ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships

Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, Wendy Mackay

► To cite this version:

Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, et al.. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. Proceedings of the ACM on Human-Computer Interaction , 2020, 4 (CSCW1), pp.1-27. 10.1145/3392850 . hal-02867303

HAL Id: hal-02867303

<https://hal.science/hal-02867303>

Submitted on 14 Jun 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships

JANIN KOCH, Aalto University, Finland

NICOLAS TAFFIN, Université Paris-Saclay, Inria, CNRS, France

MICHEL BEAUDOUIN-LAFON, Université Paris-Saclay, CNRS, Inria, France

MARKKU LAINE, Aalto University, Finland

ANDRÉS LUCERO, Aalto University, Finland

WENDY E. MACKAY, Université Paris-Saclay, Inria, CNRS, France

Professional designers create mood boards to explore, visualize, and communicate hard-to-express ideas. We present *ImageSense*, an intelligent, collaborative ideation tool that combines individual and shared work spaces, as well as collaboration with multiple forms of intelligent agents. In the collection phase, *ImageSense* offers fluid transitions between serendipitous discovery of curated images via *ImageCascade*, combined text- and image-based *Semantic search*, and intelligent *AI suggestions* for finding new images. For later composition and reflection, *ImageSense* provides semantic labels, generated color palettes, and multiple tag clouds to help communicate the intent of the mood board. A study of nine professional designers revealed nuances in designers' preferences for designer-led, system-led, and mixed-initiative approaches that evolve throughout the design process. We discuss the challenges in creating effective human-computer partnerships for creative activities, and suggest directions for future research.

CCS Concepts: • **Human-centered computing** → **Interactive systems and tools**; *Computer supported cooperative work*; • **Applied computing** → Arts and humanities.

Additional Key Words and Phrases: Creativity Support Tools, Ideation, Agency, Mood Board Design

ACM Reference Format:

Janin Koch, Nicolas Taffin, Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy E. Mackay. 2020. *ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships*. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW1, Article 45 (May 2020), 27 pages. <https://doi.org/10.1145/3392850>

1 INTRODUCTION

Ideation is the process of generating original ideas that define and examine desirable aspects of a design project [28]. Designers constantly alternate between divergent thinking, which involves exploring as many solutions as possible, and convergent thinking, which reduces alternatives to find the best solutions. Designers draw from different sources of inspiration, including collaboration with others, reflection about their own ideas, and serendipitous encounters [83]. Although creativity support tools (CST) are designed to provide computational support for ideation, they usually only address one or two of these sources of inspiration [23]. Worse, they rarely address convergent thinking, so visually-oriented designers must come up with specific, text-based terms to search for images that emerge from their divergent thinking processes. We argue that creativity support

Authors' addresses: Janin Koch, janin.koch@aalto.fi, Aalto University, Finland; Nicolas Taffin, nicolas.taffin@lri.fr, Université Paris-Saclay, Inria, CNRS, France; Michel Beaudouin-Lafon, mbf@lri.fr, Université Paris-Saclay, CNRS, Inria, France; Markku Laine, markku.laine@aalto.fi, Aalto University, Finland; Andrés Lucero, lucero@acm.org, Aalto University, Finland; Wendy E. Mackay, mackay@lri.fr, Université Paris-Saclay, Inria, CNRS, France.

© 2020 Association for Computing Machinery.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Proceedings of the ACM on Human-Computer Interaction*, <https://doi.org/10.1145/3392850>.

tools should facilitate inspiration that stems from both convergent and divergent thinking, and help designers construct a new understanding of their work.

Given recent advances in machine learning (ML) and artificial intelligence (AI), we are interested in how to incorporate intelligent assistance into the ideation process, while leaving the human designer in control. The key challenge is how to share agency. In traditional recommender systems, users provide initial input that an intelligent agent “aggregates and directs to appropriate recipients” [66]. Early ‘mixed-initiative’ approaches advocated user-centered agency, where users guide the system toward a desired goal [34]. More recent work suggests that users and machines can both have agency [6, 76] and should actively share it [41, 81]. As collaborative activities such as idea generation grow more complex, this form of shared agency becomes more appealing. However, this raises a key challenge: How can we design the interaction so that designers benefit from intelligent support, but still retain control? What does a satisfying and effective ‘human-computer partnership’ look like for complex, evolving and open-ended creative tasks such as mood board design?

Furthermore, how can we accomplish this for highly collaborative tasks, where human designers, each with different ideation needs, collaborate with each other? We argue that effective tools should offer multiple forms of agency from entirely designer-led, to mixed, to system-led. This requires finding appropriate forms of interaction that support different levels of shared agency, between human designers and with the computer. We are interested in how to enhance the ideation workflow by increasing collaboration with other designers and providing multiple types of intelligent advice.

We focus on mood boards, visual collages composed of images, text, and objects, that express concepts, ideas and emotions. Commonly used in creative fields such as design or fashion, they “stimulate the perception and interpretation of more ephemeral phenomena such as color, texture, form, image and status” [26]. Designers often collaborate in the design of physical mood boards, where the act of finding, choosing and curating visual material not only helps designers express ideas they already have, but also inspires new ideas based on their reactions to the images that emerge [36]. Mood boards let designers explore hard-to-express ideas [13], and offer the potential for innovative discovery [26]. However, providing computational support for visual ideation is difficult, since designers’ goals evolve rapidly.

In this context, we pose the following research questions: RQ1: How can we integrate contributions from both human and intelligent agents seamlessly within a digital mood board? RQ2: Which kinds of intelligent assistance are appropriate for which types of ideation challenges? RQ3: How do human collaborators differ from intelligent assistants, and how can they support each other? In order to evaluate these questions we developed *ImageSense*, a digital mood board tool that supports both divergent and convergent thinking for visual ideation. *ImageSense* provides a collaborative environment for multiple designers to create a shared mood board, while taking advantage of several different intelligent tools that support different levels of inspiration and reflection.

Specifically, designers can 1) select from a cascade of images curated by other designers, 2) search using text and images exploiting semantically enriched images, 3) take advantage of an intelligent agent that explores relevant images with the designer, and 4) receive reflection and sensemaking support. The key contributions of this paper include: 1) the design and implementation of *ImageSense*, an intelligent, collaborative digital mood board tool that supports the full ideation process, where human designers retain control of the interaction and actively choose the type and level of machine agency; and 2) a study with nine professional designers that assesses how *ImageSense* supports collaborative mood board design, and contrasts human-human and human-machine collaboration in a realistic ideation task.

We first discuss related work in creativity support tools, collaborative tools and intelligent assistants, and identify key implications for design. Next, we introduce *ImageSense* and illustrate it with a realistic use scenario. After describing the technical details, we report the results of the study,

including implications for design. We conclude with a discussion of the challenges in incorporating intelligent agents into a collaborative design task, with corresponding directions for future research.

2 MOTIVATION: TOWARDS DIGITAL IDEATION

In design ideation, designers move between analysis and synthesis of ideas or concepts to construct a potential future [74]. Abductive reasoning [59] and abstraction allow designers to “break through to the a-ha! moment of inspiration” [70]. Ideation often involves verbal, visual or tangible material, which may be intentionally ambiguous to facilitate abstraction. However, the ability to see and reason about it is fundamental to ‘designerly’ thinking. Visual material is thus considered most suitable for supporting the construction of new ideas [74].

Designers and their clients use mood boards to generate, select and communicate ideas related to a design brief [47]. First the client expresses a rough idea for a product or service. Then, the designer *collects* appropriate visual material, traditionally by browsing through physical magazines. They also draw inspiration from art [39], walking in nature, talking to other designers, and by searching for images online. Designers select and compose images on the mood board, which helps them *interpret* relationships among the images and combine concepts [67]. The process is highly iterative, switching continuously between searching for images, laying them out, and looking for missing images [50]. The completed mood board is then presented to the client [12, 26, 47].

Mood boards are innovative and fun to create [26], but also encourage designers to probe more deeply into the project’s themes and concepts. Especially in the early stages of design, mood boards help designers visualize hard-to-express ideas [13] and serve as a powerful tool for communicating a “web of seemingly unconnected ideas, difficult to express verbally with similar impact” [22]. Finding the ‘right’ image and developing the ‘right’ narrative are crucial but difficult skills, and are often the hardest to teach [13].

Shneiderman [70] argues that computational support is particularly valuable when searching for and collecting materials. However, advanced creativity support tools are rarely, if ever, deployed in early-stage design [53]. This is partly because designers seek inspiration through diverse activities, and approach design problems from different perspectives [48], which highlights the need for diverse types of inspiration support. Another issue is that, although today’s computers can perform hundreds of image searches and analyses, this is only useful if the system knows *what* the user is currently looking for in terms of color, mood, content. Finally, designers who are in what Csikszentmihalyi [16] calls ‘flow’ i.e. immersed in idea generation, do not want to be distracted or interrupted by unwanted suggestions [27, 32, 38] from an intelligent agent.

3 RELATED WORK

3.1 Creativity Support Tools

Recent reviews of creativity support tools [23, 83] argue that most tools focus on finding new material by, e.g., retrieving previously searched material [19]; suggesting related images [7]; or encouraging collaboration [30, 77]. However, very few systems support the entire creative process [23], with support for different facets of inspirational practice [83]. We briefly review relevant work for the main phases of the mood board process: collecting and reflecting on inspirational material, as well as collaborative tools for ideation.

3.1.1 Collecting Inspirational Material. In creative practices, designers alternate between ‘exploration’ and ‘exploitation’ to reach a final design idea [62]. Computational tools can be particularly effective for supporting search and collection of material [69].

Digital mood board designers often turn to commercial search engines to collect this information such as *Google Image Search* [46]. Other, more designerly search engines, such as *ImageFinder* [64],

Unsplash [60], *Muzli Search* [10] and *Pinterest* [61] offer vast curated collections of high quality images. Machine learning approaches often support image-search queries using user pre-defined or automatic rules to find images more quickly or narrow down a query [17, 20, 21, 85]. Although these text-based search approaches enable the retrieval of images from large corpora, they provide designers with little help in defining the query itself.

Another promising approach is to improve the suggested material, instead of the query. For example, semantically-based search suggests images based on a textual analysis of the image's source web pages. Such search results can offer more inspirational images than those from standard image search engines, due to their semantic relevance [7]. Although semantic analysis offers an intriguing direction for ideation support, it still requires designers to articulate vague, highly visual ideas at the beginning of each search.

Improvements in computer vision and machine learning have recently made image-based search engines possible. *Google's Reverse Image Search* lets users upload images, which are then analyzed with image recognition algorithms that generate semantic labels [57]. This third approach lets designers find images that are similar to the seed image, but limits exploration, since users can neither see nor manipulate the search term.

SemanticCollage [42] combines the last two approaches by providing a semantic search based on the DuckDuckGo search engine to support multi-object search, including multiple images or text. Unlike Google's approach, semantic labels are presented to the user and can be selected based on their relevance. This approach helps designers be more expressive and explore visual material more easily. However, the approach is limited to open source images available in the DuckDuckGo corpora, which is more limited than the designers' preferred platforms.

Design implications: Combine the flexibility of text-based approaches that use design-oriented image resources, with those using semantic search and retrieval.

3.1.2 Reflecting with Tools. Designers must make sense of their own ideas, in a constant process of acquisition, reflection, and action [43, 67]. The reciprocal interaction between the users' perception of the material and the material itself helps designers elaborate, reflect upon and question their own understanding of their ideas [40], leading to new ideas and perspectives. Recent research highlights the importance of reflection in design for helping designers contextualize their ideas in the making [18, 52, 68].

Several tools support reflection for inspiration: *InspirationWall* [1] projects text-based stimuli on a 'wall'; *IdeaExpander* [82] uses visual stimuli to actively support reflection; *Rich bookmarks* [84] let users create links to visual and semantic metadata to reflect on the collected material; *CoSense* [58] summarizes earlier web searches as tag clouds, and *SemanticCollage* [42] uses semantic labels. These descriptive representations can help designers make sense of their design process and track the history of their earlier searches. However, sense-making also involves reflecting on ideas at different levels of abstraction, from a single object's signification, to meanings that emerge from a group of juxtaposed objects, to the holistic impression of the entire mood board. Designers need to be able to synthesize ideas and generate larger concepts based on their collected material and experience [43].

Design implications: Semantic analysis of textual and visual material can support abstract reflection and the creation of better textual and visual summaries.

3.1.3 Collaborating through Tools. In visual ideation, designers use their personal experiences as well as interact with other designers to find and develop new ideas [13]. Considering and discussing more extreme ideas with other designers can expand the designer's personal idea space. The physical

nature of traditional mood boards also invites designers and stakeholders to engage in spontaneous discussions. By contrast, commercial digital mood boards are primarily intended for individual use.

The advent of more powerful, distributed digital ideation tools has increased collaboration among designers [24, 51]. Brainstorming and other ideation methods often include walls or tables to collect and sort written ideas. Hence a number of tools use wall displays [29, 31] and tabletops, e.g. [33], to provide digital support for a physical process [56]. For example, the *Funky Wall* [50] interactive wall display and *Funky Coffee Table* [49] tabletop interface each facilitate collaborative image browsing. These demonstrate the potential of collaborative creativity when collecting and reflecting on design material through the simultaneous presentation of multiple ideas, but provide only limited control over the manipulation of images.

Design implications: Support collaborative ideation with large-scale, collaborative presentation tools.

3.1.4 Crowdsourcing. Crowdsourcing is another collaborative approach that harnesses the ‘wisdom of the crowd’ to enhance brainstorming. For example, *Crowdboard* [2] is a wall display where crowd-sourced ideas enhance in-person brainstorming in real time, extending the idea space. However, crowdsourcing is less relevant for mood board design. Visual ideation requires expertise in visual abstraction, a skill usually developed in professional design education [11, 13] or over years of professional experience, which implies that members of the crowd should be experienced designers. Another issue arises when the mood board involves confidential client or product information, making crowdsourcing impossible.

Finally, unlike brainstorming sessions that are typically restricted to short time periods, mood boards usually evolve over days or weeks, as designers reflect on and update the images in their search for a common narrative. Crowdsourced mood boards would require an equivalent commitment from the crowd, and require moderation to align new ideas with the common goal. Ideation platforms such as *IdeaHound* [71] address the last issue by applying semantic modeling to identify meaningful relations among ideas suggested by the crowd. A similar approach could use artificial agents that play the roles of external contributors, similar to *brAInstorm*’s [75] intelligent moderator.

Design implications: Develop artificial agents, as a form of external observing crowd, that mediate their contribution according to semantic analysis.

3.2 Intelligent Assistants

Co-creative agents let humans and machines share initiative in creative processes [45, 86] such as brainstorming, music and dance. Systems such as *InspirationWall* [1] or *Momentum* [3] inspire designers and help them organize and present text-based brainstorming ideas. However, while many brainstorming systems focus on suggesting related ideas, so-called ‘extreme suggestions’ provide another important source of inspiration [72], especially when designers feel ‘stuck’. Suggestions closely related, however, can help designers when they are in a creative flow [14]. Just as in other creative practices e.g. improvised music [79], contemporary dance [37], drawing [55], and ideation [41], intelligent agents can lead the ideation process if the designer so desires, but only with sufficient contextual explanation to let them reason about the idea. For example, *May AI* [41] uses a cooperative contextual bandit to suggest inspirational and contextually relevant material. The system learns when to explore or exploit visual material over time, but the suggestions and reasoning are limited to visual features, such as color. Steinfeld argues that “computers cannot see the way we see, they cannot help us to reason the way we wish to reason” [74]. This suggests that color-based approaches could be extended to handle semantic information [71].

Design implications: Combine visual features with semantic information to improve contribution's relevance to the current ideation process and better express relationships between visual material and the mood board.

4 DESIGN GOAL

Our goal is to integrate intelligent tools into a collaborative mood board design process that offers designers a fluid mix of serendipitous browsing, mixed text and image search, and proactive suggestions. Schön argues that design, as a reflective practice, relies upon the interpretation of action [67]. He stresses the importance of sense-making, especially in visual ideation.

We are interested in harnessing the potential of machine learning to 1) provide text-based summaries of the designer's visual concepts, and 2) help them develop useful abstractions of their concepts. The goal is to expand the designer's idea space into an active visual and verbal discussion platform with other professional designers. None of today's current tools support this range of interaction styles. We argue that a collaborative, intelligent digital mood board tool should:

- help designers identify what interests them, expressed both visually and textually;
- offer a diverse set of relevant inspirational sources, including serendipitous encounters, search and collaboration;
- support abstract reflection on visual and other related material to synthesize ideas and generate larger concepts;
- provide detailed control mechanisms for manipulating images, beyond current wall and tabletop systems;
- enable ideas suggested by remote human designers and artificial agents; and
- provide semantic analysis and meta-data to support image search [7, 42] and reflection [84].

We are particularly interested in how these intelligent tools operate in the context of collaboration among human designers. Most earlier intelligent tools have been tested with single-user scenarios. However, collaboration with other designers is common, with many benefits both for collecting and reflecting upon the material. We are interested in how designers choose to share agency with other human designers and with intelligent systems, when all are available, under the designers' control, in a single system.

5 IMAGESENSE: SCENARIO

We designed *ImageSense* to provide an integrated, collaborative ideation space that supports multiple designers throughout the entire creative process, from collection of visual material, to image and text composition, to reflection and communication of the results. A key design goal is to offer multiple forms of intelligent assistance, while ensuring that designers retain control over when and how to access the AI assistance. The following scenario illustrates how designers use *ImageSense*:

Ann is a designer working for a French car manufacturer. She is looking for fresh ideas for an upcoming car release. Bob is a design consultant, located in another country. He has the latest information about the car's specifications and the target group the company hopes to reach.

Ann and Bob use *ImageSense* to collaborate on a shared mood board. After logging in, they brainstorm some initial ideas and start searching for images by entering search terms into the search bar on the left (Fig. 1(c)). Bob drags several interesting family car images onto the canvas (Fig. 1(a)). He also sees several additional images that Ann has added. One glows red, indicating that she is currently resizing it (Fig. 2.a).

Fig. 1. *ImageSense* includes both a central *Mood board Canvas* (a) and *Maybe Area* (b). The four image collection methods include: *Search Bar* (c) with *Search Results* (d); *Upload area* (f); and *System Suggestions* (g). Composition and reflection tools include: *Color palette* (h); *Tool palette* (i); and *Semantic Tag Clouds* (j).

As they discuss the general direction of the mood board, Bob selects one of Ann’s green-tech images. He clicks the reflection tag cloud above the image (Fig. 1(j)), which displays the tags-by-images window with the associated semantic labels. He likes the idea of the ‘ecological’ aspect associated with the image and writes it in the search field.

Next, Bob drags one of his family cars into the search bar. He examines the labels and highlights relevant aspects concerning cars. He decides to reduce the ‘family’ characteristic, by clicking on the associated ‘-’ button (see Fig. 3.c). He presses ‘enter’ to retrieve more ecological car images, and drags several results onto the canvas. He decides to drag a few other potentially interesting images into the ‘maybe’ area (Fig. 1(b)). He sees that Ann already placed several other images there and takes a look. He also notices that the AI suggestions from the ‘suggest-o-matic’ dial have produced related images about alternative energy sources and futuristic-looking cities (Fig. 1(g)). Bob hovers over the suggestions and checks out the attached semantics in the tag cloud next to it. He likes the future-oriented ideas and adds those images to the mood board canvas.

Later, Ann and Bob agree that the mood board is missing several images. Ann investigates what the tag clouds say about the current semantics of the mood board, while Bob activates the Image Cascade by pressing the button under the search bar (Fig. 1(d) top). He browses through several potentially relevant categories using the arrows on the side and stops at ‘leisure’ (see Fig. 4). He watches the images flow down and occasionally clicks on one to enlarge it. He likes the images of surfers taking out surfboards of a large car, which makes him think of cars as living spaces. He drags one surfer image onto the board and explains his thinking to Ann. He also finds several related images from his computer, uploads them (Fig. 1(f)) into the ‘maybe’ area and drags a few onto the mood board. Meanwhile, Ann has rearranged images in groups to represent several concepts they discussed. She also added titles and sticky elements using the tool bar on the right (Fig. 1(i)).

Bob crops several images using the crop tool, to focus attention on the rough fabrics that appear in the image. He likes one of the cropped images and drags it into the search bar to look for more

related material. The cropped image now has new labels attached to it, related to the materials, e.g. roughness or organic. After searching, he adds more images to the board. In the meantime, the AI suggestions include several fabric patterns that match the style of his images but with different features such as softness, which he adds. He drags several box elements from the tool bar onto the canvas (Fig. 1(i)). He then aligns them and drags several color swatches (Fig. 1(h)) that let him extract the exact color to create an outline of the overall style. Finally, Bob changes several colors within an image by selecting them and manipulating the color tool on the right. He consults the ‘Abstract notions’ cloud to find overall topics that he writes on stickies dragged from the tool bar. He discusses the overall concept with Ann and exports the mood board for a meeting, using the ‘save mood board’ button at the bottom of the tool bar (Fig. 1(h)).

6 SYSTEM DESCRIPTION

We first give a short overview of the system architecture, then introduce the three types of tools that make up the interface: design tools, collection tools, and reflection tools.

6.1 System Architecture

ImageSense is a web application that supports synchronous collaboration among designers. The user interface uses standard web technologies (*HTML*, *CSS*, *JavaScript*) while the back-end uses the *web.py* Python framework and a *PostgreSQL* database. The custom collaboration server uses *Socket.io* [65] to share and synchronize data in the design space (Fig. 1(a)(b)).

Semantic and association labels are extracted whenever a designer adds a new image to the *Mood board Canvas*, *Maybe Area*, or *Search Bar*. As in previous work [42], we retrieve the semantic labels of images with the Google Vision object recognition and semantic labeling API¹ and attach the top ten labels to the image. Each label is then forwarded to a word-association API [80] to fetch the ten most common human-based associated terms and attach them to each semantic label and image. Finally, we analyze each image for its ten major colors using the MMCQ algorithm² [5].

6.2 Design Tools

The design tools include the main work area, called the *Mood board Canvas* (Fig. 1(a)), the *Maybe Area* (Fig. 1(b)), and the *Text and Graphical Object Tools* (Fig. 1(h) and (i)).

Mood board Canvas: Designers can freely drag images onto the shared *Mood board Canvas* (or *Canvas* for short), and resize them. Each image or graphical object embeds its own tools for changing the z-order, resizing, and deleting. Images also support cropping, recoloring, rotating, and flipping (Fig. 2.c). New semantic labels are generated when an image is transformed.

The content of the *Canvas* is shared among collaborators. Images and objects glow red when a remote designer manipulates them, which temporarily deactivates all editing functions for the other designers to avoid unintended conflicts, and glow blue when designers manipulate their own images (Fig. 2.a-b). Designers can activate a *detailed* collaboration mode to see image manipulations by remote collaborators in real time rather than just the end result.

Maybe Area: Designers can save potentially useful images in the *Maybe Area*, a scrollable list of images under the *Canvas* (Fig. 1(b)) that is also shared with collaborators. Images can be dragged in or out of this area at any time, and are grayed out to minimize distraction.

¹<https://cloud.google.com/vision/>

²<https://github.com/fengsp/color-thief-py>

Text and Graphical Object Tools: The tool panel (Fig. 1*i*) features standard tools for creating shapes (circles and boxes), text, and sticky notes. It also contains color tools (Fig. 1*h*): a color selector with opacity slider and a color palette. Dragging color swatches to the *Canvas* changes the color of the underlying object or background.

a. Images glow red when a collaborator selects or edits them.

b. Selected images glow blue. Tools include crop, resize, delete, back and front.

c. Additional crop tools include: recolor, rotate and flip horizontally or vertically.

Fig. 2. Collaborative image editing

6.3 Collection Tools

ImageSense provides four tools for collecting images: *Upload*, *ImageCascade*, *Semantic Search*, and *AI Suggestions*. The last two are ‘intelligent’ tools that use AI techniques.

Upload: Personal photos, art, or self-created sketches can be added to the *Canvas* by dragging them into the *Upload* area (Fig. 1*f*). The system tags them with semantic, associative, and color meta-data as well as adds them to the shared *Maybe Area*.

Image Cascade: Clicking the *Image Cascade* button under the search bar replaces the search result panel with a cascade of categorized photos that drop slowly from the top (Fig. 4). We use 2700 photos from *Death to Stockphoto* [73], manually organized into nine common categories (architecture, art, food, leisure, living spaces, nature, objects, people, and work). The images drop at different speeds, creating a cascade effect with at most 15 simultaneous images. Clicking on an image enlarges it (Fig. 4.b) so that it can be inspected and dragged into the *Canvas* or *Maybe Area*.

Semantic Search: Designers can initiate a search using text and/or images and a selection of search engines (Fig. 3.a). Images can be dragged from anywhere in the interface onto the search bar

a. Changing the search engine preference

b. Dragging an image to the search bar

c. Editable semantic labels help to clarify the current idea

Fig. 3. *Semantic Search* supports both text and images as input and uses semantic information

Fig. 4. The ImageCascade displays a collection of stock images, organized in categories

(Fig. 3.b). When adding an image, the attached semantic labels appear underneath it (Fig. 3.c) with their confidence score: most relevant (in **bold**), relevant (in normal text), or not relevant (strike-through). Designers can adjust these weights with $\oplus$ and $\ominus$ buttons (Fig. 3.c).

When triggered, the search replaces the images with the corresponding relevant labels. Search results appear in the left-hand panel, and can be dragged to the *Canvas*, *Maybe Area*, or back to the search bar. Previous searches are available in the *Search History*.

AI Suggestions: The AI system is triggered when a designer adds or removes an element from the *Canvas* or clicks the *Request more* button (Fig. 5©). We developed a cooperative contextual bandit as a collaborative agent, similar to Koch et al. [41]. We extended the algorithm to include semantic and cognitive features that provide more semantically relevant contributions.

In order to analyse the colors, we convert a screenshot of the mood board into numeric values and allocate it to a bandit agent. Each bandit agent is a uniformly sliced [78] combination of color (c), saturation (s), luminance (l), image orientation (o), and color distance (d) [41]. However, we found that the original slicing of luminance into bright, dark, or color was too rough, and thus halved the slicing parameter for saturation and luminance, resulting in smaller but more accurate categories. Depending on which bandit agent's probability is higher, exploitation or exploration, the bandit suggests a (c,s,l,o,d) vector similar to the current mood board or diverging from it [41].

In order to select relevant semantics for suggestions, we sample from an Ebbinghaus forgetting curve [54] over the chronological order of images added to the mood board, to gradually forget early images. This reflects the decreasing relevance of semantics over time as the designers' ideas

Fig. 5. AI features include: suggested images (a); semantics derived from images (b) triggered when hovering; Request more button (c) for additional suggestions according to the setting of the Suggest-o-matic Dial (d): 'more like this', 'not this one', or 'surprise me'

evolve. By converting the suggested image-color vector into human understandable words and combining it with the selected semantics, the system can retrieve new suggestions using any of the supported search engines.

Each resulting image is downloaded and compared to the selected semantics, their associations and color vector. The AI then suggests images that match both color and semantic criteria (Fig. 5④). Semantic labels appear as a tag cloud visualization (Fig. 5⑤) to help designers make sense of the suggestions. Designers can steer the exploratory nature of the *AI Suggestions* with the ‘suggest-o-matic’ dial (Fig. 5④): ‘*More like this*’ encourages the AI to exploit the current color strategy; ‘*Not this one*’ enforces the learned model prediction; and ‘*surprise me*’ encourages additional exploration.

6.4 Reflection Tools

ImageSense features reflection tools in the form of dynamically generated tag clouds based on semantic information extracted from the images. We extend earlier work on summary clouds for reflection [42, 58] by providing five clouds representing different levels of abstraction of the mood board content. Based on pilot tests, we provide two identical drawers so designers can select two tag clouds at the same time (Fig. 6).

Tag by image shows the semantic labels of the currently selected image, with $\oplus/\ominus$ buttons to let designers adjust their weights. Designers can also add new labels to customize the clouds. All changes are attached to the image and affect, e.g., the search tool.

Mood board tags shows all the labels from the images on the mood board. Font size is linked to the number of images with that label and the relative surface they cover on the mood board. Resizing an image hence dynamically updates the cloud. Clicking a label highlights all images on the mood board that share that label.

Association tags shows associations related to the current mood board labels. These associations are ranked by their frequency and size dominance from the related images in the mood board. To avoid over crowding, we only show the top twenty highest-ranked associations. Selecting a label highlights all the images on the mood board that share it.

Abstract notions shows the text classification of current mood board labels ordered by their size dominance on the board. We apply a text classification algorithm [80] to find larger categories that encompass the selected labels. The tag displays the twenty most relevant categories (Fig. 6-right). Due to the abstraction level, there is no direct link between a category and individual images.

Color moods shows emotional states and feelings associated with the colors on the board. We extended Berlin & Kay’s eleven basic color terms [4] with cyan and magenta. Based on practitioner literature [15, 25] and online sources [35], we created a color-to-feeling dictionary mapping colors and their list of meanings. The color dominance on the mood board determines how many words are selected randomly from the color-to-feeling lists and displayed. To facilitate understanding, we color each label according to that mapping (Fig. 6-left).

Fig. 6. Two of the five *Semantic Tag Clouds*: *Color moods* and *Abstract notions*

7 STUDY

7.1 Method

We are interested in how *ImageSense* can help professional designers progress through a typical ideation process, including collection, composition, reflection, and final presentation. Our goal is to assess whether and how *ImageSense* helps them collaborate with a remote human designer, and the circumstances under which different AI tools are most effective. We do not focus on the quality of the mood boards created, but rather on how and when participants choose to collaborate with the remote human designer and the various *ImageSense* tools. The study is modeled after Bousseau et al. [8] structured observation of designers as they go through a complete design process.

ID	Age	Sex	Area of design	Years of practice
1	29	F	Fine Arts, Design Research	5
2	34	M	Interaction Design	10
3	32	M	Interactive Media Design	6
4	37	F	Textile Design	8
5	27	M	Architecture	4
6	32	F	Mechanical Engineering, Drama Instructor	7
7	44	F	Architecture, Experience Design	18
8	29	F	International Trade, Fashion Design	2
9	40	F	Fine Arts, Interaction Design	11

Table 1. Professional designers have diverse backgrounds, with 2-18 years of experience.

7.2 Participants

Effective mood board design requires advanced visual abstraction skills, hence we limited our study to designers with at least two years of professional experience. We recruited nine professional designers (6 women, 3 men, ages 27–44). We targeted professional designers with prior mood board experience, recruited via email lists to local design and media students and personal networks. The recruitment text highlighted the importance of practice experience in design and previous mood board experience. Participants were informed about the procedure, their rights and data usage by an informed consent form and received one movie ticket for their participation.

7.3 Setup

Our goal was to create a realistic design setting where each participant collaborates with an experienced, remote designer. We simulate a designer’s office-based work environment, where designers have free access to white boards, personal devices such as phones, and drinks. We chose a confederate study design, as is common for complex or safety critical scenarios [44] where participants need potential guidance. It further increases comparability across participants and conditions. The confederate is a professional graphic designer with 25 years of design experience, including extensive use of mood boards for both design and marketing.

The confederate plays the role of Pierre, a French designer working for an eco-friendly car company, who ensures that participants progress through the entire design process. He acts as both a resource for explaining how *ImageSense* works, and also as a remote collaborator who participates in the creation of the mood board, according to the design brief. As suggested in the literature, we scripted his role, so that he provided equivalent instructions and assistance about *ImageSense*, and ensured that participants spent roughly equal amounts of time on the different activities. He created a similar rapport with each participant, taking an advisory role rather than acting as lead

designer; but also to respond to the participants' requests for suggestions or help. Introducing him as a representative of the client's company added ecological validity to the process, making him more expert in the tool but also actively seeking their new ideas for the marketing campaign. We pilot-tested the interaction between the confederate and another designer, to establish appropriate timing for the task and to ensure similar levels of interaction across participants.

An experimenter sets up the equipment, welcomes the participant, and conducts briefing and debriefing interviews. The participant sits at a desk with two monitors, a keyboard and a mouse. The larger monitor displays the *ImageSense* interface at full screen and a smaller screen displays FaceTime to establish a separate audio and video connection with the remote designer (see Fig. 7). The confederate sits at a remote location in another country, with a similar laptop and Monitor setup with the *ImageSense* interface displayed full screen, and a FaceTime audio and video connection on an iPhone. The experimenter was *not* present during the design and reflection phases, to avoid giving the designer the feeling of 'being watched', but returns to act as the jury for the final presentation of the mood board by the participant, and the final interview.

Fig. 7. Study setup: Participant (left) and remote confederate (right)

7.4 Procedure

Each session lasts approximately one hour. After setting up the system, the experimenter welcomes the participant, and describes the study, the participant's rights and which data will be collected and stored according to GDPR guidelines. After the participant signs the informed consent form, the experimenter starts recording audio data. The initial briefing interview asks for demographic data and current mood board practice, including how often, which type (physical/digital) and which materials are commonly used. Afterwards, the experimenter explains the study procedure and hands over the design brief. The participant has time to read through the design brief and ask questions (see below).

Structure: The confederate appears on the small laptop using a FaceTime audio and video connection, and follows a standard script to explain the basic features of the system. The participant tries each feature for about five minutes or when they feel comfortable with the *ImageSense* interface. The experimenter then resets *ImageSense* to create a new, empty mood board and leaves the room. The participant and the confederate have 35 minutes to collect relevant images, compose them on the Mood Board, and reflect on how to communicate the ideas to an independent juror (the experimenter). The confederate acts as a true design partner, but follows the participant's lead. After 10 minutes, the experimenter reminds the designers that they need to think about the composition phase, if they have not already done so. After 20 minutes, the experimenter reminds the designers that they need to think about the reflection phase, if they have not already done so. After 30 minutes, the experimenter reminds the designers that the participant will have to present the Mood

Board in 5 minutes. Throughout, designers are free to work on any aspect of the mood board at any time. After 35 minutes, the experimenter returns and asks the participant to explain the current mood board design. After the participant speaks, the confederate may add additional comments. The two designers are encouraged to react to each other's comments. The experimenter concludes with a semi-structured interview with the participant. The goal is to obtain specific examples of when the participants' interaction with the human designer or one of the four intelligent assistants was particularly helpful (or not). Participants are also urged to reflect on situations in their current practice where particular types of collaboration with the designer or *ImageSense* would be helpful. The participants rank each tool according to their assessment of its usefulness and which tool they would 'take home' for everyday use and why. Finally, participants fill out a brief questionnaire, with rankings and Likert-style questions.

Design Brief: The design brief focuses on an eco-friendly car from a French car manufacturer. After a short introduction about eco-friendly cars and their increasing role in Europe, the brief describes the client as a "French car manufacturer, who wants to stay anonymous, and who plans to reveal a new electric car at the upcoming 'Paris Mondial de l'Automobile' in Paris." Core features of the car include zero emissions and non-toxic, sustainable materials to improve the air quality inside the car. The primary target group is the family. However, the car also includes a built-in car sharing system for friends, family, or other registered drivers that indicates possible shared routes. The design task is described as a design competition, where the participant will suggest a new visual campaign for the car. The brief combines traditional images of comfortable and safe French cars, but with an ecological twist.

7.5 Data Collection and Analysis

We screen-recorded each session, and audio-recorded the participants' discussions with the confederate. In addition we saved the final mood boards, a video recording of the mood board presentation, and the answers to the questionnaire. Further hand-written notes from the experimenter and the confederate were collected. To analyze the data we conducted a thematic analysis [9] for the semi-structured interview data, using a mixed top-down and bottom-up approach. Specifically, we first read all transcribed interviews and developed an initial coding for all topics relevant to our research questions. We then grouped closely related topics in order to generate initial themes. Descriptive statistics of the participants' answers to the Likert scales are reported.

8 RESULTS

Nine professional designers successfully completed mood boards that met the requirements of the design brief, each with innovative ideas for the eco-friendly car visual campaign, shown in Fig. 8. The following results first focus on how *ImageSense* enables and integrates in a fluid ideation process, followed by a closer look on how *ImageSense* supports design collaboration as well as collecting and reflection on inspirational material. We conclude with comments regarding the differences between human and machine support and opportunities for future studies.

8.1 *ImageSense*, a realistic Ideation Tool

Participants liked the tool: they "would love to have one" (P6) and see "a huge potential for this collaboration direction" (P9). In total 5/9 participants reported that they would find it useful in their everyday design practice. Designers described three aspects that make *ImageSense* useful: as a communication tool, as a collaboration tool and as an inspirational tool.

Fig. 8. Mood boards created by designers in the study (participants 1-9: left to right; top to bottom)

Communicate your Design to Others: most designers found *ImageSense* helpful to communicate their design (8/9) and to reflect on their composition (6/9). P9 reported: “Especially for communication, it was very easy and it makes communication more easy.” One fashion designer explained that “when you work with smart textiles it’s a lot about communicating something that doesn’t exist yet.” (P4) The high quality of the images help communicate such a future to the clients. Similarly, an architect (P7) mentioned that *ImageSense* would have encouraged her to build more mood boards in the past to better communicate unknown styles to her clients, especially when time and money is short for more elaborate ideation processes.

Collaborate visually: participants also highlighted the visual collaboration feature of *ImageSense*, especially when setting up a project that requires remote collaboration (P2). P2 described that he usually uses tools such as “storm boards”, an online mood boarding tool for workshops, when he cannot get everyone in the same location. Unlike existing tools, with *ImageSense* “you can have the pictures and the conversation at the same time”. Similarly, P8 mentioned that *ImageSense* is particular helpful as “a very fast way of communication with the team” for ideation.

A new way to find Inspiration: designers appreciated *ImageSense* for exploring different idea alternatives (8/9) and felt more creative (5/9). This applied to creating mood boards or “just ordering my thoughts”, for “more artistic work” or “com[ing] up with stories, to come up with living worlds” (P3). P3 saw as the main benefit of *ImageSense* “that there are so many ways to become inspired and to reflect and evolve your thoughts and that’s why I would definitely use it”. Several designers mentioned its potential for “the initial phase of the design process, if you’re choosing materials or if you’re trying to come up with a concept” (P5) or as a visual tool “for brainstorming to explore different directions” at the beginning of a project (P9). It was deemed “super useful because if I would do a mood board like I used to, then it would be just Photoshop [...] and then I just drag in photos that I find somewhere, but then it’s much harder in the process to refine them and to be inspired.” (P3) This ability could also help new designers “to generate or talk about their concept, which a lot of people struggle with. They make fantastic ideas but they cannot talk about it so much. So this is like a tool that you could train students [...] which then carry out actual practice.” (P7)

8.2 Supporting Collaboration with *ImageSense*

The overall consensus among participants was that *ImageSense* supports rather than hinders collaboration with a remote designer: “there is this shared space for this cooperative design. I think that’s something that I would actually like to use in the future” (P4), because “working together with another person brings unpredictable situations, which is very good” (P9). Consistent with previous literature, designers appreciated the involvement and discussion with a design collaborator to bounce ideas off someone (P1, P4), get inspired (P6) and overcome the problem of starting with a ‘blank sheet’, which can be very hard (P2).

Starting and picking up momentum: many participants (5/9) stated that collaborating with another designer helped start the process and gather momentum: “there’s always this process where you first have to find your thoughts” (P3). Participants’ ideation strategies varied greatly: some began by proposing a general idea (P2), which Pierre used to gather initial images and ‘kickstart’ the process. Other participants asked for Pierre’s opinion (P5), and said that discussing the topic with him helped them ideate. Since Pierre also explained how to use *ImageSense*, he was also usually the first to bring in “a stream of images” (P3), which both inspired them and encouraged them to react and search for their own images.

Collaborative Ideation: some participants (3/9) mentioned that each designers’ role in ideating or searching for images switched throughout the process: “one guides and the other follows, but then that changes through time depending on what you like in different states” (P1).

Participants and Pierre offered diverse feedback to each other, ranging from “complimenting the other person” (P1) to “negotiating” (P6) and “[validating] the relevance of [a] picture” (P5), to simply “dealing with it” (P3). This was facilitated by both the audio/video feeds and the synchronized design space. Designers reported they could better “understand these pictures he’s putting in” (P3) and discuss “a lot about writing [the final text labels]” (P9). The back-and-forth between the two designers also encouraged the idea generation process: participants mentioned that it is “pushing me to think faster” (P7), which enabled them to create concepts that “would have not been possible in that fast time” alone (P3).

8.3 Exploring and Finding Inspiration from AI Tools

Most participants (8/9) reported that *ImageSense* helped them to explore design alternatives during the process, and to collect useful images (6/9). As described in our design goals, *ImageSense* supports a diverse set of inspirational sources commonly used by designers. In line with the literature, *ImageSense* supports discovery by aimless browsing, active searching, and through collaboration in a digital manner. In the following we present comments regarding these intelligent tools.

8.3.1 Browsing with ImageCascade. Getting inspiration or starting points can sometimes come from unrelated places. *ImageCascade* allows designers to extend their thought by using unrelated, but not random, images. When asked which tool they would like to take home and use in their everyday practice, 5/9 participants reported the *ImageCascade* tool.

Extending One’s Thoughts: designers described the main benefit of *ImageCascade* as the ability to steer inspiration towards new directions. Designers found it more useful to browse “instead of going too deeply into what we already have” (P3), especially when feeling stuck and when they “might not have found what we want to express yet” (P3). Browsing through high-quality images could “take you to many directions away from a specific task and open things up and kind of explore a wider quantity of inspiration” (P6). By observing the stream of images, designers sometimes found new, unexpected perspectives: “‘Oh that’s great’, and then I would type something up here that

fits the picture, but it is not necessarily [that] picture” (P1). The floating animation supported this serendipity by providing several images that “are moving and disappearing automatically” (P9), which was a helpful “interaction to get a lot of imagery and inspiration” (P9).

Inspiration with Direction: designers described that the presented categories gave them a certain level of control over the suggested images: “Okay, let us search for art’ or ‘Let’s search for people’ and then I get inspired in a certain direction.” (P3) While these broad categories were appreciated by most participants, some designers (2/9) were interested in creating their own category “that would suit this particular task” (P4), because otherwise “it’s a bit of luck. It might be the thing that you’re looking for but also it may not be.” (P5)

8.3.2 Semantic Search. P3 stated that “[he] personally think[s] that the mood board is a lot [of] search process” and concluded that *Semantic Search* was hence “the strongest tool, because that’s what I have to do anyway. I have to refine the search terms, I have to refine semantically what I’m actually doing there, because it is visual” (P3). Most participants (8/9) agree that the *Semantic Search* is the most useful tool overall, and four added that they would love to use it in their everyday work.

Searching for Hard-to-Express Ideas: the “many ways or points of access or angles of attack” (P3) to find images were described as the main benefit of *Semantic Search*. He continued that with *Semantic Search* “you can put in an image, you can refine using the words that come up or you can use and change the semantics as well” to redefine it to your to your imagination. The ability to combine images and text “to specify the results that I wanted to get” (P5) was deemed useful especially when it is hard to express ideas and search terms (P2,P3,P4,P6,P9). P4 put it nicely as follows: “If you don’t know how to search for something like ‘this’ or some kind of feeling you would like to have [...] then I think it was easier that you don’t need to know what you are searching, you don’t need to know the words, but you kind of know that I want to search for something that is [...] this interior textile and then having also [...] this futuristic car and then getting search results about this car interiors” (P4).

A different approach to search. Some designers (2/9) mentioned that searching with images, text and semantic labels requires designers “to change the way you think, like your entire approach to getting images” (P2). This might not be as “intuitive [as] ‘I type a word and now I drop it there’” (P1) at first, even though it is beneficial to the result. When asked about the limits of such a system, designers mentioned that descriptive labels might not be sufficient “when you don’t know what you are doing” (P4) or when your thoughts are very abstract, e.g. “I want to take the ocean and find some similar fabric” (P6).

8.3.3 AI Suggestions. Half of the participants ranked the *AI Suggestions* from the ‘suggest-o-matic’ as most useful when feeling stuck or looking for inspiration, even more than the human collaborator.

Collaborative Ideation: designers felt that the *AI Suggestions* were particularly helpful when exploring new ideas in the beginning (P4), throughout the process (P3,P5,P6), or in later phases (P2). P8 described how such an agent broadens the idea space as follows: “It gives you a little bit of a push if you need it. [...] when you think about something very much, you still stay focused on the points and you don’t see outside of it. So it’s good to have something that is shocking or something that is out of reach. Or so I think this is very out of reach but also could be very good for ideation like new things and daring things just to try and see what happens” (P8) Many designers commented on the AI’s ability to inspire them when ideas were sparse and they were “in a dire situation” (P3) with suggestions directly related to the current mood board, unlike *ImageCascade* (P1). Even if “it took a bit of time to really find the right picture” (P5), designers “felt like [the system] understood me” (P5) and the suggestions were surprising (P1, P3, P6) but accurate (P8).

Inspiration source vs. finding the ‘right’ image: some designers (3/9) commented on the trade-off between finding sources of inspiration and finding the ‘perfect’ image. Negative feedback ranged

from suggestions being “too close to what I do” (P2) to being “too random” (P6), which reveals a major challenge for designing AI systems that support creative activities. Designers found the *AI Suggestions* less useful when they “wanted the perfect thing” (P7). This contrasts with some earlier comments about the accuracy and usefulness of the suggestions, and reflects the participants’ changing needs and assumptions: “It sort of gives me some related pictures but I felt like it got less and less useful over time” (P2); “I requested more and more but [...] it didn’t fit I guess” (P6).

8.4 Articulating Ideas with Semantic Clouds

Nearly all participants (8/9) reported in the questionnaire that *ImageSense* helped them to ‘communicate their design’ to others and some felt it helped them to better ‘articulate what I was looking for’ (4/9) and improved their ability to ‘compose images and text’ better (3/9).

A new way to express oneself: designers described the clouds using terms such as “useful” (P6), “interesting”, “very inspiring” (P3), and “helpful” (P3, P8), and that “it just responds and reads [the mood board] so well” (P7). The clouds also provided adjacent concepts that helped opening up the process: “it helped me to find some words I didn’t think about it, like a neighborhood [which] can be the place, but it can be also this idea of being together in it and doing things together” (P3). The clouds also helped summarize the concepts in an existing moodboard: “And then it actually showed at one point what I was really looking for in the whole concept” (P7). Incidentally, non-native speakers also found it “helpful to find words for the thing that you have in your mind” (P8).

Getting inspiration for search: some designers (3/9) found the reflection clouds useful to “get inspiration for keywords for the search tool” (P5). They found the *Tags by image* especially important to understand what collaborators, human or machine, have added to the mood board. It helped them analyze “what is this image about” (P3) or “what kind of messages or different kind of associations do we have here” (P4), and use it as a new starting point for future exploration.

Clouds as a form of verbal representation: designers appreciated the tag clouds and suggested that “they could actually be a nice way of presenting the work as well. So then towards the end having it like a word cloud would be quite nice” (P2). However, the clouds should be configurable so that designers could “remove some of them” (P2) as well as use them as “a way to tag your photos” (P6) to highlight important aspects.

8.5 Role of Human and System Support in Ideation

In order to get a better understanding of the role and function that intelligent tools can play in collaborative ideation, we asked participants to rank their preferred tool for various use cases (multiple answers were allowed).

The first use case described a situation where ‘I knew exactly what image I was looking for’. Most participants (7/9) felt that the human designer was most helpful, followed by *Semantic Search* (6/9). Only two designers mentioned *ImageCascade* and *Reflection Clouds* and only one mentioned the *AI Suggestions*. The second use case referred to a situation where ‘I had a vague idea of an image in my head’. Most participants preferred working with a human designer (6/9), followed by *Semantic Search* (5/9) and *Reflection Clouds* (3/9). Only one designer preferred *ImageCascade* and another the *AI Suggestions*. The third use case described a situation where ‘I was stuck and needed new ideas’. All participants ranked the *AI Suggestions* as either most useful (4/9) or second most useful (5/9). Next tools that were ranked either most or second most useful were *Semantic Search* (5/9), the other designer (3/9) and *ImageCascade* and *Reflection Clouds* (2/9).

Participants commented that different features of *ImageSense* support different aspects of the mood board process. One dividing line is searching for something specific vs. reflecting on the collected material. When looking for material that they know how to find, some designers focused on “searching for pictures [and] were not collaborating that much” (P5) with Pierre. P7 even

commented that she would have liked to “take a little more time, more calmer on my own time” to find the image she liked. By contrast, when designers knew exactly what to look for, but had a hard time expressing it as a query, they often turned to Pierre to discuss potential search terms. Some even asked Pierre to find images ‘with’ them, by explaining their intended image. This helped especially in the beginning, because “starting from a blank sheet is very, very hard” (P2) and “there’s always this process where you first have to find your thoughts” (P3). Working together with Pierre was perceived as easier by some participants. Some participants also reported that later on in the process, collaborating with Pierre triggered new ideas: “sometimes you think that words are enough to express what I really want, but then he comes up with some image like these bottles and then it’s like ‘Oh, that I didn’t think of’ and then the jump in the thinking process was quite fast” (P7).

In situations where the participants felt stuck or needed ideas, they were more likely to use the intelligent tools. For example P4 used several images from the *AI Suggestions* in the beginning and explained her approach as first “getting these requested pictures then I can start by using these and then let’s see what’s going to happen” (P4). Others spend more than five minutes browsing through the *ImageCascade*, as they felt something was missing, but could not explain what (P1). Some participants also spent long periods of time looking at the semantic labels and clouds as “something [where] I can analyze what I’ve done, something [where] I can get deeper into ‘ok, what is this visual thing actually and what makes it interesting’” (P3).

Many participants appreciated the discussion with Pierre. Two designers mentioned that they mainly reflected “through conversation” (P1) with the human collaborator. Others found the task “quite clear because there is a lot of description” (P9), but mentioned that “if this is more like a brainstorming phase I might refer to [reflection clouds] more.” We observed a common pattern whereby after searching for material, both users were “re-evaluat[ing] the things that we put here” (P5) and validating “the relevance of [each other’s] picture or how can this represent the idea” (P5). However, when it was hard to make sense of the material, designers turned to intelligent tools. For example, some participants (3/9) used *Semantic Search* to interpret their own visual thoughts, or the other designer’s, “to get a similar associated image” (P9). In one case a participant used the image search to investigate images they had in mind but were added by the collaborator, to “search what he had searched for” (P3). This “traceability, [helps] understand how did he end up with this” (P3), and allowed designers to better interpret each other’s thought processes.

8.6 Opportunities for Future Studies

Participants expressed interest in exerting more control over the images suggested by the *ImageCascade* and the *AI Suggestions*. They would like to define their own categories (P4, P6), e.g. by selecting reference pictures in the mood board (P2) or by name (P4). This is similar, and an interesting complement, to the keywords control already available in the Image Search. However this would effectively limit the exploration capabilities of the AI system, which is one of its central features. Future work should investigate how to provide more control while keeping the benefits of the AI system. One participant (P3) contemplated the possibility and implications of giving a more prominent role to the AI system, including adding and deleting images directly in the mood board. “Like let’s say, if I don’t have time for Pierre, or Pierre doesn’t have time for me, at least I could go online and maybe the AI does this role”. P3 weighted the potential helpfulness of this augmented role against a ‘sorcerer’s apprentice’ scenario in which the *AI Suggestions* overloads the mood board with images. While contributing directly to the process is feasible, more research is needed to develop suitable reasoning and conflict management approaches between human and machines in creative processes. Participants also suggested adding input and output channels, from voice commands (P8) to interfacing the virtual mood board with a physical one (P5) and automatically extracting words from the design brief to use them in the search (P3). The latter could easily be

added to *ImageSense* by filtering out unimportant words and adding important ones as zero-priority keywords until they are activated by the user. One participant mentioned that “every image talks. So when I see another image [...] I’m switching from one idea to another” (P8). She suggested having a hardware setup with two screens, one with the mood board and the other to explore other ideas. While this is technically feasible, it also raises a larger issue of supporting visual inspiration without overloading visual perception.

Finally, one participant (P6) suggested that previous mood boards could also be accessible and searchable while working on a new one, and in particular be queried using the most common semantics of the current mood board to retrieve images and combinations of images used in previous work. This could easily be added to *ImageSense* since all images used in a mood board are saved for quick retrieval, along with their labels and associations.

8.7 System Limitations

Several participants commented on the size of images and text in the interface and would have liked to have an additional preview function (P1,P2,P6,P9). Some participants would like standard features of professional tools such as Photoshop, e.g. modifiable shapes (P1) and color pickers (P2). Some designers (2/9) commented on the differences between the physical and digital mood board design process and experience. P9 explained: with a physical mood board, “we actually printed many images and then posted them on a wall. While I did like it, this system [*ImageSense*] is much easier because we don’t need scissors or stickies and of course we can collaborate with the people who are away from here. [...] So this is a not a limitation but the difference between this experience.” (P9) Fashion designers mentioned that one “constraint would be if I would like to add physical materials” (P5), because textiles play a crucial role in this domain. However she continued, that “the positive thing about this [*ImageSense*] is that you can cooperate with the person that you’re working with, you can collaborate with multiple people [...] because you can’t do this in the physical environment.” This raises the question of how to incorporate tangible material into a digital mood board, while maintaining the right level of control.

9 DISCUSSION

The goal of this work is to find a suitable integration of intelligent tools into a collaborative mood board design process that offers designers a diverse set of inspirational sources and reflection material, while encouraging the collaborative nature of the process. Designers appreciated the kind of visual collaboration enabled by *ImageSense* and the new ways in which it supports inspiration. In addition, *ImageSense* was considered a very useful tool to communicate one’s design concepts.

9.1 Collecting Ideas with *ImageSense*

Our goal is to help designers identify interesting material by supporting the main sources of inspiration [39, 63, 83] through serendipitous browsing, mixed text and image search, and a proactive AI. Our results show that each tool was found useful for image retrieval, but in different contexts. *ImageSense* was mainly used to inspire completely new ideas, when designers struggled to express what they were missing on the board. Designers found the fluid visualization and control inspiring.

Semantic Search was mentioned as the most useful tool. It helps designers find what they are looking for, which was one of our design goals. Similar to previous work on semantics-supported searches [7, 42], designers found the semantics-based results more inspiring than usual search engines. Due to our shared agency approach, designers felt empowered to better search for hard-to-express ideas, as the system took the initial cognitive burden of defining what might be interesting in an image, while the designer could choose which images were most relevant in the current context. It was further appropriated by designers for investigating material that the human collaborator

had searched for. This traceability of others' thoughts was mentioned as a good way to see and align ideas among the collaborators.

The *AI Suggestions* were reported to be the most useful tool when feeling stuck or in need of new ideas. Designers found it crucial that unlike *ImageSense*, which presents unrelated images, the *AI Suggestions* were related to the current mood board. They were used throughout the process, from getting early ideas to finding alternative perspectives later in the process. Designers found the suggested material surprising but relevant and they appreciated its diversity, similar to that provided by a human participant.

All designers highly appreciated

the ability to choose which form of interaction made most sense for their particular needs at any given moment. This confirms Wang et al.'s criticism [83] that creativity support tools focus on specific perspectives without considering the overall process. Our work illustrates how a combination of tools can support the creative practice more holistically.

9.2 Supporting Reflection using Semantic Analysis and Abstractions

Our goal is to extend the action- and object-related perspectives of reflection support to more abstract tag clouds of relevant words. Designers said that *ImageSense* helped them to find more useful search terms to find new images and better articulate the meaning behind their mood board designs. The abstraction levels also allowed designers to extend their reflection beyond the current (descriptive) mood board content, which they found inspiring. Providing semantic and abstracted content offered designers new ways to express themselves, and some felt that the cloud visualizations worked as a useful communication medium that translates visual content into a textual form. However, further work is required to find the right levels of control to support shared agency in reflection clouds.

9.3 Collaboration with Human Partner

Participants felt that remote collaboration with a human design partner was a key benefit of *ImageSense*. Although they appreciated the intelligent tools in specific situations, designers enjoyed the overall experience of having natural discussions about concepts, image choice and meaning with another person. As shown in earlier work on in-person collaboration with mood boards [49] and brainstorming [1], designers found interacting with Pierre particularly useful for expanding the idea space. Our results highlight the benefits of integrating human-human collaboration into creativity support tools, where designers share both common and a personal ideation spaces, which remains rare in today's professional tools.

9.4 Comparing Human and AI-based Creativity Support

The study design lets us compare how designers interact both with a remote human designer and intelligent creativity support tools. We found that designers preferred discussions with a human collaborator when they want feedback about the quality of their ideas, and when they want to make sense of the mood board as a whole. Interaction between human designers works best when the designers already have words to express themselves, but they turn to intelligent tools when they lacked these words. Designers spent considerable amount of time looking for things: they might think something is missing, but do not know how to express it. Designers also clearly distinguish among the tools. They prefer the *AI suggestions* when they need an 'external' source of inspiration. Some felt that the system 'understood' them and made suitable suggestions in that particular context. *Semantic Search* and the *Semantic Tag Clouds* helped designers to reflect on their current set of images, and reduced the burden of interpretation by articulating words that capture the essence of the images. Communicating with the human collaborator was both engaging, and offered an

additional source of inspiration and clarification. Designers enjoy browsing through images from the *ImageSense* when they want to disconnect themselves from the current set of images on the mood board, and find serendipitous inspiration for a new direction of design ideas.

9.5 Limitations of the Study

Even though our confederate approach allowed us to observe how participants collaborate remotely with another designer on a real-world design task, and how they use each tool in a realistic ideation process, it does not offer strong quantitative measures, which is always a challenge for open-ended design tasks. That said, the clear next step is to assess *ImageSense* in the context of a longitudinal study with professional designers.

Interim tasks in the study followed the traditional design process, but were designed to enhance comparability within and across participants guided through the confederate. While this allowed us to compare results of tool use among designers at all phases within a short, but realistic ideation process, it limits the generalizability of our results to these short experiences. A future study is necessary to identify if a free discovery of the tool by designers without a confederate would evolve our initial results into stronger personal preferences over a longer time.

The goal of this study was not to provide a quantitative, post-hoc assessment of quantitative measures of performance, but rather to see which tools users need and when, and also whether they can move fluidly between interaction with remote human design partners and across different types of tools. Nevertheless, visual ideation has a strong impact on the following design process, and is embedded into an often tight project schedule. Analysing the novelty of the visual ideas based on the designers' presentation with external experts compared for using and not-using *ImageSense*, would have allowed us to make more quantitative impact measures on the design process itself. This however, was outside the scope of this work and presents potential for further analysis.

9.6 Larger Implications for this Work

Most research in creativity support tools emphasize how they support for one, or perhaps two aspects of the ideation process. Yet real-world design processes involve multiple designers, who use a combination of physical and digital tools. Their diverse needs change both at a global level, in the different phases of the design process, and in their minute-by-minute requirements within any design task, as they are inspired by a new image to follow a new path, or feel stuck and seek new inspiration. Many AI-based tools focus on providing users with suggestions, which is only useful if the designer is, at that moment, open to and in need of suggestions. Otherwise, it is simply a distraction or interruption.

We see large potential in semantic analysis for creativity support tools. The approach in this work presents how to successfully apply semantic analysis to ideation for explorative search, artificial agent support and reflection. The easy implementation and powerful results can spark new ideas about how to develop new sophisticated AI systems that combine feature-oriented analyses, like color values here, with larger contextual knowledge in the form of semantics.

Artificial agents become more integrated in creative and artistic contexts. Most work, however, focuses on either active participation of artificial agents in dyadic interactions, or help/support tools for collaborative interactions. Our goal was to evaluate the role of a diverse set of AI tools in the context of a human-human creative collaboration. Our results show how creativity can benefit from human-human-**agent** interactions, complementing rather than replacing the designers' roles. Our work shows that different artificial tools play different roles in the context of ideation, from each other and from human participants. We hence argue that the traditional notion of collaborator might be extended to artificial agents, which add their own value and their own needs to collaborative work and settings. This opens new opportunities for future work on collaborative creative systems.

10 CONCLUSION

We presented *ImageSense*, a collaborative, digital mood board design tool that supports collaboration between remote designers and offers intelligent tools for: serendipitous discovery of curated images from the *ImageCascade*, combined text- and image-based *Semantic search*, and intelligent *AI suggestions* to find new images. Our goal is to support human-computer partnership, allowing designers to share agency both with each other and with intelligent tools, creating diverse forms of collaboration that remain under the user's control. *ImageSense* uses state-of-the-art machine learning algorithms to provide active intelligent agents that contribute and adapt to the ideation process, under the designer's control. By enriching images with semantic labels, *ImageSense* facilitates both semantic visual search and reflection, and provides opportunities for serendipitous exploration.

Our research questions were to RQ1) explore new ways to seamlessly integrate contributions from *both* human and intelligent agents into a digital mood board; RQ2) identify the appropriate kinds of intelligent assistance for different types of ideation challenges; and RQ3) assess some of the differences in how designers interact with remote human collaborators and intelligent assistants, and how they can support each other.

We addressed the first goal by creating a digital mood board tool that supports the entire ideation process, from early collection of ideas, to layout of visual material, to reflection and final presentation to stakeholders. The system is fully collaborative: remote designers can work together in a shared space, but each retain their own individual spaces and access to design tools. They can see each other's work progress, and propose and react to the remote designer's images and reactions to their choices. *ImageSense* also incorporates three visual inspiration tools, each with a different type of agency. Designers control when to propose or seek suggestions or reactions, from both their human partners and intelligent agents. Designers can use these different tools without disrupting their creative flow. Most designers (8/9) reported that *ImageSense* helped them explore novel alternatives and to effectively communicate their designs to stakeholders. They also reported that *ImageSense* helped them reflect on both the composition and traceability of the process. Most agreed that *ImageSense* helped them to 'collect useful images' (6/9) and 'be more creative' (5/9).

To address the second and third research goals, our study identified situations under which participants preferred to work with a human designer or with each type of tool. Most designers said they would 'love to own' *ImageCascade* and *Semantic Search* when they were actively collecting images. *ImageCascade* was most useful when they needed to extend their ideas with images that are not directly related, but also not random. They particularly appreciated *Semantic Search* when they struggled to find search terms for hard-to-express visual ideas. Designers rated the *AI suggestions* as most suitable when they felt stuck or as an additional push to get started and extend their idea space, ignored it when they were browsing images with *ImageCascade*, or actively searching for images with *Semantic Search*. By contrast, designers found the human collaborator particularly helpful when they knew what they were looking for and for discussing content representation in the end. In each case, designers enjoyed being able to choose, minute-by-minute, the types of interactions they had with both human design partners and intelligent assistants.

When we consider how best to integrate intelligent systems into any creative process, we need to consider just what it means for the user to remain in control of the interaction. Clearly, this does not have a single answer, but depends greatly on the user's situation in the moment. When designers have a clear idea of what they want and how to express it, they require different types of interaction with intelligent tools than if they are completely at a loss and need surprising, but relevant new ideas. *ImageSense* demonstrates how to integrate different forms of intelligent assistance to create varied human-computer partnerships within the same collaborative, creativity support platform, and, we hope, serves as inspiration for future creativity support tools.

ACKNOWLEDGMENTS

We thank our colleagues, especially a big thank to Felix Epp and Ali Mughal, for their input on this project, and to all participants for their ideas described in this paper. 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) and CREATIV: Creating Co-Adaptive Human-Computer Partnerships (grant agreement No 321135).

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] Salvatore Andolina, Hendrik Schneider, Joel Chan, Khalil Klouche, Giulio Jacucci, and Steven Dow. 2017. Crowdbboard: augmenting in-person idea generation with real-time crowds. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. 106–118.
- [3] 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.
- [4] Brent Berlin and Paul Kay. 1991. *Basic color terms: Their universality and evolution*. Univ of California Press.
- [5] Dan Bloomberg. 2019. Color quantization using modified median cut. (03 2019), 6.
- [6] Susanne Bødker. 1991. Through the interface-A human activity approach to user interface design. *DAIMI Report Series* 224 (1991).
- [7] Carole Bouchard, Jean-francois Omhover, Celine Mougenot, Ameziane Aoussat, and Stephen J Westerman. 2008. TRENDS: a content-based information retrieval system for designers. In *Design Computing and Cognition'08*. Springer, 593–611.
- [8] Adrien Bousseau, Theophanis Tsandilas, Lora Oehlberg, and Wendy E. Mackay. 2016. How Novices Sketch and Prototype Hand-Fabricated Objects. In *Conference on Human Factors in Computing Systems (CHI)*. ACM, San Jose, United States. <https://doi.org/10.1145/2858036.2858159>
- [9] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [10] Search by Muzli. 2019. *Muzli* (2019). <https://search.muzli>
- [11] Hernan Casakin and Gabriela Goldschmidt. 1999. Expertise and the use of visual analogy: implications for design education. *Design studies* 20, 2 (1999), 153–175.
- [12] Tracy Cassidy. 2011. The mood board process modeled and understood as a qualitative design research tool. *Fashion Practice* 3, 2 (2011).
- [13] Tracy Diane Cassidy. 2008. Mood boards: Current practice in learning and teaching strategies and students' understanding of the process. *International Journal of Fashion Design* 1, 1 (2008), 43–54.
- [14] 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* (Singapore, Singapore) (*C&C '17*). ACM, New York, NY, USA, 93–105.
- [15] Jean Chevalier. 1996. *The Penguin dictionary of symbols*. (1996).
- [16] Mihaly Csikszentmihalyi and Keith Sawyer. 2014. Shifting the focus from individual to organizational creativity. In *The systems model of creativity*. Springer, 67–71.
- [17] 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* (Copenhagen, Denmark) (*SIGIR '92*). ACM, New York, NY, USA, 318–329.
- [18] 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*. ACM, 17.
- [19] Volodymyr Dziubak. 2019. Supporting graphic designers' online inspiration seeking practices via software tools. (2019).
- [20] 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.
- [21] 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* (Florence, Italy) (*CHI '08*).

- ACM, New York, NY, USA, 29–38.
- [22] Charles Freeman, Sara Marcketti, and Elena Karpova. 2017. Creativity of images: using digital consensual assessment to evaluate mood boards. *Fashion and Textiles* 4, 1 (2017), 17.
 - [23] Jonas Frich, Lindsay MacDonald Vermeulen, Christian Remy, Michael Mose Biskjaer, and Peter Dalsgaard. 2019. Mapping the Landscape of Creativity Support Tools in HCI. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, Scotland Uk) (CHI '19). ACM, New York, NY, USA, Article 389, 18 pages. <https://doi.org/10.1145/3290605.3300619>
 - [24] Alex Gabriel, Davy Monticolo, Mauricio Camargo, and Mario Bourgault. 2016. Creativity support systems: A systematic mapping study. *Thinking Skills and Creativity* 21 (2016), 109–122.
 - [25] John Gage. 1999. *Color and meaning: Art, science, and symbolism*. Univ of California Press.
 - [26] 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.
 - [27] 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* (Montreal QC, Canada) (CHI '18). ACM, New York, NY, USA, Article 121, 11 pages.
 - [28] Milene Gonçalves, Carlos Cardoso, and Petra Badke-Schaub. 2016. Inspiration choices that matter: the selection of external stimuli during ideation. *Design Science* 2 (2016).
 - [29] François Guimbretière, Maureen Stone, and Terry Winograd. 2001. Fluid interaction with high-resolution wall-size displays. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*. 21–30.
 - [30] Sonja Gutwein. 2013. Computer support for collaborative creativity. *Beyond the Desktop* (2013), 15.
 - [31] Joshua Hailpern, Erik Hinterbichler, Caryn Leppert, Damon Cook, and Brian P Bailey. 2007. TEAM STORM: demonstrating an interaction model for working with multiple ideas during creative group work. In *Proceedings of the 6th ACM SIGCHI Conference on Creativity & Cognition*. 193–202.
 - [32] 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.
 - [33] Otmar Hilliges, Lucia Terrenghi, Sebastian Boring, David Kim, Hendrik Richter, and Andreas Butz. 2007. Designing for collaborative creative problem solving. In *Proceedings of the 6th ACM SIGCHI conference on Creativity & cognition*. 137–146.
 - [34] Eric Horvitz. 1999. Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*. 159–166.
 - [35] Incredibleart.org. 2019. Color Symbolism and Culture. <https://www.incredibleart.org/lessons/middle/color2.htm>
 - [36] Nanna Inie and Peter Dalsgaard. 2017. A Typology of Design Ideas. In *Proceedings of the 2017 ACM SIGCHI Conference on Creativity and Cognition*. ACM, 393–406.
 - [37] Mikhail Jacob, Alexander Zook, and Brian Magerko. 2013. Viewpoints AI: Procedurally Representing and Reasoning about Gestures.. In *DiGRA Conference*.
 - [38] Ben Jonson. 2005. Design ideation: the conceptual sketch in the digital age. *Design studies* 26, 6 (2005), 613–624.
 - [39] AI Keller, GJ Pasman, and Pieter Jan Stappers. 2006. Collections designers keep: Collecting visual material for inspiration and reference. *CoDesign* 2, 01 (2006), 17–33.
 - [40] Gary Klein, Brian Moon, and Robert R Hoffman. 2006. Making sense of sensemaking 1: Alternative perspectives. *IEEE intelligent systems* 21, 4 (2006), 70–73.
 - [41] Janin Koch, Andrés Lucero, Lena Hegemann, and Antti Oulasvirta. 2019. May AI?: Design Ideation with Cooperative Contextual Bandits. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 633.
 - [42] Janin Koch, Nicolas Taffin, Andrés Lucero, and Wendy Mackay. 2020. SemanticCollage: Enriching Digital Mood Board Design with Semantic Labels. In *Proceedings of the Designing Interactive Systems Conference 2020 (DIS '20)*. ACM.
 - [43] John Kolko. 2010. Sensemaking and framing: A theoretical reflection on perspective in design synthesis. *Design Research Society* (2010).
 - [44] Rebecca Leis and Lauren Reinerman-Jones. 2015. Methodological implications of confederate use for experimentation in safety-critical domains. *Procedia Manufacturing* 3 (2015), 1233–1240.
 - [45] Antonios Liapis and Georgios N Yannakakis. 2016. Boosting computational creativity with human interaction in mixed-initiative co-creation tasks. In *Proceedings of the ICCG workshop on Computational Creativity and Games*, Vol. 310.
 - [46] Google LLC. 2018. Google Images Search. (2018). <https://images.google.com/>
 - [47] Andrés Lucero. 2012. Framing, aligning, paradoxing, abstracting, and directing: how design mood boards work. In *Proceedings of the designing interactive systems conference*. ACM, 438–447.
 - [48] Andrés Lucero. 2015. Funky-Design-Spaces: Interactive Environments for Creativity Inspired by Observing Designers Making Mood Boards. In *Human-Computer Interaction*. Springer, 474–492.
 - [49] Andrés Lucero, Dima Aliakseyeu, and Jean-Bernard Martens. 2007. Augmenting mood boards: flexible and intuitive interaction in the context of the design studio. In *Second Annual IEEE International Workshop on Horizontal Interactive*

- Human-Computer Systems (TABLETOP'07)*. IEEE, 147–154.
- [50] 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*. 11–18.
 - [51] Lena Mamykina, Linda Candy, and Ernest Edmonds. 2002. Collaborative creativity. *Commun. ACM* 45, 10 (2002), 96–99.
 - [52] Ine Mols, Elise Van den Hoven, and Berry Eggen. 2016. Informing design for reflection: an overview of current everyday practices. In *Proceedings of the 9th Nordic Conference on Human-Computer Interaction*. ACM, 21.
 - [53] Philippaa MOTHERSILL and V. Michaela BOVE. 2018. An Ontology of Computational Tools for Design Activities. In *Proceedings of Design Research (DRS)*. DRS.
 - [54] Jaap MJ Murre and Joeri Dros. 2015. Replication and analysis of Ebbinghaus' forgetting curve. *PloS one* 10, 7 (2015).
 - [55] 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.
 - [56] Alex F Osborn. 1953. Applied imagination. *Charles Scribner, New York* (1953).
 - [57] Rajan Patel. 2018. World around you with Google Lens and the Assistant. (2018). <https://www.blog.google/products/google-lens/google-lens-real-time-answers-questions-about-world-around-you/>
 - [58] Sharoda A Paul and Meredith Ringel Morris. 2011. Sensemaking in collaborative web search. *Human-Computer Interaction* 26, 1-2 (2011), 72–122.
 - [59] Charles Sanders Peirce. 1974. *Collected papers of charles sanders peirce*. Vol. 5. Harvard University Press.
 - [60] Unsplash: Beautiful Free Images & Pictures. 2019. *Unsplash* (2019). <https://unsplash.com/>
 - [61] Pinterest. 2019. *Pinterest* (2019). <https://www.pinterest.com>
 - [62] Peter Pirolli and Stuart Card. 2005. The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of international conference on intelligence analysis*, Vol. 5. McLean, VA, USA, 2–4.
 - [63] Martin Porcheron, Andrés Lucero, and Joel E. Fischer. 2016. Co-Curator: Designing for Mobile Ideation in Groups. In *Proceedings of the 20th International Academic Mindtrek Conference (Tampere, Finland) (AcademicMindtrek '16)*. Association for Computing Machinery, New York, NY, USA, 226–234. <https://doi.org/10.1145/2994310.2994350>
 - [64] Image Finder High quality free stock images. 2019. *ImageFinder* (2019). <https://imagefinder.co/>
 - [65] Guillermo Rauch. 2019. Socket. io. *Socket IO.[Online]* <http://socket.io> (2019).
 - [66] Paul Resnick and Hal R. Varian. 1997. Recommender Systems. *Commun. ACM* 40, 3 (March 1997), 56–58. <https://doi.org/10.1145/245108.245121>
 - [67] Donald A Schön. 2017. *The reflective practitioner: How professionals think in action*. Routledge.
 - [68] Moushumi Sharmin and Brian P Bailey. 2011. I reflect to improve my design: investigating the role and process of reflection in creative design. In *Proceedings of the 8th ACM conference on Creativity and cognition*. ACM, 389–390.
 - [69] Ben Shneiderman. 2000. Creating creativity: user interfaces for supporting innovation. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7, 1 (2000), 114–138.
 - [70] Ben Shneiderman. 2007. Creativity support tools: Accelerating discovery and innovation. *Commun. ACM* 50, 12 (2007), 20–32.
 - [71] Pao Siangliulue, Joel Chan, Steven P Dow, and Krzysztof Z Gajos. 2016. IdeaHound: improving large-scale collaborative ideation with crowd-powered real-time semantic modeling. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*. 609–624.
 - [72] 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 (Glasgow, United Kingdom) (C&C '15)*. ACM, New York, NY, USA, 83–92.
 - [73] Shaun Singh. 2019. Death to Stock - When stock dies art thrives. <https://deathtothestockphoto.com/>
 - [74] Kyle Steinfeld. 2017. Dreams may come. In *Acadia*. 590–599.
 - [75] Timo Strohmman, Dominik Siemon, and Susanne Robra-Bissantz. 2017. brAInstorm: Intelligent Assistance in Group Idea Generation. In *International Conference on Design Science Research in Information System and Technology*. Springer, 457–461.
 - [76] Lucy A Suchman. 2007. *Human-machine reconfigurations: Plans and situated actions*. Cambridge university press.
 - [77] Hillevi Sundholm, Henrik Artman, and Robert Ramberg. 2004. Backdoor Creativity: Collaborative Creativity in Technology Supported Teams.. In *COOP*. 99–114.
 - [78] 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.
 - [79] Belinda Thom. 2000. BoB: an interactive improvisational music companion. In *Proceedings of the fourth international conference on Autonomous agents*. ACM, 309–316.

- [80] Inc. Twinword. 2018. Natural Language Processing APIs - Twinword API. <https://www.twinword.com/api/word-associations.php>
- [81] Peter-Paul Verbeek. 2015. COVER STORY Beyond interaction: a short introduction to mediation theory. *interactions* 22, 3 (2015), 26–31.
- [82] Hao-Chuan Wang, Dan Cosley, and Susan R Fussell. 2010. Idea expander: supporting group brainstorming with conversationally triggered visual thinking stimuli. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*. 103–106.
- [83] Kai Wang and Jeffrey V Nickerson. 2017. A literature review on individual creativity support systems. *Computers in Human Behavior* 74 (2017), 139–151.
- [84] Andrew M Webb, Rhema Linder, Andruid Kerne, Nic Lupfer, Yin Qu, Bryant Poffenberger, and Colton Revia. 2013. Promoting reflection and interpretation in education: Curating rich bookmarks as information composition. In *Proceedings of the 9th ACM Conference on Creativity & Cognition*. ACM, 53–62.
- [85] 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.
- [86] Georgios N Yannakakis, Antonios Liapis, and Constantine Alexopoulos. 2014. Mixed-initiative co-creativity.. In *FDG*.

Received October 2019; revised January 2020; accepted March 2020