CrowdCritter: Strategies for Crowdsourcing Visual Design Critique

Wei Wu
Kurt Luther
Amy Pavel
Björn Hartmann
Steven Dow
Maneesh Agrawala

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2013-95
http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-95.html

May 17, 2013

Acknowledgement

This work was partially funded by NSF grant IIS-1210836.

CrowdCritter: Strategies for Crowdsourcing Visual Design Critique

By Wei Wu

Research Project

Submitted to the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, in partial satisfaction of the requirements for the degree of **Master of Science, Plan II**.

Approval for the Report and Comprehensive Examination:

Committee:

Professor Björn Hartmann
Research Advisor

May 16, 2013

* * * * * * *

Professor Maneesh Agrawala
Second Reader

Abstract
Crowdsourced labor has successfully tackled a variety of rote tasks like image tagging and transcription, but has difficulty accomplishing complex tasks. Visual design critique is a crucial part of the design process for improving upon existing work, but it is a complex task that requires some implicit knowledge of design concepts on the part of the critiquer. In this report, we propose a model for crowdsourcing visual design critique to enable online workers who lack background knowledge to provide useful feedback on designs. Our model organizes the "crit" around a set of design principles, and asks workers to select relevant statements about design principles and provide freeform comments within the scope of the selected statements. We evaluate two variants of a critiquing system we built based on this model (CrowdCritter) through a study with 73 participants. We found that the integration of a set of design principles allowed novice workers to provide more critiques than freeform feedback alone. We were unable to identify significant differences between the two variants of CrowdCritter in the quality and level of agreement of critiques produced through each interface.

Table of Contents

Introduction

Paid online crowdsourcing where employers post short-term jobs to online platforms for small monetary incentives is a growing job sector in the digital economy. In the past decade, over one million workers have earned $1-2 billion through crowdsourced labor markets [1]. These socio-computational marketplaces appeal to employers because they can complete high volumes of work without committing to long-term contracts. Workers appreciate the crowdsourced model because they can select the tasks they wish to complete and perform work on demand. Today's crowdsourcing markets like Amazon's Mechanical Turk [17], CrowdFlower [18], and Microtask [19] usually only leverage workers' basic cognitive skills to complete rote tasks like audio transcription, business address verification, or image labeling. More complex tasks such as writing an essay or designing an advertisement require employers to recruit more permanent online contractors who have the requisite skills for the assignment.

In this report, we explore approaches to completing complex tasks in micro-task crowdsourced markets. Specifically, we focus on the task of providing critiques for visual designs. Traditionally, a "crit" is a co-located communication event where someone presents his work and critics, usually teachers and peers, provide feedback to improve the design [27]. Online, visual design critique has evolved from its offline form into two broad categories: community critique and crowd feedback. On community critique sites, members share their designs and provide feedback on each other's work without financial incentives. Crowd feedback, on the other hand, refers to paid services where a designer pays to have the crowd tell their impressions of a design. These existing services are both limited in the granularity of feedback that is returned to the original designer.

In this paper, we present CrowdCritter, a system for crowdsourcing visual design critique that allows designers of all experience levels to receive useful feedback about their designs from workers in a crowdsourced labor market. The provision of useful feedback requires a shared vocabulary for expressing ideas about the design domain between critique seekers and providers. This is challenging in a crowdsourced labor market, where workers typically lack specialized knowledge and skills related to design aesthetic.

To make the complex creative task of design critique feasible for crowdworkers, we propose the following model: (1) critique is organized around a set of design principles exposed in the user (2) workers select relevant statements about design principles and provide comments within the scope of selected principles, (3) multiple workers' responses are aggregated to get full coverage of critiques for a design.

In this report, we first draw on existing theories of critique and feedback to arrive at a definition of "useful" feedback and frame our approach of incorporating structured design principles into the critiquing process. We then review related work in crowdsourcing and critique systems. Next, we present data from several pilot studies that suggest that novice online workers do not provide useful critiques in the absence of structured critiques. We then describe the design and implementation of the CrowdCritter system. Finally, we discuss a formal study we conducted to evaluate the effectiveness of CrowdCritter with 73 participants. Our evaluation compares two versions of the CrowdCritter interface that offer different strategies to elicit critique drawn from

the learning literature; we also compare both to the status quo of freeform feedback. We find that the integration of structured design principles into critiquing interfaces allows inexperienced workers to provide more critiques than freeform feedback. However, we are unable to identify significant differences in the quality and level of agreement of critiques provided through the two CrowdCritter interfaces.

Background: From Traditional Critiques to Online Critiques

The goal of conducting critiques in studio art education is to improve the art making of the students [29]. A "crit" is a co-located communication event where someone presents his work and critics, usually teachers and peers, provide feedback to improve the design [27]. Critique not only leads to knowledge sharing, it helps inculcate the values and aesthetics important to an organization. Critique providers also learn by developing a vocabulary for a design domain [30].

Researchers in art and design education have established how critiques can aid the design process [9, 10, 29]. Feldman [9] proposed a four-part framework for carrying out critiques which involves description, analysis, interpretation, and judgment. Teo and Chai built on this framework and applied it to collaborative critiquing of video projects and found that the critiquing process helped facilitate novice videographers' progress towards expertise [11].

Outside of aesthetic education, Sadler has reviewed feedback mechanisms for student assessment [12]. He argues that good feedback must incite a student to possess a concept of a standard, compare the actual level of performance with this standard, and engage in action that closes this gap [12]. Applying this to visual design critique, we argue that a good critical critique statement must (1) identify a specific element in which a certain design principle that is violated, (2) recognize why the work falls short with respect to the principle, and (3) offer some means of addressing the issue. Only feedback that encompasses all three components will be useful to the original designer. Without (1), the designer cannot clearly identify what the problem is; without (2), she cannot validate that the element indicated is actually a problem; without (3), she has no direction in how she can address the issue. These theories about critique in offline contexts inform our design of a system to facilitate visual design critique online by crowdsourced workers.

Participants in a traditional crit usually already share a set of terminologies to effectively communicate feedback under Sadler's framework with each other because they have background in visual design. In a crowdsourced labor market, however, most people lack visual design experience. They cannot bring to mind the vocabulary necessary to frame a critique in visual design theory because they do not deal with the domain-specific lexicon frequently. We hypothesize that while novices are incapable of providing effective critique along Sadler's guidelines when given a purely freeform format, they can provide feasible design critiques when visual design theory is made salient to them in the form of structured design principles. In this paper, we propose an interface where workers select statements from a holistic list of design principles that are applicable to a given design.

CrowdCritter uses the task decomposition paradigms established by existing work to accomplish the complex task of critiquing visual design, but is novel in that the interface introduces unskilled workers to design principles they need to know to make useful critiques. To inform our interface

design, we draw on learning theory literature that explores how best to teach new concepts to students. Schwartz et al. reviews a large body of work that looks at how people transfer in knowledge to tackle new situations and how they transfer out knowledge from the process [2]. In an empirical study, Schwartz tests how two different teaching techniques, the "tell and practice" and the "invent with contrasting cases" affects how eighth graders learn the concept of density. The former approach is a traditional technique in which students are shown a concept and asked to apply it in subsequent problems. The latter is an alternative teaching method where students are guided through a concept by exploring on their own first. We adapt these two techniques in two prototypes of the CrowdCritter interface to test which works better in the context of teaching visual design principles and critique statements.

Related Work

Much of recent work in crowdsourcing has centered on accomplishing complex, multi-step tasks in a micro-task environment like Amazon's Mechanical Turk. Kittur et al.'s work with CrowdForge develops a partition-map-reduce framework for dividing up separate components in a complex task's workflow amongst many workers and combining them into a cohesive result at the end, which they apply to article writing, decision-making, and scientific journalism [4]. Soylent, a system devised by Bernstein et al., achieves paper editing in a crowdsourced micro-task context [5]. Others have harnessed crowdsourced workers to provide complex services like nutritional analysis and email task management through task decomposition [6, 7]. These applications, while complex, do not require workers to have prerequisite skills or domain-specific knowledge besides basic language proficiency to complete the tasks.

Work related to visual design, however, necessitates an understanding of design concepts and their manifestation [14]. Xu and Bailey proposed a model for crowdsourcing design critique that applies CrowdForge's task decomposition paradigm to Feldman's process for carrying out critiques [13]. The output of their model, while better than those received in a design feedback community, is not framed by principles of visual design theory.

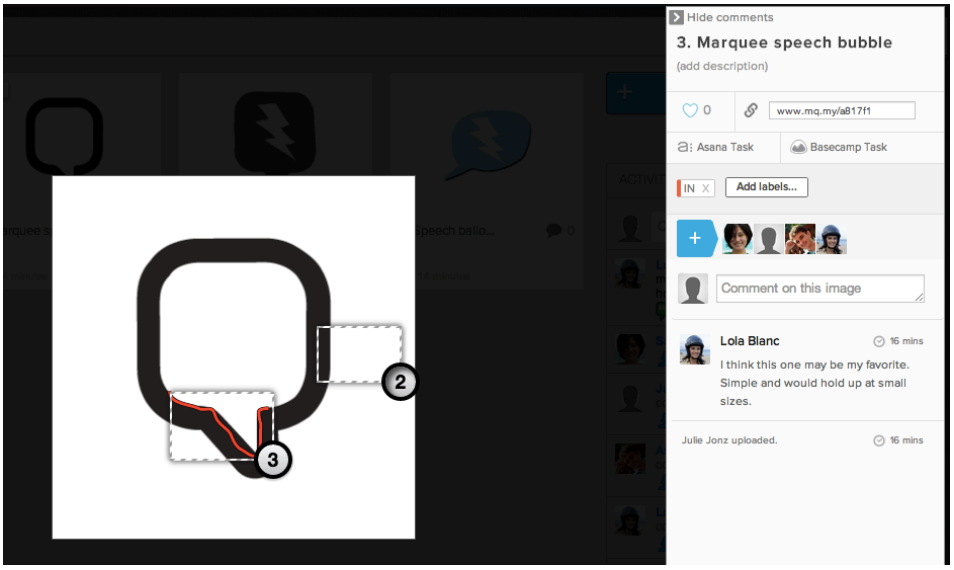

Figure 1: Community feedback site Marqueed lets you upload designs and invite others to annotate your work and add comments about the piece. In the example above, people have contributed three comments total.

There are several commercial solutions to online visual design critique that generally fall into two broad categories: community critique, and crowd feedback. On community critique sites, members upload their design projects and provide feedback on each other's work (Figure 1). Members receive no financial rewards for their critiques, but are driven instead by reputation. In our review of such sites, we found that only established designers whose calibre of work is already very high put their work up for critique on community critique sites. Existing community critique sites like Marqueed [15], and Dribbble [16] are not geared for amateur designers or non-designers who are tasked with a creative assignment, like making a slide presentation or a flyer for an organization, to receive feedback on their work. In addition, responses tend to be superficial statements such as I like it" or "I don't like it" [13]. In a survey of a digital photography community, the median number of critiques per photo uploaded was only 2 [28].

Crowd feedback refers to services where a designer or site owner pays for the crowd to provide impressions of their design. The services conduct usability tests with the crowd, and aggregate the feedback into word clouds, charts, and heatmaps (Figure 2). While designers can receive evaluations back quickly through crowd feedback, the information they receive is primarily a collection of reactions, not critiques. These services capture people's overall impressions of a design, but do not provide fine-grained, actionable critiques. Crowd feedback includes sites like Five Second Test [21], Usabilla [22], and Feedback Army [23].

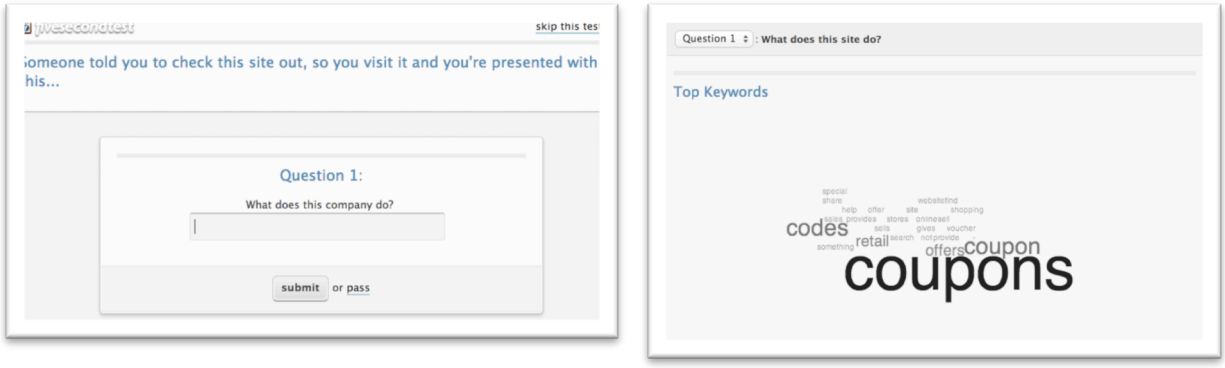

Figure 2: Crowdfeedback service FiveSecondTest shows a user a screenshot of a site for 5 sec. and asks her questions about her impression of it (left). It presents word-cloud reports of the responses to the requestor (right).

Unlike existing visual design sites now, CrowdCritter aims to elicit more specific, actionable feedback from workers who do not necessarily have design experience at all. Our solution is also amenable to designers of all levels and mediums who would like their work evaluated.

Formative Pilot Studies

To understand how novice crowd workers approach tasks related to visual design, we first ran several exploratory studies on Mechanical Turk. The tasks centered around providing feedback on an individual slide from a Powerpoint slide deck or a poster. These pilots informed our approach in designing the final CrowdCritter interface to incorporate structured design principles as part of the critiquing process.

Pilot #1: Naïve visual design feedback
For the first pilot, we selected a pool of eight visual designs (six slides and two posters) for workers to critique. Workers were shown an image of a design, and given the following prompt and a single text box to write their response:

> Write down the problems you see with this design. If something doesn't look right or could be improved, please explain here.

Three workers provided responses for each design. This pilot shed light on the types of feedback that workers would provide when prompted in a freeform manner. Typically, responses fell into a combination of three components: (1) identifying a problem, (2) explaining why it's a problem, and (3) proposing a solution. These three parts correspond to Sadler's guidelines for good feedback and confirm their applicability to visual design critique.

Figure 3a shows an example response from the pilot that addresses all three components. The worker stated that (1) the background is a problem because (2) it makes it hard to concentrate on the foreground and (3) suggests using a plainer background instead. Figure 3b shows another worker's response to the same slide, which only proposes an improvement ("More images of the people should be included in the design.") without citing the problem or justifying it. This feedback is less helpful to a designer because the suggestion has no context—without comprehending a specific problem from the critique, she has no reason to fix it. Critique statements should include all three components to be useful.

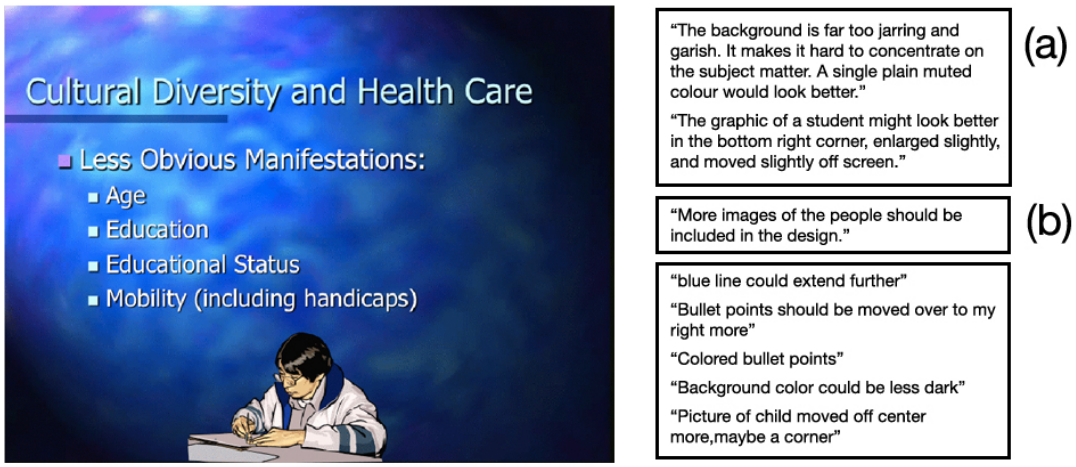

Figure 3: Three workers' responses from Pilot 1, which asked Mechanical Turk workers to write down problems with the design on the left. (a) is a high quality useful critique while (b) is an unsubstantiated opinion and more representative of most responses in the study

Three workers contributed a total of 8.125 statements on average for each design. Out of 65 separate statements, only 5 of the responses (7%) we received in this pilot encompassed all three of the criteria for useful critique. In fact, 20% of the statements are similar to Figure 1b, which, lacking reasoning to support the statement, are merely subjective comments. The remaining 73% of responses have some combination of the three components but do not encompass all three. In total, we see that 93% of the responses are not useful to the designer for improving her work.

This initial naïve approach to crowdsource visual design critique suggests that novices need more guidance to provide useful feedback that situates suggestions in design principles.

Pilot #2: Structured freeform visual design feedback
Because critiques provided by novices on Mechanical Turk in the first pilot were largely unsubstantiated by visual design principles, we ran a second pilot to see if asking specific questions would guide workers to express and justify a design's shortcomings (Figure 4a). In this study, three workers are asked to identify three problems in a design (a poster or a slide) for four designs total. Unlike the first pilot, the interface prompts workers to both describe the problem and explain why it is a problem in two separate textboxes.

Example results from this task (Figure 3b) show that this breakdown of a critique statement still fails to elicit clear identification of and reasons for problems with a visual design. Each worker interpreted what she was supposed to say in the boxes differently. One worker simply stated characteristics of a design as problematic without clarifying why (problem: "color", explanation: "the picture is not clear"). Another worker justifies the issues he identifies with reasons that lack visual design foundations (problem: "logo", explanation: "it should be on the top left/right corner"). The results of this pilot, along with pilot #1, suggested that novices or workers with no prior knowledge about visual design would be largely unsuccessful at providing useful design critiques when just prompted for freeform opinions.

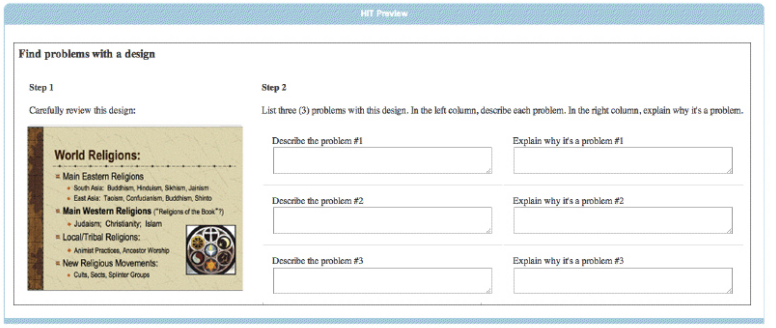

(a)

	Problem	Explanation
Worker 1	Logo (picture)	It should be on the top left/right corner
Worker 1	Highlighting of Main Western Religions	Why only " Main Western religions" are highlighted?
Worker 1	Inconsistency in "Colons" after religion subtitle	Local/Tribal Religions have 'colon' at the end. While other sub-headers don't have
Worker 2	Font size	Some fonts are bold and some are not , hence the bold ones takes all the attention
Worker 2	Image is small.	Image is always a attention seeker hence it has to be made large.
Worker 2	Background.	Background does not suit the content, it will be good if relative background is used.
Worker 3	format	the format is bad
Worker 3	picture	the picture is not good
Worker 3	information	the information details is low

(b)

Figure 4
(a): The second pilot displays a design and asks the worker to describe up to 3 problems and explain why they are a problem separately.

(b) Actual responses by the three workers provided responses to the slide shown in (a). The specificity of the questions still fails to elicit responses that clearly state the problem and justify it with design theory.

Pilot #3: Structured checklist visual design feedback
The first two pilots led us to explore a more structured paradigm for crowdsourcing visual design critique that would better guide novices who lack design experience. For this pilot, we devised a checklist of generalized visual design principles pertaining to alignment and font usage.

Participants were shown designs from a pool of nine slides, and asked to indicate whether statements in the checklist pertained to the slide. They were also given the option to elaborate on the statement by describing what and where the problem was on the design in a text box.

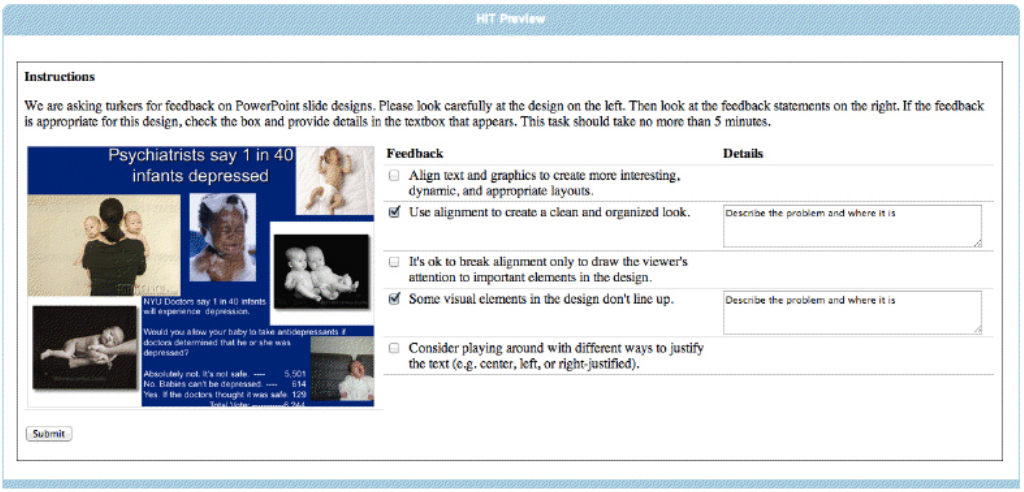

Figure 5: Workers in the third pilot checked relevant statements from a list pertaining to either alignment or font usage for a given design. There was 65% unanimous agreement between 3 workers about alignment-related statements and 52% agreement about font usage.

The pilot drew from a set of 4 slides for alignment, and 5 slides for font usage. Three workers looked over each design with the same checklist. We then reviewed the degree of consensus between their responses. For alignment principles and font usage principles, the three workers unanimously selected a statement or left it unchecked 65% and 52% of the time respectively. The details text box was also generally used to justify their decision for checking a statement and offer suggestions to a problem. A possible explanation for the fair agreement rate amongst workers is that the statements presented to workers in this pilot were inconsistent in phrasing. Some lacked a clear assertion, which made the meaning of checking the statement ambiguous. We believe that a set of clearer, more consistent design principle statements can yield higher agreement rates amongst workers.

While novices may have been bad at incorporating design principles to rationalize their critiques, this pilot demonstrates that they are able to identify visual design problems and generate solutions. An interface where workers choose from structured guidelines for their critique allows those who lack knowledge of design-specific vocabulary to provide relevant, useful feedback. Structured output is also easier to aggregate for the designer receiving the feedback to interpret.

The CrowdCritter System

From our pilots, we determined that the open-ended critique format analogous to "crits" conducted in person for design students would not work in an online crowdsourced context. Though workers in crowdsourced labor markets can recognize visual design problems, they lack specialized design knowledge to express their feedback in a useful, actionable way. To make the complex creative task of design critique feasible for crowdworkers, we propose the following model: (1) critique is organized around a set of design principles exposed in the user (2) workers select relevant statements about design principles and provide comments within the scope of

selected principles, (3) multiple workers' responses are aggregated to get full coverage of critiques for a design.

Constructing a set of structured design principles
To test our hypothesis that novices would be able to provide useful critiques when visual design theory is made salient to them in the form of structured design principles, we first compiled a holistic set of design statements. We conducted a comprehensive survey of design books and other design-related resources [24, 25, 26] to get a list of universally applicable visual design principles, and grouped them into overarching principles (Appendix A). The set encompasses a wide range of design principles, from low-level issues like spelling errors and font usage to higher-level concepts like visual flow and appropriateness for the intended audience.

All of the statements are phrased in the same three-part format shown in Figure 6, intended to encompass all three components of Sadler's guidelines for good feedback criteria—specific, conceptual, actionable [12]. This guarantees that if a worker selects a relevant critique statement, it will be presented to the original designer in a useful, actionable way. Because traditional critiques involve both positive and negative statements [9], we included both general positive statements, and more specific negative critiques in our pool.

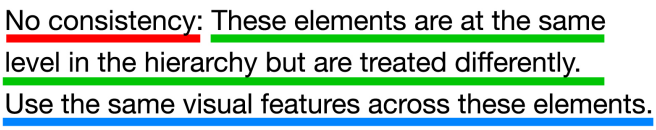

Figure 6: All negative critiques from our set of compiled design principles are formatted as shown on the left with 3 parts: (1 – red) a summarization of the problem, (2 – green) an explanation of why this is a problem using visual design theory, (3 – blue) a way to address the issue

A fixed set of critique statements also makes aggregation of critiques between multiple workers straightforward since individual critiques do not have to be coded. For example, we can easily extrapolate glaring problems with a design by analyzing how many workers selected the same statement for a given design, or the breadth of problems in a design by examining how many larger principles were covered by worker responses.

Interface workflow
We hypothesized that selecting relevant statements from a list of design principles would enable novices to provide useful feedback, but it is unclear what type of tool or workflow would be best for supporting crowdsourced design critique. Drawing on Sadler's theory on feedback-providing techniques [12] and Schwartz' tell-and-practice versus guided discovery models of teaching new concepts [3], we developed two different workflows with different feedback-giving methodologies. These interfaces contain the same core set of features:

- Design: The visual design to be critiqued is displayed on the left.
- Design context: The design's title and some context about the purpose of the design are displayed above the design. It is important to communicate a design's intended audience and message to critiquers so they can consider whether the design choices made are appropriate—knowing the context helps them make better-informed critiques.
- Annotation tools: In both interfaces, workers visually indicate problem areas or aspects that work well on the design using a set of annotation tools. Users can place markers on the design, draw boxes or polygons, and circle the whole slide. Existing communities for

design feedback offer these similar marking tools for their users to comment on designs [15]. For the worker, visually indicating areas is easier than describing the areas with words. For the designer reviewing the feedback, interpreting where notable regions are visually is also easier than having to infer from reading a textual description.

- Comment provision: In both interfaces, critiquers associate any annotation they make to a textual description of what they want to say about the region they have marked.

We hypothesized that selecting relevant statements from a list of design principles would enable novices to provide useful feedback, but an open research question remains about how structured principles are best integrated into the critiquing workflow to yield the best feedback from novice critiquers. Sadler argues that techniques for making complex judgments fall into two broad categories: the analytic approach, and the configurational approach [12]. The analytic approach involves identifying relevant criteria and measuring the amount present of each criterion in the object being judged. One makes an overall judgment by breaking down the multifaceted task into separate criteria and following explicit rules. On the other hand, in the configurational approach, the evaluator reacts to the work as a whole, making an entire assessment first and substantiating it by referencing separate criteria.

Another way to frame the question is to ask what kind of critiquing workflow will best teach novice critiquers to learn how to provide feedback the best. In this vein, Schwartz discusses two different approaches to teaching in school [3]. The first, called "teach and practice," involves the student first being shown an explanation of the concepts, then practicing on a set of well-designed problems. Alternatively, the guided discovery model involves withholding didactic teaching to allow the student to explore the problems themselves. This discovery phase allows students to render the knowledge they later learn useful [34]. The two CrowdCritter interfaces we built reflect these two approaches to teaching as well as Sadler's methods for providing feedback.

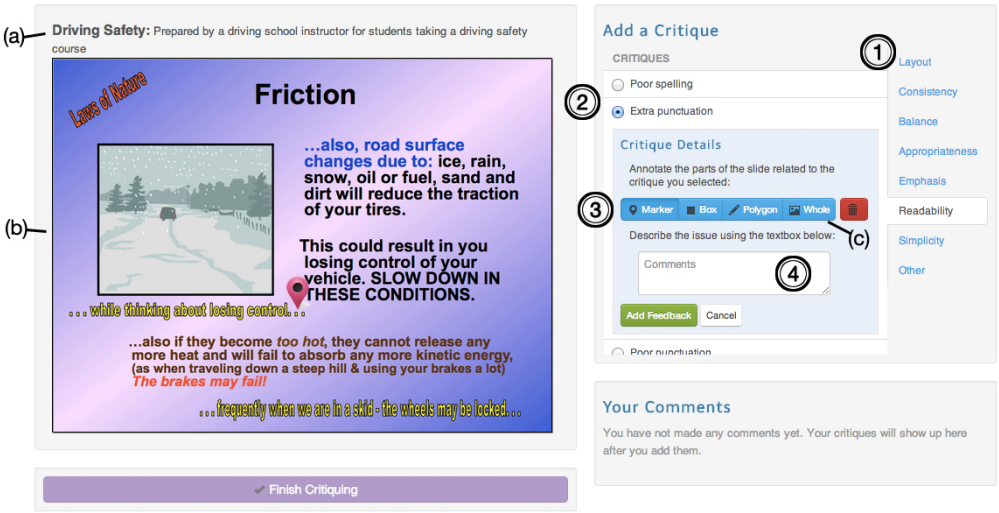

Figure 7 (Critique-first interface): (a) design context, (b) design, (c) annotation toolbar. To make a critique, the user first (1) browses through a list of design principles organized by theme, (2) selects a relevant statement, (3) makes an annotation for the part of the design relevant, (4) types a freeform comment to further clarify the critique

Critique-first (CF) interface—the analytic approach

In the critique-first interface, users are shown a list of design principles, organized by category. To make a critique, they first select the statement they would like to make. Then, they are prompted to use the annotation tools to visually indicate where on the slide the statement is relevant. They can also type their own comment to associate with the statement they selected. They repeat this process until they have made all the critique statements they wish. This interface takes an analytic approach to giving feedback because it prompts user to identify relevant criteria first in order to make an overall assessment of the design. It also applies the "teach and practice" method of teaching novice critiquers because it shows them the principles upfront.

Freeform-categorize (FC) interface—the configurational approach

In the freeform-categorize interface, users are shown the annotation tools and a textbox. They are asked to type a critique idea and indicate on the design where it is relevant. They repeat this process until they have made all the critique statements they wish. After they submit, they are taken to a second page where they are asked to categorize the freeform statements they made one-by-one. Users are prompted to check all the relevant statements from a list of structured design principles that best encompass each of the freeform comments they made on the previous page. This interface elicits feedback from users in a configurational manner—they first give an overall impression of a design, then select which critique statements are salient to their initial appraisal. From a teaching perspective, this reflects guided discovery because it asks users to identify problems themselves first without being shown the principles upfront.

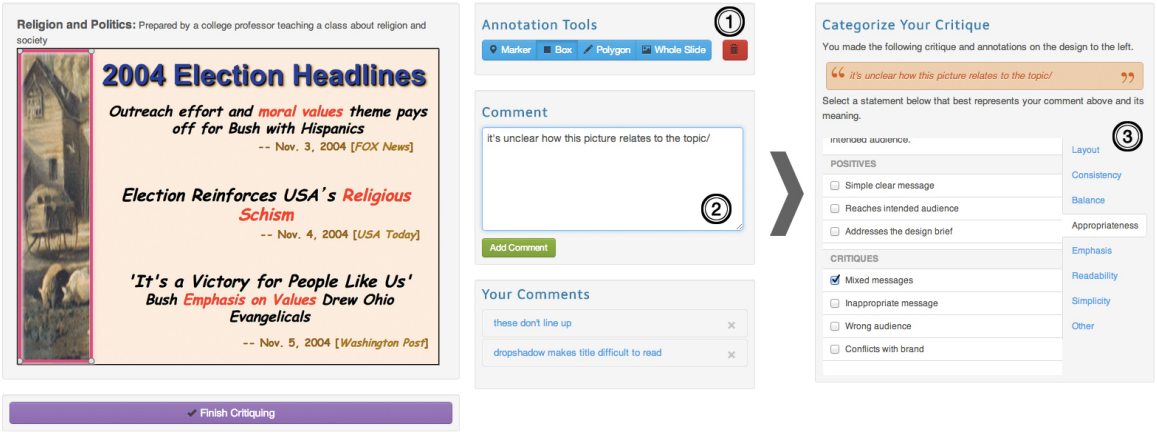

Figure 8 (Freeform-categorize interface): In the freeform-categorize interface, users first (1) annotate a part of the design they would like to address, and (2) type in a freeform comment to express that problem. They repeat this until they have expressed all their thoughts about the design. Afterwards, (3) they are asked to select statements from a list of design principles organized by theme that best characterize each of the freeform comment they made.

Implementation

CrowdCritter is built as a web application with the Django web framework [32], which facilitates the saving and retrieval or worker critiques. All frontend user interactions including the annotation tools were implemented with jQuery [33] and the KineticJS library [34], which facilitates building rich interactions on top of the HTML5 canvas.

Evaluative Study
We devised a study to evaluate how the two interfaces we built would allow both expert and novice workers to critique visual design. For this study, we were interested in measuring the critiques made by the two groups using these interfaces along four different dimensions:

- Quality: Did certain interfaces levels lead to better quality critiques?
- Quantity: Did certain interfaces lead to more critiques for a given design?
- Breadth: Did certain interfaces lead to wider variety of critiques for a given design?
- Agreement: Did certain interfaces lead to greater agreement about the issues that exist in a design among multiple raters?

Study Design
We conducted a between-subjects 3x1 study where participants used one of three interfaces to carry out visual design critique tasks. The conditions were (1) critique-first (CF), (2) freeform-categorize (FC), and (3) freeform (FR). The freeform interface was a baseline condition representative of the community critique interfaces that existing sites like Marqueed [15] use today. Users made annotations on the design of issues or strong points, and wrote freeform comments to accompany the annotations. They were never exposed to the structured list of principles that the other interfaces involve. This interface was intended to clarify the differences in the types of responses that an approach using structured design principles would yield.

Participants were randomly assigned to use one of the three interfaces, and asked to critique three visual designs in a row. The first of the three designs was randomly drawn from a large pool of 51 slides from presentations and 33 posters scraped from Google searches. The second and third designs participants critiqued were randomly drawn from a smaller pool of 3 slides and 3 posters. The larger pool of designs covers more breadth in terms of quality and problems that can exist in a visual design. Critiques done on designs in the larger pool provide qualitative insight on the coverage of the set of design principles that we compiled and how the annotation tools are used. We perform our primary summative analysis measuring the dimensions mentioned above with participants' responses on the designs from the smaller pool. We chose to have participants critique a design in the larger pool first to eliminate possible learning effects from using the interface for the first time.

Participants were recruited via email online, and conducted the task remotely on their own computers. They were first asked to fill out a survey about their design background and demographics. After the survey, participants were shown a video introducing them to the interface they would use to do the design critique tasks. Then, they critiqued three designs using that interface in the order described above. Finally, they completed a post-task survey that asked them about their experience using the interface. Subjects were paid a $10 Amazon gift card at the end of the study.

Results
The study was conducted with 73 participants total. Out of the 73 participants, 37 reported having zero design experience. In our reporting of results below, we label to the zero design experience participants as novices and everyone else as experts. We refer to a specific comment-

annotation-design-principle (if any) cluster made by a user as an "individual critique," and all individual critiques a user made for a given design as a "submission."

Quality

To measure quality, we first devised a rubric for rating individual critiques (see Appendix B). A piece of critique is defined as a combination of (1) an annotation, (2) a user-provided freeform comment, if any, and (3) a selected principle, and is rated holistically on a 3-point scale where 1 is poor, 2 is fair, and 3 is good. Ratings were divided between three experts, who reached consensus in their grading through several rounds of piloting with the rubric. To minimize any implicit bias and ensure that the judges would be blind to the condition in which a piece of critique was provided, we built a grading interface that only displays a singular principle, freeform comment, and annotation to the grader. For freeform interface condition responses that are not associated with any structured principles, the interface would display an "Other" principle with the user's comment. This was a viable disguise because users in the critique-first and freeform-categorize conditions could also conceivably associate their freeform comment with the "Other" principle.

	Novice	Expert	All users
Critique first	2.321 (0.692)	2.505 (0.626)	2.430 (0.661)
Freeform-categorize	2.331 (0.645)	2.345 (0.653)	2.339 (0.649)
Freeform	2.442 (0.586)	2.471 (0.571)	2.458 (0.578)

Table 1: Average quality ratings given to novice, expert, and all participants' critiques in the three conditions across all designs in the small set. Std. dev. is shown in parenthesis.

Table 1 shows an overview of the average rating for critiques in each condition provided by novices, experts, and all users. For novices and across all users, critiques in the freeform condition were rated highest quality, critiques in the critique-first condition the second highest, and critiques in the freeform-categorize condition the lowest. Separate Kruskal-Wallis tests were performed for novices and experts, with interface as a factor and critique quality as the dependent variable. No significant difference in quality rating was found across conditions for novices ($\chi^2(2)=0.827$, $p=0.662$). For experts, the Kruskal-Wallis test revealed a significant effect of interface on quality rating ($\chi^2(2)=10.06$, $p < 0.01$). A post-hoc test using Mann-Whitney tests with Bonferroni correction showed significant differences in quality between the critique-first and the freeform conditions ($p < 0.01$), but differences between other pairs of conditions were not significant.

Quantity

To measure quantity, we looked at the number of individual units of critique users made per design. In total, there were 603 responses provided by the 73 participants for the 6 designs in the small pool. Figure 9 shows the average number of critiques in each condition provided by novices, experts, and all users across all designs. An ANOVA was performed with interface and design experience as factors, and number of critiques as the dependent variable, revealing significance for both design experience and interface ($F(1, 140)=13.18$, $p<0.001$, $F(2, 140)=5.822$, $p<0.01$). We performed separate post-hoc Tukey HSD tests for novices and experts. Novices produced significantly more individual critiques using the critique-first and freeform-categorize interfaces compared to the freeform interface ($p=0.05$ for both). However, the

difference between the number of critiques they made using CF and FC was not significant. Experts made significantly more critiques in the critique-first condition over the freeform-categorize and freeform conditions (p=0.05, p=0.1), but the difference in number of critiques between the latter two conditions was not significant.

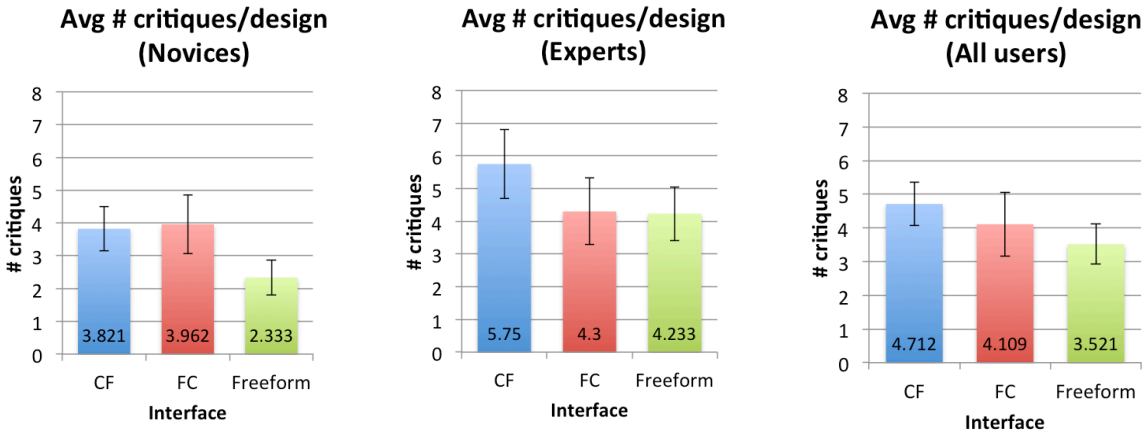

Figure 9: Average number of critiques per design with 95% confidence intervals for each condition made by novices, experts, and all users.

Breadth

To look at whether the critique-first or freeform-categorize interfaces would yield a wider variety of critiques, we measured the number of unique structured design principles statements used in a participant's submission. This captures the breadth of visual design-related ideas that a user employed to critique the design. We also examined how many unique design principle categories were encompassed per submission across all designs. We analyzed only the responses from the critique-first and freeform-categorize conditions, since there was no data about the design statements from the freeform condition.

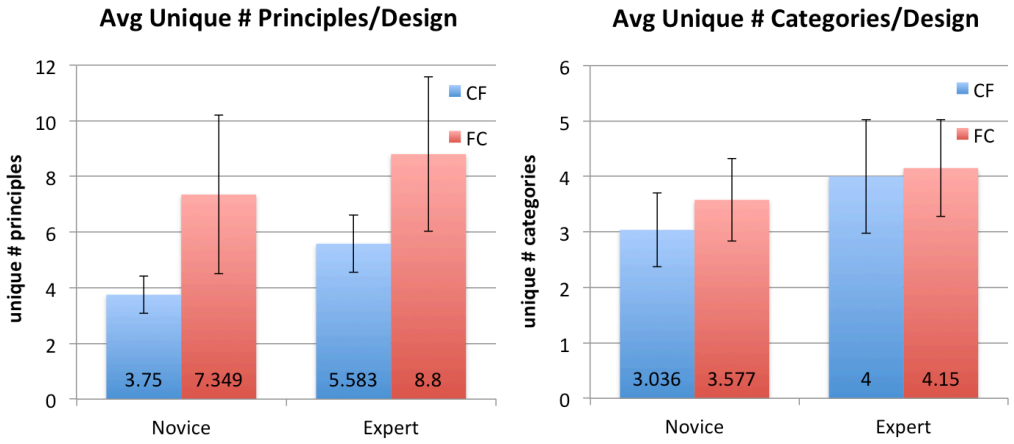

Figure 10: Average number of unique principles and higher level principle categories encompassed by participants' submissions across all designs in the smaller set. 95% conf. intervals are shown.

Figure 10 shows the average number of unique principles and categories identified by novices and experts in each condition. In general, we note that both novices and experts capture a larger breadth of principles and categories using the freeform-categorize interface over the critique-first interface. We performed t-tests to confirm the significance of these differences (p=0.05 in both cases). Novices selected almost two times more unique principles in FC than in CF, while experts selected over 1.5 times more. However, the difference in breadth at the category level is not statistically significant.

Agreement

To examine the degree of agreement between participants critiquing a given design using an interface, we perform Fleiss' kappa analysis on all six designs in the smaller pool for participants in the critique-first and freeform-categorize conditions (Table 2, 3) at the design category level. Appendix C displays the six designs that are referenced in the tables. We constructed a matrix M where M[i, j] = {*1 if user i selected a design principle in category j, 0 otherwise*}. A negative kappa is considered poor agreement, while a kappa value between 0.01-0.20 is considered to be slight agreement. Overall, we see that for most designs, all users are only able to reach slight agreement, if any. The results are more haphazard when we break down agreement by design experience, but this may be due to data scarcity—conditions were not balanced for design experience so some designs were rated by many experts and few novices, and vice versa.

CF	Design #1	Design #2	Design #3	Design #4	Design #5	Design #6
Novice	0.022	-0.094	0.258	0.111	0.179	-0.067
Expert	-0.067	-0.067	N/A	0.080	0.111	-0.101
All users	0.054	-0.042	0.222	0.157	0.198	-0.048

Table 2: Fleiss' kappa of for six designs amongst novices, experts, and all CF condition users at the category level. Cells with N/A mean that only there was only one user of that design experience level who critiqued the design, so kappa value cannot be calculated.

FC	Design #1	Design #2	Design #3	Design #4	Design #5	Design #6
Novice	-0.026	0	0.133	0.04	-0.129	-0.270
Expert	-0.067	-0.185	0.123	0	0.167	0.314
All users	0.044	-0.030	0.193	0.129	0.119	-0.053

Table 3: Fleiss' kappa for six designs amongst novices, experts, and all FC condition users at the category level

Discussion

Structured design principles enabled novices to produce more critiques.

Novices using the critique-first and freeform-categorize interfaces produced significantly more critiques per design on average than novices using the freeform interface. A major difference between the two interfaces and the freeform interface is the incorporation of structured design principles in the former. This suggests that the presence of the design statements guided novices to consider aspects of the design that they would not have considered otherwise. The presence of design principles can serve as a checklist for novices who, being unfamiliar with visual design concepts, have difficulty recalling design theory on their own. One user who used the critique-first interface noted in the post-survey, "I liked that the interface was somewhat guided…This was really useful because I am not familiar with the jargon for design critique and made it easier to find the right language to express my thoughts on the designs."

Interestingly, the freeform-categorize interface, which does not reveal the list of structured design principles to the user until after they have submitted their ideas, still produced more critiques than the freeform interface. This suggests that seeing the principles after critiquing one design can carry over when a user critiques a subsequent design. A participant who used the freeform-categorize condition wrote in the post-survey comments, "The options to categorize each comment helped remind me of other elements I should look at on top of the ones I already noticed." Seeing the principles reminds users what they should consider when critiquing subsequent designs. This may explain why freeform-categorize also led to a higher average number of critiques per design than the freeform interface.

Freeform-categorize yielded critiques encompassing a wider breadth than critique-first.
The freeform-categorize interface (configurational approach) covered much more breadth in terms of unique number of principles covered than the critique-first interface (analytic approach). This can be explained by the design of the interfaces themselves—in critique first, users can only associate one design principle with any individual critique they make, while in freeform-categorize, users can pick any number of design principles to associate with an individual critique. Because the average number of individual critiques made by users in each condition was about the same, we can conclude that the inherent design of the interfaces lead to this difference in breadth.

Results from measures of critique quality from the three conditions are inconclusive.
The critique quality results for the three conditions were surprising and opposite of our original hypothesis. Based on our pilot studies, we expected critiques from the interfaces that use structured principles to be higher quality than those from the freeform interface because we thought that the associating principles with participants' freeform thoughts. We thought this would be especially true for novices, because the structured statements would help clarify their ideas. However, experts rated novice critiques from the freeform interface the highest on average instead. The differences were not statistically significant, so we cannot discount that the difference was due to chance.

One possible explanation is that our grading scheme of rating a critique on a 1-2-3 scale was not fine-grained enough to differentiate the quality of the critiques. The scale was designed for multiple graders to reach a clear consensus about the category (1, 2, or 3) of a piece of critique, but these buckets may have been too broad to capture the spectrum of user responses across the conditions. Future work will entail a refinement of the rating process that will better separate the range of quality.

It is difficult to reach agreement about visual design problems, even amongst experts.
Our Fleiss-kappa analysis on the level of agreement was largely inclusive, which in part demonstrates the complexity of visual design critique. During the grading process, the raters noticed a reoccurrence of the same problems pointed out by different users for a given design. However, there are many valid structured principles one can choose to express a particular idea. Though participants may have agreed on the same conceptual problems, the structured principles they associated with their critiques may vary. An example for such ambiguity in the critiquing process is shown in Figure 11. Because of this flexibility in selecting principles to associate with

problems, we may not be capturing actual agreement very well if we only compare the design concept categories that participants' critiques encompassed.

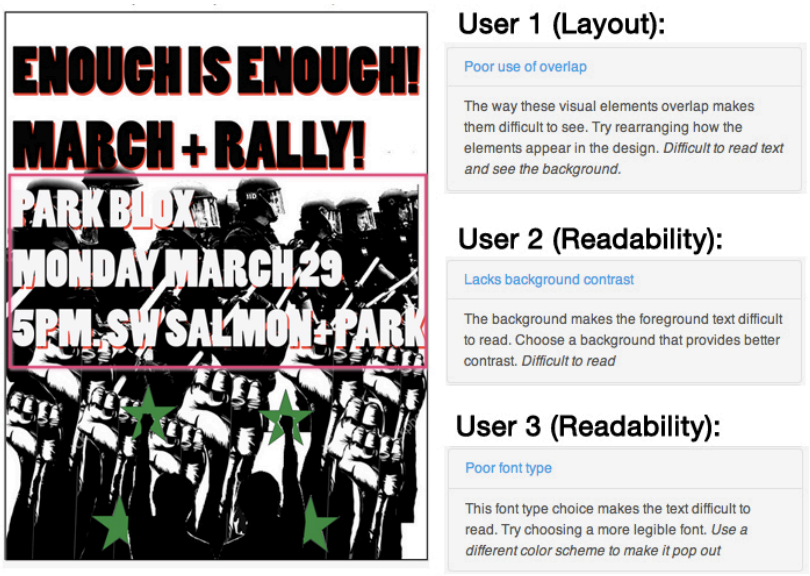

Figure 11: Examples of critiques made by different users that all express the same problem of illegibility of the white text annotated by the pink box. The design principles these users selected are all applicable, but belong in different higher-level categories.

Another plausible explanation is that users only critique a subset of all problems existent in a design based on how much effort they exert to complete the task. Different participants' critiques may include only certain higher-level categories and all of these critiques may be equally valid. However, these critiques will not appear to agree with each other under our current method of analyzing Fleiss' kappa. It is clear that more sophisticated methods are necessary to more accurately measure participants' level of agreement within a condition. These methods should consider the annotation region of individual critiques, as well as the underlying idea of a critique.

Conclusions and Future Work

Our findings did not provide valuable insights about whether the "teach or practice" or guided discovery model would be more effective for teaching novices the complex task of visual design critique. However, the evaluative study did reveal that incorporating structured design principles into the critiquing workflow is a promising approach to eliciting critique from people with no design experience. At the very least, it enables workers to provide more critiques for a given design. One direction for future work is to empirically evaluate the usefulness of critiques provided through the CrowdCritter system with designers actually seeking to improve their work. While measures of quality, quantity, breadth, and agreement of critiques can provide insight on how effectively the system enables workers to provide critiques, the raw numbers are difficult to compare and place in a practical context. Ultimately, it is important to establish how the aggregated feedback from these systems impacts the kinds of changes that designers make to their work.

Another area for further exploration is different approaches to quickly teaching those with no design experience the skills required to critique visual design. CrowdCritter provides very little training to workers about design concepts—it relies on their implicit knowledge of design-related terms and simply helps them recall this knowledge through the presentation of structured design principles. Low-level design concepts like clashing colors or overlapping elements may be

universally understood, but higher-level ideas like appropriateness and balance may require training. Follow-up studies may explore just-in-time learning models to instruct novices on a particular concept of visual design before they are asked to critique a design along that concept.

References

[1] Frei, B. (2009) Paid Crowdsourcing: Current State & Progress toward Mainstream Business Use. *Smartsheet. http://bit.ly/smartsheet_report* (Accessed March 15, 2011)

[2] Schwartz, D. L., Bransford, J. D., & Sears, D. (2005). Efficiency and innovation in transfer. In J. P. Mestre (Ed.) Transfer o flearning from a modern multidisciplinary perspective (pp. 1-52). Greenwich, CT:Information Age Publishing.

[3] Schwartz, Daniel L., Catherine C. Chase, Marily A. Oppezzo, and Doris B. Chin. "Practicing Versus Inventing with Contrasting Cases: The Effects of Telling First on Learning and Transfer." Journal of Educational Psychology 103, no. 4 (2011): 759–775. doi:10.1037/a0025140.

[4] Kittur, Aniket, Boris Smus, Susheel Khamkar, and Robert E. Kraut. "CrowdForge: Crowdsourcing Complex Work." In Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology, 43–52. UIST '11. New York, NY, USA: ACM, 2011.

[5] Bernstein, Michael S., Greg Little, Robert C. Miller, Björn Hartmann, Mark S. Ackerman, David R. Karger, David Crowell, and Katrina Panovich. "Soylent: a Word Processor with a Crowd Inside." In Proceedings of the 23nd Annual ACM Symposium on User Interface Software and Technology, 313–322. UIST '10. New York, NY, USA: ACM, 2010.

[6] Noronha, Jon, Eric Hysen, Haoqi Zhang, and Krzysztof Z. Gajos. "Platemate: Crowdsourcing Nutritional Analysis from Food Photographs." In Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology, 1–12. UIST '11. New York, NY, USA: ACM, 2011.

[7] Kokkalis, Nicolas, Thomas Köhn, Carl Pfeiffer, Dima Chornyi, Michael S. Bernstein, and Scott R. Klemmer. "EmailValet: Managing Email Overload through Private, Accountable Crowdsourcing." In Proceedings of the 2013 Conference on Computer Supported Cooperative Work, 1291–1300. CSCW '13. New York, NY, USA: ACM, 2013. doi:10.1145/2441776.2441922.

[8] Barrett, Terry. "A Consideration of Criticism." Journal of Aesthetic Education 23, no. 4 (December 1, 1989): 23–35.

[9] Feldman, E. B. Practical Art criticism. Englewood Cliffs, NJ: Prentice Hall. 1994

[10] Dutton, T.A. Design and Studio Pedagogy. *Journal of Architectural Education, 41,* 1, (1987), 16-25

[11] Teo, Y. H., and Chai, C.S. Scaffolding Online Collaborative Critiquing for Educational Video Production. Knowledge Management & E-Learning: An International Journal (KM&EL) 2009.

[12] Sadler, D. Royce. "Formative Assessment and the Design of Instructional Systems." Instructional Science 18, no. 2 (June 1, 1989): 119–144.

[13] Xu, Anbang, and Brian P. Bailey. “A Crowdsourcing Model for Receiving Design Critique.” In CHI ’11 Extended Abstracts on Human Factors in Computing Systems, 1183–1188. CHI EA ’11. New York, NY, USA: ACM, 2011.

[14] Trumbo, Jean. “The Process of Critique in Visual Communication.” Journalism & Mass Communication Educator 52, no. 2 (June 1, 1997): 15–23.

[15] http://marqueed.com

[16] http://dribbble.com

[18] http://mturk.com

[19] http://crowdflower.com

[20] http://www.microtask.com

[21] http://www.fivesecondtest.com

[22] https://usabilla.com

[23] http://feedbackarmy.com

[24] Reynolds, Garr. Presentation Zen: Simple Ideas on Presentation Design and Delivery. 1st ed. New Riders, 2008.

[25] Williams, Robin. The Non-Designer’s Design Book: Design and Typographic Principles for the Visual Novice. 1st ed. Pearson Education, 1995.

[26] Lidwell, William, Kritina Holden, and Jill Butler. Universal Principles of Design, Revised and Updated: 125 Ways to Enhance Usability, Influence Perception, Increase Appeal, Make Better Design Decisions, and Teach through Design. Second Edition, Revised and Updated. Rockport Publishers, 2010.

[27] Dannels, D.P. and Martin, K.N. Critiquing Critiques: A Genre Analysis of Feedback Across Novice to Expert Design Studios. Journal of Business and Technical Communication 22, 2 (2008), 135-159.

[28] Xu, Anbang, and Brian Bailey. “What Do You Think?: a Case Study of Benefit, Expectation, and Interaction in a Large Online Critique Community.” In Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, 295–304. CSCW ’12. New York, NY, USA: ACM, 2012. doi:10.1145/2145204.2145252.

[29] Barrett, Terry. “A Comparison of The Goals of Studio Professors Conducting Critiques and Art Education Goals for Teaching Criticism.” Studies in Art Education 30, no. 1 (October 1, 1988): 22–27.

[30] Dannels, Deanna. Beyond Content, Deeper than Delivery: What Critique Feedback Reveals about Communication Expectations in Design Education. International Journal for the Scholarship of Teaching and Learning 2, 2 (2008).

[31] Bransford, J. D., Franks, J. J., Vye, N. J. & Sherwood, R. D. (1989). New approaches to instruction: Because wisdom can't be told. In S. Vosniadou & A. Ortony (Eds.), Similarity and analogical reasoning (pp. 470-497). New York, NY: Cambridge University Press.

[32] https://www.djangoproject.com

[33] http://jquery.com

[34] http://kineticjs.com

Appendix A – List of Design Principle Statements

Readability: The text on the design should be free of errors. The text's visual treatment and formatting should make it easy to read and respect rules of typography.

Positives:

- Spelling error free: The text is understandable and free of spelling/grammatical errors.
- Very readable: The paragraph formatting supports readability.
- Visually legible. The selection of visual features helps to make the text legible.

Critiques

- Poor spelling: There are spelling errors here. Double-check the spelling.
- Extra punctuation: These punctuation marks are unnecessary. Remove them.
- Poor punctuation: This is difficult to parse. Include appropriate punctuation.
- Poor grammar: This phrasing here is confusing. Work on improving the grammar.
- Orphan: This single word should not have its own line. Shorten the line above or choose a slightly different font.
- Widow: Do not start a new text block with a partial sentence or phrase. Shorten the text block above or use a slightly different type size.
- Poor line length: The length of these lines makes it difficult to read. Format the paragraph to a more appropriate length (50-75 characters per line).
- Poor line spacing: The spacing between these lines of text looks awkward. Change the line spacing to about 120% of the type size.
- Poor kerning: The space between characters looks awkward. Choose more appropriate character spacing.
- Poor character width: This character width feels awkward. Choose a more appropriate character width.
- Poor paragraph spacing: The spacing between these paragraphs looks awkward. Change the spacing to a more appropriate distance between paragraphs.
- Justify text: These lines of text create a ragged appearance. Try using justified paragraphs to get a uniform, rectangular visual appearance.
- Remove justification: This paragraph justification introduces awkward word spacing. Try removing the justification to improve readability.
- Lacks background contrast: The background makes the foreground text difficult to read. Choose a background that provides better contrast.
- Poor color scheme: This choice of colors makes it difficult to read. Try choosing colors with better contrast.
- Poor font type: This font type choice makes the text difficult to read. Try choosing a more legible font.
- Poor font appearance: The use of font modifications makes the text difficult to read. Try using different visual features to improve legibility.

Layout: Every element should have some connection to another element on the page. Nothing should be placed on the page arbitrarily.

Positives

- Good alignment: The effective use of a grid structure helps create a clear and organized design.
- Well organized: The alignment of visual elements create a well organized design.

Critiques

- Wrong size: The size of this visual element makes it difficult to see. Try changing the size to make it more legible to the viewer.
- Overuse of rotation: The use of rotation makes this visual element difficult to see. Try aligning it along a horizontal and vertical layout.
- Poor scaling: Do not change the aspect ratio of this visual element if it distorts the content. Try cropping the visual element instead of scaling it.
- Poor cropping: Important information gets cropped from this visual element. Try re-cropping the visual element or select a replacement.
- Poor use of overlap: The way these visual elements overlap makes them difficult to see. Try rearranging how the elements appear in the design.
- Poor placement: These elements are not aligned in any organized way. Try aligning the elements along a common line to create a clear and organized look.
- Uneven margins: The uneven margins around the edges look odd. Try to balance the spacing around the border of the design.
- No margins: The lack of margins around the edges looks awkward. Try to leave a little space between elements and the edge of the design.
- Uneven gutters: The uneven spacing between these elements looks odd. Use similar amounts of space around the elements to frame the content.
- No gutters: The lack of space between these elements looks does not leave breathing room for the viewer's eye. Use space around the elements to help frame the content.

Simplicity: The design should use as few elements as possible to achieve its goals. Each visual element should contribute to the overall message; all non-essential elements and should be omitted.

Positives

- Simple and clean: The simplicity of this layout gives the design a clear and organized appearance.

Critiques

- Too much content: The design tries to include too much content. Try removing content or splitting the content between multiple designs.
- Overuse of images: An abundance of these images may overwhelm the viewer. Try removing content that does not help convey the primary message.
- Clashing elements: These elements clash. Try choosing more complementary elements.
- Complicated background: This background makes the design feel cluttered and overly complicated. Select a different background or remove it.
- Too much text: The abundance of text makes this difficult for viewers to comprehend. Consider condensing this text by focusing on the essential message.

- Not enough text: These are too abstract and short for the viewer to understand your argument. Consider adding critical detail.
- Lacks white space: This design feels cluttered. Try using more space—the absence of text and graphics—to provide visual breathing room for the eye.

Emphasis: The design should match the importance of content to its visual prominence - make the most important information visually dominant. Use clear contrast to distinguish different levels of information.

Positives

- Good visual emphasis: This design successfully gives visual prominence to the most important information.
- Good visual flow: This design provides a successful visual flow to guide viewers through the content.
- Strong focal point: This design successfully guides viewers through the content by using a strong focal point.
- Good visual contrast: The visual contrast between these elements helps create a strong point of emphasis.

Critiques

- Elements lack contrast: This element does not have enough contrast with other elements. Consider using different visual features to create emphasis on this element.
- Design lacks contrast: This design lacks visual emphasis. Consider using different visual features to draw the viewer's attention to the most important content.
- Overemphasis: These elements are overemphasized making the design appear busy. Consider using fewer visual features to more subtly draw a contrast between these elements.
- Lacks hierarchy: The design fails to create a visual flow through the content. Try emphasizing the most important element and then guide the viewer using different visual features.
- Poor hierarchy: This design creates to too many focal points. Try to create visual emphasis for only one focal point and then use different visual features to create a visual flow.
- Lacks proximity: These related elements seem to have no visual relationship to each other. Consider placing them together to better express their logical relationship.
- False proximity: These elements should not be visually grouped together. Try spacing elements in a way that suggests their logical relationship.

Consistency: Elements that occupy similar positions in the information hierarchy should be given similar graphic treatment. This develops the organization and promotes unity.

Positives

- Good repetition: These repeating visual features successfully create an organizational scheme that guides the design.
- Consistent design language: The design language is consistent and expresses a clear organization for viewers.

Critiques

- No consistency: These elements are at the same level in the hierarchy but are treated differently. Use the same visual features across these elements.
- Poor consistency: These elements do not need to be visually different. Consider making their visual appearance similar.
- Lacks repetition: The design lacks visual cohesiveness. Consider repeating this element to create visual interest and to give the design a more complete look.
- Lacks visual unity: The design fails to come together as a whole. Try squinting at the design for an overall feel and then strengthen or add repetitive elements to create unity.
- Element disrupts unity: This element takes away from the visual unity of the design. Try selecting elements that come together as a whole.

Balance: Use the positioning of elements relative to each other to deliberately achieve an active or restive appearance.

Positives:

- Good use of symmetry: This design successfully uses symmetry to create a sense of order and stability.
- Good use of asymmetry: This design successfully uses asymmetry to sense of movement and change.
- Good overall balance: The design is well balanced providing an overall sense of unity.

Critiques

- Lacks balance: This design feels unbalanced and awkward. Try aligning these elements to a common axis to create a sense of order and stability.
- Lacks movement: This design needs asymmetry to create a sense of movement and change. Try moving these elements to one side to create unbalance.

Appropriateness: The design should address the design brief and should appropriately communicate the content to its intended audience.

Positives:

- Simple clear message: The design successfully communicates a single clear message.
- Reaches intended audience: The design successfully reaches out to the intended audience.
- Addresses the design brief: The design successfully addresses the specified design brief.

Critiques

- Mixed messages: This design does not convey a clear message. Try consolidating your design to make a single unifying point.
- Inappropriate message: This choice of design elements is inappropriate for the overall message. Try choosing more suitable visual features to carry the intended message.
- Wrong audience: These design elements are not appropriate for the intended audience. Try choosing visual features that will resonate better the target audience.
- Conflicts with brand: These visual features conflict with the overall identity of the brand. Try choosing visual features that better represent the brand.

Appendix B – Critique Quality Rubric

Poor (1)	Fair (2)	Good (3)
(No effort) Annotation, comment, or critique statement does not indicate effort or does not exist Any: • Spam • no effort • appears random • single word answers • no answer	(Somewhat acceptable, significant problems) Annotation, comment, or critique statement shows consideration but significant problems exist Any: • opinion statements without justification • not reasonably justified by design principle or common knowledge • not specific • indicates misunderstanding • relation between comment and critique not clear	(Acceptable) Annotation, comment, or critique statement reasonably adheres to "specific, reasonable" All: • arguably supported by design principle or common knowledge • reasonably specific • does not indicate misunderstanding

Appendix C – Designs in Evaluative Study Small Pool

Design #1

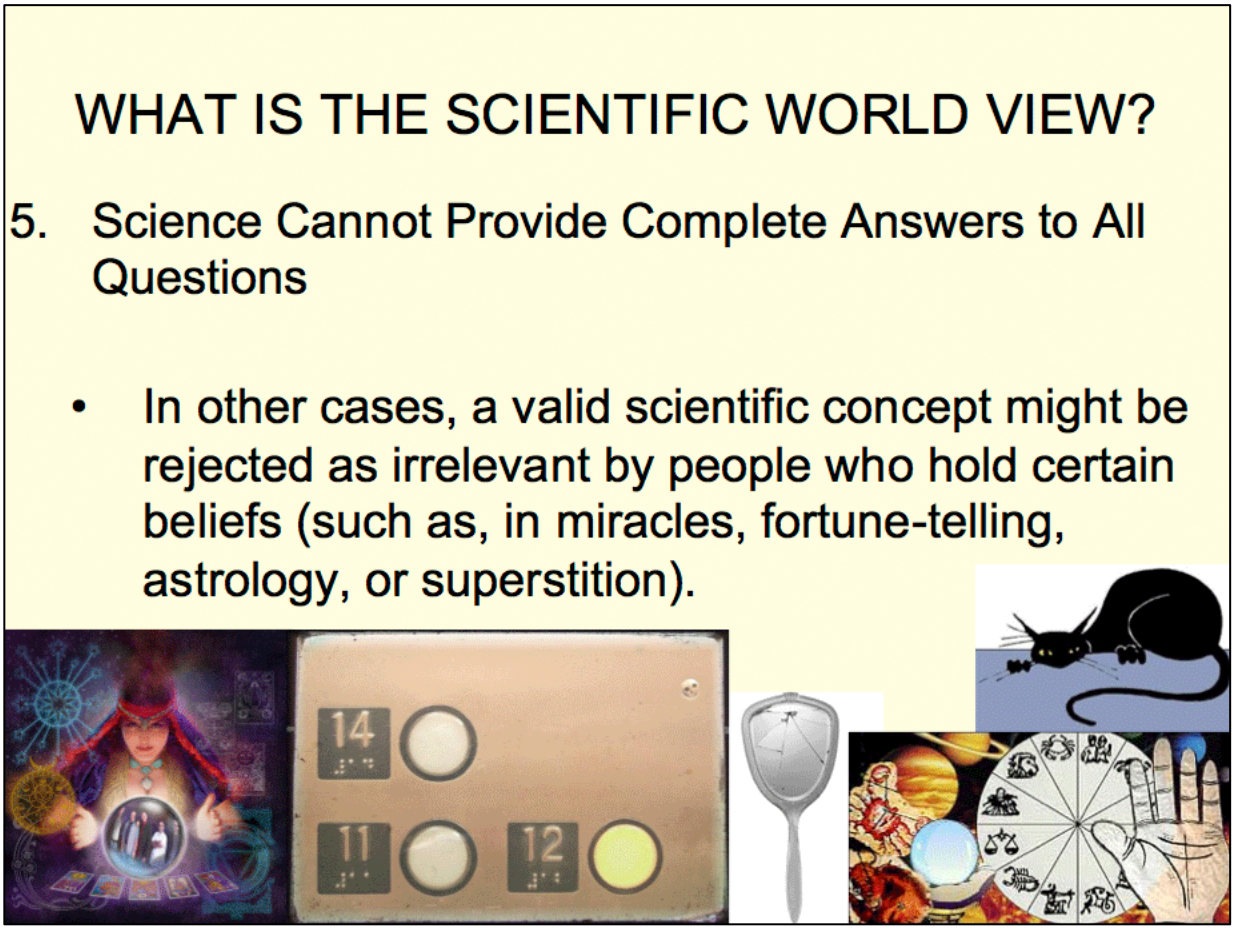

Design #2

Laws of Nature

Checking Tire Pressure

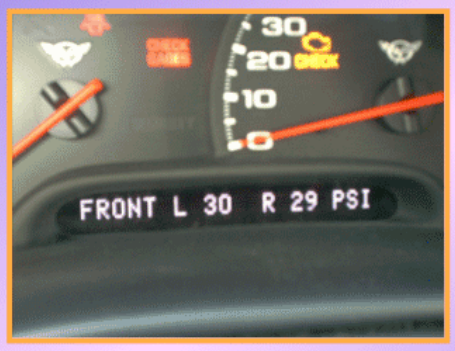

Benefits of Proper Tire Inflation:

- optimal tire life (even wear of tires)
- better gas mileage
- more comfortable ride

- It is also a very convenient way to check your tire pressure.

Let's finally consider forces in a crash ...

Design #3

Akara

(Blackeyed Pea Cakes)

2 cups dried blackeyed peas
1 medium onion
1/2 tsp. red pepper or Tabasco to taste
1 egg, beaten
1 tsp. salt
Oil to deep fry

Soak peas 30 minutes or longer, then rub briskly to remove skins (Some people don't remove the skins, but all Nigerian cooks would). Grind peas in blender with just enough water for the blender to operate smoothly. Grind onion and pepper. Add to the pea mixture with egg and seasoning. Mix thoroughly and drop by teaspoon into hot oil. Fry until golden. Remove from oil and drain on paper towels. Tastes best while hot. Makes an excellent hors d'oevres or side with meal.

Submitted by Lou Daniel

http://www.wfeca.net/clfonline/May04/recipe.html

Suya

http://photocom.gozaru.jp/gallery2004/pages/011suya.html

Some other traditional favorite foods mentioned in the book include Akara, Yams and Suya.

Yams

http://food.oregonstate.edu/images/fruitveg/yam/

Design #4


ENOUGH IS ENOUGH!
MARCH + RALLY!
PARK BLOX
MONDAY MARCH 29
5PM. SW SALMON + PARK

Design #5

BCEC CALLOUT

MAKE MOVIES

PLAN CONCERTS

HOST EVENTS

Purdue Business Careers in Entertainment Club invites you to come to our fall callout Tuesday October 2 at 7:00 pm in BCHM room 105. We will have plenty of fun waiting for you!

Design #6

SAE® Callout

Wednesday Aug 31, 2011 at 6PM
Location ME 1052

- All Majors Accepted
- No Previous Knowledge Required
- Free Pizza

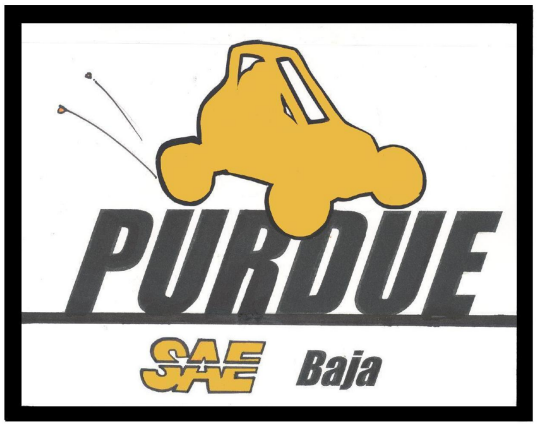

DESIGN, BUILD, RACE