May AI?

Design Ideation with Cooperative Contextual Bandits

Janin Koch, Andrés Lucero, Lena Hegemann, Antti Oulasvirta

Aalto University, Finland

ABSTRACT

Design ideation is a prime creative activity in design. However, it is challenging to support computationally due to its quickly evolving and exploratory nature. The paper presents cooperative contextual bandits (CCB) as a machine-learning method for interactive ideation support. A CCB can learn to propose domain-relevant contributions and adapt their exploration/exploitation strategy. We developed a CCB for an interactive design ideation tool that 1) suggests inspirational and situationally relevant materials (“may AI?”); 2) explores and exploits inspirational materials *with* the designer; and 3) explains its suggestions to aid reflection. The application case of digital mood board design is presented, wherein visual inspirational materials are collected and curated in collages. In a controlled study, 14 of 16 professional designers preferred the CCB-augmented tool. The CCB approach holds promise for ideation activities wherein adaptive and steerable support is welcome but designers must retain full outcome control.

CCS CONCEPTS

• **Human-centered computing** → **Interactive systems and tools**; • **Computing methodologies** → *Machine learning*; • **Applied computing** → *Arts and humanities*.

KEYWORDS

Interactive machine-learning; ideation support; creativity support tools; mood board design

ACM Reference Format:

Janin Koch, Andrés Lucero, Lena Hegemann, Antti Oulasvirta. 2019. May AI?: Design Ideation with Cooperative Contextual Bandits. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland UK. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3290605.3300863>

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

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-5970-2/19/05...\$15.00

<https://doi.org/10.1145/3290605.3300863>

1 INTRODUCTION

This paper discusses a machine-learning based interactive support for *design ideation*: the process of generating and curating original and useful ideas so as to define and explore what is desirable in a design project [16]. In design ideation, designers move between analysis and synthesis of ideas or concepts to construe a potential future [42]. Abductive reasoning [38] and abstraction are argued to allow designers to “break through to the a-ha! moment of inspiration” [40].

Scholars have suggested that computational support holds particular use potential in searching and collecting of materials [40]. However, advanced creativity support tools are rarely, if ever, deployed in early stage design [33]. The hurdle is to make creative contributions without distracting design thinking [15, 17, 21]. This is a challenge for non-interactive approaches in machine-learning, or for any approach assuming *pre-defined* objectives, which may yield irrelevant proposals [46]. Hence, it is important to study methods that allow designers to work *with* an algorithm rather than *for* it.

Ideation often involves verbal, visual or tangible material, which may be intentionally ambiguous to facilitate abstraction. However, the ability to ‘see’ and reason on it is fundamental to designerly thinking. Hence, visual material is considered to be most suitable to support the construction of new ideas [42]. In this paper, we look at *mood board design* as a representative and challenging area of ideation. A mood board is a visual collage composed of images, text, and objects. Its construction “stimulates the perception and interpretation of more ephemeral phenomena such as color, texture, form, image and status” [14]. They are used in the earlier stages of a design project for visualizing hard-to-express ideas for further inspiration-seeking and decision-making. The ideation process itself is dynamic and iterative in which designers switch between searching and making, going back to find the missing image that fits [30]. Designers engage here in both problem-defining and problem-solving [5]. The final collage can assist in the transmission of a new mindset, story, or vision to stakeholders [27].

Thus far, work on interactive mood board design focused mainly on collaboration and collocation. Lucero [28] identified six stages in mood board-making: *defining*, *collecting*, *browsing*, *connecting*, *building*, and *presenting*. The *Funky Coffee Table* [29] is a tabletop system that supports *browsing* by storing images in virtual layers. The *Funky Wall* [30]

Figure 1: The CCB-based interactive mood board design tool: image suggestion (1), verbal explanation (2), steering controls (3), history panel of previous suggestions (4), image search (5), and editing tools (6).

is an interactive wall display that supports *presenting* with multimodal and multi-stakeholder feedback.

This paper contributes to the complementary problem of how to best help designers *collect and curate* material. Traditionally, designers browse through physical magazines, explore art or colleagues’ work [22]. While search engines and dedicated online services have become prominent sources, they rely on verbalization of ideas via sequential queries, which may counter the visual and abstractive nature of ideation [24]. It further limits “serendipitous encounters”, crucial to the original mood board method [22]. Today’s computers have the capability to perform hundreds of image searches and analyses in parallel, which could provide valuable support. To take full advantage of this power, the system needs to know *what* to search for in terms of color, mood, content, etc., which are subject to changing objectives. This is where AI can help, in steering this search power according to the designer’s evolving constraints and interpretations.

We therefore focus on a central technical problem in this context: *how to identify and provide inspirational materials to a designer in a situationally appropriate manner, and how to support their exploration* (“May AI?”)? We build on a known class of machine-learning methods called *bandit systems*. We apply a variant called *cooperative contextual bandits* (CCBs) [43], with the goal of a “co-creative system”, where the system works more like a partner or assistant [46]. The CCB learns about the problem at hand, searches the space with the designer, and adapts to their style. Our CCB can 1) autonomously transition between exploration and exploitation while 2) taking into account the style and content of an evolving design by being *steerable* using control widgets. It can also support interpretation by asking for the designer’s rationale for his choices, while offering verbal justifications for its own suggestions. Figure 1 shows our mood board design tool.

In the following, we present related work, the tool concept, the method, and results from a controlled study.

2 RELATED WORK

Our approach builds on several ideas presented in previous work on interactive and computational support for brainstorming, dance, music, and visual collages.

Brainstorming has gained considerable attention in HCI and AI research. Systems such as *InspirationWall* [1], *Momentum* [2], and *V8 Storming* [23] are designed to collect, organize, and present ideas during brainstorming. Many systems focus on suggesting related ideas, from crowds, user-trained association models [23], knowledge graphs [1], etc. However, so-called far suggestions too are important [41]: they can help exploring when “stuck”, while near suggestions can aid in exploiting when one is “on a roll” [6]. Bandit systems in general are appropriate for striking a balance between exploration and exploitation. The CCB approach presented here can further adapt its near–far strategy over time.

In applications for the dance and music fields, *Viewpoints AI* is an AI agent projected on a surface that improvises and explores movements jointly with a dancer [20]. It uses rule-based reasoning to react to spatial and time-related factors. *BoB* is an AI agent for supporting jazz improvisation [44]. As a “believable agent,” it learns a generative model from data that can impro-play believably. Both *BoB* and *Viewpoints AI* try to avoid “heavy use of pre-created instantial knowledge and rather focus on procedural expression” [20]. They watch the user improvise, to configure themselves in a musically appropriate manner [44]. Our CCB does not require rule-based architecture, and can be pre-trained with domain-related data. Then, when interacting with a user, it can continuously update its beliefs, which over time will better reflect personal preferences and strategies.

There is increasing interest in co-creative agents in drawing. Oh et al. [34] presented an AI-assisted drawing tool that can give instructions to users and explain its intentions when needed. We aim for an approach that, similarly, is able to lead – if the designer so desires – and has explanations available upon request. Other exploratory drawing agents, such as the *Drawing Apprentice* [10], use a turn-taking approach to draw with an artist, while exploration/exploitation is steered by means of sliders. We also draw from the idea of active participation via instant feedback. Further, to support interpretability, our algorithm can explain how its suggestions are related to the features of the visual collage created.

Considering visual collages in particular, related work has focused on two main subtasks: 1) finding visual materials and 2) laying them out on a canvas/board. Machine-learning methods can assist users in finding specific images – e.g., with user-specified rules [13], preferences, colors, or patterns [12] or via user-specified [8] and dynamic [45] clustering. As Fogarty et al. [12], we use a feature-based method for searching relevant images. A key aspect of our work, however, is the ability to switch between exploitation and exploration strategy. Regarding the laying out of visual materials on a canvas, most scholars have attempted to automate or support collaging [3, 12, 39, 45] by letting users specify abstract areas [12], adapt sizes to their actions [45], or automate it in line with preference models [3] or areas of interest [39]. While most previous papers use pre-defined aesthetics criteria that drive optimization, we assume that these criteria evolve during the process, and we aim to recognize and adapt to them without actively interfering: e.g., through image size and visibility or letting designers create their own spatial representation. With a concept similar to free-form curation [31], we aim to enable “*elements to be spontaneously gathered from the web [...], manipulated, and visually assembled in a continuous space*” to encourage the evolving of ideas and relations among the objects.

3 WALKTHROUGH

The cooperative contextual bandit system¹ (see next section) was integrated into a design tool for mood boards. Figure 1 shows an overview of the tool from a designer’s standpoint. The UI is divided into three main regions: canvas (middle), tool panel (left), and “AI” panel (right). One starts a project by providing a login name and a short description of the general theme of the mood board (e.g., “vegan” or “urban entrepreneurs”) that can be changed later on.

Image Search and Editing. The designer can search for images (Fig. 1: 5) by using *DuckDuckGo* Image Search [18] and drag-and-drop of images to the workspace. In the image search panel, every search produces 25 results, divided over

five pages. There are four regular functions available in this panel: editing background and element color, adding shape primitives such as squares and circles, changing the z-order (front or back), and removing items (Fig. 1: 6).

AI panel. The AI panel displays images suggested by the CCB. The user can ask for more images, using three buttons (Fig. 1: 3): “More like this,” “Not this one,” and “Surprise me,” which impact the CCB’s exploration/exploitation behavior. All unused images can be browsed via the History panel (Fig. 1: 4). This panel follows the metaphor of a physical magazine or image library, where the designer can go back to earlier pages and revisit images that are suitable later on. Our tool also permits text elements and gradient backgrounds, but these were turned off in the main experiment to reduce total time by giving less priority to finer editing of images.

System Perspective

Initialization. The system first loads a general and a personal prior from a *Postgres* database. The personal prior contains every choice the designer made; if there are none, only the general prior is loaded. The general prior is based on sample mood board designs from *Pinterest* (see next section), intended to reflect contemporary design styles. The specified theme (see above) is forwarded to a word-associations API [19], which fetches associated terms the system then stores in an association list, to explore related themes on its own. Every time the designer adds an image from the image search panel, the corresponding query word is added to this list.

Suggestions. Every image added to or removed from the canvas triggers a screenshot of the current mood board, which is submitted for analysis of features (for a list of image features, see the next section). The color values of the mood board are obtained via dynamic clustering [32], and the dominant color features are used to define the *context* of the CCB.

The feature-based notion of context allows the CCB to exploit (similar features) and explore (dissimilar features) different design strategies. It selects a suggestion vector, containing the image features that have the highest probability of a good fit. Given this vector and the query words in the association list, either a new image is retrieved from a local database or a new online image query is made in real time (using *DuckDuckGo*). In this case, we query the verbalized feature vector in combination with each word in the association list, one after another, until a suitable image is found. For each query, we analyze the first 40 images. To exclude explicit images, we apply face (“Haar” cascades [36]) and text detection (EAST detector [35]), using *OpenCV 4*. The remaining images are dynamically clustered for dominant feature retrieval and added to the image database, with metadata.

The CCB updates its beliefs when the designer 1) selects or rejects a suggested image, 2) deletes an image from the

¹Code available at <https://userinterfaces.aalto.fi/ccb>

mood board (not suitable anymore), or 3) retrieves one from the suggestions history. This reflects the idea that an image can be a good fit in one context but unnecessary in another.

Steering. Designers can express one out of three preferences (Fig. 1: 3): “More like this” to favor exploitation; “Surprise me” for exploration; and “Not this one” for default suggestions. These do not overrule the CCB’s learned behavior.

Explanations. The difference between the suggested image features and the mood board is used to create a verbal justification, displayed above it (Fig. 1: 2). Also, the system can ask for justifications: when the user adds an image from the search engine whose features are significantly different from the current mood board’s, (s)he is asked whether the image was selected for “content,” “harmony,” or “contrast.” For example, if the designer selects “content,” the current association list is replaced with a new one, based on the current search term, to enable a shift of focus within the ideation process. After that, all terms searched that lead to selected images will be added to the association list again.

4 COOPERATIVE CONTEXTUAL BANDITS

The term “bandit” originates from from so-called one-armed bandit machines in casinos. When a ‘lever’ (arm) is pulled, it triggers a symbol combination, some of which provide a payout. The question is whether one should pull a ‘lever’. *Multi-armed bandits* are a generalization to several levers: limited to pulling one lever at a time, the problem is to estimate which lever produces the highest payout (reward). Bandit systems are commonly employed in such applications as marketing and recommendation engines [9, 26]. A standard multi-armed bandit solution is insufficient for our purposes, since it lacks the ability to accommodate the variations in designs, design strategies, and user-specific objectives.

Contextual Bandits

Contextual bandits extend multi-armed bandit algorithms by considering the context of use or users. A contextual bandit observes a context vector x_a of each arm $a \in A$. Working from actions observed in previous trials, it selects arm $a_t \in A$ and receives reward r_{a_t} from the user, whose expectation depends on arm a_t . The algorithm then improves its arm-selection strategy with the new observation, (x_{a_t}, a_t, r_{a_t}) .

The goal of any contextual-bandit algorithm is to maximize the expected total reward [4]. The usual process starts with the untrained algorithm, updating the probability distribution of relevant arms in every trial t . It should be noted that this type of algorithm does not need a pre-specified definition of the goal or optimal pre-learned values. However, it requires learning for each possible suggestion to find an overall successful decision. In our case, that would require more than 13,000 agents to be trained (see below).

Cooperative Contextual Bandits

We adapted an online learning cooperative contextual bandit algorithm presented by Tekin et al. [43] to address our objectives to 1) identify and propose material that is novel and relevant for a designer yet also 2) adapt to the designer’s changing strategy of diversification and intensification. Further, it should 3) support reflection on decisions.

The CCB coarsely partitions the context space and assigns the partitions to agents, called *strategy agents*, which unlike contextual bandits can cooperate with each other. As in [43], each strategy agent can refer a suggestion to its immediate neighboring agents in each dimension (Fig. 2: a). This allows exploring alternative strategies without throwing in overly eccentric ideas. Each strategy agent is then partitioned into multiple subagents, called *suggestion agents* – a contextual bandit’s arms. Every strategy agent A_i has a probability function for relevance of each of its own suggestion subagents a_{in} , and for each of its neighboring strategy agents A_j , independently of A_j ’s probabilities for its own subagents a_{jn} . These probabilities are updated with every iteration.

This cooperation allows the algorithm to diverge from and exploit current strategies that remain controllable by the user and the system even in very large context spaces, unlike contextual bandits. The partitioning is crucial in our task since it allows us to abstract the huge context space, representing all possible mood boards, to a few partitions that roughly represent design strategies that are visually understandable by humans. Tekin et al. [43] describe two slicing approaches: a uniform one, where all dimension are sliced into equal parts, and an adaptive one, where the number of slices increases progressively in regions of the contextual space with higher densities. The latter lets one learn more details about frequent design strategies but comes with the risk of slicing these regions too finely. The resulting, fine-granularity slices can end up being hard to distinguish, and therefore to control, by a user. Furthermore, our approach for exploration relies on referring to neighboring slices; overly fine slicing would limit the explorative power of our algorithm, because neighboring strategies would remain very similar to each other. Therefore, we applied a uniform slicing approach.

Overview of the Algorithm

We first slice the potential mood board space into partitions handled by strategy agents, each responsible for recommendations by its suggestion agents. In every discrete trial:

- (1) a mood board is transformed into a five-dimensional vector in the context space and is assigned to the strategy agent of the corresponding partition;
- (2) the agent queries its own suggestion agents for similar suggestions (exploitation) and nearby strategy agents for alternative moods (exploration);

- (3) each suggestion agent within the current strategy, and each nearby strategy agent, provides probabilities for making a good suggestion (Fig. 2: b);
- (4) the agent with the highest probability is selected with respect to an exploration/exploitation criterion, c ;
- (5) if a suggestion agent is selected (Fig. 2: c), it describes the next image suggestion feature vector; otherwise (Fig. 2: d), the corresponding strategy agent queries its own suggestion agents to identify this vector;
- (6) this vector, in combination with the association list, is used to query a suitable image in the local database; if not successful, it will be translated into human-readable features to query images online in real time;
- (7) the user accepts or rejects the suggested image; and
- (8) that feedback is used to update the probability distributions of the corresponding suggestion agents and, in case of referral, of the neighboring strategy agent.

Below, we will go through the details of the CCB and examine how this structure can be used to justify suggestions.

Context Partitioning

The algorithm considers the context space as a 5-dimensional vector describing the dominant values of the mood board. Each vector consists of the dominant *color value* (C), *saturation* (S), *color lightness* (L), *image orientation* (O), and *color distance* (D). The space of all possible mood boards (MB) can be described as: $MB = \langle C, S, L, O, D \rangle$. Selection of these features was based on perceivable differences in mood board designs as defined by two authors working in the field.

We applied a uniform slicing to each dimension of the context space, dividing the space in 96 partitions according to (roughly) human-perceivable increments. We divided the color space (C : 360° of Hue) into six fundamental colors, i.e. slices of 60° for strategy agents, and in slices of 5° for suggestion subagents. Saturation (S : $[0,1]$ based on the HSL color space) is sliced into Low $[0, 0.5[$ and High $[0.5, 1]$ for strategy agents and into slices of 0.25 for subagents. Lightness (L : $[0,1]$ based on the HSL color space) is sliced similarly: Dark $[0, 0.5[$ and Light $[0.5, 1]$ for strategy agents and in slices of 0.25 for subagents. Context orientation (O : $\{horizontal, vertical\}$) is the most prevalent orientation of the images in the mood board, with no subdivision for subagents. Color distance (D : $[0, 180]$) is the hue difference between the two most dominant colors in the mood board. We divide it into three slices: Similar $[0, 60[$, Neutral $[60, 120[$, and Colorful $[120, 180]$.

Agents

Each partition in the context space is represented by a strategy agent. Given the current context vector allocated to partition n , strategy agent A_n is assigned for recommending the next suggestion to the user (Fig. 2: b). A_n consists

Figure 2: The mood board is best described by strategy A_i . (a) Simplified 2D context space with possible strategies. (b) A_i selects the best relevance probability, either (c) one of its suggestion agents a_{in} or (d) one of its neighboring strategies $A_{e..m}$, which queries its own suggestion agents. Selected distributions are individually updated based on feedback.

of suggestion subagents $\{a_{n1} \dots a_{nm}\}$ representing uniform sub-slices of n . A strategy agent updates probability distributions describing the relevance of each of its suggestion subagents, as well as of its neighboring strategy agents.

Decision-Making

For every observed context, the corresponding strategy agent A_i has to decide whether to refer the task of selecting a suitable image feature vector to its own suggestion agents $a_{in} \in A_i$ at a cost $c_{a_{in}}^i$ (with some abuse of notation) or refer the task to another strategy agent A_j at a cost of c_j^i (Fig. 2: b). Strategy agent A_i can evaluate the expected probability for only its neighboring strategy agents A_j (Fig. 2: a) and has no access to their suggestion subagents. Each probability is based on a standard Thompson sampling approach with a beta prior on the binomial distribution learned.

Exploitation. If its own suggestion agent $a_{in} \in A_i$ provides the highest probability for a suitable image feature vector (Fig. 2: c), the CCB sticks to the current strategy.

Exploration. If a neighboring strategy agent A_j has the highest probability of suggesting a good image (Fig. 2: d), A_i refers the task to A_j , which selects one of its selection agents $a_{jn} \in A_j$ that yields the highest probability.

User Action

For each suggested feature vector a_{nm} we observe a binary reward r_{nm} : whether the designer accepts or rejects a suggestion. This feedback updates the learned probability distribution of the selected suggestion agent accordingly. In case a neighboring strategy agent is referred (exploration), the feedback will influence the learned relation of A_i to A_j and also the learned relation of A_j to a_{jn} . The user can steer the suggestions with three buttons in the AI panel (Fig. 1: 3) that affect exploration cost c_m^n : “More like this” gives it a positive value, “Surprise me” applies a negative value, and “Not this one” resets it to 0. The cost is added to the probability value provided by neighboring strategy agents.

Explanation

Making Justification. Verbal justifications are built by selecting one feature of the suggested feature combination, in order to keep the justification simple. To make the justification relevant and easy to understand, we select a feature that is easy to see in both the image and the mood board. Once a feature is selected, its numeric values are translated into text such as color names or descriptions of luminance, saturation, and contrast. If no feature is meaningful in relation between image and mood board, the system explains itself via its associations, using the word from the image query.

Requesting Justification. Depending on the features of newly added images and context, the system may prompt the designer to indicate whether the image was added for “content,” “harmony,” or “contrast.” The prompt is triggered when the image came from the image search and when the saturation, luminance, or color contrast of the image differs by a certain threshold from the context. A large difference in color between the context and image triggers a prompt only if the colors in the context are otherwise homogeneous. The prompt is presented in the mood board as a small window with an arrow pointing to the new image.

To respond, the user can click one of the buttons: “Content,” “Harmony,” or “Contrast.” “Harmony” increases the cost c_m^n for selecting a different strategy agent, which favors exploitation. “Contrast” reduces this exploration cost c_m^n , which favors exploration. When “Content” is chosen, the current association list is replaced with new word associations obtained from the current search term. If the designer clicks in the background or simply continues to work, adding further images to the mood board, the prompt disappears.

Adapting to Changing Criteria: Simulation Data

To assess how quickly CCBs *adapt* to changing design criteria, we created a synthetic task. A CCB presents a (simulated) designer with one suggestion at a time, and the designer responds to either “include” or “exclude,” using criteria unknown to the CCB. For example, the designer may start favoring similar colors, then switch after approx. 60 selections to favoring contrasting colors. To make the task more realistic, we added noise to the designer’s choices (10% random choices). We analyzed regret (optimal expected reward

Figure 3: CCB adapts to a change of design criterion (in red).

minus total reward per trial) over time. On average, when one-dimensional criteria were considered (here color), it took around 100 guesses to recover a previously unseen intention. However, given prior exposure to that criterion, much less time was needed, around 20 guesses (Fig. 3). We therefore carried out training with a large dataset of real mood boards.

Constructing the Prior

Typically, only user feedback is used to train a contextual bandit system during interaction. In our case, the number of interactions per user is limited. To help with initial suggestions, and to enable domain-relevant suggestions, we constructed a general prior used for every participant. To get a wide range of examples, we collected 1,024 mood boards from online sources, reflecting numerous uses. The images of the sample mood boards were retrieved via OpenCV’s shape descriptors [37]. For each mood board, the images retrieved were then ordered randomly to simulate their successive addition. That was used to build a prior for the probability distributions of the suggestion and strategy agents.

In a contrast against general contextual bandits, with CCBs each strategy agent only has to know the general success of referring to its neighbors’ suggestions (i.e., one distribution per neighbor), rather than each individual suggestion agent of each neighbor. The probability distribution from A_i to A_j is updated every time the mood board best described by A_i successfully receives an image from A_j , irrespective of which suggestion agent (a_{jn}) was responsible. That approach reduces the training required for very large context spaces.

5 EVALUATION

Our evaluation methodology follows established practices in empirical research on creativity. In particular, we aimed for 1) a representative sample of end-users, who in our case are professional designers; 2) a mixed-methods approach that is able to gauge both the process and the outcome of ideation [46], including designers’ subjective views; 3) realism in design briefs over a larger number of observations per participant [11]; 4) comparison of AI² (*with-AI*) against a baseline with the same functionality (*without-AI*), which allows us to learn about the *effects* of AI without confounding them with the design tool itself; and 5) use of standardized measurements that support both user experience (i.e., AttrakDiff [25]) and perceived creativity (i.e., Creativity Support Index [7]). To obtain balanced feedback from designers and to avoid order effects, we followed a within-subjects design with counter-balancing. The without-AI condition was tested with the design tool shown in Figure 1, excluding the AI suggestion panel (Fig. 1: 4).

²The CCB was introduced to the participants as an “AI method,” so we use the term “AI” from here on when describing their viewpoint.

Participants

We recruited 16 professional designers (12 F, 4 M), with a mean age of 34 years and an average of five years' experience. Their expertise covered architecture, interaction, textile, fashion, and graphic design. Most were enrolled in a PhD program at a local university. See Table 1 for an overview.

Each experiment took about one hour and was audio-, screen-, and video-recorded. All volunteered under informed consent and agreed to recording and anonymized publication of results. European privacy law (GDPR) was followed throughout. They were compensated with a cinema voucher.

Tasks and Materials

We created two realistic briefs for the task of proposing a new visual identity for a sub-brand of a known company: (B1) a bank and (B2) a grocery store. The briefs were expressed in the form of a one-page client description with background and goals. The briefs are included in *Supplementary Materials*.

Procedure

Firstly, the designer was shown a video of the basic functions of the tool. They then received the first design brief. After creating a mood board, the designer filled out the questionnaire and was instructed to present it as if the experimenter were the customer. The designer was then asked to assess the mood board's quality for hypothetical use in a real setting, in the context of a semi-structured interview (see below). We repeated this process for the other condition and brief.

Questionnaires

AttrakDiff. AttrakDiff [25] measures perceived attractiveness and usability of a tool, distinguishing between pragmatic and hedonic types. It considers four dimensions: *Pragmatic Quality* (PQ), or the tool's ability to support the achievement

of behavioral goals; *Hedonic-Stimulation* (HQ-S), or the tool's ability to stimulate personal growth; *Hedonic-Identification* (HQ-I), or its ability to be appropriated by the user; and *Attractiveness* (ATT), an aggregate of PQ and HQ. The questionnaire entails rating 28 opposite-adjective pairs on these four dimensions on a seven-point scale (-3 to 3).

Creativity Support Index. CSI [7] is a standardized psychometric tool for assessing the perceived creativity support of a tool, looking at 1) collaboration, 2) enjoyment, 3) exploration, 4) expressiveness, 5) immersion, and 6) worthiness of effort.

Semi-structured Interviews

At the end of each experiment, we conducted a semi-structured interview (outline given in *Supplementary Materials*) focusing on experience, perceived issues, and the value of the tool and the AI support. We asked also about general satisfaction with the outcomes produced. The final designs and selected intermediate screenshots were used to aid recollection.

6 RESULTS

We report results from statistical testing and observations from interview data. Examples of mood boards created in the study are shown in Figure 4. All mood boards from the study are provided in *Supplementary Materials*.

Quantitative Results

We compare the two conditions (with and without AI) via data on four dependent variables: 1) usage of CCB suggestions, 2) AttrakDiff, 3) CSI, and 4) outcome appraisal. For statistical comparison of quantitative dependent variables, we use repeated-measures ANOVA.

Inclusion of CCB Suggestions. Most participants (13 out of 16) utilized at least one suggestion made by the bandit system. On average, those 13 included 2.3 CCB-provided images per final mood board (25.5%). While self-searched images were included more commonly, the probability of removing a CCB-suggested image after insertion in the mood board was only 3.7%, vs. 5.8% for self-searched images.

AttrakDiff. Hedonic-dimension scores increased with the AI (Table 2) but did not reach $\alpha = 0.05$ statistical significance. *Pragmatic Quality* (PQ) was significantly greater in the without-AI condition. In contrast, the value for the aggregate metric *Attractiveness* was significantly greater in the with-AI condition, from 9.9 to 14.4.

Looking more closely at PQ, we found that *Simplicity* ($F(1, 15) = 6.51, p < .05$) was significantly higher in without-AI than with-AI (means 2.31 vs. 1.63), as was *Clear Structure* ($F(1, 15) = 9.92, p < .01$; means 1.88 vs. .5). Predictability ($F(1, 15) = 7.06, p < .05$) also was significantly higher in the without-AI condition (mean .88) than in with-AI (mean -.13).

ID	Age	Sex	Area of design	Years of practice	Education
1	35	F	Fashion	10	PhD student
2	28	F	Architecture	6	PhD student
3	26	M	Architecture	2	PhD student
4	34	F	Fashion, textile	3	PhD student
5	29	F	Graphic, textile, industrial	2	MA
6	31	F	Textile, fashion	3	PhD student
7	34	F	Textile, industrial, material	6	PhD student
8	36	F	Furniture, industrial	5	PhD student
9	36	F	Textile	8	PhD student
10	33	M	Interaction	4	PhD student
11	33	F	Urban, graphic, digital, service	10	PhD student
12	39	F	Industrial	2.5	PhD student
13	38	M	Industrial, strategic	10	PhD student
14	28	F	Industrial, product, STS	2	PhD student
15	39	F	Web, fine arts, interaction	6	PhD
16	31	F	Industrial, interaction	13	Postdoctoral

Table 1: Participants' demographics and expertise.

Figure 4: Example mood boards designed in the study with the support of a CCB (a and b) and without (c).

In contrast, the effect of AI on HQ-S resulted mostly from an effect on Novelty ($F(1, 15) = 5.84, p < .05$; with-AI mean .69; without-AI mean 0). AI's effect on HQ-I stems mainly from a significant effect on Connectiveness ($F(1, 15) = 4.75, p < .05$; with-AI mean .19; without-AI mean -.44).

	Without-AI		With-AI		Sig.
	Score	SD	Score	SD	p
Pragmatic quality	21.14	13.67	10.57	12.41	.006*
Hedonic-Identification	1.86	7.15	5.57	7.59	.130
Hedonic-Stimulation	-3.14	7.54	4.14	7.63	.060
Attractiveness	9.86	2.61	14.43	4.12	.036*

Table 2: AttrakDiff results (* denotes significant difference).

CSI. Users rated all CSI the with-AI condition as more creativity-supporting on all dimensions except *Immersion*. However, none reached statistical significance (see Table 3).

Outcome Ratings. We asked the participants to rate their preference for the final designs on a scale of 1 to 7 (1 denotes strong preference for the without-AI result, 7 for the with-AI result). Their average preference has a median of 5 (mean 4.7), indicating a clear tendency to prefer the with-AI condition. Fourteen (of 16) reported preferring the with-AI condition.

In addition, we asked them to rate the final mood board in terms of 1) *usefulness* for presenting their cases to hypothetical customers and 2) perceived level of *surprise*. With both

metrics, the two systems were nearly equal. For perceived usefulness, both conditions showed a median of 5 (avg. 5.4 with AI vs. 5.1 without). For perceived surprise, the with-AI condition had a median of 3 (avg. 3.8) and the without-AI a median of 4 (avg. 3.7).

Qualitative Results

Capability of the Tool. In the interviews, 15 of the 16 participants stated that they would use the tool in their mood board process with some changes, especially “when you quickly want to create something,” since “the atmosphere [of] this is more convenient than in Photoshop or InDesign” (P6). Four of them highlighted effectiveness (P2, P9, P10, P12). Most designers reported the (non-AI parts of the) tool to be efficient, calling it “very fast to learn” (P1) and “quite straightforward” (P7) and saying that it “has everything needed” (P13, P15). Some reported missing functions such as cropping or color adjustments (P10), but most appreciated the simplicity of the tool and mentioned that it forces one to focus on the task itself (P12). The image search facility was commended. One participant found it a “very good idea to have it integrated into the system” (P1); another said, “Because every image that has to go to Photoshop has to be downloaded first, in that sense it is great [to have it integrated]” (P7). However, some were not always satisfied with the result quality from DuckDuckGo (P5, P16) and asked for a larger set of results (P2, P13).

Perceived AI Capability and Effectiveness. Fourteen designers deemed the AI version definitely more interesting for their work: e.g. “I obviously prefer having AI – this stimulated my brain more. Without AI, what I should do is very obvious” (P15). Help with “tricky topics” (P5) was highlighted especially. The system could help “if I feel I am stuck in existing solutions – I don’t generate anything new” (P11). One stated, “I didn’t see the [AI] part [in the non-AI condition] [...] – I was looking for that, because I got a bit stuck somewhat and thought it can suggest to me some other things” (P3). Six of the eight participants exposed to the with-AI condition first mentioned missing the AI’s suggestions afterwards (P3, P7–8, P14–16). We also asked what could *change*. Although some asked for

	Without-AI		With-AI		Sig.
Factor	Score	SD	Score	SD	p
Collaboration	13.52	3.54	13.60	4.33	.920
Enjoyment	12.89	3.73	13.98	3.39	.064
Exploration	10.39	3.38	11.72	2.73	.060
Expressiveness	9.38	3.20	9.84	3.32	.549
Immersion	11.80	4.08	11.25	5.57	.512
Results Worth Effort	12.11	3.05	13.67	3.89	.179
CSI	55.61	12.57	59.67	15.87	.134

Table 3: Creativity Support Index results.

more support to understand the AI (P10, P13, P16), most focused on adding functions similar to Adobe products, like zooming and cropping (P4, P13) or touch (P2, P3).

Agency and Adaptivity. The designers had a wide range of experiences related to the quality of the suggestions and the AI's role as a collaborator. Eight ascribed a sort of agency to the AI but appreciated that it leaves decisions to them, as in “[i]t has its own agenda, and it was making suggestions, but I had the choice to not follow that, so I did not feel that kind of obligation” (P3). For some, this meant that it “was trying to help me by showing the images that might inspire me, but it did not or it did not end up giving me what it wanted to give me” (P13). Others described the system as independent and stated, “I think it was a ‘she,’ and she maybe heard me but she had her own opinions as well, I think” (P14). One reflected, “I feel like I don’t work alone, I feel like there is another person [pointing to AI side]. It’s like having brainstorming with two people or in a workshop” (P15). Another, who followed many suggestions presented (P16), noted, “I cannot say [the mood board] is totally from me [...] it is also from ‘her,’ so it is a kind of collaboration between me and the system.”

Six other participants mentioned noticing that the system was adapting (P2: “The very first time it was very random, and then like the system starts to follow the colors or in accordance with things that I pick”) or that the system was “following what I was doing but not exactly following” (P8). Two reported having a feeling that the system was only following their guidance, which resulted in the perception either that it was “following me too literally; I thought it didn’t understand my direction at all” (P13) or that it was “definitely assisting me rather than on its own” (P5). The latter participant described the interaction thus: “I think it was trying to suggest stuff that could fit with mine, and when I started to try the ‘Surprise me’ they were related somehow to what is presented here [points to the mood board].” One was critical of the suggestions’ effect, though: “The suggestion panel was good, but now I am thinking: could it be also forcing me to become lazier, because it brings the images itself? Well, it is actually good for the outcome but maybe not the best for my designer self” (P14).

Characterization. We asked the designers to characterize the AI via some metaphor – e.g., an animal. We got responses ranging from a teenager to a companion or even an eccentric. P11 said, “it would be a bit like a teenager, because the images are not really clichés or anything but they are really specific in terms of blueness and colors they suggest, a bit like a teenager is looking for images”; P14 described the system as independent and as a “she” with “her own opinions,” and P15 called it an “eccentric collaborator,” as if there were another person. While P16 characterized the AI “a kind of collaboration,” P15 was critical and termed it “a very nice colleague who is not helpful.” In turn, P6 saw it more as a companion and P5 as “kind of a

helper [...], like a horse; in a way, ready to help if needed but fine on its own if not.” Only one participant (P12) criticized the AI for interrupting the workflow, “which would be fine, when I am getting stuck [...], but I didn’t see the role as really meaningful.” P11 made a very interesting remark about the broader influence of AI: “Basically, people just go to Behance or Pinterest and copy each other’s designs or whatnot. There are design inspiration websites all over the Web. But this, because you don’t have that, actually I think it is good, because it is your authentic stuff rather,” adding, “It was quite exciting because you don’t usually control when you get an inspiration – it is quite a process that you can’t really force, and being able to produce something that you didn’t know you knew before is, I think, always a good process.”

Novelty and Surprise. Participants reported being surprised by the suggestions: “I was surprised with this apple image [...]. That was my a-ha moment” (P3) and “[felt] ‘Oh!’, and I could go for something like that” (P7). Some (P8, P10, P12, P16) also used the suggestions to reflect on their work like “There was a couch. I did not even understand why, but [...] afterwards I actually thought ‘ah, café,’ also someplace where you like to spend time, so it was interesting but in a good way.” (P14)

A few participants observed that the suggestions sometimes pointed in very different directions from the current mood board: “there was some [suggested] grid image here that [could] have given a completely different graphical layout direction to the mood board [...], but I just didn’t take it, because I didn’t have the time to realign whatever I was doing” (P4).

Some surprises were also seen as interruptive: “I got a lot of bluish/purple and I found that a bit annoying. So yes, I know blue does inspire a bit of trust, but I would want it a little bit more happy” (P10). However, off-topic suggestions could also be positively disruptive: “This was a funny pic [points at AI’s history]. I found it more like a random throw, like ‘wake up your brain!’ and I think it is really pretty cool” (P10).

Explainability and Reflection. While only six participants noted the passive explanation (Fig. 1: 2), most mentioned the proactive questioning feature. Opinions were divided. Some said that it forced them “to think if in the next picture I should follow on content or follow on harmony; at least it indicated that I need to balance” (P16). Some said it helped them understand and reflect on why pictures were chosen (P3, P6, P9), and on “what [they] actually want” (P16).

However, it also raised doubts. Some felt criticized and were not sure whether “I was in a right direction or am I out of the context” (P3). In a surprise to us, some felt that this feature was meant not for supporting them but to train the AI (P10), and it was therefore found to be disturbing. In line with P13’s thinking, P10 would have preferred marking features themselves on images instead of the limited dialogue our

tool offered. That said, not all participants received proactive questions during the study, and some received only a few.

7 DISCUSSION

Overall, the results are positive. Of the 16 designers, 13 included CCB-made suggestions in their final mood boards (25.5% of their images, on average). Results indicate that the CCB-equipped tool improved *attractiveness*, the ability to express oneself and to support the achievement of one's goals. We attribute much of this positive feedback to CCBs' ability to both exploit the user's current strategy and explore alternative routes. Importantly, the explorative suggestions enable some serendipitous encounters, but do not wander too far from the current style, since constantly throwing in eccentric ideas would quickly lead to thwarting of the system. Users reported the suggestions as novel without being too eccentric or useless. They also told of being surprised by some suggestions, even reporting "a-ha" moments and insights that led them to change their approach to the task.

Unsurprisingly, this benefit came with an increase in perceived complexity (see PQ), in that the CCB added five elements to the UI, most of which required designers to assess a suggestion or reflect on their thought process. However, considering the positive and encouraging feedback, the added complexity did not deter participants from using the AI-augmented tool. We found some evidence of CCBs' ability to align suggestions with users' styles also. While alignment is not surprising in light of the extensive uses of bandit systems in personalization, it is valuable to know that in rapidly evolving activities such as ideation, a CCB can adapt to a designer's style in an acceptable timeframe.

Interactions with CCBs were commented on with somewhat surprising attributions of agency (8 participants) and even personality, such as "*an eccentric collaborator*" or "*a 'she'.*" Strong ascribing of collaborative and helping behaviors led to perceptions of mixed agency, such as P16's "*I cannot say [the final mood board] is totally from me [...] it is also from 'her'.*" Three participants even felt as if they were criticized or judged by the system when it asked for justification for the image choices they made. This calls for careful design of the interaction between the system and the designer, to facilitate and not hinder creative exploration.

In contrast to earlier ideation support tools [1, 20], our system offers verbal explanations for suggestions. However, most designers considered them unnecessary, because they formed their own criteria – often more nuanced – related to the fit of an image to the mood board. Designers use visual material mainly for abstracting ideas from the current concepts at hand to visualize an intention [40]. This might explain why our verbal justifications focused on low-level visual features were considered less meaningful even though those features were the reason for the suggestion.

Supporting the collection of ideation material might require more abstract explanations of relatedness and context. The system also asked the designer to reflect on the currently chosen images and their relation to the current mood board. These proactive questions were perceived as disrupting by some, similarly to the slider manipulations presented by the Drawing Apprentice [10]. However, we also received positive feedback indicating that these active questions can support the reflection on design choices – e.g., as a reminder that there are more dimensions that one might consider.

Application of CCBs

Our CCB-based approach showed promising results in a real-world design task. Being feature-driven, it requires defining the smallest meaningful features that together describe an inspirational motif or artifact – in our case, an image. We believe this approach has potential to be applied to other creative domains, such as dance or music, provided that similar meaningful descriptive aspects can be identified.

In choreography, these could be small movements described by posture, velocity, direction, and acceleration as expressed by a dancer. From an observed movement and ongoing choreography, the CCB could suggest continuing with a similar style or breaking from the current pattern. This could inspire choreographers to new creations, similarly to Viewpoints AI [20] but without requiring pre-defined rules. A CCB could allow more flexible exploration and exploitation that follow the flow of the choreography, by adapting to the preferences of the choreographer through online learning.

In music, a motif (e.g., a short sequence of notes) could be described by pitch, tempo, key, and so on. From such a feature vector, a CCB could either suggest a continuation with similar features or diverge from one or even several of them. The CCB would be able to adapt to the ongoing piece and to the musician's style, rather than rely exclusively on pre-training as Bob does [44]. In effect, it would allow the musician to more effectively *explore* music on the fly.

8 CONCLUSION

Supporting early stages of the design process is challenging for most machine-learning approaches. Accordingly, we have described a bandit-based method that shows promise as a technical basis for supporting design ideation, especially when it can be interfaced in a manner that neither insists on systematic explicit feedback nor compromises the designer's agency. We hope this work can inspire others to explore bandit approaches for visual and other creative processes.

9 ACKNOWLEDGEMENTS

This work has been funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 637991).

REFERENCES

- [1] Salvatore Andolina, Khalil Klouche, Diogo Cabral, Tuukka Ruotsalo, and Giulio Jacucci. 2015. InspirationWall: Supporting Idea Generation through Automatic Information Exploration. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 103–106.
- [2] Patti Bao, Elizabeth Gerber, Darren Gergle, and David Hoffman. 2010. Momentum: Getting and Staying on Topic During a Brainstorm. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 1233–1236.
- [3] Simone Bianco and Gianluigi Ciocca. 2015. User preferences modeling and learning for pleasing photo collage generation. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)* 12, 1 (2015), 6.
- [4] Sébastien Bubeck and Nicolo Cesa-Bianchi. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning* 5, 1 (2012), 1–122.
- [5] Tracy Cassidy. 2011. The mood board process modeled and understood as a qualitative design research tool. *Fashion Practice* 3, 2 (2011).
- [6] Joel Chan, Pao Siangliulue, Denisa Qori McDonald, Ruixue Liu, Reza Moradinezhad, Safa Aman, Erin T. Solovey, Krzysztof Z. Gajos, and Steven P. Dow. 2017. Semantically Far Inspirations Considered Harmful?: Accounting for Cognitive States in Collaborative Ideation. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition (C&C '17)*. ACM, 93–105.
- [7] 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), 21.
- [8] Douglass R. Cutting, David R. Karger, Jan O. Pedersen, and John W. Tukey. 1992. Scatter/Gather: A Cluster-based Approach to Browsing Large Document Collections. In *Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '92)*. ACM, 318–329.
- [9] Pedram Daei, Joel Pyykkö, Dorota Glowacka, and Samuel Kaski. 2016. Interactive Intent Modeling from Multiple Feedback Domains. In *Proceedings of the 21st International Conference on Intelligent User Interfaces (IUI '16)*. 71–75.
- [10] Nicholas Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. 2016. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*. ACM.
- [11] Kees Dorst and Nigel Cross. 2001. Creativity in the design process: co-evolution of problem-solution. *Design Studies* 22, 5 (2001), 425–437.
- [12] James Fogarty, Jodi Forlizzi, and Scott E Hudson. 2001. Aesthetic information collages: generating decorative displays that contain information. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*. ACM, 141–150.
- [13] James Fogarty, Desney Tan, Ashish Kapoor, and Simon Winder. 2008. CueFlik: Interactive Concept Learning in Image Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, 29–38.
- [14] Steve Garner and Deana McDonagh-Philp. 2001. Problem interpretation and resolution via visual stimuli: the use of ‘mood boards’ in design education. *Journal of Art & Design Education* 20, 1 (2001), 57–64.
- [15] Karni Gilon, Joel Chan, Felicia Y. Ng, Hila Liifshitz-Assaf, Aniket Kittur, and Dafna Shahaf. 2018. Analogy Mining for Specific Design Needs. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, Article 121, 11 pages.
- [16] Milene Gonçalves, Carlos Cardoso, and Petra Badke-Schaub. 2016. Inspiration choices that matter: the selection of external stimuli during ideation. *Design Science* 2 (2016).
- [17] Noe Vargas Hernandez, Jami J. Shah, and Steven M. Smith. 2010. Understanding design ideation mechanisms through multilevel aligned empirical studies. *Design Studies* 31, 4 (2010), 382–410.
- [18] DuckDuckGo Inc. 2018. DuckDuckGo Search engine. <https://duckduckgo.com>
- [19] Twinword Inc. 2018. Natural Language Processing APIs - Twinword API. <https://www.twinword.com/api/word-associations.php>
- [20] Mikhail Jacob, Alexander Zook, and Brian Magerko. 2013. Viewpoints AI: Procedurally Representing and Reasoning about Gestures. In *DiGRA Conference*.
- [21] Ben Jonson. 2005. Design ideation: the conceptual sketch in the digital age. *Design Studies* 26, 6 (2005), 613–624.
- [22] Al Keller, GJ Pasman, and Pieter Jan Stappers. 2006. Collections designers keep: Collecting visual material for inspiration and reference. *CoDesign* 2, 01 (2006), 17–33.
- [23] Yui Kita and Jun Rekimoto. 2018. V8 Storming: How Far Should Two Ideas Be?. In *Proceedings of the 9th Augmented Human International Conference*. ACM, 14.
- [24] Janin Koch, Magda Laszlo, Andrés Lucero, Antti Oulasvirta, et al. 2018. Surfing for Inspiration: Digital Inspirational Material in Design Practice. In *Proceedings of DRS2018 International Conference*.
- [25] Bettina Laugwitz, Theo Held, and Martin Schrepp. 2008. Construction and evaluation of a user experience questionnaire. In *Symposium of the Austrian HCI and Usability Engineering Group*. Springer, 63–76.
- [26] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A Contextual-bandit Approach to Personalized News Article Recommendation. In *Proceedings of the 19th International Conference on World Wide Web (WWW '10)*. 10.
- [27] Andrés Lucero. 2012. Framing, aligning, paradoxing, abstracting, and directing: how design mood boards work. In *Proceedings of the Designing Interactive Systems Conference*. ACM.
- [28] Andrés Lucero. 2015. Funky-Design-Spaces: Interactive Environments for Creativity Inspired by Observing Designers Making Mood Boards. In *Human–Computer Interaction*. Springer, 474–492.
- [29] Andrés Lucero, Dzmitry Aliakseyeu, and Jean-Bernard Martens. 2007. Augmenting Mood Boards: Flexible and Intuitive Interaction in the Context of the Design Studio. In *Horizontal Interactive Human–Computer Systems 2007*.
- [30] Andrés Lucero, Dzmitry Aliakseyeu, Kees Overbeeke, and Jean-Bernard Martens. 2009. An Interactive Support Tool to Convey the Intended Message in Asynchronous Presentations. In *Proceedings of the International Conference on Advances in Computer Entertainment Technology (ACE '09)*. ACM, 8.
- [31] Nic Lupfer, Andruid Kerne, Andrew M. Webb, and Rhema Linder. 2016. Patterns of Free-form Curation: Visual Thinking with Web Content. In *Proceedings of the 24th ACM International Conference on Multimedia (MM '16)*. ACM, 12–21.
- [32] Aliaksei Miniukovich and Antonella De Angeli. 2014. Quantification of Interface Visual Complexity. In *Proceedings of the 2014 international working conference on advanced visual interfaces*. ACM, 153–160.
- [33] Philippa Mothersill and V. Michael Bove. 2018. An Ontology of Computational Tools for Design Activities. In *Proceedings of DRS2018 International Conference*.
- [34] Changhoon Oh, Jungwoo Song, Jinhan Choi, Seonghyeon Kim, Sungwoo Lee, and Bongwon Suh. 2018. I Lead, You Help but Only with Enough Details: Understanding User Experience of Co-Creation with Artificial Intelligence. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. ACM, 649.
- [35] OpenCV. 2018. EAST: An Efficient and Accurate Scene Text Detector (EAST text detector). <https://www.pyimagesearch.com/2018/08/20/opencv-text-detection-east-text-detector/>
- [36] OpenCV. 2018. Face Detection using Haar Cascades. https://docs.opencv.org/3.4/d7/d8b/tutorial_py_face_detection.html

- [37] OpenCV. 2018. Structural Analysis and Shape Descriptors. https://docs.opencv.org/2.4/modules/imgproc/doc/structural_analysis_and_shape_descriptors.html
- [38] Charles Sanders Peirce. 1974. *Collected Papers of Charles Sanders Peirce*. Vol. 5. Harvard University Press.
- [39] Carsten Rother, Lucas Bordeaux, Youssef Hamadi, and Andrew Blake. 2006. AutoCollage. In *ACM SIGGRAPH 2006 Papers (SIGGRAPH '06)*.
- [40] Ben Shneiderman. 2007. Creativity support tools: Accelerating discovery and innovation. *Commun. ACM* 50, 12 (2007), 20–32.
- [41] Pao Siangliulue, Joel Chan, Krzysztof Z. Gajos, and Steven P. Dow. 2015. Providing Timely Examples Improves the Quantity and Quality of Generated Ideas. In *Proceedings of the 2015 ACM SIGCHI Conference on Creativity and Cognition (C&C '15)*. ACM, 83–92.
- [42] Kyle Steinfeld. 2017. Dreams may come. In *Acadia*. 590–599.
- [43] Cem Tekin and Mihaela van der Schaar. 2015. Distributed Online Learning via Cooperative Contextual Bandits. *IEEE Transactions on Signal Processing* 63, 14 (2015), 3700–3714.
- [44] Belinda Thom. 2000. BoB: an Interactive Improvisational Music Companion. In *Proceedings of the fourth international conference on Autonomous agents*. ACM, 309–316.
- [45] Jun Xiao, Xuemei Zhang, Phil Cheatle, Yuli Gao, and C Brian Atkins. 2008. Mixed-Initiative Photo Collage Authoring. In *Proceedings of the 16th ACM international conference on Multimedia*. ACM, 509–518.
- [46] Georgios N Yannakakis, Antonios Liapis, and Constantine Alexopoulos. 2014. Mixed-Initiative Co-Creativity. In *FDG*.