Juxtapoze: Supporting Serendipity and Creative Expression in Clipart Compositions

William Benjamin,[1] Senthil Chandrasegaran,[1] Devarajan Ramanujan,[1]
Niklas Elmqvist,[2] SVN Vishwanathan,[3,4] and Karthik Ramani[1,2]
[1]School of Mechanical Engineering, [2]School of Electrical and Computer Engineering,
[3]Department of Statistics, and [4]Department of Computer Science
Purdue University, West Lafayette, IN 47907, USA
{william, senthil, dramanuj, elm, vishy, ramani}@purdue.edu

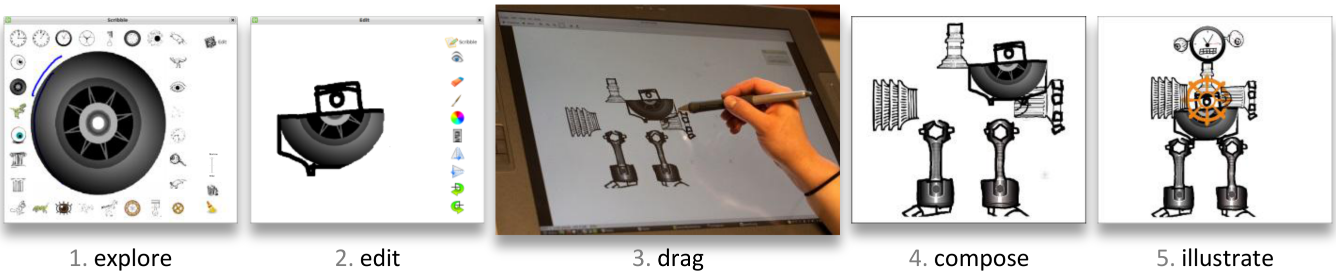

Figure 1. The Juxtapoze workflow (from left to right): (1) exploring the shape database using scribbled input; (2) edit a selected shape using standard drawing operations; (3) drag the shape into position on the canvas; (4) compose it with other shapes; and (5) repeat for a full drawing.

ABSTRACT
Juxtapoze is a clipart composition workflow that supports creative expression and serendipitous discoveries in the shape domain. We achieve creative expression by supporting a workflow of searching, editing, and composing: the user queries the shape database using strokes, selects the desired search result, and finally modifies the selected image before composing it into the overall drawing. Serendipitous discovery of shapes is facilitated by allowing multiple exploration channels, such as doodles, shape filtering, and relaxed search. Results from a qualitative evaluation show that Juxtapoze makes the process of creating image compositions enjoyable and supports creative expression and serendipity.

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation: User Interfaces—*Interaction styles*; I.3.6 Computer Graphics: Methodology and Techniques—*Interaction techniques*

Author Keywords
serendipity; sketching; shape search; clipart composition; creative expression;

INTRODUCTION
Advances in digital technology are changing the traditional notions of "creating"; the new generation of creators copy, recombine, divide, compose, remix, fork, mash up, and fuse.

CHI 2014, April 26–May 1, 2014, Toronto, Ontario, Canada.

http://dx.doi.org/10.1145/2556288.2557327

The idea of creating something from existing material is ubiquitous:[1] programmers build on existing libraries, people retweet other tweets on Twitter [4], and both amateur and professional composers take pieces of existing songs and splice them together into remixes [15]. As a result, the distance between producers and consumers of digital artifacts is decreasing [5], and non-experts are becoming more empowered and interested in means for creative self-expression.

One such means of self-expression is clipart compositing, an extension of photomontages [1], which consists of arranging multiple clipart elements into a single picture. Clipart composition is popular among novice users, as cliparts are more amenable to editing and composition. Many existing systems can compose images from individual elements [2, 7, 13, 14, 17, 20]. These systems are powerful means of creating photorealistic montages that precisely match the user inputs. Professional tools such as Photoshop, Illustrator, and CorelDraw can also be used to compose clipart into a montage. However, while these tools are powerful and widely used, they do not address two key aspects of clipart montages as an art form: (1) *diversity* as well as (2) *expression* in creative composition.

In terms of diversity, existing systems are targeted towards minimizing the number of possibly irrelevant images while maximizing the number of possibly relevant results. In contrast, the hallmark of creative artistic production is high diversity. Montages particularly use incongruous juxtapositions of an object and its positioning in a landscape that is strange to it, but with which unexpected analogies form [1]. The key essence of a montage is the inherent new analogies between diverse shapes that may not necessarily be related.

[1]`http://everythingisaremix.info/`

The second key aspect is the creative expression facilitated by the ability to not just compose existing imagery, but also modify and transform the constituent parts. Existing systems [2, 7, 14, 20] focus on *copy* and *combine* through the retrieval, cutting, and blending of photos into a photorealistic single image. However, the *transform* step, where the artist can make substantial changes to the image, is essential since it adds to the originality of the artwork [16].

We present *Juxtapoze*, a clipart composition workflow that supports serendipity and creative expression by allowing users scribble, search, and compose clipart into interesting combinations. (Figure 1). The fact that users utilize scribble input to search for desired clipart means that the workflow requires little additional cognitive overhead, as well as allows for serendipitous discoveries based on shapes without requiring premature fixation on named object classes. We call this chance discovery of shapes *shape serendipity*. We support shape serendipity by providing three pathways for exploration: (1) Scribbled sketches help constrain the shape the user is looking for. For example, an artist may be looking for a figure consisting of two side-by-side circles to fit the available space in the illustration and to balance its visual design, but searching for such a figure requires the user to prematurely fixate on a specific class of objects such as bikes, when perhaps a skull with empty eye sockets would have fit the theme better. Serendipity is further supported by a (2) *creativity slider* that allows users to control the diversity of results from exact matches to a broad variety. Finally, a (3) *shape filter* allows users to discover relevant shapes within a class of clipart visually similar to the user preferences. Finally, after settling on a shape from the search results, the user is given the opportunity to modify it, which includes not only scaling, translation, and rotation, but also erasing desired parts and adding new content using paint, fill, and airbrush tools.

We have built a prototype of the Juxtapoze technique (Figure 1) that provides a practical implementation of the workflow. Using this prototype, we performed a qualitative user study where non-artist participants were given tasks to create both free illustrations, guided, and replicated illustrations in a limited time (10 minutes per task). Feedback from this study caused us to introduce several new features to Juxtapoze, including the creativity slider and shape filter discussed above, which we then evaluated in a second qualitative study focusing on serendipity and creative expression. We found strong evidence supporting our hypothesis that Juxtapoze does indeed scaffold the creative process in clipart composition.

We thus claim the following contributions in this work: (1) a fast 2D shape search engine designed to support serendipitous discovery of parts using sketches, a shape filter, and a creativity slider; (2) a creativity workflow for searching, selecting, and modifying clipart before composing them into an illustration; (3) a practical desktop implementation of query-based sketch composition called Juxtapoze; and (4) results from two qualitative user studies confirming our goals of supporting creative expression and serendipity have largely been met. The results suggest that non-artists perform surprisingly well in creating detailed illustrations.

BACKGROUND

Here we review the opportunities and challenges in content authoring, and will then focus specifically on drawing, design, and image composition. Finally, we review the concept of serendipity and discuss ways of inducing it in our work.

Sketching and Drawing

Sketching is a ubiquitous ideation activity in nearly all creative disciplines [6], and most sketches are visual in nature. In general, the purpose of such an informal and visual medium is to generate several, potentially many, design alternatives for a particular idea to avoid premature fixation on any single one. In particular, freehand drawing is an important basic tool for engineers and designers [32].

Starting from SketchPad [30], digital drawing and painting has always been a prime target for interactive systems [29]. A recent example of a digital drawing aid is Vignette [19], which supports drawing by automatically filling the drawing canvas with pen-and-ink textures based on user input. Recently, there has be great interest in 3D shape assembly using sketch based search [34] constrainted by contextual information of component parts Methods which use sketch based retrieval [9, 10] are targeted towards minimizing the number of irrelevant objects rather than maximizing the number of possibly relevant objects. This does not support serendipitous discoveries [24]. The existing work that is most closely related to our work is ShadowDraw [21] which is an interactive drawing system that guides free form drawing by displaying a dynamically updated blended image behind the user's sketch. The blended image is created by retrieving precise matches to the current user sketch. In contrast, Juxtapoze retrieves a set of diverse images in support of serendipity.

Drawing by Composition

Constructing new digital drawings by juxtaposing existing images is a common way to create new visual content. Current graphical tools such as Microsoft PowerPoint, Adobe Photoshop, and CorelDraw excel at this interaction. Using an image search engine, such as Google Images or Bing Images, even a non-artist can use this method to create new artwork. However, this approach will not leave the user much creative expression since the component images tend to be used in their original form, if only scaled and rotated.

Very few systems use clipart images to form composites. The closest example is Sketch2Cartoon [33], which uses sketch input to enable targeted retrieval of exact matches. In contrast, Juxtapoze is designed for dynamic exploration where the primary focus is to support creativity.

Since photomontages are related by the nature of image compositing, we review some of the work relating to photomontages. Many existing research systems can compose a photomontage from photographs [2, 7, 14, 20]. These systems are powerful means of creating photorealistic montages that precisely match the user inputs. Again, these tools do not address the requirements of diversity and serendipity since their focus is to find the best matches photometrically and geometrically. For example, Lalonde et al. [20] present a system with algorithms for improving object segmentation, blending, and estimating 3D object size and orientation. Chen et al. [7] use sketches and tags to automatically select images that are closest matches and blend them to create a photomontage. Interactive photomontages [2] describes a workflow for identifying good seams and fusing gradients to compose several photographs into one. Gal et al. [13] propose a method to create 3D collages by geometrically fitting individual elements into a target shape. Similarly, very recent work on cliplets [18] explore the use of juxtaposition to combine still and dynamic imagery into miniature videos. Our Juxtapoze workflow goes beyond this prior art by integrating modification of the visual components and support for serendipitous discovery.

Serendipity and Exploration

In response to the ever-growing ocean of information that us humans combat every day, current information retrieval mechanisms such as Google Search have evolved for returning near-exact matches. However, this runs the danger of entirely losing the benefit of *serendipitous* findings: unexpected yet valuable discoveries that are tangential or completely unrelated to a query [11, 12]. Our decision to support serendipitous discovery is grounded on the fact that artists tend to browse existing repositories for creative simulation [3, 8].

There is more to serendipity than mere chance and coincidence. This could be related to *synchronicity*: simultaneous manifestations of related ideas that seem random, but become meaningful [22]. While browsing a repository, serendipity can be stimulated by maximizing the user's ability to explore [28]. Free exploration supports serendipitous discoveries more than focused tasks [11]. To achieve this, Thudt et al. [31] recommend explicitly designing for serendipity using an exploratory approach to browsing information that entices curiosity, promotes playfulness, and supports multiple pathways through a dataset. Designing for free but relevant exploration is particularly useful for serendipitous discoveries.

JUXTAPOZE

Juxtapoze is a digital drawing workflow that seamlessly integrates shape search for composition into the artistic process. In this section, we first examine the design considerations that shaped our work. We then present the Juxtapoze workflow itself, including its user interface.

Design Guidelines

One of the key challenges identified in our literature review is that most users lack the time, training, talent, and tools to create original artwork from scratch. For this reason, we made the fundamental decision to design Juxtapoze as an artistic workflow based on *composition* of existing visual components. However, drawing upon our review of existing artistic workflows, we derived a set of additional design guidelines for the novel Juxtapoze method:

- G1 **Interactive response** to enable quick exploration and iteration over a variety of shapes in the shape database [31];
- G2 **Easy personalization of shapes** to scaffold creative expression and allows for creating original content [16];
- G3 **Broad search results** to help the user to explore the database, thus supporting serendipitous discovery [11, 24];
- G4 **Immediate feedback** to communicate the behavior of the system and improve user engagement [21];
- G5 **Multiple pathways for exploration** to further support serendipitous discoveries [31]; and
- G6 **Easy juxtaposition** to achieve creative outcomes through translation and scaling [16].

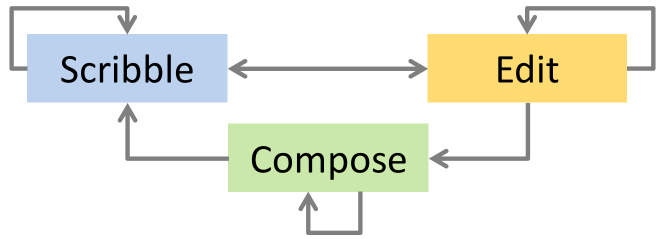

Figure 2. The Juxtapoze workflow: ***scribble, edit, and compose*****.**

General Workflow

Based on these guidelines, we designed the system to consist of a shape search engine that supports editing of the selected shapes and juxtaposes the resulting shapes to form a composite drawing. For user interface elements that satisfy a particular design goal, we indicate the goal in bold letters.

The Juxtapoze workflow can be represented by a state model consisting of three mutually exclusive states (Figure 2):

- **State 1 – Scribble:** The users scribble on the screen and the system quickly suggests relevant shapes from the database (**G1**). The suggestions panel surrounding the sketcher is populated with retrieved results. The user selects a particular shape by double-clicking on its icon. When the user finds a shape that matches their intention, they can progress to the Edit state following it. The corresponding interface element is a scribbler (Figure 3(b)).
- **State 2 – Edit:** Here the user customizes the chosen shape using conventional drawing tools such as a paintbrush, eraser, scaling, rotation, and different colors. From this state, the user can either go back to the Scribble state to select another shape from the search results, or proceed to the next state (State 3 – Compose) if satisfied.
  The editor contains tools for manipulating the selected image (Figure 3(c)) through drawing and erasing, as well as geometric transformations for scaling, rotating, and translating. This allows for easily personalizing shapes (**G2**).
- **State 3 – Compose:** The Compose state allows the user to move the finished sketch to a suitable location on the canvas (**G6**). In case the user is not satisfied with the outcome, she can switch back to the Scribble state. In this state the edited shape is shown as image floating on the canvas. As the user hovers over the shape, the draggable corners, which can be used to resize the image, are shown.

This design cycle iterates, each time generating a new scriblet, until the user is satisfied with the final illustration. While serendipitous discovery can be induced using various inputs like color, visual style and keywords, we choose to focus on *shape serendipity* because shape input is the primary means of designing shapes [6, 32].

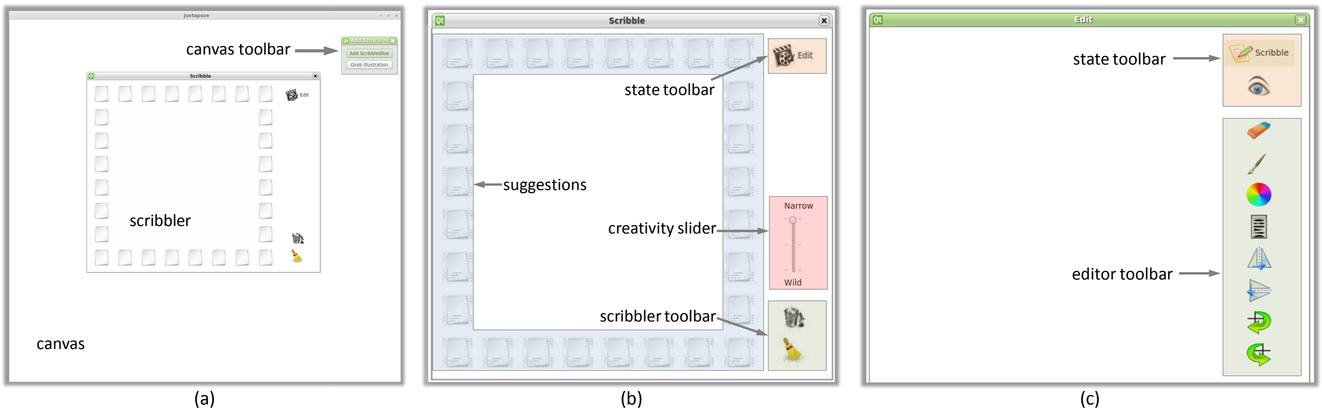

Figure 3. (a) The Juxtapoze canvas presented to the user in the beginning of the session. (b) The scribbler input panel corresponding to the *scribble* state for specifying scribble queries. Once the user scribbles, the results are updated in the suggestions panel. The *creativity slider* is used to control the diversity of the results. Double-clicking the results in the suggestions panel activates the *shape filter*. (c) The editor panel corresponding to the *edit* state with editing toolbars for operations like erasing, drawing, coloring, flipping, rotating, etc.

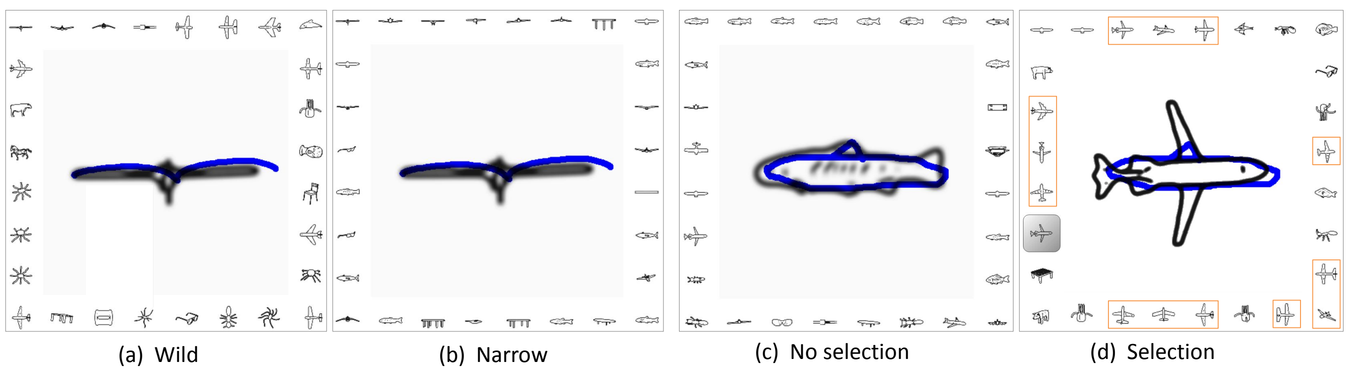

Figure 4. (a) & (b) Query results corresponding to the same query when the creativity slider is set to 'Wild' vs. 'Narrow'. Setting the creativity slider to 'Wild' produces visually diverse but relevant results, while setting it to 'Narrow' produces relevant results that closely match the input scribble. (c)& (d) Query results corresponding to the same query when the *shape filter* is inactive vs. when it is active. Selecting the 'airplane' from the results in (c) activates the *shape filter*, which suggests similar 'airplanes' matching the scribble as shown in (d).

Multiple Pathways of Querying

The Juxtapoze method supports three different pathways to support *shape serendipity* (**G5**):

1. **Scribble Search:** As the user sketches, the topmost match is shown as a blurred image in the background (Figure 5), thus providing immediate visual feedback (**G4**). A typical workflow consists of the user doodling simple strokes and the system suggesting matches (Figure 6). Scribble allow for constraining the search in the visual domain.
2. **Creativity Slider:** As discussed earlier, finding interesting shapes serendipitously can be facilitated through allowing exploration through the dataset of shapes. Exploration through the dataset can be achieved by increasing the variety of shapes displayed in the suggestion panel. The design goal of the creativity slider is to provide the user with greater control on the diversity of the results (**G3**). Given a limited space for display of the possible building blocks for the illustration, the amount of creative diversity of the outcome decreases if a similar shape is already presented to the user. For example, an additional apple when the user is looking for apple-shaped objects will have lesser value since an apple has already been presented to the user. This does not mean that the additional apple is not what the user was looking for. It points to an idea of diminishing addition to the creativity of the illustration. The attractiveness of every additional identical result decreases as the number of results already added to the suggestion panel increases. Setting the slider to 'narrow' shows the shapes that are very closely related to the sketch. These shapes have very low visual variation from the original sketch. Setting the slider to 'wild' shows shapes that have a greater variety (Figure 4 (a) and (b)). The suggestions panel is updated every time the slider is moved to a new position.
3. **Shape Filter:** The shape filter provides a way to constrain the shapes within categories of interest. During our discussion with users of the system, we found that narrowing down on a particular variety of shapes helps the user to choose from a particular category of shapes. For example, the user finds an airplane in the search results and would like to explore more airplanes instead of all the shapes that match the sketch. In order to specify that choice, the user can pin down the shapes similar to the displayed airplane

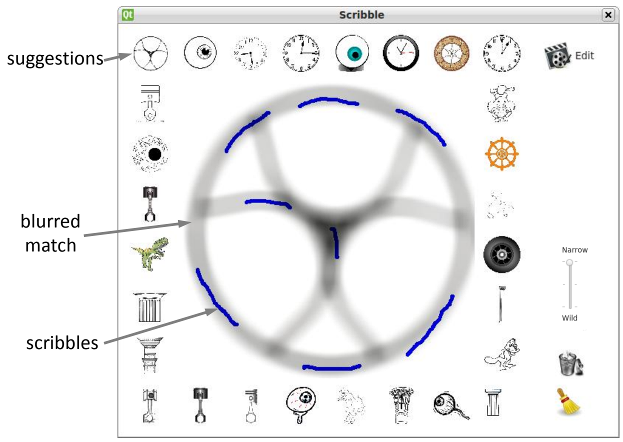

Figure 5. A scribbler displaying query results after the user has provided sketch input. The blurred background image in the center represents the currently highest matched shape in the database, whereas the images around its perimeter are currently matched search results.

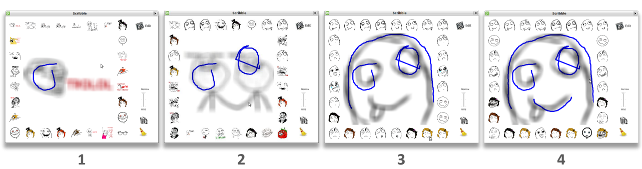

Figure 6. The progression of search results from a shape database containing images from rage comics.

by double-clicking the shape. This gives a higher preference to searching within the neighborhood of shapes visually similar to airplanes. The shape can be deselected and removed from the pipeline by double-clicking again (Figure 4 (c) and (d)).

SYSTEM IMPLEMENTATION

The Juxtapoze prototype application is implemented in C++ using the Qt Framework library and OpenCV. Because Qt and OpenCV are cross-platform, our application can run on any operating system that are both supported by Qt and OpenCV. In this section we describe our shape search engine that supports the Juxtapoze workflow presented above.

Modular Shape Database

Clipart images are easy to extract, vectorize and compose, and have been widely used for making illustrations. Thus, clipart can be easily used by non-artists for creative expression. Our database consists of clipart and line drawings of a variety of shapes, both natural and man-made. Examples of object categories in our database include fruits, animals, insects, mechanical parts, road signs, faces, etc.

Preprocessing

All shapes are preprocessed before being added to our shape database. This subsection describes this process.

Sketch-based Descriptor

In order to search images using sketches, we need to compute descriptors that encode the location, angle, and length of edges. To this end, we have implemented an optimized version of Binary Coherent Edge descriptors (BiCE) [35]. Each image is divided into patches using an overlaying grid. The final output of the BiCE algorithm is a binarized edge histogram for each patch, which encodes the position, orientation, and local linear length of an edges in the patch. The image descriptor is an aggregate of the patch descriptors. We optimized the computation of the BiCE descriptor by introducing the following optimizations:

- **Parallel Processing:** Since each image is divided into patches, we can compute the descriptors for each patch in a parallel way. We used multithreading to implement a parallel pipeline, yielding a substantial speed increase.
- **Robustness to Sketch Variations:** We introduced a 50% overlap between adjacent patches to accommodate variations in sketch input for searching.

Image-based Descriptor

To support shape filtering, we need to search in the neighborhood of the user preferences. To this end, we use the popular SIFT descriptors [23], which are scale and rotation invariant. Our goal is to be able to quickly query the database to get visually similar images to the input image. We use the Bag of Words approach, which computes a descriptor based on the visual words present in the image. Fast retrieval is achieved by finding k-nearest neighbors using FLANN. [25]

Creativity Slider

Two shapes that are equally relevant to a particular sketch input may have different effect on the overall creative benefit. This points to a notion of diminishing utility of additional identical shapes. Submodular functions are known to capture the effects of diminishing utility fairly accurately [26].

Submodularity

Given a set S, a subset D, and elements $x,y \in S \setminus D$, a submodular function $f : \{0,1\}^S \rightarrow \mathbb{R}$ satisfies the following diminishing returns property:

$$f(D \cup \{x\} - f(D) \geq f(D \cup \{x,y\}) - f(D \cup \{y\}). \quad (1)$$

In particular, let D be a collection of shapes, and for each $d \in D$ let u_d denote its query descriptor. Furthermore, let $u_D := \sum_{d \in D} u_d$. Given a query descriptor u, we define a *query coverage* function as

$$\zeta_u(D) = \sum_j u_j \rho_j(u_D) \quad (2)$$

where u_j is the j-th feature of the query descriptor, and $\rho_j(\cdot)$ is a monotonically non-decreasing concave function which operates on the j-th component of its input. One can show that the ζ_u defined above is submodular, and captures how well the features of u are covered by the same features in the collection D. For Juxtapoze, we set σ to be the probabilistic cover function given by

$$\sigma(z) = 1 - e^{-\theta z} \text{ with } \theta > 0. \quad (3)$$

We use the slider to control the value of θ. For large values of theta, $\sigma(z)$ saturates quickly and we select the shapes as they are ranked. This setting corresponds to *narrow* state of the slider, which gives most relevant and less diverse results.

For smaller values of theta $\sigma(z)$, the elements that maximize the difference in query coverage are selected in a sequence that maximizes the variety of shapes. This corresponds to the *wild* state of the slider.

Greedy Algorithm

While exact maximization of submodular functions is NP-hard, efficient greedy algorithms which obtain a $(1-1/e)$ approximation to the optimum exist [27]. The algorithm starts with an empty set D_0, and in iteration i selects the shape $a \in S \setminus D_{i-1}$ that maximizes the change in query coverage function.

$$D_i = D_{i-1} \cup \{\underset{a}{\operatorname{argmax}}\, \zeta_u(u_{D_{i-1}} + u_a) - \zeta_u(u_{D_{i-1}})\} \quad (4)$$

This process continues until all the locations in the suggestion panel are filled.

Diversity with respect to user sketch. To diversify the results, we have two choices depending on the choice of u:

- Diversifying the matching part of the sketch. This can be done by using the original descriptor u.
- Diversifying the non-matching part of sketch. In this case, we use the complement u' of u from Equation 2.

In our implementation, we diversify the results using the complement u' since we would like to have diverse results that are relevant to the sketch.

Shape Querying Pipeline

As the sketch is created, the search engine generates a BiCE decriptor for the sketch. Then, depending on whether user preferences are selected or not, the search engine choses one of the following pathways:

- **Without user preferences:** The sketch is used to query the BiCE database, retrieving the top 100 results.
- **With user preferences:** The user preferences are used to query the SIFT descriptor database, and $100 \times n$ results are retrieved, where n is the number of user preferences.

Finally, the retrieved queries are ranked based on the state of the *creativity slider*.

EVALUATION

We evaluated Juxtapoze using an experiment to study the accuracy of the algorithm as well as through two user studies. For the accuracy measurement, we simulated the scenario where a user randomly selects a database image and removes 50% of its ink. We found that in 79% of the cases, our retrieval engine was still able to rank the relevant image in one of the top 10 positions. The first user study, a formative user study, was aimed towards evaluating the scribble–edit–compose workflow described in Figure 2. Based on observations from this study, we developed the *creativity slider* and *shape filter* to support shape serendipity. The purpose of the second study was to evaluate the utility of these new features.

Methods

Both studies were performed on the same apparatus, and used similar tasks. This subsection describes the common aspects.

Apparatus and Software

We conducted both studies on a pen-based Wacom 21UX digital tablet connected to a Lenovo ThinkPad T530 running a Linux Mint 13 operating system.

Procedure

We chose our tasks with the intent of observing the progression from undirected, exploratory tasks to more directed tasks and how Juxtapoze affected the participants' workflow between the three activities of scribble, edit, and compose. Two tasks that were common to both studies were (1) a free, exploratory illustration task, and (2) a directed illustration task with a specific goal. In addition, the summative user study had a "creative expression" task that allowed for some exploration, yet was goal-oriented.

Participants were initially trained in Juxtapoze through a guided demonstration (approximately 5 minutes), followed by free practice of scribbling, editing, and composing. Participants used 2-10 minutes for this "free play" session. We logged the participants' activity using event triggers that recorded their transitions between scribble, edit, and compose activities for all the specified tasks. Our intent was to analyze these activities among users, and observe any patterns in behavior defined by a switching from one activity to another. After each task, participants were required to respond to a survey that included both Likert scale ratings and open-ended questions.

Study 1: Juxtapoze Flow

This study focused on participant activity and their use of the scribble–edit–compose loop when presented with both undirected and directed tasks.

Participants

For the formative study we recruited 9 paid participants (6 male, 3 female), aged between 19 and 26 years. All participants were either graduate or undergraduate students, and 8 of them were right-handed. Additionally, we recruited a professional industrial designer as an expert participant for a think-aloud protocol for both studies. The designer was proficient with Adobe Illustrator and Adobe Photoshop.

Tasks

For the formative study, participants used an earlier version of Juxtapoze *without* the search filtering and the creativity slider, and linked to a repository of 900 images.

- **Task 1 – Free illustration:** Participants were asked to create any illustration that they wished, sketching strokes to query and display images that they could use.
- **Task 2 – Interpretation of an existing illustration:** Participants were shown an illustration that had been created using Juxtapoze, and were asked to use the system to create their own interpretation of this illustration.

Results and Discussion

Juxtapoze scored high on user engagement, with a majority of participants (7 out of 9) finding the interface enjoyable for both tasks, and useful for serendipitous discovery for Task 1. One participant reported that Juxtapoze *"was great for suggesting images I hadn't thought to use and getting my ideas and creativity flowing."*

However, most participants felt that Juxtapoze offered little control over the diversity of the search results. Especially for Task 2, we observed that most of the participants were frustrated while searching for something specific, and 5 out of 9 reported finding the interface difficult to use. As one participant put it, *"the shapes and the inventory were too diverse and it would be better if a context could be defined before the shape creation process"*. Other usability issues included better control over the composition, especially the placement and scaling of images.

The formative study helped identify the less intuitive interactions that users had trouble with in Juxtapoze, and led us to make significant improvements such as more intuitive drag handles to move individual images, zoom controls for the canvas, and scaling images. However, the chief insight obtained from the participants was that while the diversity of images made the exploratory illustration in Task 1 fun, the same diversity of images made the more focused Task 2 difficult and frustrating. To address this issue, we added the *creativity slider* to provide user control over the diversity of query results, and added the *shape filter* to constrain these results within categories of interest. The next iteration of Juxtapoze sought to address this issue through these new features.

Study 2: Supporting Serendipitous Creativity

This study focused on the effect of the selection filter and the creativity slider in providing opportunities for serendipitous inspirations that would support creativity, embedded in the established scribble–edit–compose loop.

Participants

We recruited 21 paid participants (16 male, 5 female), aged between 18 and 30 years. Participants were either graduate or undergraduate students, one of whom was left-handed. Additionally, we recruited two sketching experts as participants for a think-aloud protocol based on the same three tasks as the normal users. One expert was the same industrial designer used for Study 1. The other, a Ph.D. student in design, is highly proficient in sketching for design.

Tasks

Participants used the new version of Juxtapoze for this study, which included the shape filter and creativity slider. For this study, Juxtapoze was linked to a larger repository of 3,000 images. We used the below tasks:

- **Task 1: Free illustration:** As with Study 1, participants were asked to create any illustration that they wished, with the creativity slider set to "wild", and the shape filter disabled, to provide a broader set of search results.
- **Task 2: Interpretation of an existing illustration:** The same as study 1, participants were shown the same illustration created with Juxtapoze, and asked to create their own interpretation of the illustration. The creativity slider was set to "narrow" to provide closer matches to queries.
- **Task 3: Creative Expression:** Participants were given a task to "create an illustration for a device to help ease your morning chores". They were given complete control of the creativity slider, and could use the shape filter as required.

Results and Discussion

Serendipity. A majority of participants found Juxtapoze enhanced serendipitous discovery (T1: 67%, T2: 71%, T3: 76%). For example, one participant said, *"What I intended to draw was not what I ended up drawing (in a good way)."; "As you draw, suggestions kept popping up that helped me add new elements and ideas to my illustration."*

Users specifically appreciated the multiple pathways for serendipitous discovery, i.e. sketches, the creativity slider, and the shape filter. A majority (62%) of the users agreed that the inclusion of the creativity slider offered more control over image diversity for Task 3. Users commented that *"I liked the need to create only a simple shape in order to obtain complex related figures."*; *"Using the creativity slider can lead to pictures that the user is not looking for but would like to appear in future selections or alternatively help in coming up with creative ideas."*; and *"I really liked the filter option, it gives you multiple options of similar images to choose from, that makes the task a lot easier."*

Creative Expression. For all of the three tasks a significant majority (T1: 76%, T2: 76%, T3: 67%) of the participants indicated that using the Juxtapoze system positively affected their ability for creative expression. For example, participants stated *"I had the freedom to draw what I had in mind and then match it with a similar image"*; *"I liked the idea behind the program and it helps inspire and increase creativity by giving suggestions"*; and *"Predicting images while drawing gave me new ideas to include in my illustration."*

Browsing Behavior. Overall, a notable majority of the users found Juxtapoze enjoyable (T1: 86%, T2: 81%, T3: 76%). Based on the survey responses, we observed a significant difference in the agreement levels of Task 1 (T1) and Task 2 (T2) with regards to efficacy (T1: 62%, T2: 33%), ease of use (T1: 71%, T2: 48%) and variety (T1: 48%, T2: 38%). These results are shown in Figure 7 and suggest that the Juxtapoze system is currently more suited for loosely-defined exploratory tasks than focused tasks.

Due to the high variation in usage patterns between participants, we constructed timeline plots of their activities from the interaction logs. Three such participant activity plots are shown in Figure 8. These were chosen based on the participant response to the surveys: P1 on the chart was a participant who responded favorably to Juxtapoze, while P2 had an overall negative response. The last row corresponds to the industrial designer mentioned earlier.

From the plots, we can see that P1 and the expert user had fewer transitions between activities, while P2 had more frequent transitions between scribble, edit, and compose. The staggered use of the creativity slider within relatively longer 'scribble' activities could indicate frustration with the query results. When compared to the expert's performance in Task 3, we observe far fewer activity transitions, and fewer, more focused use of the creativity slider toward the end of respective scribble activities. Interestingly, the highest use of the creativity slider among the three is P1, who indicated that he was searching for one specific object: an image of a battery. Such usage patterns indicate the need for resolving ambiguity in sketch-based queries by augmenting the system with metadata such as image tags and categories.

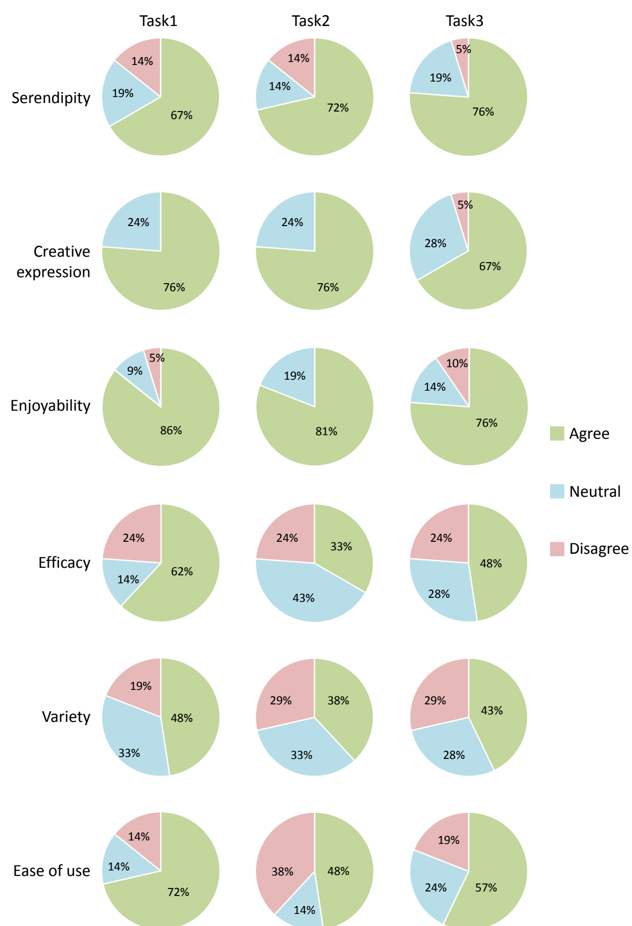

Figure 7. Distribution of agreeable, neutral and disagreeable user responses to post-task survey questions. The results have been aggregated into these three categories from a five point Likert scale response (strongly agree, agree, neutral, disagree, strongly disagree).

We provide a sampling of some of the compositions created by the participants in Figure 9.

DISCUSSION

The Juxtapoze workflow was created to support shape serendipity and creative expression for exploratory tasks. User performance and responses for Task 1 indicate that the current implementation of Juxtapoze addresses these design goals. Our hypothesis was confirmed by the two sketching experts, who agreed that the Juxtapoze system was particularly helpful for tasks related to undirected creativity. One expert commented *"Juxtapoze is enjoyable and a fun way of creating clip art scenes."*

Analyses of more focused tasks (i.e. Task 2 and Task 3) using our system reveal significant differences in usage patterns which leads us to believe that the definition of creativity takes a wholly different meaning for such tasks. We noticed that for more loosely defined tasks, users tend to be creative while using the Juxtapoze interface. They were more open towards diversity of query results and readily diverged from existing ideas through extensive image modification. On the other hand, for more defined tasks, users were mentally establishing a reference solution and attempting to illustrate this idea using very specific clipart. The ability of our system to be useful for creative expression in these tasks directly corresponded to the success that the users had in illustrating this reference solution. Users that were successful, moved on towards modifying their idea while the others chose to remain in the scribble state re-querying the clipart database.

This dichotomy in the workflow is illustrated in contrasting results and user comments for these tasks. While some users had comments similar to *"sometimes it was not possible to obtain the shapes I had in mind,"* others noted that *"suggesting images helped me be more creative for Task 3."* This insight motivates further refinement of the Juxtapoze system targeted towards better supporting both undirected and focused tasks. For this, we will look at (1) conducting user studies related to focused tasks with a significantly larger database of clipart, and (2) incorporating a tagged database that supports semantic queries and filtering.

Our evaluation for Juxtapoze in this paper was by necessity a qualitative and exploratory one rather than controlled, quantitative, and comparative. Attempting to pin down artistic activities and its outcome is notoriously difficult, if not impossible, and any results from such an analysis would have questionable value. While our evaluation does not allow us to make statements on the superiority (or inferiority) of Juxtapoze over other creative workflows, we nevertheless feel that the results speak to the expressiveness of the system.

While our current 2D shape search engine in Juxtapoze has been tested with a shape database of up to around 3,000 images, it is clear that scalability will continue to be a concern as the database grows in size. At the same time, it is also clear that the larger the shape database is, the larger the space of sketch suggestions available to the user becomes, thereby increasing the potential for creative diversity and expression. For example, a long-term, possibly unrealistic, vision for the would be to be able to use the entire Google Image Search database in the system. However, managing truly big data of this scale is left for future work.

While overall successful, we found several ways to improve creative expression and serendipity for Juxtapoze in the future; for example, by supporting other forms of querying such as color, visual style, and keywords, by providing advanced tools for shape editing such as shape deformation, and by providing support for multiview searching.

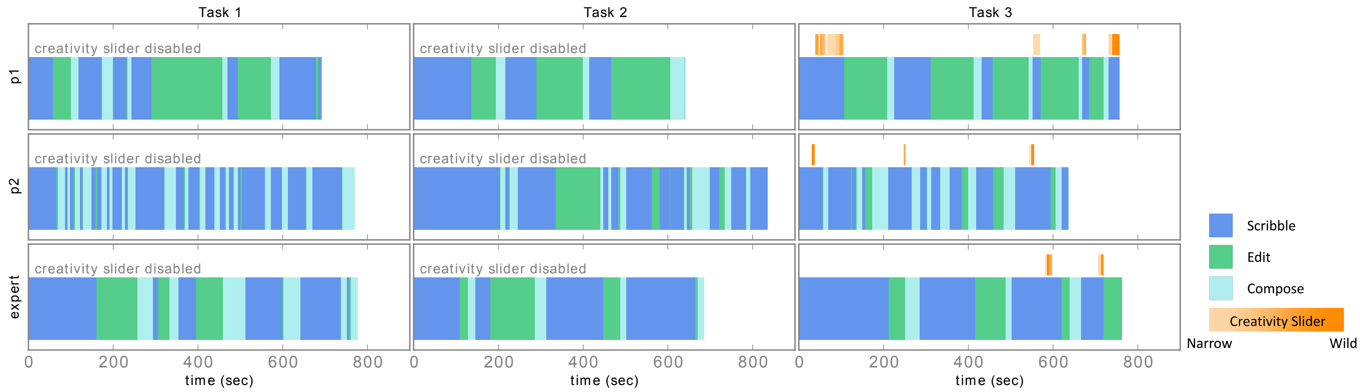

Figure 8. Timeline view of user activity, indicating transition between scribble (dark blue), edit (green), and compose (cyan) activities. Usage of the creativity slider is indicated by orange striations over corresponding scribble activities. Slider position is indicated by the lightness of the orange, with a lighter orange indicating the slider set to 'narrow' and darker orange indicating it set to 'wild'. Participants P1 and P2 were selected from the pool of 21 participants based on their diverging responses to Juxtapoze. The expert user is an industrial designer.

Figure 9. Sample compositions created by participants using Juxtapoze, arranged by task.

CONCLUSION

We have presented Juxtapoze, a creative workflow for stimulating serendipitous discoveries of shapes matching user sketches. Although serendipitous discoveries is an important part of learning, ideation, and creativity, most existing systems aim towards photometrically and geometrically correct content. This means that creative diversity and expression, which are key ingredients of artistic production, are neglected. The ease with which user can scribble, edit, and compose using Juxtapoze puts back the control and creative freedom in the hands of the users. Two qualitative user studies informed our work and showcased the utility of our ideas by letting both novices and expert designers create digital imagery using our workflow implementation. The participants found that Juxtapoze is enjoyable to use and helped them come up with creative, high-quality, and imaginative results.

User-generated content is on the rise like never before, and this movement is particularly strong for visual content. The HCI and creativity academic communities must respond by creating tools and techniques that are geared towards unleashing the raw creativity of users. We think that tools such as Juxtapoze are uniquely positioned to meet these future challenges, but much more work is needed in this domain.

ACKNOWLEDGEMENTS

This work is partly supported by the U.S. National Science Foundation on grants IIS-1227639, IIS-1249229, and IIS-1253863, as well as the Donald W. Feddersen Chaired Professorship and the Purdue School of Mechanical Engineering. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

REFERENCES

1. Ades, D. Photomontage. *Society 14*, 1 (1976), 13–18.
2. Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D., and Cohen, M. Interactive digital photomontage. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 294–302.
3. Bawden, D. Information systems and the stimulation of creativity. *Information Science 12*, 5 (1986), 203–216.
4. Boyd, D., Golder, S., and Lotan, G. Tweet, tweet, retweet: Conversational aspects of retweeting on twitter. In *Proceedings of the Hawaii International Conference on System Sciences* (2010), 1–10.
5. Bruns, A. Produsage. In *Proceedings of the ACM Conference on Creativity and Cognition* (2007), 99–106.

6. Buxton, B. *Sketching User Experiences*. Morgan Kaufmann, 2007.

7. Chen, T., Cheng, M.-M., Tan, P., Shamir, A., and Hu, S.-M. Sketch2Photo: internet image montage. *ACM Transactions on Graphics 28*, 5 (2009).

8. Cobbledick, S. The information-seeking behavior of artists: Exploratory interviews. *The Library Quarterly* (1996), 343–372.

9. Eitz, M., Richter, R., Boubekeur, T., Hildebrand, K., and Alexa, M. Sketch-based shape retrieval. *ACM Transactions on Graphics 31*, 4 (2012).

10. Eitz, M., Richter, R., Hildebrand, K., Boubekeur, T., and Alexa, M. Photosketcher: Interactive sketch-based image synthesis. *IEEE Computer Graphics and Applications 31*, 6 (2011), 56–66.

11. Erdelez, S. Information encountering: It's more than just bumping into information. *Bulletin of the American Society for Information Science 25*, 3 (1999), 25–29.

12. Foster, A. E., and Ford, N. Serendipity and information seeking: an empirical study. *Journal of Documentation 59*, 3 (2003), 321–340.

13. Gal, R., Sorkine, O., Popa, T., Sheffer, A., and Cohen-Or, D. 3D collage: expressive non-realistic modeling. In *Proc. Symposium on Non-Photorealistic Animation and Rendering* (2007), 7–14.

14. Gavilan, D., Saito, S., and Nakajima, M. Query-by-sketch based image synthesis. *IEICE Trans. Inform. and Syst. E91-D*, 9 (2008), 2341–2352.

15. Hawkins, E. *The complete guide to remixing*. Berklee Press, 2004.

16. Hemphill, C. S., and Suk, J. Remix and cultural production. *Stanford Law Review 61* (2008), 1227.

17. Huang, H., Zhang, L., and Zhang, H.-C. Arcimboldo-like collage using internet images. *ACM Transactions on Graphics 30*, 6 (2011).

18. Joshi, N., Mehta, S., Drucker, S. M., Stollnitz, E. J., Hoppe, H., Uyttendaele, M., and Cohen, M. F. Cliplets: juxtaposing still and dynamic imagery. In *Proceedings of the ACM Symposium on User Interface Software and Technology* (2012), 251–260.

19. Kazi, R. H., Igarashi, T., Zhao, S., and Davis, R. C. Vignette: interactive texture design and manipulation with freeform gestures for pen-and-ink illustration. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (2012), 1727–1736.

20. Lalonde, J.-F., Hoiem, D., Efros, A. A., Rother, C., Winn, J., and Criminisi, A. Photo clip art. *ACM Transactions on Graphics 26*, 3 (2007), 3.

21. Lee, Y. J., Zitnick, C. L., and Cohen, M. F. ShadowDraw: real-time user guidance for freehand drawing. *ACM Transactions on Graphics 30*, 4 (2011).

22. Liestman, D. Chance in the midst of design: approaches to library research serendipity. *Renaissance Quarterly 31*, 4 (1992), 524–532.

23. Lowe, D. G. Object recognition from local scale-invariant features. In *Proceedings of the IEEE international Conference on Computer Vision*, vol. 2 (1999), 1150–1157.

24. Marchionini, G. Exploratory search: from finding to understanding. *Communications of the ACM 49*, 4 (2006), 41–46.

25. Muja, M., and Lowe, D. G. Fast approximate nearest neighbors with automatic algorithm configuration. In *International Conference on Computer Vision Theory and Application* (2009), 331–340.

26. Nehring, K., and Puppe, C. A theory of diversity. *Econometrica 70*, 3 (2002), 1155–1198.

27. Nemhauser, G. L., Wolsey, L. A., and Fisher, M. L. An analysis of approximations for maximizing submodular set functions-i. *Mathematical Programming 14*, 1 (1978), 265–294.

28. O'Connor, B. Fostering creativity: enhancing the browsing environment. *International Journal of Information Management 8*, 3 (1988), 203–210.

29. Smith, A. R. Digital paint systems: An anecdotal and historical overview. *IEEE Annals of the History of Computing 23*, 2 (Apr. 2001), 4–30.

30. Sutherland, I. Sketchpad: a man–machine graphical communication system. In *Proceedings of the Spring Joint Computer Conference* (1963), 329–346.

31. Thudt, A., Hinrichs, U., and Carpendale, S. The Bohemian Bookshelf: supporting serendipitous book discoveries through information visualization. In *Proceedings of the ACM Conference on Human Factors in Computing Systems* (2012), 1461–1470.

32. Ullman, D. G., Wood, S., and Craig, D. The importance of drawing in the mechanical design process. *Computers & Graphics 14*, 2 (1990), 263–274.

33. Wang, C., Zhang, J., Yang, B., and Zhang, L. Sketch2Cartoon: composing cartoon images by sketching. In *Proceedings of the ACM Conference on Multimedia* (2011), 789–790.

34. Xie, X., Xu, K., Mitra, N. J., Cohen-Or, D., Su, Q., Gong, W., and Chen, B. Sketch-to-design: Context-based part assembly. *Computer Graphics Forum 32*, 8 (2013), 233–245.

35. Zitnick, C. Binary coherent edge descriptors. In *Computer Vision*, K. Daniilidis, P. Maragos, and N. Paragios, Eds., vol. 6312 of *Lecture Notes in Computer Science*. Springer, 2010, 170–182.