

Designing with Interactive Example Galleries

Brian Lee, Savil Srivastava, Ranjitha Kumar, Ronen Brafman, Scott R Klemmer

Stanford University HCI Group
Computer Science Department
Stanford, CA 94305-9035
srk@cs.stanford.edu

ABSTRACT

Designers often use examples for inspiration; examples offer contextualized instances of how form and content integrate. Can interactive example galleries bring this practice to everyday users doing design work, and does working with examples help the designs they create? This paper explores whether people can realize significant value from explicit mechanisms for designing by example modification. We present the results of three studies, finding that independent raters prefer designs created with the aid of examples, that users prefer adaptively selected examples to random ones, and that users make use of multiple examples when creating new designs. To enable these studies and demonstrate how software tools can facilitate designing with examples, we introduce interface techniques for browsing and borrowing from a corpus of examples, manifest in the Adaptive Ideas Web design tool. Adaptive Ideas leverages a faceted metadata interface for viewing and navigating example galleries.

ACM Classification: H.1.2. [Models and Principles]: User/Machine Systems—*Human factors*.

General terms: Design, Human Factors

Keywords: Design thinking, examples.

INTRODUCTION

Many users learn Web design by viewing and modifying the source code from other pages on the Web. For its ability to scaffold learning, the “view source” option in Web browsers is a pinnacle of interface design. Leveraging examples of previous work is an established technique in design [6]. Many design education programs teach students to think like experts by exposing and encouraging them to make use of existing examples [39]. Exposure to examples provides awareness of the design space of potential options, and adapting past solutions to fit the current context can help people creatively address new situations [23, 24]. Design compendiums such as *The Big Book of Logos* [7] serve as valuable resources for inspiration, and the advent of prolific, searchable Web content has provided ready access to a broad array of work created by other designers. When appropriate, example designs can offer pragmatic

value as well as inspirational value. Starting with an existing design and modifying it can provide a lower barrier to entry than starting with a blank slate. Amateurs, prototypers, and those trying to create a new design quickly find reusing examples especially valuable [5, 19, 34].

Current practices for working with design examples are largely informal and ad hoc [22, 35]—especially for non-professionals. This paper explores whether structured corpus navigation can help users find inspirational examples and facilitate design by example modification. It comprises three main sections. First, it presents a conceptual perspective on the role of examples in design. Second, it describes the Adaptive Ideas system and the approach it introduces for selecting and displaying examples. Third, it presents three experiments that explore the value of explicit mechanisms for design by example modification. These studies found that independent raters prefer designs created with the aid of examples, that examples benefit novices more than experienced designers, that users prefer adaptively selected examples to random ones, and that users make use of multiple examples when creating new designs. While the software tool and empirical results in this paper examine the specific context of Web page design, the intuitions this work employs—most notably, the importance of analogy in creative cognition [15, 46]—suggests these findings likely have broader import.

The Existing and Potential Role of Examples

While it sometimes seems like ideas arise out of thin air, creativity is necessarily the result of applying existing knowledge [4]. Our prior experiences provide the scaffold upon which we create new ideas [15, 36, 43], and copying someone else’s successful actions is more efficient than reinventing them from scratch. Ironically, given the centrality of experience to creativity and insight, people often neglect to make use of relevant knowledge, even when encouraged to do so through summarizing the relevant experience, stating the principle it embodies, or creating a diagram [17, 18].

Comparing multiple examples can help overcome people’s limited transfer from a single example. People are much more likely to make use of analogous experiences and infer the underlying principle when provided with multiple examples, or when presented with both a principle and a case, and asked to compare them [18]. As Gentner writes, “comparison processes can reveal common structure and combine partial structures and thus promote transfer, even early in learning when neither example is fully understood”

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.

Copyright 2010 ACM 978-1-60558-929-9/10/04...\$10.00.

[16]. The benefits of principle-case comparison may help explain the value of design patterns in domains like architecture [2], software engineering [13], and Web design [48]. Similarly, software developers frequently copy example software code from Web tutorials, and the tutorials' descriptions help developers locate and comprehend relevant examples [5].

While working with examples can provide important benefits for both learning and outcome, there is (at least in theory) the danger of over-relying on at-hand examples. Sometimes, people interpret the current situation largely in terms of recently cued information, even when it is to their detriment [38]. Smith *et al.* have demonstrated that people's ideas conform to presented examples in creative generation tasks [41]. On this view, conforming designers are like the mimiknik birds in an Al Capp cartoon, whose songs are a facsimile of whomever they hear [1].

Whether, when, and how much conformity induced by prior knowledge negatively impacts creativity in the “real world” is a matter of active debate. Weisberg, among others, champions the view that more experience is nearly always beneficial, and that out-of-the-box thinking is generally better explained by the *presence* of relevant valuable knowledge rather than the *absence* of inhibiting or misleading knowledge [49].

Marsh *et al.* studied how exposure to examples affects both conformity and novelty in a creative task. They found that subjects did conform toward examples. However, participants who saw many examples did not create less novel work because conforming elements “replaced” *mundane* aspects of designs rather than *novel* ones [31]. Importantly, neither Smith *et al.* nor Marsh *et al.* assessed any type of *quality*, a piece of the puzzle this paper seeks to help fill in.

How can one reconcile the potential benefits and dangers of analogical reasoning? Expertise minimizes the conformity bias: while novices can be overly enticed by proximally appealing options (or overly discouraged by proximal roadblocks), experts often see past them [8]. Additionally, examples from other domains are also valuable for insights and breakthroughs. For example, Johannes Kepler extensively used analogical reasoning in his astronomical discoveries [15]. When effective, “local analogies fill in an established framework; distant analogies create new frameworks” [14].

In short, this literature suggests that people are wise to seek out examples, that creative thought benefits from diverse experiences, and that using principles and examples together can be particularly valuable. However, there are important limitations and open issues in developing a theory of example usage.

First, most prior empirical work studied domains with “correct” answers and/or explicitly articulated principles. However, many design practitioners and theorists believe there are fewer explicit principles in design, or at least they are less crisp, and that judgment and intuition play a larger

role [10, 39]. Exemplifying this perspective is the mantra attributed to Ansel Adams that, “There are no rules of composition in photography, there are only good photographs.”

Second, the literature has mostly studied participants' ability to leverage a carefully selected example, rather than the ability to recognize a good example from a set of real-world examples. If tools harvest and structure naturally occurring examples, does that benefit users? Tohidi *et al.* [47] demonstrated that exposing users to multiple alternative designs aids their *discussions*. This paper explores the complementary question of whether examples strengthen users' *creations*.

THE ADAPTIVE IDEAS WEB DESIGN TOOL

To better understand the role that examples—and example-augmented tools—can play in design, we created and evaluated the Adaptive Ideas Web design tool. Adaptive Ideas augments a direct-manipulation Web page editor with an interface for parametrically browsing a corpus of example pages. Users can sample elements from these pages for use in their own work (Figure 1).

Adaptive Ideas is an extension to the Firefox browser's built-in graphical HTML editor. The prototype's example corpus comprises 250 manually harvested Web pages; most are student home pages found from university Web sites. To reflect the diversity of content on the Web, quality was not a selection criterion.

The Adaptive Ideas interface comprises an editing pane, an examples gallery, and a preview pane (see Figure 1). The user can begin with a complete example by selecting one from the gallery. The preview pane enables users to see an example in greater detail and select a feature of it to borrow. Additionally, the preview example can serve as a pivot



Figure 1 The Adaptive Ideas Web design tool extends a direct-manipulation editing pane (*top*), introducing the ability to browse an adaptively-generated examples gallery (*lower right*), and borrow elements from examples through a preview pane. (*lower left*).

for updating the gallery set; users can request to see examples that are similar to the preview example along a specified dimension. The dimensions used in the prototype were: background color, primary font, number of columns, and visual density.

Adaptive Ideas introduces an optimization-based approach to selecting, presenting, and browsing design material. The technical core of this approach is a *subset selection algorithm* that chooses examples from a corpus to display so as to maximize estimated design value.

To create an example-based design tool, two main technical building blocks are required: an interface for working with examples, and an algorithm for selecting which examples to display and how.

Interfaces for working with examples

Several existing design tools—such as Apple’s iWeb—enable users to create Web sites based on pre-defined templates. d.mix introduced the idea of graphical interaction with Web sites as a means for “sampling” code [20], and CopyStyler introduced an interface for transferring CSS styles between sites [11]. Design Galleries [30] introduced two insights relevant to this paper’s goals. First, a small-multiples presentation of alternatives enables users to rapidly view many options and gain intuitions about the design space. Second, to present users with a visually diverse set of designs, select examples using a perceptually-based distance metric. Similarly, Side Views [45] introduced a small multiples interface for parameter setting; enabling simulta-

neous comparison of alternatives and serendipitous discovery. Additionally, some tools provide explicit interface mechanisms for parameterizing content [21, 40, 44].

The Adaptive Ideas system introduced in this paper (see Figure 2) offers four contributions beyond prior work. First, existing systems offer either templates specifically created for modification or computationally synthesized examples. This paper introduces mechanisms for browsing and borrowing elements from examples created *in the wild*. Second, template-based tools only allow one template to be used, and the normative model is that the template is selected at the outset. There is significant value in sampling *different elements* from multiple designs created by others, rather than simply selecting one option from a fixed set. Third, while templates abstract form from use, examples demonstrate the *in situ* use of a design. This is valuable because the aesthetic and functional goals of design are interdependent [32]. Finally, this paper introduces navigation techniques for users to select a design and see examples that are similar to or different than it.

Structuring the Space of Examples

To decide what examples to display, Adaptive Ideas introduces a subset-selection algorithm that attempts to maximize estimated design value. It operationalizes design value through two proxy measures: the *content-value* of the example content for the user’s current task, and the *size-value* of presenting an example at a particular size.

Examples are organized using faceted metadata [33]. The example pages harvested from the Web are assigned the

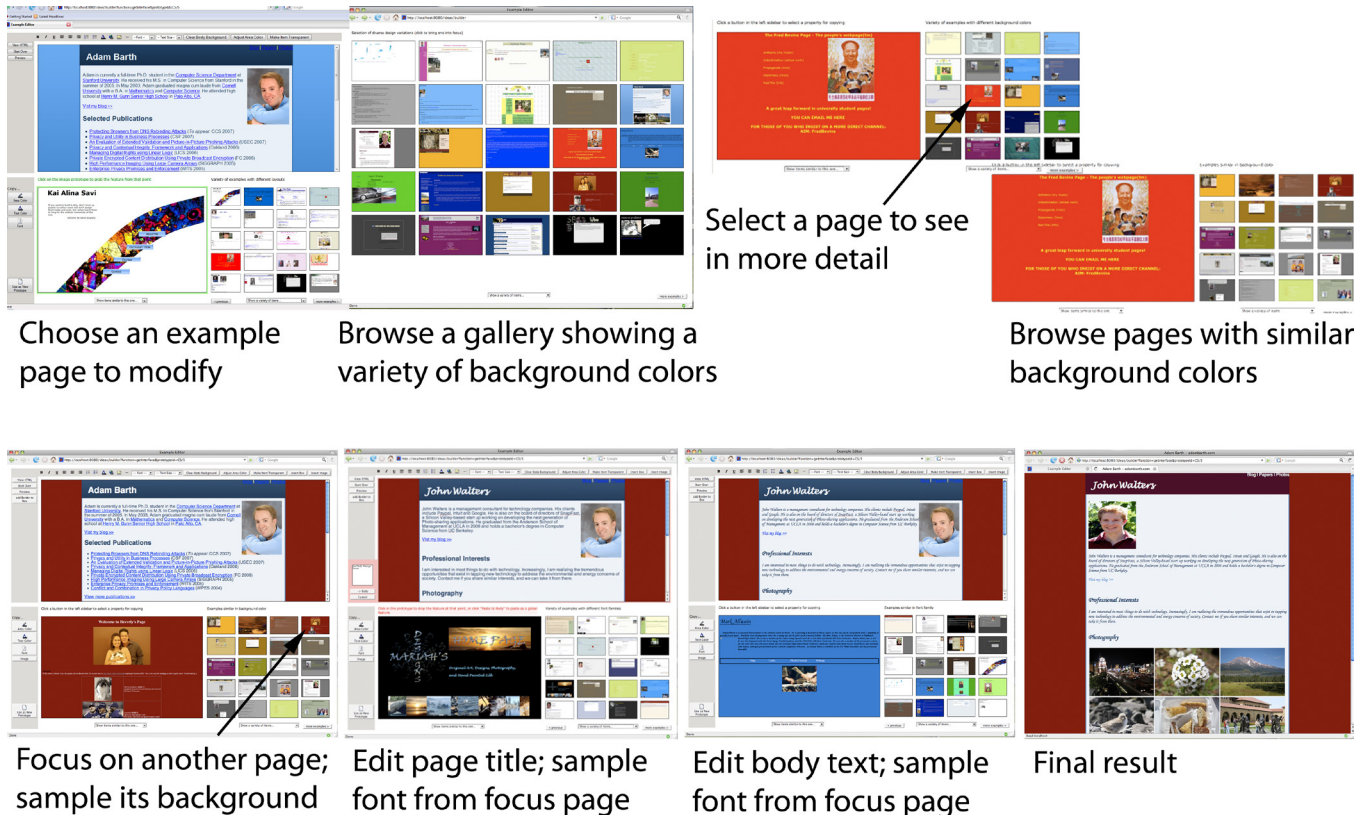


Figure 2 With Adaptive Ideas, users can modify and combine elements from multiple example pages to create a Web page.

following facet attributes: background color, primary font, number of columns, and visual density (see Appendix B). In the prototype, we manually assigned attributes to pages. In particular, visual density was subjectively assessed by one of the authors. We believe a production implementation could automatically assign density using *e.g.*, the clutter metrics proposed by Rosenholz *et al.* [37].

For each pair of pages in the corpus, Adaptive Ideas computes a pairwise distance metric for each attribute. For example, the current prototype defines the distance between two colors to be their Euclidean distance in the HSB color space. (One could use a more perceptually based space, such as LUV. Informal exploration suggests the practical difference would be negligible.) Adaptive Ideas uses a ternary function for fonts: 0 if the fonts are the same; 1 if they are both serif or both sans serif; 2 if the fonts have different serif style. (A production version might want to include font weights and styles in the distance metric.)

The gallery thumbnails are pre-rendered images. When the user selects a location in an example image to sample from, the Adaptive Ideas server looks up the appropriate HTML fragment by rendering the corresponding Web page on a server-side browser.

Deciding Which Examples Are Currently Most Valuable

Users can request a set of examples that are either *similar* to a specified focus element, or that represent a *variety* of options along a particular dimension. Our intuition is that showing similar examples will be useful when looking for subtle design variations, and that a variety of examples will be more valuable when looking for broad inspiration.

To create a subset of similar items, Adaptive Ideas calculates the distance of pages from the specified focus, sorts them in ascending order, and selects the first n items.

Creating a good variety is less straightforward. Should it convey as broad a range as possible (as in Design Galleries [30]), represent the distribution of the underlying dataset, or follow some other formula? For instance, the majority of Web sites in our dataset have a white background. A representative subset would contain mostly white Web pages. A broader range might be more interesting, but might inappropriately emphasize outliers or unusual points in the design space.

To balance these concerns, the Adaptive Ideas framework takes a *spaced stochastic* approach to selecting a representative variety: “spaced” meaning that examples need to be at least a certain distance apart, and “stochastic” meaning that elements are selected randomly from the set that meets the distance constraint. It starts by randomly selecting a seed example. Iteratively, random examples are selected from the remaining elements in the dataset which are at least ϵ distance away from all of the elements selected thus far, where ϵ is the normalized result of a spacing function defined on a per-attribute basis. Similar selection methods

are also used in other domains, such as computer graphics [9] and chemical compound clustering [42].

The choice of ϵ significantly influences the behavior of the spaced stochastic algorithm. When ϵ is zero or small relative to the design space, this algorithm degenerates to the completely random case. As ϵ gets larger relative to the space, the algorithm has fewer elements from which to choose, and thus risks not filling up the space. The Adaptive Ideas prototype uses a large ϵ , such that the theoretical maximum number of elements chosen is close to n .

The algorithm continues picking elements until either n elements have been selected or no legal elements remain, *i.e.*, every unselected element is less than ϵ distance away from an element in the selected subset. If more elements are needed, the system selects elements at random from the full set of remaining elements until n have been chosen. On balance, ϵ guarantees that distinctly different values for the given attribute will be represented in the variety set, while filling out remaining elements randomly implies that some of the underlying distribution of values will be reflected. A variant would be to iterate over successively smaller values of ϵ until enough legal elements are found; this would further emphasize the breadth of the design space.

Laying out the Interface

How many example elements should be shown on the screen? Answering this question trades off showing a smaller number of items at larger sizes and showing a larger number of items at smaller sizes. To encapsulate this trade-off and efficiently evaluate candidate interfaces, Adaptive Ideas uses constraint-based optimization [12, 50]. We briefly summarize the approach here; for a thorough description, see [26].

How much “value” does the presence of a particular element contribute? The presentation value of elements generally increases with size, however this relationship need not be linear. For this paper, element value is defined as a monotonically increasing function that starts at zero, scales up quickly at small sizes, and approaches an upper limit at larger sizes. (Past a certain point, further increasing size yields diminishing returns.)

The estimated value of an element is defined as the product of its presentation value at the given size and its relevance to the current focus element. Thus, a low *presentation* or *relevance* score yields a low value, even when the other is high: a highly relevant item is of little value if it is unrecognizable; a large but irrelevant item also has low value.

Adaptive Ideas defines the estimated value of a candidate layout to be the sum of the values of its constituent elements. This presumes that the contributions of a given element are *independent* of the presence or absence of other elements. There are, of course, cases where the total value may be supermodular (the whole is greater than the sum of the parts) or submodular (diminishing returns).

Selecting an optimal size and layout for a set of elements can be viewed as a two-dimensional variant of the knapsack problem [29], where the constraining resource is display real estate. As knapsack exhibits the optimal-substructure property, we use dynamic programming. Here, partial interface layouts are the subproblem result being cached. To boost performance, we take a branch-and-bound approach, where subproblem solutions estimated to be worse than the current best are immediately discarded.

Optimizing the layout of discrete elements—like thumbnails—is more efficient than for continuous elements because the algorithm need only search the relatively small number of sizes that yield an exact integer number of elements either across or down for a given size. For the examples pane, we vary the number of example thumbnails shown and the utility of the selected examples. Finding the optimal size and layout for a particular display size is linear in the number of examples shown in a row. Once the overall size and layout are determined, finding the best set of content items to display then becomes a greedy search, linear in the number of elements shown in a row. In practice, the optimized algorithm allows Adaptive Ideas to render interfaces at interactive speeds on current hardware.

CAN EXAMPLES SCAFFOLD DESIGN ABILITY?

To understand if and how users would benefit from explicit mechanisms for designing with examples, we conducted three studies. The first tested whether participants produced better pages when designing *with* an example-augmented editor than without one. The second tested whether *adaptively* selecting examples for display yielded a different user experience than *randomly* selecting them. The third studied *how* participants make use of examples when there is explicit tool support for doing so.

Study 1: Designing With vs. Without Examples

This experiment comprised two parts. The first part asked participants to design Web pages for a specified persona. In the second part, a separate set of participants rated the pages designed in the first part.

Method, first part: designing

Twenty-six students from our university participated in the first part of the study for course credit. They were roughly evenly divided between engineering and non-engineering majors. All were frequent Web users. Half had little-to-no Web design experience (here, we refer to them as “novices”), half had some prior Web design experience (here, we refer to them as “experienced”).

Participants created two Web pages: one using the *Examples* editor; one using the *Control* editor. Editor order was counterbalanced. In both conditions, users were presented with 12 templates to choose from as the initial layout; they were free to modify these initial designs as they saw fit. We evenly assigned the novice and experienced participants across the two groups. All participants first created a page for an *Elaine* persona, then created a page for a *Bob* per-

sona. Appendix A presents the Elaine persona; the Bob persona was similar.

The *Examples* condition used a version of Adaptive Ideas with two attributes disabled. First, the examples gallery was set to always display a random subset of the examples; the rationale for this was to first test whether a straightforward example augmentation helped. (Experiment 2 looks at the benefits of adaptive display.) Second, we disabled the mechanism for directly borrowing examples’ elements; users had to manually specify the desired attributes in the editor. The rationale here was similar: direct borrowing might provide a performance enhancement, and so requiring manual example borrowing is more conservative with respect to the hypothesis. The *Control* condition did not display the example pane and focus pane to the user. Screen size was held constant across conditions. As such, the design pane had more real estate in the control condition.

Users were presented with a questionnaire before and after the task. Some questions—aimed at tracking self-confidence in Web design and perceived need for external assistance—were repeated before and after. In the post-task questionnaire, users were also asked to report on their satisfaction with the Web page they had made.

After completing the *Examples* condition, users answered additional questions about whether viewing examples made the design process more engaging; the effects of examples on a user’s evaluation of his or her own Web page; the influence of examples on the participant’s resulting design, and the types, number, and size of desired examples.

Method, second part: rating

Forty-six students from our university participated in the second part of the study for course credit in an HCI or communications course. None had participated in the first part. These students included a mix of engineering and non-engineering majors. This second part was conducted over the Web. Each participant rated a subset of the 50 pages created in the first part. (Two pages for each of 25 participants; one participant did not successfully complete the study.) Participants were shown 7 randomly selected pages. At the beginning, participants were shown the persona description, and the site explained they would see several pages that people had created for that persona. For each page, the survey asked five questions about the quality of the page and suitability for that persona. Raters responded to each question using a 7-point Likert scale; Figure 3 lists the questions and summarizes the scale used.

Results

To avoid ordering confounds, the following analysis looks exclusively at each participant’s first page: the Elaine persona. This yields a between-subjects comparison. The analysis compares the average rating of each page. Pages created in the *Examples* condition were rated more highly than those in the *Control* condition ($M=4.04$ vs. 3.37 , $F=4.98$, $p < 0.05$; see Figure 4). Also, experienced participants created more highly rated pages than novices

Rate on a scale of 1 to 7, where 1 = not at all, 4 = moderately, and 7 = a lot

1. How suitable is the color scheme of the Web page for the persona depicted above?
2. How suitable is the layout of the Web page for the persona depicted above?
3. How suitable is the overall design of the Web page for the persona depicted above?

Below, the scale is 1 = strongly disagree and 7 = strongly agree

4. This is a sophisticated Web page.
5. This is a visually attractive Web page

Figure 3 The five questions raters were asked about each of the pages created in Experiment 1.

($M=4.00$ vs. 3.41 ; $F=3.59$, $p<0.05$). In this study, there was no significant interaction between expertise and manipulation; *i.e.*, experienced participants and novices benefited equally from examples.

The design participants also provided input on how the examples interface may be improved. The study presented 20 example pages in the example pane, laid out in five rows of four examples at 103 pixels \times 64 pixels each. Participants generally reported that they would prefer fewer examples, displayed larger. When told that the pane showed 20 example Web pages and asked what would be the ideal, the average response was 10.

Study 2: Adaptive vs. Random Example Selection

Study 1 showed that users created higher-rated pages when presented with examples. The second study investigated users' views on different interfaces for browsing examples. Specifically, it explored the relative merits of adaptively selected examples and randomly selected examples. Nine subjects participated in this study; they were compensated with a US\$15 gift certificate. Participants' ages ranged from 24 to 30. All were frequent Web users; two self-rated as experienced Web designers; the others had little-to-no Web design experience.

Method

Participants were seated at a workstation with the Adaptive Ideas editor, where they could browse example pages and borrow elements from them. Sessions began with a demonstration of its functionality. In this within-subjects study, participants created Web pages for the Elaine and Bob personas; the interface in both conditions was identical except for how examples were structured. The *standard* condition disabled the similarity and variety features and sorted the examples randomly: users could view all examples, but could only browse them using the next and previous page controls. The *adaptive* condition enabled browsing by similarity and variety. Personas and interfaces were counterbalanced. At the end of the study, after both tasks, participants completed an eighteen-question survey. This questionnaire asked about their background and about the user experience of the two conditions. It referred to the

alternatives as “circle” and “square”—*not* whether it was adaptive.

Results

Participants reported that “it was easier to navigate with [adaptively presented examples] than [random ones]” ($M=3.7$, $SD=1.22$, on a 5-point scale), and that both the variety tool ($M=4.5$, $SD=0.53$) and similarity tools ($M=4.4$, $SD=1.01$) helped them find examples. Participants disagreed with the statement, “examples distract from the design task” ($M=2.2$, $SD=0.83$). Several users expressed a desire to browse along aesthetic or social attributes, such as formality.

Participants selected half as many examples for larger viewing in the adaptive condition than the random condition ($M=96$ vs. 196 , $p<0.05$). This may be because the adaptive tools facilitated more directed exploration. When working with randomly presented examples, several participants resorted to long stretches of clicking on many examples in a row in order to examine them. For larger corpora, the benefits of presenting examples adaptively rather than randomly will likely be even greater.

Overall, novices rated the value of examples more highly than experienced designers, though not significantly so. The two self-rated experienced users differed in their opinions about the use of examples for Web design. One thought examples wasted screen real estate and distracted from the task. The other commented that the browsing of examples




Figure 4 On all five scales, independent viewers rated pages created in the *Examples* condition significantly more highly than pages in the *Non-examples* condition. In each pair, the left-hand bar is the *Examples* condition; the right-hand bar is the *Non-examples* condition. The sections of each bar represent the fraction of ratings of that value.

worked well with her personal strategy for this type of design task: “That’s my philosophy of designing Web sites: I like to find a template or exemplar that I think is good and then tweak it by hand.” This suggests that, to effectively use limited screen real estate, people should be able to hide the example pane when not using it.

Study 3: Common Usage Patterns

The first study disabled the ability to directly borrow examples so as to isolate the value of *seeing* examples. Would users benefit even more if they could also directly *use* example elements? Study 3 investigates this question.

Method, first part: designing

Nine students from our university participated in this study for course credit. Prior experience varied widely: one participant had never made a Web page before, most had prior experience with visual authoring tools, the most-experienced had several years of professional experience. As in the previous experiment, participants were seated at a workstation with the Adaptive Ideas editor and sessions began with a demonstration of its functionality. Participants were then asked to create Web sites for the Elaine persona (see Appendix A). Each user was presented with a fully functioning Adaptive Ideas editor, and the system logged participants’ major interactions. The same questionnaire as the previous study was given both before and after the task.

Method, second part: rating

One hundred participants participated in the second part of the study for US\$1 each; participants were recruited using Amazon’s Mechanical Turk [3]. This second part was conducted over the Web, and asked participants to rate the nine pages created in the first part. Raters responded to each question using a seven-point Likert scale.

Results

There was a modest positive correlation between the amount of time a participant spent browsing for examples and the Turkers’ mean rating of the page created ($r^2=0.185$). Usage logs show that participants predominantly explored examples early on, and nearly always borrowed one of the examples they saw (see Figure 5). In this study, users could browse the corpus *adaptively* to look for items similar to a focus element or for a variety of items. They could also browse *randomly*. When users browsed with adaptive display, they more quickly found an example to work from than when they browsed with random display.

All 7 instances of search by similarity were directly followed by closer examination of at least one or more examples in the resulting gallery; this was also true for 8

out of the 10 instances of search by variety. When searching by similarity, participants viewed an average of 3.3 pages from each gallery ($SD=2.3$). When searching by variety, participants viewed 6.1 pages on average ($SD=6.3$). We hypothesize this is because when searching by similarity designers are engaged in a more directed searching task, while searching by variety is indicative of more open-ended browsing.

Approximately half the time, users clicked through many iterations of random search to find a new focus example. They also viewed examples less frequently. 13 of the 27 instances of random search took multiple iterations to find an example that piqued the user’s interest; in 2 cases it took more than 15 iterations.

The results suggest that people use the media *content* of example pages—in addition to the formal aspects of layout and design—in deciding what to select. When selecting a starting example for the design of Elaine Marsh’s page, 6 of the 9 users only selected home pages of research students, and 3 of the participants independently selected the same Web page to begin with. The Web page had a photo of someone who appeared “studious and reserved.” One even left the photo in the final design.

DISCUSSION AND CONCLUSION

This paper presented results from three experiments demonstrating that explicit interface mechanisms for working with corpora of examples can improve design work. These experiments isolated the effects of *presenting* examples, *adaptively browsing* examples, and *borrowing* example elements, finding that each of these pieces aided design. To understand these questions, we introduced the Adaptive Ideas system. Adaptive Ideas introduces techniques for dynamically selecting content and generating layouts of examples using a combination of decision-theoretic selection, designer specification, and end-user preference. This work raises a number of important questions.




Figure 5 This graph shows when users are *browsing* through examples (*blue*), when they *borrow* elements (*orange*), and when they *start over* by replacing their work with a new example (*pink*).

First, in this paper, the Web pages were manually harvested and the metadata manually applied. Can this be automated? The modern Web, through CSS and application frameworks, has begun to encourage a separation of content and layout. Is it possible to automatically adapt the style of one page to the content of another? Kumar *et al.* have begun exploring this direction using machine learning [25].

Second, this paper focuses on the Web, and specifically on individual pages. Could similar approaches be applied to multi-page examples and Web *sites*? What tools could enable other design domains to exploit an example-based approach?

Third, an important difference between examples and design patterns is the level of curation. Example galleries, as used in this paper, are naturally occurring and thereby abundant. Want to see a hundred different pages with a black background, or that use a particular typeface? No problem. Design patterns, by contrast, contain carefully selected examples (sometimes even constructed ones) and the examples accompany a significant description. From a user's perspective, what are the relative benefits of a large (but uncured) corpus and a smaller (but curated) one? Do different situations benefit from different approaches? From both a technical and a conceptual perspective, how can (and should) design patterns and example galleries be merged?

Fourth, returning to our inspiration of the importance of examples in art and design education, what are the potential learning benefits of example-based tools? Can this approach help novices acquire and internalize the intuitions of experts? If so, can these users eventually abandon the scaffold that examples provide and work just as effectively, or do examples remain valuable across a broad spectrum of expertise levels?

Fifth, can example-based design tools benefit the very best and brightest in a field to produce higher quality work? The psychological literature suggests more "distant" examples help experts achieve creative leaps [14]; it remains an open question how to design such tools. Even if example-based tools "only" raised the quality of average design, the impacts would be substantial.

Finally, the ethical and legal issues around working with examples are actively evolving. Plagiarism clearly runs afoul of both morality and the law. By contrast, leveraging prior ideas to create new ones enjoys significant legal support because of the value that fair-use adaptations provide to society [28]. As the writer Jonathan Lethem describes it, "apprentices graze in the field of culture" [27].

ACKNOWLEDGEMENTS

We thank the National Science Foundation (IIS-0745320 and IIS-0534662) and the Hasso Plattner Design Thinking Research Program for supporting this work.

REFERENCES

1 Li'l Abner, 2009.

http://en.wikipedia.org/wiki/Li'l_Abner

- 2 Alexander, C., S. Ishikawa, and M. Silverstein, *A Pattern Language*: Oxford University Press. 1977
- 3 Amazon, *Mechanical Turk*, 2009. <http://requester.mturk.com>
- 4 Boden, M. A., *The creative mind: Myths and mechanisms*: Taylor & Francis. 2003
- 5 Brandt, J., P. J. Guo, J. Lewenstein, M. Dontcheva, and S. R. Klemmer. Opportunistic Programming: Writing Code to Prototype, Ideate, and Discover. *IEEE Software* **26**(4), 2009
- 6 Buxton, B., *Sketching user experiences: getting the design right and the right design*: Morgan Kaufmann. 2007
- 7 Carter, D. E., *The Big Book of Logos*: Watson-Guptill. 2001
- 8 Chi, M. T. H., P. J. Feltovich, and R. Glaser. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 1981
- 9 Cook, R. L. Stochastic sampling in computer graphics. *ACM Transactions on Graphics (TOG)* **5**(1). pp. 51–72, 1986
- 10 Cross, N., *Designerly Ways of Knowing*: Birkhäuser. 2006
- 11 Fitzgerald, M., *CopyStyler : Web design by example*, Unpublished MS, MIT, Electrical Engineering and Computer Science, Cambridge, MA, 2008. <http://dspace.mit.edu/handle/1721.1/46032>
- 12 Gajos, K. and D. S. Weld. SUPPLE: automatically generating user interfaces. In *Proceedings of IUI 2004: Proceedings of the 9th international conference on Intelligent user interface*. pp. 93–100, 2004
- 13 Gamma, E., R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. 1st ed: Addison-Wesley Pub Co. 1995
- 14 Gentner, D. Analogy in scientific discovery: The case of Johannes Kepler. *Model-based reasoning: Science*, 2002
- 15 Gentner, D., K. J. Holyoak, and B. N. Kokinov, *The Analogical Mind: Perspectives from Cognitive Science*: MIT Press. 2001
- 16 Gentner, D., J. Loewenstein, and L. Thompson. Learning and Transfer: A General Role for Analogical Encoding. *Journal of Educational Psychology* **95**(2). pp. 393-408, 2003
- 17 Gick, M. L. and K. J. Holyoak. Analogical problem solving. *Cognitive psychology*, 1980
- 18 Gick, M. L. and K. J. Holyoak. Schema induction and analogical transfer. *Cognitive psychology*, 1983
- 19 Hartmann, B., S. Doorley, and S. R. Klemmer. Hacking, Mashing, Gluing: Understanding Opportunistic Design. *IEEE Pervasive Computing* **7**(3). pp. 46-54, 2008
- 20 Hartmann, B., L. Wu, K. Collins, and S. R. Klemmer. Programming by a Sample: Rapidly Creating Web Applications with d.mix. In *Proceedings of UIST 2007: ACM Symposium on User Interface Software and Technology*. pp. 241-50, 2007

- 21 Hartmann, B., L. Yu, A. Allison, Y. Yang, and S. R. Klemmer. Design as exploration: creating interface alternatives through parallel authoring and runtime tuning. *UIST: Proceedings of the ACM symposium on User interface software and technology*, 2008
- 22 Herring, S. R., C.-C. Chang, J. Krantzler, and B. P. Bailey. Getting inspired!: understanding how and why examples are used in creative design practice. In *Proceedings of Proceedings of the 27th international conference on Human factors in computing systems*: ACM, 2009
- 23 Kim, H. and W. C. Yoon. Supporting the cognitive process of user interface design with reusable design cases. *International Journal of Human-Computer Studies* **62**(4). pp. 457-86, 2005
- 24 Kolodner, J. L. and L. M. Willis, Case-Based Creative Design, in *AAAI Spring Symposium on AI and Creativity*. 1993: Stanford, CA
- 25 Kumar, R., J. Kim, and S. R. Klemmer, Automatic retargeting of web page content, in *Proceedings of the 27th international conference extended abstracts on Human factors in computing systems*. 2009, ACM: Boston, MA, USA
- 26 Lee, B. A., *Adaptive interaction techniques for sharing and reusing design resources*, Stanford University, Computer Science, Stanford, CA, 2007.
- 27 Lethem, J. The ecstasy of influence: A plagiarism, *Harper's Magazine*(2): pp. 59-71, 2007
- 28 Leval, P. N. Toward a Fair Use Standard. *Harvard Law Review* **103**, 1990
- 29 Lodi, A. and M. Monaci. Integer linear programming models for 2-staged two-dimensional Knapsack problems. *Mathematical Programming* **94**(2-3). pp. 257-78, 2003
- 30 Marks, J., B. Andalman, *et al.*, Design galleries: a general approach to setting parameters for computer graphics and animation, in *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. 1997, ACM Press/Addison-Wesley Publishing Co.
- 31 Marsh, R. L., J. D. Landau, and J. L. Hicks. How examples may (and may not) constrain creativity. *Memory and Cognition* **24**. pp. 669-80, 1996
- 32 Mullet, K. and D. Sano, *Designing visual interfaces: communication oriented techniques*: Prentice-Hall, Inc. 1995
- 33 Naaman, M., S. Harada, Q. Wang, H. Garcia-Molina, and A. Paepcke. Context data in geo-referenced digital photo collections. In *Proceedings of MM2004: ACM International Conference on Multimedia*. pp. 196-203, 2004
- 34 Nardi, B. A., *A Small Matter of Programming: Perspectives on End User Computing*: MIT Press. 1993
- 35 Newman, M. W. and J. A. Landay. Sitemaps, Storyboards, and Specifications: A Sketch of Web Site Design Practice. In *Proceedings of DIS: Symposium on Designing Interactive Systems* ACM Press. pp. 263-74, August 17-19, 2000
- 36 Novick, L. Analogical transfer, problem similarity, and expertise. *Journal of experimental psychology. Learning, memory, and cognition* **14**(3). pp. 510-20, 1988
- 37 Rosenholtz, R., Y. Li, J. Mansfield, and Z. Jin, Feature congestion: a measure of display clutter, in *Proceedings of the SIGCHI conference on Human factors in computing systems*. 2005, ACM: Portland, Oregon, USA
- 38 Schank, R. C. and P. G. Childers, *The creative attitude: learning to ask and answer the right questions*. 1988
- 39 Schön, D. A., *The Design Studio: An Exploration of its Traditions & Potential*: RIBA Publications Limited. 1985
- 40 Shah, J. J. and M. Mantyla, *Parametric and Feature Based CAD/Cam: Concepts, Techniques, and Applications*: John Wiley & Sons, Inc. 1995
- 41 Smith, S. M., T. B. Ward, and J. S. Schumacher. Constraining effects of examples in a creative generation task. *Memory and Cognition*, 1993
- 42 Snarey, M., N. K. Terrett, P. Willett, and D. J. Wilton. Comparison of algorithms for dissimilarity-based compound selection. *Journal of Molecular Graphics and Modelling*, 1997
- 43 Spearman, C. E., *The nature of "intelligence" and the principles of cognition*. Macmillan. 1923
- 44 Talton, J. O., D. Gibson, L. Yang, P. Hanrahan, and V. Koltun. Exploratory modeling with collaborative design spaces. *ACM Transactions on Graphics (TOG)* **28**(5). pp. 1-10, 2009
- 45 Terry, M. and E. Mynatt. Side views: persistent, on-demand previews for open-ended tasks. *Proceedings of the 15th annual ACM symposium on User interface software and technology*. pp. 71-80, 2002
- 46 Terry, M. and E. D. Mynatt. Recognizing creative needs in user interface design. *Proceedings of the 4th conference on Creativity & cognition*. pp. 38-44, 2002
- 47 Tohidi, M., W. Buxton, R. Baecker, and A. Sellen. Getting the right design and the design right. In *Proceedings of CHI: Conference on Human Factors in Computing Systems*: ACM Press, 2006
- 48 Van Duyne, D. K., J. A. Landay, and J. I. Hong, *The Design of sites: Patterns for creating winning web sites*: Prentice Hall PTR. 2007
- 49 Weisberg, R. W., Creativity and Knowledge: A Challenge to Theories, in *Handbook of Creativity*, R.J. Sternberg, Editor. Cambridge University Press. pp. 226-50, 1999
- 50 Zhou, M. X. and V. Aggarwal. An optimization-based approach to dynamic data content selection in intelligent multimedia interfaces. In *Proceedings of UIST 2004: ACM Symposium on User Interface Software and Technology*. pp. 227-36, 2004

APPENDIX A: PERSONA AND TASK

Elaine's persona was described as follows:

Elaine Marsh is a 21-year-old economics student, starting her senior year. Studious and reserved by nature, Elaine spends much of her time outside the classroom serving as vice president of the student business association. She also volunteers as a tutor at a local high school. Elaine wants to make a homepage that details her undergraduate activities, including class projects, research papers, and leadership positions.

In her personal life, Elaine enjoys making origami and traveling to new places. During her time at college, she has traveled to New York, Boston, Paris and London. She lives in the arts-themed dorm, Kimball.

Her vision for the page includes a mature, sophisticated design and a somewhat professional feel.

Participants were asked to make a Web page for Elaine with the following minimum requirements:

- Choose a layout suitable for Elaine's vision of her page.
- Place an image that is relevant to Elaine's activities. This image may be from an example page, a Web page on the Internet or somewhere on the computer.
- Customize a vertical or horizontal (or any other fancier variation) menu bar.
- Set the Color of the page such that the color scheme is true to Elaine's vision of her page.
- Set the fonts of the page.
- Enter textual information about Elaine.

Beyond these requirements, the users were free to customize the page as much as desired.

APPENDIX B: WEB PAGE METADATA

Each Web page was manually tagged with the following metadata.

<u>Property</u>	<u>Range of Values</u>
Font-Family	Name
Font-size	As found on page
Number of Images	As found on page
Image:Text:Whitespace Area	Ratios add up to 100
Background Color	RGB values
Use of Background-Image	YES or NO
Placement of Personal Picture	1 to 9 (Divide page into 3 rows and 3 columns)
Formality of Personal Picture	1 to 10
Vertical or Horizontal Navigation Bar	Vertical or Horizontal or curved
Number of Side Bars	As found on page
Use of Header	YES or NO
Visual Density of Page	1 to 100 (subjective value)
Date of Creation	As found on page
Use of Pictures as Links	YES or NO
Is Page Centered	YES or NO