

Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation

Sangho Suh*
University of California San Diego
La Jolla, CA, USA
sanghosuh@ucsd.edu

Meng Chen*
University of Notre Dame
Notre Dame, IN, USA
mchen24@nd.edu

Bryan Min
University of California San Diego
La Jolla, CA, USA
bdmin@ucsd.edu

Toby Jia-Jun Li
University of Notre Dame
Notre Dame, IN, USA
toby.j.li@nd.edu

Haijun Xia
University of California San Diego
La Jolla, CA, USA
haijunxia@ucsd.edu




Figure 1: Our approach, structured multi-output (C), is shown with two current interaction paradigms (A & B). We use *structured* to denote the presence of dimensions relevant to the task / domain in guiding the response generation and *unstructured* to denote their absence. Specifically, in our approach, users' prompt triggers (c1) generation of dimensions and subsequently the (c2) generation of responses using the dimensions from the previous step. Users can (c3) select dimension(s) to organize the responses in one- or two-dimensional space for exploration. The implementation and interaction details for each step are presented in the remainder of this paper.

ABSTRACT

Thanks to their generative capabilities, large language models (LLMs) have become an invaluable tool for creative processes. These models have the capacity to produce hundreds and thousands of

visual and textual outputs, offering abundant inspiration for creative endeavors. But are we harnessing their full potential? We argue that current interaction paradigms fall short, guiding users towards rapid convergence on a limited set of ideas, rather than empowering them to explore the vast latent design space in generative models. To address this limitation, we propose a framework that facilitates the structured generation of design space in which users can seamlessly explore, evaluate, and synthesize a multitude of responses. We demonstrate the feasibility and usefulness of this framework through the design and development of an interactive system, Luminate, and a user study with 14 professional writers. Our work advances how we interact with LLMs for creative tasks, introducing a way to harness the creative potential of LLMs.

* Contributed equally



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

CHI '24, May 11–16, 2024, Honolulu, HI, USA
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0330-0/24/05
<https://doi.org/10.1145/3613904.3642400>

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**; **Information visualization**.

KEYWORDS

Large language models; human-AI interaction; human-AI co-creation; creativity support; dimensional exploration; design space

ACM Reference Format:

Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminate: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems (CHI '24)*, May 11–16, 2024, Honolulu, HI, USA. ACM, New York, NY, USA, 26 pages. <https://doi.org/10.1145/3613904.3642400>

1 INTRODUCTION

A dictum from Linus Pauling, a two-time Nobel prize winner, states: “if you want to have good ideas, you must have lots of ideas and learn to throw away the bad ones” [41]. Indeed, creative processes typically begin with generating and exploring multiple ideas rather than refining a single one. Without first exploring diverse ideas, one can hastily converge to a *sub-optimal* idea and spend time iterating on it without considering other options — a phenomenon called fixation that hinders creative processes [19, 24, 37]. Besides helping people avoid fixation, helping them understand the *design space* — the space of possible ideas and solutions to a task and problem — has been found to play an essential role in creative processes [23, 34, 45, 57]. Thus, researchers have voiced the need to support *design space thinking* in creativity support tools, stating various benefits of offering designers the ability to leverage an understanding of the design space in creative processes [7, 10, 23, 32, 55, 57].

On the other hand, large language models (LLMs) today offer an unprecedented opportunity to empower creative processes. In addition to utilizing LLMs’ generative power to automate and accelerate creation tasks [27, 40, 69], the capacity to instantly produce tens to hundreds of outputs is especially conducive to assisting users in developing a comprehensive understanding of the design space [32, 39], a potential that current user interaction paradigms for LLMs (Fig. 1A & B) are not fully harnessing. Instead, the current interaction paradigms help users *converge* to an idea and refine the prompts (prompt engineering) to produce a more polished output from initial idea(s) they may have *prematurely* converged to.

For example, recent text-to-image generation systems that generate multiple images for exploration focus on helping users refine their prompts to find the image that fits their desired quality or characteristics [11, 26]. Specifically, they do not provide a design space for reflection and exploration to facilitate and encourage *design space thinking*. Instead, after users add their prompt (e.g., subject), they are quickly presented with a set of similar results (e.g., images about the subject), confining them within a narrow design space. Although this approach, illustrated in Fig. 1B, allows users to explore alternative ideas, the generation of these outputs is *unstructured* (i.e., the method does not employ a systematic approach, such as structuring the generation of responses using key dimensions and values that make up the design space) and the exploration of

these outputs encourages users to converge rather than diverge [24] — an argument we further elaborate on in Section 3.2.1.

The goal of this work, therefore, is to explore a new interaction paradigm with LLMs that scaffolds and encourages design space thinking in the human-AI co-creation process. To achieve this, we draw insights from prior work in exploratory search [3, 51, 56], creativity support [13, 24, 37], dimensional reasoning [23, 50], and information visualizations [25, 62]. These fields repeatedly demonstrated that explicit enumeration of, visualizing data within, and interaction with the principle dimensions associated with the target domains or tasks can enhance understanding and exploration of the design space. Based on these insights, we propose a human-AI collaboration framework (Fig. 1C). Instead of generating responses directly based on the user’s prompt, we first prompt LLMs to generate key dimensions associated with the topic or task in the original prompt, and then combine the derived dimensions with the original prompt to systematically generate responses to construct the design space. Users can then leverage these derived dimensions to explore the generated design space in a structured, systematic manner.

To demonstrate the feasibility and usefulness of this framework, we instantiated it by developing Luminate, an interactive system that uses LLMs to first generate key dimensions from a user’s prompt and then generates responses using these dimensions for structured exploration of LLM outputs. Our user study with 14 professional writers specializing in creative writing shows that our approach has the potential to help users effectively explore the design space during their creative processes. Our work contributes a step towards improving the interaction and usefulness of LLMs in assisting creative processes. In summary, our contributions include:

- A new interaction framework for human-AI collaboration in creative tasks that build on the premise that LLMs should allow users to explore a space of possible responses, rather than giving a single data point in response to user input.
- Luminate,¹ an interactive system that demonstrates this idea with novel interaction techniques and features for structured generation and exploration of LLM outputs.
- A user study demonstrating that enabling dimensional exploration of LLM output space has the potential to prevent fixation and nurture design space thinking.

2 BACKGROUND

We present prior research and concepts relevant to our work, namely supporting creative process, co-creating with AI, and facilitating visuo-spatial exploration and sensemaking.

2.1 Supporting Creative Process

Divergent-convergent thinking. Creative thinking entails two modes of thought: divergent and convergent thinking [30, 66]. 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 [21]. Because divergent thinking is responsible for the generation of ideas and convergent thinking the evaluation of these ideas, prior research established that the interplay between the

¹<https://luminate-research.github.io/>

two is imperative in creative processes and noted that creativity support tools should support the constant cycle of these two modes in creative processes [12, 31]. Despite the benefits of divergent thinking, such as avoiding fixation, leading to better and more diverse outcomes, and improved self-efficacy [24], studies found that even experts often converge prematurely during creative processes [19, 37], suggesting the human tendency to converge and challenges to diverge.

Design space. With divergent-convergent thinking, design space plays a critical role in creative processes [7, 10, 64]. The term design space is generally used to refer to a conceptual ‘space of possibilities’ — broadly construed as collections of ideas, designs, concepts, or solutions [32, 49]. Design spaces have commonly been represented with Cartesian spaces where possible variations reside within the dimensions of a design, and variations reflect different values on those dimensions [7, 32, 45, 57]. For example, the design space of a narrative might include dimensions such as genre, setting, point of view, tone, and so on. Each narrative in the space would reflect the values in those dimensions (e.g., *fantasy* for genre, *modern times* for setting, etc.) and positioned accordingly in the Cartesian space.

Dimensional Reasoning. Since design space makes dimensions and their values explicit, it enables designers to reflect on the problem space, desired direction and characteristics, and to systematically evaluate different possibilities [7]. Essentially, it enables dimensional reasoning, allowing people to step back from the individual idea, artifact, or solution and view the entire problem space at a conceptual, higher level. This offers an all-encompassing view [34], which is why design space is frequently used in HCI and other design-related disciplines to classify new artifacts, as it makes it easy to compare them with existing artifacts or designs to understand what makes them different or novel (e.g., [65]).

Simultaneously, design space plays a crucial role in generating new designs, ideas, and artifacts [6, 50]. While a design space “constrains design possibilities along some dimensions, it leaves others open for creative exploration” [7]. This can be seen, for example, in research that uses design cards to generate new designs [29, 52, 58, 59]. For instance, Lomas et al. created design space cards for game designs where elements of game design theory (e.g., Action, Theme, Play Styles) were used as dimensions. They found that these design space cards can help a range of audiences — from children and design novices to experts — to effectively create new game designs [45]. Other similar research that created design cards to support design in their respective areas reported the same findings [5, 52, 58, 59].

As such, while thinking with dimensions offers many benefits, prior work found that dimensional reasoning does not happen naturally and that even experts can have difficulties answering the dimensions that characterize the problem space they work in [50]. This suggests that in order to encourage and enable dimensional reasoning, we need to explicitly surface dimensions to users’ attention and support the generation of these dimensions.

Our work builds on these lines of work and extends it in several ways. First, our framework demonstrates how we can interweave the idea of design space (thinking) into the LLM-powered workflow to enhance the usefulness of LLMs in creative processes. Second, we show the use of LLMs to generate dimensions for dimensional reasoning in creative processes, which should bring the benefits of

dimensional reasoning while solving the challenge users have with identifying or recalling key dimensions characterizing the problem space, task, and domain. (We showcase these interactions through our system in Section 4.)

2.2 Co-Creating with AI

Co-creation with AI is rapidly gaining traction as a prevalent approach in creative processes. This collaborative style of work has already found application in numerous domains, such as music composition [46], creative coding [39], writing [27, 40, 68, 69, 71], design [38, 42, 44, 47], and video authoring [63]. Indeed, much work focused on demonstrating how to co-create in these particular domains, as well as understanding the challenges in co-creating with AI [28] and enabling users to more effectively control and steer AI’s output (e.g., prompt engineering [11, 40, 67]).

Amidst the growing interest in this partnership, there has been a long-lasting theme of research to explore how AI can be effectively intertwined with humans so that AI augments — rather than fully automates — our productivity and creativity [4, 35, 43]. Licklider proposed human-AI symbiosis, where the human defines the goals and criteria, and AI prepares the tasks for the human to perform more insightful work [43]. Furthermore, to build a comprehensive perspective on human-AI co-creation, efforts have been made to chart the design space in the context of text generation tasks [16, 22]. Cheng et al. conducted a systematic analysis of the design space, culminating in the creation of a taxonomy featuring five dimensions, including human actions, types of human control, model iterations, workflow initiation, and design of interfaces between human and AI [16]. Building on this taxonomy, Ding and Chan identified a spectrum of human-AI co-creation tasks, each associated with specific patterns of human-AI interaction [22]. Within this spectrum, they placed human-AI co-creation tasks according to the degree of required human intervention, from minimal to substantial. They concluded that complex and interdependent creative tasks, such as paper and fiction writing, thrive in human-AI co-creation due to the iterative and interdependent nature of the interactions.

Our work builds on the extensive exploration of human-AI collaboration research, where they leverage the understanding of the unique strengths and weaknesses of humans and AIs. Our perspective and goal are different from these prior research in that, instead of focusing on leveraging human-AI collaboration to accomplish a creation task, our primary goal is to leverage AI to support humans in developing a more comprehensive understanding of the design space so that the human can more effectively identify the requirements, criteria, and constraints of the creativity task, which was often assumed to be the responsibility and difficulty that the human need to take on their own [43, 50].

2.3 Facilitating Exploration and Sensemaking

While LLMs’ ability to generate a large number of responses presents an opportunity for creativity support, it also poses a challenge in how to visualize the responses without overwhelming users and help them navigate the design space. To address this challenge, we build upon visualization and interaction techniques that have been explored in information visualization and HCI.

Visualization techniques for multivariate data visualization — such as interactive scatter plots, linked axes [18], and small multiples [9, 48] — enable users to inspect the distribution of data points across individual or pairs of dimensions. For example, with linked axes, each dimension of the data set forms an axis, and data points are visualized as links that connect the various axes [18]. Small multiples with scatter plots display a series of scatter plots showing the data distribution across pairs of data dimensions [9]. Such systems represent vast amounts of data in a compact 2-dimensional format and empower users with interactive features such as searching, filtering, and user-driven clustering of data points to enable flexible exploration and facilitate the understanding of the data set.

While representing rich multivariate data points as dots on a scatter plot enables users to easily inspect the distribution of data points, such abstraction also makes it challenging to dive into the details of specific data points. Semantic zooming is a technique that has been used to address this. It suggests visualizing different levels of details at various levels of zooming scales [8]. This technique enables users to flexibly switch between the overview of the space to the detailed view of specific data points, facilitating the understanding of focused information in context [14], and has been used for a variety of sensemaking tasks [33, 53, 60]. Closely related to our context is Sensecape, which leverages semantic zoom to enable users to easily overview large amounts of text responses generated by LLMs, where text responses flexibly transition across three different levels of abstraction — keywords, summary sentences, and full text — at different zooming scales [60].

Our work builds on these well-known visualization and interaction techniques formulated to facilitate the exploration and sensemaking for the large dataset of multivariate responses generated by LLMs. We leverage the interaction visualizations as graphical interfaces to enable users to request more responses with existing dimensions or re-evaluate existing responses with new dimensions to facilitate the iterative divergent and convergent creative processes — novel interaction techniques shown with Luminate (Section 4) that we believe can contribute to advancing Human-AI co-creation.

3 PROMPTING FOR DESIGN SPACE: A FRAMEWORK FOR DESIGN SPACE THINKING IN HUMAN-AI CO-CREATION

Motivated by the benefits of design space thinking in creative processes — (1) avoiding fixation [36, 37] and (2) enabling systematic generation and evaluation of ideas [23, 50], we argue that generative AI models should assist in the generation of the design space rather than individual artifacts at the early stage of the creative process to empower users and harness the creative potential of AI.

3.1 Current Interaction Paradigms (Figs. 1A & B)

The way we interact with LLMs is designed to help users converge rather than diverge. Fig. 1A and Fig. 1B — two prevalent interaction paradigms — illustrate this. In Fig. 1A (unstructured single-output), users iteratively engineer prompts to generate a new output or refine the previous output. For instance, this would be equivalent to people asking ChatGPT to write a story about a rabbit and it returns one particular story; if they are not satisfied and want to refine or tweak the story, they will change (or

refine) the prompt (Fig. 1a1) to find another version of the story. On the other hand, Fig. 1B (unstructured multi-output) represents an interaction where people ask the system to generate multiple responses (e.g., `give me five different stories about a rabbit`, `give me five different ideas on ...`) or the system generates multiple outputs for its users (e.g., [11, 26]).

In either case, if users' prompt is concise, missing enumeration of key details like setting, plot, character personality, tone, and so on, LLMs will fill in those missing details in their output. For example:

*Once upon a time, in a peaceful meadow [setting] nestled between rolling hills, there lived a small and **curious** [personality] rabbit named Flora. Flora's fluffy, white fur and bright, inquisitive eyes made her stand out among her fellow rabbits. She was known for her **adventurous spirit** [personality] ... Flora shared her stories of the underground world and the Heartbloom. From that day on, the meadow [setting] became a place of wonder and adventure [tone] for all the rabbits ...*

Upon reading, users may realize they do not wish for the character to possess curious or adventurous traits, nor desire the story to center around an underground adventure. They recognize the need for further specification of essential story elements, such as the plot and character personality. They refine their prompt, providing additional details: 'Write a story about a rabbit named Jim, renowned in his village for his reserved personality. The plot of the story should revolve around...' As they iterate, they gradually uncover various aspects crucial to storytelling. However, as these aspects naturally emerge, they may not think to explore alternative ideas along those aspects and engage in divergent thinking. This constrained exploration can result in fixation on a single narrative direction, such as narratives that vary along certain dimensions (e.g., setting) but focus solely on a shy character. Consequently, this constrains the scope of creative exploration and leaves them stagnant in a small region of the design space. This highlights two issues: (1) **unstructured generation**: the outputs are generated in a haphazardous manner; (2) **unstructured exploration**: users' exploration of the design space is analogous to rummaging through the dark without a clear sense of direction and comprehensive view of the design landscape they are treading.

3.2 Structured Multi-Output Approach (Fig. 1C)

3.2.1 Motivation. This framework is motivated and grounded in research that uncovered (1) the challenges associated with dimensional reasoning despite its advantages and (2) the benefits of steering users toward parallel exploration of ideas in creative processes.

First, people normally do not reason with dimensions, mainly due to the challenges associated with identifying (and remembering) the list of relevant dimensions. MacNeil et al. [50] found that even experts struggle when asked to compile a list of relevant dimensions within their domains of expertise. However, if available, they can serve as valuable tools for systematically comparing ideas and even catalyzing the generation of new ones. After all, dimensions serve as pathways for contemplating and exploring the design landscape. Since LLMs have shown an impressive capability, there are reasons




Figure 2: An example showing the benefits of exposing design space to users. By generating relevant dimensions for a task or topic, we can inform users about many dimensions and [values] they can consider — e.g., plot complexity: [linear, twisting, ...], genre: [fantasy, adventure, ...]. If all the possible dimensions and values are used to generate responses, the responses will cover many sub-spaces within the overall design space and reveal to users a space of possible responses they can generate with LLMs.

to believe that they are also capable of recognizing and generating relevant dimensions for constructing appropriate design spaces.

Second, the literature on creativity consistently underscores our tendency to prematurely converge [19, 24, 36, 37], unless interventions are employed, such as explicitly requiring the generation of multiple ideas at the start of the creative process [24, 54]. For instance, prior work compared the effect of mandating the generation of multiple ideas before providing feedback with providing instant feedback on a single idea. They found that divergent ideation promoted comparison and exploration, and led to better results and self-efficacy [24]. This comparison between divergent and serial methods, drawn from prior work, bears a resemblance to the contrast between current interaction paradigms and our framework, as depicted in Fig. 1. The current interaction paradigms (Fig. 1A & B) resemble the serial method, where users are encouraged to promptly evaluate the response. In contrast, our framework (Fig. 1C) initiates by generating dimensions and responses, thereby engaging users in reflective exploration of the design space before evaluation.

3.2.2 Design Goals. Motivated by the challenges hindering dimensional thinking and creative processes, we propose the following design goals (DGs):

- **DG1. Reduce the difficulty of generating dimensions by harnessing LLMs.** Although challenging even for experts, generating relevant dimensions for a specific topic, task, and domain can be effortlessly accomplished by LLMs, thanks to their extensive training on terabytes of data. Thus LLMs should be leveraged to ease the process of generating dimensions.
- **DG2. Leverage dimensions to support structured generation of responses.** The goal of our framework is to not only help generate responses that reside in many sub-spaces within the design space but also do so in a structured, less haphazardous

manner. By using dimensions, users can gain a better understanding of why and how certain responses were generated and what additional responses can be generated. Users should be able to use dimensions and their corresponding values to generate responses with those specific values.

- **DG3. Support flexible transition between divergent and convergent thinking.** While our framework places significant emphasis on the divergent phase of the creative process, it was developed with the understanding that it should seamlessly integrate into the entire creative process. Therefore, its implementation should enable a flexible and iterative interplay between divergent and convergent thinking. Users should be able to quickly transition from exploring ideas to evaluating specific ideas and vice versa, as needed.
- **DG4. Enable users to engage in parallel exploration of ideas at every opportunity.** As previously mentioned, parallel exploration of ideas can reduce fixation, lead to better and more diverse outcomes, and enhance self-efficacy [24]. Even when user is in the convergent phase, there should be no friction for users to switch to divergent phase and initiate parallel exploration of ideas. In fact, the system should encourage users to engage in creative exploration by initiating parallel idea exploration whenever users prompt it.
- **DG5. Support efficient and structured navigation of the design space.** Nonlinear canvases are more challenging to navigate than linear conversational interfaces [60], as users have a greater degree of freedom in their movement — i.e., they can move not only vertically but also horizontally. However, given that the design space is represented as 2-dimensional, nonlinear space with ideas positioned within it, we should make it easy for users to locate, navigate to, and examine responses of interest.




Figure 3: Luminate interface consists of two main sections: (A) text editor and (F) exploration view. Users can (B) type and use various (C) text styles (e.g., TITLE, normal text). They have the option to (D) input a prompt (e.g., Give me a poem about ocean) to LLM and view (E) one of the generated responses in the text editor. (F) In the exploration view, users can observe (J) a cluster of responses and (H) dimensions. At first, no dimension is selected, to invite users to explore the design space. Users can add dimensions to the axes (I) and arrange responses by their dimension values (see Fig. 5). For optimal experience, we promptly display (G) dimensions as soon as they are generated, allowing users to examine them while responses are being generated. Users can click (K) semantic level icons to adjust the zoom scale and view responses at different levels of detail (see Fig. 8).

4 LUMINATE

In this section, we first present an overview of the Luminate interface, detailing its individual features, and demonstrate its workflow with a creative story writing task. Finally, we provide a concise overview of its implementation.

4.1 Interface & Features

The Luminate interface has two main sections: text editor (Fig. 3A) and navigation canvas (Fig. 3B). In the text editor, users can type and format text. They can also use the chat input box (Fig. 3D) at the bottom to ask LLM for ideas. On the right side of the interface (Fig. 3B) is the exploration view that displays generated responses. As shown, the responses are initially presented as a single cluster (Fig. 3G). When users select a node (●), this adds the full text to the text editor and highlights it in yellow, as shown (Fig. 3E). The selected node is also highlighted in yellow to indicate the mapping with text (Fig. 3J) and show which node is currently selected. To see other responses, users can click other nodes, replacing the text highlighted in yellow with the response the selected node matches. Once users edit this text, the yellow highlight disappears to indicate that it is now a part of the users' writing. We explain each feature and corresponding interactions next.

4.1.1 Dimension Generation (Fig. 4, DG1). Users can initiate structured exploration by leveraging dimension generation in Luminate. After users submit prompt (e.g., Give me a poem about ocean) in the chat input box (Fig. 3D), Luminate takes that prompt and instructs an LLM to generate a pre-defined number (e.g., 5) of categorical dimensions and corresponding values (e.g., Wave Intensity: [rolling waves, gentle ripples]). To extract all possible dimensions, we instruct LLMs to generate ordinal and nominal dimensions for categorical dimensions. We exclude numerical dimensions such as word length, as our testing and pilot studies showed they do not deliver meaningful interactions and benefits.

4.1.2 Dimension-Guided Response Generation (Fig. 4, DG2). After dimensions are generated, Luminate uses those dimensions and their values to guide the response generation. For example, suppose two of the generated dimensions and their [values] are: Narrative Perspective: [first-person, third-person] and Poetic Style: [haiku, sonnet, free verse]. Luminate randomly selects a value for each dimension to create a *requirement* — a list of dimensions and their corresponding values (e.g., Narrative Perspective: [first-person], Poetic Style: [haiku]) — and instructs the LLM to generate responses using these requirements, as illustrated in Fig. 4. Fig. 2 and Table 6




Figure 4: Technical pipeline of Prompting for Design Space: The pipeline consists of two LLM prompting steps. (A) The first LLM step feeds the inputs from the text editor into the LLM API call to generate dimensions and their values as a JSON object. (B) Luminate transforms the generated object to form a list of requirements per response by randomly selecting a value from each dimension. The second LLM step feeds each list into separate LLM API calls and generates all responses in parallel. (C) Generated responses are then visualized in Luminate’s response space.




Figure 5: Dimension selection: upon (1) selecting a dimension (Mood) for the x-axis, responses (2) reposition to vertically align to their dimension values (Romantic, Somber, Cheerful, Vengeful). (3) Selecting a dimension for the y-axis (Tone) additionally repositions the responses to align horizontally to their dimension values (Motivating, Hopeful, Inspiring, Peaceful).

in Appendix A show few examples of the dimensions and responses generated for different creative writing tasks.

4.1.3 Dimension Selection (Fig. 5, DG5). Users can cluster and arrange generated responses by selecting dimensions. This allows them to quickly locate the right responses but also explore the design space in a structured manner. For example, as shown in Fig. 5, once users select a dimension (e.g., Mood) or two, the layout of the nodes changes, from (1) a cluster to (2) a one-dimensional or two-dimensional scatter plot.

4.1.4 Generation of New Responses (Fig. 6 & 7, DG4). Luminate offers three novel techniques for steering the generation of new responses. Specifically, users can add additional responses or add additional attribute(s) to existing responses. The first two support the generation of responses with specific attributes within existing dimension(s). Users can select the **More Like This** button on the node to generate responses containing the same attributes (Fig. 6A). Or they can first select label(s) to specify a subspace in the design

space and then the **Add More** button to generate new responses in the subspace, as shown in Fig. 6B-C. The last technique is adding a user-defined dimension. As shown in Fig. 7, after users add a new dimension, Luminate generates its values and then updates existing responses with new requirements (Fig. 4B) containing these values.

4.1.5 Semantic Zoom (Fig. 8, DG5). Motivated by recent work that demonstrated the benefits of offering users the ability to control the level of detail for LLM-generated text [8, 20, 60], Luminate uses *semantic zoom* to give users control over the granularity of detail, facilitating their navigation within the subspace. As shown in Fig. 8, there are five levels of detail: **DOT**, **TITLE**, **KEYWORD**, **SUMMARY**, and **FULL TEXT**. Users can manually zoom in and out (by scrolling on the mouse or pinch-to-zoom) or click the semantic level icons (Fig. 3K) to adjust the view to the right level of detail. Users can also double-click on an empty space in the canvas to fully zoom in to view the full response. Double-clicking from the **FULL TEXT** level brings them back to the **DOT** level.




Figure 6: Two methods for generating similar responses: Method#1: (A) Users can click the More Like This button on the node to generate similar responses. Method#2: (B) When nodes are filtered on the canvas, the user can click the Add More button to generate more nodes to be added to the filtered subspace; (C) newly generated nodes join the canvas and align themselves to their dimension values.




Figure 7: Generating new responses via user-defined dimension: Users can generate new responses by (A) adding new dimensions to their exploration space. Upon entering their desired dimension in the filter bar, Luminate generates (B) values (e.g., *Victorian*, *Medieval*, *Colonial*, *Apocalyptic*) for the desired dimension (Time Period) and revises all existing responses to include one of the value for that response. Once complete, users can use the new dimension to search, filter, and cluster the revised responses.

4.1.6 Search & Filter (Fig. 9, DG5). Users can use the search box to search for responses containing certain keywords. Once users type a keyword into the search box, Luminate *hides* responses that do *not* contain those keywords by lowering their opacity. Users can also use the filter bar to filter responses by dimension values. These search features help users quickly find responses of interest.

4.1.7 Contextual Generation (Fig. 10, DG3-4). During our pilot studies, we found participants desired to generate responses that align with the writing they had added. To support this, Luminate provides the LLM with the content present in the text editor as background context. However, there can be cases where users may want a more targeted generation — e.g., expand on specific parts of the content as opposed to follow up on the entire content. Thus Luminate allows users to highlight text in the editor when prompting for responses. We add the highlighted content into the prompt along with any other content present in the text editor (see Table 5). This provides users with more contextualized responses and a way to precisely steer the generation of responses.

4.1.8 Design Space Exploration (Fig. 11, DG4). When users engage in creative tasks, it is natural for them to go back and forth between a set of ideas. Thus, Luminate supports a space-switching feature that enables users to change their selected text and explore different design spaces, which are sets of responses from different prompts, at anytime they want. As shown in Fig. 11, for example, a user has recently added a new stanza to the ocean-themed poem. In order to enhance the alignment of length and lexical level between the previously generated stanza and the newly added one, the user clicks the *show space* () button and searches for an alternative stanza.

4.2 Example Workflow: Writing Short Story

Below we present an example workflow to demonstrate some of the features described above. Chris is a professional writer who writes children's books. He is attending a workshop for writers interested in learning about ways AI tools can help them in their creative writing process. The workshop provides Chris access to Luminate and suggests he write a short story over the weekend and share it with other writers on Monday.

4.2.1 Divergent Thinking Phase (Exploration). Chris opens Luminate on the web browser. He is greeted by the Luminate interface




Figure 8: Semantic zoom: The amount of information users see changes with zoom scale. As shown, as users zoom in, the displayed information transitions from **T** TITLE to **ABC** KEYWORD, then to **≡** SUMMARY and to **≡** FULL TEXT. Users can double click anywhere on canvas to zoom directly into the **≡** FULL TEXT level or zoom out to the **≡** DOT level from the **≡** FULL TEXT level, for efficient navigation of the space.




Figure 9: Search features: To address the challenge of *locating* the desired response with specific properties, Luminate offers two search methods. With (a), users can search for responses containing specific words such as ‘love and peace.’ As shown, instead of completely removing the non-corresponding responses, we reduce their opacity, to preserve their spatial information while making the corresponding responses stand out. With (b), users can filter for (A) bookmarked nodes or by (B) dimension values using the dropdown. Alternatively, as shown, users can also (C) click on the dimension value (**Somber**) to apply the filter.

(Fig. 3), with text editor on one side and exploration view on the other side. The text editor and exploration view are both empty. He writes a header ‘My Story’ at the top of the text editor. Without any specific idea on what he wants to write about, he first recalls recent events and encounters for inspiration. He remembers running into a rabbit outside his house in the morning. Unsure what interesting story he can write about, he decides to ask AI. He writes in the chat input box: `write a story about a rabbit`. Luminate starts to generate 30 stories. After about 5 seconds, Luminate shows several toast messages (e.g., Fig. 3G), with each specifying what dimensions LLMs found relevant: plot, setting, genre, character type, tone, and originality. Then after another 3 seconds, he sees a story highlighted

in yellow in the text editor and a cluster of dots in the exploration view.

4.2.2 Convergent Thinking Phase (Evaluation). Chris selects several dots randomly, which updates the text in the text editor, showing him the full text for each dot. After browsing a few and reading only the first one or two sentences, he zooms in further, which turns the dots into ellipses with titles of the story. They read: `The Rabbit’s Journey`, `True Love`, `Mystery Rabbit`, and so on. `Mystery Rabbit` triggers his curiosity. He zooms in further to see keywords of the `Mystery Rabbit` story, which reads: `Forest`, `Brave Young Squirrel`, `Uncover Secrets`. Zooming in further reveals the summary of the story: ‘A Mysterious rabbit lurks in a

My Poem

In the hush of twilight's swoon,
A tidal wave meets the crescent moon.
Silver light glimmers on the crest,
As the sea is put to a shimmering test.

Ask AI for ideas

My Poem

In the hush of twilight's swoon,
A tidal wave meets the crescent moon.
Silver light glimmers on the crest,
As the sea is put to a shimmering test.

A tidal wave meets the crescent moon.

|

My Poem

In the hush of twilight's swoon,
A tidal wave meets the crescent moon.
Silver light glimmers on the crest,
As the sea is put to a shimmering test.

A tidal wave meets the crescent moon.

write the next stanza focusing on this line

Figure 10: Contextual generation: Users can select text in the editor (1) to highlight a line or phrase (e.g., A tidal wave meets the crescent moon) to focus on as context. Users can then (2 & 3) prompt Luminate to generate responses given the focused context. This allows users to have a more fine-grained control over the generation of responses.




Figure 11: Switching to other design space: (A) Clicking the *Show Information* (ⓘ) button displays the current design space prompt ('add one stanza about life in the ocean') along with its unique space ID. (B) The *Show Space* (📍) button allows users to switch back to the design space for prompt 'write a poem about ocean'. (C) When the design space switches, the exploration view is updated, to show a cluster of responses, along with (c1) a node that corresponds to the (c2) text block in the editor highlighted, and (c3) dimensions in this design space (Space ID = 2). (D) Clicking on another node updates the content in the (d1) block to show its full text.


forest, and a brave young squirrel sets off to uncover its secrets.' He realizes the protagonist of this story may be a squirrel, and the antagonist a rabbit. He had not considered this but realizes this might be an interesting direction he can take as well. The attributes listed below the summary confirm this, as it shows the Character

Type: Antagonist tag. The other attributes — Plot: Mystery, Setting: Forest, Genre: Fantasy, Tone: Mysterious/Suspenseful — inform him what this story is about but also make him think about what other genres or settings he could explore. Although he can select

the node and read the full text in the editor, he decides to zoom in to reveal the full text in the node. It reads:

Once upon a time, in a mystical forest, there lived a mysterious rabbit... The rabbit seemed to be up to no good, and his motives were unclear... One day, a brave young squirrel decided to confront him and ask what his mission was. Before she could get her answer, however, the rabbit disappeared into thin air! The squirrel decided to find out more about this mysterious creature and set off on an adventure to uncover his secrets...

Chris thinks this could be one starting point, but he is also curious if he can explore more stories like this one. He clicks a  button. Luminate shows  and after 3 seconds, adds 5 more stories into the space.

4.2.3 Divergent Thinking Phase (Exploration). He zooms out to the TITLE level to quickly grasp what the stories are about. He reads the titles of these new stories. The title for one of them reads, *Mystery Quest*, and the summary goes: ‘A white rabbit embarks on a journey to uncover the truth behind mysterious disappearances in the forest.’ While interesting, he is not ready to commit to this direction, because he knows he has not yet explored all the dimensions. So he clicks the bookmark icon () to save the story. Then he decides to explore more broadly by manipulating dimensions. He clicks the  button and selects Genre. The labels representing values along the Genre dimension — *Mystery*, *Fairytale/Folklore*, *Romance*, *Sci-fi/fantasy*, *Horror*, and *Adventure* — are added to the x-axis. The nodes that were in a single cluster moved to align with the corresponding labels. Intrigued by how he can think and explore in terms of dimensions, he selects another dimension, Setting, which replaces previous labels in the x-axis with labels such as *forest*, *urban area*, *underground*, and *underwater*. He is surprised to see *underwater* and zooms into the SUMMARY level to see some interesting stories about a rabbit named Tony in an underwater kingdom. He bookmarks the story, zooms out, and explores a few other dimensions as well.

4.2.4 Convergent Thinking (Synthesis). Eventually, he feels he has collected enough ideas to write a story. He clicks a bookmark icon () to find the stories he saved — some for the phrases or sentences he particularly liked and others he wanted to use as potential ideas for aspects of his story such as setting, conflict, and tone. He clicks the nodes to display the text in the text editor and edits it, which embeds that text block into the text editor. When he selects another node, it appends the text under the previously embedded text. Chris goes on to edit the text to write a draft of his story.

The scenario above demonstrates a flexible and iterative interplay between divergent and convergent thinking. This workflow, however, is one of numerous potential workflows and does not encompass all the features within Luminate.

4.3 Implementation & Prompt Engineering

Luminate was implemented using React, Editor.js, and d3.js. The text editor was implemented using Editor.js, and the exploration view and animation effects (clustering and positioning of points) with d3.js. Specifically, we used the force layout mechanism in d3.js for the clustering of points.

We used OpenAI’s text-davinci-003 API to generate dimensions, dimension values, and responses. Given the user’s prompt and context (Fig. 10), Luminate first prompts the LLM to generate a set of categorical (ordinal and nominal) dimensions followed by their dimensional values (see Table 6). Once dimensions and their values are generated, Luminate then forms requirement lists (e.g., [40]) with one random value from each dimension, and for each requirement list, prompts the LLM to generate responses that fall under those requirements. Once each response is generated, Luminate additionally prompts the LLM to generate abstractions of the response for each semantic level (Fig. 8). All subsequent generations of new responses repeat the process of generating new responses but form requirement lists from the filtered dimension values instead. All prompts used in this pipeline are listed in Table 4 - 6 in Appendix A.

5 USER STUDY

To investigate the feasibility and usefulness of our framework, we conducted a user study with Luminate. We recruited professional writers with experience using AI tools for their creative writing tasks in order to understand how Luminate compares with existing AI creativity support tools. The screen and audio were recorded for accurate transcription and analysis. For their participation in an 80-minute study, they received 40 USD.

5.1 Participants

Fourteen professional writers (9F, 5M; P1-P14; $M_{age} = 30$, $SD = 7.7$) were recruited via Upwork [2]. All of them had several years of experience ($M = 7.3$; $SD = 8.4$, range = [2, 35]) in creative writing, had experience incorporating AI-generated writing into their writing, and had exposure to various AI tools (14 ChatGPT, 4 JasperAI, 4 DALL-E, 1 Stable Diffusion, 1 frase.io, 1 QUILLBOT, 1 Bard, 1 Sage Poe, 1 Claude). They shared that approximately 26% of their final writing was AI-generated ($SD = 15.4\%$). Most of them frequently used ChatGPT (6 on a daily basis; 6 on a weekly basis; 2 on a monthly basis). They varied in terms of how they used AI tools. Almost all of them (13/14) had used them for idea generation (e.g., generate creative ideas, story concepts), 11 for editing and proofreading (e.g., grammar and style checks), 7 for assistance in descriptive writing (e.g., add descriptions of settings, scenes, or characters), and 5 for content expansion (e.g., expand on existing content), 5 for world-building (e.g., create worlds with cultures and settings, for use in novels, games, or other creative projects), 5 for getting suggestions on character names and book and chapter titles, and 1 for creating different versions of the writing.

All participants had experience engaging in more than one creative writing task professionally. All 14 participants had experience in copywriting. Eleven had written a story and 10 an email (e.g., marketing email); 6 had experience in writing a letter (e.g., cover letter) and 4 a poem; there was one participant with experience in various tasks, including screenwriting, how-to guides, travel guides, product reviews, blog writing in their respective professional experience, composing song lyrics, writing story outlines for video games, comics, article, character development, and sci-fi stories. To hear them compare their experience using Luminate with other AI tools they have previously used and investigate the scalability of this approach, we assigned creative writing tasks based on their

prior experiences: 4 participants were asked to do copywriting, 4 crafting a short story, 2 writing an email, 2 writing a poem, 1 constructing video game scene, and 1 composing song lyrics.

5.2 Procedure

Set-up and Pre-study Survey (10 min). All 14 studies were hosted online on Zoom. After receiving participants' signatures on the consent forms and their approval of being recorded video and audio for later transcription and review, the study started off with a pre-study questionnaire to collect their demographic information (i.e., name, age, gender), experience in creative writing (i.e., years, types), and previous exposure and experience using AI tools in creative writing (i.e., frequency of usage, kinds of AI tools, type of creative writing tasks).

Interface Tutorial (25 min). After completing the pre-study survey and receiving the link to Luminate, participants shared their screen and granted the researcher remote control, which enabled the latter to demonstrate how to use the system. In the tutorial session, researchers first introduced the interface to the user and then showed interactions via a practice task — `write a story about a rabbit`. The researcher explained the purpose and functionality of each interaction and demonstrated a few actions at a time, followed by users performing the same actions as practice. This ensured participants had enough practice prior to the actual task and could ask the researcher if anything was unclear. The researcher followed the script containing detailed instructions, to ensure all participants thoroughly understood all the features in the system, and undergo the same tutorial.

Creative Writing Task with AI Support (25 min). Once they were familiar with the system, participants had up to 25 minutes to use the system for a creative writing task. Participants performed think aloud throughout the whole section. Participants could choose any topic they want in a category decided by the researcher based on their previous experience (see Table 7 for the list of topics and writing tasks). During the study, Luminate was configured to generate up to 5 nominal dimensions with up to 8 different labels and up to 3 ordinal dimensions with 5 levels (least, less, neutral, more, most). Based on the feedback from our pilot studies, we limited the length of LLM-generated responses to 150 words and set the number of responses per prompt to 40, to avoid overwhelming participants but provide enough variations for creative exploration.

Post-study Survey and Exit Interview (20 min). Following the completion of the task, a post-study survey was administered. It contained: the Creativity Support Index (CSI) questionnaire, questions related to the quality of AI-generated dimensions, and the evaluation of system's usability and usefulness of individual features. The study ended with a semi-structured interview in which users provided qualitative feedback on the system regarding the quality of the dimensions, compatibility with their workflow, usability, and effectiveness of features. For example, we asked them to compare Luminate with other AI tools and how they might integrate Luminate into their workflow.

6 RESULTS

We present analysis of survey responses, participants' interactions with Luminate (e.g., how many responses they explored), observation, and interview, describing participants' general assessment of Luminate, their workflows, as well as how Luminate helps steer users toward divergent thinking and enables an understanding of the design space.

Table 1: Creativity Support Index (CSI) Results (N=14). The highest value is in bold. The second highest in underline. Since our study did not involve collaboration, we followed the practice from [15, 61], omitting the Collaboration Factor to avoid confusion.

Factor	Av.g. Score (SD)	Avg. Factor Count
Exploration	18.21 (1.76)	4
Enjoyment	<u>17.85</u> (1.81)	1.79
Expressiveness	16.71 (2.58)	3.07
Results Worth Effort	16.57 (2.08)	<u>3.57</u>
Immersion	15.07 (3.01)	2.07
Overall CSI Score	82.16 (5.14)	

Participants generally found Luminate usable (EASY TO USE: 5 Strongly Agree, 5 Agree, 3 Neutral, 1 Disagree; EASY TO LEARN: 7 Strongly Agree, 4 Agree, 3 Neutral). Five participants who answered neutral or disagreed explained there was not enough time to learn and utilize all the features to their potential. P6 said: "The program was fairly easy to use and intuitive. I feel it had a lot more to offer, but we had limited time."

All of them except one agreed they would use the system for creative writing tasks (6 Strongly Agree, 7 Agree, 1 Neutral) and for exploring ideas in general (9 Strongly Agree, 4 Agree, 1 Neutral). They also found Luminate provided creativity support, scoring CSI at 82.16 (SD = 5.14) — a high score for creativity support index [17]. As shown in Table 1, while Luminate scored relatively high in all dimensions, participants found it provided the strongest support for the exploration of ideas (M = 18.21, SD = 1.76) and regarded it as the most important aspect (Avg. Factor Count = 4).

6.1 How Can Luminate Be Used in Creative Workflows?

We report three workflows (Ws) observed from the study, demonstrating the flexibility and scalability of our system and framework.

W1. Using One or Multiple Responses from a Single Prompt (1/14). P2 found that a single batch of responses provided enough for his creative writing task. For example, tasked with copywriting on the topic of 'Latest Trend of UI Design across the Globe,' P2 selected responses to UI design trends varying from Asia to South America. He incorporated three responses covering three different continents (Asia, South America, and Oceania) into the text editor as three paragraphs, then edit and merge them. This workflow suggests that a single batch of responses can provide enough assistance for some creative writing tasks.

W2. Using Responses Generated from a Chain of Prompts (8/14). Most participants employed a chaining of prompts, whereby they progressively refined their creative output by iterating through prompts and responses. As P3 was writing a short story about ‘forgotten memories,’ P3 first found a story about friendship and then asked AI to generate a detailed profile of the group of 5 friends. Eventually, the participant prompted AI to generate a final closing paragraph on the moral insight of the story. This iterative approach enabled P3 to build upon numerous ideas and refine his initial ideas, ultimately leading to complex and nuanced creative outcomes.

W3. Using Responses Generated via Dimension-Driven Generation Techniques (5/14). Several participants leveraged novel features provided in Luminate for generating additional responses. P4, who wrote ‘Adventures of Tom Sawyer, the Time Traveller’, added label *future* in the Time Period dimension during her exploration to generate story versions that have the *future* element. P7 also used the same feature to generate new results. Tasked with writing an email where the instruction read, ‘write an email to ask a professional translator to step down from a project,’ P7 steered the system towards producing responses that aligned more closely with her evolving creative vision, allowing her to get the version that matched the desired tone and formality.

6.2 How Does Luminate Help with Divergent Thinking?

The analysis of participants’ survey, interview responses, and exploration suggest that Luminate helps users engage in divergent thinking. On questions ‘The tool was helpful in allowing me to track different ideas, outcomes, or possibilities’ and ‘It was easy for me to explore many different ideas, options, or outcomes, using this tool,’ Luminate received an average score of 9/10 (1 as Strongly Disagree, 10 as Strongly Agree; SD = 1.1).

It helps by generating multiple responses even from a simple prompt. All the participants reported that having multiple options from one prompt was beneficial for the brainstorming and ideation stage of creative writing. They found that it offers variation, serendipitous responses (P2), and the ability to mix-and-match different responses (P5) and think outside the box (P10). P2 said: “This AI gives a lot of ideas and options to choose, from just one simple prompt. I think this is very powerful.” As shown in Fig. 12a, all the participants also agreed that the system helped discover new content (11 Strongly Agree, 3 Agree). For example, P3 said having access to a diverse set of ideas inspired new ideas: “while seeing these options, another ideas come.” On average, each participant prompted 4.3 times (SD = 3.9, range = [1, 15]) and explored 13.8 responses (SD = 7.6, range = [4, 26]) during the 25-minute study session. As a result, participants suggested that Luminate can help avoid fixation. P6 said:

“It gives you a range of possible scenarios. This broadens your mind so you won’t be fixed on something. I think it’s the way. It gives me options and lets me know there are multiple pathways I can choose.”

It helps with thinking outside the box. In addition to benefiting from a sheer number of options, participants also felt the responses were creative, saying that Luminate “showed ideas [they] did not think of” (P5) and “helped think outside the box” (P2). P2 elaborated

by saying that whereas what we can see is constrained by a limited set of experiences and contexts, “AI can go 1,000 or 1,000,000 ways different than [our] context because it has all the intelligence it learned over time.” She further highlighted that it can broaden our horizon by making different perspectives more accessible, saying: “it is a good way not only to enhance creativity, but also to see other points of view from different kind of people, from different kind of situations that obviously you do not have access to.”

It helps by shifting the focus from prompt engineering to exploring and generating new ideas. The availability of options and their creativeness made it possible for participants to break free from the iterative prompting and focus instead on exploring and generating new ideas. P6 said: “I like this better than just having one option and then me having to fine tune everything or tweak it up to my liking. This way, I have more choice and more freedom. I don’t want to just get something spit out and then rely on it.”

The flexible interactions and exploration view enable creative exploration. Many found the various interactions in Luminate useful for creative exploration. P5 highlighted the ease with which she could vary the responses, saying: “there’s so many customization options and it generates even more ideas than you could have thought of. You are able to generate different elements of the story like the plot, the language, the tone, the characters, the setting. And you can keep customizing and the result it brings out actually makes sense.” In addition to flexible interactions, several mentioned the exploration view as their favorite feature as it enabled creative exploration. They elaborated, saying they enjoyed seeing “everything” (P1) and “more area of thinking” to see “variations” and “expand the story” (P3).

6.3 How Does Luminate Help with Developing an Understanding of the Design Space?

Participants generally agreed that Luminate helped them develop a better understanding of the space of: (1) possible responses (7 Strongly Agree, 3 Agree, 2 Neutral, 2 Disagree); and (2) the limitations and capabilities of the AI (7 Strongly Agree, 4 Agree, 3 Neutral). Participants who disagreed or answered neutral attributed the reason not to Luminate but to lack of time. For example, P8 said: “These are useful, but I would need more time because I am a slow learner. I would need to spend maybe a week just playing around.” Participants also agreed that using Luminate ingrained in them a sense that they can now foresee possible response space of their prompts (8 Strongly Agree, 5 Agree, 1 Neutral). On average, participants selected the dimension button 1.8 times (SD = 2.1, range = [0, 7]), exploring 1.1 unique (SD = 1.2, range = [0, 4]) 1-dimensional and 1.2 unique (SD = 1.7, range = [0, 4]) 2-dimensional design spaces.

Dimensions help with reflecting on the design space. All participants found that our approach, especially the generated dimensions, helped them reflect on what factors are important for their task and topic (6 Strongly Agree, 8 Agree). Participants mostly found the generated dimensions to be accurate (4 Strongly Agree, 6 Agree, 3 Neutral, 1 Disagree) and found them very useful (7 Strongly Agree, 7 Agree) and relevant (4 Strongly Agree, 9 Agree, 1 Neutral) for exploring ideas and reflecting on the design space. Moreover, it reminded them of aspects of the task they had never thought about.




Figure 12: Evaluation of the system and the dimensions generated from the prompt

P7 said: “it made me think about the level of empathy, the positive or negative aspect of [tone] that I never think about.”

Dimensions help with structuring thinking around the design space. Most participants found that Luminate helped them structure their thinking around the design space (3 Strongly Agree, 6 Agree, 4 Neutral, 1 Disagree). P1 and P5 explicitly attributed this to dimensions. P5 explained that dimensions helped break down different elements of the storyline and even in parts where she did not expect to help her think about the design space, saying:

“I find it fascinating that it was able to break down different elements of the storyline... the different story blocks, different elements that make up the story, like the setting, the character, the time, the tone and all of that. I was not expecting that there would be a tool that could really pull them out. You give it a prompt and it pulls out several different story elements and there are so many ideas to explore. I did not expect there would ever be a tool like that. And then it categorizes all of this into different elements and it’s really impressive... breaking down into different dimensions helps me think about the story.”

Exploration space helps understand the design space. All participants agreed that the exploration view in Luminate was also important in helping them achieve a comprehensive understanding of the design space (7 Strongly Agree, 7 Agree). They also agreed that interactively building new design spaces by selecting dimensions

was not only useful for building a comprehensive understanding of the space (3 Strongly Agree, 10 Agree, 1 Neutral) but also for exploring the space of possible responses (7 Strongly Agree, 5 Agree, 1 Neutral, 1 Disagree). P6 praised this ability to explore in the space, saying:

“The brainstorming aspect of it is very, very helpful because it shows you all bunch of options right there. It helps you navigate exactly where you want to go with a certain topic and that’s what I would predominantly use it for.”

6.4 How Does Luminate Compare to Existing Approaches? What Benefits or Challenges Does This Approach Introduce?

At the end of our interview, we asked participants to share: (1) how their workflow with Luminate differed from their workflow with other AI tools they used for creative writing, (2) which they preferred, and (3) why. While our participants had exposure to various AI tools, because ChatGPT was the tool they used the most, all of them compared Luminate with ChatGPT.

It is better for creative tasks. Most participants (10/14) expressed preference for Luminate for creative writing. P7 explained that he preferred ChatGPT, because it is simple and he did not need a long and complicated response for the writing tasks he engages in. Participants also suggested that the scenario where Luminate fits may be contingent on the flexibility of the requirement for the

writing task and the task complexity. If the project is more open-ended, Luminate is better because it gives writers the opportunity to aberrate from their normal path. For example, P1 noted that “for projects that are creative and flexible, [such as] brainstorming or working in marketing campaigns, art development, it could be amazing.” P14 was particularly impressed with the way ideas are organized for creative exploration, saying:

“Where it blew every other tool out of the water is being able to present a range of options. Not just present them, but being able to present them in a way where you’re not just sifting through pages of text. You can see everything and move around by categories. As a brainstorming tool, it’s crazy.”

The way Luminate generates many diverse responses was also seen as an approach that better aligns with how they think and work as creative professionals. P8 said: “The AI first determines some important aspects and then give you 40 different responses; iterations like this is a technique I use a lot as an artist.” P3 testified this is also the case for writing: “as a writer, I always look for more and more options, so it is good that I can explore all the options and then select one.”

It simplifies the workflow by providing dimensions as an outline. Several participants shared they generally write an outline first and then use AI to expand on it. They said Luminate simplifies this workflow, equating Luminate generating dimensions as creating an outline. P5 said: “in my current workflow, I create an outline first. A structure if I don’t already have that provided to me. Then I use AI to generate the content based on what I already have. But with a tool like this, I feel like it would save you the time of creating an outline.” P4 envisioned that dimensions would make her workflow more efficient and improve her writing, saying: “I’ll take advantage of the dimensions to simulate the best way to go about my write up. It would make it easier to finish on time.”

It empowers users as it expands their prompt and allows them to steer the response generation. Participants also appreciated the flexibility and steerability of Luminate. P5 found that she was able to customize the generation to produce a desired writing, saying: “I was able to explore more ideas and customize storyline. What was generated at the end matched everything that I had selected — every prompt and every idea.” The way Luminate expands a prompt was also seen as a unique feature that encourages deeper thinking. P10 said:

“ChatGPT just brings out one story. It doesn’t care if you want to talk about different time periods like future, modern time, or ancient time. It can be very limiting having to read a response and say, ‘I want to edit one or two things here.’ You can’t think too deeply this way. With Luminate, however, you get to explore different options, so it’ll make you think better.”

There is still room for improvements such as minimizing cognitive overload. Although many appreciated being able to see multiple responses from a prompt and felt the agency of choosing a response from multiple alternatives outperforms having only a single response — especially in the brainstorming stage, several found 40 responses overwhelming. P4 and P8 suggested showing fewer and

simpler responses at first and offering more suggestions once users have a better understanding of their needs. P4 noted that there can be two modes: basic and advanced, where users receive fewer suggestions in the basic mode but more suggestions and access to the features in Luminate in the advanced mode. P4 also mentioned a recommendation idea, stating: “Perhaps the AI can be structured to suggests options that are the best fit for the story, it will help separate the wheat from the chaff and assist users in structuring their thinking.”

6.5 Prolonged Use in the Wild: Do Participants Still Find Luminate Useful After Using It in the Wild for an Extended Period of Time?

After the study, participants showed strong interest in Luminate, asking when the tool would become publicly available. Simultaneously, we had multiple participants who said there was too little time to assess Luminate. We saw this as an opportunity to address, to some extent, the limitation of our study where our evaluation is based on participants’ limited exposure (25-50 minutes) to Luminate and under the controlled setting. Thus, we contacted the participants, asking whether they are interested in using it freely for the duration of about 2 weeks and share their experience.

Eight participants, which included those who expressed interest in using the tool and those who said there was little time to assess it, accepted our invitation. One of the authors first met with each participant to remind them about the features and give them a URL link to Luminate. Five participants met with this author twice and the other three met once over the duration of two weeks to share their experience. Participants received \$20 for attending each interview session, which took, on average, 28 minutes. Below, we share a summary of the insights from this in-the-wild study.

Participants still prefer Luminate for brainstorming and creative tasks even after prolonged use. While some preferred ChatGPT for tasks that need longer, detailed responses, all the participants still maintained that they prefer Luminate for brainstorming and creative tasks. For example, after having used it for two weeks, P4 said, “Luminate right now is actually my go to for ideas when crafting stories.” P8, who called himself “a slow learner” and needs “maybe a week” to play around with it to understand its usefulness, noted after a week of use in the wild that he finds it useful, saying: “Now I know how to use it a little better. I think it is a very useful tool. I’m now interested to find out what I can do with it.” Additionally, P14, a writer with 30+ years of creative writing experience, maintained the same enthusiasm he had for Luminate from the user study to after the in-the-wild study. He was given access to Luminate for more than two weeks, during which he extensively experimented with its limits and potential use cases. After the in-the-wild study, he was able to pinpoint the value of Luminate in supporting a specific stage in his workflow, saying:

“I think ChatGPT is useful at the mid level of designing something. It’s not good at polishing, it’s not creating an end product, but when you already have your idea and you want to build a little bit of a framework for it, that’s what’s really useful for it. [Luminate] seems useful for before that, where you’re starting from a blank page and you just want a pool of ideas

to start piecing together. It seems more like an interesting brainstorming tool that I would use during pre-production rather than during production.”

He further expressed a desire to continue using the tool and learning about other ways Luminate can be used, saying:

“I would particularly be interested in the upcoming beta tests because I’m really curious to see how other people are using it. I feel like I know the basics now. This is the point where I’d traditionally start looking to forum/Discord discussions in search of advanced applications for the tool that I hadn’t thought of yet.”

7 DISCUSSION

In addition to contributing a framework for systematically generating and exploring the design space constructed by LLMs, our work demonstrates how humans and LLMs can complement each other during creative exploration. As generative models, LLMs excel at quickly generating information. Humans, on the other hand, are slower in producing a variety of responses (that vary in tone, composition, etc.) but excel at comprehending a broader context and the requirements of the tasks at hand. They can judge which response is best suited for a given task. Additionally, the black-box nature of LLMs presents a major challenge for users when it comes to instructing the model to produce the desired output. In our approach, users gain an enhanced ability to navigate the design space and evaluate AI-generated responses using dimensions. These dimensions — along with novel features for seamlessly adding user-defined dimensions and their values to generate additional responses (and update existing responses) — provide a structured framework for users to flexibly express their preferences and guide AI towards generating responses that are not only creative but also well-suited to their task and creative goals.

7.1 Limitations

We describe several limitations in the study, to clearly define the scope of our study findings and motivate future work.

Lack of control over the length and number of responses. As explained in Section 5.2, we intentionally fixed the number of words in each response to 150 words and number of responses to 40 to support wide exploration in their writing. But we found some participants expressed a desire for longer responses, and some even wanted fewer responses. This may have affected Luminate’s perceived value in supporting writing tasks, causing participants to favor conventional AI tools such as ChatGPT for longer and more detailed outputs. Enabling users to adjust the number and length of responses could change participants to view Luminate as more versatile and broadly adoptable than shown in our studies.

Lack of testing with diverse user groups. Our study focused on creative writers with experience in using AI for their writing. Testing with diverse groups of participants such as non-professional writers would provide a more comprehensive understanding of the usefulness and scalability of our structured multi-output approach.

Lack of comparison with existing AI tools. Finally, our attempt to compare Luminate with other AI tools ended up being a comparison with ChatGPT based on participants’ self-reports. A rigorous study comparing the two with specific measures would provide deeper

insights to where Luminate stands in the landscape of AI writing tools. Admittedly, given the ever-increasing number of AI tools in the market [1], a rigorous comparison with other AI tools may, in practice, be infeasible. Regardless, it is worth acknowledging that our study does not provide an answer to this question.

7.2 Future Work

7.2.1 Addressing Study Limitations. As noted above, there are several limitations that should be addressed through: a study with a more adjustable interface, tests with diverse user groups, and comparing our approach, to the extent possible, with alternative approaches found in other AI tools.

7.2.2 Exploring New Interaction Paradigms Beyond Prompt Engineering. In light of recent research, which has predominantly focused on exploring ways to support and enhance prompt engineering, our work raises an interesting question: “Can we extend our interactions with LLMs beyond prompt engineering?” While our Prompting for Design Space framework was designed to enhance interactions with LLMs in creative processes, we believe this work can serve as a source of inspiration, stimulating new lines of research aimed at envisioning novel interaction paradigms across various workflows and contexts. Although LLMs have enabled users to execute complex actions through natural language, prompt engineering does not always provide intuitive or efficient means to perform complex actions [70]. This work shows that there could be a rich opportunity for future research in this direction.

7.2.3 Extending the Framework and Interaction Techniques to Other Domains and Tasks. In this work, we have demonstrated a framework for supporting structured generation and exploration of design space in the domain of writing support. However, we believe the proposed framework can be broadly applied to a wide range of domains and tasks, as the idea of generating multiple, diverse ideas that serves as motivation of this work can generalize [24]. For instance, our framework could go beyond text input and output to other types of media such as image and video as input and output. For example, we could have a system that takes image as an input, generates dimensions relevant to the subject or setting in the image, and use their attributes to generate diverse images for users to explore. At the same time, the interaction techniques used in the framework could also be applicable (and necessary) in other domains and tasks. For example, irrespective of the output format, semantic zoom may be essential for efficient exploration and analysis of generated responses.

7.2.4 Implementing Design and Features to Address Cognitive Overload. As mentioned, some participants found the multitude of outputs overwhelming, and suggested interaction techniques, such as initially revealing a limited number of outputs (e.g., 1-3 responses) and then allowing users to easily generate more as they move around the space. They also suggested offering basic and advanced modes, with the basic mode providing fewer outputs and the advanced mode offering more outputs, along with the ability to steer responses and explore the design space, as seen in the current version of Luminate. Another suggestion for improving the navigation and organization was to allow users to dismiss or mark the responses they no longer needed. Taking these into account, we

envision that the re-design of Luminate would involve enabling flexible transition between simpler and more complex interfaces, which is in line with prior work where they also concluded the need to support flexible transition between a simpler, familiar interface and a more complex interface to accommodate various user needs and tasks that vary in complexity during different stages of the workflow [60].

7.2.5 Generating Meaningful Dimensions. Luminate currently generates a set of categorical dimensions — ordinal and nominal dimensions. While they are useful in most cases, our own testing and user study revealed that these dimensions can vary in usefulness. For instance, creativity always appeared as an ordinal dimension with labels [least, less, neutral, more, most]. Yet, creative writers rarely, if at all, need or want a less creative artifact. Similarly, nominal dimensions occasionally demonstrated instances suggesting the need for calibration. In response to P1’s prompt `What does the world think about Peruvian food`, Luminate generated [gender] as one of nominal dimensions. Though it was *not* irrelevant to the prompt, the user did not find it useful for exploring the design space. While finding the set of dimensions that can satisfy everyone in every task may be impossible due to subtle variations in how one interprets the relevance of generated dimensions, the issue still affects the user experience and therefore should be tackled. A simple, quick solution would be, e.g., to provide users with the ability to easily delete irrelevant dimension(s) and use the option panel where they can specify dimension(s) to filter or check which dimension(s) were filtered out.

8 CONCLUSION

Our work addresses the challenge of harnessing the creative potential of large language models for creative endeavors. While LLMs offer vast generative capabilities, we argue that we are not leveraging them to their fullest potential by not systematically structuring the generation and exploration of their outputs. To tackle this challenge, we developed a framework — Prompting for Design Space — that enables the structured creation of a design space. Within this space, users can effortlessly explore, assess, and integrate numerous responses. Through the development of Luminate and a user study involving 14 professional writers, we have demonstrated the practicality and effectiveness of our framework. To our knowledge, this work is the first to explore ways to generate multiple responses from a single prompt in a systematic manner in a text-to-text context. We believe our work contributes an advancement in the way users can engage with LLMs for creative tasks, providing novel ways to harness the creative power of LLMs through structured output generation and exploration during the creative process.

REFERENCES

- [1] [n.d.]. Futurepedia. <https://www.futurepedia.io/>. Last accessed 27 August 2023.
- [2] [n.d.]. Upwork. <https://www.upwork.com/>. Accessed: September 13, 2023.
- [3] Christopher Ahlberg and Ben Shneiderman. 1994. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 313–317. <https://doi.org/10.1145/191666.191775>
- [4] 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. <https://doi.org/10.1145/3290605.3300233>
- [5] Benjamin Bach, Zezhong Wang, Matteo Farinella, Dave Murray-Rust, and Nathalie Henry Riche. 2018. Design patterns for data comics. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–12. <https://doi.org/10.1145/3173574.3173612>
- [6] Michel Beaudouin-Lafon, Susanne Bødker, and Wendy E. Mackay. 2021. Generative Theories of Interaction. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 6 (2021), 1–54. <https://doi.org/10.1145/3468505>
- [7] Michel Beaudouin-Lafon and Wendy E. Mackay. 2007. Prototyping tools and techniques. In *The human-computer interaction handbook*. CRC Press, 1043–1066. <https://www.kth.se/social/upload/52ef5ee4f2765445a466a28a/mackay-lafon-prototypes-52-HCI.pdf>
- [8] Benjamin B. Bederson and James D. Hollan. 1994. Pad++: a zooming graphical interface for exploring alternate interface physics. In *Proceedings of the 7th annual ACM symposium on User interface software and technology*. 17–26. <https://doi.org/10.1145/192426.192435>
- [9] Anastasia Bezerianos, Fanny Chevalier, Pierre Dragicevic, Niklas Elmqvist, and Jean-Daniel Fekete. 2010. Graphdice: A system for exploring multivariate social networks. In *Computer graphics forum*, Vol. 29. Wiley Online Library, 863–872.
- [10] Michael Mose Biskjaer, Peter Dalsgaard, and Kim Halskov. 2014. A constraint-based understanding of design spaces. In *Proceedings of the 2014 conference on Designing interactive systems*. 453–462. <https://doi.org/10.1145/2598510.2598533>
- [11] Stephen Brade, Bryan Wang, Mauricio Sousa, Sageev Oore, and Tovi Grossman. 2023. Promptify: Text-to-Image Generation through Interactive Prompt Exploration with Large Language Models. *arXiv preprint arXiv:2304.09337* (2023). <https://doi.org/10.48550/arXiv.2304.09337>
- [12] Dennis R. Brophy. 2001. Comparing the attributes, activities, and performance of divergent, convergent, and combination thinkers. *Creativity research journal* 13, 3-4 (2001), 439–455. https://doi.org/10.1207/S15326934CRJ1334_20
- [13] Bill Buxton. 2010. *Sketching user experiences: getting the design right and the right design*. <https://doi.org/10.1016/B978-0-12-374037-3.X5043-3>
- [14] Stuart K. Card, Jock Mackinlay, and Ben Shneiderman. 1999. *Readings in information visualization: using vision to think*. Morgan Kaufmann. <https://doi.org/10.5555/300679>
- [15] Erin A. Carroll, Celine Latulipe, Richard Fung, and Michael Terry. 2009. Creativity factor evaluation: towards a standardized survey metric for creativity support. In *Proceedings of the seventh ACM conference on Creativity and cognition*. 127–136.
- [16] Ruijia Cheng, Alison Smith-Renner, Ke Zhang, Joel R. Tetreault, and Alejandro Jaimes. 2022. Mapping the design space of human-ai interaction in text summarization. *arXiv preprint arXiv:2206.14863* (2022). <https://doi.org/10.18653/v1/2022-naacl-main.33>
- [17] Erin Cherry and Celine Latulipe. 2014. Quantifying the creativity support of digital tools through the creativity support index. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 4 (2014), 1–25. <https://doi.org/10.1145/2617588>
- [18] Jarry HT Claessen and Jarke J Van Wijk. 2011. Flexible linked axes for multivariate data visualization. *IEEE Transactions on Visualization and Computer Graphics* 17, 12 (2011), 2310–2316. <https://doi.org/10.1109/TVCG.2011.201>
- [19] Nigel Cross. 2004. Expertise in design: an overview. *Design studies* 25, 5 (2004), 427–441. <https://doi.org/10.1016/j.destud.2004.06.002>
- [20] Hai Dang, Karim Benharrak, Florian Lehmann, and Daniel Buschek. 2022. Beyond text generation: Supporting writers with continuous automatic text summaries. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–13. <https://doi.org/10.48550/arXiv.2208.09323>
- [21] Edward De Bono. 1970. Lateral thinking. *New York* (1970), 70. https://www.kioulanis.gr/rivips/images/Lateral_thinking.pdf
- [22] Zijian Ding and Joel Chan. 2023. Mapping the Design Space of Interactions in Human-AI Text Co-creation Tasks. *arXiv e-prints* (2023), arXiv–2303. <https://doi.org/10.48550/arXiv.2303.06430>
- [23] Graham Dove, Nicolai Brodersen Hansen, and Kim Halskov. 2016. An argument for design space reflection. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. 1–10. <https://doi.org/10.1145/2971485.2971528>
- [24] Steven P. Dow, Alana Glassco, Jonathan Kass, Melissa Schwarz, Daniel I. Schwartz, and Scott R. Klemmer. 2010. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Transactions on Computer-Human Interaction (TOCHI)* 17, 4 (2010), 1–24. <https://doi.org/10.1145/1879831.1879836>
- [25] Alex Endert, Patrick Fiaux, and Chris North. 2012. Semantic interaction for sensemaking: inferring analytical reasoning for model steering. *IEEE Transactions on Visualization and Computer Graphics* 18, 12 (2012), 2879–2888. <https://doi.org/10.1109/TVCG.2012.260>
- [26] Yingchaojie Feng, Xingbo Wang, Kam Kwai Wong, Sijia Wang, Yuhong Lu, Minfeng Zhu, Baicheng Wang, and Wei Chen. 2023. PromptMagician: Interactive Prompt Engineering for Text-to-Image Creation. *arXiv preprint arXiv:2307.09036* (2023). <https://doi.org/10.48550/arXiv.2307.09036>
- [27] Katy Ilonka Gero, Vivian Liu, and Lydia Chilton. 2022. Sparks: Inspiration for science writing using language models. In *Designing interactive systems conference*. 1002–1019. <https://doi.org/10.1145/3532106.3533533>
- [28] Frederic Gmeiner, Humphrey Yang, Lining Yao, Kenneth Holstein, and Nikolas Martelaro. 2023. Exploring Challenges and Opportunities to Support Designers in

- Learning to Co-create with AI-based Manufacturing Design Tools. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–20. <https://doi.org/10.1145/3544548.3580999>
- [29] Michael Golembewski and Mark Selby. 2010. Ideation decks: a card-based design ideation tool. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*. 89–92. <https://doi.org/10.1145/1858171.1858189>
- [30] Joy Paul Guilford. 1961. Three faces of intellect. (1961). <https://doi.org/10.1037/h0046827>
- [31] Joy Paul Guilford. 1967. The nature of human intelligence. (1967). <https://doi.org/10.1017/9781316817049>
- [32] Kim Halskov and Caroline Lundqvist. 2021. Filtering and informing the design space: Towards design-space thinking. *ACM Transactions on Computer-Human Interaction (TOCHI)* 28, 1 (2021), 1–28. <https://doi.org/10.1145/3434462>
- [33] Devamardeep Hayatpur, Daniel Wigdor, and Haijun Xia. 2023. CrossCode: Multi-level Visualization of Program Execution. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–13. <https://doi.org/10.1145/3544548.3581390>
- [34] Chris Heape. 2007. The Design Space: the design process as the construction, exploration and expansion of a conceptual space. (2007). <https://www.semanticscholar.org/paper/The-Design-Space%3A-the-design-process-as-the-and-of-Heape/40b912badea3b575a8f4bde95df4f83a4427ab78>
- [35] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 159–166. <https://doi.org/10.1145/302979.303030>
- [36] Irving Lester Janis. 1982. Groupthink: Psychological studies of policy decisions and fiascoes. (1982). [https://www.scrip.org/\(S\(351jmbntvnsjt1aadkposzje\)\)/reference/ReferencesPapers.aspx?ReferenceID=2122583](https://www.scrip.org/(S(351jmbntvnsjt1aadkposzje))/reference/ReferencesPapers.aspx?ReferenceID=2122583)
- [37] David G Jansson and Steven M Smith. 1991. Design fixation. *Design studies* 12, 1 (1991), 3–11. <https://doi.org/10.1017/S0890060414000043>
- [38] Ellen Jiang, Kristen Olson, Edwin Toh, Alejandra Molina, Aaron Donsbach, Michael Terry, and Carrie J Cai. 2022. Promptmaker: Prompt-based prototyping with large language models. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–8. <https://doi.org/10.1145/3491101.3503564>
- [39] Martin Jonsson and Jakob Tholander. 2022. Cracking the code: Co-coding with AI in creative programming education. In *Proceedings of the 14th Conference on Creativity and Cognition*. 5–14. <https://doi.org/10.1145/3527927.3532801>
- [40] Jeongyeon Kim, Sangho Suh, Lydia B Chilton, and Haijun Xia. 2023. Metaphorian: Leveraging Large Language Models to Support Extended Metaphor Creation for Science Writing. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 115–135. <https://doi.org/10.1145/3563657.3595996>
- [41] John Kirwan. 2017. It's good to have lots of bad ideas. *Nature* 548, 7668 (2017), 491–491. <https://doi.org/10.1038/nj7668-491a>
- [42] Tomas Lawton, Francisco J Ibarrola, Dan Ventura, and Kazjon Grace. 2023. Drawing with Reframer: Emergence and Control in Co-Creative AI. In *Proceedings of the 28th International Conference on Intelligent User Interfaces*. 264–277. <https://doi.org/10.1145/3581641.3584095>
- [43] Joseph CR Licklider. 1960. Man-computer symbiosis. *IRE transactions on human factors in electronics* 1 (1960), 4–11. <https://doi.org/10.1109/THFE2.1960.4503259>
- [44] 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. <https://doi.org/10.1145/3526113.3545621>
- [45] James Derek Lomas, Mihovil Karac, and Mathieu Gielen. 2021. Design space cards: using a card deck to navigate the design space of interactive play. *Proceedings of the ACM on Human-Computer Interaction* 5, CHI PLAY (2021), 1–21. <https://doi.org/10.1145/3474654>
- [46] 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. <https://doi.org/10.1145/3313831.3376739>
- [47] Yuwen Lu, Chengzhi Zhang, Iris Zhang, and Toby Jia-Jun Li. 2022. Bridging the Gap between UX Practitioners' work practices and AI-enabled design support tools. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–7. <https://doi.org/10.1145/3491101.3519809>
- [48] Alan MacEachren, D Xiping, Frank Hardisty, Diansheng Guo, and Gene Lengierich. 2003. Exploring high-D spaces with multiform matrices and small multiples. In *IEEE Symposium on Information Visualization 2003 (IEEE Cat. No. 03TH8714)*. IEEE, 31–38. <https://doi.org/10.1109/INFVIS.2003.1249006>
- [49] Allan MacLean, Richard Young, Victoria Bellotti, and Thomas Moran. 1991. Design space analysis: Bridging from theory to practice via design rationale. *Proceedings of Esprit* (1991). [https://doi.org/10.1016/0142-694X\(94\)90026-4](https://doi.org/10.1016/0142-694X(94)90026-4)
- [50] Stephen MacNeil, Johanna Okerlund, and Celine Latulipe. 2017. Dimensional reasoning and research design spaces. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 367–379. <https://doi.org/10.1145/3059454.3059472>
- [51] Gary Marchionini. 2006. Exploratory search: from finding to understanding. *Commun. ACM* 49, 4 (2006), 41–46. <https://doi.org/10.1145/1121949.1121979>
- [52] Simone Mora, Francesco Gianni, and Monica Divitini. 2017. Tiles: a card-based ideation toolkit for the internet of things. In *Proceedings of the 2017 conference on designing interactive systems*. 587–598. <https://doi.org/10.1145/3064663.3064699>
- [53] Jeyakumar Muthukumarasamy and John T Stasko. 1995. Visualizing program executions on large data sets using semantic zooming. <https://doi.org/10.5555/832277.834333>
- [54] Alex F Osborn. 1953. Applied imagination. (1953). <https://archive.org/details/appliedimagination00osborich>
- [55] Donald A Schön. 1992. Designing as reflective conversation with the materials of a design situation. *Knowledge-based systems* 5, 1 (1992), 3–14. [https://doi.org/10.1016/0950-7051\(92\)90020-G](https://doi.org/10.1016/0950-7051(92)90020-G)
- [56] MC Schraefel, Max Wilson, Alistair Russell, and Daniel A Smith. 2006. mSpace: improving information access to multimedia domains with multimodal exploratory search. *Commun. ACM* 49, 4 (2006), 47–49. <https://doi.org/10.1145/1121949.1121980>
- [57] Mary Shaw. 2011. The role of design spaces. *IEEE software* 29, 1 (2011), 46–50. <https://doi.org/10.1109/MS.2011.121>
- [58] Sangho Suh, Sydney Lamorea, Edith Law, and Leah Zhang-Kennedy. 2022. PrivacyToon: Concept-driven Storytelling with Creativity Support for Privacy Concepts. In *Designing Interactive Systems Conference*. 41–57. <https://doi.org/10.1145/3532106.3533557>
- [59] Sangho Suh, Martinet Lee, Gracie Xia, et al. 2020. Coding strip: A pedagogical tool for teaching and learning programming concepts through comics. In *2020 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 1–10. <https://doi.org/10.1109/VL/HCC50065.2020.9127262>
- [60] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *The 36th Annual ACM Symposium on User Interface Software and Technology* (San Francisco, CA, USA) (UIST '23). Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3586183.3606756>
- [61] Sangho Suh, Jian Zhao, and Edith Law. 2022. Codetoon: Story ideation, auto comic generation, and structure mapping for code-driven storytelling. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*. 1–16.
- [62] Robert Villa, Iván Cantador, Hideo Joho, and Joemon M Jose. 2009. An aspectual interface for supporting complex search tasks. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*. 379–386. <https://doi.org/10.1145/1571941.1572007>
- [63] Sitong Wang, Samia Menon, Tao Long, Keren Henderson, Dingzeyu Li, Kevin Crowston, Mark Hansen, Jeffrey V Nickerson, and Lydia B Chilton. 2023. Reel-Framer: Co-creating News Reels on Social Media with Generative AI. *arXiv preprint arXiv:2304.09653* (2023). <https://doi.org/10.48550/arXiv.2304.09653>
- [64] Bo Westerlund. 2005. Design space conceptual tool—grasping the design process. *Nordes* 1 (2005). <https://doi.org/10.21606/nordes.2005.048>
- [65] Mikael Wiberg and Erik Stolterman. 2014. What makes a prototype novel? A knowledge contribution concern for interaction design research. In *Proceedings of the 8th Nordic conference on human-computer interaction: fun, fast, foundational*. 531–540. <https://doi.org/10.1145/2639189.2639487>
- [66] Robin H Willemsen, Isabelle C de Vink, Evelyn H Kroesbergen, and Ard W Lazonder. 2023. The role of creative thinking in children's scientific reasoning. *Thinking Skills and Creativity* 49 (2023), 101375. <https://doi.org/10.1016/j.tsc.2023.101375>
- [67] Zhuohao Wu, Danwen Ji, Kaiwen Yu, Xianxu Zeng, Dingming Wu, and Mohammad Shidujaman. 2021. AI creativity and the human-AI co-creation model. In *Human-Computer Interaction. Theory, Methods and Tools: Thematic Area, HCI 2021, Held as Part of the 23rd HCI International Conference, HCII 2021, Virtual Event, July 24–29, 2021, Proceedings, Part I* 23. Springer, 171–190. https://doi.org/10.1007/978-3-030-78462-1_13
- [68] Daijin Yang, Yanyang Zhou, Zhiyuan Zhang, Toby Jia-Jun Li, and Ray LC. 2022. AI as an Active Writer: Interaction strategies with generated text in human-AI collaborative fiction writing. In *Joint Proceedings of the ACM IUI Workshops*. <https://hai-gen.github.io/2022/papers/paper-HAIGEN-YangDaijin.pdf>
- [69] Ann Yuan, Andy Coenen, Emily Reif, and Daphne Ippolito. 2022. Wordcraft: story writing with large language models. In *27th International Conference on Intelligent User Interfaces*. 841–852. <https://doi.org/10.1145/3490099.3511105>
- [70] JD Zamfirescu-Pereira, Richmond Y Wong, Björn Hartmann, and Qian Yang. 2023. Why Johnny can't prompt: how non-AI experts try (and fail) to design LLM prompts. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–21. <https://doi.org/10.1145/3544548.3581388>
- [71] Zheng Zhang, Jie Gao, Ranjodh Singh Dhaliwal, and Toby Jia-Jun Li. 2023. VISAR: A Human-AI Argumentative Writing Assistant with Visual Programming and Rapid Draft Prototyping. *arXiv preprint arXiv:2304.07810* (2023). <https://doi.org/10.48550/arXiv.2304.07810>

A APPENDIX

A.1 Examples of Study Participant's Creative Writing: Copywriting, Story, Email, and Lyrics

* Sentences or phrases with underline indicate AI-generated content.

** Different fonts are used to specify each prompt study participants used and its corresponding responses that were generated and used in their final text.

Copywriting Task (P1)

P1's Prompts:

1. Peruvian cuisine as a cultural and social activity
2. Peruvian cuisine
3. PROOFREAD THE TEXT AND ADD IDEAS OF SLANGS IN SPANISH
4. Which Peruvian dishes have an Asian origin

How Peruvian cuisine become a social space for foreigners

Peruvian cuisine has become a social space for foreigners due to its unique fusion of flavors from around the world. With special kind of ingredients as aji amarillo, tuna, noni and cuy; the dishes in the Peruvian gastronomy has increased the level of income for many Peruvians in the country.

Peruvian cuisine has a unique fusion of flavors from around the world, especially North America. An African-American chef could create a gourmet dish that incorporates the bitter flavors of Peru, while still staying true to its authenticity. A great example is streamer fish with bitter yucca and salsa criolla.

BESIDES, DISHES AS CHAUFA HAVE BECOME ONE OF THE MOST IMPORTANT PLATES IN THE COUNTRY. THIS IS A FRIED RICE DISH THAT COMBINES THE FLAVORS OF CHINESE FRIED RICE WITH LATIN AMERICAN INGREDIENTS LIKE PAPERS, EGGS, AND CILANTRO.

One of the main point for any chef in Peru is to strive for accessibility by making the dish affordable and accessible to all dinners, but focus on preserving its original flavor profile without introducing too much innovation. This would be an interesting way to explore Peruvian cuisine in an exciting new way!

BY INCORPORATING MODERNIST CUISINE TECHNIQUES INTO THE PREPARATION PROCESS WHILE STILL PRESERVING THE ORIGINAL FLAVOR PROFILE, a talented chef can create an accessible yet innovative fine-dining experience for all diners all over the world.

Story Writing Task (P5)

P5's Prompts:

1. Write me a story about a man from the future specifically from the year 2079. Make it fun, thrilling, and adventurous.
2. Adding 100 more words to make the story more interesting

A Man From The Future

A time traveller finds himself in a futuristic city, filled with strange architecture and mysterious technology. He is overwhelmed by everything around him and sets off to explore the unknown.

Along the way, he must battle his inner doubts and fears in order to reach his goals. His journey is full of thrills and danger, making it an exciting adventure. The time traveller will have to be charming and brave to find out what lies ahead in the world he has landed in.

He stumbled across a robotic shop that sold all kinds of gadgets and gizmos. The time traveller was amazed at all the futuristic items, some of which he had never seen before. He also found a curious-looking vending machine that sold snacks from around the world. After stocking up on snacks for later, he continued his exploration.

He soon encountered a large tower with an opening at its peak, leading to a platform where people were gathering in anticipation for something unknown. As it turned out, the tower was actually an amusement park ride! The time traveller joined in the excitement as he rode through twists and turns, feeling like an astronaut in space!

Email Writing Task (P7)

* Names that appear in the following are pseudonyms.

P7's Prompts:

1. Can you help me create an email for a professional translator who works on a project that I am coordinating to inform him that he will not be part of the project anymore. I need the email to be very kind and don't hurt his self-esteem
2. I need it to be in first person. I am informing him (Tom is his name) that he will not be receiving more work for this project but he will still be considered for further projects

Email for a Translator

Hi Tom,

I hope you are doing well. This email is to let you know that we won't be able to provide any further work for this project. However, we'll still be considering you for future projects. Thank you for your patience and cooperation; it's really appreciated. Please don't hesitate to reach out if you have any other questions or concerns.

Best regards,
Amy

Lyrics Writing Task (P8)*P8's Prompts:*

1. Write me dark lyrics with themes of death, money, politics, industrial corporation, and ambition
2. Adding a new character to the story
3. ADD A NEW VERSE AND CHORUS THAT EVOLVES THE STORY WHERE THIS DARK WORLD AND SIR WILLIAM HAVE A FINAL CONFLICT WHERE THE WORLD CHANGES BUT SO DOES SIR WILLIAM. THE WORLD BECOMES MORE ALTRUISTIC BUT NOT TOO MUCH, BUT SIR WILLIAM LOSES HIS FAITH IN THE GOOD IN THE WORLD

Ambition**Verse 1:**

The city skyline looms so tall,
Filled with ambition, money, and politics
for all.
A corporate greed that drives us all,
With power-hungry beasts controlling the
hall.

Chorus:

Ambition steals our souls away,
As we dance to the corporations' play.
We march to a somber tune,
Doomed forever in this urban gloom.

Verse 2:

Money controls our fate and dreams,
No one escapes the corporate schemes.
They own our lives and hopes of glory,
Selling false dreams of a different story.

Verse 3:

In a future, bereft of hope and light,
A global corporate reign holds us tight.
Money rules the day and ambition leads,
As politicians serve their corporate needs.
The world is dark and bleak, yet I still
strive,
For power and wealth that's sure to arrive.
My ambition guides me in this cruel age,
To turn to death my every wage.

Verse 4:

The factories loom large with smoke on
high,
A reminder of the lives that will die.
The rich keep getting richer as I go lower;
My life is nothing but a crushing cinder-
shower.
But still, I strive for more than just
money;
I want the power - it's my only honey!
So I'll take all that death can give me -
Ambition fueling this perverse economy.

Verse 5:

A dark land where ambition breeds,
Where money rules all that we need.
Politics divide us in two,
Industrial corporations take over what's
true.

Chorus:

Death lingers in this land so cold,
Money is the ruler we behold.
Politics
unleash a never-ending strife,
Ambition is what keeps us alive.

Verse 6:

The people are divided and torn apart,
Industrial corporations devour our hearts.
Money will always be our guide,
As death lingers in the air outside.

Verse 7:

Sir William, the old and stubborn one,
In a dark world spoke out what no one had done.
People thought him crazy, and avoided him like the plague,
For speaking truthfully regardless of rage.
Though he gained an uneasy reputation for trustworthiness
and influence divine,
His honesty was too much for his own good at that time.

VERSE 8:

SIR WILLIAM'S HONESTY WAS HIS CURSE,
THOUGH HIS COURAGE WAS NEVER WORSE.
HE SPOKE THE TRUTH NO MATTER THE COST,
AND FACED THE DARKNESS WITHOUT GETTING LOST.
HIS WORDS FLEW FREE FROM HIS LIPS,
AND CAUSED A STIR AMONG THOSE WHO DID NOT DIP.
THE PEOPLE THOUGHT HIM CRAZY AND MAD,
BUT AS TIME PASSED, TRUTH BECAME MORE THAN A FAD.

VERSE 9:

THOUGH HE MARCHED INTO THE FRAY,
SIR WILLIAM'S WILL BEGAN TO SWAY.
HE FOUGHT FOR TRUTH AND JUSTICE STRONGLY,
BUT DEEP WITHIN HIS HEART WAS WRONGED.
AND THOUGH SIR WILLIAM WAS FOREVER CHANGED,
THE WORLD FOUND TRUTH THROUGH HIS EXCHANGE.

Table 2: Time taken to generate different numbers of dimensions (Dims). The OpenAI API was invoked for different dimension quantities. The time taken to retrieve all responses was recorded over three trials (each marked as “Trial x”). Failed Call/Total Call is the ratio of failed API calls to the total number of calls. An API call is considered failed if an error occurs in the API call or the response is not in a well-formed valid JSON format (unit: seconds).

# of Dims	Trial 1		Trial 2		Trial 3		Avg Time	Failed / Total Call
	Time	Failed / Total Call	Time	Failed / Total Call	Time	Failed / Total Call		
6	4.91	0/2	4.19	0/2	4.63	0/2	4.58	0/6
10	6.07	0/2	8.13	1/3	5.91	0/2	6.70	1/7
14*	20.25	4/6	44.82	6/8	32.19	3/5	32.42	13/19

* Increasing the number of dimensions leads to token limits being exceeded, resulting in JSON parsing errors, such as unclosed braces.

Table 3: Time taken to generate different numbers of responses (Resps). The OpenAI API was invoked for different response quantities, and the time taken to retrieve all responses was recorded over three trials (each marked as “Trial x”). Failed Call/Total Call is the ratio of failed API calls to the total number of calls. An API call is considered failed if an error occurs in the API call or the response is not in a well-formed valid JSON format. For each response, the API is invoked twice (once for text generation and once for text summarization). Thus, the total calls are equal to twice the number of responses (unit: seconds).

# of Resps	Trial 1		Trial 2		Trial 3		Avg Time	Failed / Total Call
	Time	Failed / Total Call	Time	Failed / Total Call	Time	Failed / Total Call		
20	7.75	0/40	8.01	0/40	6.79	0/40	7.52	0/120
40	7.63	1/81	9.12	0/80	7.94	1/81	8.23	2/242
60	9.01	0/120	9.51	1/121	8.05	0/120	8.88	1/361
80*	8.72	9/169	9.42	9/169	11.11	9/169	9.75	27/507

* Appear “too many request” error when calling OpenAI API

Table 4: A list of constants defined to express the prompts used in Luminare and shown in Table 5.

Prompt Constant	Value
dimensionDef	A dimension will contain categorical dimension values (attributes) that are qualitative and subjective to the user. This means there is no right answer for selecting a dimension value. The user should be able to select any dimension value depending on their preference. The dimensions must not be an evaluation of how good the writing is. All responses are assumed to be the best writing generated by you.
dimensionConclusion	Even though I encourage you to use some of these examples if best fitting, I highly recommend that I also get unique and orthogonal dimensions.
nominalDimensionDef	dimensionDef + A nominal dimension will contain dimension values that do not have a particular order and are up to the user’s selection. Some nominal dimensions that I would want are Tone, Setting, Style, or Perspective. I do NOT want Length, Grammar, Quality, or Clarity. + dimensionConclusion
ordinalDimensionDef	dimensionDef + An ordinal dimension will contain dimension values measured in an order (least, less, neutral, more, most). The type of dimensions I want are ones that are of a single key property that the user may want more of or less of depending on their preference. Some ordinal dimensions that I would want are Concreteness, Realism, or Subjectivity. I do NOT want Quality, Creativity, or Length. + dimensionConclusion
wordLimit	Limit the response to 150 words

Table 5: A list of prompts used in Luminate. The underscored text in the ‘Prompt’ column is a placeholder for example input(s). The **bold text in the ‘Prompt’ column is a placeholder for prompt constants from Table 4. A Prefix refers to strings that can be added to the prompt to provide context, but these strings are not directly sent to the API. The **SMALL CAPS** in the ‘Prompt’ column is a placeholder for prompt prefix.**

Prompt Type	Prompt	Example Input(s)	Example Response
PREVIOUS CONTEXT PREFIX	This is the context: <u>background</u> —end context —	background = It’s full of surprises, that can make us smile or frown. But it always has something to teach us when we look around.	N/A
EXISTING DIMEN- SION PREFIX	These are the current existing di- mensions and their values: <u>current- Dimensions</u>	currentDimensions = “Set- ting(Nominal):[Campuse Stadium, Football Locker Room, Victory Parade] Engagingness(Ordinal):[least, less, neutral, more, most]”	N/A
NOMINAL DIMEN- SION GENERATION	nominalDimensionDef + list <u>cat- Num</u> nominal dimensions and asso- ciated <u>valNum</u> possible values on which we can categorize and assess the content for the prompt: <u>prompt</u> #### You MUST answer in the follow- ing JSON object format, wrapped in curly braces. Replace all strings with <...>. There must be <u>catNum</u> items in the JSON object: <pre>{ "<dimension name #1>":["<valNum values for this dimension>"], ..., "<dimension name #catNum>" : [<valNum values for this dimen- sion>] }</pre>	catNum = 5 valNum = 6 prompt = “write a story about a rab- bit”	<pre>{ "Genre": ["Fantasy", "Adventure", "Romance", "Mystery", "Comedy", "Drama"], "Tone": ["Lighthearted", "Humorous", "Moody", "Frightening", "Hopeful", "Sus- penseful"], "Setting": ["Modern Day", "Medieval Times", "Western Era", "Futuristic World", "Mythical Realm", "Urban City"], "Style": ["Narrative Poem", "Dialogue Driven Story", "Traditional Fable Tale", "Nonlinear Prose Piece", "Epic Saga", "Short Story"], "Perspective": ["First-Person POV (Protagonist)", "Third-Person Limited (Protagonist)", "Third-Person Omni- scient (Narrator)", "Second-Person POV (Reader, Audience)", "Multiple Perspec- tives/Voices", "Objective/Impersonal Narrator"] }</pre>

Continued on the next page

Table 5: A list of prompts used in Luminate. The underscored text in the ‘Prompt’ column is a placeholder for example input(s). The **bold text in the ‘Prompt’ column is a placeholder for prompt constants from Table 4. A Prefix refers to strings that can be added to the prompt to provide context, but these strings are not directly sent to the API. The **SMALL CAPS** in the ‘Prompt’ column is a placeholder for prompt prefix.**

Prompt Type	Prompt	Example Input(s)	Example Response
ORDINAL DIMENSION GENERATION	ordinalDimensionDef + list <u>catNum</u> ordinal dimensions on which we can assess the outcome for the prompt: <u>prompt</u> to what extent represents the dimensions #### answer in the following JSON format: <pre>{ "<dimension name>": ["least", "less", "neutral", "more", "most"] }</pre>	catNum = 5 prompt = "write a story about a rabbit"	<pre>{ "Creativity": ["least", "less", "neutral", "more", "most"], "Imagination": ["least", "less", "neutral", "more", "most"], "Grammatical Accuracy": ["least", "less", "neutral", "more", "most"], "Originality": ["least", "less", "neutral", "more", "most"], "Presentation Style": ["least", "less", "neutral", "more", "most"] }</pre>
RESPONSE GENERATION	wordLimit #### PREVIOUS CONTEXT PREFIX #### Prompt: <u>prompt</u> #### Requirements: <u>requirements</u>	Prompt = "write a story about a rabbit" Requirements = "Genre: Comedy Tone: Frightening Setting: Medieval Times Style: Short Story Perspective: Third-Person Omniscient (Narrator) Creativity: more Imagination: most Grammatical Accuracy: less Originality: less Presentation Style: less"	Once upon a time, in a far away kingdom, there lived a brave rabbit named Flopsy. Flopsy was determined to prove himself as the bravest rabbit in the land. One day, he decided to venture into the dark and mysterious forest that surrounded his home. As he ventured deeper into the forest, he encountered all sorts of frightening creatures and obstacles. He eventually came across an old castle with an eerie presence. He cautiously entered and soon found himself face-to-face with a giant dragon! The dragon roared and Flopsy trembled in fear but he stood his ground and bravely faced off against the beast. After a long battle, Flopsy emerged victorious! He had saved his kingdom from certain destruction and was hailed as a hero by all who heard of his courageous deed.

Continued on the next page

Table 5: A list of prompts used in Luminate. The underscored text in the ‘Prompt’ column is a placeholder for example input(s). The **bold text in the ‘Prompt’ column is a placeholder for prompt constants from Table 4. A Prefix refers to strings that can be added to the prompt to provide context, but these strings are not directly sent to the API. The **SMALL CAPS** in the ‘Prompt’ column is a placeholder for prompt prefix.**

Prompt Type	Prompt	Example Input(s)	Example Response
NEW DIMENSION GENERATION	PREVIOUS CONTEXT PREFIX #### prompt: <u>prompt</u> + EXISTING DIMENSION PREFIX + Generate a new dimension that is new and orthogonal to the existing ones above.	prompt = “write a story about a rabbit”	Tone
	Given following text, return key words and a one sentence summary, a structure , and a title of the text. #### Text is: <u>text</u> #### Don’t include any text other than the json; Word limit of the summary text is 20 words; Word limit of the title is 5 words; Maximum 5 key words #### Should be in the following JSON format: { “Key Words”: [“<key word 1>”, “<key word 2>”, ...], “Summary”: “<summary>”, “Structure”: “<part 1>-<part 2>-<part 3>...”, “Title”: “<title>” }	text = “Once upon a time, in a futuristic world, there lived a rabbit. He was an adventurous soul who loved to explore the unknown. One day, he decided to take a journey and see what the world had to offer. As he hopped along his way, he encountered many strange and wonderful things. He met robots that could talk and fly, creatures that could swim in the air, and even plants that glowed in the dark! Despite all these wonders, nothing compared to the joy of discovering new places and meeting new people. The rabbit’s journey was full of laughter and fun as he made his way through this strange yet exciting world. In the end, he returned home with stories of his travels that would be told for generations to come!”	{ “Key Words”: [“Brave”, “Adventure”, “Journey”, “Creatures”, “Love”], “Summary”: “A brave rabbit embarks on a journey to explore the world and finds true love in a magical kingdom.”, “Structure”: “Once upon a time-Journey-Encounter creatures-Finds true love-Becomes ruler of kingdom”, “Title”: “Rabbit’s Journey” }

Table 6: Dimension Generation. Ordinal dimensions and nominal dimensions generated by Luminate given example input(s) across different tasks including story writing, email writing, copywriting and song lyrics.

Task Type	Example Input	Prompt	Example Ordinal Dimensions	Example Nominal Dimensions
Story Writing	“Write a story about <u>time travelling</u> ”	NOMINAL DIMENSION GENERATION and ORDINAL DIMENSION GENERATION in Table 5	Imagination Creativity Suspense Excitement Engagement	Genre Tone Setting Point of View Time Travel Method
Email Writing	“Write an email to all tennis club members about the next <u>tournament</u> ”		Compliance Impact Importance Relevance Urgency	Age Group Format Level Location Type
Copywriting	“Write an article on the most exciting place to visit in <u>Hawaii</u> ”		Level of relaxation Level of fun Nature of activities Location amenities Quality of experience	Activity Location Price Range Size Time Frame
Song Lyrics	“Write a song lyrics about the <u>universe</u> ”		Beauty Grandeur Magnificence Mystery Power	Genre Instrumentation Length of Verse Mood Vocal Range

Table 7: A list of some topics study participants could choose from during the user study.

Writing Task	Topic
Copywriting	E-commerce and Retail
	- Product descriptions for online stores.
	- Promotions and discounts for specific products.
	- Creating a compelling shopping experience.
	Travel and Tourism
	- Promoting travel destinations.
	- Writing travel guides.
	- Highlighting the features of a hotel or resort.
	Health and Wellness
	- Advertising health supplements or products.
	- Writing about the benefits of a healthy lifestyle.
	- Creating content for fitness and wellness programs.
Short Story	Parallel Universe
	Time Travel
	Dystopian Society
	First Contact
	Family Secrets
	A Mysterious Inheritance
	Survival in the Wilderness
	An Unlikely Hero
	Love in Unexpected Places
	A World Without Technology
	The Power of Music
Email / Letter	Personal Letter Topics
	- Expressing Gratitude
	- Celebrations and Milestones
	- Updates on Life
	- Condolences
	- Friendship and Appreciation
	Professional Letter Topics
	- Cover Letters
	- Resignation Letters
	- Recommendation Letters
	- Networking
	Love and Relationships
	- Apology Letters
	Social and Advocacy Letters - Letters to Elected Officials
	- Letters to the Editor
	Educational Letters - College Application Letters
	- Scholarship Application Letters