Shöwn: Adaptive Conceptual Guidance Aids Example Use in Creative Tasks

Tricia J. Ngoon
tngoona@ucsd.edu
Design Lab, UC San Diego
La Jolla, California, USA

Joy O. Kim
joykim@adobe.com
Adobe Research
San Francisco, California, USA

Scott Klemmer
srk@ucsd.edu
Design Lab, UC San Diego
La Jolla, California, USA

Figure 1: Shöwn’s interface and features. The *drawing helper*: a) A user wants to draw a comic panel showing a table and requests the drawing helper’s assistance by typing a sticky note that says, “Show me a table”. b) The drawing helper then provides a basic icon of the requested object in the requested location. *Adaptive conceptual guidance* c) As the user draws their panel, Shöwn adaptively suggests the user to “consider the framing of this panel” on the right side of the screen in the form of a sticky note. d) The user clicks on the sticky note to view relevant examples. The user can then return to their drawing screen.

ABSTRACT

Examples are powerful tools for creativity and can provide inspiration and structure. However, novices often don’t know when to seek examples or how to apply them. Showing real-time examples targeted to the current task (adaptive conceptual guidance) may

help novices consider inspiration more often and better implement the ideas that examples illustrate. We explore this in a Wizard-of-Oz system, *Shöwn*, that presents examples based on the user’s current activity while drawing a comic strip. *Shöwn*’s design is informed by interviews with novice and expert comic artists ($n=18$). A between-subjects experiment ($n=24$) found that adaptive conceptual guidance improved the clarity and uniqueness of drawings and stories compared to non-adaptive examples. Users also found the examples more useful and inspirational than those without adaptive guidance. Our results present an initial exploration of the challenges and benefits of contextually showing examples in interactive creativity tools.

This work is licensed under a [Creative Commons Attribution International 4.0 License](https://creativecommons.org/licenses/by/4.0/).

DIS '21, June 28–July 2, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8476-6/21/06.
<https://doi.org/10.1145/3461778.3462072>

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**;
Empirical studies in HCI.

KEYWORDS

creativity, examples, wizard-of-oz, comics

ACM Reference Format:

Tricia J. Ngoon, Joy O. Kim, and Scott Klemmer. 2021. Shown: Adaptive Conceptual Guidance Aids Example Use in Creative Tasks. In *Designing Interactive Systems Conference 2021 (DIS '21), June 28-July 2, 2021, Virtual Event, USA*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3461778.3462072>

1 EXAMPLES PROVIDE STRUCTURE FOR CREATIVE WORK

A perennial challenge for novices is figuring out how to get started on a creative endeavor. Viewing and adapting others' work can help novices find traction in a mountain of possibility. Ward [54] defines using examples as "structured imagination," modifying an existing solution and applying it to something new. Good examples can give inspiration, unveil new ideas, and highlight the advantages of alternative paths [8, 48]. Examples can also help people compare and contrast different ideas across examples, drawing attention to important features (such as the composition or color scheme used by a painting) that a creator might apply to their own work [3, 11, 22, 33, 39].

Today, with the proliferation of online communities and galleries, real examples of creative work are highly accessible. Communities like Behance¹ and HitRecord² present examples in searchable galleries that enable users to browse, modify, and implement ideas from creators with diverse backgrounds and levels of expertise, bolstering exploration and inspiration [31, 36, 37]. Tools such as iMovie³ and Adobe Spark⁴ leverage expertise through provided templates or expert patterns. Novices in particular benefit from templates because they can shape their ideas and work to the high-level structure provided by the template [32, 34, 56]. For more procedural tasks, video tutorials serve as accessible examples that give guidance and important factual information for achieving concrete goals [52] such as knitting a tri-colored scarf or sketching a human face. Presenting examples enables people to follow in the footsteps of others, giving a point of comparison and achievement for their own work.

1.1 The Challenges of Applying Examples

In order to use examples effectively, users must be able to assess the context and concept of the suggested example being presented. For example, if a novice photographer sees an example of a sunset photo with an off-center composition, they may not understand how the framing or composition affects the overall photo quality. Presenting examples shares similarities to search interfaces where a primary challenge is how to select and present the most relevant example or result [25, 40]. Even when examples are present and

available, novices especially may not understand what examples are most relevant or how to implement the ideas within them, instead replicating surface details [29]. Without requisite domain knowledge, discerning how to build upon examples without simply replicating them can be difficult.

While examples can give inspiration, they can also paradoxically constrain creativity. Examples that are too semantically different from the target goal might actually be harmful for ideation [8, 9]. Timing of examples also matters; examples presented too late are no more beneficial for creative outcomes than not viewing examples at all [33, 48]. Additionally, when given specific examples, novices tend to conform to salient surface features of examples and apply them to their own outcomes, causing potential fixation and decreasing novelty of ideas [28, 33, 39, 49]. Similar to work comparing experts and novices, novices tend to focus on the details of examples rather than on the underlying concepts [13]. An important consideration for presenting examples is choosing the right examples at the right moments.

1.2 Examples as Reflective Dialogue

Schön [47] presents a scenario in which an expert guides a novice architect to reframe the given problem, using example sketches to demonstrate key concepts tailored to where the novice needs the greatest assistance. The expert in this scenario uses examples to support concepts that fit the novice's situation and where they need the most help. Presenting the right concepts and supporting examples at the right time can lead to a reflective dialogue between a creator and their work in which different alternatives are explored and evaluated for their relative benefits. Similarly, human tutoring studies show how tutors use interactive dialogue to identify knowledge gaps and tailor using examples to help learners better understand problems [12, 42]. Adaptively using examples as a form of support can not only help novices solve a specific problem, but also understand the conceptual reasoning underlying the example.

1.3 Computational Aid for Applying Examples

Given these challenges of examples presentation and use, how might creativity support tools better help novices understand how to implement the ideas that examples can provide? Computationally-generated guidance can take a contextual approach, presenting examples tailored to a user's current task or creative intent [2, 30, 34, 36]. Two exemplar systems demonstrate this approach for both visual and text-based work. CritiqueKit uses past expert critique as examples and suggests these examples as a reviewer gives feedback based on whether the feedback fits certain criteria [43]. This approach uses text-based context to determine whether feedback fits structural characteristics to adapt what type of examples are most helpful for a novice reviewer. In visual work, adaptive templates for photography suggest gridlines to help novice photographers apply composition concepts to their photos, using visual context as a basis for suggesting composition options [16]. We extend computational guidance approaches by adapting not only the examples shown but also the *concepts* they embody to help novices see past the details and better understand and apply the insights within examples.

¹<https://www.behance.net/>

²<https://hitrecord.org/>

³<https://www.apple.com/imovie/>

⁴<https://spark.adobe.com/>

1.4 Adaptive Conceptual Guidance: Examples in Context

This paper investigates *adaptive conceptual guidance*, which presents in-situ conceptual examples and helps users apply insights from them. We differentiate adaptive conceptual guidance from example galleries and template libraries [31, 36, 38, 50] by taking a proactive approach to showing high-level domain principles (such as framing and transitions) at specific points in time as a user works to suggest relevant options that they may not notice or consider. Similar to Schön's scenario of an expert giving both concepts and specific examples, we hypothesize that contextually presenting examples alongside the concepts they encompass will aid example use and creative outcomes.

We investigate this hypothesis through implementing a Wizard-of-Oz prototype called *Shōwn* that uses adaptive conceptual guidance to help novices explore and apply examples in the domain of comic drawing. We choose comics as a domain because it involves multimedia storytelling and does not require special equipment, making it a useful exemplary domain for studying the challenges related to creative decision-making when multiple media (such as writing, drawing, and story pacing) are involved. To unearth concrete novice and expert differences for this domain, we conducted participant observations with nine expert comic artists and nine novices. We found that novices, in contrast to experts, focus on low-level detail when asked to create a comic strip and struggle to explore alternative compositions and story ideas.

We evaluated adaptive conceptual guidance in a between-subjects experiment where participants created a comic strip. This experiment ($n = 24$) compared comics created by participants using one of two versions of *Shōwn*: a) a version that provided an examples gallery for users to browse on their own and b) a version that provided adaptive conceptual guidance to illustrate possible ways of framing or composing comic panels. Adaptive conceptual guidance enabled participants to create comics that were rated by readers as having more unique and clear stories and drawings than comics created without adaptive guidance (Figure 4). Raters also preferred comics created with adaptive conceptual guidance overall. Participants stated in interviews that guidance inspired them to consider different ways of showing their story, lending support to the hypothesis that adaptive conceptual guidance assists exploration. Most participants without adaptive conceptual guidance either did not view the examples or did not find them useful, further supporting the hypothesis that novices need guidance for knowing how and when to implement examples to their own work.

To summarize, this paper contributes the following:

- an interview study that demonstrates how novices focus on details and struggle with exploring alternative ways to show their comic story,
- a Wizard-of-Oz system that contextualizes and adaptively suggests examples and concepts at particular moments based on a user's activity, and
- empirical evidence that adaptive conceptual guidance can help novices better utilize examples to improve their creative outcomes.

2 INTERVIEWS: WHERE DO NOVICES STRUGGLE MOST IN DRAWING COMICS?

2.1 Method

Differences in experts and novices have been well-studied in various domains from chess [10] to physics [13]; experts are able to contextually apply different high-level strategies to their work while novices often focus on details. We wanted to understand where novices struggle most in the domain of drawing comics to learn about the types of examples that would be most beneficial to support in an adaptive creativity support tool. We interviewed and observed nine expert artists and nine novices as they performed a comic-drawing task. We chose comics as a domain because most people would likely benefit from examples when engaging in the intersection of verbal and visual storytelling [1, 18, 41] and because drawing comics requires only pen and paper at a minimum.

We recruited experts from the freelance site Upwork, and novices through social media posts. Experts had formal art training and professional experience in making comic strips. We recruited novices who had a familiarity with comics, but no experience making comics and no formal education or professional experience in art. Over the course of an hour, participants drew a four-panel comic in a think-aloud protocol [20] for one of three randomly given prompts:

- PENGUIN: a man befriends a penguin because of shared interests,
- ALIENS: aliens invade Earth during the 2020 global pandemic quarantine, and
- CHARACTER: a woman is upset at her date because he doesn't know who her favorite television or movie character is.

As participants drew, we asked questions about their process. After participants finished drawing, we asked them what they found easiest and hardest about making their comics, what the feel is the most important part of comics, and what type of help they feel would be most useful. All interviews took place on the Zoom video conferencing platform, and participants were allowed to draw digitally or use pen and paper. Experts received USD \$60 payment for their participation; novices received a \$25 gift card.

2.2 Results

2.2.1 Experts Focus on Story to Yield Interesting & Readable Comics. All nine expert artists expressed that story and composition are the most important parts of a comic: "*It's all about composition because if you can't understand how a comic moves, the narrative kind of decomposes*" (Expert 4). Experts stressed that good drawing skills alone do not make a comic interesting, explaining that writing is more critical to a good comic than the fidelity of its drawings: "*I feel like the most important part is the composition of each panel because you can have great illustration..., but if the composition leads to just a random object in the background that isn't important to the story..., the audience is going to miss something*" (Expert 2). Expert 8 similarly said, "*The most important thing is the writing because if you're a good artist, but your writing stinks, you're not going to make it. But you can have awesome writing and stick figures and make a million bucks, so to speak.*"

To achieve their narrative goals, experts mentioned using composition concepts and rules to help establish the flow and message

Figure 2: Examples of an a) expert and b) novice comic for the ALIENS prompt: aliens invade Earth during the 2020 global pandemic quarantine.

of the comic. Expert 5 stated, “There’s pretty simple rules of composition, you know, thirds, rhythm, all that stuff. They kind of help lead you to the subject’s eyelines.” As in prior work on expert drawing strategies [17, 45], experts drew rough sketches to rapidly outline their ideas. These quick sketches helped them evaluate their work in terms of visual concepts like composition, as well as narrative concepts like flow and pacing. For example, Expert 6 said, “So I just plan out how I want it to look. And, you know, once you’re done with something you’re always like, ‘Oh, I can make that better. I can make it nicer.’ To get the general gist, we’re going to see what they’re doing, and etcetera.” Expert 1 similarly explained, “And at this point, I’m still kind of planning stuff out. So I’m not really drawing what might be there, just kind of a reference. So I guess laying it out.” Three experts talked about typically creating thumbnail sketches to ideate panel compositions: “I’ll do like four or five different little sketches or thumbnails of what I think are good for the theme that I chose in my brainstorming session” (Expert 7). After a rough sketch, experts refine their sketch by adding in details, then move into a final inking stage to finish drawing their comic. Expert 5 said that doing a final drawing was the quickest part of the process because the idea in mind was already on paper. In general, experts spent most of their time on exploration and iteration while focusing on high-level elements like composition and flow, enabling them to finish drawing their comics quickly and simply.

2.2.2 Novices Focus on Details Rather Than Story. Verbally, novices also expressed the view that narrative arc is paramount: e.g., “[The most important thing is] conveying the story I want to tell in as clear a way as possible” (Novice 8). Despite this, novices time allocation favored detailed drawings over composition and flow. For example, one novice drew the details of a character before deciding on story plot or composition. Another spent nearly 10 minutes drawing a

person starting from the feet up in one panel, leading them to rush through subsequent panels: “I’ve spent like 10 minutes on this, and I’m only on the feet!” (Novice 2). In contrast to experts, novices often went straight into drawing their final comic with little conceptual planning or exploration. They were often hesitant to use sketches as exploratory tools and instead wrote out their ideas when conceptually planning their work: “I know how to communicate my idea through words. That’s what I’m familiar with. I’m not as comfortable with drawing” (Novice 1). Conveying a clear story in the comic was the most important goal for all participants, but novices did not have an understanding of how to show their story visually. Figure 2 shows an example of an expert and novice comic drawn for the ALIENS story prompt.

Interestingly, despite novices noting that narrative was the most important aspect of a comic, all novices attributed their inability to clearly tell a story to their lack of drawing skills. Novice 2 said, “I can have an idea in my head, but actually translating it onto a comic is difficult.” Similarly, Novice 3 expressed that, “I am able to visualize all of this very easily in my mind’s eye but everything else is kind of blank... It’s hard to translate that into a specific thing to draw because it’s not told to me what that might be.” When asked what help would be most useful for making their comic, seven of the nine novices said drawing help: “I think generating the figures is something that could be automated, and [a computer] could have done a much better job than me” (Novice 7). Another expressed that drawing help would reduce tedium: “I find drawing to be actually pretty tedious...if I could just come up with the panel and the script and give it someone or the computer that would be great” (Novice 6). The other two novices expressed a desire for templates or reference images to help them show their comic story.

The perceived lack of skill seemed to be a major factor in how novices chose compositions for their comic panels. Novice 2 said, *“I’m thinking about what I want - like the image in my mind is literally what is on the square. And the way I decide what is on the square is because of my lack of drawing ability, it’s sort of like what is the simplest way that I can convey the message that I’m trying to set.”* Exploring beyond the most achievable composition for a panel was also difficult: *“After you’ve decided on how to present the information or the previous information, and I feel like I don’t want to be repetitive..., and I didn’t know what would be the best way to draw my panel”* (Novice 4). These observations highlighted two primary challenges for novices: 1) exploring ways to visually show their story idea, and 2) executing on that idea through drawing. Experts are able to quickly evaluate and decide upon alternatives through rough sketching whereas novices often fixate on making their single idea look an precise way. As Jodi Picoult states, *“You can edit a bad page, but you can’t edit a blank page.”*

3 GUIDING COMICS WITH SHOWN

Previous research in expert and novice differences find that novices focus on surface details and fixate on ideas [10, 13, 28, 56]. Likewise, our observations showed that novices struggle with exploring alternative ideas and focusing excessively on details while drawing comics. We hypothesize that a system that helps novices consider and choose among options for *what* to draw by presenting timely and relevant concepts and examples can help novices overcome these challenges as they work.

Schön states that reflective practice is a continuous reframing and experimentation of ambiguous problems [47]. Inspired by this articulation of reflective practice, we designed Shown to explore our hypothesis. The Shown Wizard-of-Oz system augments Google Jamboard, a web-based collaborative drawing tool, with two scaffolds for presenting examples: a drawing helper and adaptive conceptual guidance (Figure 1). We chose a Wizard-of-Oz method for two reasons: to study interaction recognition and contextual heuristics for presenting examples without the cost of full implementation, and to better observe how people used and applied conceptual examples [53]. Because Google Jamboard is collaborative, the Wizard and the user can simultaneously see the other’s activity by visiting the same link. This allows the Wizard to recognize the user’s activity and provide assistance accordingly.

The Shown interface comprises a drawing area for each comic panel, alongside selected examples. Users can flesh out panels in any order. Following the drawing screens are a corpus of examples, adapted from McCloud’s book [41], that show comic-specific concepts related to comic planning such as framing, transitions, and image and word combinations (Figure 1d). To communicate with the system, users can use voice commands or Jamboard’s sticky notes feature to place notes on drawing screens. Sticky notes as an input mechanism allowed natural language text descriptions and deictic, “put that there” interactions [19].

3.1 Adaptive Conceptual Guidance: Showing Relevant Examples for Exploration

We also found in interviews that, while novices enjoyed coming up with their own ideas and stories, they often did not know how to

best visualize their ideas. To address this, Shown provides adaptive conceptual guidance by first presenting high-level comic principles as a suggestion and then specific examples as the user works to inspire ideas for *how* to portray their story. The guidance that Shown currently provides is based on our interviews with experts and three concepts used by experts to decide how to convey a comic story [1, 18, 41]: choice of framing, choice of moment or transition, and choice of images and words. McCloud [41], a comic expert and artist, writes that these concepts are part of the *“rough planning stage where a story’s events are first broken down into readable chunks”* (pg. 11). When presenting guidance, Shown pairs an explanation of the concept with examples from McCloud’s book that show how the concept can be applied to a comic story. (The concept of narrative flow is also mentioned in the book, but this refers to panel layout and spacing, which were not applicable to our study as we asked all participants to use the same four-panel layout.) Highlighting the concept behind examples reflects the scenario of an expert guiding a novice through a problem by explaining high-level concepts and providing specific instantiations of the concepts [47].

The timing of guidance is determined by the following set of recognition heuristics with three principles in mind based on prior work: examples should be related to what the user is currently doing [2, 30, 34, 36, 43], examples should be shown as early in the creative process as possible [33, 48], and examples should be paired with an explanation of the underlying concept they illustrate [29, 43].

The three types of guidance that Shown supports are:

Consider the framing for this panel. Framing refers to composition or use of camera angle and distance to show pertinent actions or details in the panel. This guidance is often presented on the first panel before a user begins drawing to help them get started with their comic. If the user is drawing a subsequent panel, Shown presents this guidance if the panel’s framing is repeated from any previous panels (e.g. if the user begins to use a medium shot in the second panel after using the same camera angle/shot in the first panel).

Consider the moment or transition between panels. Moment or transition refer to the continuity between actions in panels and the specific moments shown as part of the story in each panel. This guidance is presented when a user moves to a new panel to help them think of transitions between panels.

Consider the combination of images and words in this panel. The combination of images and words refers to how words and images together can tell story more powerfully than either alone. This guidance is presented if the images between panels are similar, if a panel contains a lot of or no text or dialogue, or if the user is attempting to draw facial expressions.

3.2 Adaptive Conceptual Guidance in Action

While the user is drawing their comic, the Wizard is simultaneously using the same prototype link to provide guidance to the user. As the user is drawing their comic, Shown adaptively displays one of the three guidance suggestions based on the user’s current task and context. The Wizard displays guidance ambiently [40] on the right side of the panel page without interruption to the user via the sticky notes tool (Figure 1c). For instance, a user is drawing the

same composition for two panels, in which two people are sitting at a table in the center of the page. Because of the repeated framings, the Wizard presents guidance for the user to “*consider the framing of this panel*” in the form of a sticky note that is added to the right side of the drawing screen (Figure 1c) to provide inspiration for alternative ways of composing the panel. When the user sees the suggestion and clicks on the sticky note, the Wizard displays a specific example screen that shows examples corresponding to the concept being presented (Figure 1d). The user can then return to their drawing screen and modify their drawing if they choose. When the user moves to the next panel, the Wizard presents guidance for the user to “*consider the moment or transition between panels*” before the user begins to draw their panel. After clicking on the sticky note, the Wizard presents examples showing different ways to use pacing in a comic narrative. Users can also proactively ask for guidance at any point by clicking the “Show me a suggestion” sticky note on the right side of any drawing screen. Users can also ignore guidance if they choose.

3.3 The Drawing Helper: Freeing Focus from Detail

In our interview study, novices experienced difficulty in executing their ideas and overly focused on details due to a lack of confidence in their drawing skills. To mitigate this and enable a larger focus on story, Shown provides a drawing helper that can automatically render a basic sketch of a requested object using Google Jamboard’s Autodraw feature, which presents icons that match the sketch input. We differentiate these icons from examples presented in adaptive conceptual guidance as these are basic renderings of objects rather than conceptual examples. If the options provided by Autodraw were not sufficient, the human Wizard provided a simple drawn sketch. The drawing helper was designed to draw basic objects such as a table or person rather than complex characters or scenes. For instance, a user wants to draw a comic story about two people having a dinner date. They want to draw the couple sitting at a table, but don’t want to focus on the details of the table itself. The user asks the drawing helper for a table by making a sticky note with the text “*show me a table,*” and placing it in the center of the drawing screen (Figure 1a). The Wizard then produces an icon of the table in the location specified by the user (Figure 1b). Later, the user wants a drawing of a chair and this time requests an example via voice command by saying “show me a chair.” Shown renders an icon of a chair, and the user can move the chair to where they desire on the screen. The drawing helper provides simple drawing objects so the user can then focus their attention on utilizing the conceptual examples that support the high-level narrative aspects of their comic story.

4 EXPERIMENT: (HOW) DOES ADAPTIVE CONCEPTUAL GUIDANCE HELP?

4.1 Method

We wanted to test the hypothesis that a creativity support tool that offers adaptive conceptual guidance would help novices better utilize examples and improve creative outcomes than a tool that only provides non-adaptive examples.

4.1.1 Participants. We recruited 24 novice participants (9 male, 15 female, median age = 22) through social media, mailing list advertisements, and the site UserTesting to evaluate our Wizard-of-Oz system. Similar to our interviews, we recruited novices that had familiarity with comics, but no formal education in art or professional experience in making comics. We also required participants to draw using a digital tablet with a stylus for ease of drawing. Half the participants were randomly assigned to the *non-adaptive* condition while the other half was assigned to the *adaptive* condition. The task asked participants to make a rough draft of a four-panel comic based on a randomly-assigned prompt from the interview study (PENGUIN, ALIENS, or CHARACTER prompts). Four participants from each condition saw each of the three prompts. We encouraged black-and-white comics for simplicity, but using color was allowed. All one-hour sessions took place on the Zoom video conferencing platform, and participants received a USD \$25 gift card for their participation.

4.1.2 Procedure. At the start of each study session, participants received a link to Shown on Google Jamboard and were given a tutorial from the experimenter about Shown’s features. *Adaptive* condition participants created their comic in a version of Shown with adaptive conceptual guidance where conceptual examples were presented selectively based on the heuristics described earlier. *Non-adaptive* condition participants used an otherwise identical version of Shown without adaptive conceptual guidance (the same conceptual examples were available via the example screens, but not selectively presented, similar to previous work with example galleries). All participants had access to the drawing helper because our main research question was about adaptive conceptual guidance versus non-adaptive examples. We chose to include the drawing helper for both conditions because Google Jamboard’s Autodraw feature, which is the engine behind the drawing helper’s icons, already exists within the tool itself. Participants spent up to an hour drawing their comic while using a think-aloud protocol [20]. At the end of the study session, participants answered post-interview questions about their process and thoughts on Shown’s scaffolds. Following the study, we debriefed participants to inform them about the Wizard-of-Oz nature of the study.

4.1.3 Measures. We measured participants’ behaviors and interactions with Shown’s features, such as the number of times participants used the drawing helper, viewed examples, and verbally expressed implementing concepts from the examples. We determined whether participants implemented examples either by their verbal expression of doing so or if they changed their panel after viewing examples. We also measured creative outcomes by gathering quality ratings of each comic from lay readers. We chose lay readers as raters (rather than experts) because lay readers are a large audience for comics, and because we wanted to capture overall story impact rather than technical quality, in line with the focus of both experts and novices in our interview study on narrative clarity. Specifically, we measured comics in terms of clarity, creativity, and overall preference. In pairwise comparisons, raters from Mechanical Turk (61 unique Mechanical Turk workers) were sequentially presented with a pair of comics (two comics for the same prompt, one from each condition). For each pair, raters selected which comic they believed had more unique drawings and a

Figure 3: Mechanical Turk workers rated *adaptive* and *non-adaptive* comics in pairwise comparisons in our between-subjects experiment. This figure shows which comic received the most preference ratings for: a) clear drawings, b) unique drawings, c) clear story, and d) unique story and overall preference.

more unique story (creativity), which comic had clearer drawings and a clearer story (clarity), and which comic they preferred overall. We separated drawing measures from story measures to account for potential individual differences in drawing ability. Pairing one comic from each condition for each prompt yielded 48 unique pairs (four *non-adaptive* comics $\times$ four *adaptive* comics $\times$ three prompts). Workers were paid USD \$1 to rate five pairs in sequence. Every pair received evaluations from at least five different raters, resulting in 306 total ratings.

4.2 Results

4.2.1 Raters Preferred Adaptive Condition Comics Overall and for Most Other Measures. Raters preferred *adaptive* comics significantly more often than *non-adaptive* comics overall (*adaptive*=57.8%, *non-adaptive*=42.2%, $\chi^2 = 7.87, df = 1, p < .01$) (Figure 1). In terms of clarity, raters preferred *adaptive* comics more often when asked which had clearer drawings (*adaptive*=55.6%, *non-adaptive*=44.4%, $\chi^2 = 3.91, df = 1, p < .05$) and which had clearer stories (*adaptive*=56.2%, *non-adaptive*=43.8%, $\chi^2 = 5.12, df = 1, p < .05$). *Adaptive* comics were also more often preferred when raters were asked which had more unique stories (*adaptive*=56.2%, *non-adaptive*=43.8%, $\chi^2 = 5.12, df = 1, p < .05$). There was no significant difference in preference ratings for which comic had more unique drawings (*adaptive*=53.9%, *non-adaptive*=46.1%, $\chi^2 = 2.19, df = 1, p = .14$).

Figure 3 shows the most preferred comic for each measure as rated by the MTurk workers.

4.2.2 Adaptive Conceptual Guidance Made Examples More Useful. Shown presented guidance at least three times to each *adaptive* participant per session with a maximum of six times ($M = 3.92, SD = 0.90$). All *adaptive* participants viewed at least one example screen at least one time while drawing, with an average of three times per session ($SD = 1.13$). *Adaptive* participants verbally expressed implementing concepts from examples an average of 2.17 times per session ($SD = 1.11$). In contrast, three *non-adaptive* condition participants viewed examples, with an average of 0.67 examples viewed per session ($SD = 1.23$) ($t = 4.84, df = 21.8, p < .001$).

Adaptive participants mentioned that guidance served as inspiration, helping them think of new ways of showing their story they would not have otherwise realized. For example, P5 (*adaptive*) noted, “As I was going through I think I got into the same rhythm of what I wanted to show, and I think the [guidance] made me more cognizant of ways that I could switch things up. Like I don’t think I would have thought of this last panel if I didn’t get a nudge from those.” P16 (*adaptive*), whose comic was the most preferred overall and for the most unique story, similarly stated, “I think the [guidance] helped me consider different things. For example, for the 3rd panel, I wouldn’t have considered doing a shot from the back of his head to show [the woman’s] point of view” (Figure 5). For the *non-adaptive* participants who did view the examples, they also served as inspiration. P21

Figure 4: Adaptive condition comics received significantly more preference ratings for more unique stories, clearer drawings, clearer story, and overall preference. $*p < .05$, $p < .01$**

(non-adaptive), whose comic was most preferred for having unique drawings, also cited the examples as helpful for their comic: “[The examples] were kind of like a guide to give me a vocabulary for my intentions.”

However, some non-adaptive participants did not believe the examples were applicable. For example, P18 (non-adaptive) stated, “I already had an idea in mind so I didn’t really need the examples.” P15 (non-adaptive) forgot the examples were available but wished he had used them after finishing his comic: “I think I should have looked at [the examples] before I started drawing.” Similarly, although participants could proactively ask for guidance in the adaptive condition, only two participants did so. One participant mentioned that he “just wanted to see what suggestions were available” (P14, adaptive).

4.2.3 Participants Wanted More Understandable Heuristics for Adaptive Conceptual Guidance. The Wizard uses specific heuristics for displaying guidance. However, the heuristics shown used to decide when to show guidance, even if slightly off, could diminish the effectiveness of examples. Most felt the real-time guidance was useful: “[The suggestions] were like a guidance for me. I don’t think I would go look at [the examples], so it was nice that they were shown throughout the drawing part” (P10, adaptive). However, one participant wished guidance was presented at the beginning of drawing: “I think these come a bit too late for me to use them because they’re sparsely in between while I’m drawing. For example, the perspective one, I thought ‘oh, I should’ve used this earlier,’ but the information just came a bit too late for me to use” (P6, adaptive). This suggests that guidance heuristics should not be “one-size-fits-all” but rather allow users to adjust them based on perceived usefulness.

Some adaptive participants also thought the timing and reasoning for presenting certain guidance was opaque. They wanted to know more about why particular advice was being given at a particular moment. P24 (adaptive) stated, “The suggestions were helpful, but it wasn’t clear why that suggestion was being given. Like for the transition one, it could be that my transition isn’t the best one I could use for my viewer or it could just be helping me get started so it wasn’t clear why it was giving me those suggestions.” While adaptive conceptual guidance helped participants better understand examples, clarity around what the system is responding to would help users decide between many possible valid uses of those examples.

4.2.4 The Drawing Helper Shifted Focus from Details to Story. We tracked user interactions with the drawing helper to understand how people used this feature and how it possibly interacted with example use. Participants in both conditions used the drawing helper frequently, with an average of 4.67 times ($SD = 2.87$) per session for adaptive participants and 5.67 times ($SD = 4.68$) for non-adaptive participants ($t = -0.63$, $df = 18.3$, $p = .54$). In some cases, the drawing helper made drawing objects—particularly repeated objects—less tedious. For example, P24 (adaptive) used the drawing helper to draw a phone screen that would be consistent across all panels. P11 (non-adaptive) felt the drawing helper did indeed help him focus less on details: “I thought [the drawing helper] was really cool, especially for like tables or chairs or to hash out the overall setting so you can just focus on the other stuff.” P21 (non-adaptive) thought the drawing helper was useful for showing her intended vision: “I could just tell [the drawing helper] my vision, and it would bring up things that are relevant.” Our findings suggest that the concreteness

Figure 5: a) This comic was created using Shown's adaptive conceptual guidance. b) The participant cited that seeing a relevant example inspired the composition of the third panel of the comic.

provided by tools like the drawing helper may complement the efficacy of strategies like adaptive conceptual guidance.

However, a few participants felt limited in their creativity with the drawing helper. One participant noted an over-reliance on the drawing helper rather than sketching, stating, "I kinda became too dependent on [the drawing helper]. Normally, I would just draw it out quickly, but I want to see what the system can do...With the tool, I think if the system can do this for me, and it's perfect the first time, I try to see how I can fit my idea to what the system can do" (P15, *non-adaptive*). Another participant felt that the simplicity of the drawing helper's sketches did not allow them to completely show their vision: "It was pretty limited in what it could show. Like I couldn't do perspectives or angles with the basic icons I was given" (P12, *adaptive*). These observations corroborate the fact that we found no significant difference in visual creativity across the two conditions ($\chi^2 = 2.20, df = 1, p = .14$). Three participants did not use the drawing helper or only used it once because, as one participant stated, "It might have been more helpful if I was drawing something more complex, but it was just easier to draw it on my own" (P13, *non-adaptive*).

5 DISCUSSION

This paper investigated whether providing adaptive conceptual guidance aids example use more effectively than non-adaptive examples through a Wizard-of-Oz prototype and empirical study.

We next discuss implications of our findings for future creativity support tools and generalization to other domains.

5.1 System-Directed Versus User-Directed Creativity Support

In our experiment, we observed that novices in both conditions rarely sought examples without being prompted, perhaps because they were unaware of needing help or did not know what help they needed. Only two *adaptive* condition participants proactively asked for guidance, and only three *non-adaptive* condition participants viewed the examples gallery on their own. While exploring alternatives is a useful problem-solving strategy, meaningfully evaluating them requires expertise and the comparison induces a hefty cognitive load [51]. In line with prior work [6, 23, 48], this suggests that timing is relevant to giving the most appropriate help as novices may not know what type of help is most helpful and when. By incorporating concepts with examples, we found that Shown's real-time adaptive conceptual guidance helped to expand participants' field-of-view to concepts they would not have known to explore on their own. Such proactive system-directed approaches seem to benefit open-ended creative problems (like figuring out how to draw a comic story) because they help users explore beyond known possibilities.

Similar to other systems that present galleries of examples to provide broad inspiration [31, 37, 38], Shown provides high-level examples meant for conceptual exploration rather than examples

that are particular to the user's specific work. For example, if a user wants to draw two people talking, Shown presents general examples of different panel framings or combinations of text and images rather than specific instances of two people talking in a panel. These context-agnostic examples are framed as considerations (rather than statements that impose specific approaches as the "correct" ones). This open-ended approach is analogous to human tutoring, where less didactic interactions allow the learner to form their own explanations and hypotheses in evaluating ideas [14]. In addition, recommending guidance in-situ allowed participants to better understand the relevance of guidance and evaluate examples in context [47]. One participant explicitly mentioned that showing guidance in the moment while drawing encouraged them to view the examples when they would not have otherwise. Participants cited that even if they did not change their initial idea, the provided examples served as inspiration or as a way to help them explain their choices. In these instances adaptive conceptual guidance seemed to serve as a nudge for users to consider different options or evaluate their own decisions.

In contrast, user-directed approaches (such as query-based systems and other help-seeking tools [2, 27, 46]) seem to work best when the user has a concrete idea they want to execute. We saw this with how participants made use of Shown's drawing helper, which produced basic icons to quickly visualize ideas based on a user's request. Some used voice commands to direct the drawing helper to provide images. However, the drawing helper had its limits: it could not generate drawings in response to complex, ambiguous requests. For example, P15 (*non-adaptive*) asked the drawing helper to lay out an entire panel: "Show me a man holding a phone with a dating app on it and lists for interests." This participant attempted to shape their ideas around what the drawing helper could generate, rather than using it to show their own ideas. This reflects how novices may form an over-reliance on help-tools or conform to salient features rather than high-level considerations [28, 29, 39, 49]. Our findings with Shown suggest that user-directed guidance works best for small and concrete tasks where the user already has an idea in mind.

Both system-directed and user-directed guidance can support the creative process in complementary ways. Context-agnostic conceptual aid allows for the ambiguity of exploration while specific user-directed guidance can give a concrete direction towards accomplishing goals. In many ways, using examples as guidance is similar to coaching or tutoring, where a coach or tutor gives both open-ended hints or suggestions and concrete examples depending on a person's progress. Shown provides a demonstration of using context from user actions to leverage such guidance for creative work. Future research could further examine how computational systems can provide more adaptive coaching or tutoring for creative endeavors.

5.2 Designing Adaptive Conceptual Guidance Systems: Limitations & Opportunities

We found two areas where adaptive conceptual guidance could be improved. The first is in determining the timing of guidance. One participant felt that guidance was presented too late and would be more helpful or more relevant earlier. While guidance was shown

adaptively, the heuristics for determining when to show certain concepts and examples were static for each participant, which is both a limitation of the Wizard-of-Oz implementation as well as potential technical implementations. One open question is how to better anticipate when to provide guidance when even the user does not know when to ask for help. Potential effective timing mechanisms might be measuring how long a user is idle [7, 48] or taking a mixed-initiative approach in incorporating user input, such as refining or pruning suggestions [30] to determine the best moments to present examples.

The second area of improvement is providing transparency around why guidance was being given. Three participants wanted further explanation for why certain pieces of guidance were being shown at specific times. One participant asked the experimenter whether the system presented guidance randomly or not. Another thought the presentation of guidance meant the system did not "like" their current drawing. This sentiment reflects prior work showing that people prefer explainable and transparent interactions with AI systems so they not only understand what help is being given but also why [4, 26, 44]. One possibility is to display guidance alongside interactive checkboxes or dynamic rubrics to help users understand how to better situate guidance within the user's context [5, 43].

We also found potential in improving user-directed support. While participants did not often explicitly seek out examples while working on their comics, many participants in both our interviews and experiment asked ideation questions aloud such as "How do I show this place is empty?" or "How do I show that this character is angry?" as part of the think-aloud protocol. These questions could serve as queries to indicate user intent and aid in tailoring the type of examples shown. During exploratory tasks like brainstorming ideas or deliberating ways to translate an idea visually, users may have more of these open-ended queries in mind rather than concrete goals. Tools that adaptively display help for just-in-time learning as well as incorporate contextual search mechanisms like natural language and deictic instructions may be useful ways of enabling user-directed conceptual support [21, 24, 35, 55].

As a Wizard-of-Oz prototype, Shown utilizes a human Wizard to use a user's actions as heuristics to determine what conceptual guidance to provide. Some of the heuristics were based simply on starting on new panels (*i.e.* the Wizard shows the guidance to consider moments or transitions when the user moves on to a new panel). Other heuristics were based on capturing the user's past and current actions (*i.e.* the Wizard presents guidance on images and words if the panel contained too much or too little text). Existing tools already use sketch-based actions as context such as Procreate's Quick Shape tool⁵ to automatically complete shapes or straighten a user's lines or Google Jamboard's Autodraw tool that uses object recognition to infer what icons to suggest to the user. Co-creative intelligent agents can also use object and line recognition to improvise collaborative drawing [15]. Our Wizard-of-Oz evaluation shows how adaptive systems might use visual, sketch-based context to provide conceptual guidance beyond automated drawing help or example galleries. Open remaining questions are how systems might be trained to learn such context as well as what

⁵<https://procreate.art/handbook/procreate/guides/quickshape//>

types of context and concepts are most appropriate for adaptive and automated assistance beyond sketching.

One limitation of Shown was that the example screens were separate from the drawing screens, requiring users to switch between drawing and example screens in order to view them. Tools incorporating adaptive conceptual guidance or examples presentation should consider more ambient displays that show examples without requiring the user to change the context of their work. Because of the remote nature of the study, one limitation was that we could not control for the screen size participants used. This may have affected drawing ability, though this limitation may have been mitigated since all participants used a stylus and had the drawing helper available.

Tools that help novices overcome skill barriers to execute their creative goals are powerful, but without seeing potential alternatives and possibilities, novices remain bounded by their limited conceptual expertise. Shown evaluates simple recognition heuristics through a Wizard-of-Oz implementation to contextually present concepts and examples for better exploration and execution. Future work could examine feasibility and implementation of such heuristics for creativity systems across domains. We find that combining user-led support for concrete queries and adaptive support for open-ended exploration of alternatives helps novices better understand and utilize alternatives for their own work. An interleaving of system-led support and user-led agency is a promising direction for the development and evaluation of future creative tools [26]. We provide a demonstration of this direction that can apply to creative endeavors across domains.

6 CONCLUSION

We present Shown, a Wizard-of-Oz system that provides adaptive conceptual guidance by suggesting relevant concepts and examples for a user to explore based on their current task. We hypothesized that adaptive conceptual guidance would guide the application of examples and improve creative work more than simply providing static examples alone. Through interviews and a between-subjects experiment in the domain of comics, we found that adaptive conceptual guidance led to comics with more unique and clear stories and clearer drawings. Users also found adaptive conceptual guidance to be timely and useful for inspiration while users without guidance did not find the examples as relevant and were less likely to take advantage of them. We argue for tools that help novices explore ideas by providing the right kinds of assistance in the right situations. In this way, we may better expand novices' vantage points to explore concepts more broadly and improve creative work. This work provides a direction for the future of adaptive creativity support tools and presentation of examples.

ACKNOWLEDGMENTS

We thank all our expert and novice participants for the time and effort. We also thank Jane E and Mira Dontcheva for their insightful feedback on this work. This work was funded in part by Adobe Research.

REFERENCES

- [1] Jessica Abel and Matt Madden. 2008. *Drawing Words and Writing Pictures: Making Comics: Manga, Graphic Novels, and Beyond*. Macmillan.
- [2] C Ailie Fraser, Tricia J Ngoon, Mira Dontcheva, and Scott Klemmer. 2019. RePlay: Contextually Presenting Learning Videos Across Software Applications. (2019). <https://doi.org/10.1145/3290605.3300527>
- [3] Louis Alfieri, Timothy J Nokes-Malach, and Christian D Schunn. 2013. Learning through case comparisons: A meta-analytic review. *Educational Psychologist* 48, 2 (2013), 87–113. <https://doi.org/10.1080/00461520.2013.775712>
- [4] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, Jaime Teevan, Ruth Kikin-Gil, Eric Horvitz Microsoft, Paul G Allen, and Adam Fourney. 2019. Guidelines for Human-AI Interaction. (2019), 13. <https://doi.org/10.1145/3290605.3300233>
- [5] Aditya Bharadwaj, Pao Siangliulue, Adam Marcus, and Kurt Luther. 2019. Critter: Augmenting Creative Work with Dynamic Checklists, Automated Quality Assurance, and Contextual Reviewer Feedback. (2019), 12. <https://doi.org/10.1145/3290605.3300769>
- [6] Merim Bilalić, Peter McLeod, and Fernand Gobet. 2008. Why good thoughts block better ones: The mechanism of the pernicious Einstellung (set) effect. *Cognition* 108, 3 (2008), 652–661.
- [7] Joel Chan, Steven P Dow, and Christian D Schunn. 2018. Do the best design ideas (really) come from conceptually distant sources of inspiration? (2018). <https://doi.org/10.1016/j.destud.2014.08.001>
- [8] Joel Chan, Katherine Fu, Christian Schunn, Jonathan Cagan, Kristin Wood, and Kenneth Kotovsky. 2011. On the benefits and pitfalls of analogies for innovative design: Ideation performance based on analogical distance, commonness, and modality of examples. *Journal of mechanical design* 133, 8 (2011). <https://doi.org/10.1115/1.4004396>
- [9] Joel Chan, Pao Siangliulue, Denisa Qori Mcdonald, Ruixue Liu, Reza Moradinezhad, Safa Aman, Erin T Solovey, Krzysztof Z Gajos, and Steven P Dow. 2018. Semantically Far Inspirations Considered Harmful? Accounting for Cognitive States in Collaborative Ideation. (2018). <https://doi.org/10.1145/3059454.3059455>
- [10] William G Chase and Herbert A Simon. 1973. *Perception in Chess*. Technical Report. [https://doi.org/10.1016/0010-0285\(73\)90004-2](https://doi.org/10.1016/0010-0285(73)90004-2)
- [11] Min Chi, Ilsa Dohmen, Jonathan T Shemwell, Doris B Chin, Catherine C Chase, and Daniel L Schwartz. 2012. *Seeing the Forest from the Trees: A Comparison of Two Instructional Models Using Contrasting Cases*. Technical Report.
- [12] Michelene T.H Chi, Stephanie A Siler, Heisawn Jeong, Takashi Yamauchi, and Robert G Hausmann. 2001. Learning from human tutoring. *Cognitive Science* 25, 4 (jul 2001), 471–533. [https://doi.org/10.1016/S0364-0213\(01\)00044-1](https://doi.org/10.1016/S0364-0213(01)00044-1)
- [13] Michelene T. H. Chi, Paul J. Feltovich, and Robert Glaser. 1981. Categorization and Representation of Physics Problems by Experts and Novices*. *Cognitive Science* 5, 2 (apr 1981), 121–152. https://doi.org/10.1207/s15516709cog0502_2
- [14] Michelene T H Chi and Marguerite Roy. 2010. How Adaptive Is an Expert Human Tutor? *LNC5* 6094 (2010), 401–412.
- [15] Nicholas Davis, Chih Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *International Conference on Intelligent User Interfaces, Proceedings IUI*, Vol. 07-10-March-2016. Association for Computing Machinery, New York, New York, USA, 196–207. <https://doi.org/10.1145/2856767.2856795>
- [16] Jane L E, Ohad Fried, Jingwan Lu, Jianming Zhang, Radomir Měch, Jose Echevarria, Pat Hanrahan, and James A Landay. 2020. Adaptive Photographic Composition Guidance. (2020). <https://doi.org/10.1145/3313831.3376635>
- [17] Betty Edwards. 1979. *Drawing on the Right Side of the Brain* (4th ed.). Penguin Books Ltd.
- [18] Will Eisner. 2008. *Comics and sequential art: Principles and practices from the legendary cartoonist*. WW Norton & Company.
- [19] Douglas C Engelbart and William K English. 1968. A research center for augmenting human intellect. In *Proceedings of the December 9-11, 1968, fall joint computer conference, part I*. 395–410. <https://doi.org/10.1145/1476589.1476645>
- [20] K Anders Ericsson and Herbert A Simon. 1984. *Protocol analysis: Verbal reports as data*. the MIT Press.
- [21] C Ailie Fraser, Julia M Markel, N James Basa, Mira Dontcheva, and Scott Klemmer. 2020. ReMap: Lowering the Barrier to Help-Seeking with Multimodal Search. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*. 979–986. <https://doi.org/10.1145/3379337.3415592>
- [22] Dedre Gentner, Jeffrey Loewenstein, and Leigh Thompson. 2003. Learning and transfer: A general role for analogical encoding. *Journal of educational psychology* 95, 2 (2003), 393. <https://doi.org/10.1037/0022-0663.95.2.393>
- [23] Arthur C Graesser and Natalie K Person. 1994. Question Asking During Tutoring. *American Educational Research Journal Spring* 31, 1 (1994), 104–137. <https://doi.org/10.3102/00028312031001104>
- [24] Arthur C Graesser, Kurt Vanlehn, Carolyn P Rosé, Pamela W Jordan, and Derek Harter. 2001. Intelligent Tutoring Systems with Conversational Dialogue. *AI Magazine* 22, 4 (dec 2001), 39–52. <https://doi.org/10.1609/AIMAG.V22I4.1591>
- [25] Marti Hearst. 2009. *Search user interfaces*. Cambridge university press.
- [26] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. 116, 6 (2019). <https://doi.org/10.1073/pnas.1807184115>

- [27] Forrest Huang and John F Canny. 2019. *Sketchforme: Composing Sketched Scenes from Text Descriptions for Interactive Applications*. Technical Report. <https://doi.org/10.1145/3332165.3347878> arXiv:1904.04399v1
- [28] David G. Jansson and Steven M. Smith. 1991. Design Fixation. *Design Studies* 12, 1 (1991), 3–11. [https://doi.org/10.1016/0142-694X\(91\)90003-F](https://doi.org/10.1016/0142-694X(91)90003-F)
- [29] Elahe Javadi. 2012. *The impact of user interface design on idea integration in electronic brainstorming: an attention-based view*. Ph.D. Dissertation. University of Illinois at Urbana-Champaign.
- [30] Sean Kandel, Andreas Paepcke, Joseph Hellerstein, and Jeffrey Heer. 2011. Wrangler: Interactive visual specification of data transformation scripts. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3363–3372. <https://doi.org/10.1145/1978942.1979444>
- [31] Hyeonsu B. Kang, Gabriel Amoako, Neil Sengupta, and Steven P. Dow. 2018. Paragon: An Online Gallery for Enhancing Design Feedback with Visual Examples. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems - CHI '18*. ACM Press, New York, New York, USA, 1–13. <https://doi.org/10.1145/3173574.3174180>
- [32] Joy Kim, Mira Dontcheva, Wilmot Li, Michael S Bernstein, and Daniela Steinsapir. 2015. Motif: Supporting Novice Creativity through Expert Patterns. (2015). <https://doi.org/10.1145/2702123.2702507>
- [33] Chinmay Kulkarni, Steven Dow, and Scott Klemmer. 2012. Early and Repeated Exposure to Examples Improves Creative Work. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Vol. 34. https://doi.org/10.1007/978-3-319-01303-9_4
- [34] Ranjitha Kumar, Jerry O Talton, Salman Ahmad, and Scott R Klemmer. 2011. Bricolage: example-based retargeting for web design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2197–2206. <https://doi.org/10.1145/1978942.1979262>
- [35] Gierad P Laput, Mira Dontcheva, Gregg Wilensky, Walter Chang, Aseem Agarwala, Jason Linder, and Eytan Adar. 2013. Pixeltone: A multimodal interface for image editing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2185–2194. <https://doi.org/10.1145/2470654.2481301>
- [36] Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, and Scott R Klemmer. 2010. Designing with Interactive Example Galleries. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. <https://pdfs.semanticscholar.org/7c0c/1bef1e605fd5b8929ef623861ec59f930b35.pdf>
- [37] Martin Lee and Daniel L Schwartz. 2009. Prospective Adaptation in the Use of External Representations. *Cognition and Instruction* 27, 4 (2009), 370–400. <https://doi.org/10.1080/07370000903221775>
- [38] Joe Marks, Brad Andalman, Paul A Beardsley, William Freeman, Sarah Gibson, Jessica Hodgins, Thomas Kang, Brian Mirtich, Hanspeter Pfister, Wheeler Ruml, et al. 1997. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 389–400.
- [39] Richard L Marsh, Joshua D Landau, and Jason L Hicks. 1996. How examples may (and may not) constrain creativity. *Memory & cognition* 24, 5 (1996), 669–680. <https://doi.org/10.3758/BF03201091>
- [40] Justin Matejka, Tovi Grossman, and George Fitzmaurice. 2011. Ambient help. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2751–2760. <https://doi.org/10.1145/1978942.1979349>
- [41] Scott McCloud. 2006. *Making comics: Storytelling Secrets of Comics, Manga, and Graphic Novels*. Harper Collins.
- [42] Bruce M McLaren, Sung-Joo Lim, and Kenneth R Koedinger. 2008. When is assistance helpful to learning? Results in combining worked examples and intelligent tutoring. In *International conference on intelligent tutoring systems*. Springer, 677–680. https://doi.org/10.1007/978-3-540-69132-7_75
- [43] Tricia J. Ngoon, C. Ailie Fraser, Ariel S. Weingarten, Mira Dontcheva, and Scott Klemmer. 2018. Interactive Guidance Techniques for Improving Creative Feedback. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. <https://doi.org/10.1145/3173574.3173629>
- [44] Changhoon Oh, Jungwoo Song, Jinhan Choi, Seonghyeon Kim, Sungwoo Lee, and Bongwon Suh. 2018. I Lead, You Help But Only with Enough Details: Understanding the User Experience of Co-Creation with Artificial Intelligence. (2018). <https://doi.org/10.1145/3173574.3174223>
- [45] Henry Rankin Poore. 1967. *Pictorial Composition: An Introduction*. Sterling Publishing Co.
- [46] Arvind Satyanarayan and Jeffrey Heer. 2014. Authoring Narrative Visualizations with Ellipsis. *Computer Graphics Forum* 33, 3 (jun 2014), 361–370. <https://doi.org/10.1111/cgf.12392>
- [47] Donald A Schon. 1984. *The reflective practitioner: How professionals think in action*. Vol. 5126. Basic books.
- [48] Pao Siangliulue, Joel Chan, Krzysztof Z Gajos, and Steven P Dow. 2015. Providing Timely Examples Improves the Quantity and Quality of Generated Ideas. (2015). <https://doi.org/10.1145/2757226.2757230>
- [49] Steven M Smith, Thomas B Ward, and Jay S Schumacher. 1993. Constraining effects of examples in a creative generation task. *Memory & cognition* 21, 6 (1993), 837–845. <https://doi.org/10.3758/BF03202751>
- [50] Michael Terry and Elizabeth D Mynatt. 2002. Side views: persistent, on-demand previews for open-ended tasks. In *Proceedings of the 15th annual ACM symposium on User interface software and technology*. 71–80. <https://doi.org/10.1145/571985.571996>
- [51] Juhani E Tuovinen and John Sweller. 1999. A comparison of cognitive load associated with discovery learning and worked examples. *Journal of educational psychology* 91, 2 (1999), 334. <https://doi.org/10.1037/0022-0663.91.2.334>
- [52] Hans van der Meij, Peter Blijleven, and Leanne Jansen. 2003. What makes up a procedure. *Content & Complexity. Information design in technical communication* (2003), 129–186.
- [53] Jagoda Walny, Bongshin Lee, Paul Johns, Nathalie Henry Riche, and Sheelagh Carpendale. 2012. Understanding pen and touch interaction for data exploration on interactive whiteboards. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2779–2788. <https://doi.org/10.1109/TVCG.2012.275>
- [54] Thomas B Ward. 1994. Structured imagination: The role of category structure in exemplar generation. *Cognitive psychology* 27, 1 (1994), 1–40. <https://doi.org/10.1006/cogp.1994.1010>
- [55] Dongwook Yoon, Nicholas Chen, Bernie Randles, Amy Cheatle, Corinna E Löckenhoff, Steven J Jackson, Abigail Sellen, and François Guimbretière. 2016. RichReview ++ : Deployment of a Collaborative Multi-modal Annotation System for Instructor Feedback and Peer Discussion. (2016). <https://doi.org/10.1145/2818048.2819951>
- [56] Alvin Yuan, Kurt Luther, Virginia Tech, Markus Krause, Sophie Vennix, Steven P Dow, and Björn Hartmann. 2016. Almost an Expert: The Effects of Rubrics and Expertise on Perceived Value of Crowdsourced Design Critiques. (2016). <https://doi.org/10.1145/2818048.2819953>