

Smart Sketchpad: Using Machine Learning to Provide Contextually Relevant Examples to Artists

Michael Fischer^(✉) and Monica Lam

Computer Science Department Stanford, Stanford University,
Stanford, CA 94305-9035, USA
{mfischer, lam}@cs.stanford.edu

Abstract. Sketching is a way for artists to generate ideas quickly, explore alternatives with less risk, and encourage discussions. How might computational tools amplify the abilities of artists? This paper introduces Smart Sketchpad, a digital sketchpad that uses machine learning to identify what is being sketched. The sketchpad then shows example images, color palettes, and subject information. The goal of Smart Sketchpad is to increase an artist’s ability to get ideas down with higher fidelity by making it easier to reference and include existing example works. Our study compares traditional sketching on a phone to Smart Sketchpad. We found that introducing examples during the sketching process leads to higher satisfaction of the sketch by the artist and an external expert.

Keywords: Sketching · Machine learning · Content-aware software

1 Introduction

There is a marketing legend about an instant cake mix that only required the addition of water but sold poorly because home cooks felt guilty about not having to do much work. The marketers found there was an “emotional imperative to bake cakes from scratch”. To fix the problem they released “ads at persuading bakers to view the plain cake layers as merely ‘step number one’”. “You can put your effort into glorifying your cake with frosting, dreaming up an exciting trim that puts your own label on it.” [1] The art, they argued, began after the cake came out of the oven. The cake layers were elements of a larger goal, an overall cake creation with frosting, filling, and decorations that allowed for a higher level of creativity from the baker. Similarly, in this paper, we examine how a computer can assist with “step number one” of the sketching. The goal is to allow the user more time to add the creative details that make a sketch their own.

Typically, to create interesting and good-looking work, an artist must think about content, color, and composition, given the constraint of how much time or effort that they can spend working on the project [2]. We seek to address this problem by providing a sketch interface where the artist can make use of example content and color palettes that are targeted at what the artist is sketching. The system works as a search system. The input to the search system is a sketch that is labeled by a machine learning algorithm based on prior examples. The results of the search are:

- (1) Line drawn images from an online image search.
- (2) Color pallets from a repository of user generated color pallets.
- (3) A reference to an encyclopedia article on the subject being drawn.

By referring to the related images and colors, the user can get inspiration from an example design, use parts of the example design, or directly add the example design by dragging and dropping it onto their composition.

In this paper we describe the user interface, how it was constructed, and the results of our study. We found that users and an external expert rated drawings produced with the help of the smart interface higher than those produced with the control sketch interface (Fig. 1).

2 Related Work

2.1 Examples

Our future ideas and creations are based on our prior experiences [3, 4]. Recognizing exemplary content is easier (and perhaps a precursor) to creating exemplary content. To this end, research has looked into how to provide examples within a computer interface [5, 6] to improve outcomes in creativity and productivity. There is still debate though that looking at examples may limit forms of creative thinking [7].

2.2 Sketching

Sketching is a commonly accepted practice of design and is how some designers convey their work [8, 9]. Algorithms have been developed that can identify simple sketches of objects about as accurately as humans [10].

2.3 Color

Color theory is closely linked to the physiology of our visual system and perception. Color is known to be matched to context [11, 12]. This research was then extended into the data visualization field by introducing an algorithm for automatic selection of semantically-resonant colors to represent data [13].

3 System

3.1 User Flow

The artist starts with a blank canvas. The control canvas has a button to change the brush color, an undo button, an erase all button, and a button to save the images to the server. After every stroke, the image is sent to the server.

When the server is more than 80 % confident it has found a match, the interface goes into smart mode. In smart mode, the artist has the option to tap a button at the top of the screen to display example images. The example images can be dragged and resized on



Fig. 1. Users were asked to draw a cat that could be put on a birthday card to a friend. On the left are sketches that were drawn using the control interface. On the right are sketches that were drawn with the help of examples images and colors that were suggested by the Smart Sketchpad.

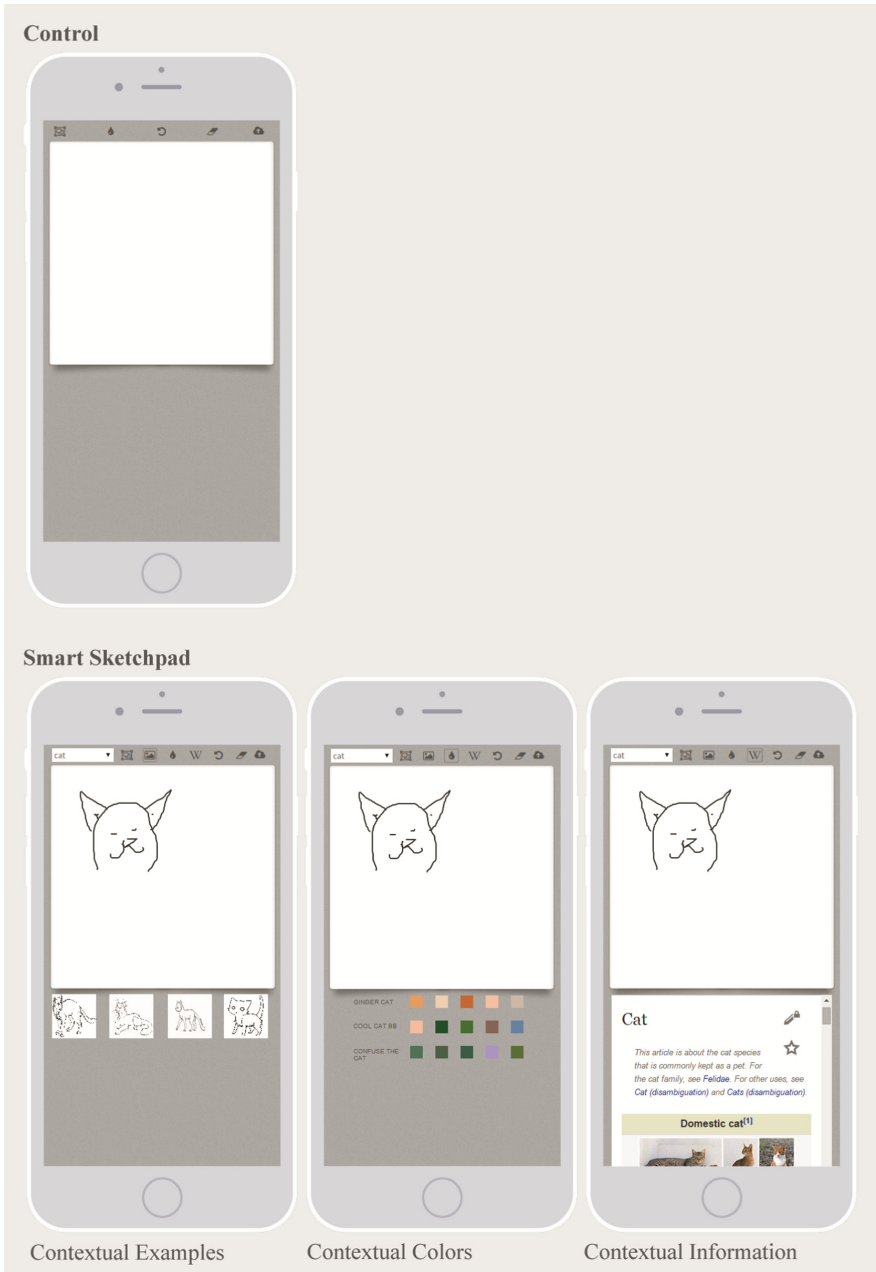


Fig. 2. The control sketchpad interface (top). The Smart Sketchpad (bottom). Smart Sketchpad provides contextual information about what is being sketched. Example images can be dragged onto the canvas (bottom-left). Color pallets related to the content being sketched are displayed (bottom-center). Contextual information about the subject is available from Wikipedia (bottom-right).

the drawing area or viewed in-place. There is a drop down box in the upper-left corner with other possible matches that were returned by the classification algorithm. The color palette button displays colors related to the item being sketched. Finally, there is a button to view the Wikipedia page for the item being sketched (Fig. 2).

3.2 Implementation

The system was built to run in a web browser. The drawing area on the client side was built using the HTML5 canvas object. We used a condensed nearest neighbor model on MetaMind [14] that was trained on hand drawn images of cats automatically downloaded from Google Image [15] search. The server was written in Node.JS [16]. Suggested images were show to the user by automatically searching Google Images using the tag returned from the classifier with the word “line drawing” appended. Color pallets were retrieved from the website COLOURlovers [17] by using their API to search for the top three most popular color pallets tagged with the name of the classifier (Fig. 3).




Fig. 3. The Smart Sketchpad system architecture. The user begins by sketching in the sketch input area. The sketch acts as a search query that the machine learning algorithm classifies. Based on the classification, example images, color pallets, and an encyclopedia entry are found. The information from the lookup is then displayed in the box below the sketch input area.

4 User Study

We conducted a user study to evaluate the effectiveness of the proposed system. Using the prompt:

You forgot to get a present for your friend’s birthday party. The party starts in 5 min. Your friend likes cats. Draw a cat for their present.

4.1 Participants

Six university students participated in the study (Table 1). Three self-identified as “experienced” sketchers and three self-identified as “novice” sketchers.

Table 1. The profile of the participants. Half were novice user and half were experienced users. An external evaluator always chose the sketches from the Smart Sketchpad as being better.

#	User’s Sketching Experience	User’s Rating of Image (Control)	User’s Rating of Image (Smart)	Expert’s Choice of Image
1	Experienced	1	4	Smart
2	Novice	2	4	Smart
3	Experienced	3	5	Smart
4	Experienced	3	3	Smart
5	Novice	4	5	Smart
6	Novice	4	6	Smart

4.2 Apparatus

A Motorola G, 3rd generation mobile phone was used for the evaluation. An optional capacitive stylus was provided and all participants chose to use it.

4.3 Procedure

At the beginning of the test we let participants test the interface and to get familiar with the system for approximately 2 mins.

Using the prompt above, each participant created two sketches on a mobile phone. In both cases they were encouraged to make the sketch as best as they could and to notify the examiner when they were done.

One sketch was with the smart system disabled. The other sketch was with the smart system enabled. The conditions were counter-balanced to control for ordering effects.

Using within-subject testing may have accounted for some learning effects. We were aware of this and made the decision that doing between-subject testing would make evaluating our experimental condition difficult because of the high degree of variability in subject’s artistic ability and style.

During the session we used a think aloud protocol to gather data about how the user thought about the design challenge. After completing the sketch the users answered a questionnaire about their experience. At the end, the user and examiner had an informal talk to find out more about questions not in the questionnaire or to examine a particular area in more depth.

We had an outside expert review the pairs of images produced. They were told to choose an image based on which they would be more excited to receive as a birthday gift from a friend. Within a pair, the control and smart sketchpad images were randomly ordered. The expert was not told about the context of the experiment.

5 Results

5.1 Quantitative Assessment

The result shows that in all but one case users were more satisfied with their result when using Smart Sketchpad. That user, who had advanced artistic skills, felt the same about both of their images.

In all cases, the expert chose sketches done using the Smart Sketchpad as the one they would prefer to receive.

Every user looked at the examples in the Smart Sketchpad condition. Four users used the example as a template.

Only two users picked the same image of the four that were offered. There was no clear definition of what a “best” cat template looked like between users.

All users said they sketch to capture an idea, two users said they also sketch for enjoyment.

All users completed the task in less than five minutes and none asked for more time. All users also said that they would like to use the system again. All users except for one experienced user were satisfied with the example images provided. Two experienced users asked for more color choices.

5.2 Qualitative Assessment

The results of our qualitative assessment show that users used the examples because they thought that they provided a better foundation than what they could do. Some sample comments:

The reference gave me something I could work from. I didn't have to figure out how to draw a cat by myself.

I liked that I could get a pre-drawn outline because it was way better than anything I could have done.

When I saw the pre-drawn cat outlines I decided to just use that because it was better than anything I could have drawn, and I got lazy.

Users then took the opportunity to add to the existing image their own additions.

I used the image itself and drew some lines on it for texture and color so that it would look nicer.

I just put additions onto the photo of a cat that I liked best.

Two of the more experienced sketchers used the examples as counter examples or to correct misconceptions that they had.

I drew something different from that in order to avoid drawing a stereotyped image.

The reference images were useful in correcting some misconceptions I had about anatomy.

Using examples fit into the existing workflow of some of the users but having it built into the interface made it easier to access.

Typically I have a pretty strong idea, but I'll look around Pinterest or flip through drawing books to develop it more.

One user wanted the examples to be less concrete and allow for more creativity.

It feels like cheating to use an image of a cat for drawing in the drawing. Potentially would be better if I could use it to outline then remove the image.

The color pallet was popular with users. Having color pallets that other people had validated that went together was found to be helpful. Although some felt restricted and wanted more color options.

The colors weren't as complementary as the ones in the smart interface.

The colors generally went together better because I didn't have to chose them.

I wanted a color that wasn't in the color list and I didn't know what to do.

I think the colors kind of restricted me - I originally envisioned a blue cat but I had to make a normal, brownish-orange cat because that's the colors that were given to me.

Only one user looked at the Wikipedia page. Although three others mentioned that they would have looked at the page if they were sketching something that they were not familiar with or wanted to find information about other things that they could add to the sketch.

6 Discussion

The feedback that was received about the system was positive and users were interested in continuing to use the smart interface.

An area that we found interesting was that three of the four users that used a template in the smart condition also included other elements (grass, a ball, clouds, or a flower) to embellish their drawing. This was done without any prompt, whereas the people that did not use the template did not embellish their work with anything else. A possible explanation of this behavior could be that users who used the example images as a template had more time left to focus on the other areas of their sketch.

Based on this observation, we will follow up this work with a test to determine how the smart interface is used when composing a larger, more complex work that involves multiple sub-elements. We are interested to investigate how the usage of the tool changes as the drawing goes from a small sketch to a larger composition. We'd be looking to see if the usage of the tool increases as the complexity of the final image increases.

Based on user feedback, another area that we would like to further explore is to give the user the option of producing their own search query for the system with either text or speech instead of having to use an image. The original idea was to be able to match the type of image being drawn; a cat in a certain pose generates an example of a cat in that pose. However, we found that users typically start drawing with a concept in their mind as opposed to a specific image of a cat in their mind, although we expect this to vary across use cases.

Bringing examples into the sketching interface supported the natural workflow that some users already did on a desktop computer. One remarked how this system would be useful when they only had their phone available to them.

We can also improve the training set of the images that can be recognized with the machine learning algorithm by adding all images sketched using the system back into the training set. Perhaps the system could then be used as a semantic authoring system for sketches.

Overall, we found that contextual suggestions improve the output of users as rated by both themselves and by an outside expert.

References

1. Shapiro, L.: *Something from the Oven: Reinventing Dinner in 1950's America*. Viking, New York (2004)
2. Sketching: the Visual Thinking Power Tool. <http://alistapart.com/article/sketching-the-visual-thinking-power-tool>
3. Smith, S., Ward, T., Schumacher, J.: Constraining effects of examples in a creative generation task. *Mem. Cogn.* **21**(6), 837–845 (1993)
4. Gentner, D., Holyoak, K.J., Kokinov, B.N.: *The Analogical Mind: Perspectives from Cognitive Science*. MIT Press, Cambridge (2001)
5. Kumar, R., Talton, J., Ahmad, S., Klemmer, S.: Bricolage: example-based retargeting for web design. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2011)*, pp. 2197–2206. ACM, New York (2011)
6. Lee, B., Srivastava, S., Kumar, R., Brafman, R., Klemmer, S.: Designing with interactive example galleries. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI 2010)*, pp. 2257–2266. ACM, New York (2010)
7. Schank, R., Childers, P.: *The Creative Attitude: Learning to Ask and Answer the Right Questions*. Macmillan, New York (1988)
8. Design thinking. (2015, October 13). In Wikipedia, The Free Encyclopedia. Retrieved 17:54, October 14, 2015, from https://en.wikipedia.org/w/index.php?title=Design_thinking
9. Newman, M. W., Landay, J.A.: Sitemaps, storyboards, and specifications: a sketch of Web site design practice. In: Boyarski, D., Kellogg, W.A. (eds.) *Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS 2000)*, pp. 263–274. ACM, New York (2000)
10. Eitz, M., Hays, J., Alexa, M.: How do humans sketch objects? *ACM Trans. Graph.* **31**(4), Article 44, 10 p., July 2012
11. van de Weijer, J., Schmid, C., Verbeek, J., Larlus, D.: Learning color names for real-world applications. *Trans. Img. Proc.* **18**(7), 1512–1523 (2009)
12. Schauerte, B.; Stiefelhagen, R.: Learning robust color name models from web images. In: *2012 21st International Conference on Pattern Recognition (ICPR)*, pp. 3598–3601, November 11–15, 2012
13. Lin, S., Fortuna, J., Kulkarni, C., Stone, M., Heer, J.: Selecting semantically-resonant colors for data visualization. In: *Proceedings of the 15th Eurographics Conference on Visualization (EuroVis 2013)*, pp. 401–410. The Eurographs Association & John Wiley & Sons, Ltd., Chichester (2013)
14. MetaMind. <https://www.metamind.io>
15. Google Images. <https://www.google.com/imghp>
16. Node.JS.. <http://nodejs.org/>
17. COLOURlovers. <http://www.colourlovers.com/>