VisiBlends: A Flexible Workflow for Visual Blends

LYDIA B. CHILTON and SAVVAS PETRIDIS, Columbia University
MANEESH AGRAWALA, Stanford University

1. INTRODUCTION

Visual blends are an advanced graphic design technique to draw attention to a message [Barry 2016]. They combine two objects in a way that is novel and useful in conveying a message symbolically. This paper presents VisiBlends, a flexible workflow for creating visual blends that follows the iterative design process [Norman 2002]. We introduce a design pattern for blending symbols based on principles of human visual object recognition [Sternberg 2011]. The workflow decomposes the process into both computational techniques and human microtasks. It allows users to collaboratively generate visual blends with steps involving brainstorming, synthesis, and iteration. An evaluation of the workflow shows that decentralized groups can generate blends in independent microtasks, co-located groups can collaboratively make visual blends for their own messages, and VisiBlends improves novices' ability to make visual blends by a factor of ten.

1.1 Approach: Design patterns to scaffold collaborative creation

Visual blends pose a challenging design problem because the designer must find two symbolic objects that satisfy conflicting goals - the two objects must be blended into one object, yet still be individually identifiable. Although there is no simple formula for solving design problems, many can be scaffolded around design patterns - abstract structure that specifies typical relationships between elements of the solution. [Alexander et al. 1977; Gamma et al. 1995; Duyne et al. 2002; Merrell et al. 2011; Agrawala and Stolte 2001; Agrawala et al. 2011; Sauppé and Mutlu 2014; Yu et al. 2014; Kim et al. 2015].

We introduce a design pattern for making visual blends: blends must combine two objects that have the same main shape (the Starbucks logo and sun are both mainly circles), but one object has identifying parts outside of the shape (the sun has rays outside the circle). See Figure 1. This shape-based constraint guides users in finding symbolic objects, and allows the system to automatically synthesize prototypes that combine contributions from multiple users.

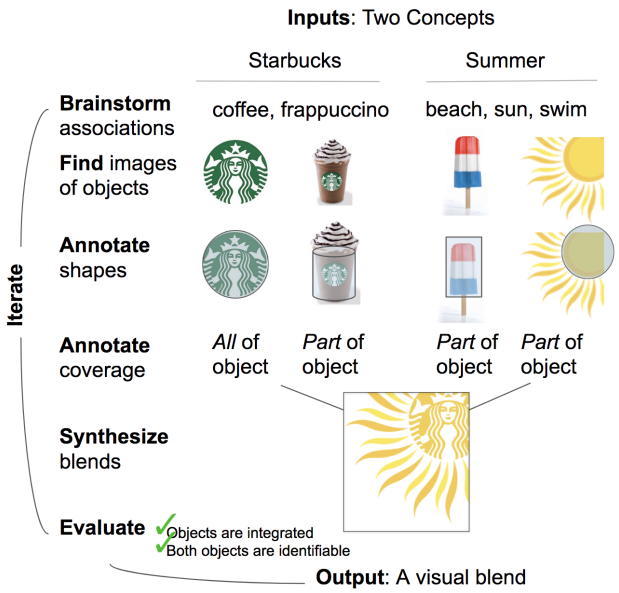

Fig. 1. An illustration of the VisiBlends workflow to create a visual blend for the concepts *Starbucks* and *summer*.

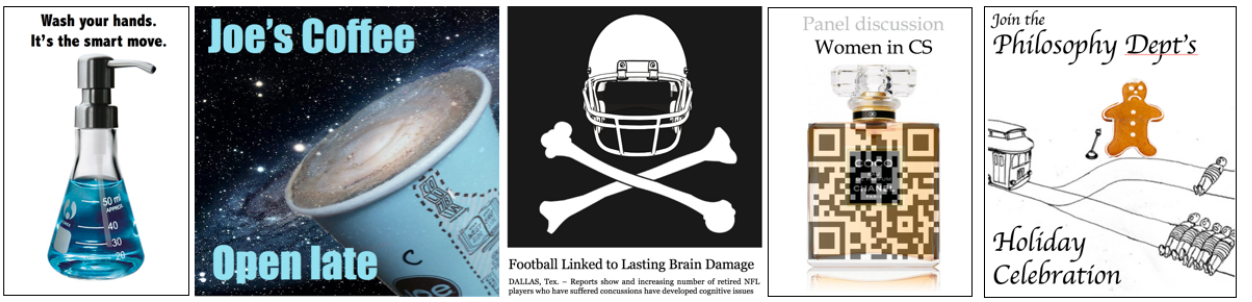

Fig. 2. Blends produced for the 5 messages in Study 2, with aesthetic improvements done by the users.

1.2 VisiBlends Workflow

VisiBlends is a flexible workflow that uses a hybrid of human intelligence and computational techniques to create visual blends. The workflow scaffolds the design process with separate interfaces for each type of microtask: brainstorming, annotation, and evaluation. The workflow is implemented as a website where each type of microtask is contained on its own page.

To enable collaboration, users' contributions are synchronized to each other in real time. To be flexible, users are allowed to move between the tasks and see others' work [Zhang et al. 2012]. Traditional workflows are more rigid [Little et al. 2010; Bernstein et al. 2010; Chilton et al. 2013]— they assign tasks to workers and do not allow users access all the data. By allowing access to data in a structured fashion, VisiBlends allows users to adapt to emerging constraints [Chilton et al. 2014].

The input to the workflow are two concepts such as *Starbucks* and *summer*, and the output of the workflow is one or more visual blends that synthesize two images together based on their shared shape (such as the Starbucks logo placed in the sun). Figure 1 contains an illustration of the workflow. In the first 3 steps, people brainstorm associations with each concept, find simple iconic images that represent it, and annotate the main shape within the image. Next, the matching and blending algorithm finds all pairs of objects that can be blended and synthesizes a prototype of the blend. People evaluate whether the prototype meets the requirements of the blend. If not, users can iterate to find more images that address the problems in the blends or improve the aesthetics of the blend. The system outputs a prototype, which is sometimes ready to use right away, but often a small amount of editing can increase the aesthetic quality of the blend.

1.3 Evaluation

Decomposing a creative task into independent microtasks is notoriously difficult [Retelny et al. 2017]. To validate the task decomposition in the VisiBlends workflow, we test it on three populations: decentralized workers, co-located groups, and individuals. We address the following research questions:

1. **Decentralized collaboration**: *Can users create blends in independent steps with no central planner or coordinator?* We ran a study where each person did one step of the workflow for randomly generated concept pairs like *autumn* + *bicycle* or *New York City* + *fashion*. In initial attempts, we only trained workers on the step they were assigned to. This failed to result in blends because workers were confused about how to select images and selected overly complex images that were hard to blend. However, when we trained workers on the entire workflow, they found simpler objects because they understood that simpler objects were easier to blend. They did much better at their individual tasks because they were able to better anticipate how their choices would affect later stages of the workflow.

Any decomposition of a problem will leave gaps between the steps - no matter how good the instructions are. This will lead to errors that accumulate later in the pipeline. If people are aware of the overall workflow, they can reason about the effects downstream and make better local decisions.

2. **Group collaboration for messages**: *Can co-located teams effectively collaborate to create blends for their own messages?* We selected 5 student groups with messages they wanted to convey: a public service announcement (PSA) "Washing your hands is smart", an ad: "Joe's Coffee is open late", a news article: "Football is dangerous", and two announcement "Panel discussion: Women in CS", and "Philosophy Department's Holiday Celebration". Figure 2 contains the outputs. Minor aesthetic edits to the blends were done by the users in PowerPoint.

The groups were co-located and used VisiBlends to scaffold their work and synchronise their images, annotations, and blends in real time. We found that the tool facilitated collaboration by automatically synthesizing surprising blends combining work from multiple users. For example one person's coffee cup image was automatically blended with another persons's galaxy image based on a shared ellipse shape. Groups noted that iteration was useful to improve every blend. The original images blended could always be improved to have a better color, style, or shape fit.

3. **Individual novices**: *Does VisiBlends help novices create more successful blends than their innate process?* We ran a study where 13 people with no graphic design experience made blends with and without VisiBlends. They were asked to make as many blends as they could in 5 minutes. They did this three time for three different concept pairs, like *New York City + night* and *bicycle + smart* (Figure 3).

Fig. 3. Examples of images made with and without VisiBlends for *New York City + night* and *bicycle + smart*. The images made without the system are not visual blends.

For participants who started without VisiBlends, they made on average 0.56 blends per concept pair without VisiBlends and they made 5.56 blends per concept pair with VisiBlends ($t(18)=4.88$, $p<0.001$). This is a 10-fold improvement in their ability to make blends. However, it could be biased by a warm-up effect. For participants who saw VisiBlends first, they made on average 5.67 blends per concept pair with VisiBlends, and an average of only 0.67 blends per concept pair without it ($t(21)=5.84$, $p<0.001$). Thus, using VisiBlends has a large and significant effect on performance in both conditions, demonstrating the utility of the system beyond just having knowledge of the process.

Even for individual users, the VisiBlends system is valuable because it scaffolds the iterative design process and helps them meet multiple constraints. Without the system, users focus on one constraint, like finding good symbols, but sacrifice the other constraints like shape fit by forcing the images together 3. Additionally, the system saves time by reusing images across designs and mocking up blends automatically. This helps users prototype more designs.

In conclusion, flexible workflows structured around a design pattern can help people collaboratively solve a creative problem. The design pattern is useful in automatically synthesizing individual contributions; the flexible workflow allows users to see all contributions and iterate towards better results.

REFERENCES

Maneesh Agrawala, Wilmot Li, and Floraine Berthouzoz. 2011. Design Principles for Visual Communication. *Commun. ACM* 54, 4 (April 2011), 60–69. DOI:http://dx.doi.org/10.1145/1924421.1924439

Maneesh Agrawala and Chris Stolte. 2001. Rendering Effective Route Maps: Improving Usability Through Generalization. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '01)*. ACM, New York, NY, USA, 241–249. DOI:http://dx.doi.org/10.1145/383259.383286

Christopher Alexander, Sara Ishikawa, Murray Silverstein, Max Jacobson, Ingrid Fiksdahl-King, and Shlomo Angel. 1977. *A Pattern Language - Towns, Buildings, Construction*. Oxford University Press.

Pete Barry. 2016. *The Advertising Concept Book: Think Now, Design Later (Third)*. Thames & Hudson, London, UK. 296 pages.

Michael S. Bernstein, Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. 2010. Soylent: A Word Processor with a Crowd Inside. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 313–322. DOI:http://dx.doi.org/10.1145/1866029.1866078

Lydia B. Chilton, Juho Kim, Paul André, Felicia Cordeiro, James A. Landay, Daniel S. Weld, Steven P. Dow, Robert C. Miller, and Haoqi Zhang. 2014. Frenzy: Collaborative Data Organization for Creating Conference Sessions. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1255–1264. DOI:http://dx.doi.org/10.1145/2556288.2557375

Lydia B. Chilton, Greg Little, Darren Edge, Daniel S. Weld, and James A. Landay. 2013. Cascade: Crowdsourcing Taxonomy Creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 1999–2008. DOI:http://dx.doi.org/10.1145/2470654.2466265

Douglas K. Van Duyne, James Landay, and Jason I. Hong. 2002. *The Design of Sites: Patterns, Principles, and Processes for Crafting a Customer-Centered Web Experience*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.

Erich Gamma, Richard Helm, Ralph E. Johnson, and John Vlissides. 1995. *Design patterns: elements of reusable object-oriented software*. Vol. 206. 395 pages. DOI:http://dx.doi.org/10.1093/carcin/bgs084

Joy Kim, Mira Dontcheva, Wilmot Li, Michael S. Bernstein, and Daniela Steinsapir. 2015. Motif: Supporting Novice Creativity Through Expert Patterns. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1211–1220. DOI:http://dx.doi.org/10.1145/2702123.2702507

Greg Little, Lydia B. Chilton, Max Goldman, and Robert C. Miller. 2010. TurKit: Human Computation Algorithms on Mechanical Turk. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology (UIST '10)*. ACM, New York, NY, USA, 57–66. DOI:http://dx.doi.org/10.1145/1866029.1866040

Paul Merrell, Eric Schkufza, Zeyang Li, Maneesh Agrawala, and Vladlen Koltun. 2011. Interactive Furniture Layout Using Interior Design Guidelines. In *ACM SIGGRAPH 2011 Papers (SIGGRAPH '11)*. ACM, New York, NY, USA, Article 87, 10 pages. DOI:http://dx.doi.org/10.1145/1964921.1964982

Donald A. Norman. 2002. *The Design of Everyday Things*. Basic Books, Inc., New York, NY, USA.

Daniela Retelny, Michael S. Bernstein, and Melissa A. Valentine. 2017. No Workflow Can Ever Be Enough: How Crowdsourcing Workflows Constrain Complex Work. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW, Article 89 (Dec. 2017), 23 pages. DOI:http://dx.doi.org/10.1145/3134724

Allison Sauppé and Bilge Mutlu. 2014. Design Patterns for Exploring and Prototyping Human-robot Interactions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1439–1448. DOI:http://dx.doi.org/10.1145/2556288.2557057

Robert J Sternberg. 2011. *Cognitive Psychology*.

Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2014. Searching for Analogical Ideas with Crowds. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems (CHI '14)*. ACM, New York, NY, USA, 1225–1234. DOI:http://dx.doi.org/10.1145/2556288.2557378

Haoqi Zhang, Edith Law, Rob Miller, Krzysztof Gajos, David Parkes, and Eric Horvitz. 2012. Human Computation Tasks with Global Constraints. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 217–226. DOI:http://dx.doi.org/10.1145/2207676.2207708