DesignScape: Design with Interactive Layout Suggestions

Peter O'Donovan[1,2]
[1]Dept. of Computer Science
University of Toronto
odonovan@dgp.toronto.edu

Aseem Agarwala[2]
[2]Adobe Research
Seattle, WA, USA
asagarwa@adobe.com

Aaron Hertzmann[1,3]
[3]Adobe Research
San Francisco, CA, USA
hertzman@adobe.com

ABSTRACT
Creating graphic designs can be challenging for novice users. This paper presents *DesignScape*, a system which aids the design process by making interactive layout suggestions, i.e., changes in the position, scale, and alignment of elements. The system uses two distinct but complementary types of suggestions: *refinement suggestions*, which improve the current layout, and *brainstorming suggestions*, which change the style. We investigate two interfaces for interacting with suggestions. First, we develop a suggestive interface, where suggestions are previewed and can be accepted. Second, we develop an adaptive interface where elements move automatically to improve the layout. We compare both interfaces with a baseline without suggestions, and show that for novice designers, both interfaces produce significantly better layouts, as evaluated by other novices.

Author Keywords
Graphic design, suggestion interfaces, adaptive design

ACM Classification Keywords
D.2.2 Software Engineering: Design Tools and Techniques User Interfaces

INTRODUCTION
Graphic design is ubiquitous in modern life. Unfortunately, creating designs can be difficult, particularly for novices, who often wish to create simple posters, cards, or social media designs. Starting from a blank canvas can be overwhelming, and exploring alternatives is time-consuming. Novice designers also make a variety of mistakes, from misalignment to incorrect emphasis of elements. Existing tools range from simple template-based interfaces like PowerPoint, to complex systems like Illustrator. However, these tools provide no suggestions when modifying templates or designs.

This paper presents a novel system for graphic design using layout suggestions, i.e., changes in the size, position, and alignment of elements. Our system proposes two complementary types of suggestions: *refinements* which improve the current layout, and *brainstorming* suggestions which explore alternative layouts with large changes in style (see Fig. 1). Exploration and refinement are critical and complementary tasks in design. However, exploration is difficult since a designer must imagine possible layouts, and modify many elements. Refinement is also difficult, since a single modification can necessitate many other changes. Our system includes both types, allowing users to easily switch between exploring alternative layouts and refining the current layout.

CHI 2015, April 18 - 23, 2015, Seoul, Republic of Korea

http://dx.doi.org/10.1145/2702123.2702149

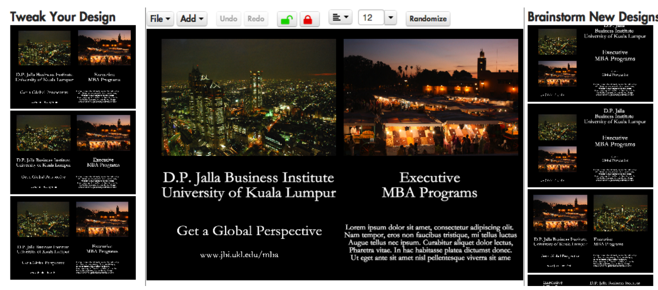

Figure 1. DesignScape Interface. The central canvas allows the user to create layouts in a simple editor. On the left, the system provides *refinement suggestions*, layouts similar to the canvas, but slightly improved. On the right, the system provides *brainstorming suggestions* large-scale layout changes in a variety of styles. Photos courtesy of Wilhelm Joys Andersen and Martin Fisch.

We use an energy-based model to generate designs that encode design principles such as symmetry, alignment, and overlap. User constraints are used to infer the designer's intent, and to make refinement suggestions on the current layout. We also learn a "style space" from examples, which can be used to generate new layouts in a variety of styles, providing starting points for design. The system can also retarget layouts, allowing the user to easily modify the design size.

We also investigate different ways users can interact with suggestions. First, we develop a suggestive interface, where suggestions are previewed and accepted. Second, we develop an adaptive interface which moves elements automatically. The two modes are compared to a baseline without suggestions by novice users on Mechanical Turk, and the quality of the resulting layouts are also evaluated. Both modes produce significantly better designs than the baseline on average. Lastly, we demonstrate the system's use for tablet-based design.

RELATED WORK
Exploring alternatives is a vital part of the design process. Gross and Do [2] present a prototyping interface which allows users to sketch drawings and store alternatives. Terry et al. [7] present an interaction technique which allows users to save and embed alternatives during the design process, and easily manipulate alternatives at a later point. Dow et al. [1] find that forcing users to create multiple design alternatives, instead of refining a single design, leads to improved results. Lee

Figure 2. Refinement Suggestion. Left: the user sets the current layout by positioning and scaling the elements. Note the user has locked the top-right graphic, and left-aligned the selected text block. Right: the system suggests left-aligning the three elements, and increasing the bottom graphic.

et al. [3] present a web-design interface for browsing related designs with simple attributes like background color; viewing these alternatives produced higher-quality designs. Merrel et al. [4] demonstrate interactive suggestions for furniture layouts. However, no previous work provides interactive suggestions for single-page designs like posters/advertisements.

There is little work automating single-page graphic design layouts. Most relevant is O'Donovan et al. [5], who present an energy-based design model. Unfortunately this approach takes 40 minutes to synthesize a single layout. We expand this model, increasing its efficiency substantially and introducing user constraints and style suggestions, enabling novel *interactive* interfaces.

DESIGNSCAPE OVERVIEW

We next provide a high-level view of the design interface (Fig. 1). See the Supplemental Video for demonstrations. The key goal of our interface is to provide automatic assistance during the design process, both for large-scale layout changes, as well as small improvements in position, scale, alignment, and line breaks.

Suggestive Interface. In this mode, the system shows three refinement suggestions, which vary in their similarity to the current layout. The top layout is the most conservative, only making slight modifications, whereas the bottom layout will suggest larger changes. Suggestions are easy to view and accept; the user mouses-over the suggestion to see a full-size preview in the main canvas, then clicks to accept. These suggestions also adapt to inferred user constraints. For example, if the user center-aligns two multi-line text blocks in the canvas, the interface will suggest internally center-aligning the text blocks. Fig. 2 shows an example refinement suggestion.

Adaptive Interface. The system also works in a separate adaptive mode where elements are changed automatically. This interface provides more fluid interaction, as the user does not need to view and accept changes. However, the adaptive interface is potentially frustrating if the suggestions do not match the user's desired goals.

Brainstorming Suggestions. To help users explore alternatives, the system also shows example layouts in a variety of styles. Each layout is distinct from the others, and the styles vary according to symmetry, text and graphical size, alignment preferences, etc. We describe the sampling procedure later, and provide examples. Brainstorming suggestions appear in both suggestive and adaptive modes.

Retargeting. The system can also retarget layouts to different sizes, an increasingly common task for designers. If the user modifies the canvas size, the system will move the elements to match the previous relative locations and scales, while respecting design principles such as alignment or overlap.

User Interactions

The system implements a common set of design interactions. The user can move, scale, add, or remove elements, and change text line breaks, font, colour, and alignment. The interface provides smart guidelines, which appear when two elements are close to aligned, and will snap elements together. Along with undo and redo functionality, intermediate layouts can also be saved, and are shown to the right of the canvas.

A user views the proposed layouts to the left and right of the canvas. Mousing-over the proposal previews the suggestion in the main canvas, and the user clicks to accept. The user can lock elements, which fixes the position and scale of the elements in the suggestions. However, the suggestions use the canvas layout as a soft constraint, and will try to remain similar to the canvas layout. In Fig. 2, the user has locked the top-right graphic, so this element is fixed in the suggestion.

SYNTHESIZING LAYOUTS

To synthesize layouts, we adapt the approach of O'Donovan et al. [5] which uses an energy function to model design principles. We adapt this model for the GPU, and represent elements as bounding boxes, which permits efficient computation of distances, area, and overlap. That model also segmented designs, detected alignment groups, and measured the perceptual importance of elements. Unfortunately, those techniques are time-consuming so we use simpler heuristics. We use an element's size as its perceptual importance; we find alignment between elements, but do not detect alignment groups. We measure symmetry by the difference in distances to the nearest left-most and right-most element or border.

User Constraints

To help the user specify their intent, we introduce additional constraints in the model, both inferred from user behaviour and explicitly specified. A soft constraint prefers that elements remain close to their current position and scale. However, the user can specify a hard constraint which locks an

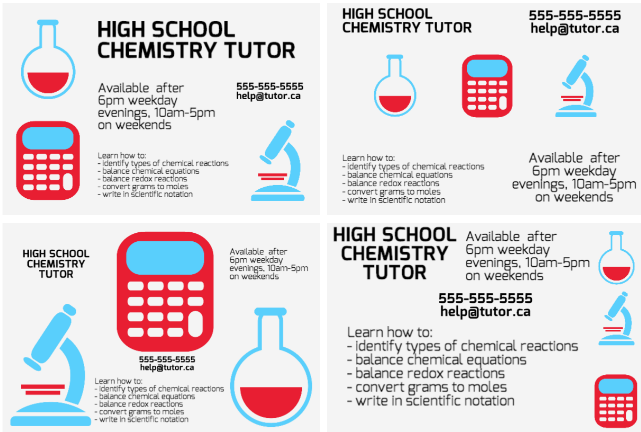

Figure 3. Brainstorming Suggestions. The system generates layouts in a variety of styles by sampling a low-dimensional parameter space learned from examples.

element's position and scale. The system also infers alignment constraints using the guide lines which automatically snap nearby elements. These lines then act as temporary constraints, with the elements in the suggested layouts matching the given alignment. The internal alignment of multi-line text blocks also changes to match these alignment lines. The user may also specify explicit size constraints, i.e., that specific elements must have the same size.

Optimization

We optimize layouts on the GPU (Nvidia GTX 770) based on the approach of Merrel et al. [4], using parallel tempering [6]. Convergence for the refinement suggestions is usually less than a second, since the canvas layout is used for initialization. Optimizing the brainstorming suggestions takes an average of 5 seconds. The system also saves brainstorming suggestions, and if pre-computed layouts are available, will display a random layout every 2 seconds. Such pre-computation is reasonable for testing; parallelization on multiple GPUs would eliminate the need for pre-computation.

Style Sampling

The brainstorming suggestions provide layouts in a variety of styles and arrangements. To generate different styles, we require different model parameters. As in O'Donovan et al. [5], we use Nonlinear Inverse Optimization (NIO) to estimate parameters given an example, but we still require a method to generalize to new styles. To accomplish this goal, we construct a low-dimensional subspace of the parameter space. We first use NIO to estimate the parameter vector for 18 example layouts, then use PCA to form a new basis. We then sample randomly in this new basis, producing layouts within the "style space" of the examples (see Fig. 3).

EVALUATION

To evaluate the adaptive and suggestive interfaces, we conducted several studies comparing them with a baseline interface without suggestions. First, novices created layouts in each interface. We then conducted two studies using AB testing: one comparing designs created using the adaptive interface to those with the baseline, and a second comparing suggestive and the baseline. Note that the baseline included snapping alignment lines and resizing constraints.

We then conducted two studies analyzing user preferences for each interface. In the first study, users used both the baseline and adaptive interfaces, then gave ratings and comments. In the second, users compared the baseline and suggestive.

Design Quality Studies

To evaluate design quality, we first had workers on Mechanical Turk (MTurk) create layouts with one of the three interfaces. Users created three layouts for a design, with no time restrictions, and could then provide comments. We used 2 designs, with 20 workers each. Workers only used a single interface (a between-subjects study), taking 2.41 minutes on average to create a layout in the baseline, versus 2.43 and 2.47 for suggestive and adaptive, respectively. Over 100 layouts were collected from each interface.

To compare layouts, we used AB testing on MTurk. For the adaptive interface study, we compared 60 adaptive layouts (randomly selected) with 4 baseline layouts (also randomly selected), and vice-versa for the baseline, producing 480 comparisons (2 interfaces $\times$ 60 layouts $\times$ 4 alternatives). Each worker evaluated a set of 20 pairs and chose their preferred layout in each pair; 10 workers completed each set. Duplicates were added and inconsistent users removed. For the suggestive interface study, we created a similar set of 480 pairs comparing suggestive and baseline designs.

Designs created by the new interfaces are generally preferred to those from the baseline. For the adaptive vs. baseline study, users voted for adaptive designs 63.4% of the time (2239 votes for adaptive designs vs. 1291 baseline votes). A χ^2 test comparing the counts to an expected frequency of 50% (1765 votes) produces $\chi^2(df = 1)$ = 254.59, with $p < 0.01$. In the suggestive vs. baseline study, users voted for suggestive designs 57.7% of the time (2193 votes for suggestive vs. 1605 votes for baseline), with a $\chi^2(df = 1)$ = 91.03, with $p < 0.01$. Therefore, both new interfaces help novices to create better layouts than the baseline.

We can also compare the adaptive and suggestive interfaces using a χ^2 test on the 2×2 contingency table of votes from both studies. This test shows $\chi^2(df = 1)$ = 24.51, with $p < 0.01$, indicating a significant difference between the new interfaces. This difference is likely because the adaptive interface fixes many basic layout problems for novices such as overlap or misalignment. The suggestive interface is more subtle, and the user can always ignore the suggestions and treat the interface as a standard layout tool. In fact, in our study, 34.17% of the suggestive interface layouts used no suggestions. Note that for AB testing, we only used layouts where the user accepted at least one suggestion. Figure 4 shows the worst and best layouts from each interface.

Preference Studies

We next examined user preferences for the different interfaces. In a within-participant study, MTurk workers used both

Figure 4. Worst and best rated layouts from each interface. Top: baseline without suggestions. Middle: suggestive interface. Bottom: adaptive interface. We report the mean fraction of votes from AB testing.

the adaptive and baseline interface (called the "Direct" interface in the study). Users performed a short tutorial, then created two layouts with that interface. The user then switched to the other interface, completed that tutorial, and created two other layouts. The order of interfaces was random. After using both interfaces, workers provided a 5-point Likert scale rating for each, a binary preference (either baseline or adaptive), and optional comments. 40 workers completed the test.

The mean of the Likert ratings for the baseline and adaptive interfaces were 3.75 ± 0.15 and 3.81 ± 0.15 respectively, with a median of 4 for both. A Wilcoxon signed-rank test indicated no statistical difference between paired ratings. 60% of people preferred the baseline to the adaptive interface.

To compare suggestive and baseline, we completed an identical study with those interfaces. The mean ratings for the baseline and suggestive interfaces were 3.73 ± 0.14 and 3.67 ± 0.14 respectively, with a median of 4 for both; A Wilcoxon test showed no difference between the ratings. 59% preferred the baseline to the suggestive. To explain these ratings, we next examine participant feedback.

Participant Feedback

Many users reported enjoying the new interfaces: *"I really like this interface and how it automatically lines things up." "Great to have a decent starting point," "The suggestion interface acted as a supplement to my own thought. I could find creativity and inspiration in it."* However, several users disliked the adaptive shifting of elements, e.g., *"I did not like when the picture would resize the words or pictures without my permission."* Many users preferred the baseline interface because of the greater control: *"I liked having a lot more control over the elements with the Direct Interface."* Furthermore, some users felt the suggestions constrained their creativity: *"I feel I am more creative if I use my ideas," "I felt as if the suggestions were taking the human element out a little," "I preferred the direct interface because I felt like the suggestions were so good, I did not feel I could make something significantly different. "*

DISCUSSION

In this paper, we present a novel system for graphic design using layout suggestions. Our key contributions are interfaces and interaction techniques allowing suggestions to aid the design process. Our adaptive interface allows users to create layouts using automatic improvements in a fluid manner. We also present a suggestive interface where users must actively accept changes. We evaluate these interfaces and find that users create better designs than a baseline interface without suggestions. However, many open problems remain, including improved exploration tools for style suggestions, more stylistic variety, and collaboration tools between users. In the Supplemental Video, we also demonstrate the system on a tablet. Suggestion-based tools are well-suited for touch interfaces where precise control of elements is difficult.

Current tools often make design tedious and time-consuming, particularly for novices. This paper presents a small step towards automating parts of the design process. However, we found that suggestions were not always desired. Some users preferred to have complete control over the design process, and found that automatic suggestions took away from their creativity. This intriguing result suggests more work investigating how suggestion-based tools are perceived by users. How can we develop tools to eliminate the tedious parts of designing, while still encouraging creative exploration?

REFERENCES

1. Dow, S. P., Glassco, A., Kass, J., Schwarz, M., Schwartz, D. L., and Klemmer, S. R. Parallel Prototyping Leads to Better Design Results, More Divergence, and Increased Self-efficacy. *ACM Trans. Comput.-Hum. Interact. 17*, 4 (2010), 18:1–18:24.
2. Gross, M. D., and Do, E. Y.-L. Ambiguous Intentions: A Paper-like Interface for Creative Design. In *Proc. UIST* (1996), 183–192.
3. Lee, B., Srivastava, S., Kumar, R., Brafman, R., and Klemmer, S. R. Designing with Interactive Example Galleries. In *Proc. CHI* (2010), 2257–2266.
4. Merrell, P., Schkufza, E., Li, Z., Agrawala, M., and Koltun, V. Interactive Furniture Layout Using Interior Design Guidelines. *ACM TOG* (2011).
5. O'Donovan, P., Agarwala, A., and Hertzmann, A. Learning Layouts for Single-Page Graphic Designs. *IEEE TVCG*, 20, 8 (2014), 1200–1213.
6. Swendsen, R. H., and Wang, J. S. Replica Monte Carlo Simulation of Spin-Glasses. *Physical Review Letters 57*, 21 (Nov. 1986), 2607–2609.
7. Terry, M., Mynatt, E. D., Nakakoji, K., and Yamamoto, Y. Variation in Element and Action. In *Proc. CHI* (2004), 711–718.