capture regions containing sketches, references, and other board content, then combine them with text prompts. Both conditions used GPT-Image-2 for image generation.

Participants completed two object-design tasks, one per condition: an ocean-inspired office workspace corner and a forest-inspired café corner. Each scenario included example objects such as tables, chairs, storage, planters, and lamps. Theme-to-condition assignment and condition order were counterbalanced. Each starting board contained two design ideas from a hypothetical collaborator, represented as sketch cards with references and design history in *MoMix*, and as sketches, generated images, and lightweight annotations in the baseline.

5.2.2 Study Procedure. The two-hour study used a within-subjects design with counterbalanced condition and theme order **(Figure 9)**. After consent and a demographic survey, participants completed both conditions. Each condition began with a five-minute tutorial and hands-on practice, followed by five minutes of inspecting the existing board while thinking aloud about the collaborator's design direction and how surrounding references might have shaped it.

[3]Adobe Firefly Boards. https://www.adobe.com/products/firefly/features/boards.html, accessed 2026-08-31.
[4]Flowith. https://flowith.io, accessed 2026-08-31.
[5]Lovart. https://www.lovart.ai/, accessed 2026-08-31.

Participants then completed a 30-minute reference-driven experimentation task. They could refine the collaborator's direction, develop it differently, or start their own. Participants were encouraged to explore possibilities and progressively develop two to four object designs until the direction was sufficiently formed to communicate. Each condition concluded with a five-minute communication task, in which participants treated the experimenter as the collaborator who had created the initial design. They explained their design direction, the references and alternatives that shaped it, and how it related to the collaborator's original direction. When needed, the experimenter asked follow-up questions about reference influence and continued collaboration. Participants could use any features available in the condition.

After both conditions, participants completed a 10-minute semi-structured interview comparing the two workflows on their design process, reference experimentation, interpretation of the collaborator's intent, and communication.

5.2.3 Measures. After each condition, participants completed four questionnaires. The first three assessed the design workflow: the System Usability Scale (SUS) [5], NASA Task Load Index (NASA-TLX) [24], and Mixed-Initiative Creativity Support Index (MICSI) [32]. MICSI extends the Creativity Support Index [9] to human–AI co-creation, measuring creativity support (exploration, expressiveness, immersion, enjoyment, and results worth effort), human–machine collaboration (communication, alignment, agency, and partnership), and perceived outcomes of authorship, satisfaction, surprise, and novelty.

The fourth questionnaire measured communication self-efficacy in the study's two-way handoff. Following Bandura's format [3], participants rated six items from 0 (*cannot do at all*) to 10 (*highly certain can do*): three on communicating one's own design and three on understanding and continuing an inherited design. Items drew on prior work on design communication [6], moodboards' communicative roles [38], and shared grounding [11]. The full survey appears in Appendix **Table 5**.

5.2.4 Data Analysis. We used a mixed-methods analysis. Think-aloud and interview data were analyzed inductively using open coding and constant comparison [8], informed by grounded-theory techniques [7].

For questionnaire data, we compared conditions using paired tests. SUS was scored from 0–100; NASA-TLX as the unweighted mean of six dimensions (performance reverse-scored); MICSI as five creativity-support factors, four collaboration dimensions, and four bipolar outcome items; and communication self-efficacy as two subscales. Multi-item scales used paired t-tests (Cohen's d_z) when Shapiro–Wilk tests supported normality of paired differences, and Wilcoxon signed-rank tests (rank-biserial r) otherwise; single items used Wilcoxon tests. All tests were two-sided at $\alpha = .05$. Given the exploratory study and small sample, we report effect sizes without multiple-comparison correction and interpret quantitative results alongside qualitative and behavioral data.

At the output level, we measured how far high-fidelity outputs departed from their source references using Learned Perceptual Image Patch Similarity (LPIPS) [62]. We computed both mean distance across references and minimum distance to the nearest reference; higher values indicate less reference-derivative outputs.

We also derived behavioral measures from interaction logs and saved boards. For each design, *iteration depth* was the number of editing steps in *MoMix* (e.g., applying color, material, or shape, or redrawing a part) and the number of generated images in the baseline. We separately counted *high-fidelity generations* per design and *distinct references* used across the session.

Table 2. Similarity of each participant's high-fidelity outputs to the references used to create them, averaged across participants (n = 12; paired Wilcoxon signed-rank tests). Higher LPIPS indicates outputs that are less derivative of their references.

	MoMix	Baseline	p
LPIPS to references (mean)	0.69	0.66	.052
LPIPS to the nearest reference (min)	0.64	0.56	.11

Table 3. Process measures per design (n = 12). The baseline offers no editing below the level of a generation, so the first row compares each condition's *smallest unit of iteration* rather than the same operation. The conditions did not differ significantly in generations per design or in references used.

Process measure	*MoMix*	Baseline	p
Iteration depth (smallest unit of iteration)[a]	13.1	1.9	$< .001$
High-fidelity generations per design	2.7	1.9	.052
References used over the process	7.9	6.4	.15

[a]Editing steps in *MoMix* (applying a color, material, or shape, or redrawing a part); high-fidelity generations in the baseline, which has no sub-generation editing.

6 Study Results

6.1 How the Sketch as a Mutable Anchor Supported Reference-Driven Experimentation

6.1.1 Designer-Originated Variation Supported More Creative Outcomes. Nearly all participants perceived the final high-fidelity outcomes produced with *MoMix* as more creative than those produced with the baseline **(Figure 10)**. Their MICSI ratings showed the same pattern: *MoMix* outcomes were rated significantly more novel than baseline outcomes (median 3.5 → 5.5, p = .019, r = .79) **(Figure 11A)**. A perceptual-distance analysis pointed in the same direction. *MoMix* outputs' mean LPIPS to the reference set was higher (mean 0.66 → 0.69, p = .052, r = .59; higher for 9 of 12 participants), and their minimum LPIPS to the nearest reference was also higher (mean 0.56 → 0.64, p = .11, r = .46). Neither difference reached significance, but the pattern suggests that baseline outputs sometimes stayed closer to a single reference whereas *MoMix* outputs stayed farther from any individual one **(Table 2)**. As P7 described, the baseline results were "*kind of generic and stick strongly with the reference*," while with *MoMix* she arrived at an idea she "*never would have thought of*," such as making a lamp arm resemble jellyfish tentacles.

Participants attributed this difference to the role of the sketch as a mutable anchor. Rather than generating directly from visual references, they first introduced their own forms and structures, then progressively incorporated reference properties into them. As P5 explained, "*The output feels more novel when I do the manual sketch, because the sketch is actually creating something, whereas prompting mostly brings out things the model already knows.*" P8 similarly described the sketch as adding "*a layer of harness to the AI output*," grounding the process in human-authored ideas rather than "*AI generative stuff.*" Thus, more creative outcomes arose not simply from using references, but from using the sketch as a mutable anchor to introduce designer-originated variation beyond reference images and model priors alone.

6.1.2 Decomposed Reference Properties Supported Direct Visual Expression. Participants described reference-driven design as an inherently visual process. They already attended to properties such as color, material, and shape when interpreting references (P5, P6, P10), but conventional GenAI workflows often required these visual intentions to be translated into language. As P2 explained, design was "*a fundamentally visual process*," whereas text required her to "*verbalize the visual changes rather than just working with the visuals.*" This became difficult when participants knew

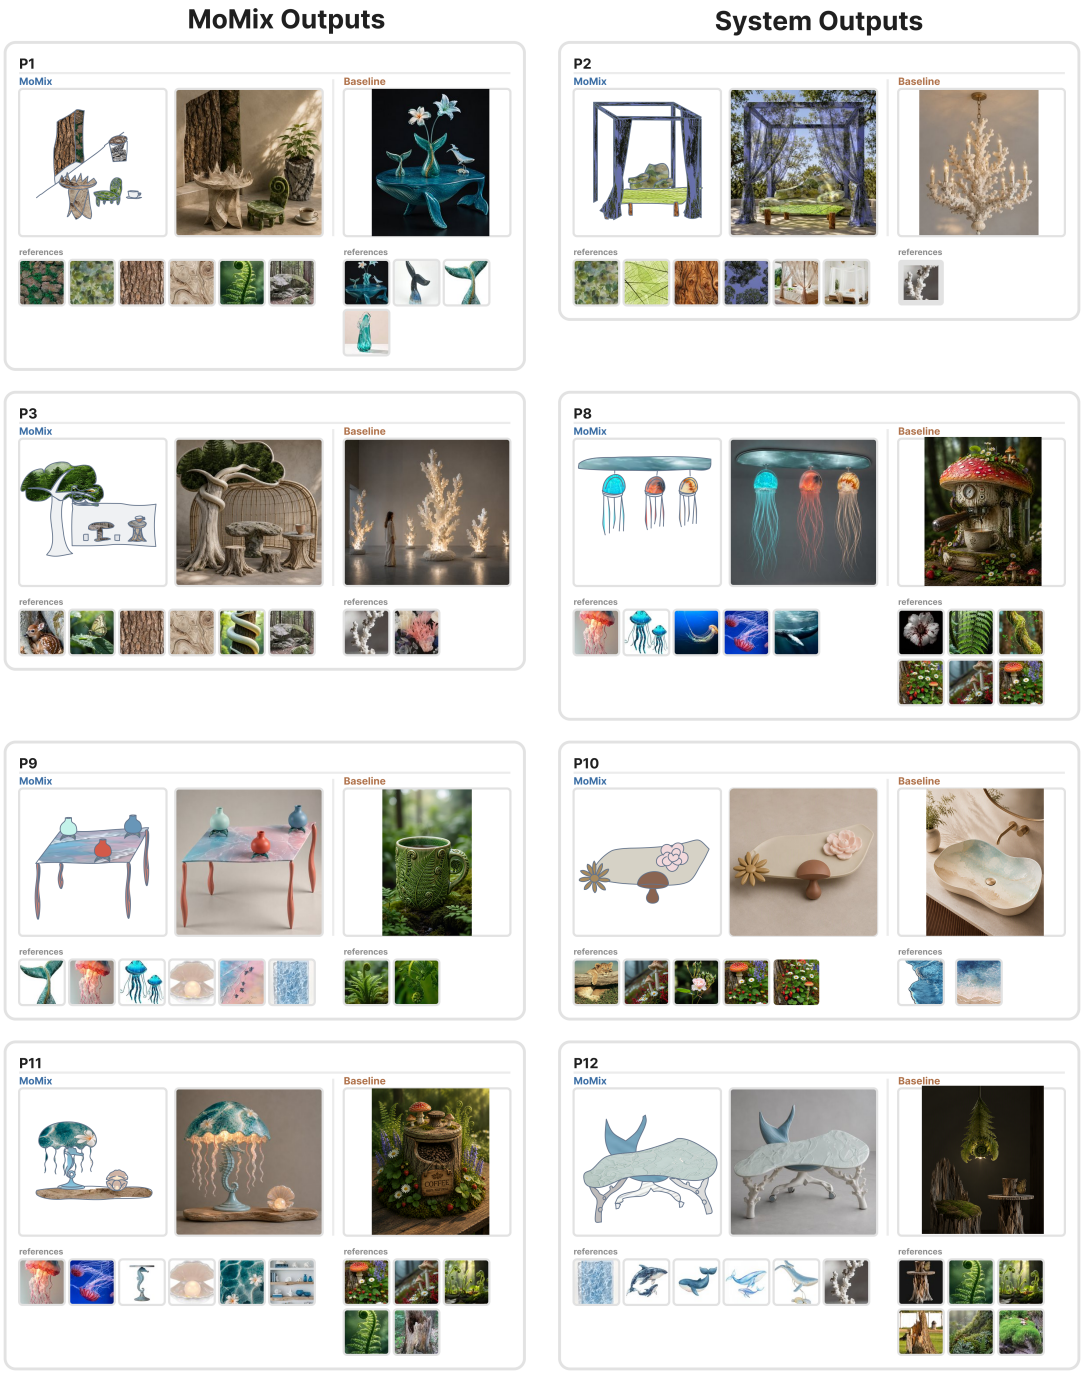
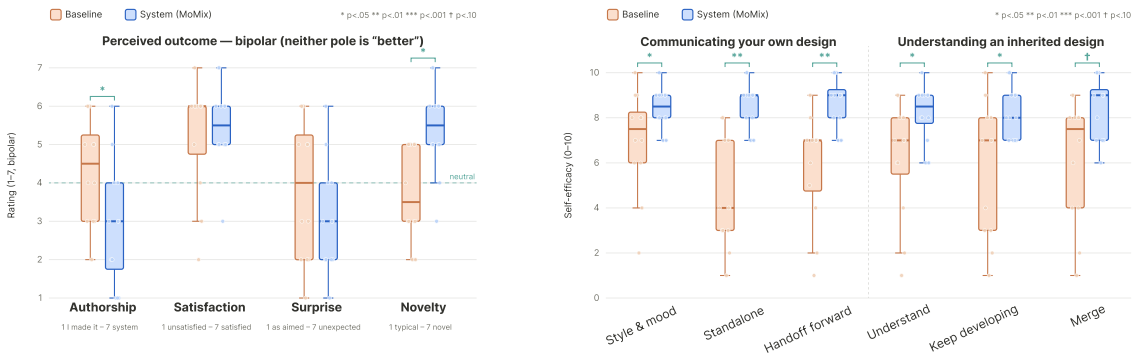


Fig. 10. Participants' sketches and high-fidelity outputs from *MoMix* (left column) and the baseline (right column).



(A) MICSI perceived outcomes **(B) Design-communication self-efficacy**

Fig. 11. (A) The MICSI perceived-outcome results on four bipolar items. Authorship and novelty shifted with *MoMix*, while satisfaction and surprise did not. (B) The communication self-efficacy results on their two aspects: communicating one's own design to someone who did not watch it being made, and understanding and continuing a design inherited from a collaborator.

what they wanted visually but lacked a precise linguistic description (P5, P6, P7, P11). P7 noted that "*sometimes … the language is limited and you cannot describe … what you actually imagine visually in your head*," and described a lamp shape as "*difficult to describe in text*" **(Figure 12)**. Participants also struggled to specify which reference property should affect which design part; P1 sometimes ignored shape because it was "*too tedious to describe the shape in the language form*." With *MoMix*, participants could instead extract color, material, or shape and apply it directly to the relevant sketch part. Shape extraction also helped when they were unsure how to draw a form themselves (P1, P6, P9, P10, P11). As P11 explained, "*It's hard for me to actually draw a peel or a lamp directly… But here I can actually just go from a picture and get the shape and then make changes.*" The last clause matters: what landed on the part was not the reference form copied over it, but one of three variants between that form and the part's existing geometry, and it stayed editable afterwards. Participants used this to carry a form across categories rather than to reproduce it, giving a chair back the curve of a shell **(Figure 13)** and a pair of chairs the profile of a rock **(Figure 14)**. A prompt can ask for a shell-like chair back, but it cannot say which of the designer's own parts should keep its proportions while taking on that curve, which is why P1 found shape "*too tedious to describe in the language form*" at all. Consistent with these accounts, participants reported more effective communication of intent to *MoMix* on the MICSI communication dimension (median 4.5 → 6.0, $p = .020$, $r = .89$) **(Figure 15)**. Thus, decomposed reference properties reduced the need to translate visual intent into language by making it directly actionable on the evolving design.

6.1.3 Mutable Ideas and Recombinable Reference Properties Supported Cognitive Flexibility. Using the sketch as a mutable anchor kept emerging ideas open to revision rather than prematurely fixing them in high-fidelity outputs. P4 noted that in the baseline she moved quickly to high-fidelity generation and had "*already kind of become fixed*" because "*there's no … in-between.*" With *MoMix*, she was instead "*forced to flesh out the idea more before I jumped to the high-fidelity version*," allowing her to "*change and develop my idea more*" without being "*siloed into … a specific direction from the very beginning*" **(Figure 13)**. The sketch as a mutable anchor therefore provided an intermediate representation in which emerging ideas could progressively accumulate, change, and be reconsidered.

Interaction logs reflected this process **(Table 3)**. The two conditions did not differ significantly in high-fidelity generations per design (mean 1.9 → 2.7, $p = .052$, $r = .64$) or in the number of references drawn on over the session

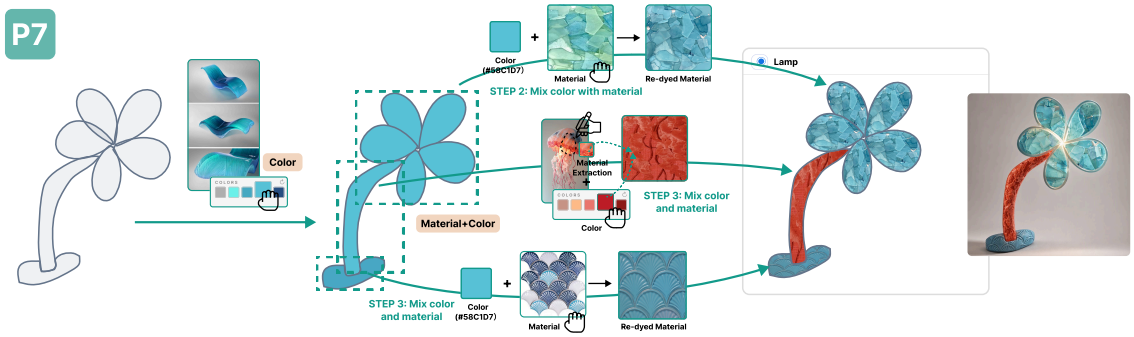


Fig. 12. P7's lamp, built from color and material alone. Each part takes a reference material fused with a chosen color, giving the shade its translucent, jellyfish-like look.

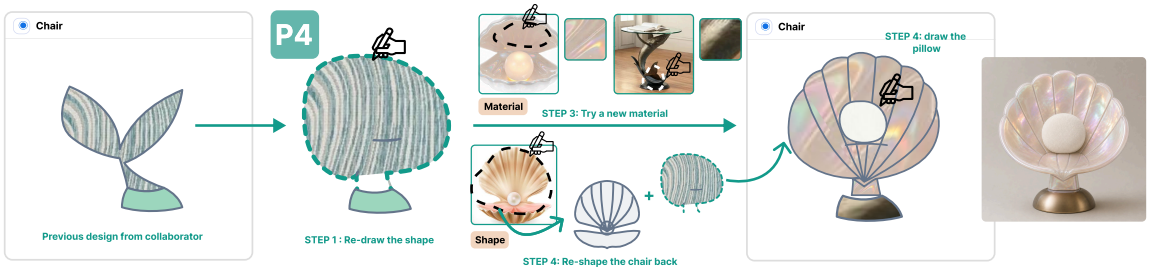


Fig. 13. P4 continuing an inherited chair: she redraws its shape, tries material references, then re-shapes the back into a shell and adds a pillow where the back meets the seat.

(mean 6.4 → 7.9, $p = .15$, $r = .50$). What differed was the process between those renderings. Each *MoMix* design developed through a mean of 13 incremental editing steps, such as applying a reference's color, material, or shape to a part or redrawing a part. The baseline had no comparable sub-generation editing, so a design advanced only through its roughly two high-fidelity generations. Both workflows thus reached high-fidelity images, but *MoMix* supported a far more incremental, step-by-step process of developing a design before and between renderings.

Cognitive flexibility also came from how designers could engage with references. Rather than treating each reference as a fixed combination of color, material, and shape, participants could selectively draw on and recombine these properties across references. P11 compared this with thinking about "*the whole picture*," which drew more on "*my priors of what I've seen*"; considering properties separately, by contrast, could lead to something "*out of distribution*" or "*that you have not intentionally thought about before.*" P4 similarly felt that treating color, material, and shape "*all together as a packet*" restricted her ideas, whereas separating them "*encourages more creative directions because you can mix and match many combinations of these extracted elements to make new things.*" P5 described this exploratory space as "*a mixing pot*" with "*endless combinations.*"

Together, using the sketch as a mutable anchor and decomposing references into recombinable properties supported cognitive flexibility in two complementary ways: designers could revise an evolving idea without committing prematurely, while also exploring reference properties without remaining bound to any single exemplar.

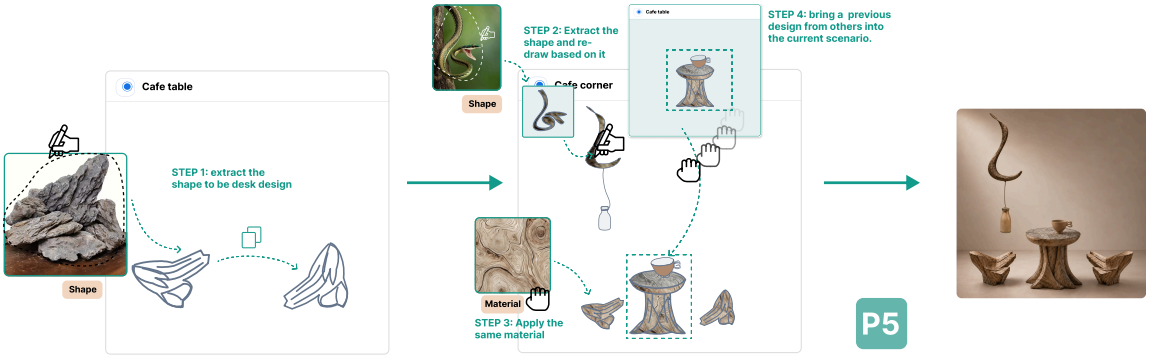


Fig. 14. P5's café corner. A rock's profile becomes two chairs, a second reference becomes a hanging decoration, one wood material is shared across the furniture, and an inherited café table joins the scene.

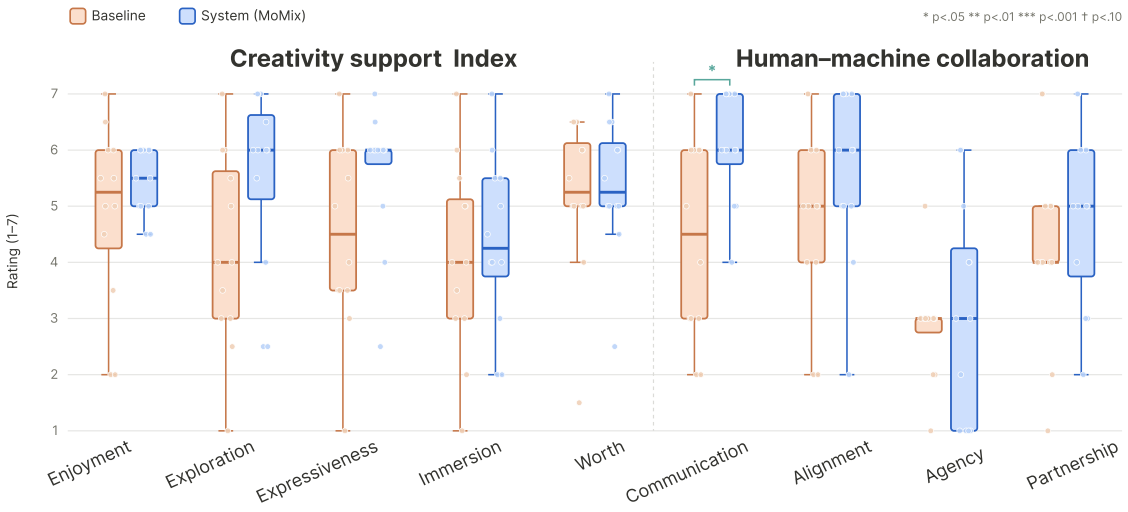


Fig. 15. How the 12 designers rated *MoMix* and the baseline on the MICSI creativity-support and collaboration items.

6.1.4 Part-Level Control Strengthened Sense of Agency and Ownership. Direct, part-level editing gave participants greater control over how their designs developed (P1, P3, P4, P5, P8, P10, P11, P12). They could set visual properties directly on the sketch and its individual parts. For example, P5 noted that sketching could show directly how tall she wanted an element to be. Colors or materials could be dragged onto specific parts, and the rotation angle could be set directly rather than by repeatedly adjusting a prompt. P3 contrasted this with the baseline, where "*every single time I had to repeat myself to explore different ideas*" and "*it was kind of hard to steer the output towards my goal.*" With *MoMix*, she could instead "*quickly create the shape and then drag it on, and then you can iterate really quickly.*" Participants could therefore make local changes while keeping previously established parts and decisions intact, which allowed more precise, step-by-step steering of the evolving design.

Part-level control also changed how participants understood what the design was. With *MoMix*, they treated the anchor itself (the sketch they kept building) as the design idea, and the high-fidelity rendering as only one depiction

of it. In the baseline, where they worked solely with the generated outputs, the rendering was the design itself. The difference in what counted as "the design" underlay their stronger sense of agency and ownership with *MoMix* (P2, P3, P5, P8, P9). When using *MoMix*, P3 viewed the final AI-generated image as "*just a re-rendering for us*," because "*the novelty comes from the choices that I made with the system*." P2 described this ownership as coming from constructing with design materials. Compared with simply editing a generated image, *MoMix* reminded her of "*playing with Legos, like you get to think of like the different shapes*," and "*painting a picture, like I had a blank canvas and I was painting a picture*." Participants therefore located their agency and ownership in the design decisions they accumulated on the anchor, rather than in who produced the final pixels. Consistent with this, on the MICSI authorship item, which runs from 1 (*I made it*) to 7 (*the system made it*), they felt more strongly that they, rather than the system, had made the result (median 4.5 $\rightarrow$ 3.0, $p = .039$, $r = -.86$) **(Figure 11A)**.

6.2 How the Sketch as a Mutable Anchor Supported Navigation

Participants reported higher communication self-efficacy with *MoMix* across both communicating their own designs and continuing designs inherited from an earlier stage **(Figure 11B)**. For communicating one's own design, the overall score increased (mean 5.78 $\rightarrow$ 8.61, $p < .001$, $d_z = 1.42$), with significant improvements in capturing the intended style, making the board understandable on its own, and enabling others to carry the direction forward. For inherited designs, the overall score increased (mean 6.03 $\rightarrow$ 8.28, $p = .016$, $r = .77$), with significant improvements in understanding the collaborator's original intent and continuing it without losing its direction; integrating one's own ideas also showed a marginal improvement ($p = .064$, $r = .63$). Overall communication self-efficacy increased (mean 5.90 $\rightarrow$ 8.44, $p = .003$, $d_z = 1.07$).

These results suggest that communication in this setting was not only about explaining a final artifact, but about making an evolving design direction navigable by someone arriving at it later. The following findings show how region-based history and provenance timelines supported this navigation by helping designers recover, communicate, and continue the shared design imaginary.

6.2.1 Navigable Design Histories Supported Recovery and Reflection. For the designers themselves, the timeline made earlier decisions, rationales, and abandoned alternatives easier to revisit along an evolving design trajectory. P4 rediscovered an earlier material while navigating the timeline, remarking, "*I didn't realize I had used this material. It's pretty cool*," and then switched the design back to that previous material. P8 similarly viewed earlier timeline nodes as return points when a design direction had developed too far. Almost all participants also navigated through design steps and used playback during the final communication stage. When explaining their designs, they explicitly described how particular references had shaped specific decisions, whereas in the baseline, their explanations focused more on the final design ideas themselves and seldom mentioned how references had influenced particular designs.

Navigating design history also supported reflection on the broader trajectory rather than merely recovering past states. P3 described the timeline as useful when "*I'm cutting down*," helping her recognize moments such as "*I need to stop trying to fix this and move on to the next*" or "*I need to ask other people for advice now*." Thus, the timeline provided a navigable overview for revisiting, pruning, and explaining one's evolving design decisions.

6.2.2 Navigable Provenance Helped a Later Designer Reconstruct a Shared Design Imaginary. For a designer arriving at a board later, provenance made it easier to navigate how references and intermediate artifacts contributed to an evolving design direction, helping that designer reconstruct a shared design imaginary rather than infer intent from the final artifact alone. This difference was clearest when P6 inspected the collaborator's board. In the baseline, she could

recognize some spatial progression between images but remained uncertain about how individual artifacts related to one another: "*I'm not sure… maybe this is just like three different variations of the same thing.*" With *MoMix*, she could instead trace a specific sequence of decisions: "*She started with a cup, made it bigger, got inspiration from butterflies, the color of a deer… Very clear. I don't have any questions.*" Provenance therefore turned relationships that otherwise had to be inferred from spatial proximity into an explicit trajectory that a later designer could follow.

This mattered most when participants needed to continue a more complex inherited design. P8 described the timeline as showing "*how they started and what are the middle points and what are the pivot points.*" P6 similarly explained that navigating the history exposed the previous designer's internal logic—which parts were connected, where ideas came from, and how the design developed—making it easier to "*work on her version.*" In this sense, process history made the evolving design imaginary navigable across stages of the design, helping a later designer recover how a direction emerged and where they could continue from it.

6.2.3 Replayable Histories Made Others' Design Processes Learnable. Participants also saw replayable histories as learning materials. P6 compared them to physically watching experienced users work in AutoCAD or Photoshop, where observing intermediate operations and small "tricks" helped her learn how experts worked. P2 and P4 similarly imagined that, as novice creators, seeing an experienced user's workflow and how they used different references could help them "*improve my own creative workflow.*" P3 further suggested that this value depended on the relationship between collaborators: novices may benefit from seeing an expert's detailed design choices, whereas experienced peers may need less process detail. Thus, provenance could support observational learning by exposing not only what others produced, but how they worked through a design.

6.2.4 Navigable Histories Could Become Reusable Design Material. Navigating another designer's history also surfaced materials for subsequent experimentation. Almost all participants opened the timeline to look for ideas from the prior creator; seeing an earlier attempt at a chair, for example, prompted P4 to explore her own chair **(Figure 13)**, while P5 extended a collaborator's wooden coffee-table direction into a stylistically related chair **(Figure 14)**. In these cases, navigation did not end with understanding the past process: earlier states and decisions could re-enter the current exploration as design material. However, turning design history into design material was selective. Only P1, P4, and P5 incorporated elements from a collaborator's prior design into their final *MoMix* outcomes, typically by reworking a useful object, form, or material into their own exploration. For others, history was more useful for understanding how a design had developed than for direct reuse. P2 drew more inspiration from the finished artifact and references, while P3 focused on collaborators' final design choices rather than the process behind them. Thus, provenance became reusable design material mainly when it surfaced elements relevant to the designer's current goals.

6.3 Constraints and Design Implications

6.3.1 Added Flexibility Came with a Higher Workload. These gains came with additional interaction costs. *MoMix* scored lower on the SUS than the baseline (mean 77.7 → 65.8, $p = .027$, $d_z = -0.74$), and participants reported higher mental demand (median 3.0 → 6.0; mean 3.75 → 4.92, $p = .039$, $r = .86$), while other NASA-TLX dimensions did not differ significantly **(Table 6)**. Participants partly attributed this burden to the greater number of choices and operations available in *MoMix*; for example, P10 associated having "*more freedom*" with greater workload. The flexibility a mutable anchor affords therefore brings its own decisions, about when and how strongly to constrain an emerging design.

6.3.2 Adapting Mutable Anchors for Design Convergence. Participants found sketch-based exploration particularly useful once an emerging direction had begun to form and they wanted to prototype and refine variations before committing to a final output. However, introducing a sketch anchor too early could narrow more open-ended exploration. P3 described *MoMix* as "*more helpful for maybe convergent exploration than divergent exploration*," because her usual early-stage process involved generating multiple independent alternatives before comparing and pruning them. P8 similarly noted that giving AI fewer constraints could expose "*more of the possibilities*," including unexpected generations that might lead beyond the designer's initial direction. P1 therefore favored a hybrid interaction, noting that "*it's helpful to have both, like, very granular control, precise control, and also have some heuristics that we can tell the system about*."

These accounts suggest supporting fluid shifts between open-ended AI-assisted exploration when designers are uncertain about what to pursue and sketch-based control as their intentions become more concrete.

Participants also wanted sketch representations to remain mutable throughout exploration, including shapes extracted from references. P12 suggested vectorizing extracted forms into adjustable control points so designers could directly reshape, resize, and refine them rather than redraw them. P6 further emphasized that extracted forms should match the fidelity of the current sketch: designers often begin with rough representations and progressively add detail, so inserting a highly resolved form too early could feel inconsistent and prematurely stabilize the design. Shape extraction could therefore support multiple levels of fidelity, from loose sketch-like forms to more refined vectors, while keeping each representation editable. More broadly, these findings suggest that both the degree of constraint and the fidelity of a sketch should adapt as a designer's emerging intent becomes more concrete.

7 Discussion

7.1 Balancing Sketch-Based Control and Open-Ended Exploration

Participants described a Lego-like and mosaic-like experience of building designs from editable sketch parts and selectively borrowed reference properties. Moving, rotating, and copying parts across sketch cards supports *combining*, while redrawing, regrouping, and transforming parts with reference shapes creates opportunities for *restructuring*: discovering different parts and relationships within a configuration [56, 57]. Polymorphic reference experimentation extends this flexibility by allowing color, material, and shape to be explored beyond their original combinations. These local changes can have implications beyond the properties being manipulated. As Schön and Wiggins explain, a move motivated by one concern can reveal unanticipated consequences that inform subsequent design decisions [48]. A borrowed shape or material can thus change how designers understand the emerging sketch and what they explore next. The sketch remains a mutable anchor through which these discoveries can be incorporated into further construction.

This interplay between combining and restructuring offers an explanation for the greater perceived novelty and ownership in our findings. Designer-initiated sketches bring forms, arrangements, and reference combinations into generation, all of them developed through visual thinking. These choices emerge through drawing and inspecting, as designers discover what they want to pursue [20]. The model therefore develops a configuration already shaped by the designer's exploration and judgment. The Lego-like experience locates ownership in this accumulated construction: designers determine what to introduce, transform, and retain. Their sense of ownership encompasses both the rendered outcome and the process through which the design becomes their own.

However, our findings suggest that requiring a sketch before a direction has emerged can constrain exploration to forms designers can already articulate visually, which may be particularly demanding for less experienced sketchers [57]. Unexpected details in generated output can prompt designers to reinterpret an initial intention [18], which Tversky et

al. describe as linking perceptual reorganizations to new conceptual meanings [55]. Future systems could therefore pair granular sketch editing with less constrained generation: when unsure what to draw next, a designer could describe a direction in language and leave its visual realization to the model, keeping control over what returns to the sketch.

7.2 Sketch and Provenance as Resources for Creative Agent Harnesses

Our findings suggest that a mutable sketch and its provenance could provide working state and contextual memory for a creative agent harness. The sketch is a concrete design object through which intent is expressed and revised, and CLEO shows why an agent needs to read direct manipulation in context, distinguishing corrections and demonstrations from independent work [50]. A pixel workflow leaves little to read. Intent lives in a prompt, the result is a finished image, and the only way to intervene is to rewrite the prompt and generate again, discarding the previous result rather than revising it. Applying a property from a reference to a named part is a different kind of act: it has arguments a model can fill, a target it can name, a result it can inspect, and a record it leaves behind. The operations *MoMix* already gives designers are, in other words, shaped like tool calls.

That shape admits three degrees of involvement. An agent could watch which references a designer keeps returning to and surface one when it becomes relevant, or propose a property for a part they have left untouched. It could go further and read an exploratory pattern out of the records themselves: which property a designer varies first, how far they usually travel from a reference, when they branch rather than overwrite. Such a pattern is not a single edit but a way of working, and once named it could be offered back as a reusable skill, carried to a new board or borrowed from someone who explores differently. Hammad et al.'s account of prospective uses of creative activity traces points the same way, toward turning these records into new experiments rather than only into an archive [23]. At the furthest reach, an agent could take a turn in the mixing itself, applying a property and leaving the designer to accept, adjust, or revert it. Building on JarvisHub's shared agent state [35], a harness could coordinate history retrieval, part-level editing, and feedback. Because every part stays editable and every application is recorded, an agent's move is no more final than a designer's.

The risk is that an agent grounded in a designer's own history is best at continuing what they were already doing, which is the opposite of what the open-ended exploration above calls for. Authoring tools for AI-bridged narrative meet the same tension and answer it by treating an author's material as a way to shape a space of possibilities rather than a template to copy, whether by deriving that space from example stories at a chosen level of abstraction [37] or by making it visible enough to see where generation is drifting [60]. A creative agent on a mutable anchor would need the same move: to read a designer's accumulated choices as a description of the space they have been working in, and then to propose somewhere in that space they have not yet been.

7.3 Limitations and Future Work

Our study was conducted in a short-term lab setting and therefore could not fully examine the longer-term value of *MoMix*'s provenance features, particularly the region-based history and branching timeline. Future longitudinal studies could investigate whether sketches as mutable anchors continue to support designers when they return to a reference-rich board after days or weeks, for example by helping them recover relationships between references and evolving design ideas, revisit prior decisions, and navigate the board more effectively.

Our study also primarily simulated asynchronous collaboration rather than examining real-time co-design. Future work could explore how mutable anchors support synchronous negotiation and collaborative experimentation, where designers may contribute different reference properties to the same evolving design. For example, one designer may

propose a shape while another introduces a color or material. Because *MoMix* provides immediate visual feedback on these part-level transformations, it may support rapid negotiation and quick prototyping around a shared design object. Such settings could further reveal how mutable anchors help reduce the semantic ambiguity of moodboarding canvases [38], make collaborators' design ideation traces more legible, and support collaborative creativity through a shared design imaginary.

Our evaluation focused on two object-based design scenarios, an ocean-inspired office and a forest-inspired café, which may not capture the full range of design practices where mutable anchors could be useful. Future work could examine more complex design settings and further explore the potential of SVG as a portable and programmable design form. In *MoMix*, sketch parts remain editable SVG geometry and preserve links to their reference-driven transformations, making them suitable for reuse across boards and collaborators. *MoMix* already generates shape as SVG code rather than as pixels. Beyond static visual properties such as color, material, and shape, future systems could extend that code generation to dynamic properties, such as extracting motion from video references and applying it to design elements. SVG may also support experimentation across different scales of granularity, since the same representation can be manipulated from whole compositions down to fine-grained parts while retaining provenance.

Our evaluation also relied largely on participants' self-assessments of communication effectiveness. Although participants reported greater confidence in explaining their own designs and understanding and continuing collaborators' designs, we did not independently evaluate how clearly these intentions were actually communicated to others. Future work could involve external designers or expert judges to assess the clarity, interpretability, and continuability of communicated design ideas. Similarly, while our evaluation of final outcomes combined participants' perceived novelty with LPIPS-based measures of visual distance from references, perceptual distance does not directly measure creativity. Future studies could complement these measures with independent expert or crowdsourced evaluations of outcome novelty, quality, and overall creativity.

8 Conclusion

This work explored how sketches as mutable anchors can support designers in experimenting with visual references and navigating evolving design ideas on infinite canvases. Through *MoMix*, designers can engage in polymorphic, part-level reference experimentation while preserving provenance across design states. Our study shows how mutable anchors can support designer-originated variation, cognitive flexibility before commitment, and ownership over the construction process, and how they can carry a shared design imaginary across stages of an evolving design. More broadly, this mosaic-like design form can serve as a harness for AI-assisted visual creation, supporting more designer-initiated visual thinking with references.

References

[1] Adobe. 2026. Firefly Boards: Your infinite visual workspace. Retrieved August 27, 2026 from https://www.adobe.com/products/firefly/features/boards.html.

[2] Rahul Arora, Ishan Darolia, Vinay P. Namboodiri, Karan Singh, and Adrien Bousseau. 2017. SketchSoup: Exploratory Ideation Using Design Sketches. *Computer Graphics Forum* 36, 8 (2017), 302–312. doi:10.1111/cgf.13081

[3] Albert Bandura. 2006. Guide for Constructing Self-Efficacy Scales. In *Self-Efficacy Beliefs of Adolescents*, Frank Pajares and Timothy C. Urdan (Eds.). Information Age Publishing, 307–337.

[4] Karim Benharrak and Amy Pavel. 2026. HistoryPalette: Supporting Exploration and Reuse of Past Alternatives in Image Generation and Editing. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 689, 18 pages. doi:10.1145/3772318.3790742

[5] John Brooke. 1996. SUS: A 'Quick and Dirty' Usability Scale. In *Usability Evaluation in Industry*, Patrick W. Jordan, Bruce Thomas, Bernard A. Weerdmeester, and Ian L. McClelland (Eds.). Taylor & Francis, London, 189–194.

[6] Adam R. Carberry, Hee-Sun Lee, and Matthew W. Ohland. 2010. Measuring Engineering Design Self-Efficacy. *Journal of Engineering Education* 99, 1 (2010), 71–79. doi:10.1002/j.2168-9830.2010.tb01043.x

[7] Kathy Charmaz. 2008. Constructionism and the Grounded Theory Method. In *Handbook of Constructionist Research*, James A. Holstein and Jaber F. Gubrium (Eds.). Guilford Press, New York, 397–412.

[8] Kathy Charmaz. 2014. *Constructing Grounded Theory* (2 ed.). SAGE Publications.

[9] Erin Cherry and Celine Latulipe. 2014. Quantifying the Creativity Support of Digital Tools through the Creativity Support Index. *ACM Transactions on Computer-Human Interaction* 21, 4 (2014), 21:1–21:25. doi:10.1145/2617588

[10] 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 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 1055, 25 pages. doi:10.1145/3613904.3642794

[11] Herbert H. Clark and Susan E. Brennan. 1991. Grounding in Communication. In *Perspectives on Socially Shared Cognition*, Lauren B. Resnick, John M. Levine, and Stephanie D. Teasley (Eds.). American Psychological Association, 127–149. doi:10.1037/10096-006

[12] Anil R. Doshi and Oliver P. Hauser. 2024. Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content. *Science Advances* 10, 28 (2024), eadn5290. doi:10.1126/sciadv.adn5290

[13] David H. Douglas and Thomas K. Peucker. 1973. Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica* 10, 2 (1973), 112–122. doi:10.3138/FM57-6770-U75U-7727

[14] Boris Ewenstein and Jennifer Whyte. 2009. Knowledge Practices in Design: The Role of Visual Representations as 'Epistemic Objects'. *Organization Studies* 30, 1 (2009), 7–30. doi:10.1177/0170840608083014

[15] Ronald A. Finke and Karen Slayton. 1988. Explorations of creative visual synthesis in mental imagery. *Memory & Cognition* 16, 3 (1988), 252–257. doi:10.3758/BF03197758

[16] Kexue Fu, Jingfei Huang, Long Ling, Sumin Hong, Yihang Zuo, RAY LC, and Toby Jia-Jun Li. 2026. Vistoria: A Multimodal System to Support Fictional Story Writing through Instrumental Image-Text Co-Editing. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 440, 26 pages. doi:10.1145/3772318.3790400

[17] Kexue Fu, Jiaye Leng, Yawen Zhang, Jingfei Huang, Yihang Zuo, Runze Cai, Zijian Ding, RAY LC, Shengdong Zhao, and Qinyuan Lei. 2026. Spatial Balancing: Designing an LLM-Powered Spatial Externalization Interface for Iterative Science Communication Writing. In *Proceedings of the 2026 Designing Interactive Systems Conference (DIS '26)*. Association for Computing Machinery, New York, NY, USA, 2935–2954. doi:10.1145/3800645.3812998

[18] Kexue Fu, Ruishan Wu, Yuying Tang, Yixin Chen, Bowen Liu, and RAY LC. 2024. "Being Eroded, Piece by Piece": Enhancing Engagement and Storytelling in Cultural Heritage Dissemination by Exhibiting GenAI Co-Creation Artifacts. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 2833–2850. doi:10.1145/3643834.3660711

[19] Vinod Goel. 1995. *Sketches of Thought*. MIT Press, Cambridge, MA.

[20] Gabriela Goldschmidt. 1991. The Dialectics of Sketching. *Creativity Research Journal* 4, 2 (1991), 123–143. doi:10.1080/10400419109534381

[21] Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2010. Chronicle: Capture, Exploration, and Playback of Document Workflow Histories. In *Proceedings of the 23rd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. Association for Computing Machinery, New York, NY, USA, 143–152. doi:10.1145/1866029.1866054

[22] Alicia Guo, David Ledo, George W. Fitzmaurice, and Fraser Anderson. 2026. Protosampling: Enabling Free-Form Convergence of Sampling and Prototyping through Canvas-Driven Visual AI Generation. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 383, 27 pages. doi:10.1145/3772318.3791884

[23] Noor Hammad, David Chuan-En Lin, Amy Smith, Max Kreminski, Erik Harpstead, and Jessica Hammer. 2026. Tracing Creativity: A Design Space for Creative Activity Traces in HCI. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 1644, 24 pages. doi:10.1145/3772318.3791263

[24] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9

[25] Paul Heckbert. 1982. Color image quantization for frame buffer display. *ACM SIGGRAPH Computer Graphics* 16, 3 (1982), 297–307. doi:10.1145/965145.801294

[26] Scarlett R. Herring, Chia-Chen Chang, Jesse Krantzler, and Brian P. Bailey. 2009. 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 (CHI '09)*. Association for Computing Machinery, New York, NY, USA, 87–96. doi:10.1145/1518701.1518717

[27] Emmanuel Iarussi, Adrien Bousseau, and Theophanis Tsandilas. 2013. The Drawing Assistant: Automated Drawing Guidance and Feedback from Photographs. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology (UIST '13)*. Association for Computing Machinery, New York, NY, USA, 183–192. doi:10.1145/2501988.2501997

[28] David G. Jansson and Steven M. Smith. 1991. Design fixation. *Design Studies* 12, 1 (1991), 3–11. doi:10.1016/0142-694X(91)90003-F

[29] Janin Koch, Andrés Lucero, Lena Hegemann, and Antti Oulasvirta. 2019. May AI?: Design Ideation with Cooperative Contextual Bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19)*. Association for Computing Machinery, New York, NY,

USA, Article 633, 12 pages. doi:10.1145/3290605.3300863

[30] Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy E. Mackay. 2020. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1, Article 45 (2020), 27 pages. doi:10.1145/3392850

[31] Janin Koch, Nicolas Taffin, Andrés Lucero, and Wendy E. Mackay. 2020. SemanticCollage: Enriching Digital Mood Board Design with Semantic Labels. In *Proceedings of the 2020 ACM Designing Interactive Systems Conference (DIS '20)*. Association for Computing Machinery, New York, NY, USA, 407–418. doi:10.1145/3357236.3395494

[32] 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 (IUI '23)*. Association for Computing Machinery, New York, NY, USA, 264–277. doi:10.1145/3581641.3584095

[33] Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen I. Brafman, and Scott R. Klemmer. 2010. Designing with Interactive Example Galleries. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. Association for Computing Machinery, New York, NY, USA, 2257–2266. doi:10.1145/1753326.1753667

[34] David Chuan-En Lin, Hyeonsu B. Kang, Nikolas Martelaro, Aniket Kittur, Yan-Ying Chen, and Matthew K. Hong. 2025. Inkspire: Supporting Design Exploration with Generative AI through Analogical Sketching. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 427, 18 pages. doi:10.1145/3706598.3713397

[35] Yunlong Lin, Zixu Lin, Zhaohu Xing, Biqiang Li, Chenxin Li, Haonan Wang, Haitao Wu, Hengyu Liu, Jianghai Chen, Kaituo Feng, Kaixin Li, Shawn Chen, Shijue Huang, Sixiang Chen, Tsung-Yi Ho, Wenxuan Huang, Xiangyan Liu, Xiaomeng Hu, Xuanhua He, Yan Sun, Yunqing Zhao, Zhiqin Yang, Zehan Wang, Zhengyang Tang, Tianyu Pang, and Xiangyu Yue. 2026. JarvisHub: An Open Harness for Canvas-Native Multimodal Creative Agents. arXiv:2607.23588 [cs.CV] https://arxiv.org/abs/2607.23588

[36] Lovart. 2026. Lovart: The World's First AI Design Agent. Retrieved August 27, 2026 from https://www.lovart.ai/.

[37] Zhuoran Lu, Qian Zhou, and Yi Wang. 2025. WhatELSE: Shaping Narrative Spaces at Configurable Level of Abstraction for AI-bridged Interactive Storytelling. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, 1–18. doi:10.1145/3706598.3713363

[38] Andrés Lucero. 2012. Framing, Aligning, Paradoxing, Abstracting, and Directing: How Design Mood Boards Work. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. Association for Computing Machinery, New York, NY, USA, 438–447. doi:10.1145/2317956.2318021

[39] Nicolai Marquardt, Asta Roseway, Hugo Romat, Payod Panda, Michel Pahud, Gonzalo A. Ramos, Steven Mark 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* (Busan, Republic of Korea) *(UIST '25)*. Association for Computing Machinery, New York, NY, USA, Article 159, 19 pages. doi:10.1145/3746059.3747631

[40] Catherine C. Marshall, Frank M. Shipman III, and James H. Coombs. 1994. VIKI: Spatial Hypertext Supporting Emergent Structure. In *Proceedings of the 1994 ACM European Conference on Hypermedia Technology (ECHT '94)*. Association for Computing Machinery, New York, NY, USA, 13–23. doi:10.1145/192757.192759

[41] Xiaohan Peng, Debanjana Haldar, Wendy E. Mackay, and Janin Koch. 2026. DesignTrace: Exploring, Iterating and Tracking Design Alternatives with GenAI. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*. Association for Computing Machinery, New York, NY, USA, Article 434, 22 pages. doi:10.1145/3772318.3791036

[42] Xiaohan Peng, Janin Koch, and Wendy E. Mackay. 2024. DesignPrompt: Using Multimodal Interaction for Design Exploration with Generative AI. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference (DIS '24)*. Association for Computing Machinery, New York, NY, USA, 804–818. doi:10.1145/3643834.3661588

[43] Xiaohan Peng, Janin Koch, and Wendy E. Mackay. 2025. FusAIn: Composing Generative AI Visual Prompts Using Pen-based Interaction. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 874, 20 pages. doi:10.1145/3706598.3714027

[44] Miquel Prats, Sungwoo Lim, Iestyn Jowers, Steve W. Garner, and Scott Chase. 2009. Transforming Shape in Design: Observations from Studies of Sketching. *Design Studies* 30, 5 (2009), 503–520. doi:10.1016/j.destud.2009.04.002

[45] Urs Ramer. 1972. An iterative procedure for the polygonal approximation of plane curves. *Computer Graphics and Image Processing* 1, 3 (1972), 244–256. doi:10.1016/S0146-664X(72)80017-0

[46] Nathalie Riche, Anna Offenwanger, Frederic Gmeiner, David Brown, Hugo Romat, Michel Pahud, Nicolai Marquardt, Kori Inkpen, and Ken Hinckley. 2025. AI-Instruments: Embodying Prompts as Instruments to Abstract & Reflect Graphical Interface Commands as General-Purpose Tools. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 1104, 18 pages. doi:10.1145/3706598.3714259

[47] Donald A. Schön. 1992. Designing as Reflective Conversation with the Materials of a Design Situation. *Research in Engineering Design* 3, 3 (1992), 131–147. doi:10.1007/BF01580516

[48] Donald A. Schön and Glenn Wiggins. 1992. Kinds of seeing and their functions in designing. *Design Studies* 13, 2 (1992), 135–156. doi:10.1016/0142-694X(92)90268-F

[49] Xinyu Shi, Yinghou Wang, Ryan Rossi, and Jian Zhao. 2025. Brickify: Enabling Expressive Design Intent Specification through Direct Manipulation on Design Tokens. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery,

New York, NY, USA, Article 424, 20 pages. doi:10.1145/3706598.3714087

[50] Kihoon Son, Hyewon Lee, DaEun Choi, Yoonsu Kim, Tae Soo Kim, Yoonjoo Lee, John Joon Young Chung, HyunJoon Jung, and Juho Kim. 2026. "When to Hand Off, When to Work Together": Expanding Human-Agent Co-Creative Collaboration through Concurrent Interaction. arXiv:2603.02050 [cs.HC] https://arxiv.org/abs/2603.02050

[51] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 644, 26 pages. doi:10.1145/3613904.3642400

[52] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. Association for Computing Machinery, New York, NY, USA, Article 1, 18 pages. doi:10.1145/3586183.3606756

[53] Masaki Suwa and Barbara Tversky. 1996. What Architects See in Their Sketches: Implications for Design Tools. In *Conference Companion on Human Factors in Computing Systems (CHI '96)*. Association for Computing Machinery, New York, NY, USA, 191–192. doi:10.1145/257089.257255

[54] Masaki Suwa and Barbara Tversky. 1997. What do architects and students perceive in their design sketches? A protocol analysis. *Design Studies* 18, 4 (1997), 385–403. doi:10.1016/S0142-694X(97)00008-2

[55] Barbara Tversky, Masaki Suwa, Maneesh Agrawala, Julie Heiser, Chris Stolte, Pat Hanrahan, Doantam Phan, Jeff Klingner, Marie-Paule Daniel, Paul Lee, and John Haymaker. 2003. Sketches for Design and Design of Sketches. In *Human Behaviour in Design: Individuals, Teams, Tools*, Udo Lindemann (Ed.). Springer, Berlin, Heidelberg, 79–86. doi:10.1007/978-3-662-07811-2_9

[56] Ilse M. Verstijnen, Cees van Leeuwen, Gabriela Goldschmidt, Ronald Hamel, and James M. Hennessey. 1998. Creative discovery in imagery and perception: combining is relatively easy, restructuring takes a sketch. *Acta Psychologica* 99, 2 (1998), 177–200. doi:10.1016/S0001-6918(98)00010-9

[57] Ilse M. Verstijnen, Cees van Leeuwen, Gabriela Goldschmidt, Ronald Hamel, and James M. Hennessey. 1998. Sketching and creative discovery. *Design Studies* 19, 4 (1998), 519–546. doi:10.1016/S0142-694X(98)00017-9

[58] Samangi Wadinambiarachchi, Ryan M. Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The Effects of Generative AI on Design Fixation and Divergent Thinking. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 380, 18 pages. doi:10.1145/3613904.3642919

[59] Qian Wan and Zhicong Lu. 2023. GANCollage: A GAN-Driven Digital Mood Board to Facilitate Ideation in Creativity Support. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference (DIS '23)*. Association for Computing Machinery, New York, NY, USA, 136–146. doi:10.1145/3563657.3596072

[60] Yi Wang, John Joon Young Chung, Melissa Roemmele, Yuqian Sun, Tiffany Wang, Shm Garanganao Almeda, Brett A. Halperin, Yuwen Lu, and Max Kreminski. 2026. Elsewise: Authoring Open-ended Interactive Narrative with Possibility Space Visualization. arXiv:2601.15295 [cs.HC] https://arxiv.org/abs/2601.15295

[61] Jun Xie, Aaron Hertzmann, Wilmot Li, and Holger Winnemöller. 2014. PortraitSketch: Face Sketching Assistance for Novices. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology (UIST '14)*. Association for Computing Machinery, New York, NY, USA, 407–417. doi:10.1145/2642918.2647399

[62] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The Unreasonable Effectiveness of Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*. IEEE, 586–595. doi:10.1109/CVPR.2018.00068

A Appendix

A.1 Demographic Information of Participants

Table 4. Background of participants in the user study. GenAI/Visual AI reports participants' frequency of using generative AI generally for creative or design activities and specifically for generating visual inspiration, respectively. Moodboard reports familiarity with moodboards and frequency of use.

ID	Gender	Age	Current Role	Design Exp.	Design Areas	GenAI / Visual AI Use	Moodboard Familiarity / Use
P1	M	28	Design enthusiast	1–3 yrs	Graphic, Illustration	Daily / Weekly	Very familiar / Rarely
P2	F	31	Design researcher	4–6 yrs	UI/UX, Graphic	Monthly / Monthly	Very familiar / Monthly
P3	F	29	Design researcher	7–10 yrs	Graphic, Illustration	Daily / Rarely	Extremely familiar / Almost every project
P4	F	25	Design student	1–3 yrs	UI/UX	Daily / Monthly	Slightly familiar / Monthly
P5	F	24	Design student	4–6 yrs	Illustration, Motion, Spatial	Monthly / Monthly	Extremely familiar / Almost every project
P6	F	33	Design researcher	>10 yrs	UI/UX, Product, Spatial	Monthly / Rarely	Very familiar / Almost every project
P7	F	33	Design enthusiast	7–10 yrs	Graphic, Illustration	Rarely / Rarely	Moderately familiar / Rarely
P8	M	29	Design enthusiast	1–3 yrs	UI/UX	Monthly / Monthly	Moderately familiar / Monthly
P9	F	25	Design enthusiast	<1 yr	UI/UX	Monthly / Rarely	Slightly familiar / Rarely
P10	F	26	Design researcher	4–6 yrs	UI/UX, Product	Daily / Weekly	Very familiar / Rarely
P11	F	24	Design enthusiast	<1 yr	UI/UX, Motion	Daily / Daily	Slightly familiar / Rarely
P12	M	28	Design enthusiast	4–6 yrs	Graphic, Illustration	Monthly / Monthly	Slightly familiar / Rarely

A.2 Design Communication Survey

Part I: Communicating your own design

Think about the board you developed as something you would show to a teammate or client who did not watch you make it.

1 The board and the design output capture the visual style and mood I intended.

2 Someone who did not watch us work could understand my design from the board alone, even when it also contained my collaborator's work.

3 Someone else could pick up this board and carry my design direction forward.

Part II: Understanding and building on an inherited design

Now think about the board a colleague had started, which you picked up and developed.

4 I can understand the design intent behind the board and design output that was handed to me.

5 I can keep developing someone else's board without losing the direction they set.

6 I can merge my own ideas with the existing design into one coherent board.

Table 5. Design-communication self-efficacy questionnaire. Each item was rated on a scale from 0 (*cannot do at all*) through 5 (*moderately certain can do*) to 10 (*highly certain can do*), reporting how confident participants were, right now, that they could do it [3].

A.3 User Study Results Appendix

Table 6. NASA-TLX and SUS results ($N = 12$).

		MoMix		Baseline		p
Measure	Dimension	Mean	SD	Mean	SD	
NASA-TLX	Mental	4.92	2.54	3.75	2.14	**.039**
	Physical	3.08	2.68	2.08	2.07	.063
	Temporal	2.58	1.31	2.92	1.51	.219
	Performance	4.67	2.81	4.33	2.67	.711
	Effort	4.25	2.83	3.42	2.15	.374
	Frustration	2.42	2.07	2.25	2.26	.906
SUS	Total	65.83	16.66	77.71	13.79	**.027**