SemanticCollage: Enriching Digital Mood Board Design with Semantic Labels

Janin Koch, Nicolas Taffin, Andrés Lucero, Wendy E. Mackay

► To cite this version:

Janin Koch, Nicolas Taffin, Andrés Lucero, Wendy E. Mackay. SemanticCollage: Enriching Digital Mood Board Design with Semantic Labels. DIS '20 - Designing Interactive Systems Conference 2020, Jul 2020, Eindhoven, Netherlands. pp.407-418, 10.1145/3357236.3395494 . hal-03047463

HAL Id: hal-03047463

<https://hal.inria.fr/hal-03047463>

Submitted on 8 Dec 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.

SemanticCollage: Enriching Digital Mood Board Design with Semantic Labels

Janin Koch
Aalto University
Espoo, Finland

Nicolas Taffin
Univ. Paris-Saclay,
Inria, CNRS
Saclay, France

Andrés Lucero
Aalto University
Espoo, Finland

Wendy E. Mackay
Univ. Paris-Saclay,
Inria, CNRS
Saclay, France

Figure 1: *SemanticCollage*: a semantically enriched digital mood board tool for image collection and interpretation. Blue tools support image and text manipulation; Red tools provide semantic labels.

ABSTRACT

Designers create inspirational mood boards to express their design ideas visually, through collages of images and text. They find appropriate images and reflect on them as they explore emergent design concepts. After presenting the results of a participatory design workshop and a survey of professional designers, we introduce *SemanticCollage*, a digital mood board tool that attaches semantic labels to images by applying a state-of-the-art semantic labeling algorithm. *SemanticCollage* helps designers to 1) translate vague, visual ideas into search terms; 2) make better sense of and communicate their designs; while 3) not disrupting their creative flow. A structured observation with 12 professional designers demonstrated how semantic labels help designers successfully guide image search and find relevant words that articulate their abstract, visual ideas. We conclude by discussing how *SemanticCollage* inspires new uses of semantic labels for supporting creative practice.

Author Keywords

Creativity support tools; Sensemaking; Ideation; Semantics

CCS Concepts

•Human-centered computing → Interactive systems and tools; •Applied computing → Arts and humanities;

INTRODUCTION

Designers often express and explore visual ideas through *mood boards*, which are collages composed of images, text, and object samples. Mood boards act as a source of inspiration in creative fields such as fashion and design, and are especially helpful when ideas are hard to express verbally [9]. Mood boards are not only innovative and fun to create [21], but also encourage designers to probe more deeply into the project's themes and concepts, serving as a powerful tool for communicating a “web of seemingly unconnected ideas, difficult to express verbally with similar impact” [19].

The process begins when the client expresses rough ideas for a product or service. The designer then engages in a highly dynamic and iterative [21] process of first *collecting* visual inspirational material, then *composing* the mood board by selecting and arranging these images, while constantly *interpreting* the material [45] by identifying missing images [32] and trying to “to understand connections, [...] anticipate their trajectories and act effectively.” [25] The final mood board communicates visual ideas to the client or other stakeholders [30, 8, 21]. The intrinsic visual nature of mood boards encourages creation of new ideas [48], with wide potential for innovative discovery and problem solving [21].

However, educators stress that visual abstraction, or finding the right images and interpreting their meaning, is a critical but difficult-to-learn skill [9] usually requiring professional design education [7, 9], or years of professional experience. Visual abstraction lets designers translate abstract ideas into structured visual representations [7]. Unfortunately, finding images that reflect abstract ideas is challenging and designers often spend significant time searching for the ‘right’ image. Today, designers turn to image search engines, e.g. *Google*, or curated inspirational platforms, such as *Behance* or *Pinterest*. However, most search engines only support text queries, forcing designers to find relevant, searchable words that capture each visual abstraction. Although *Google Reverse Image Search* lets users use a picture to find related images from the web, it offers little control over unsuitable results.

Another key challenge is why certain selected images are interesting: Designers must articulate their ideas, to themselves or others, or synthesize larger concepts, based on their collected material and experience [27]. This sense-making process is crucial for reflection on new concepts. Unfortunately, current tools offer little support for reflecting upon visual material.

Research Questions

We are interested in helping designers:

1. express and explore vague, not-yet-developed ideas for retrieving inspirational material;
2. make sense of and reflect upon their chosen material and communicate it to stakeholders; and
3. benefit from computational support without disrupting their creative flow.

One approach is to take advantage of advances in computer vision and machine learning, which can extract semantic information from images. Yet, as Steinfield argues, if “computers cannot see the way we see, they cannot help us to reason the way we wish to reason” [48]. This paper explores how to create digital mood boards that benefit from semantic labels.

After discussing related work, we describe the results of two preliminary studies that explored professional designers’ everyday experiences in creating mood boards. We then introduce *SemanticCollage*, a digital mood board that extracts semantic labels from images, and uses them to support designers’ search and reflection activities. We describe the system with an illustrated scenario, followed by a detailed technical description. Next, we describe the results of a structured observation with 12 professional designers who performed two pairs of comparable, ecologically grounded design tasks related to composition and reflection. We conclude with a discussion of the implications for design and directions for future research.

RELATED WORK

The mood board design process involves three key activities: collecting material, constructing the board and reflecting upon it [30]. Here, we focus on studies and tools related to the collection and reflection processes.

Collection

Designers alternate between ‘exploration’ and ‘exploitation’ [39]: They start with a known ‘anchor point’, e.g. objects

or associations, and refine the search, step by step, to narrow down the possibilities until they reach the desired result [50].

Text-Based Search is the most common strategy for finding images in large online collections, such as *Google Image Search* [28]. Others, such as *ImageFinder* [40], *Unsplash* [37], *Muzli Search* [5] and *Pinterest* [38], provide vast collections of curated images. Machine learning often supports image-search queries with user-specified rules [18], pre-defined preferences, e.g. colors or patterns [17], or via user-specified [11] and dynamic [53] clustering. By contrast, Bouchard et al. [2] show how to retrieve images that better fit search queries semantically: A textual analysis of the image’s web page produces semantic labels, and the resulting images are selected based on how well their semantics match the query. Designers found semantic-based search results more inspiring than those from standard image search engines. However, while promising, such systems still require designers to formulate precise text queries or visual specifications. Verbalizing vague, visual ideas is difficult, which limits the potential of this approach.

Image-Based Search has benefited from recent advances in computer vision and machine learning. *Google Image Search* uses *Google Lens* [35] for ‘reverse image retrieval’: When the user drags an image into the search bar, an image recognition algorithm uses neural networks to translate it into semantic labels [28]. These reference the image’s original web page, if applicable, as well as images deemed similar to the primary object. *Pinterest* also lets users select parts of a Pinterest image to suggest similar objects. Designers can highlight an object’s importance by specifying three to five semantic labels. *Nice* [34] lets users search for images via user-provided tags or dominant colors from a previously selected image. Image-based search reduces the need for finding the ‘right’ search terms, but current systems only handle one image at a time, and users cannot control which features are deemed important. We argue that users should be able to search for interesting features from any combination of images and text.

Interpretation

Sense making is an immersive process that involves discovery and learning, also known as ‘reflection-in-action’ [44]. Mood board designers synthesize ideas and concepts based on collected material, contextual understanding and their own experience [27]. Russell et al. [43] explore sense making with respect to large document collections, while Klein et al. [25] consider the reciprocal interaction between the user’s envisioned material and the material itself. Making sense of visual material helps designers elaborate, reflect on and question their current vision. Reflection methods that encourage such behavior in non-digital settings focus on filtering, generalizing and sharing meaning for idea convergence [14, 46].

Interpretation Tools help designers articulate the meaning of their work. For example, *CoSense* [36] is a web search tool that encourages sense making among participants using search summaries in the form of search-term tag clouds. The *Rich bookmarks* [52] interface lets users create links to visual and semantic metadata to reflect on collected material. These approaches share the goal of grouping materials together, where links and images highlight internal relationships,

e.g. tag clouds. Although designers can record annotations, these approaches do not support the reflection process itself. The key challenge in visual ideation is how to provide reflective material that helps designers make sense of their images.

Schön's [45] concept of reflection-on-action refers to post-design processes that revisit earlier design processes, decisions, and activities. Tools such as *ReflectionSpace* [47] and *Maps for design reflection* [13, 12] let designers visualize their ideation activities over time, both to reflect upon earlier alternatives, and to develop a holistic view of the whole design process. This suggests that capturing the semantic meanings of images, especially given how they evolve over time, may offer new insights to mood board designers.

Creativity Support Tools

Creativity support tools focus mainly on finding new material [51] by retrieving previously searched material [16]; suggesting related material [2]; encouraging collaboration [22, 49]; or combining existing physical and digital material [24]. These systems offer a wide variety of solutions for collecting relevant material. However, they do not fully support the expression of vague, visual ideas via text-based search terms to find relevant images, nor do they help users articulate the meaning of their mood boards.

In design practice, professional mood board designers rely primarily on rich design tools, e.g. *Adobe Illustrator*, to compose their material and construct meaning. However, mood board design requires iterative exploration and construction of meaning through structure and search. We thus apply Lupfer et al.'s free-form curation concept, which enables "elements to be spontaneously gathered from the web [...], manipulated, and visually assembled in a continuous space" [33], thus encouraging the evolution of ideas and semantic relationships.

INFORMING THE DESIGN

The research literature suggests that mood board design is an iterative, cyclic process that involves collecting images from magazines and online sources, and reflecting upon these on a mood board. To gain a more nuanced understanding of the mood board design process and inspire ideas for tools that avoid disrupting creative flow, we conducted two preliminary studies: 1. a participatory design workshop to better understand the mood board design process, and 2. a survey about digital mood board design with professional designers.

Participatory Design Workshop

The participatory design workshop (Fig. 2) included two groups of three designers, each with an author and at least one practicing mood board designer. The three-hour workshop began with a brief mood board presentation based on [9, 19]. Each group chose an initial topic, then looked for images, negotiated their selection, and arranged the results on a mood board. We provided foamcore boards, diverse magazines, and standard paper prototyping supplies. Each group presented their work and reflected on the overall process. We video recorded the workshop and took hand-written notes.

Results

One group explored a 'stretchable materials store'; the other created a 'sustainability lab'. Both spontaneously followed

Figure 2: Each group created a physical mood board

a standard design process [8, 31], which involves *collecting* material and *interpreting* it during the construction of the board, each with different tools and communication styles.

Collection: Participants sought images that 'fit' the group's evolving narrative, but were happy to discover surprises and happy accidents. Some images 'matched' the concept "colorwise" or because it was "the image I looked for". Others were chosen even though the reason was hard to express in words: "I don't know – I just like this image". Some material was selected for its general appearance: "This is really cute" or "I don't know what to do with it but I will cut it out anyway". We noticed that they often lacked adequate words to express why they chose particular images. After selecting their material, both groups laid out the images, often stacking or clustering them thematically, e.g. saying: "These are professionals!" then adding them to the mood board.

Interpretation: Participants cropped and adjusted images, then clustered and separated them and discussed their meaning: "This is something like a digital texture". They arranged them to support the group's narrative: "I'm going to cut it so it looks like a tool tree" or "This image connects both these images very nicely." As new ideas arose, participants sought additional images. One group removed and stacked their images, then reassembled them in a new layout, adding new concepts, deleting others. Both groups added notes and labels for the final presentation.

Survey of Professional Mood Board Designers

We conducted a digital survey to gain insights into challenges faced by professional mood board designers [41].

Participants: 15 professional designers from Europe and North America (8 men, 7 women, age 23–45) responded to the survey. Professional experience ranged from 1 to 15 years (mean 5.8), primarily in User Experience, Graphic and Interaction design. Responses were anonymous, and participants were not paid.

Survey instrument: Questions (27 open-ended, 5 Likert-style, and 5 multiple-choice) focused on search and sense-making strategies. The first nine revisited their own recent project; the rest asked them to analyze an early- and a late-stage mood board and say how they would continue, given the design brief.

Data Analysis: We ran a mixed-approach thematic analysis [4]. Top-down themes included 'collection', 'composition' and 'reflection' from the research literature. One author coded the participants' answers based on both top-down themes and

themes that emerged (bottom-up), with particular focus on themes related to our research questions. The same author summarized closely related themes and identified patterns; a second author reviewed the themes. We also calculated descriptive statistics for answers to Likert-scale questions.

Results

Most designers (9/15) use mood boards for at least half their projects; the majority (12/15) said they only create digital mood boards. A common reason for creating digital mood boards (5/15) was “sorting my thoughts and feelings relating to the project”. Designers said that mood boards help them explore different styles (P8, P12) and communicate them through visual material (6/15). Like the participatory design groups, participants follow the standard mood board design process, i.e. collecting images, designs and stories (P7), adding details e.g. “defining words” or “sorting and composing” (P11, P14, P9). Respondents use a variety of professional tools, including Adobe (5/15), mainly Illustrator (3/5), Sketch/ InVision (4/15), online Platforms such as Pinterest/ Behance (3/15), and dedicated tools such as Arena (P2) or Milanote (P9).

Collecting Material: Designers find inspirational material via search engines (9/15), platforms such as Behance/Pinterest (3/15) or magazines (3/15). However, they mentioned that “sometimes it is painful to find the right search keyword to find the image” (P7) and it can take minutes (6/15) to hours (5/15) to find it. They describe this as an iterative process: Even though the first image “was found in half an hour, it didn’t depict everything I wanted, so I kept looking for three more hours” (P9) to find the right image.

Finding the right initial image was often challenging. One designer wanted to “search images by color on Dribbble and Google, [and] use the first image as input for Google image search, but also look for new images by just browsing content on Behance to add some versatility to the mood board” (P5). Another would “jot down what is still missing [and] search with these keywords and with some research on what these keywords mean in the target group’s lives” (P9). Both designers actively avoided defining a query (1) or searching for inspiration. Some select suitable material based on color (P5, P13) or “emotions an image evokes, especially regarding senses but also associations and yet again semantic meanings of objects, materials and symbols” (P9).

Implications for Design

Transforming visual ideas into textual search terms: Workshop participants had difficulty articulating why particular images were ‘right’, and survey designers had problems expressing visual ideas as text search terms. *Design implication:* Help designers transform visual ideas into effective search terms.

Alternating visual and text-based search: In both studies, designers spent significant time searching for images. Survey designers described first browsing digital inspirational platforms for new images and then searching; searching for initial relevant keywords; or using inspirational images to begin a search. *Design implication:* Let designers use any combination of image- and text-based material as search expressions.

Guiding image search: Even if designers chose an appropriate initial query, they still struggled to guide the remaining search, since most image platforms return highly similar or duplicate images. They needed ‘anchor points’ or in-between steps to refine the search towards the desired result [50]. *Design implication:* Let designers specify either visual or textual anchor points to guide the search.

Reflecting on visual ideas: Workshop participants combined images into larger abstract concepts, which they then discussed. One group questioned and re-evaluated their whole mood board and launched a new search process. Survey designers also reported using their mood boards to reflect on their ideas and evoke emotions. All designers valued mood boards as a means of communicating their ideas. *Design implication:* Help designers reflect on abstract visual concepts by suggesting meanings for relevant clusters of images.

SEMANTICCOLLAGE SCENARIO

The following scenario illustrates how Tom, a designer, would use *SemanticCollage* to develop a new concept for a ‘coffice’ that combines a Café atmosphere with an office space. Tom starts with an empty *SemanticCollage* canvas for his mood board, a search panel that accepts both images and text as tool, and a ‘maybe’ stack for potentially useful images (Fig. 3).

Figure 3: *SemanticCollage* interface: Tool panel, Mood board canvas, and ‘Maybe’ stack.

Tom writes “cozy looking coffeehouses” into the search bar and scrolls through the results to see more images. He drags interesting ones to the main canvas. Next, he searches for the phrase ‘office spaces’, but finds the results rather ‘cold’, not really what he was looking for. He decides to drag a ‘cozy-looking’ coffeehouse image into the search field. When it lights up yellow, he knows it has been registered and added to the existing search term (Fig. 4.1).

SemanticCollage displays semantic labels below the image in the search bar (Fig. 4.2), with ‘+’ and ‘-’ buttons that let him adjust their influence. Tom removes ‘room’ as irrelevant, then presses ‘enter’ to search for the revised combination of his initial phrase and relevant semantics associated with his ‘cozy’ image. He drags several interesting possibilities to the canvas.

Tom forms clusters of images that represent both concepts and contrasts he is contemplating. When he hovers over an image, a menu with various tools appears (Fig. 1.9 & 1.11). He cuts and recolors several images, then clicks on the ‘Reflection’ tab to see a tag cloud of the current semantic labels and a palette of the images’ colors (Fig. 5). The first cloud (Fig. 1.5) displays the full set of semantic labels in the mood board. The size of

Figure 4: Mixed Media Search: 1) The search field accepts images. 2) Semantic labels related to the image appear and their relevance can be adjusted.

each word is dictated by the number of images with that label, and the sizes of those images in the mood board.

Tom selects words he finds surprising, which highlights the images with the corresponding semantic label. He re-positions several images, but still feels something is missing. He shift-clicks on several images of ‘cozy coffeehouses’ to create a group of images. He then generates a new tag cloud by clicking on the ‘common labels’ button, which extracts semantic labels as if it were a single image (Fig. 1.6). This cloud contains new semantics that do not appear in any of the original images. He realizes that “architectural styles” is an implicit trait the images share. He searches again, looking for ‘cozy’ business-related architectural styles. Tom adds a color palette and text that explains key elements of the final collage (Fig. 5).

Figure 5: Tom’s mood board: Tag clouds show semantic labels. The color palette is derived from the images.

SYSTEM AND USER INTERFACE

SemanticCollage is an easy-to-use semantically enhanced tool for creating mood boards. It uses *HTML/JS/Jquery* (front-end), *Python* (back-end) connected to a *Postgres* database and *Web.py* as a web-framework. It is lightweight and can be deployed on any local computer¹. The interface contains three main areas: a tool panel, a mood board canvas, and a ‘maybe’ stack for keeping images (Fig. 3). Designers can switch between the collection and interpretation panels using the tabs on the left.

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

The mood board construction tools are integrated into the canvas and can be accessed by the designer at any time (Fig. 6). Unlike existing tools, *SemanticCollage* connects all the design phases interactively. Designers can drag searched images directly to the mood board. All images can also be used as search terms (Fig. 4). Every change to the mood board affects the reflection tools in real time. *SemanticCollage*’s two main panels are called ‘Compose’ and ‘Reflect’.

Semantic Labels

Semantic labels play an important role in *SemanticCollage*, whether they are visible or hidden from the user. The system uses a two-stage retrieval process: 1) an object recognition algorithm detects image attributes, e.g. objects, faces or text, and passes them to 2) a classification algorithm that maps these attributes to classes in an ontology. Training these machine-learning classifiers requires a large quantity of labeled data and is sensitive to domain-dependent patterns that may result in incorrect classification [42]. Since our goal is to let designers use any visual material, either from online sources or their own uploaded material, we apply a state-of-the-art object recognition and semantic labeling algorithm provided by Google Vision [29]. Because it is trained with several thousand labels from diverse contexts, we can retrieve semantic labels quickly, without being limited to specific domains or sources.

We retrieve new semantic labels every time an image is cropped or dragged onto the canvas, the ‘maybe’ stack or the upload area. For each image, we select the ten labels retrieved from the Vision API with the highest confidence scores. Labels are linked to an image and assigned an integral weight based on their confidence scores: the label with the highest score receives a weight of 2 (most relevant), the next two labels 1 (relevant), and the other labels 0 (not relevant). Designers can modify these relevance levels by hand and add labels using the *Inspector* menu attached to every image (Fig. 1.4).

Collection

Semantic Search

We chose *DuckDuckGo* [15], an open source search engine, in order to release the *SemanticCollage* code to the research community. We appreciate its limited tracking behavior, which lets us isolate the impact of using semantic labels from the ‘learning’ of the search algorithm. The result quality of *DuckDuckGo* was perceived as sufficient in pre-studies, unlike the large differences in semantic labeling algorithms we reviewed. However, other search engines can also be used.

Designers can use text and images as search objects. Images can be dragged into the search area. They are then visually added to the search bar and their semantic labels are displayed underneath (Fig. 4). If more than one image is added, the order of the labels follows the order of the images in the search field. Designers can modify the weight of these search labels with the ‘+’ and ‘-’ buttons next to them.

The search term reflects the order of the searched elements, text or images, and sorts the semantic labels for each image by relevance, excluding irrelevant ones. Search results are displayed as a scrollable list of 30-50 images, and images can be dragged to the canvas, ‘maybe’ stack or search bar directly.

Images added to the canvas are stored in full resolution for later use in the local file system. Each search query creates a new *Search History* item displayed underneath the search bar in a scrollable list containing text and image information. Designers can reuse previously searched terms by clicking an item in this list, which is then added to the search bar.

Uploading Material

Designers frequently use previous work or photos as sources of inspiration. In *SemanticCollage*, designers can upload their own material (Fig. 1.3), which is then made available in the ‘maybe’ area. Semantic labels are queried and attached to the uploaded image and can be used in the design process.

Interpretation

During the design process, designers interpret images and groups of images to make sense of larger concepts. *SemanticCollage* supports this by providing overview clouds, requested clouds and extracted color samples (Fig. 1).

Semantic Clouds

SemanticCollage’s interactive semantic clouds display labels for all mood board images (Fig. 1.5). We present two tag clouds: semantic frequency, and semantic frequency combined with dominance in the mood board. The first cloud ties word size to the number of images being labeled. The second cloud adds weights representing the sizes of those images in the mood board. Resizing an image dynamically updates the cloud. Clicking labels within a cloud highlights all images on the mood board that share that semantic label.

Requested Semantic Clouds

Designers can shift-click to select multiple images to retrieve the semantic labels for the whole group, which is then treated as a single image. The ‘common labels’ button takes a screenshot of the selected images, preserving their location, size and order, and requests new semantic labels. The resulting labels appear in the field below the button, ordered by confidence score (Fig. 1.6). Treating a group of images as a whole often produces semantic labels that diverge from the individual semantics, making it easier to understand their similarities.

Color Palette

Since color plays an important role in visual inspiration, *SemanticCollage* uses the MMCQ algorithm² [1] to analyze each image for its ten major colors, which are then added to the color palette (Fig. 1.7). Designers can select non-aggregated color samples for image manipulation or color patches (Fig. 5).

Mood Board Construction Tools

Designers require basic image manipulation tools to successfully design a professional-looking mood board, e.g. to modify size, color, and orientation (Fig. 1.8–11). For easier access, these tools surround the image currently selected (Fig. 6).

Editing Tools

Designers can freely drag and scale images on the canvas, as well as change the z-order (front and back), rotate or flip images horizontally or vertically (crop tool) (Fig. 6). Selecting the Color item opens a color editor with a color wheel

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

Figure 6: Image and graphical Menus

and color picker that uses *CSS* filters (Fig. 1.11) to modify dominant image colors. Finally, designers can crop images to highlight important elements (Fig. 1.9), which then receive new, locally saved semantic labels.

Adding Graphical, Textual Elements

Double clicking on the canvas opens a menu (Fig. 6, right) that lets the user add shapes (circles, rectangles), text elements and sticky notes. Objects and images share the same editing tools. The text editor supports changing text color, weight, font and size (Fig. 1.10).

Maybe Stack

Designers can store images that they are unsure of in the ‘maybe’ stack, a scrollable list of images located below the mood board canvas (Fig. 3). Images can be dragged in or out of this area, and are greyed out to minimize distraction from the canvas. When an image is selected on the mood board canvas, a filter is applied to the maybe stack that highlights images with similar semantics.

STUDY METHOD

We ran a study to investigate the potential of semantic labels for mood board ideation. We wanted to see if *SemanticCollage* helps designers 1) find appropriate images, even if they cannot clearly articulate what they should be; 2) make sense of the images they gather and communicate their meaning to an external stakeholder; and 3) allow them to maintain control of intelligent suggestions, without disturbing their creative flow.

We conducted a structured observation [20, 3] with 12 professionally trained designers. This type of quasi-experiment [6] is designed to enhance ecological validity, by combining controlled conditions derived from empirical observations, thus facilitating comparisons within and across real-world tasks.

Participants

We recruited 12 trained designers (6 women, 6 men; age 24–40) with 1–16 years of professional experience (Table 1). All provided informed consent and agreed to the recording of the session and anonymized publication of the results. European privacy law (GDPR) was followed throughout.

Setup

Participants sit at a desk with a Macintosh laptop, one external monitor, a mouse and keyboard. The experimenter launches *SemanticCollage*, which appears on the large screen,

ID	Age	Sex	Area of design	Design Practice
1	31	M	Interaction Design, Digital Media	3 yrs.
2	40	M	Industrial Design	6 yrs.
3	33	M	Interaction Design	4 yrs.
4	39	F	Web, fine arts, interaction	6 yrs.
5	29	F	Industrial, product, STS	2 yrs.
6	28	F	Fine arts, IT product design	5 yrs.
7	31	F	Industrial, interaction	10 yrs.
8	26	M	Architecture	2 yrs.
9	40	F	Graphic Design, Arts School	16 yrs.
10	25	M	Information Systems	1 yrs.
11	33	M	Architecture and information systems	9 yrs.
12	24	F	Interior design	1 yrs.

Table 1: Professional designers backgrounds and experience.

and changes the tool state according to each of the five conditions. The introduction videos for each condition, as well as the post-condition questionnaires, are displayed on the laptop and are launched by the participants.

Procedure

The experiment takes approximately one hour, and includes two controlled tasks to evaluate *Compose* and *Reflect* (described below), followed by an open task. The first two tasks each involve two scenarios, with separate but equivalent design briefs. Each task is presented with and without semantic labels resulting in four controlled conditions. We used a within-participant design, counter-balanced for order, so that each participant experienced all four conditions. In the final, open-ended task, the designer can use any of *SemanticCollage*'s features to create a new concept.

Protocol: Each condition begins with a video describing relevant *SemanticCollage* functionality. Participants read a card with a short description of the brief, goals and time available. After completing the tasks, they present their mood boards to a potential client, played by the experimenter, and complete a questionnaire. The experimenter conducts a final semi-structured interview to probe for details about the participant's experience. Participants are encouraged to reflect upon *SemanticCollage*'s strengths and weaknesses, and think how they might use it in their own work.

Composition tasks: Participants are asked to create two new mood boards that express a new visual concept for a 'coffice store' and a 'MEETbar'. The 'coffice store' combines the concepts of a cozy Café and an office space, where employees engage in formal and informal meetings in a cozy environment. The 'MEETbar' is a place where teams can meet after work, with the casual atmosphere of a bar. For each brief, participants have *seven minutes* to find and compose appropriate images into a mood board to be presented later.

Reflection tasks: Participants prepare two mood boards that present a concept for a high-end shoe store and a children's fashion store. To save time, participants are told to start with a concept previously prepared by a colleague who could not attend the meeting. Participants have *five minutes* to reflect upon, modify, and present the final mood boards.

Open task: Participants were asked to create a new visual sales concept for the upcoming summer holiday season, and retrieved more details about target audience and context. They had ten minutes to build a mood board with enabled semantic labels, and present the final mood board.

Semantic Label conditions: The *Compose* and *Reflect* tasks both involved two separate design tasks, with and without semantic labels. The semantic conditions include *SemanticCollage*'s tools for revealing and editing semantic labels (blue features in Fig. 1). The non-semantic conditions hide this functionality, but still permit text and image queries. Note that even without semantic labels, this version of *SemanticCollage* exceeds the capabilities of, e.g., Google Image Search or Pinterest, making it a more fair comparison to assess the benefits of semantic labels.

Questionnaire: Participants filled out two sections:

1. Creativity Support Index [10], which measures 1) collaboration, 2) enjoyment, 3) exploration, 4) expressiveness, 5) immersion, and 6) worthiness of effort; and
2. Five Likert-style questions that compare participants preferences with respect to the semantic and non-semantic conditions for each scenario.

Data Collection and Analysis

We collected audio and screen recordings, questionnaires, and hand-written notes. We ran a mixed-approach thematic analysis, with top-down themes, *Collection* and *Interpretation* from the literature; as well as use of *Semantic Labels*, the main contribution of the system, and overall *Usability*. The coding and verification process matched that of the earlier studies.

RESULTS

We report results from statistical testing and observations from our interview data. Examples of mood boards created in the study appear in Figure 7.

Quantitative Results

We compared the two pairs of conditions, with and without semantic labels, according to the Creativity Support Index, and preferences about each condition's perceived usefulness. We used a repeated-measures ANOVA, treating 'participant' as a random factor using SAS JMP's REML procedure.

Creativity Support Index

We find a significant preference for *Compose* ($F(1,33) = 10.29, p = .003$; means 65.36 vs. 56.64) over *Reflect*, which suggests the primary importance of collecting and constructing in the ideation process. When interacting with semantic labels, users rate creativity support higher in terms of *Enjoyment* and *Exploration*, which are both crucial to collecting and reflecting upon inspirational material. The additional interaction with the provided semantics, though, also led to a small decrease in *Immersion*, and caused additional *Effort to use*. No measures were significant (see Table 2), which was expected since the non-semantic condition already exceeds the functionality of currently available common tools.

Preferences

Participants reported their perceived usefulness of semantic labels in the two conditions. Due to a logging error, answers from one participant's compose preferences were not saved. The rest of the data is consistent and was not affected.

Participants preferred the semantic version of each tool, independent of the task and condition. Specifically, they preferred

(a) Coffice

(b) MEETbar

(c) Shoe store

Figure 7: Example mood boards designed in the study

semantic labels for the *Compose* task (64%) and *Reflect* task (58%), to the non-semantic versions: *Compose* (0%) and *Reflect* (25%). The other participants were neutral: *Compose* (36%) and *Reflect* (17%). This suggests that providing interactive semantic labels can help designers throughout the inspiration process.

We also asked participants to rate the perceived usefulness of semantic labels. Most designers rated semantic labels as more useful for the *Compose* (64%) than the *Reflect* condition (33%). However, a few rated the labels as irrelevant for mood board design: *Compose* (18%) and *Reflect* (17%). The rest were neutral: *Compose* (18%) and *Reflect* (50%).

Qualitative Results

Providing interactive semantic labels for the inspiration phase helped designers better express their ideas. Participants used them to find “the proper word for what you are already thinking”, because they often “have some inspiration maybe already pictured, but how to find this picture?” (P1). Designers were able to search for vague or partially expressed ideas: “This picture is kind of nice, but I want something more similar to that one” (P3). Providing semantic analysis for reflection encouraged sense making, where designers “think through the concepts” (P7), and remain in control of the current process.

Semantic Enriched Search

Designers expressed a clear preference for the semantic condition when using interactive labels.

Discovering material: Designers appreciated the ability to “mix text and images in the search” (P10), a feature that the most experienced designers deemed “unprecedented” (P9). Some found “a lot of new images [that I] wouldn’t probably find if I didn’t drag the visuals [to the search]” (P5), and were happy about the many “options to choose from” (P6). The interactive semantic labels helped designers find the ‘right’ images faster (P5, P8), and provided a new source of inspiration

(P1, P6, P8). Labels also suggested alternatives, especially when words were missing, e.g. when searching for abstract objects such as texts and fonts. Some designers added system-generated semantic words to the search just to “see what comes out” (P6). Most found the process enjoyable: “[it’s] fun to collect images and that I can use them immediately” (P4).

Expressing ideas: Designers felt that the primary innovation was the ability to say: “This picture is kind of nice, but I want something more similar to that one.” (P3), and to manipulate the semantics accordingly (P2, P3, P6, P7): “To combine the pictures with words, to express it both with words and the images.” Designers found the labels useful and missed them in the non-semantic condition (P6). They said the labels helped them “narrow down” (P8) what they want and “easily get more out of a specific topic.” (P2).

Feeling in control: Semantic labels in the search area also helped participants better understand and control the system: “Because I feel like that is how the system works, this is also how I would search for stuff with the semantics” (P1). This was true not only of the non-semantic conditions, but also commercial tools. The ability to adjust with “plus and minus helped me to get the picture I wanted more quickly” (P8). P6 said that the semantic search interface “gave me more choices and also it was more accurate. I have the feeling in terms of color choices or the images that it was showing me. I was kind of overwhelmed actually with the information. I had so many options to choose from.”

Reflection-in-action: One designer said that what he “really liked in the inspect tool was that there were some keywords, like these are the main keywords, and they have a hierarchy.” He added: “If these are the main keywords, but that is actually not the thing that I need, then I go to the other ones and select them [which] helped to align my thinking.” (P1). Participants also said that semantic labels help “explain why I am interested in the picture” and “somehow helps me to do this artistic work” (P8). Finally, P1 reflected on his ability to use existing material for search: “You can always go from where you have been before, so you don’t have to remember what you have searched before”. When combined with the search history, it offers a “tracking along your line of thinking”.

Semantics Clouds support Sense making and Reflection

Most designers preferred the semantic reflection condition (10/12), except two designers who preferred the simplicity of the non-semantic conditions.

	No Semantic		Semantic		Sig.
Factor	Score	SD	Score	SD	<i>p</i>
Collaboration	13,08	4,45	13,17	3,81	.8057
Enjoyment	13,21	4,99	14,38	3,85	.1459
Exploration	10,63	4,14	11,42	4,38	.483
Expressiveness	10,88	4,03	10,92	3,37	.5945
Immersion	13,17	4,68	12,46	4,85	.544
Results Worth Effort	13,17	4,54	12,96	4,28	.325
CSI	57	19.17	60.93	18.03	.9596

Table 2: Creativity Support Index results for sem/no sem.

Sense making: The most beneficial aspect of using labels is to support sense making. Participants said that the labels helped them “think through the concepts” (P7, P3) and “help me to explain why I like it” (P8). The semantic overview clouds also helped “clear...up my thinking, so what do I want to talk about” (P1), helped “organise my ideas” (P6) and “gives me a very good, very fast understanding of what has been there” (P7). With respect to their own, self-collected material, it helped them “categorize these and sort of see the connections between these pictures” (P6). Without semantic clouds “it was not as easy” (P7) to reflect on the mood board and harder “to make up my own words” (P3).

Reflection-on-action: As with semantic search, participants said they were more aware of the mood board content (P4, P5). It helped “in the moment where I basically had the same things there and then I looked at the tags and I realised: [...] maybe I can just address something completely different” (P1); and encouraged exploration of other elements (P11). Just “seeing and checking the meaning” (P4) and “different words in different sizes made you think through” (P7). The designers reported that semantics helped them “articulate the reason why I chose it, because the gut feeling is somehow very difficult to explain” (P8). The labels made designers feel “a bit more in control of my design, because I was able to see these keywords” (P6) and reflect upon them. Finally, this reflection was not limited to their own creation process: “This software helped to explain the idea to someone else” (P4).

Using SemanticCollage

Impact on mood board design: Almost all participants (11/12) appreciated the simplicity of the interface and found it ‘fun’ (P4,P7) and fast to use and learn (P2, P6, P11). This includes control of the features (P1, P4, P7): “You can move the pictures” (P1), “adjust them easily” (P2) and things are “easy to find” (P5). Ten participants wanted to use the tool in their design practice, especially because of the increased speed over creating a physical mood board. One mentioned that he hoped “this could be a professional tool for designers, then I would like to use it” (P8) while others said they “would definitely use it, because there are strong benefits: it is fast and it is usable and in many ways it follows the natural workflow” (P2).

Comparison with commercial tools: The experience of creating visual collages encouraged many participants to compare *SemanticCollage* to existing digital and physical tools. They described it as “really fast and easy to use compared to, let’s say, Photoshop” (P2), which is “time consuming to get the pictures there from the websites” (P2). This holds despite *SemanticCollage*’s more limited feature set compared to Photoshop (P7), which might reduce its usefulness (P12). Participants also said that *SemanticCollage* was “definitely much faster than just having the google search” (P5, P8). They appreciated having access to all the tools in a single system, which lets them “follow the natural workflow” (P2), while accelerating the process (P6, P11) because “[I] don’t have to look for tools” (P12) and can use images immediately (P4).

Semantics beyond SemanticCollage

Multiple participants suggested innovative ideas for using *SemanticCollage*. Several said that semantic search would be

useful for “anything visual ... not maybe just mood boards but the idea of finding pictures or finding visuals that portray a certain message” (P1). They also felt that semantic labels could help designers reflect on their previous work: “If I make a mood board and then look at it again in one week, having these labels would be super useful.” (P5). They also highlighted its potential for collaborative work (P4, P5, P7, P9), where semantic labels could “track [] your co-workers’ opinions, and also your own opinions” (P5). P4 felt that *SemanticCollage* could be the ‘communication board of the future’, when working on a mood board “with a co-worker” or for “group discussion or brainstorming” (P4), P7 said that collaborators would benefit by developing “their own way and understanding of those words, so there would be a language that they use you cannot display at the moment”. Others suggested using it to communicate with external stakeholders: *SemanticCollage* would help to “get a better brief, and better understand [his client’s] ideas, mood, tastes, like colors” (P9).

DISCUSSION

This paper investigates the potential of semantic labels to support digital mood board design. We are particularly interested in whether they address our key research questions, i.e. do they help designers 1) better express and explore vague, visual ideas; 2) make sense of, reflect upon and communicate the meaning of their visual concepts; and 3) take advantage of semantic labels without disrupting their creative flow?

Semantic Labels Help Express Visual Ideas

Since designers think visually, revealing the underlying semantics of images reduces the burden of coming up with the right search terms. They can begin with images, or combinations of images and text as search objects. *SemanticCollage* uses each image’s semantic labels to digitally represent visual objects for search. This retains the flexibility of text-based searches, with free choice of search engine. Note that, although we chose *DuckDuckGo* for our studies, *SemanticCollage* also works well with *Google Image Search* and curated collections such as *Unsplash*. *SemanticCollage* also better aligns search results to the designer’s intentions. Some image search engines use similar descriptive semantic labeling algorithms to label their collections, but they cannot handle vague requests, such as “I want something more similar to that one” (P3). Furthermore, *SemanticCollage* lets users adapt the semantics to improve iterative search, providing “choices that are more accurate [...] I had so many options to choose from.” (P6)

Participants said that semantic-based search tools helped them find images they would not otherwise have found, similar to [2]. However, retrieving semantic labels from the image itself, rather than from its original web page, lets *SemanticCollage* stay independent from the source of the material. Users can choose any source, as well as generating or modifying their own images. *SemanticCollage* also helps designers find better words and arguments for explaining their ideas to others. Semantic labels and clouds are especially useful for articulating the reasoning behind a particular decision, because: “a gut feeling is somehow very difficult to explain” (P8). Earlier research uses time or activity-based information to support design reflection, whereas *SemanticCollage* draws semantic

content directly from the images on the mood board. Designers find these content-centered clouds very useful for reflecting on relationships among images, and seeking better descriptions of the overall mood board.

SemanticCollage also helps address ‘design fixation’ [23], where designers become stuck on certain ideas, thus harming creativity. Even if designers are tempted to reuse the semantics provided, instead of constructing their own ideas, *SemanticCollage* retrieves semantics from a large set of training data which generates diverse, but accurate labels. Providing the top ten semantic labels balances the trade-off between exploiting existing search items, and expanding the designer’s inspiration space. Our results indicate that *SemanticCollage* can support diverse forms of mood boards, including unstructured collections, focused design spaces, and communicative layouts.

Semantic Labels Help Sense Making and Communication

SemanticCollage demonstrates a new form of computational support for reflecting upon visual material. It supports *reflection-in-action* by letting designers take advantage of system-generated semantic clouds and requested clouds to synthesize ideas and concepts. These text-based tag clouds offer an enriched overview of selected image sets, increasing designer’s understanding both of the mood board’s current content, and what it lacks. *SemanticCollage* also helps designers visualize relationships across images based on content, e.g. by highlighting images with similar semantics and letting users request ‘common labels’ for selected groups of images.

Designers also found semantic labels helpful when seeking missing words. Deciding which image to include can be difficult: “I don’t know – I just like this image”. However, Teevan et al. [50] argue that designers need these ‘anchor points’ to let them refine their search in order to eventually obtain the desired result. Study participants found semantic labels helpful “to explain why I like it” (P8). The material itself thus enriches and alters the designer’s evolving vision [25], and serves as a new starting point for further exploration.

SemanticCollage avoids disrupting creative flow

Designers found the interaction with *SemanticCollage* to be fun, easy and powerful. They could easily personalize semantic labels, creating a highly dynamic, iterative form of interaction, but only when it made sense within the creative process. Study participants especially enjoyed their level of control, and appreciated the sense that they understood “how the system works.” (P1)

Directions for Future Research

SemanticCollage offers an example of how to apply semantic analysis to improve ideation and reflection on visual material in the context of mood board design. However, many creative practices, including sketching, crafting and sculpting, could also benefit from exploring open-ended image and shape associations. Enriching system knowledge with semantic labels can improve the suitability of user’s contributions, and would also open new semantic spaces to explore. Our results show that semantic analysis of images helps designers make sense of their own work, suggesting that systems such as *SemanticCollage* could be designed for other creative practices, such as

web, game character or industrial product design. *SemanticCollage* could also enrich the collaborative aspects of creative ideation and sense making. A first step is ImageSense [26], a collaborative mood board design tool where designers and machine learning algorithms share agency in the exploration, sharing, and communication of ideas.

Limitations of the Study

The study is designed to let us compare two common design tasks, collection and reflection, with and without semantic labels, which balances the trade-offs between a lab study that offers control at the expense of ecological validity, and a real-world observation study that provides the reverse. However, the resulting tasks are only five-seven minutes each, and some participants felt under time pressure during the study. Also, our goal was to understand if and how semantic labels support collection and reflection activities, rather than to evaluate *SemanticCollage* as a complete mood board design tool. The latter would require a follow-up longitudinal study, ideally with professional designers, to study the impact of *SemanticCollage* in a real world design setting.

CONCLUSION

We present *SemanticCollage*, a digital mood board design tool that exploits existing computer vision algorithms to enrich images with semantic labels. These interactive computer-generated semantic labels are accessible to designers, enhancing both search and reflection capabilities. *SemanticCollage*’s features and layout are closely aligned with the design process, combining collection, composition and reflection phases within a single tool. This synergy produces a visual ideation process where designers naturally express visual ideas, create relationships among images, and reflect on image collections. *SemanticCollage* supports human-computer partnerships where designers remain in control by personalizing and extending the influence of the semantic labels, as needed.

Our study showed that *SemanticCollage*’s semantic features increase exploration and enjoyment; and most designers find semantic labels useful throughout the design process. Designers felt they were in control, and created highly diverse types of mood boards, including unstructured collections, design spaces, and communicative layouts. Semantic labels also helped designers ‘reflect in action’, helping them to transform their vague ideas into expressive search queries. Reflecting on semantic labels further increases awareness of mood board content, including identifying missing elements on the board, and helps designers discover new relationships among images, and find words to communicate their ideas to external stakeholders. *SemanticCollage* demonstrates how we can create a fluid, intuitive tool that takes advantage of state-of-the-art semantic labeling algorithms to offer designers better support for ideation and sense making.

ACKNOWLEDGMENTS

This work was 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] Dan Bloomberg. 2019. Color quantization using modified median cut. (03 2019), 6.
- [2] 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.
- [3] 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. DOI: <http://dx.doi.org/10.1145/2858036.2858159>
- [4] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative research in psychology* 3, 2 (2006), 77–101.
- [5] Search by Muzli. 2020. *Muzli* (2020). <https://search.muz.li>
- [6] Donald Thomas Campbell and Thomas D Cook. 1979. *Quasi-experimentation: Design & analysis issues for field settings*. Rand McNally College Publishing Company Chicago.
- [7] Hernan Casakin and Gabriela Goldschmidt. 1999. Expertise and the use of visual analogy: implications for design education. *Design studies* 20, 2 (1999), 153–175.
- [8] Tracy Cassidy. 2011. The mood board process modeled and understood as a qualitative design research tool. *Fashion Practice* 3, 2 (2011).
- [9] 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.
- [10] Erin Cherry and Celine Latulipe. 2014. Quantifying the creativity support of digital tools through the creativity support index. *ACM Transactions on Computer-Human Interaction (TOCHI)* 21, 4 (2014), 21.
- [11] Douglass R. Cutting, David R. Karger, Jan O. Pedersen, and John W. Tukey. 1992. Scatter/Gather: A Cluster-based Approach to Browsing Large Document Collections. In *Proceedings of the 15th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '92)*. ACM, New York, NY, USA, 318–329.
- [12] Peter Dalsgaard and Kim Halskov. 2012. Reflective Design Documentation. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, New York, NY, USA, 428–437. DOI: <http://dx.doi.org/10.1145/2317956.2318020>
- [13] Peter Dalsgaard, Kim Halskov, and Rune Nielsen. 2008. Maps for design reflection. *Artifact* 2, 3–4 (2008), 176–189.
- [14] Alanah Davis, Gert-Jan de Vreede, and Robert Briggs. 2007. Designing thinkLets for convergence. *AMCIS 2007 Proceedings* (2007), 358.
- [15] DuckDuckGo. 2020. DuckDuckGo Search engine. (2020). <https://duckduckgo.com>
- [16] Volodymyr Dziubak. 2019. Supporting graphic designers' online inspiration seeking practices via software tools. (2019).
- [17] 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.
- [18] James Fogarty, Desney Tan, Ashish Kapoor, and Simon Winder. 2008. CueFlik: Interactive Concept Learning in Image Search. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*. ACM, New York, NY, USA, 29–38.
- [19] 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.
- [20] Jérémie Garcia, Theophanis Tsandilas, Carlos Agon, and Wendy E Mackay. 2014. Structured observation with polyphony: a multifaceted tool for studying music composition. In *Proceedings of the 2014 conference on Designing interactive systems*. ACM, 199–208.
- [21] 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.
- [22] Sonja Gutwein. 2013. Computer support for collaborative creativity. *Beyond the Desktop* (2013), 15.
- [23] David G Jansson and Steven M Smith. 1991. Design fixation. *Design studies* 12, 1 (1991), 3–11.
- [24] Ianus Keller, Froukje Sleeswijk Visser, Remko van der Lugt, and Pieter Jan Stappers. 2009. Collecting with Cabinet: Or how designers organise visual material, researched through an experiential prototype. *Design Studies* 30, 1 (2009), 69–86.
- [25] Gary Klein, Brian Moon, and Robert R Hoffman. 2006. Making sense of sensemaking 1: Alternative perspectives. *IEEE intelligent systems* 21, 4 (2006), 70–73.
- [26] Janin Koch, Nicolas Taffin, , Michel Beaudouin-Lafon, Markku Laine, Andrés Lucero, and Wendy Mackay. 2020. ImageSense: An Intelligent Collaborative Ideation Tool to Support Diverse Human-Computer Partnerships. In *Proceedings of the ACM Human-Computer Interaction (CSCW '20)*. ACM.
- [27] John Kolko. 2010. Sensemaking and framing: A theoretical reflection on perspective in design synthesis. *Design Research Society* (2010).

- [28] Google LLC. 2020a. Google Images Search. (2020). <https://images.google.com/>
- [29] Google LLC. 2020b. Vision AI, Cloud Vision API. (2020). <https://cloud.google.com/vision/>
- [30] 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.
- [31] Andrés Lucero. 2015. Funky-Design-Spaces: Interactive Environments for Creativity Inspired by Observing Designers Making Mood Boards. In *Human-Computer Interaction*. Springer, 474–492.
- [32] Andrés Lucero, Dzmitry Aliakseyeu, Kees Overbeeke, and Jean-Bernard Martens. 2009. An Interactive Support Tool to Convey the Intended Message in Asynchronous Presentations. In *Proceedings of the International Conference on Advances in Computer Entertainment Technology (ACE '09)*. ACM, New York, NY, USA, 8.
- [33] Nic Lupfer, Andruid Kerne, Andrew M. Webb, and Rhema Linder. 2016. Patterns of Free-form Curation: Visual Thinking with Web Content. In *Proceedings of the 24th ACM International Conference on Multimedia (MM '16)*. ACM, New York, NY, USA, 12–21.
- [34] shared visual library Niice-A smart. 2020. *Niice* (2020). <https://niice.co/>
- [35] 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/>
- [36] Sharoda A Paul and Meredith Ringel Morris. 2011. Sensemaking in collaborative web search. *Human-Computer Interaction* 26, 1-2 (2011), 72–122.
- [37] Unsplash: Beautiful Free Images & Pictures. 2020. *Unsplash* (2020). <https://unsplash.com/>
- [38] Pinterest. 2020. *Pinterest* (2020). <https://www.pinterest.com>
- [39] 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.
- [40] Image Finder High quality free stock images. 2020. *ImageFinder* (2020). <https://imagefinder.co/>
- [41] Yvonne Rogers. 2004. New theoretical approaches for human-computer interaction. *Annual review of information science and technology* 38, 1 (2004), 87–143.
- [42] Natalia Rümmele, Yuriy Tyshetskiy, and Alex Collins. 2018. Evaluating approaches for supervised semantic labeling. *arXiv preprint arXiv:1801.09788* (2018).
- [43] Daniel M Russell, Mark J Stefik, Peter Pirolli, and Stuart K Card. 1993. The cost structure of sensemaking. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems*. ACM, 269–276.
- [44] Donald A Schön. 1987. Educating the reflective practitioner. (1987).
- [45] Donald A Schön. 2017. *The reflective practitioner: How professionals think in action*. Routledge.
- [46] Isabella Seeber, Gert-Jan De Vreede, Ronald Maier, and Barbara Weber. 2017. Beyond brainstorming: Exploring convergence in teams. *Journal of Management Information Systems* 34, 4 (2017), 939–969.
- [47] Moushumi Sharmin and Brian P Bailey. 2013. ReflectionSpace: an interactive visualization tool for supporting reflection-on-action in design. In *Proceedings of the 9th ACM Conference on Creativity & Cognition*. ACM, 83–92.
- [48] Kyle Steinfeld. 2017. Dreams may come. In *Acadia*. 590–599.
- [49] Hillevi Sundholm, Henrik Artman, and Robert Ramberg. 2004. Backdoor Creativity: Collaborative Creativity in Technology Supported Teams.. In *COOP*. 99–114.
- [50] Jaime Teevan, Christine Alvarado, Mark S Ackerman, and David R Karger. 2004. The perfect search engine is not enough: a study of orienteering behavior in directed search. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 415–422.
- [51] Kai Wang and Jeffrey V Nickerson. 2017. A literature review on individual creativity support systems. *Computers in Human Behavior* 74 (2017), 139–151.
- [52] 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.
- [53] 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.