ICONATE: Automatic Compound Icon Generation and Ideation

Nanxuan Zhao[1,2] **Nam Wook Kim**[3] **Laura Mariah Herman**[4] **Hanspeter Pfister**[1]
Rynson W.H. Lau[2] **Jose Echevarria**[5] **Zoya Bylinskii**[5]
[1]Harvard University [2]City University of Hong Kong [3]Boston College [4]Adobe Inc.
[5]Adobe Research
nanxuanzhao@gmail.com nam.wook.kim@bc.edu {lherman, echevarr, bylinski}@adobe.com
pfister@seas.harvard.edu rynson.lau@cityu.edu.hk

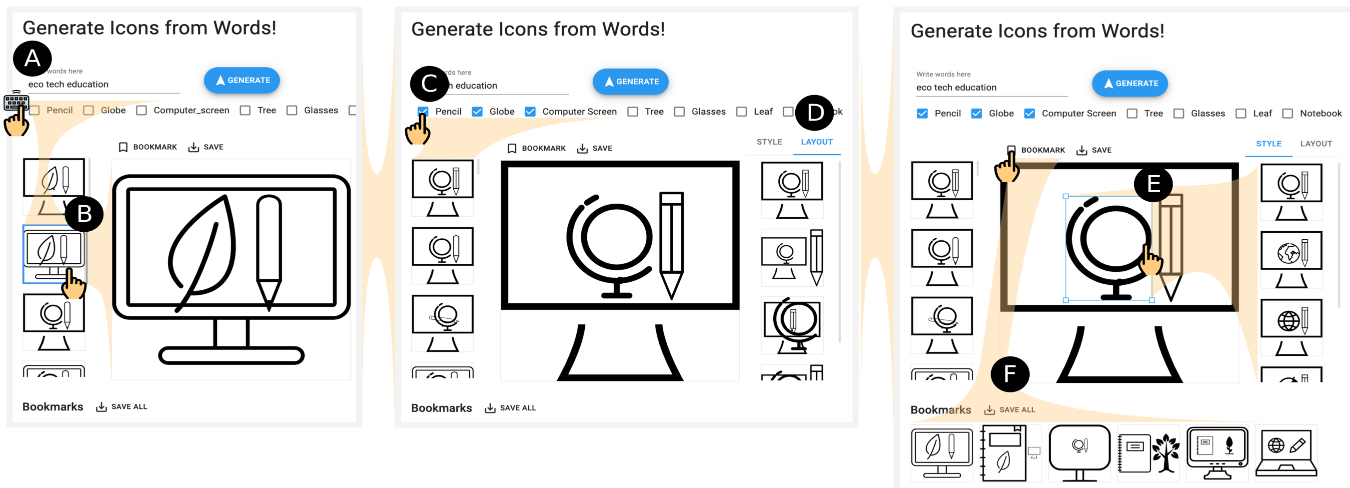

Figure 1. ICONATE is a system, including an interactive interface and a computational pipeline, for compound icon generation and ideation. Given an input text query such as "eco tech education" (A), our system will provide a list of diverse, automatically generated compound icon suggestions (B). Suggestions can then be customized through the interface, according to semantics (C), space/layout (D), and style (E) features. The system facilitates iterative design by bookmarking in-progress designs (F). The beige overlays in this figure highlight the interface elements relevant to each interaction.

ABSTRACT
Compound icons are prevalent on signs, webpages, and infographics, effectively conveying complex and abstract concepts, such as "no smoking" and "health insurance", with simple graphical representations. However, designing such icons requires experience and creativity, in order to efficiently navigate the semantics, space, and style features of icons. In this paper, we aim to automate the process of generating icons given compound concepts, to facilitate rapid compound icon creation and ideation. Informed by ethnographic interviews with professional icon designers, we have developed ICONATE, a novel system that automatically generates compound icons based on textual queries and allows users to explore and customize the generated icons. At the core of ICONATE is a computational pipeline that automatically finds commonly used icons for sub-concepts and arranges them according to inferred conventions. To enable the pipeline, we collected a new dataset, *Compicon1k*, consisting of 1000 compound icons annotated with semantic labels (i.e., concepts). Through user studies, we have demonstrated that our tool is able to automate or accelerate the compound icon design process for both novices and professionals.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

ACM ISBN 978-1-4503-6708-0/20/04 ...$15.00.
https://doi.org/10.1145/3313831.3376618

Author Keywords
Compound Icon; Ideogram; Pictogram; Icon Design; Graphic Design; Design Tools

CCS Concepts
•Human-centered computing → Graphical user interfaces; Graphical user interfaces; •Computing methodologies → Visual content-based indexing and retrieval;

INTRODUCTION
Icons, as a universal language, appear everywhere in our daily lives – from human-computer interfaces to human-to-human

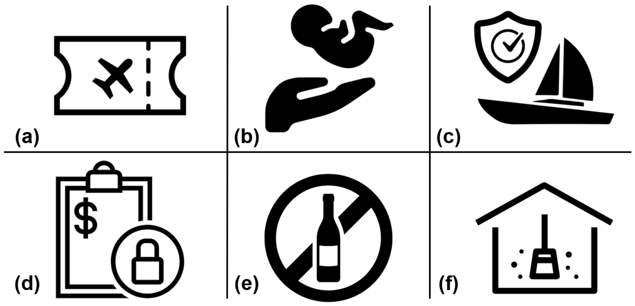

Figure 2. Examples of compound icons: combinations of constituent icons to convey a more complex idea. (a) flight ticket, (b) baby care, (c) boat insurance, (d) business contract, (e) no drinking, (f) house cleaning.

communication. They are one of the most simple and efficient ways to convey a message [19, 20, 31]. To transmit complex information, an icon is often composed of simpler sub-icons, corresponding to basic, easily-recognizable concepts. We refer to such an icon as a *Compound Icon*. For example, the icon for "business contract" in Figure 2(d) is composed of a "dollar sign" for business, "lock" for security, and "clipboard" for contract[1]. Road signs and maps use compound icons to point drivers in the direction of food and other necessities, while webpages and infographics use compound icons to guide viewers quickly to the topics that they care about.

However, creating effective compound icons is a challenging task, since a well-designed icon should be both understandable and aesthetically pleasing. It involves considerations of how to represent each sub-component, how to lay them out, and how to arrive at a cohesive design. Novices can get overwhelmed by hundreds of options. Even professionals with design knowledge and extensive experience often turn to diverse resources (to find inspirational or referential icons) and create many iterations before coming up with the final icon.

In this work, our goal is to provide intelligent design support for creating compound icons given input compound concepts (e.g., "flight ticket" or "baby care"), while facilitating rapid exploration and ideation. To gain insights from existing icon design processes, we conducted semi-structured interviews with professional designers with extensive icon design experience. We learned that designers begin with an inspiration and ideation phase, searching for existing visual metaphors semantically associated with the target concepts ("Semantics"). Once the visual metaphors have been identified, designers start iterating on multiple icon designs to create a final representation to convey the compound concept. They do so by repurposing or combining existing icons that are stylistically compatible ("Style") and arranging them in a cohesive layout ("Space") to convey the intended meaning. They gather feedback from teammates to continue the iterative design process.

[1]The icons included in this paper were obtained from the Noun Project (https://thenounproject.com/) by paying for a NounPro subscription, which provides a royalty-free license to the downloaded icons. While we plan to share the icon vocabulary we curated, including all the meta-data and icon URLs, we will not share the icon images themselves, to abide by the terms of use of the website.

To mimic the professional design process, we introduce ICONATE: a novel system to support compound icon design, which given a compound text query, automatically generates a set of diverse compound icon proposals, and allows a user to iterate on the final design. The core idea behind this system lies in our **3S** icon design framework consisting of **S**emantics, **S**tyle, and **S**pace components. We implemented a computational pipeline based on this framework that combines constituent icons together to represent compound concepts.

To enable the computational pipeline, we constructed new datasets. First, to search for semantically compatible constituent icons given a text query (**S**emantics), we curated a 152-concept constituent icon vocabulary (*IconVoc152*), by consolidating concepts from multi-language picture-naming studies [49], universal Emoji conventions [1], and existing icon collections [2]. For the pipeline to learn to combine constituent icons together in a stylistically and spatially compatible way (**S**tyle and **S**pace), we also constructed a dataset of annotated compound icons (*Compicon1k*). We evaluated the feasibility of ICONATE by testing it on professionals and novices. Our user studies suggest that ICONATE succeeds as an ideation tool for professionals and is an efficient icon creation tool for novices.

The contributions of this work include[2]:

1. ICONATE, a system that facilitates ideation and iteration for compound icon design. The system design is based on our "3S" framework (i.e., Semantics, Style, and Space), inspired by interviews with professional designers.
2. A computational pipeline that powers the system by automatically generating icon proposals given a user-defined text query.
3. A data collection process and annotation tool that enable the construction of datasets for the computational pipeline:
    (a) *IconVoc152*: consisting of 152 basic universal concepts, as the building blocks for compound icons.
    (b) *Compicon1k*: consisting of 1000 compound icons corresponding to compound concepts, each one segmented into constituent icon parts that map to basic concepts in *IconVoc152*.

RELATED WORK

Visual Language

In this paper, we see icons as an expressive visual language for communicating abstract concepts. We provide an automatic approach through which text queries are converted into icons. Related prior work that has considered abstract visual symbols as an expressive language includes: Khandekar et al.'s investigation of Emoji as a form of communication in a mobile social app [22], Zitnick et al.'s abstract clipart scenes to represent text stories [55, 56], and Chilton et al.'s crowd-driven pipelines for creating visual metaphors starting with text input [12]. Icons have also been shown to be effective supplements to text and data, making information visualizations more attention-grabbing and memorable [3, 4, 5, 18].

[2]The project website: http://nxzhao.com/projects/ICONATE/

Design Exploration
Design tools are essential for creative professionals, enabling them to express ideas in a variety of forms such as illustrations, typography, logos, and icons. Traditional graphic design tools such as Adobe Illustrator provide a slew of features, ensuring flexibility in design to enhance the user's creative freedom. However, design processes in these tools still remain mostly manual, hampering rapid design iterations.

Quickly iterating through diverse alternatives is key to a better design outcome [6]. To aid this exploration process, many research projects have been devoted to providing computational support in design tools. This includes supporting low-level design decisions such as intelligent snapping [13], selection of complex nested elements [52], content-aware color palettes [21], and font search by semantic attributes [40].

Other research investigates ways to support high-level design tasks from ideation to layouts and design feedback. "Example galleries" can serve as a source for design inspiration such as looking for similar styles [42], adapting components from existing examples [26, 48], or automatically repurposing the examples with new content [24]. Another line of work centers on supporting the exploration of alternative layouts based on spatial and contextual constraints [35, 39, 51]. Others have also studied task-level action suggestions to rapidly explore different visual effects [16] or providing design feedback based on visual saliency [7] or through crowdsourcing [32].

The design of the ICONATE interface drew inspiration from the existing tools. However, we specifically focus on icon design rather than general graphic design and also consider multiple design factors (i.e., 3S icon design framework: Semantics, Style, and Space) rather than a single one.

Design Synthesis
Key to our 3S framework and the resulting computational pipeline is the semantics-driven selection of components that are compatible in layout and style. Here, we discuss the work most relevant to these design challenges.

Semantics. Early studies on using visual concepts [43] focused on the generation and usage of icons in graphical interfaces. Chang et al. [8, 9, 10] defined formal specifications for icon generation systems. Lin [29] studied design styles to successfully go from shape, image or function features to representational, abstract, or arbitrary icons. In comparison, our goal is not to come up with a formal system, but rather one that supports the icon design workflow. Cheng et al. [11] introduced Semantical Visual Templates for querying image repositories based on visual concepts. These templates emerge from a two-way interaction between the user and the system, capturing the most successful queries for a concept. Setlur and Mackinlay [45] used natural language processing to automatically propose sets of coherent and meaningful icons for data visualizations. Madan et al. [34] learned a mapping between text tags and icons in infographics for summarization into visual hashtags.

Style. Style transfer for images is currently an active research area [27]. However, icons show higher degrees of abstraction, remaining an open and challenging problem [34]. Style similarity metrics for vector illustrations based on handcrafted features [17] and for icons based on deep features [25] have been presented. Unlike these approaches which group icons by semantics (e.g., house icons together with similar house icons), our approach specifically disregards the semantic categories to learn a notion of style that can match different semantics together (e.g., a house icon with a style-compatible dog icon, for a professional-looking "dog house" icon). We thus train a new metric using deep features to find compatible icons to combine in our system.

Space. Qiang et al. [41] introduced a method to composite graphical elements in scientific posters. Zheng et al. [54] proposed a deep generative model to synthesize layouts based on the visual and textual semantics of user inputs. Li et al. [28] solved the layout problem in a more generally applicable space without being constrained to a single type of design. More similarly, Setlur et al. [44] extracted the context of a file to generate a representative icon using heuristic rules. Machida et al. [33] analyzed song lyrics to extract different topics and sentiments, and assigned to them existing icons and images for visualization. In contrast, ICONATE targets more abstract and universal icons, learned from universally recognizable semantics and current conventions. Liu et al. [30] focused on generating and editing the same types of icons as ours but without considering the semantic relationships between parts. This work could complement our system in terms of generating new icons from images, and suggesting more complex edits to the resulting compound designs.

UNDERSTANDING ICON DESIGN PRACTICE
In order to understand designers' current processes and design considerations in creating novel compound icons, we conducted rapid ethnographic interviews [37] with professional designers who had icon design experience.

Procedure and Tasks
We interviewed three professional designers: P1 (Senior Experience Designer, 11 years of experience), P2 (Senior Product Designer, 9 years), and P3 (User Experience Designer, 7 years of experience). P1 is a male in the 25-34 age range, works at a large company (> 10,000 employees), and makes icons on a daily basis. P2 is a male in the 35-44 age range, works at a small company (< 100 employees), and makes icons on a weekly basis. P3 is a male in the 25-34 age range, works at a large company (> 10,000 employees), and makes icons on a weekly basis. Each interview, conducted via video call, lasted approximately 45 minutes and was moderated by a user experience researcher. Each participant was paid $75 (for two sessions) and completed a semi-structured interview, which probed the designer's workflow, motivations, tools, experiences, and concluded with a step-by-step walkthrough of a recent icon design project of theirs.

Observations
Ideation. Designers often create compound icons by combining multiple simple icons together. How often? P3: *"Frequently!"*; P2: *"Most of our icons have this duality, of trying to convey two different concepts at once."*; P1: *"We remix, look*

and get inspirations from different metaphors, and bring them together to create a new metaphor."

When creating a new compound icon, all three designers begin with (i) ideating on semantics that may be associated with the concepts they are trying to convey (e.g., a rocket to represent launching a venture, a lightbulb to represent a bright intern, etc.) and (ii) searching existing image and icon repositories to get inspiration for how those concepts have been previously visualized (2/3 designers mentioned using the Noun Project, 3/3 designers mentioned using Google). P1: *"If a metaphor comes up many times, it may indicate that it works for people."* (e.g., using an icon of "hands" to represent "care").

Iteration. Once the visual metaphors have been identified, the next step is to create a single representation. P1: *"Making something new rarely happens. Existing icons are frequently repurposed."* P2 commented that getting an icon to convey multiple concepts at once can be tricky. Feedback from others might be that the icon is *"too vague on one topic over another"*. P3: *"the goal is to create recall for the user, yet do something different."* Designers consider style compatibility, style consistency for brand identity, layout, and layering. For example, P2 tries to *"make things balanced"* when combining different icons. All the designers work to make/find icons with consistent style features (e.g., stroke width and fill color).

Validation. Ultimately, the icon needs to be correctly understood by its target audience. To gather feedback, designers turn to their teams or other members of their company. There may be some back-and-forth with the client. If designing for universal use, cross-cultural and language differences need to be accounted for. In some cases there might be some basic testing online with 100 different people (P2), but in other cases the design team votes on the final design (P1), or the voting moves beyond the design team, but remains internal to the company (P3). All three designers commented that most of their external validation comes not from user testing, but from comparing their work to how other expert designers have represented related concepts elsewhere (i.e., common conventions in existing icon repositories).

Lessons

Interestingly, most icon design falls into the hands of generalists that work on different kinds of graphic and user experience design, rather than specialized icon designers. Professional designers turn to existing icon repositories and books for guidance and inspiration (e.g., [14, 15, 19]). Our main lessons from the interviews were:

L1. Compound icons are common. To create a compound icon, designers begin by brainstorming a set of associated basic concepts to compose together ("Semantics").

L2. Designers turn to existing icon repositories to get a sense of how various concepts have typically been visually represented in the past. This serves as external validation for their final design choices. Simpler icons, corresponding to basic, easily-recognizable concepts are remixed together.

L3. A few variants are created after several iterations, and the final selections are made based on multiple considerations: whether the intended meaning is evoked, whether the icon maintains brand identity, and how the icon will look in the final context and at different scales ("Style" and "Space").

3S ICON DESIGN FRAMEWORK

Based on the lessons from the expert interviews as well as our literature survey, we developed the 3S framework, a fundamental guideline for building tools to support the icon design process. Three key design considerations in this framework include semantics, style, and space, each of which can be tackled separately. Below, we outline each design consideration, why it is important, and how it might be challenging to implement. In the following sections, we will present our own particular implementation of this framework, called ICONATE.

Semantics. The semantics of an icon are the set of visual metaphors that can be used to represent the icon [23]. In order to ensure that a compound icon conveys its intended meaning, it is imperative to find visual metaphors that accurately portray the relevant concepts, as highlighted in our expert interviews. This step is particularly challenging because the semantic concept an icon represents may not be a direct translation of the visual object itself. For instance, a hand icon may represent “protection" or “insurance" (as in Figure 2b) rather than the physical body part, and the dollar sign may represent “business" (Figure 2d) rather than the type of currency. A design tool should support users in mapping abstract concepts (e.g., "secure") to visual metaphors (e.g., a lock or a shield).

Style. The style of a compound icon is the set of visual features characterizing the representation of the icon. Icon style can span a spectrum from abstract to realistic representations and can also vary in lower level features, such as stroke or fill types. An icon is rarely used alone but appears together among other icons to establish a visual language in the intended domain (e.g., brand marketing or a menu in a computer system). As a result, it becomes crucial to maintain style consistency. A design tool should support generating style-compatible icons.

Space. The space of a compound icon includes the spatial juxtapositions (layout) of visual elements, which are often also icons. Conventional layout design principles such as symmetry and alignment can apply to icon design. However, the relative positions of visual elements within a compound icon can convey a meaning. For instance, a hand positioned below a house can indicate house insurance while a different arrangement may have an alternative meaning or lack thereof. Designers must consider intricate spatial interactions between different icons to create a meaningful compound icon. A design tool should provide scaffolding for layout decisions based on both spatial semantics and conventional design principles.

ICONATE

We use the 3S framework as a basis for the design of our ICONATE system. Our main goal is to facilitate the icon design process by providing automatic assistance for creating compound icons to represent compound concepts. We mainly focus on addressing the challenging and tedious problem of generating and rapidly exploring design alternatives. In this

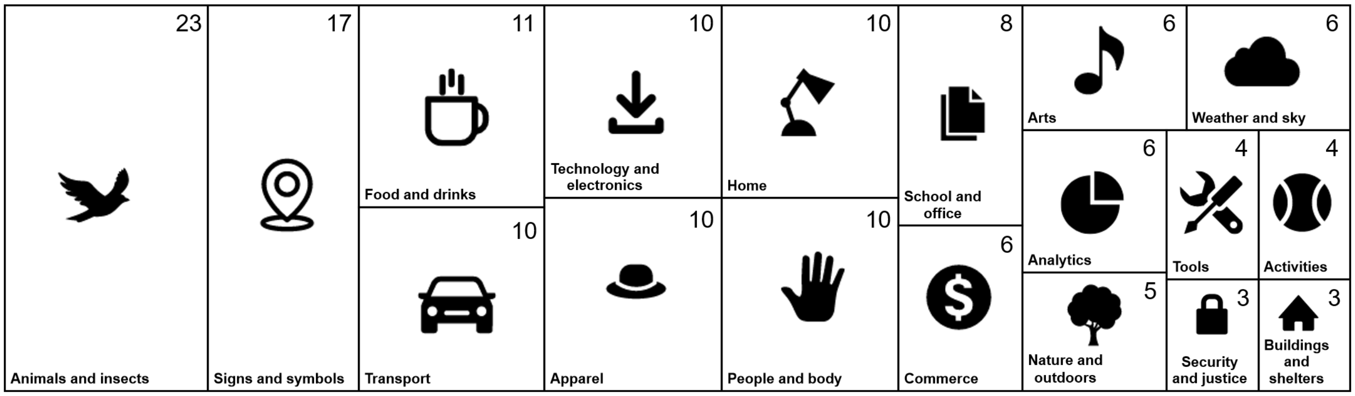

Figure 3. The 18 categories covering the 152 icon concepts from *IconVoc152*, with an icon representing the most commonly-found concept in each category (e.g., "bird" is the most frequently occurring icon in the Noun Project among the 23 icon concepts of animals and insects). The number at the corner of each cell indicates the number of concepts found in the category.

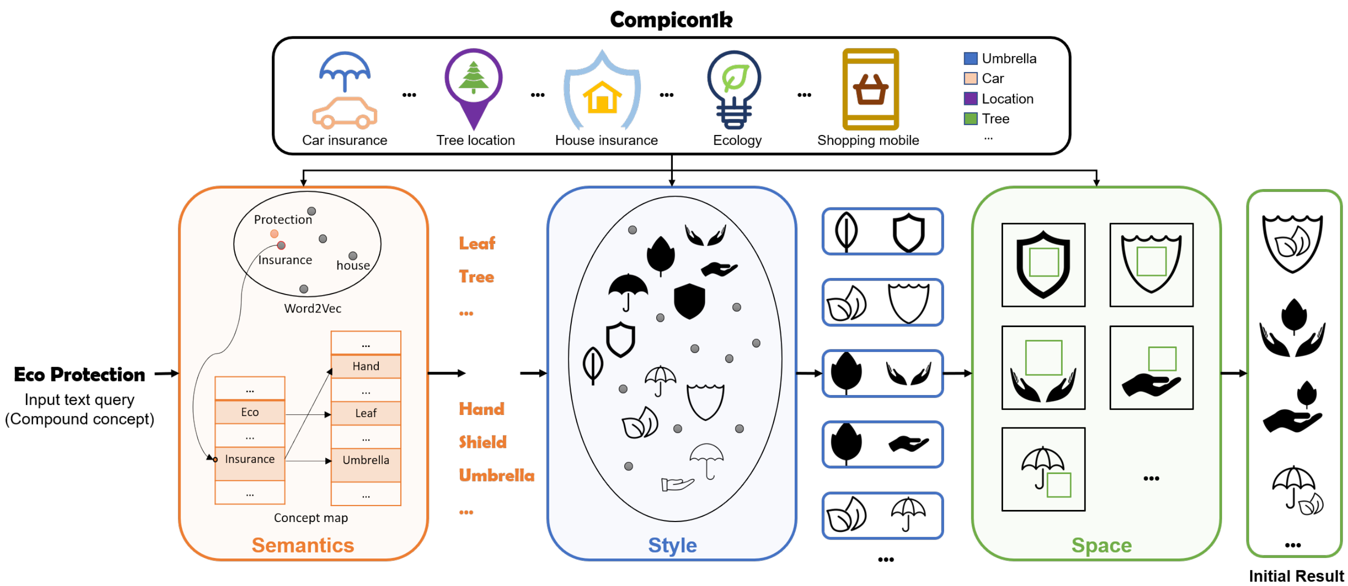

Figure 4. Our computational pipeline. With the support of our *Compicon1k* dataset, our pipeline is able to generate a list of diverse suggestions automatically by considering the 3S factors: Semantics, Space, and Style.

section, we provide an overview of our system. We will discuss the dataset and computational pipeline that power the ICONATE system in the following sections.

The ICONATE interface (Figure 1) has three main components: 1) a search interface for ideating on semantically-meaningful icon candidates, 2) a canvas for customizing and refining the style and space of the resulting icons, and 3) a bookmarking capability for capturing design variations, following professional workflows and the principle of iterative design.

Ideating on icon semantics. To retrieve candidate icons, a user types a textual query corresponding to a compound concept like "eco tech education" (Figure 1A). The interface then suggests compound icons generated by the computational pipeline (Figure 1B), as well as a textual list of basic concepts associated with the compound concept. The user can toggle which basic concepts they would like to be included/excluded from the final compound icon suggestions (Figure 1C), and the interface updates the suggestions accordingly. The user can retrieve additional suggestions by scrolling down to the bottom of the suggestion list. The user can select any of the suggested icons for further customization and refinement.

Refining style and space. On the canvas, a user can select a different arrangement of the constituent icons inside the final compound icon – keeping the same icons, but only modifying their relative positions and sizes (Figure 1D). Alternatively, a user can select a different style for one of the constituent icons, while keeping the other constituent icons unchanged (Figure 1E). On the main canvas, the user selects a constituent icon, and the pipeline suggests alternative style variants. Once the user accepts a new style from the suggestion panel, the icon on the canvas updates. The user can also directly manipulate the positions and sizes of the constituent icons, for greater control and freedom.

Iterating through design variants. Throughout the ideation and refinement process, a user can bookmark icons for later reference (Figure 1F). This functionality can be used to capture the evolution of a particular design, as it is refined, or to save a few variants (as an "idea board"). The user can retrieve a bookmark to continue iterating on that design at any point. For any advanced icon refinements that our tool does not currently support (e.g., stroke or color editing), the user can easily export the icons in SVG format, to be imported into other graphic design tools.

DATA COLLECTION

In this section, we introduce the datasets collected to enable our computational pipeline.

Constituent Icon Vocabulary (*IconVoc152*)

Our first step was to find a manageable but comprehensive set of concept "building blocks" (e.g., house, umbrella, person) that can be recombined to form a large diversity of compound icons (e.g., "house insurance", "rainy day"). To do this, we turned to multiple sources, including psycholinguistic studies and existing icon repositories.

Initializing from picture-naming datasets. To make sure our basic concepts map to recognizable icons across different languages, we turned to datasets that have previously been used for picture-naming tasks in the psycholinguistic literature. We started with the UCAL corpus [50] by selecting the concepts with 100% visual categorization accuracy, and at least 90% object naming accuracy across seven different languages in the International Picture Naming Project (IPNP) [49]. This produced an initial icon vocabulary of 93 concepts.

Growing to include icon-specific concepts. The IPNP and UCAL datasets built on much earlier psycholinguistic work and are missing some emerging concepts related to technology (e.g., computer screen, cell phone) and common symbols (e.g., checkmark, prohibited). To test and extend our initial icon vocabulary, we ran several rounds of icon annotation, by asking human participants to label the constituent icons within randomly sampled compound icons (study details described below). We grew our initial vocabulary by including new concepts that participants used repeatedly, while merging synonyms. Our vocabulary doubled in size to 211 concepts.

Shrinking down to commonly-used icon concepts. In the final step, to ensure that the concepts we include are useful as icons, we kept only the concepts that could be found in either the full Unicode Emoji list (v12.0) or in Microsoft's PowerPoint (v16.27), two popular repositories of icons. Our final icon vocabulary, *IconVoc152*, contains 152 concepts. This icon vocabulary along with all of the meta-data inherited from previous studies and other resources is provided together with this paper, made openly available to the research community.

Using the categories found in PowerPoint's icon lists and the Emoji taxonomy as inspiration, we classified our 152 icon concepts into 18 categories, each containing 3 to 23 concepts shown in Figure 3. Please find more details in the appendix.

Compound Icons Dataset (*Compicon1k*)

For our computational pipeline to be able to assemble constituent icons together in semantically, stylistically and spatially consistent ways, we reuse design knowledge (i.e., conventions) encoded in existing compound icons. Towards this purpose, we manually curated a set of compound concepts and collected the corresponding compound icons. We then annotated these icons using an interface developed specifically for this purpose.

Compound Concepts. Three authors curated a lit of compound concepts composed of concepts from the constituent icon vocabulary. Each compound concept was searched on the Noun Project and 1–5 different icons representing the compound concept were selected for download. We collected a total of 1000 compound icons for our *Compicon1k* dataset (some examples are provided in Figure 2). Each of our *IconVoc152* basic concepts appeared in at least three, and an average of 13, compound icons (with the exception of 8 concepts like "camel" and "gorilla", which are not commonly part of existing compound icons and had fewer instances). The most common concept, "person", appeared in 190 compound icons.

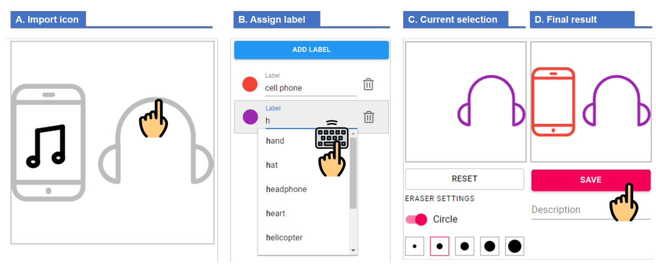

Figure 5. Our icon annotation interface.

Compound Icon Annotation. Each icon in *Compicon1k* was segmented into constituent icons using a novel icon annotation interface (interested readers may find additional details in the appendix and supplementary video). To label the *Compicon1k*, two external annotators (recruited through mailing lists) and one author annotated the full *Compicon1k* dataset. Annotators watched a few short videos with a demo of the annotation interface. The annotation process included two rounds. The first round of annotation was open-vocabulary, and annotators were free to type any word when annotating a constituent icon. This annotation round helped to refine the constituent icon vocabulary from the previous section. After the constituent icon vocabulary was finalized, the second annotation round constrained annotators to only existing vocabulary words via an auto-complete functionality added to the interface. The annotators were able to label miscellaneous, undefined areas as "others". Another author double-checked all of the annotations to remove inconsistent labeling.

As a resource for the research and design communities, we will make our list of 1000 compound icons (URLs corresponding to the Noun Project icons) available, as well as the human-generated segmentations of all the icons (in CSV form).

COMPUTATIONAL PIPELINE

Given a compound concept as text input, our computational pipeline aims to generate a ranked list of compound icons by

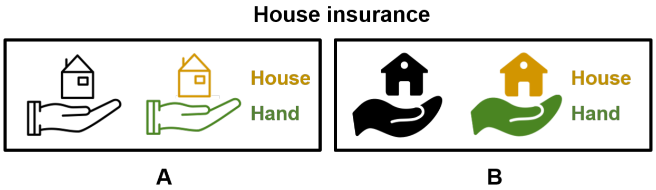

Figure 6. Our computational pipeline learns about semantics, style, and space from annotated data. (1) *Semantics:* hand can be used to represent insurance; (2) *Style:* hand in *A* is more style-compatible to the house in *A* than the hand in *B*; (3) *Space:* hand is often put below other objects to indicate protection and insurance.

considering three main factors: semantics, style, and space. More specifically, we assume that each compound concept (e.g., house insurance) can be decomposed into basic concepts (house, insurance), and each basic concept can be represented by a single constituent icon (house, hand). We collected at least 5 icons for each basic concept in *IconVoc152* from the Noun Project. To learn to combine these constituent icons in compatible ways, we turn to examples of existing compound icons to guide the semantics, style, and space of our generated results. By treating annotated compound icons in *Compicon1k* as templates (Figure 6), our pipeline can leverage existing conventions to create new compound icons outside of the existing dataset. The pipeline is shown in Figure 4.

Semantic Compatibility

A single icon may be used to represent multiple concepts, while the mappings from concepts to icons requires additional knowledge and familiarity with icons. Therefore, we manually curated a *concept map*: a dictionary of 500 concept mappings (e.g., the concept "eco" can be represented with "leaf" or "tree", "insurance" with a "shield", "umbrella", or "hand", etc.). This way, when the user types query concepts in the ICONATE interface (Figure 1A), if any of those concepts do not exist in *IconVoc152*, we can perform a look-up in the concept map to determine which icons should be suggested to the user. We further extended our concept map using the WordNet [38] hierarchy by associating the basic concept to its parent and grandparent nodes. For example, we connected "offspring" and "child" to "baby". If the query does not exist in the concept map, we look for the nearest neighbor (NN) concept that exists in our dictionary, using cosine distance D in the Word2Vec embedding space [36]. We define the following semantics compatibility score for a generated compound icon C:

$$\mathscr{S}_{semantics}(C) = \frac{\sum_{i \in C} E(i)}{|C|}, \tag{1}$$

$$E(i) = \begin{cases} 1, & \text{if } i \text{ exists in the concept map} \\ 1 - D(T(i), NN), & \text{otherwise} \end{cases}, \tag{2}$$

where $|C|$ represents the number of constituent icons in C and $T(i)$ is the concept of a constituent icon i.

Style Compatibility

Style compatibility is key to the aesthetics of compound icons. To measure style compatibility, we propose a model to learn an embedding space, where each icon is projected into an embedding vector. The closer the distance between two icons in the embedding space, the more style compatible the two icons are. Instead of using hand-crafted features, we chose to learn the style representation automatically using a convolutional neural network (CNN), inspired by successes of CNNs on related computer vision and graphics problems [25, 46, 53]. The novel insight of our training procedure is that constituent icons within a compound icon are more likely to be style compatible with each other than constituent icons from different compound icons (Figure 6). We used the *Compicon1k* dataset for training, since the constituent icons have been annotated for 1000 compound icons. We used a ranking formulation for training and designed a loss function to force constituent icons in the same compound icon to have a smaller distance in our embedding space than constituent icons from different compound icons.

Given two different compound icons C_1 and C_2, we randomly sample two constituent icons (i.e., i_1^1 and i_1^2) from C_1, and one icon (i.e., i_2) from C_2. For each constituent icon i, the style network S aims to predict a style embedding vector $S(i)$, where the distance between $S(i_1^1)$ and $S(i_1^2)$ is constrained to be smaller than $S(i_1^1)$ and $S(i_2)$. More formally,

$$D(S(i_1^1), S(i_1^2)) < D(S(i_1^1), S(i_2)), \tag{3}$$

where D is a cosine distance function. Then, we define the following loss function over the Compicon1k dataset P as:

$$L(P) = \sum_{C_1, C_2 \in P, \, C_1 \neq C_2} \; \sum_{i_1^1, i_1^2 \in C_1, i_2 \in C_2} H(S(i_1^1), S(i_1^2), S(i_2)), \tag{4}$$

$$H(s_1^1, s_1^2, s_2) = max(0, m - D(s_1^1, s_2) + D(s_1^1, s_1^2)), \tag{5}$$

where $S(i_1^1) = s_1^1$, and $m = 0.2$ is a margin hyperparameter. At test-time, we can query our trained embedding space to measure the style compatibility of constituent icons i_1, i_2 in a compound icon C:

$$\mathscr{S}_{style}(C) = \frac{\sum_{i_1, i_2 \in C, i_1 \neq i_2} (1 - D(S(i_1), S(i_2)))}{|C|^2}. \tag{6}$$

Space Compatibility

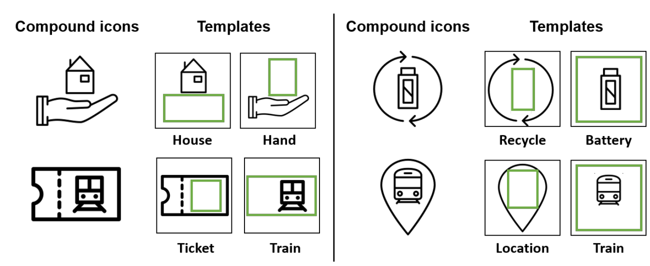

Figure 7. Space templates derived from the *Compicon1k* dataset. We treat each constituent icon as an anchor, and all other icons within the same compound icon as placeholders (i.e., green bounding box).

ICONATE arranges constituent icons based on their shapes and semantics (e.g., putting a hand below something to represent "care"). To composite a given set of constituent icons,

Figure 8. Automatic results generated by our computational pipeline for a few sample compound concepts which otherwise have few or no results in existing icon repositories.

we first derive the templates from the *Compicon1k* dataset, and arrange constituent icons according to these templates. We generate templates for each constituent icon in the *Compicon1k* dataset by regarding the current constituent icon as an anchor and other icons within the same compound icon as placeholders (Figure 7). We normalize all templates by fitting the anchor into a fixed-size bounding box in the center of the template. By considering all templates that have the same anchor concept (e.g., hand), we compute the mean IOU (i.e., Intersection Over Union) between all the placeholder bounding boxes. This gives us a confidence score for different placeholder locations (templates) given the anchor concept (e.g., other icons are commonly found *above* the hand icon). When generating new icons, we use the templates with the highest confidence. If an icon does not have a space template, we use a neighbor retrieval approach [25] to find the most similar-looking icon that does have a space template, and we adopt its template. Finally, given a pair of constituent icons i_1, i_2 in a compound icon C, we can compute a space compatibility score to represent the compatibility of the generated layout:

$$\mathscr{S}_{space}(C) = \frac{\sum_{i_1,i_2 \in C, i_1 \neq i_2} IOU(i_1^A, i_2^P)}{|C|^2}, \tag{7}$$

where i^A is the bounding box of the anchor and i^P is the bounding box of the placeholder in i's normalized space template.

Final Ranking of Generated Suggestions

To produce the final ranking over the generated compound icons in the left-most suggestion list in ICONATE (Figure 1), we calculate the overall compatibility score for a compound icon C as:

$$\mathscr{S}_{semantics}(C) + \mathscr{S}_{style}(C) + \mathscr{S}_{space}(C). \tag{8}$$

Sample Results

Figure 8 shows some compound icons that were automatically generated by our computational pipeline without human intervention. We only show top-ranked icons for each compound concept. Note that our pipeline is able to generate compound icons with more than two parts (e.g., "eco tech education"). Please find more results in the supplementary video.

EVALUATING WITH PROFESSIONALS

We invited back the three professional designers from our previous interviews (section *Understanding Icon Design Practice*) to experiment with ICONATE for compound icon creation. The consistent set of participants helps us evaluate whether our system succeeds in facilitating the specific workflows we sought to model.

Procedure and Tasks

As before, each session was conducted via video call, lasted approximately half an hour, and was moderated by a user experience researcher. We conducted a formal concept test, probing the users' expectations, comprehension, and perceived use of the tool. Each designer was asked to use the interface to generate icons for 3 compound concepts: "dog insurance", "fast download", and (if time allowed) "secure folder". We used a thinkaloud [47] methodology to gain insights into users' mental models for icon generation using ICONATE. In addition to real-time usage feedback, we also conducted a post-task debrief interview, which probed the expert users' perception of the tool's usefulness, and at what workflow junctures it would be most useful.

Observations

The feedback we received was very positive. All of the designers commented on the ease of ideating and iterating on complex and abstract concepts using ICONATE. Typical design workflows require manual searches for individual words in available icon repositories (e.g., the Noun Project) for inspiration, and then manual iterations combining icons within a vector design application like Adobe Illustrator. In contrast, ICONATE speeds up the ideation process by supplying multiple concepts and combining the icons automatically. P2: *"It's convenient for sketching out quick little ideas on the fly. The value of this tool is how little thought it takes."* P3 enjoyed the *Style* panel: *"I like component-wise variations, that is something that totally blew me away."* (Figure 9). P3 mentioned this fits directly into his existing workflow, as he would be similarly iterating on different styles for different component icons, but it would take significantly more effort to do so using existing icon repositories. Please find example results in the appendix. Our interface provides an efficient tool for ideation. All of the designers also commented on the usefulness of having the bookmarking option, to save snapshots of the design process, or to save a few different ideas. This is similar to "idea boards" that designers create when beginning to ideate on novel icon concepts.

While ICONATE can streamline the ideation phase as well as the process of identifying visual metaphors for queried concepts, it is not sufficient for generating the end product. Designers mentioned preferring to use familiar tools like Adobe's Illustrator to create the final icons (P1: *"I would see this more as an idea generator"*). Designers often work within the style guides of specific products or clients. Future iterations of the tool can also consider style constraints specified by the designer. P2 also mentioned it would be useful to allow a designer to input their own assets/icons and iterate on compound icons with them. Although not currently supported, our computational model could be used to recalculate compatibility

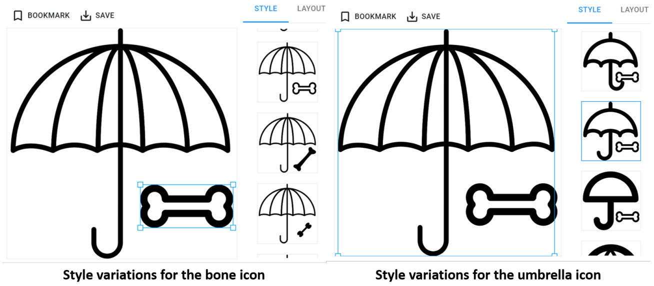

Figure 9. ICONATE allows selecting a constituent icon to explore its style variations, keeping the other constituent icons fixed. The variations are ranked by the style compatibility predicted by our computational pipeline. This was noted by professionals as being a very helpful feature for quick iteration.

for compound icon suggestions using any input icons that the designer provides.

EVALUATING WITH NOVICES

Our system is designed in accordance with the workflow of design professionals. To examine how ICONATE could support novices, we conducted a user study in which we compared our system to a traditional method more familiar to novice users.

Procedure and Tasks

We recruited 10 participants via university mailing lists, varying in age (23 on average) and design experience (poster or slide design, minimal icon design experience). Participants were provided with a presentation slide template (Figure 10) and given the prompt: "Imagine that you are using this slide to pitch a new product that you are developing." Each slide contained placeholders for three icons, with text describing a compound concept below each placeholder. Participants were asked to design the icons using two different approaches. One approach ("traditional pipeline") was to use online icon repositories (e.g., Google, the Noun Project) and a familiar tool such as PowerPoint to lay out the icons. The other approach was to use our ICONATE system. We randomized which approach participants were instructed to use first. Each session was self-guided and remotely conducted, in which participants were asked to record their comments on their overall experience.

Observations

Novices' ideation workflows using existing tools involved looking for icons by querying the given concepts. They directly used existing icons if their search query through standard icon repositories returned results (such as for "secure cloud"). Most participants selected icons based on style, choosing their favorite styles (N1, N6), or styles compatible with the rest of the slide (N7, N8, N9). If they were unable to find existing icons for the text query (e.g., the "eco tech education" in Figure 10c), some participants were stuck (N1), while others settled for a related icon that didn't perfectly capture the intended meaning. They initiated searches for semantically-related concepts to look for inspiration or reference material. They then combined the related icons manually to produce their final results. The icons being combined were not guaranteed to match in style. Novices prioritized an icon's semantics over its aesthetics. N2: *"since an icon is used for communication, we should make sure it conveys the correct meaning."*. N1: *"I know that the icons I designed are not good-looking, but I have no idea of what to do next, especially for layout."*

In contrast, ICONATE provided an easy-to-use tool without requiring extensive instruction or on-boarding. Participants felt our system would be generally useful for icon design. N6: *"I like the tool overall, I like being able to design an icon from components"*. N7: *"I thought the intent behind the tool is very helpful, being able to customize which visual features show up in the image ..."*. N9: *"This website is much better than the Noun Project, because I can customize my icon using this tool. It is really practical."* Participants were able to quickly generate icons they were satisfied with (< 1 min on average). Figure 11 shows example icons they created using our system, compared with those created using a traditional pipeline. Novices, who may not be adept at professional design tools, struggle to combine different icons together, especially if there are style differences. ICONATE can automate these difficult design steps for novice users, while giving them extra opportunities for customization by making iteration through style and space variations one click of a button away.

DISCUSSION AND LIMITATIONS

Scaling up the repository. Icons generated by ICONATE may not be sufficient for all user needs due to the limited size of our datasets. Occasionally users struggled to find a relevant icon suggestion. Users prefer to see diverse suggestions, not only for associated concepts, but also for style and layout. For example, a user (N2) wanted to use a bitcoin to represent "digital rental", but ICONATE could not support this query. However, our computational pipeline for synthesizing icons is not limited to a particular repository or icon style used. Many of our system components can be reused by substituting in a new set of training data. With larger annotated datasets, more diverse results are possible. With the help of the growing icon community (e.g., the Noun Project) and our icon annotation interface, more richly annotated datasets may be available in the near future.

Beyond additive operations. ICONATE generates compound icons via an additive operation on line icons. However, designer-generated icons can be the result of substractive operations, or more complex manipulations such as morphing the shape of one icon to fit another icon (Figure 12a). Furthermore, designers often adjust the rotation and color of the constituent icons, which our system does not currently support.

Beyond current conventions. ICONATE treats icons from the Noun Project as conventions since professional designers regularly use this website for their design inspiration. However, blindly reusing icons or templates from one compound icon in another can lead to unrecognizable concepts (e.g., reusing the "sun" from "hot weather" to represent "hot" in "hot food" in Figure 12b). In the last examples of Figure 12b, by blindly reusing space templates from existing icons, these icons becomes visually unbalanced. Automatically judging whether an icon is valid and aesthetically pleasing is a hard

(a) (b) (c)

Figure 10. User study results generated by novices using PowerPoint and the Noun Project. We asked participants to design icons for presentation slides with the prompt: "Imagine that you are using this slide to pitch a new product that you are developing."

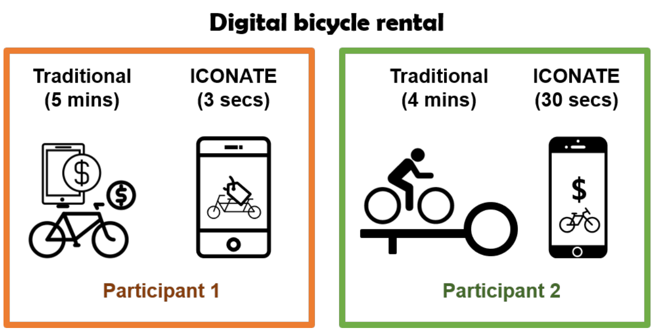

Figure 11. Results generated by novices using a traditional pipeline (i.e., PowerPoint + online repositories) and our ICONATE system. Time spent on the task is indicated above each icon. Our system can help generate satisfactory results more efficiently.

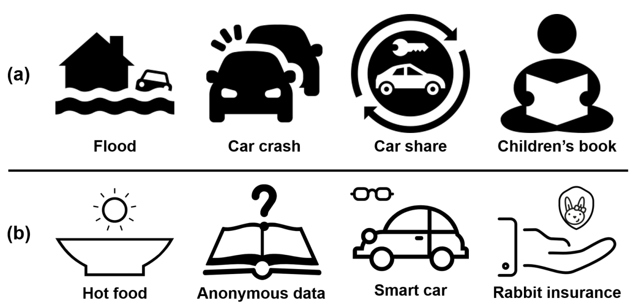

Figure 12. Limitations of our computational pipeline. (a) Icons that our pipeline can not generate because they involve deformation or negative space. (b) Icons that our pipeline did generate, but with problems with semantics (e.g., using "sun" to represent "hot" in the first icon) or layout (e.g., the last two examples).

and subjective task, requiring a contextualized understanding of semantics, and sensitivity to style and layout. Our work offers a step in this direction, but there is much future work to be done.

Towards adaptive suggestions. Another limitation of our pipeline is the inability to provide adaptive suggestions based on user input. For example, when a user adjusts an icon on the canvas by scaling or shifting, our system does not apply these modifications to the other generated variants. Future iterations of such a tool should respond to even more sophisticated user input, for instance by generating icons that are style-compatible with user-specified icons (e.g., a particular brand). Similarly, users may have company-sanctioned logos that they would like to incorporate into their icon design process. By accepting user-defined input and providing adaptive suggestions, future icon design systems may drive wider adoption of such automated tools.

CONCLUSION

ICONATE enables the automatic generation of icons for novel or unique compound concepts with minimal human effort (Figure 8). This can not only accelerate the design process of professionals, especially for ideation, but it can also lead to new workflows within existing design tools for novices, which currently offer only fixed libraries of icons. Logo and brand design tools, user interfaces, and slide presentation applications can all benefit from giving the user more opportunities for customization and content generation on the fly. Instead of struggling to find an icon from a library that is somewhat related to the concept the user is trying to convey, our tool can directly generate new icons from user queries. To the best of our knowledge, this is the first project that solves the problem of automatic compound icon generation. Together with other recent work [12, 22, 34, 45] it paints the way forward for machine learning-driven approaches to help people visualize complex and abstract concepts, create visual shades of meaning, and communicate in an effective, memorable, and universally-understandable ways. Used appropriately, icons are a visual language that can transcend spoken language and facilitate accessibility.

ACKNOWLEDGEMENTS

We thank Spandan Madan for some of the initial inspiration for this project, Anelise Newman for some helpful discussions, and anonymous reviewers for their insightful comments.

REFERENCES

[1] Accessed: 2019-09-17. Full Emoji List, v12.0. `https://unicode.org/emoji/charts/full-emoji-list.html`. (Accessed: 2019-09-17).

[2] Accessed: 2019-09-17. Microsoft PowerPoint(v16.27). (Accessed: 2019-09-17).

[3] Scott Bateman, Regan L Mandryk, Carl Gutwin, Aaron Genest, David McDine, and Christopher Brooks. 2010. Useful junk?: the effects of visual embellishment on comprehension and memorability of charts. In *Proceedings SIGCHI*. ACM, 2573–2582.

[4] Michelle A Borkin, Zoya Bylinskii, Nam Wook Kim, Constance May Bainbridge, Chelsea S Yeh, Daniel Borkin, Hanspeter Pfister, and Aude Oliva. 2015. Beyond memorability: Visualization recognition and recall. *IEEE TVCG* 22, 1 (2015), 519–528.

[5] Michelle A Borkin, Azalea A Vo, Zoya Bylinskii, Phillip Isola, Shashank Sunkavalli, Aude Oliva, and Hanspeter Pfister. 2013. What makes a visualization memorable? *IEEE TVCG* 19, 12 (2013), 2306–2315.

[6] Bill Buxton. 2010. *Sketching user experiences: getting the design right and the right design*. Morgan kaufmann.

[7] Zoya Bylinskii, Nam Wook Kim, Peter O'Donovan, Sami Alsheikh, Spandan Madan, Hanspeter Pfister, Fredo Durand, Bryan Russell, and Aaron Hertzmann. 2017. Learning Visual Importance for Graphic Designs and Data Visualizations. In *Proceedings UIST (UIST '17)*. ACM, New York, NY, USA, 57–69. DOI: http://dx.doi.org/10.1145/3126594.3126653

[8] SHI-KUO CHANG. 1987. Icon Semantics a Formal Approach tO Icon System Design. *IJPRAI* 01, 01 (1987), 103–120. DOI: http://dx.doi.org/10.1142/S0218001487000084

[9] SHI-KUO CHANG, MICHAEL J. TAUBER, BING YU, and JING-SHENG YU. 1988. The Sillicon Complier an Icon-oriented System Generator. *IJPRAI* 02, 02 (1988), 241–273. DOI: http://dx.doi.org/10.1142/S0218001488000169

[10] SHI-KUO CHANG, G. TORTORA, BING YU, and A. GUERCIO. 1987. Icon Purify Toward a Formal Theory of Icons. *IJPRAI* 01, 03n04 (1987), 377–392. DOI: http://dx.doi.org/10.1142/S0218001487000254

[11] S. . Cheng, W. Chen, and H. Sundaram. 1998. Semantic visual templates: linking visual features to semantics. In *Proceedings ICIP*. 531–535 vol.3. DOI: http://dx.doi.org/10.1109/ICIP.1998.727321

[12] Lydia B Chilton, Savvas Petridis, and Maneesh Agrawala. 2019. VisiBlends: A flexible workflow for visual blends. In *Proceedings SIGCHI*. ACM, 172.

[13] Marianela Ciolfi Felice, Nolwenn Maudet, Wendy E. Mackay, and Michel Beaudouin-Lafon. 2016. Beyond Snapping: Persistent, Tweakable Alignment and Distribution with StickyLines. In *Proceedings UIST (UIST '16)*. ACM, New York, NY, USA, 133–144. DOI: http://dx.doi.org/10.1145/2984511.2984577

[14] Aaron James Draplin. 2016. *Draplin Design Co.: Pretty Much Everything* (1 ed.). Abrams.

[15] Michael Evamy. 2015. *Logo: The Reference Guide to Symbols and Logotypes*. Laurence King Publishing.

[16] C. Ailie Fraser, Mira Dontcheva, Holger Winnemöller, Sheryl Ehrlich, and Scott Klemmer. 2016. DiscoverySpace: Suggesting Actions in Complex Software. In *Proceedings DIS (DIS '16)*. ACM, New York, NY, USA, 1221–1232. DOI: http://dx.doi.org/10.1145/2901790.2901849

[17] Elena Garces, Aseem Agarwala, Diego Gutierrez, and Aaron Hertzmann. 2014. A similarity measure for illustration style. *ACM TOG* 33, 4 (2014), 93.

[18] Steve Haroz, Robert Kosara, and Steven L Franconeri. 2015. Isotype visualization: Working memory, performance, and engagement with pictographs. In *Proceedings SIGCHI*. ACM, 1191–1200.

[19] Jon Hicks. 2011. *The Icon Handbook*. Five Simple Steps Ltd.

[20] William K Horton. 1994. *The icon book: Visual symbols for computer systems and documentation*. John Wiley & Sons, Inc.

[21] Ghita Jalal, Nolwenn Maudet, and Wendy E. Mackay. 2015. Color Portraits: From Color Picking to Interacting with Color. In *Proceedings SIGCHI (CHI '15)*. ACM, New York, NY, USA, 4207–4216. DOI: http://dx.doi.org/10.1145/2702123.2702173

[22] Sujay Khandekar, Joseph Higg, Yuanzhe Bian, Chae Won Ryu, Jerry O Talton Iii, and Ranjitha Kumar. 2019. Opico: A Study of Emoji-first Communication in a Mobile Social App. In *Companion Proceedings of The 2019 World Wide Web Conference*. ACM, 450–458.

[23] Nathalie Cindy Kuicheu, Ning Wang, Gile Narcisse Fanzou Tchuissang, Francois Siewe, and De Xu. 2012. Description logic based icons semantics: An ontology for icons. In *IEEE ICSP*, Vol. 2. IEEE, 1260–1263.

[24] Ranjitha Kumar, Jerry O. Talton, Salman Ahmad, and Scott R. Klemmer. 2011. Bricolage: Example-based Retargeting for Web Design. In *Proceedings SIGCHI (CHI '11)*. ACM, New York, NY, USA, 2197–2206. DOI:http://dx.doi.org/10.1145/1978942.1979262

[25] Manuel Lagunas, Elena Garces, and Diego Gutierrez. 2019. Learning icons appearance similarity. *Multimedia tools and applications* 78, 8 (2019), 10733–10751.

[26] Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, and Scott R. Klemmer. 2010. Designing with Interactive Example Galleries. In *Proceedings SIGCHI (CHI '10)*. ACM, New York, NY, USA, 2257–2266. DOI:http://dx.doi.org/10.1145/1753326.1753667

[27] Jianan Li, Tingfa Xu, Jianming Zhang, Aaron Hertzmann, and Jimei Yang. 2019a. LayoutGAN: Generating Graphic Layouts with Wireframe Discriminator. In *ICLR*. https://openreview.net/forum?id=HJxB5sRcFQ

[28] Jianan Li, Jimei Yang, Aaron Hertzmann, Jianming Zhang, and Tingfa Xu. 2019b. LayoutGAN: Generating Graphic Layouts with Wireframe Discriminators. *arXiv:1901.06767* (2019).

[29] Rungtai Lin. 1994. A Study of Visual Features for Icon Design. *Design Studies* 15, 2 (1994), 185 – 197. DOI: http://dx.doi.org/https://doi.org/10.1016/0142-694X(94)90024-8

[30] Yiming Liu, Aseem Agarwala, Jingwan Lu, and Szymon Rusinkiewicz. 2016. Data-driven iconification. In *Proceedings of the Joint Symposium on Computational Aesthetics and Sketch Based Interfaces and Modeling and Non-Photorealistic Animation and Rendering*. Eurographics Association, 113–124.

[31] Logo Design Love. A Guide to Creating Iconic Brand Identities. (????).

[32] Kurt Luther, Amy Pavel, Wei Wu, Jari-lee Tolentino, Maneesh Agrawala, Björn Hartmann, and Steven P. Dow. 2014. CrowdCrit: Crowdsourcing and Aggregating Visual Design Critique. In *Proceedings CSCW (CSCW Companion '14)*. ACM, New York, NY, USA, 21–24. DOI:http://dx.doi.org/10.1145/2556420.2556788

[33] Wakako Machida and Takayuki Itoh. 2011. Lyricon: A visual music selection interface featuring multiple icons. In *2011 15th International Conference on Information Visualisation*. IEEE, 145–150.

[34] Spandan Madan, Zoya Bylinskii, Matthew Tancik, Adrià Recasens, Kimberli Zhong, Sami Alsheikh, Hanspeter Pfister, Aude Oliva, and Fredo Durand. 2018. Synthetically Trained Icon Proposals for Parsing and Summarizing Infographics. *arXiv preprint arXiv:1807.10441* (2018).

[35] Nolwenn Maudet, Ghita Jalal, Philip Tchernavskij, Michel Beaudouin-Lafon, and Wendy E. Mackay. 2017. Beyond Grids: Interactive Graphical Substrates to Structure Digital Layout. In *Proceedings SIGCHI (CHI '17)*. ACM, New York, NY, USA, 5053–5064. DOI: http://dx.doi.org/10.1145/3025453.3025718

[36] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv:1301.3781* (2013).

[37] David R. Millen. 2000. Rapid Ethnography: Time deepening strategies for HCI field research. *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques* (2000), 280–286.

[38] George A Miller. 1995. WordNet: a lexical database for English. *Commun. ACM* 38, 11 (1995), 39–41.

[39] Peter O'Donovan, Aseem Agarwala, and Aaron Hertzmann. 2015. DesignScape: Design with Interactive Layout Suggestions. In *Proceedings of SIGCHI (CHI '15)*. ACM, New York, NY, USA, 1221–1224. DOI: http://dx.doi.org/10.1145/2702123.2702149

[40] Peter O'Donovan, Jānis Lībeks, Aseem Agarwala, and Aaron Hertzmann. 2014. Exploratory Font Selection Using Crowdsourced Attributes. *ACM TOG* 33, 4, Article 92 (July 2014), 9 pages. DOI: http://dx.doi.org/10.1145/2601097.2601110

[41] Yuting Qiang, Yanwei Fu, Yanwen Guo, Zhi-Hua Zhou, and Leonid Sigal. 2016. Learning to generate posters of scientific papers. In *Proceedings AAAI*.

[42] Daniel Ritchie, Ankita Arvind Kejriwal, and Scott R. Klemmer. 2011. D.Tour: Style-based Exploration of Design Example Galleries. In *Proceedings UIST (UIST '11)*. ACM, New York, NY, USA, 165–174. DOI: http://dx.doi.org/10.1145/2047196.2047216

[43] Gabriele Rohr. 1986. *Using Visual Concepts*. Springer, Boston, MA, 325–348. DOI: http://dx.doi.org/10.1007/978-1-4613-1805-7_14

[44] Vidya Setlur, Conrad Albrecht-Buehler, Amy A. Gooch, Sam Rossoff, and Bruce Gooch. 2005. Semanticons: Visual metaphors as file icons. In *Computer Graphics Forum*, Vol. 24. Wiley Online Library, 647–656.

[45] Vidya Setlur and Jock D. Mackinlay. 2014. Automatic Generation of Semantic Icon Encodings for Visualizations. In *Proceedings SIGCHI*. ACM, 541–550. DOI:http://dx.doi.org/10.1145/2556288.2557408

[46] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

[47] Maarten Someren, Yvonne Barnard, and Jacobijn Sandberg. 1994. *The Think Aloud Method - A Practical Guide to Modelling CognitiveProcesses*.

[48] Amanda Swearngin, Mira Dontcheva, Wilmot Li, Joel Brandt, Morgan Dixon, and Andrew J. Ko. 2018. Rewire: Interface Design Assistance from Examples. In *Proceedings SIGCHI (CHI '18)*. ACM, New York, NY, USA, Article 504, 12 pages. DOI: http://dx.doi.org/10.1145/3173574.3174078

[49] Anna Szekely, Thomas Jacobsen, Simona D'Amico, Antonella Devescovi, Elena Andonova, Daniel Herron, Ching Ching Lu, Thomas Pechmann, Csaba Pléh, Nicole Wicha, and others. 2004. A new on-line resource for psycholinguistic studies. *Journal of memory and language* 51, 2 (2004), 247–250.

[50] Alex Taikh, Ian S Hargreaves, Melvin J Yap, and Penny M Pexman. 2015. Semantic classification of pictures and words. *The Quarterly Journal of Experimental Psychology* 68, 8 (2015), 1502–1518.

[51] Kashyap Todi, Daryl Weir, and Antti Oulasvirta. 2016. Sketchplore: Sketch and Explore with a Layout Optimiser. In *Proceedings DIS (DIS '16)*. ACM, New York, NY, USA, 543–555. DOI: http://dx.doi.org/10.1145/2901790.2901817

[52] Haijun Xia, Bruno Araujo, and Daniel Wigdor. 2017. Collection Objects: Enabling Fluid Formation and Manipulation of Aggregate Selections. In *Proceedings SIGCHI (CHI '17)*. ACM, New York, NY, USA, 5592–5604. DOI: http://dx.doi.org/10.1145/3025453.3025554

[53] Nanxuan Zhao, Ying Cao, and Rynson WH Lau. 2018. What characterizes personalities of graphic designs? *ACM TOG* 37, 4 (2018), 116.

[54] Xinru Zheng, Xiaotian Qiao, Ying Cao, and Rynson W. H. Lau. 2019. Content-aware Generative Modeling of Graphic Design Layouts. *ACM TOG* 38, 4, Article 133 (July 2019), 15 pages. DOI: http://dx.doi.org/10.1145/3306346.3322971

[55] C Lawrence Zitnick and Devi Parikh. 2013. Bringing semantics into focus using visual abstraction. In *IEEE Proceedings CVPR*.

[56] C Lawrence Zitnick, Ramakrishna Vedantam, and Devi Parikh. 2016. Adopting abstract images for semantic scene understanding. *IEEE TPAMI* 38, 4 (2016), 627–638.