Exploration and Optimization of Generative Variability in Future Work: A Mixed-Initiative Analysis

Michael Muller
IBM Research
Cambridge, MA, USA
michael_muller@us.ibm.com

Jessica He
IBM Research
Seattle, WA, USA
jessicahe@ibm.com

Justin D. Weisz
IBM Research
Yorktown Heights, NY, USA
jweisz@us.ibm.com

Abstract

Exploration and Optimization are ubiquitous activities in human work. Exploration allows us to consider multiple design or solution possibilities. Optimization allows us to narrow and refine diverse possibilities. Given that generative AI produces diverse outcomes for the same input – an attribute known as generative variability – there is now a need to re-examine Exploration and Optimization in the context of generative AI and understand their implications for the future of work. The relationship between Exploration and Optimization has traditionally been portrayed as independent, inverse, or even contradictory. Guided by concepts from mixed initiative interfaces, we conducted a conceptual analysis of three future-work-centered generative AI use cases. We show that the relationship between Exploration and Optimization is not "one-size fits all." We show how Exploration and Optimization may be combined into several integrated and inter-dependent work processes. Our use cases indicate that generative variability will increase the importance of Exploration and Optimization in future work tasks and applications.

CCS Concepts

• **Human-centered computing** → **Interaction design process and methods**; **Interaction design theory, concepts and paradigms**; • **Computing methodologies** → **Artificial intelligence**.

Keywords

Exploration, Optimization, AI, Generative variability, Future of Work, Design guidelines, Human-Centered AI, Generative AI

ACM Reference Format:
Michael Muller, Jessica He, and Justin D. Weisz. 2025. Exploration and Optimization of Generative Variability in Future Work: A Mixed-Initiative Analysis. In *CHIWORK '25: Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work (CHIWORK '25), June 23–25, 2025, Amsterdam, Netherlands.* ACM, New York, NY, USA, 13 pages. https://doi.org/10.1145/3729176.3729195



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

CHIWORK '25, Amsterdam, Netherlands
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-1384-2/25/06
https://doi.org/10.1145/3729176.3729195

1 Introduction

AI, and specifically generative AI (GenAI), have become important topics for human work and human workplaces. A recent summary report from the US National Academies suggests that GenAI automation has potential for future work productivity gains, job losses, and less well-defined job gains [69]. A second report from the University of Oxford proposes a tension between automation and augmentation ([83]; see also [81]), where augmentation implies "that the use of AI enhances human intelligence, rather than operating independently of, or replacing it" [40], potentially as a collaborative or even symbiotic human-AI "system" [108]. Giaccardi [29] focuses on the inclusion of "things" (i.e., AI agents) as partners in design, being capable of creating other "things" (i.e., design instances).

In this paper, we consider use cases toward the design of AI-supported augmentation of human work through the concept of *generative variability*. Weisz et al. [101, p.4] define generative variability as what happens when GenAI "generates artifacts as outputs and those outputs may vary in character or quality, even when a user's input does not change" (see also [65]). These attributes can be seen as a weakness - i.e., *indeterminacy* of outcomes - or as a strength - i.e., tunable or "steerable" [17, 57] *multiplicities* of outcomes. Generative variability challenges understandings of HCI and human factors because users may be faced with multiple outcomes without obvious criteria for what is "best." Part of our task is to turn these challenges into design strengths. Through the lens of generative variability, we consider three use cases for future workplaces: How can generative applications augment human work to produce, curate, select, and adapt instances from the large design spaces that are now possible through generative applications [46, 84, 101]? Among the challenges in this curatorial work, two activities stand out:

- Exploration of diverse exemplars produced from generative algorithms, and
- Optimization of selected exemplars.

Exploration and Optimization[1] have long been discussed in human cognition [35], neurobiology [4], and now the HCI and HCAI literatures [1, 8, 14, 39, 44, 59, 64, 95, 101, 109]. In cognition, Hills et al. [35] describe Exploration and Exploitation (a version of Optimization) as a trade-off that traces back to humans seeking sustenance for survival, and now applies to diverse problem spaces including searching for cures for disease or recommending songs in a music library. Traditionally, these concepts were treated as distinct [39, 95, 106, 107], or even antagonistic processes [18, 55].

[1] In this paper, we use ten words to identify formal concepts and we indicate our formal usage through capitalization. These concepts are: Exploration, Optimization, Ideate, Constrain, Produce, Suggest, Select, Curate, Assess, and Adapt (see Table 1).

Using formalisms from frameworks for mixed-initiative user interfaces [26, 30, 64, 67], we argue that Exploration and Optimization are not antagonistic, but intertwined. They are both necessary phases of diverse creative processes, but they may appear in different configurations and relationships with one another or with other processes. We discuss three use cases (based in the research literature) to explore that diversity using concepts from mixed-initiative frameworks to clarify the internal and interrelational actions and dynamics of Exploration and Optimization [26, 30, 67, 90].

We briefly note that the intended contribution of this paper is analytic rather than empirical. We don't analyze "experimental data" as such. Instead, we base our argument on recent design-oriented work in generative AI. Therefore, in place of a conventional "Results" section, we provide a series of three use cases, based in the research literature, which we analyze using concepts Exploration and Optimization. To illustrate these diverse relationships, we make use of an analytic framework from mixed-initiative creative interfaces [30, 64, 67, 90]. Our paper makes the following contributions:

- Analysis of the impact of generative AI on the future of work, focusing on the phenomonon of generative variability
- Analysis of the newly-important concepts of Exploration and Optimization in relation to generative variability, and in relation to future work generative applications and domains
- Further development of action vocabularies for mixed initiative creative interfaces, that can analyze and explicate Exploration, Optimization, and generative variability

2 Related Work

We begin with a more detailed exploration of themes of human augmentation and/vs. automation, followed by a survey of Exploration and Optimization across the HCI, Computer Science, and AI literatures, complemented by a discussion of those themes in the algorithmic literature. We then particularize these concepts in design frameworks, and add a more formal analytic approach based on action-vocabularies from mixed initiative interfaces.

2.1 Augmentation, not Automation

GenAI has been described as a kind of informal production system. Nielsen [70], for example, asserts that the new paradigm of generative AI allows the user to state their goal, rather than the steps required to achieve that goal. However, Bouschery et al. [8, p.140] object to that command-oriented use pattern:

> "Will designers be reduced to "prompt engineers", that is, just designing the task given to an AI instead of the product or component? Augmenting human innovation teams with AI hence calls for investigating many aspects of jointly working with AI..."

Raisch and Krakowski [81] also discuss the tension of automation and augmentation from a management science perspective, concluding that augmentation is associated with superior organizational performance. Zysman and Nitzberg [110, p.2] concur about the importance of human augmentation through AI:

> "projections based solely on automating existing tasks fail to capture the complex reorganizations [of human work] that are likely to happen... Adaptations will vary across contexts.. Maintaining the centrality of humans is likely to prove crucial..."

However, Nielsen's *what*-rather-than-*how* command-oriented usage pattern under-utilizes GenAI interactions. Nielsen comments that "I doubt we should even describe this user experience as an 'interaction' because there is no turn-taking or gradual progress." [70, unpaged]. His proposal seems to substitute a new form of command-line for an earlier form.

We believe that this type of interaction is only part of the promise of generative AI. Applications that leverage generative variability can also be used iteratively - i.e., conversationally - to augment human abilities with greater user control over what the algorithm produces. Bouschery et al. [8], Houde et al. [37], Muller et al. [63], and Weisz et al. [101] advocated for highly interactive methods to populate a design space of diverse exemplars, and then assist one or more humans to curate and refine selected exemplars from that design space [30, 37, 62, 64, 66]. For example, Johnsson et al. [43] used an AI agent as a part of a double-diamond innovation exercise with multiple analysis teams. Louie et al. [57] used an earlier form of generative AI to create a collection of musical passages following user requirements, and to refine selected one or sets of those passages. Bouschery et al. [9], Muller et al. [64], and Shaer et al. [87] examined brainstorming-related methods in which an AI agent augmented the conversational work of groups of humans, generating ideas (exploration), curating ideas, selecting "best" or "final" ideas, and in some cases collaboratively refining ideas (optimization).

However, like other new sociotechnical capabilities, generative AI may be used both more and less effectively to achieve human goals. In contrast with the successes in the preceding paragraph, Wadinambiarachchi et al. [96] call attention to possibilities that design fixation may occur when using AI tools (see also [15]). By contrast, Bordas et al. [6] emphasized the importance of human-AI dialogic processes to reduce the likelihood of design fixation. Lee et al. [48] show that the relationship of AI to one or more humans depends on multiple factors, and that the relationship itself can be tuned through those factors (see also [37]). These dialogic paradigms emphasize fine-grained conversational augmentation of humans through the patterns of their use of generative AI applications. In this paper, we apply augmentation concepts through analysis of human-AI usage patterns of exploration and optimization.

2.2 Generative Variability and Exploration-Optimization Tradeoffs in the HCI/HCAI literature

The preceding section described how the creation of diverse innovation concepts [43], musical passages [57], and brainstorming ideas [9, 64, 87] depended on the generative variability [101] of the AI application. In different ways, each project used generative algorithms to produce not one outcome, but a collection of varied outcomes. Further actions were needed to work with this variability as a resource for innovation, music, or idea-creation and idea-refinement. In other words, researchers used generative variability to create and populate a design space of possibilities. In this paper, we consider how people have used these design spaces

of varied instances and artifacts through different processes of Exploration and Optimization.

In an influential paper originally published in 1991, MacLean et al. [59] proposed that the product of design is a single artifact, but that there is also a discussion-like process that leads to that artifact. They proposed to formalize that discussion in what they termed *Design Space Analysis*, a structured form of argumentation involving Questions, Options, and Criteria (QOC). Their design space analysis was intended to provide a well-understood multi-dimensional space of *possible* designs. We consider their *design space* as a representation of an exploratory process, and their *single artifact* as an optimized outcome. Correspondingly, their proposed methods for optimizing the artifact were quite different from their methods for exploring the design space. Using a Research through Design framework, Giaccardi [29, p.144] offers a related analysis of the exploration of a set of generated of physical artfacts to "open up an imaginative space."

Vo et al. [95, p.2] updated the MacLean conceptual analysis, moving towards more powerful explorations of design spaces: "After generating initial objects, users can further refine them using convenient import and edit prompt features" - i.e., Exploration followed by Optimization. Zhou and Lee [107] proposed a *sequential* transition from Exploration ("generating new ideas") to subsequent Optimization ("selecting the most promising ones"). Hwang [39] made similar distinctions regarding early-stage Exploration ("Hands-on Stage") and late-stage Optimization ("Camera-Ready Stage") in human creative work with AI (see also [8, 64]). In the HCI literature, Optimization is also sometimes referred to as Exploitation (e.g., [14, 44]), or occasionally as Evaluation [32].

Exploration and Optimization/Exploitation/Evaluation have been discussed as being mutually exclusive, strictly sequential, or as trade-off concepts in these literatures. Zhong et al. [106] explicitly designed for Exploration and Exploitation as opposing principles, proposing a *TranSlider* as a human-controlled specification along that trade-off dimension; however, their experiments were limited to simulations with existing datasets. Zhou et al. [109] similarly provided an interactive method for users to control their mix of Exploration and Exploitation content in a music composition task. By contrast, Medlar et al. [60] and Athukorala et al. [2] proposed to model and adapt the UI to what they believed would be the appropriate point along an Exploration/Exploitation trade-off dimension.

Lee et al. [49] reported a similar contrast between Exploration and Optimization; however, in their case, Exploration and Optimization were described as distinct user-selected strategies. Urban Davis et al. [94] similarly provided for flexible user strategies in their *Calliope* system, in which the user could adopt an optimized view of a single potential outcome, or a more globalized view of multiple potential outcomes. Song et al. [89] examined the impact of AI assistance on the design of unmanned aerial vehicles and found that when people worked with AI assistance, their design strategies focused more on Exploration. By contrast, when people worked without AI assistance, their strategies focused more on Optimization. Finally, we note that Zhang et al. [105] speculated that users might want to adjust a "slider" along a dimension that ran between Exploration at one extreme, and Optimization at the opposite extreme.

One familiar example of choosing an Exploration-Optimization strategy occurs in creative ideation and human-AI collaborative brainstorming and brainwriting [7, 33, 64, 87]. In Osborn [73]'s analysis, the core of brainstorming begins with a period of Exploration-like *divergent thinking*, in which many ideas are suggested. This activity is followed by a period of Optimization-like *convergent thinking*, in which participants develop consensus on which idea(s) are best. Participants may also modify the selected idea or ideas [31, 63, 64, 91, 105, 107]. For example, Suh et al. [91, p. 2] note that "Divergent (lateral) thinking is the process of generating multiple, unique ideas or solutions to a problem or task, whereas convergent (vertical) thinking is the process of integrating, synthesizing, and evaluating ideas and possibilities to arrive at (dive into) a single (creative) outcome..." While these processes may be used in cycles, Suh et al. assume that divergence/Exploration is a necessary antecedent to convergence/Optimization. As noted above, Bouschery et al. [8], Muller et al. [64], and Shaer et al. [87] applied this Exploration-Optimization sequence of activities for groups of humans engaging with an AI agent while brainstorming [8, 64] or brainwriting [87] activities.

At a more "meta-" level, Liu et al. [54] analyzed AI-assisted generation of research questions. In their CoQuest project, they drew an analogy between breadth-first search as Exploration and depth-first search as Optimization. Similarly, Liu et al. [54] proposed that users should be able to choose their strategy (depth vs. breadth) according to their personal preferences and the stage of their work in generating research questions.

Generative AI provides new capabilities and creates new challenges [28, 65, 101]. In this paper, we focus on generative variability as an enabler of new approaches to working with design and with outcomes. These new capabilities are already affecting how people work with design and with desiged artifacts. Exploration and Optimization will be important aspects of how we manage these future work opportunities.

2.3 Exploration-Optimization Tradeoffs in the Algorithmic Literature

Exploration and Optimization have also been discussed in the CS and AI literatures. Bettinger et al. [5] describe Exploration and Exploitation as two independent processes in an abstract account of metaheuristic-based recommendations (where "exploitation" and "optimization" have sometimes been used interchangeably):

> "Exploring means making sure that many regions of the search space are examined... so as not to get trapped in a local optimum. Exploiting means that if a promising region of the search space is found, it is examined more carefully... to see if it contains good local optima." [5, p. 475]

Similarly, Cuevas et al. [19, p. 1] write of the "balance between exploration and exploitation" in the evolutionary search of spaces of physical attributes (see also [25]). In a brief formal analysis of evolutionary algorithms, Sun et al. [92] also called for balancing Exploration and Exploitation. Ponzio et al. [79] advocated for an even more structured process for their AI-assisted architectural Design Intelligence process, in which an exploration-like process necessarily preceded an optimization-like process. Similarly, Patel

et al. [75] specified an initial step of exploration of multiple design possibilities, and a subsequent step of focused Optimization, in a potentially iterative process.

Long et al. [55] argued that Exploration and Optimization are *contradictory*:

> "It is well-known that both exploration and exploitation are necessary for population-based algorithms. In practice, the two aspects contradict each other." [55, unpaged]

Similarly, Črepinšek et al. [18, p. 1] referred to Exploration and Exploitation as "antagonistic cornerstones" in evolutionary algorithms.

2.4 Exploration and Optimization within Human-AI Design Guidelines

Considerations for the design of AI-infused user experiences have been discussed in five different sets of design guidelines, but the inclusion of topics of Exploration and Optimization has proven to be challenging. Amershi et al. [1] wrote design guidelines for a broad set of AI applications. They stated that they considered a guideline of "explore vs. exploit in moderation," but omitted it because they found few AI-infused UIs where it was relevant. Four years later, Chung and Adar [17] noted that most text-to-image appplications were "lacking support for iterative shaping of the image," and they began to remedy this shortcoming in their PromptPaint prototype. In a set of design guidelines for generative AI applications, Weisz et al. [101] described Exploration and Optimization as two distinct and important *user goals*. While they proposed six design principles and 24 design strategies for generative AI applications, they observed that

> "[W]e could not... clearly delineate different generative AI applications as supporting exploratory versus optimization usage, because many applications supported both, and users might even alternate between the two kinds of usage when using the application." [101, p. 13]

Muthusamy et al. [68] proposed a set of guidelines for the design of "digital workers," emphasizing ethical design concerns, but topics of Exploration and Optimization were not relevant to those concerns. Xu and Gao [103], Xu et al. [104] offered a set of 28 design guidelines. Their guidelines were relatively high-level, without details on how to achieve them. To the extent that they mentioned Exploration and Optimization, they appear to assume a strictly linear sequence of Exploration followed by Optimization. Herrmann [34] used socio-technical principles to derive an ethics-oriented set of design guidelines for AI applications, but did not discuss the relationship of Exploration and Optimization. Finally, Liu and Chilton [52] proposed detailed guidelines regarding prompts for text-to-image projects. Five of their guidelines focused primarily on rapid iterations to populate a space of design candidates (i.e., Exploration), with limited emphasis (one separate guideline) on detecting and correcting problems that may occur with image content (a medium-specific form of Optimization).

We seek to better understand the contrasts and overlaps identified in the preceding sets of guidelines by analyzing three representative use cases involving Exploration and Optimization. These use cases show Exploration and Optimization as connected and multi-directional, rather than opposing or unidirectional, processes.

2.5 Mixed Initiative Interfaces

Name	Description
Ideate	Create or propose an instance
Constrain	Describe or apply limitations to further generative activities
Produce	Create multiple instances, generally based on a prior exemplar or description; populate a space of instances
Suggest	Propose or recommend an instance
Select	Choose one of the instances for further attention or modification
Curate*	Collect, organize, and/or categorize multiple instances
Assess	Evaluate one or more instances
Adapt	Change, tailor, or customize one or more instances

Table 1: Mixed Initiative Action Vocabulary from the *Library of Mixed Initiative Creative Interfaces* [90]. *The *Curate* action was added in this paper (see Section 3.2).

For a more formal, analytic examination of relationships between Exploration and Optimization, we adopted an analytic vocabulary from the field of human-AI co-creation, and specifically, mixed initiative interfaces (e.g., [36]). Spoto and Oleynik [90] built on the earlier mixed-initiative work [27, 74], to propose a framework for analyzing Mixed Initiative Creative Interfaces (MICIs) using seven core actions derived from a workshop on creative interfaces [26]: Ideate, Constrain, Produce, Suggest, Select, Assess, and Adapt. We highlight several important aspects of their analytic framework:

- Each action might be taken by either a human or an AI ("computer," in their formulation). Table 1 summarizes their seven actions, and adds one more action – Curate – that we will justify in Section 3.2 of this paper.[2]
- Analyzing across the 74 creative interfaces, Spoto and Oleynik found that certain actions were used in only a subset of applications. Thereby, their action vocabulary provides a way to analyze applications in terms of the inclusion or exclusion of certain actions or groups of actions.
- In their analysis of 74 mixed-initiative creative applications from the Deterding et al. [26] workshop, Spoto and Oleynik found very different sequential patterns of utilization of the actions in their vocabulary (i.e., in Table 1) across those applications. Thereby, their action vocabulary also provides

[2]Datta and Woodbury [21] developed a more abstract, symbolic notation that was oriented toward depth-ful analyses of the content of specialized mixed initiative interfaces. Their work did not clarify issues of Exploration and Optimization. While solving a different problem, Liu and Chilton [52] discussed Exploration in ways similar to this paper. However, their approach to Optimization was as a formal algorithmic parameter, rather than as a design activity.

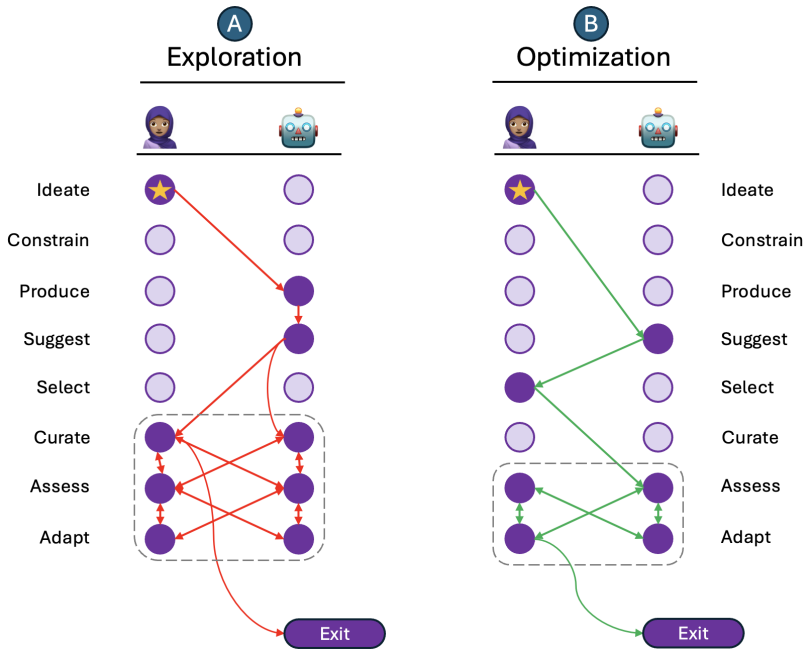


Figure 1: Exploration and Optimization as two separate processes. The two processes are potentially composed from the same action-primitives, which can be performed by either human or AI. A. Exploration. B. Optimization. In practice, the actions inside the dash-line rounded rectangles are often ad hoc and variable from use case to use case, and from instance to instance.

temporal or sequential differentiations among different applications.

We will use this action vocabulary as a rigorous method for understanding the relationships of Exploration and Optimization across diverse AI applications.

3 Preliminary: Exploration and Optimization as Separate Processes

3.1 Background

We now explore the relationships of Exploration and Optimization by applying the Mixed Initiative Action Vocabulary (Table 1). Similarly to the 74 cases analyzed by Spoto and Oleynik [90], we found examples from the available literature [18, 39, 55, 95, 106, 107], and we then iterated on our analysis until we were in full agreement about which actions were involved, and in what sequence, resulting in Figures 1–4. To establish a baseline of how the traditionally opposing relationship between Exploration and Optimization can be represented through the Mixed Initiative Action Vocabulary, we first provide Figure 1 to explicate Exploration and Optimization as separate processes. We illustrate examples of actions taken for these processes, drawing across HCI literature in which the two were separate.

Among the examples of Exploration, Muller et al. [63] described cases of generative AI support for analogy-based design, resulting in a series of ordered lists of ideas. Skorupski et al. [88] developed a more sophisticated use case, with a hierarchical list of higher- and lower- level categories. It is also possible that the instances could be displayed in a multidimensional graphical space [53]. Among the examples of Optimization, the Image-to-Image project [41] and the Drawing Apprentice project [22] provide examples human-AI activities leading to a single preferred or recommended outcome [90].

3.2 Exemplary Workflows

This paper focuses on diverse ways in which human(s) and AI(s) may interact using strategies and activities of Exploration and Optimization. To clarify these core activities, we use an action vocabulary from mixed initiative creative interfaces as our analytic framework. The analysis of Deterding et al. [26] and of Spoto and Oleynik [90] tended to focus on how one human and one AI could co-produce a creative outcome. While reflecting state-of-the-art projects in 2017, this work was done before recent developments in generative AI. We make one addition to their pathbreaking work, adding the Curate action in Table 1 based on [30, 67], to include generative variability phenomena in which people and AI collaborate on designs, products, or other outcomes that produce multiple outcomes (see also Sarkar [85]'s description of curation of AI outcomes as a creative process).

Figure 1 shows two discrete, unrelated processes - one for Exploaration and one for Optimiztion. In Exploration (Figure 1A), the user creates a description of their desired outcome (Ideate), and then asks the AI to generate variations on that description (Produce). The AI suggests one or more examples from among the produced variations (Suggest). The human and AI then sort and categorize those examples (Curate), evaluating the most promising ones (Assess). Either human or AI may change the design parameters of one

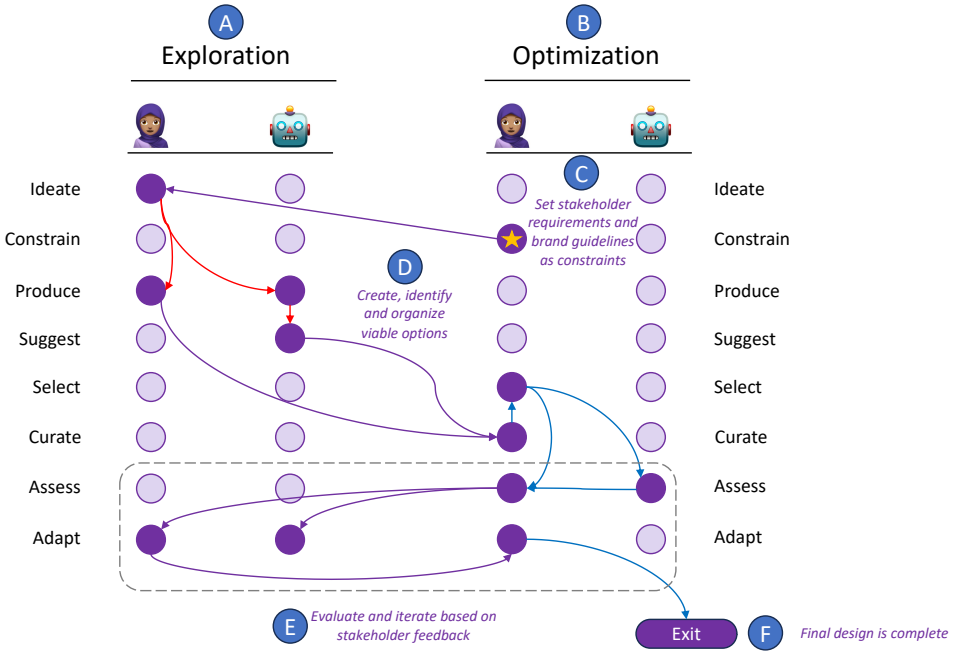


Figure 2: Use Case 1: Explorating and Optimizing graphic design artifacts with AI. A. Exploration actions. B. Optimization actions. C. Transition from Optimization to Exploration following an initial Constrain step. D. Transition from Exploration to Optimization. E. Transitions between Optimization and Exploration during the feedback and iteration stage. F. The workflow ends when the design satisfies all criteria.

or more of the examples (Adapt). When the human is satisfied that the curated collection of examples is sufficiently complete, then the work is over (Exit). For example, Kreminski et al. [46] curated a design space with multiple outcomes, and Liu and Chilton [52] proposed design guidelines for organizing a collection of generated images.

In Optimization (Figure 1B), the user inputs an example or initial instance as a starting point for human and AI to improve the design (Ideate). The AI suggests one or more changes (Suggest). The human then chooses a subset of the recommended changes (Select). The AI evaluates the changes (Assess), and the human and AI adapt the design as needed (Adapt) before the human accepts the design as ready for implementation (Exit). For example, Demirel et al. [24] used generative methods and simulations to optimize the aerodynamic design of automobiles.

These descriptions correspond to the traditional view of Exploration and Optimization as separate [5, 19], perhaps contradictory [18, 55] processes. The next three sections show that this simple dichotomy does not necessarily hold in practice.

4 Use Case 1: Graphic Design

4.1 Background

Schenk [86] identified three phases of a typical graphic design workflow: preparation, the main creative phase, and production. Activities in the creative and production phases may include ideating, seeking inspiration [47], synthesizing ideas, producing and refining visual assets, collecting and responding to feedback, and assessing adherance to business and visual criteria [99].

In recent years, researchers have proposed AI support for the Exploration and Optimization portions of this workflow. Graphic designers sometimes explore a broad range of ideas in search of innovative designs. In response to this need, Lopes et al. [56] created *EvoDesigner*, a tool powered by an earlier form of generative variablity - a genetic algorithm - that automatically evolves design elements and assesses the dissimilarity of resulting artifacts. In their CoPaint project, Jiang et al. [42, p.229] described the use of two related images to generate a "series of high-quality colored images of consistent color" and "a coherent sketch painting dataset." Going a step further, S. Wang et al. [100] used two *unrelated* images to generate multiple combinations of two different image concepts (see also D. Wang et al. [98] for a combination of image category and style). Thus, these applications used generative variability to create multiple instances (Exploration). G. Wang et al. [99] discussed the importance of selecting design elements that meet client requirements. They proposed a deep element selection network to support designers in identifying elements that fulfill visual aesthetics and consistency criteria. Other automated tools have been proposed

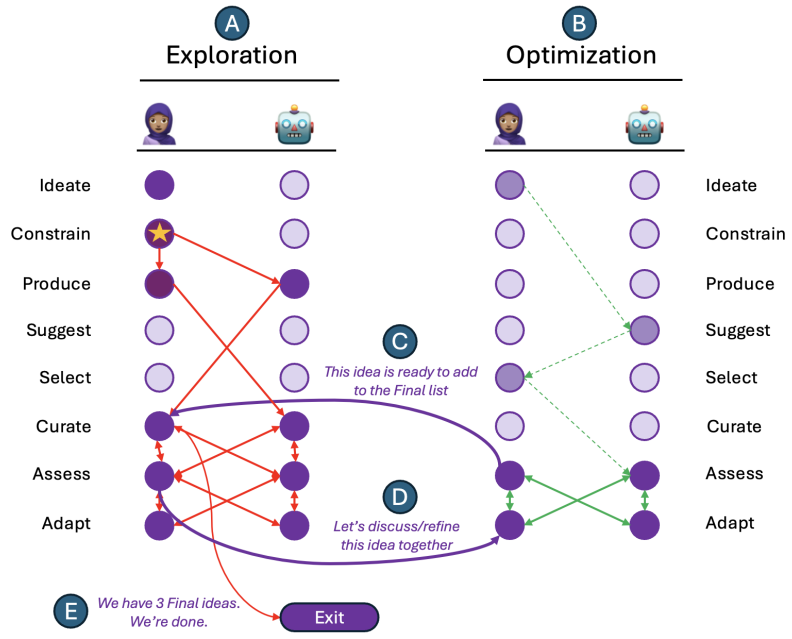


Figure 3: Use Case 2: Exploration and Optimization as iterations in group brainstorming with an AI. The two processes from Figure 1 are in iterative conversation during group brainstorming - from Exploration (Assess) to Optimization (Adapt), and from Optimization (Assess) to Exploration (Curate). A. Exploration actions. B. Optimization actions. C. Transition from Exploration to Optimization. D. Transition from Optimization to Exploration. E. When criteria are met (i.e., three brainstorming ideas to be recommended as "Final ideas"), the brainstorming task is completed.

for generating visual assets [93], suggesting harmonious fonts [16], and providing constructive feedback on designs [20].

4.2 Exemplary Workflow

Based on common actions identified across this literature [42, 56, 98–100], we synthesize a workflow for AI-supported graphic design in Figure 2. The workflow begins with the designer limiting the space of possibilities with stakeholder requirements (Constrain). The designer and AI can then ideate on possible designs – both can create artifacts, and the AI can make a variety of suggestions by leveraging generative variability (Ideate, Produce, Suggest). Once a range of possibilities has been identified, they move to Optimization to narrow down (Select, Curate). An evaluation and iteration stage follows, which may involve soliciting feedback from stakeholders and AI (Assess), returning to Exploration to respond to feedback (Adapt), and finalizing the design (Exit).

In this workflow, Optimization is primarily human-driven, whereas Exploration is a human-AI collaborative effort. Unlike the simple workflow described in Section 3.2, this use case exemplifies how one might cross back and forth between Exploration and Optimization multiple times, and how practical, business-driven design work may begin with Optimization rather than Exploration. Here, transitions between Exploration and Optimization are driven by stakeholder requirements and feedback – the designer begins by constraining the space to these requirements, then dynamically shifts between processes to identify options and iterate towards artifacts that meet stakeholders' criteria.

5 Use Case 2: Brainstorming

5.1 Background

Brainstorming has been used in human-to-human configurations for hundreds of years [77]. Osborn [73] provided an influential description of the practice, which has since been modified and extended by others (e.g., [102]). Online versions of brainstorming (often called "e-brainstorming", e.g., [7, 50, 61, 76]) have been shown to be an effective way of overcoming some of the problems observed in face-to-face brainstorming among humans.Recently, AI agents have been added to online brainstorming and similar AI-supported human activities in diverse media, either for one human and one AI [63], or for a group of humans with an AI [7, 9, 33, 64, 87].

5.2 Exemplary Workflow

Based on studies of human-AI e-brainstorming [7, 9, 37, 64, 87], Figure 3 describes common iterations of Exploration and Optimization that we identified across workflows described in these studies. The work begins with the instructions to humans and AI (Constrain), which sets limits on the brainstorming topic and on the outcome. Both humans and AI begin to engage in divergent thinking [102], producing many ideas (Produce).

Humans (optionally in collaboration with AI) organize the developing list of ideas (Curate). During ad hoc Curate-Assess-Adapt work with multiple ideas, humans may choose to focus the group's work to improve or refine a single idea (transition from Exploration to Optimization at Figure 3C), in ad hoc actions of Assess and Adapt

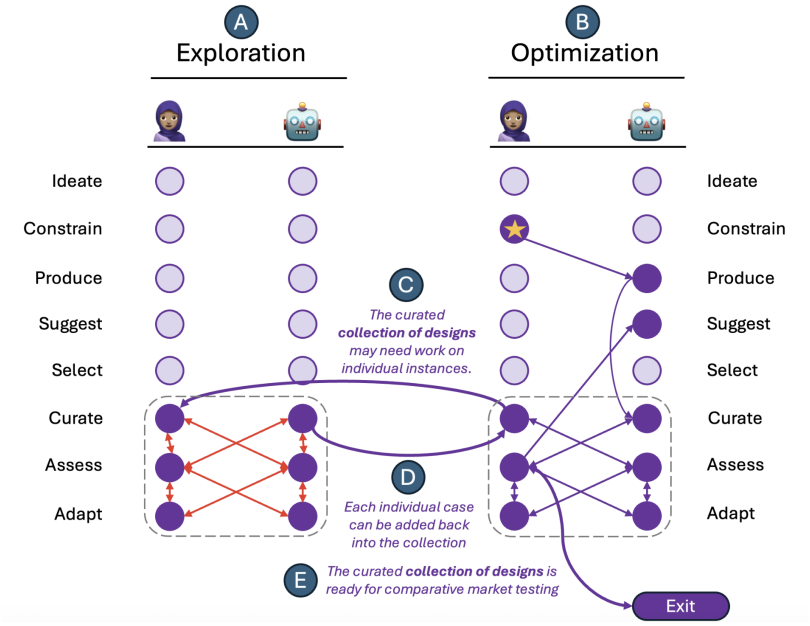


Figure 4: Use Case 3: Exploration and Optimization on a Curated Collection of Instances. In the design of perfume bottles, the focus is on curated collections of bottle designs, rather than individual instances. A. Exploration actions. B. Optimization actions. C. Tranfer of a collection of bottle designs to be first Explored, then Optimized as a curated collection. D. Delivery of the collection of designs for comparative market testing.

by humans and/or AI. When satisfied with the refinements, the group may transition back to Exploration at Figure 3D.

When humans conclude that the group has produced and refined novel ideas, then the work is done, and the session ends (Exit). Hence, uUnlike traditional descriptions of brainstorming that describe a one-way flow from divergent (Exploration) to convergent (Optimization) thinking, studies grounded in real-world brainstorming have shown that *multiple, bi-directional* transitions between Exploration and Optimization can occur within the same workflow.

6 Use Case 3: Package Design

6.1 Background

Quanz et al. [80] investigated co-creative systems for package design. In their case, the problem was to explore diverse candidate designs for perfume bottles, constrained by descriptions of the particular scent of each perfume. Similar projects have been undertaken for the design of aircraft interiors [71], automotive design [11], and image variations [52]. We focus here on the complexity of the Quanz et al. study, because they asked human-AI teams to create a *collection* of candidate bottle designs for each type of perfume.

6.2 Exemplary Workflow

Following Quanz et al. [80] and supplemented by [11, 52, 71], Figure 4 shows an initial Optimization step driven by constraining possible designs to specified scent profiles, followed by Exploration of both individual designs and curated collections of designs, which are then Optimized as curated collections. Contraints are introduced through a description of the attributes of the perfume (Constrain). The AI generates a collection of varied designs (Produce), which the AI and the human evaluate according to multiple evaluative factors (Assess). Some designs may be modified (Adapt), while others may be removed. Eventually a Curated collection of designs is transferred to the Optimization process (Figure 4C).

Unlike the preceding use cases, this use case involves the opposite order of activities, from Optimization to Exploration. The Optimization step operates on the curated *collection* of bottle designs (see also [85]). Thereby, the transfer is from the Curate step in Exploration, to the Curate step in Optimization. Note that, unlike the other use cases, there is no "return" to Exploration. Through additional evaluations and modifications of the collection of designs, the human and the AI Assess that the collection is ready for comparative market testing at Figure 4D (Exit).

7 Discussion

In this section, we show how the use cases can address issues of human augmentation vs. automation (Section 2.1), generative variability (Section 2.2), the relationship of Exploration and Optimization (Sections 2.3 and 2.4), and the descriptive utility of mixed initiative action vocabularies (Section 2.5).

7.1 AI Augmentation of Human Work

In contrast to command-oriented production-like systems (e.g., [70]), we have explored three use cases for iterative generative AI support as an augmentation to humans and their work (for related use cases, see [40, 62, 81, 83].

(1) In the Graphic Design use case (Section 4, Figure 2), the human Constrains the space of AI-produced designs and human-produced designs. The human then Selects instances for optimization, invoking the AI for services such as Assessment and Adaptation.
(2) In the Brainstorming use case (Section 5, Figure 3), the human Constrains the problem space, and the the human and the AI both Produce ideas during divergent thinking. Human and AI collaborate to Curate, Asesss, and Adapt ideas, eventually collecting the best ideas for reporting to a potential client.
(3) In the Package Design use case (Section 6, Figure 4), the human Constrains the design space, and then the AI Produces candidate package designs. The human and the AI collaborate intensely to Curate, Assess, and Adapt those designs, reporting-out their collective Assessment of the best designs.

Consistent with principles of mixed-initiative interfaces [30, 36, 64, 67], we showed that both human actions *and* AI actions were necessary components of each use case. Beyond traditional concepts of humans as reactive elements in an AI-dominated human-in-the-loop paradigm [3, 13, 23, 45], with risks of over-reliance on AI [38], these three use cases highlighted humans as initiators (e.g., Ideate, Constrain), as co-organizers of content (e.g., Select, Curate), and as decision-makers (e.g., Assess, Adapt). Thereby, along with most of the projects in the *Library of Mixed Initiative Creative Interfaces* [26, 90], we propose these use cases as strong arguments for AI augmentation of human work [8, 63, 81, 110], rather than AI automation that displaces humans from their work [69].

7.2 Generative Variability

The three use cases engaged with generative variability [65, 101] as a resource for design. As discussed in Section 2.2, generative variability refers to the use of generative algorithms to produce diverse variants of a design, artifact, or concept [101]. Designers often use multiple sources and configurations of families of design concepts as starting-points or references in their own design activities. Design Space Analysis [59, 95] provided a formal way to consider such variability among design instances. Less formally, designers have used mood boards as sources of information and inspiration for more than a decade (e.g., [58, 72, 82]), leading to recent work using generative AI to populate digital versions of mood boards [12, 78, 97]. In our example use cases, we similarly demonstrated how users can leverage generative variability to expand the design space during Ideate, Produce, and Suggest steps by augmenting their own ideas with possibilities that they may not have identified on their own.

Generative variability is a relatively recent concept, derived from the bahavior of generative algorithms. In the next section, we use concepts of Exploration and Optimization to ground the idea of generative variability in more specific human and AI actions.

7.3 Intertwining of Exploration and Optimization

The three use cases demonstrate how users may make use of generative variability by transitioning between Exploration and Optimization at multiple points in their workflow. During Exploration, users are actively engaged in Curating the variability among instances as a resource for design or for further inquiry. During Optimization, users may identify one or a subset of ideas for focused additional work, and may enter into a more Constrained round of Exploration and further Curation. Another way for users to manage variability, especially undesired variability, is by imposing Constraints as a first step (e.g., in the graphic design and package design workflows). In these scenarios, Optimization becomes the first step, demonstrating how generative variability necessitates a shift away from traditional workflows in which Exploration always comes first. Additionally, since the AI may again provide multiple, varied outputs or feedback that expands the design space in subsequent iterations, users may need to continuously shift between Exploration and Optimization to narrow down the design space. This iterative method of working with AI illustrates another deviation from traditional workflows that depict a single, uni-directional transition from Exploration to Optimization necessitated by generative variability.

These intertwined representations of Exploration and Optimization necessitate a shift away from applications geared towards one process or a uni-directional flow between processes (e.g., from Exploration to Optimization), and instead towards applications that support diverse work-configurations that take advantage of generative variability through selective use of these two processes. Amershi et al. [1] noted the tension between generality and specificity, and propose that designers may need to create additional guidelines for specific domains or contexts. Weisz et al. [101] indicated that designers should identify principles and strategies relevant to their users' needs and apply them accordingly. Building on these ideas, we highlight the importance of providing users with capabilities that can be adapted for either process. For example, providing users with control over the number of outputs would allow them to increase outputs when they want to explore a broader design space (e.g., [11, 51, 59, 95]) and constrain outputs when working on granular refinements. The small number of studies that allowed users to tune or steer the behavior of AI agents [17, 37, 57] may suggest a multi-dimensional approach in which users can indicate which directions or part(s) of the design space they wish the AI (or the human-AI team) to explore further.

Similarly, surfacing a *temperature* control allows users to dynamically control the unexpectedness of outputs as they transition between Exploration and Optimization, in line with Weisz et al. [101]'s strategy to "Provide controls relevant to the use case and technology." Weisz et al. [101] also recommended a strategy to "Visualize the user's journey," a feature that could be useful for documenting choices made during Selection and Curation, and for identifying both unexplored options and a direction to optimize in. Rost and Andreasson [84] and Suh et al. [91] show how these features can be surfaced to users to support structured Exploration.

To help users leverage these adaptible capabilities, it may also be necessary to teach designers how these principles can be applied to both Exploration and Optimization, in line with Amershi et al. [1]'s guideline to "Make clear what the system can do" and Weisz et al. [101]'s proposed strategy to "Teach effective use." For example, an LLM application that supports brainstorming may show users how they can increase the temperature, number of outputs, and other aspects of how the AI presents its contributions to explore

more creative ideas [37], then craft increasingly specific prompts to optimize.

Most of the applications of guidelines [1, 52, 101] assume a default workflow-like process structure. For example, Amershi et al. [1] organize their guidelines into four phases of interaction: initially, during interaction, when working, and over time. By contrast, Lee et al. [49], Liu et al. [54], and Urban Davis et al. [94] emphasized that users might select Exploration and Optimization as strategic options, rather than as a pre-specified sequences of operations. Liu and Chilton [52] speculated that "more controls for steering the generation at intermediate stages could be exposed," similarly to the steering concepts explored by Louie et al. [57] in the exploration-focused control of complex generated-music design spaces. Future work should examine how users should be able to choose among Exploration and Optimization at any time, and whether certain tasks may require more a priori process-structures.

7.4 Mixed Initiative Interfaces

The use cases also help us to rethink the mixed-initiative action vocabularies in relation to Exploration and Optimization. In relation to the individual actions, Exploration and Optimization are higher-level constructs. How do the actions inform our analyses of those higher-level concepts?

We designed an eight-action analytic vocabulary, based on the seven-action framework of Spoto and Oleynik [90]. In response to new generative variability that is already being pursued for future work, we reconsidered the frameworks of Grabe et al. [30] and Muller et al. [67], and we added their Curate concept to the original seven actions in Table 1. Across the use cases in Sections 3 to 6, we found that each type of action was used at least twice, and five actions were used in all four of the cases. Thereby, the actions appear be useful to describe Exploration and Optimization. Across those same sections, we perceived partially different roles for humans vs. AI. Humans Constrained and Selected, but AIs did not. AIs Suggested, but humans did not. Thereby, the actions also may be useful to differentiate the roles of human and AI in co-creative workflows. Nonetheless, there are open questions about future work with mixed initiative analyses. We describe two such questions in the following sections.

7.4.1 Additional Actions? The eight-action framework appeared to be useful to describe the three use cases involving generative variability in this paper. We speculate that other use cases might require additional actions. These new actions might be media-specific, such as simultaneous and inter-responsive actions in co-creating drawings [22, 86] or music [10]. Examples of additional media may be found in the 74 applications analyzed by Spoto and Oleynik [90].

The use cases we examined tended to focus on the creation of a single instance through Optimization. There may also be future projects that co-create *families* of drawings or *suites* of music, in which Exploratory Curation may play a greater role. Through study of additional use cases, we may learn that Curation be accomplished through different strategies, such as unstructured collections, simple lists, ordered lists, dimensionalized categories, or networked relationships. We look forward to analyses of other and future use cases.

7.4.2 Additional Patterns? We organized this paper in terms of two patterns - Exploration and Optimization. By contrast, Grabe et al. [30] proposed five higher-level *interaction patterns* that recur across applications: Exploring, Evolving (i.e., Optimizing), Conditioning, Curating, and Rewriting. We chose to add Curate as one of the action-primitives in this paper. Could we have gained more explanatory power if we had treated Curation in Grabe et al.'s more "modular" way, as suggested by Quanz et al. [80]'s goal to curate entire *collections* of perfume bottles [80] in Section 6? We may learn that the Curate action should be reformulated as a Grabe-like interaction pattern, especially as generative variability elevates its importance, and that different curatorial strategies (mentioned above) may require different Exploration- or Optimization-related steps within their respective curatorial patterns.

8 Conclusion

Generative AI will have profound impact on the future of human work with AI [28, 65, 69, 81, 83]. We have used the concept of generative variability to re-examine the seemingly distinct activities of Exploration and Optimization, as discussed across the HCI, CS, and AI literatures, and as proposed in design guidelines for AI applications [1, 52, 101]. Using analytic frameworks from mixed initiative interfaces [30, 67], we have shown diverse combinations and interweavings of those two concepts. On that basis, we propose novel design and technology implications, and we speculate on future mixed initiative analytic approaches and structured protocols.

References

[1] Saleema Amershi, Dan Weld, Mihaela Vorvoreanu, Adam Fourney, Besmira Nushi, Penny Collisson, Jina Suh, Shamsi Iqbal, Paul N Bennett, Kori Inkpen, et al. 2019. Guidelines for human-AI interaction. In *Proceedings of the 2019 chi conference on human factors in computing systems*. 1–13.

[2] Kumaripaba Athukorala, Alan Medlar, Antti Oulasvirta, Giulio Jacucci, and Dorota Glowacka. 2016. Beyond relevance: Adapting exploration/exploitation in information retrieval. In *Proceedings of the 21st international conference on intelligent user interfaces*. 359–369.

[3] Bárbara C Benato, Cristian Grosu, Alexandre X Falcão, and Alexandru C Telea. 2024. Human-in-the-loop: Using classifier decision boundary maps to improve pseudo labels. *Computers & Graphics* 124 (2024), 104062.

[4] Oded Berger-Tal, Jonathan Nathan, Ehud Meron, and David Saltz. 2014. The exploration-exploitation dilemma: a multidisciplinary framework. *PloS one* 9, 4 (2014), e95693.

[5] Alexandre Bettinger, Armelle Brun, and Anne Boyer. 2022. Independent influence of exploration and exploitation for metaheuristic-based recommendations. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*. 475–478.

[6] Antoine Bordas, Pascal Le Masson, and Benoit Weil. 2024. Switching perspectives on generative artificial intelligence: a design view for humans-generative AI co-creativity. In *R&D Management Conference 2024*.

[7] Sebastian Gregor Bouschery, Vera Blazevic, and Frank T Piller. 2023. AI-Augmented Creativity: Overcoming the Productivity Loss in Brainstorming Groups. In *Academy of Management Proceedings*, Vol. 2023. 11938.

[8] Sebastian G Bouschery, Vera Blazevic, and Frank T Piller. 2023. Augmenting human innovation teams with artificial intelligence: Exploring transformer-based language models. *Journal of Product Innovation Management* 40 (2023), 139–153.

[9] Sebastian Gregor Bouschery, Vera Blazevic, Deborah Roberts, and Frank T Piller. 2022. Artificial intelligence as an actor in innovation teams: An assessment of the GPT-3 language model. In *Academy of Management Proceedings*, Vol. 2022. 11505.

[10] Nick Bryan-Kinns, Berker Banar, Corey Ford, Courtney N Reed, Yixiao Zhang, Simon Colton, and Jack Armitage. 2023. Exploring XAI for the arts: Explaining latent space in generative music. *arXiv preprint arXiv:2308.05496* (2023).

[11] Alexander Burnap, Ye Liu, Yanxin Pan, Honglak Lee, Richard Gonzalez, and Panos Y Papalambros. 2016. Estimating and exploring the product form design space using deep generative models. In *International Design Engineering Technical Conferences*, Vol. 50107.

[12] Alice Cai, Steven R Rick, Jennifer L Heyman, Yanxia Zhang, Alexandre Filipowicz, Matthew Hong, Matt Klenk, and Thomas Malone. 2023. DesignAID: Using generative AI and semantic diversity for design inspiration. In *Proceedings of The ACM Collective Intelligence Conference*. 1–11.

[13] Chengliang Chai and Guoliang Li. 2020. Human-in-the-loop Techniques in Machine Learning. *IEEE Data Eng. Bull.* 43, 3 (2020), 37–52.

[14] Liwei Chan, Yi-Chi Liao, George B Mo, John J Dudley, Chun-Lien Cheng, Per Ola Kristensson, and Antti Oulasvirta. 2022. Investigating positive and negative qualities of human-in-the-loop optimization for designing interaction techniques. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–14.

[15] Liuqing Chen, Yaxuan Song, Chunyuan Zheng, Qianzhi Jing, Preben Hansen, and Lingyun Sun. 2025. Understanding Design Fixation in Generative AI. *arXiv preprint arXiv:2502.05870* (2025).

[16] Saemi Choi, Kiyoharu Aizawa, and Nicu Sebe. 2018. Fontmatcher: font image paring for harmonious digital graphic design. In *23rd International Conference on Intelligent User Interfaces*. 37–41.

[17] John Joon Young Chung and Eytan Adar. 2023. Promptpaint: Steering text-to-image generation through paint medium-like interactions. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. 1–17.

[18] Matej Črepinšek, Shih-Hsi Liu, and Marjan Mernik. 2013. Exploration and exploitation in evolutionary algorithms: A survey. *ACM computing surveys (CSUR)* 45, 3 (2013), 1–33.

[19] Erik Cuevas, Alonso Echavarría, and Marte A Ramírez-Ortegón. 2014. An optimization algorithm inspired by the States of Matter that improves the balance between exploration and exploitation. *Applied intelligence* 40 (2014), 256–272.

[20] Maria Cutumisu, Kristen P Blair, Doris B Chin, and Daniel L Schwartz. 2017. Assessing whether students seek constructive criticism: The design of an automated feedback system for a graphic design task. *International Journal of Artificial Intelligence in Education* 27 (2017), 419–447.

[21] Sambit Datta and Robert F Woodbury. 2002. A graphical notation for mixed-initiative dialogue with generative design systems. In *Artificial Intelligence in Design'02*. Springer, 25–40.

[22] 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 *Proceedings of the 21st International Conference on Intelligent User Interfaces*. 196–207.

[23] Maria De-Arteaga, Riccardo Fogliato, and Alexandra Chouldechova. 2020. A case for humans-in-the-loop: Decisions in the presence of erroneous algorithmic scores. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[24] H Onan Demirel, Molly H Goldstein, Xingang Li, and Zhenghui Sha. 2024. Human-centered generative design framework: An early design framework to support concept creation and evaluation. *International Journal of Human–Computer Interaction* 40, 4 (2024), 933–944.

[25] Thomas Desautels, Andreas Krause, and Joel W Burdick. 2014. Parallelizing exploration-exploitation tradeoffs in gaussian process bandit optimization. *Journal of Machine Learning Research* 15 (2014), 3873–3923.

[26] Sebastian Deterding, Jonathan Hook, Rebecca Fiebrink, Marco Gillies, Jeremy Gow, Memo Akten, Gillian Smith, Antonios Liapis, and Kate Compton. 2017. Mixed-initiative creative interfaces. In *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*. 628–635.

[27] Paul M Fitts. 1951. Human engineering for an effective air-navigation and traffic-control system. (1951).

[28] Werner Geyer, Mary Lou Maher, Justin D Weisz, Daniel Buschek, and Lydia B Chilton. 2024. HAI-GEN 2024: 5th Workshop on Human-AI Co-Creation with Generative Models. In *Companion Proceedings of the 29th International Conference on Intelligent User Interfaces*. 122–124.

[29] Elisa Giaccardi. 2019. Histories and futures of research through design: From prototypes to connected things. *International Journal of Design* 13, 3 (2019), 139–155.

[30] Imke Grabe, Miguel Gonzalez Duque, and Jichen Zhu. 2022. Towards a framework for human-ai interaction patterns in co-creative gan applications. In *Proceeding of the 3rd Workshop on Human-AI Co-Creation with Generative Models (HAI-GEN '22) at ACM IUI Workshops*.

[31] Luca Grilli and Mattia Pedota. 2024. Creativity and artificial intelligence: A multilevel perspective. *Creativity and Innovation Management* 33, 2 (2024), 234–247.

[32] Chao Guo, Yue Lu, Yong Dou, and Fei-Yue Wang. 2023. Can ChatGPT boost artistic creation: The need of imaginative intelligence for parallel art. *IEEE/CAA Journal of Automatica Sinica* 10, 4 (2023), 835–838.

[33] Jessica He, Stephanie Houde, Gabriel E Gonzalez, Darío Andrés Silva Moran, Steven I Ross, Michael Muller, and Justin D Weisz. 2024. AI and the Future of Collaborative Work: Group Ideation with an LLM in a Virtual Canvas. In *Proceedings of the 3rd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. 1–14.

[34] Thomas Herrmann. 2024. Comparing Socio-technical Design Principles with Guidelines for Human-Centered AI. In *International Conference on Human-Computer Interaction*. Springer, 60–74.

[35] Thomas T Hills, Peter M Todd, David Lazer, A David Redish, and Iain D Couzin. 2015. Exploration versus exploitation in space, mind, and society. *Trends in cognitive sciences* 19, 1 (2015), 46–54.

[36] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 159–166.

[37] Stephanie Houde, Kristina Brimijoin, Michael Muller, Steven I Ross, Dario Andres Silva Moran, Gabriel Gonzalez, Siya Kunde, Morgan Foreman, and Justin D Weisz. 2025. Controlling AI Agent Participation in Group Conversations: A Human-Centered Approach. In *Proceedings of IUI 2025*.

[38] Rosco Hunter, Richard Moulange, Jamie Bernardi, and Merlin Stein. 2024. Monitoring Human Dependence On AI Systems With Reliance Drills. *arXiv preprint arXiv:2409.14055* (2024).

[39] Angel Hsing-Chi Hwang. 2022. Too late to be creative? AI-empowered tools in creative processes. In *CHI conference on human factors in computing systems extended abstracts*. 1–9.

[40] IBM. 2023. Best practices for augmenting human intelligence with AI. https://www.ibm.com/think/insights/ai-best-practices.

[41] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.

[42] Shiqi Jiang, Chenhui Li, and Changbo Wang. 2021. CoPaint: Guiding Sketch Painting with Consistent Color and Coherent Generative Adversarial Networks. In *Advances in Computer Graphics: 38th Computer Graphics International Conference, CGI 2021, Virtual Event, September 6–10, 2021, Proceedings 38*. Springer, 229–241.

[43] Mikael Johnsson, Helena Blackbright, Kristina Swenningsson, and Elia Mörling. 2024. Towards Human-machine Innovation Teams. In *Conference: R&D Management Conference 2024. At: Stockholm, Sweden.*

[44] Mary Beth Kery and Brad A Myers. 2017. Exploring exploratory programming. In *2017 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 25–29.

[45] Bongjun Kim and Bryan Pardo. 2018. A human-in-the-loop system for sound event detection and annotation. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 8, 2 (2018), 1–23.

[46] Max Kreminski, Isaac Karth, Michael Mateas, and Noah Wardrip-Fruin. 2022. Evaluating Mixed-Initiative Creative Interfaces via Expressive Range Coverage Analysis.. In *IUI Workshops*. 34–45.

[47] Simon Laing and Masood Masoodian. 2016. A study of the influence of visual imagery on graphic design ideation. *Design Studies* 45 (2016), 187–209.

[48] Soohwan Lee, Seoyeong Hwang, and Kyungho Lee. 2024. Conversational Agents as Catalysts for Critical Thinking: Challenging Design Fixation in Group Design. *arXiv preprint arXiv:2406.11125* (2024).

[49] Seung Won Lee, Tae Hee Jo, Semin Jin, Jiin Choi, Kyungwon Yun, Sergio Bromberg, Seonghoon Ban, and Kyung Hoon Hyun. 2024. The Impact of Sketch-guided vs. Prompt-guided 3D Generative AIs on the Design Exploration Process. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–18.

[50] Lassi A Liikkanen, Kai Kuikkaniemi, Petri Lievonen, and Pauli Ojala. 2011. Next step in electronic brainstorming: collaborative creativity with the web. In *CHI'11 Extended Abstracts on Human Factors in Computing Systems*. 2029–2034.

[51] Richard Lin, Rohit Ramesh, Parth Nitin Pandhare, Kai Jun Tay, Prabal Dutta, Bjoern Hartmann, and Ankur Mehta. 2024. Design Space Exploration for Board-level Circuits: Exploring Alternatives in Component-based Design. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–14.

[52] Vivian Liu and Lydia B Chilton. 2022. Design guidelines for prompt engineering text-to-image generative models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*. 1–23.

[53] Vivian Liu, Han Qiao, and Lydia Chilton. 2022. Opal: Multimodal image generation for news illustration. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–17.

[54] Yiren Liu, Si Chen, Haocong Cheng, Mengxia Yu, Xiao Ran, Andrew Mo, Yiliu Tang, and Yun Huang. 2024. How ai processing delays foster creativity: Exploring research question co-creation with an llm-based agent. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–25.

[55] Wen Long, Jianjun Jiao, Ximing Liang, and Mingzhu Tang. 2018. An exploration-enhanced grey wolf optimizer to solve high-dimensional numerical optimization. *Engineering Applications of Artificial Intelligence* 68 (2018), 63–80.

[56] Daniel Lopes, João Correia, and Penousal Machado. 2023. EvoDesigner: Aiding the Exploration of Innovative Graphic Design Solutions. In *International Conference on Computational Intelligence in Music, Sound, Art and Design (Part of EvoStar)*. Springer, 383–398.

[57] Ryan Louie, Andy Coenen, Cheng Zhi Huang, Michael Terry, and Carrie J Cai. 2020. Novice-AI music co-creation via AI-steering tools for deep generative models. In *Proceedings of the 2020 CHI conference on human factors in computing systems*. 1–13.

[58] Andrés Lucero and Kirsikka Vaajakallio. 2008. Co-designing mood boards: creating dialogue with people. *Proc. of IASTED HCI* (2008), 254–260.

[59] Allan MacLean, Richard M Young, Victoria ME Bellotti, and Thomas P Moran. 2020. Questions, options, and criteria: Elements of design space analysis. In *Design rationale*. CRC Press, 53–105.

[60] Alan Medlar, Joel Pyykkö, and Dorota Glowacka. 2017. Towards fine-grained adaptation of exploration/exploitation in information retrieval. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*. 623–627.

[61] Paul A Mongeau and Mary Claire Morr. 1999. Reconsidering brainstorming. *Group Facilitation* 1 (1999), 14.

[62] Michael Muller. 2023. Exploring Human-AI Co-Creativity under Human Control: Framing, Reframing, Brainstorming, and Future Challenges (workshop keynote). https://sites.google.com/view/ethaics-2023.

[63] Michael Muller, Heloisa Candello, and Justin Weisz. 2023. Interactional Co-Creativity of Human and AI in Analogy-Based Design. In *International Conference on Computational Creativity*.

[64] Michael Muller, Stephanie Houde, Gabriel Gonzalez, Kristina Brimijoin, Steven I Ross, Dario Andres Silva Moran, and Justin D Weisz. 2024. Group Brainstorming with an AI Agent: Creating and Selecting Ideas. In *International Conference on Computational Creativity*.

[65] Michael Muller, Anna Kantosalo, Mary Lou Maher, Charles Patrick Martin, and Greg Walsh. 2024. GenAICHI 2024: Generative AI and HCI at CHI 2024. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–7.

[66] Michael Muller and Justin Weisz. 2023. Analogies-based design using a generative AI application: A play in three acts. https://www.researchgate.net/publication/372060579_Analogies-based_Design_using_a_Generative_AI_Application_A_Play_in_Three_Acts.

[67] Michael Muller, Justin D Weisz, and Werner Geyer. 2020. Mixed initiative generative AI interfaces: An analytic framework for generative AI applications. In *Workshop Proceedings The Future of Co-Creative Systems at International Conference on Computational Creativity (ICCC 2020)*.

[68] Vinod Muthusamy, Merve Unuvar, Hagen Völzer, and Justin D Weisz. 2020. Do's and Don'ts for Human and Digital Worker Integration. *arXiv preprint arXiv:2010.07738* (2020).

[69] National Academies of Sciences, Engineering, and Medicine. 2024. *Artificial Intelligence and the Future of Work*. National Academies Press.

[70] Jakob Nielsen. 2023. AI: First New UI Paradigm in 60 Years. *Nielsen Norman Group* 18, 06 (2023), 2023.

[71] Wedyn Noronha, Vedant Joshi, and S Jeyanthi. 2021. Design of 3D printed aircraft seat structure using latticing in combination with topology optimization and generative design. In *IOP Conference Series: Materials Science and Engineering*, Vol. 1128. IOP Publishing, 012008.

[72] Trond Are Øritsland and Jacob Buur. 2003. Interaction styles: An aesthetic sense of direction in interface design. *International Journal of Human-Computer Interaction* 15, 1 (2003), 67–85.

[73] Alex F Osborn. 1953. *Applied imagination*. Scribners.

[74] Raja Parasuraman, Thomas B Sheridan, and Christopher D Wickens. 2000. A model for types and levels of human interaction with automation. *IEEE Transactions on systems, man, and cybernetics* 30, 3 (2000), 286–297.

[75] Kaushikkumar Patel, Sandeep Kumar, Madhavi Najana, and Anandaganesh Balakrishnan. 2024. AI-powered creativity and data-driven design. *International Journal of Artificial Intelligence Research and Development* 2, 2 (2024), 20–39.

[76] Paul B Paulus. 2015. Electronic brainstorming research and its implications for e-planning. *International Journal of E-Planning Research (IJEPR)* 4, 1 (2015), 42–53.

[77] Paul B Paulus and Vincent R Brown. 2003. Enhancing ideational creativity in groups. *Group creativity: Innovation through collaboration* (2003), 110–136.

[78] 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*. 804–818.

[79] Angelica Paiva Ponzio, Sara Regiane Chornobai, Cristian Vinicus Machado Fagundes, Ricardo Cesar Rodrigues, and Gustavo Cunha Hafez Jose. 2023. Exploring creative AI thinking in the design process: The design intelligence strategy. In *Proc. CAADRIA 2023*. 49–58.

[80] Brian Quanz, Wei Sun, Ajay Deshpande, Dhruv Shah, and Jae-eun Park. 2020. Machine learning based co-creative design framework. *arXiv preprint arXiv:2001.08791* (2020).

[81] Sebastian Raisch and Sebastian Krakowski. 2021. Artificial intelligence and management: The automation–augmentation paradox. *Academy of management review* 46, 1 (2021), 192–210.

[82] Vincent Rieuf, Carole Bouchard, Vincent Meyrueis, and Jean-François Omhover. 2017. Emotional activity in early immersive design: Sketches and moodboards in virtual reality. *Design Studies* 48 (2017), 43–75.

[83] L Rodel, E Hertog, C Unruh, and D Susskind. 2024. AI and the Future of Work Workshop Report. *Institute for Ethics in AI, University of Oxford TABLE OF CONTENTS* (2024), 4.

[84] Mattias Rost and Sebastian Andreasson. 2023. Stable Walk: An interactive environment for exploring Stable Diffusion outputs. In *HAI-GEN 2023: 4th Workshop on Human-AI Co-Creation (Joint Proceedings of the ACM IUI Workshops 2023). Sydney, Australia*. 89–97.

[85] Advait Sarkar. 2023. Exploring perspectives on the impact of artificial intelligence on the creativity of knowledge work: Beyond mechanised plagiarism and stochastic parrots. In *Proceedings of the 2nd Annual Meeting of the Symposium on Human-Computer Interaction for Work*. 1–17.

[86] Pamela Schenk. 1991. The role of drawing in the graphic design process. *Design Studies* 12, 3 (1991), 168–181.

[87] Orit Shaer, Angelora Cooper, Osnat Mokryn, Andrew L Kun, and Hagit Ben Shoshan. 2024. AI-Augmented Brainwriting: Investigating the use of LLMs in group ideation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–17.

[88] James Skorupski, Lakshmi Jayapalan, Sheena Marquez, and Michael Mateas. 2007. Wide ruled: A friendly interface to author-goal based story generation. In *Proceedings of 4th ICVS Conference*. Springer, 26–37.

[89] Binyang Song, Nicolás F Soria Zurita, Hannah Nolte, Harshika Singh, Jonathan Cagan, and Christopher McComb. 2022. When faced with increasing complexity: the effectiveness of artificial intelligence assistance for drone design. *Journal of Mechanical Design* 144, 2 (2022), 021701.

[90] Angie Spoto and Natalia Oleynik. 2017. Library of Mixed-Initiative Creative Interfaces. http://mici.codingconduct.cc/.

[91] 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*. 1–26.

[92] Jianyong Sun, Hu Zhang, Qingfu Zhang, and Huanhuan Chen. 2018. Balancing exploration and exploitation in multiobjective evolutionary optimization. In *Proceedings of the genetic and evolutionary computation conference companion*. 199–200.

[93] Steven Taylor, Andrew Thomas, Kenneth Moore, Joshua Jackson, and Kevin White. [n. d.]. Automating Graphic Design: Exploring Logotype Synthesis with Generative Models. https://www.researchgate.net/profile/Andrew-Thomas-81/publication/376582242_Automating_Graphic_Design_Exploring_Logotype_Synthesis_with_Generative_Models/links/657dff778e2401526ddc330b/Automating-Graphic-Design-Exploring-Logotype-Synthesis-with-Generative-Models.pdf.

[94] Josh Urban Davis, Fraser Anderson, Merten Stroetzel, Tovi Grossman, and George Fitzmaurice. 2021. Designing co-creative ai for virtual environments. In *Proceedings of the 13th Conference on Creativity and Cognition*. 1–11.

[95] Dinh-Khoi Vo, Duy-Nam Ly, Khanh-Duy Le, Tam V Nguyen, Minh-Triet Tran, and Trung-Nghia Le. 2024. iCONTRA: Toward Thematic Collection Design Via Interactive Concept Transfer. In *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–8.

[96] 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 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.

[97] 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*. 136–146.

[98] Da Wang, Jiaqi Li, Zhen Ge, and Ji Han. 2021. A computational approach to generate design with specific style. *Proceedings of the Design Society* 1 (2021), 21–30.

[99] Guolong Wang, Zheng Qin, Junchi Yan, and Liu Jiang. 2020. Learning to select elements for graphic design. In *Proceedings of the 2020 International Conference on Multimedia Retrieval*. 91–99.

[100] Sitong Wang, Savvas Petridis, Taeahn Kwon, Xiaojuan Ma, and Lydia B Chilton. 2023. PopBlends: Strategies for conceptual blending with large language models. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.

[101] Justin D Weisz, Jessica He, Michael Muller, Gabriela Hoefer, Rachel Miles, and Werner Geyer. 2024. Design Principles for Generative AI Applications. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–22.

[102] Chauncey Wilson. 2013. *Brainstorming and beyond: a user-centered design method*. Newnes.

[103] Wei Xu and Zaifeng Gao. 2024. Enabling Human-Centered AI: A Methodological Perspective. In *2024 IEEE 4th International Conference on Human-Machine Systems (ICHMS)*. IEEE, 1–6.

[104] Wei Xu, Zaifeng Gao, and Marvin Dainoff. 2023. An HCAI methodological framework: Putting it into action to enable human-centered AI. *arXiv preprint arXiv:2311.16027* (2023).

[105] Kejun Zhang, Shixuan Cai, Wenjing Yang, Wenqi Wu, and Hanshu Shen. 2024. Exploring Optimal Combinations: The Impact of Sequential Multimodal Inspirational Stimuli in Design Concepts on Creativity. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 2788–2801.

[106] Kuo Zhong, Ying Wei, Chun Yuan, Haoli Bai, and Junzhou Huang. 2020. Translider: Transfer ensemble learning from exploitation to exploration. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 368–378.

[107] Eric Zhou and Dokyun Lee. 2023. Generative ai, human creativity, and art. *Available at SSRN* (2023).

[108] Lina Zhou, Souren Paul, Haluk Demirkan, Lingyao Yuan, Jim Spohrer, Michelle Zhou, and Julie Basu. 2021. Intelligence augmentation: Towards building human-machine symbiotic relationship. *AIS Transactions on Human-Computer Interaction* 13, 2 (2021), 243–264.

[109] Yijun Zhou, Yuki Koyama, Masataka Goto, and Takeo Igarashi. 2021. Interactive exploration-exploitation balancing for generative melody composition. In *26th International Conference on Intelligent User Interfaces*. 43–47.

[110] John Zysman and Mark Nitzberg. 2024. Generative AI and the Future of Work: Augmentation or Automation? *Available at SSRN 4811728* (2024).