

Parallel Prototyping Leads to Better Design Results, More Divergence, and Increased Self-Efficacy

STEVEN P. DOW, ALANA GLASSCO, JONATHAN KASS, MELISSA SCHWARZ, DANIEL L. SCHWARTZ, and SCOTT R. KLEMMER

Stanford University

Iteration can help people improve ideas. It can also give rise to fixation, continuously refining one option without considering others. Does creating and receiving feedback on multiple prototypes in parallel, as opposed to serially, affect learning, self-efficacy, and design exploration? An experiment manipulated whether independent novice designers created graphic Web advertisements in parallel or in series. Serial participants received descriptive critique directly after each prototype. Parallel participants created multiple prototypes before receiving feedback. As measured by click-through data and expert ratings, ads created in the Parallel condition significantly outperformed those from the Serial condition. Moreover, independent raters found Parallel prototypes to be more diverse. Parallel participants also reported a larger increase in task-specific self-confidence. This article outlines a theoretical foundation for why parallel prototyping produces better design results and discusses the implications for design education.

Categories and Subject Descriptors: H.1.m. [Information Systems]: Models and Principles

General Terms: Experimentation, Design

Additional Key Words and Phrases: Prototyping, iteration, feedback, juxtaposition, comparison, design, divergence, exploration, critique, self-efficacy

ACM Reference Format:

Dow, S. P., Glassco, A., Kass, J., Schwarz, M., Schwartz, D. L., and Klemmer, S. R. 2010. Parallel prototyping leads to better design results, more divergence, and increased self-efficacy. *ACM Trans. Comput.-Hum. Interact.* 17, 4, Article 18 (December 2010), 24 pages.

DOI = 10.1145/1879831.1879836 <http://doi.acm.org/10.1145/1879831.1879836>

The Hasso Plattner Design Thinking Research Program financially supported this work. Intel Corporation donated computers.

Authors' addresses: S. P. Dow and S. R. Klemmer, Department of Computer Science; A. Glassco, J. Kass, M. Schwarz, and D. L. Schwartz, School of Education, Stanford University, Stanford, CA 94305; email: {spdow, alana11, jkass, schwarz1, danls, srk}@stanford.edu.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 1073-0516/2010/12-ART18 \$10.00
DOI 10.1145/1879831.1879836 <http://doi.acm.org/10.1145/1879831.1879836>

ACM Transactions on Computer-Human Interaction, Vol. 17, No. 4, Article 18, Publication date: December 2010.

1. INTRODUCTION

Iteration is central to learning and motivation in design [Dow et al. 2009; Hartmann et al. 2006; Schon 1995; Schrage 1999]. Yet, its primary virtue—incremental, situated feedback—can also blind designers to other alternatives, steering them to local, rather than global, optima [Buxton 2007; Dow et al. 2009]. To combat this, creating multiple alternatives in parallel may encourage people to more effectively discover unseen constraints and opportunities [Cross 2006], enumerate more diverse solutions [Buxton 2007], and obtain more authentic and diverse feedback from potential users [Tohidi et al. 2006]. While a parallel approach has potential benefits, it can take time away from refinement.

Effectively educating a more creative workforce requires understanding how and why design practices affect results. Towards that goal, this article investigates the relative merits of parallel and serial prototyping under time constraints. In a between-subjects experiment, thirty-three participants designed Web banner advertisements for a magazine. In both conditions, participants created five prototype ads and then a final ad. They received a descriptive critique on each prototype. Participants worked independently and were given equal time to create each prototype and read each critique; the structure of the process differed across conditions. In the *serial* condition, participants received feedback after creating each prototype. Participants in the *parallel* condition created three prototypes, received feedback on all three, then made two more prototypes, and received feedback again before creating a final ad design (see Figure 1).

The study measured design performance by running a MySpace.com advertising campaign with all participants' final ad creations and measuring click-through analytics. Independent experts also rated ad quality. To measure the diversity of each participant's ad creations, independent online raters judged pairwise similarity between each of the participants' six ad prototypes. A self-report assessment measured participants' pre- and post-task view of task-specific self-efficacy [Fredrickson 2001; Hall 2008] (see Appendix B). The study concluded with an open-ended interview (see Appendix C).

Parallel participants outperformed serial participants by all performance measures: click-through rates, time spent on the target client Web site, and ratings by the clients and ad professionals. Further, independent raters found that the diversity of each participant's prototypes was greater in the parallel condition. Parallel participants reported a significant gain in self-efficacy, a measure of task-oriented confidence. Serial participants did not. In post-task interviews, nearly half of serial participants reported negative reactions to critique of their prototypes, while no parallel participants reported this. About half the participants had prior graphic or ad design experience. Participants with prior experience outperformed novices.

The study found that a parallel prototyping approach yields better results, more divergent ideas, and that parallel prototypers react more positively to critique. The results could significantly impact both how people approach creative problems and how educators teach design.



Fig. 1. The experiment manipulates when participants receive feedback during a design process: in serial after each design (top) versus in parallel on three, then two (bottom).

2. THEORETICAL BENEFITS OF PARALLEL DESIGN

Research on human problem solving traditionally examines problems with an optimal solution and a single path to reach that solution [Newell 1972]. In design, problems and solutions co-evolve [Dorst and Cross 2001], constraints are often negotiable [Schon 1995], subproblems are interconnected [Goel and Pirolli 1992], and solutions are not right or wrong, only better or worse [Rittel and Webber 1973]. How and when to explore or refine solutions to open-ended problems remains an active debate in design research and education [Ball et al. 1997; Cross 2006; Nielsen and Faber 1996]. Without exploration, designers may choose a design concept too early and fail to identify a valuable direction [Cross 2004]. Without refinement, ideas may not reach their full potential [Ball and Ormerod 1995]. Navigating a design space may come easier as designers develop intuition, however even experts can exhibit fixation [Cross 2004] and groupthink behaviors [Janis 1982]. The architect Laseau posits an idealized conceptual model for exploring and refining, where designers iteratively diverge and converge on ideas, eventually narrowing to a best-fit concept [Laseau 1988]. This article investigates the hypothesis that parallel prototyping increases learning, exploration, and design task confidence. More broadly, this research seeks a richer theoretical understanding of creative work to help practitioners and students design more effectively.

2.1 Parallel Prototyping Promotes Comparison

Throughout life, people learn interactively, trying different actions and observing their effect in the world [Gopnik et al. 2001; Piaget 2001]. Life

provides a corpus of experiences from which to draw comparisons in new learning situations [Kolodner 1993; Simon 1996]. Examples can aid problem solving [Alexander et al. 1977; Herring et al. 2009; Lee et al. 2010; Smith et al. 2008], especially when people explicitly extract principles [Gick and Holyoak 1983; Thompson et al. 2000]. Comparison helps people focus on key relations [Gentner and Markman 1997], aiding the acquisition of underlying principles [Colhoun et al. 2008; Gentner et al. 2003] and sharpening categorical boundaries [Boroditsky 2007]. This article hypothesizes that parallel prototyping better enables people to compare feedback on multiple prototypes, leading to a better understanding of how key variables interrelate.

Hypothesis 1. Parallel prototyping leads to feedback comparison and produces higher quality designs.

In the ad-design study, quality is measured with click-through analytics and expert ratings.

2.2 Parallel Prototyping Encourages Exploration

The open-ended nature of design problems often requires designers to imagine and try out alternative solutions [Buxton 2007; Kelley 2002]. Without sufficient exploration, design teams may fixate on potential solutions [Duncker 1945; Jansson and Smith 1991], overlook key insights [Kershaw and Ohlsson 2004], make poor choices to justify prior investments in money or time [Arkes and Blumer 1985], and exhibit *groupthink*, a “deterioration of mental efficiency, reality testing, and moral judgment that results from in-group pressures” [Janis 1982]. Numerous interventions have been proposed to help designers think divergently, laterally, or “outside the box” [de Bono 1999; Dym et al. 2005; Torrance 1974].

Osborn posited *premature evaluation* as a major block to organizational creativity and proposed “rules” for brainstorming: think broadly early on and save the critique for later [Osborn 1963]. Immediate feedback sets the focus on refinement, whereas postponing a critique until after creating multiple designs encourages more divergence.

Hypothesis 2. Parallel prototyping results in more divergent concepts.

In the ad-design study, independent raters judge the diversity/similarity of participants’ sets of prototypes, providing a measure of design divergence.

2.3 Parallel Prototyping Fosters Design Confidence

Self-efficacy is a person’s belief about their capabilities to perform towards a specific goal [Bandura 1997]. High self-efficacy improves one’s ability to learn [Dweck 2007], perform tasks [Bandura 1997], exert agency and persist [Mele 2005], and find enjoyment in challenges [Csikszentmihalyi 1991]. People with strong self-efficacy respond less negatively to failure and focus on strengths [Dodgson and Wood 1998]. Critique, setback, and risks make creative work extremely challenging [Schrage 1999], and high self-efficacy provides an

important robustness. With low self-efficacy, people are more likely to construe a critique as an assessment of them rather than as an assessment of the concept [Kosara 2007]. Recognizing this, the studio model of art and design education emphasizes critiquing the work rather than the person [Schon 1990].

Tohidi et al. [2006] revealed that potential users of interactive systems withhold harsh critique when presented with a single prototype; the users were concerned about offending the designer. More importantly, Tohidi et al. showed that the presence of multiple alternative concepts gave users license to be more critical with their comments. This article explores the other side of the coin: how critiquing designer-generated alternatives affects the designer's self-efficacy. This article hypothesizes that parallel prototyping changes the investment mindset: it encourages investment in a creative process rather than in a particular idea. Serial prototyping may lead people to fixate on a single concept, causing them to construe critique as a rebuke of their only option.

Hypothesis 3. Parallel prototyping leads to a greater increase in design task-specific self-efficacy.

In the ad study, self-efficacy is measured with a multi-question self-report assessment, administered before and after the design task.

3. METHOD

The study described in this article manipulates the structure of the prototyping process. Web advertising was chosen because it fulfilled the following criteria.

- Quality can be measured objectively and subjectively.
- Participants need minimal artistic or engineering ability.
- Individuals can complete tasks within a single lab session.
- Solutions demonstrate creative diversity and a range of performance quality.
- The study procedure could generate consistent and useful feedback during iteration.

3.1 Study Design

The experiment employed a between-subjects design with one independent variable: the structure of the prototyping process. The study held constant the number of prototypes created, the amount of feedback provided, and the overall time allotted. In the parallel condition, participants created 3 prototypes and got feedback, then made 2 more and got more feedback, and then made a final version. In the serial condition, participants created 5 prototypes in a series, received feedback after each prototype, then made a final version. Parallel participants were instructed to start subsequent prototypes at the same intervals as serial participants.

3.2 Participants

Thirty-three participants were recruited with fliers and assigned to one of two conditions. (Of 36 recruited, 3 dropped out before the end.) Participants'

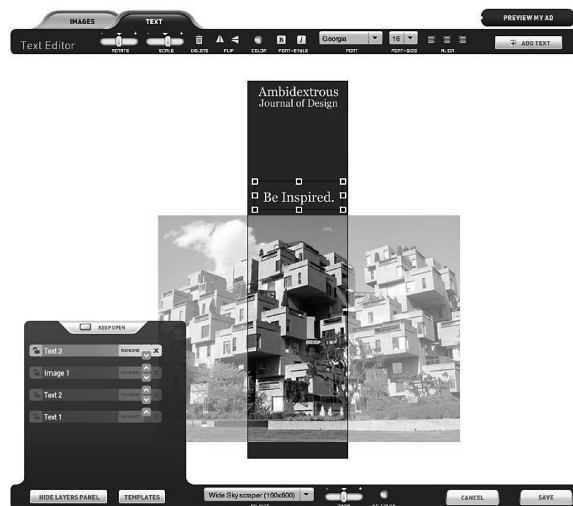


Fig. 2. The ad design study used MySpace's *AdBuilder*, a browser-based graphic design tool.

average age was 22; three-fourths were students. Using a stratified randomization approach, the study balanced gender (19 females) and prior design experience across conditions. Fourteen participants reported some prior experience in ad or graphic design; none were professional designers.

3.3 Materials

3.3.1 Graphic design tool. Participants designed a 160×600 pixel banner advertisement hosted on the social networking site MySpace.com. Ads were created using MySpace's Flash-based *AdBuilder* tool (see Figure 2). This simple graphic design tool was easy to learn, and none of the participants had used it before. Selecting a novel tool removes the confound of fluency with particular software. To insure a base level of competence, all participants had to successfully replicate a sample graphic in less than ten minutes.

3.3.2 Advertising client. Participants all created ads for the same client, *Ambidextrous* magazine, a student-led design publication. A design brief described the magazine's purpose and the kind of advertising desired by the client (Appendix D).

3.3.3 Prototype critique system. Prior to the experiment, a team of three advertising and graphic design professionals developed a list of about 50 statements that could serve as a critique for banner ads (see Appendix A). The list included three categories of statements—overall theme, composition and layout, and surface elements. Each category contained 12 to 20 statements intended to provide high-level direction without using explicitly positive or negative language. These statements express basic graphic design principles. During the study, the experimenter chose three statements, one from each category, to attach to each ad prototype (see Figure 3).

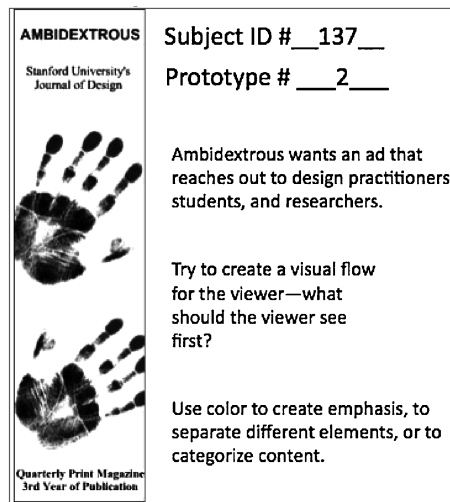


Fig. 3. Example critique.

The experimenter chose critiques relevant to each prototype and never repeated statements for the same participant. This process was identical for both conditions. Neither condition explicitly compared a participant’s ads, such as, “The color in this ad is better than that one.” In parallel, the experimenter reviewed each ad sequentially so that the process was equivalent in both conditions. The discussion section provides an in-depth treatment of the potential for bias in the study critique system.

3.4 Dependent Measures

3.4.1 Performance. After the experiment, all 33 final ad designs were uploaded to MySpace for a 15-day campaign targeted to users interested in design-related activities. This study’s total advertising costs were under \$200. Design performance was determined through two objective measures:

- MySpace click-through rates (CTR), the daily number of clicks divided by the number of impressions (number of appearances on MySpace), and
- Google Analytics¹ on the target client Web site the number of visitors, time spent, and number of pages visited daily from each ad.

Moreover, ads were independently judged by the magazine editors and by ad professionals. Editorial staff and ad professionals represent two important and different stakeholder perspectives. Four magazine editors and three advertising professionals rated the participants’ ad designs from 0 to 10 across five dimensions: adherence to the client’s theme, creativity/originality, visual appeal, tastefulness, and adherence to graphic design principles. Raters were blind to condition and rated ads individually, with no knowledge of other raters’ scores.

¹<http://www.google.com/analytics> (accessed 12/10).




Fig. 4. Example of pairwise ad similarity rating, a measure of design divergence.

3.4.2 Divergence. Creating a diverse set of ideas helps people understand the space of designs and their relative merits [Buxton 2007]. To obtain a measure of idea diversity, independent raters assessed pairwise similarity of all combinations of each participant’s ads (see Figure 4). Raters were recruited from Amazon Mechanical Turk,² a crowdsourcing system for paying workers for short online tasks. For each ad, raters assessed similarity on a scale from 1 to 7 (not similar to very similar). Each rater assessed a randomly ordered set of at least 50 ads. Rating a large number of ads helped raters calibrate their assessments. This measure generated 14,850 judgments (30 worker assessments on each of the 15 pairwise comparisons for 33 participants).

3.4.3 Self-Efficacy. Questions on self-efficacy assessed participants’ views of their graphic design ability (adopted from self-efficacy assessments in education [Fredrickson 2001; Hall 2008]). The assessment asked participants to rate their ability to: create advertisements, understand design problems, detect problems in a design idea, and incorporate feedback into a design idea (see Appendix B). Each question solicited a 7-point Likert scale response. The

²<http://www.mturk.com/mturk> (accessed 12/10).




Fig. 5. Procedure for serial and parallel conditions, with timing.

same questions were administered before and after the design task, creating a difference measure (the time between the pre- and post-test was $1\frac{3}{4}$ hours). Comparing the change in self-efficacy measures how the process manipulation (parallel/serial) influenced an individuals' belief in their design abilities.

3.5 Procedure

The experiment had the following steps (see Figure 5): consent form, pre-task questions and demographics, practice task, design brief/instructions, prototyping periods (10 minutes per prototype), critique reports (5 minutes per prototype), final design period (15 minutes), post-task questions, an open-ended interview, and a final debriefing to reiterate the consent details. The practice task required participants to replicate a graphic (unrelated to the main task). The design brief detailed the ad campaign's client, *Ambidextrous*³ magazine and outlined three goals: increase traffic to the *Ambidextrous* Web site, impress the editors, and create ads with effective graphic design.

Participants were instructed they would receive a critique from an ad expert on each prototype. As experimenters prepared critique reports in a separate room, participants were allowed to navigate the client Web site, search for images, or sketch on paper. After a set amount of time (two minutes per ad), participants received an envelope containing the printed ad prototype with feedback statements. As part of the final questionnaire, participants filled out the "Creativity Achievement Questionnaire" developed by Carson et al. [2005] to assess creative achievement across ten domains (visual arts, music, dance, architecture, writing, humor, inventions, scientific discovery, theater, and culinary arts).

For 150 minutes of participation, subjects received \$30 cash. Experiment proctors only entered the participant room to introduce the tool and task, to deliver feedback envelopes, and to conduct the open-ended interview.

4. RESULTS

Participants generated a wide variety of ad concepts. The most successful ads (high click-through rates and ratings) tended to be simple, visually balanced, professional, creative, matched the theme of the magazine, and contained some sort of intriguing hook, such as the face made of hands in the highest click-through performer (see Figure 6).

³<http://ambidextrousmag.org/> (accessed 12/10).




Fig. 6. Example ads: (Left) parallel ad, 1st in click-through rate, sixth in expert rating; (Middle), parallel ad, ninth in CTR, first in expert rating; (Right) serial ad, fourth in CTR, 32nd in expert rating.

The study supported all three hypotheses. Participants in the parallel condition produced higher quality designs (better click-through rates and higher subjective ratings) and more divergent prototypes. They also reported a greater increase in task-specific self-efficacy. Participants with prior experience in ad or graphic design outperformed complete novices, however, the prototypes created by experienced participants were less diverse than novices.

4.1 Parallel Ads Outperformed Serial Ads

4.1.1 Online click-through rates. Performance data on each ad was extracted from MySpace and Google Analytics on the *Ambidextrous* Web site (see Table I). MySpace reports that over the 15-day campaign, the 33 participant ads received 501 total clicks on 1,180,320 total impressions (i.e., number of ad appearances), giving an overall average click-through rate (CTR) of 0.0424% or 424 clicks per million impressions. The top two click-through rates were both parallel ads, with 735 and 578 clicks per million impressions, respectively. The bottom two ads were both from the serial condition; neither received any clicks.

MySpace users clicked parallel ads more than serial ads. Counting clicks can be misleading because some ads are shown more than others: when an

Table I. Summary of Campaign Data from MySpace and Google Analytics (standard deviation in parentheses)

	Parallel	Serial
Advertising host data (MySpace.com)		
Total impressions	665,133 (43968)	515,187 (36373)
Total clicks	296 (22.8)	205 (19.1)
Clicks per million impressions	445.0 (18.3)	397.9 (19.6)
Data from client site (Google Analytics reports)		
Total visitors	264 (19.9)	158 (15.3)
Average time (sec) per visitor	31.3 (143)	12.9 (79.9)
Pages visited on site	394 (31.6)	198 (21.1)
Pages visited per visitor	1.49 (0.48)	1.25 (0.41)

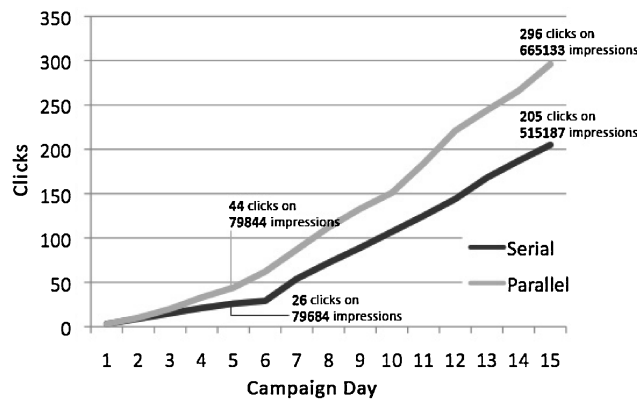


Fig. 7. Parallel ads received more clicks—and more clicks per impression—than serial ads during a 15-day campaign.

ad performs well, the host often shows it more.⁴ There are two approaches for measuring comparable performance. The first is to measure clicks per impression. The second is to hold impressions constant and compare clicks. In this study, ads received an approximately equal number of impressions for the first five days. A chi-squared analysis examines performance through day five. Parallel ads had 79,800 impressions with 44 clicks and serial ads had 79,658 impressions with 26 clicks (see Figure 7). At this early stage, parallel ads had a significantly higher click-through rate ($\chi^2 = 4.60$, $p < 0.05$).

⁴Like many advertising hosts, MySpace varies the number of impressions based on prior performance of the ad. MySpace does not publish their algorithm for determining the frequency of impressions, but a repeated-measures general linear model with the Day 5 CTR as a factor and impressions on each subsequent day as a dependent measure shows the CTR for days 1–5 to be a significant predictor of the number of impressions for the final 10 days of the campaign ($F(1,29) = 23.2$ and $p < 0.01$). MySpace receives payment on each click; intuitively, it is in their interest to show high-CTR ads more often.




Fig. 8. Parallel ads received higher average independent ratings (0–50 scale) than serial ads.

Over the entire campaign, an analysis of variances was performed with condition (serial/parallel) and creativity test scores (high/low) as factors and final click-through rates for each ad as the dependent variable. Parallel outperformed serial, 445.0 and 397.9 clicks per million impressions, respectively ($F(1,30) = 4.227$, $p < 0.05$) (see Table I).⁵ Also, high creativity scorers had a higher average click-through rate (352 clicks per million) than low scorers (305); this difference is not significant ($F(1,30) = 3.812$, $p = 0.06$).

4.1.2 Visitor behavior on client site. One common measure of ad effectiveness is the time on site [Sterne 2002]. The average time on site for parallel ads (31.3 seconds) was greater than serial ads (12.9 seconds) ($t(493) = 1.781$, $p < 0.05$). The result suggests that parallel ads were more likely to reach people genuinely interested in the product offered by the clients. The number of pages visited per visitor was about the same: 1.49 for parallel and 1.25 for serial. Visitor’s navigation behavior did not show a statistical difference: 71 of 264 visitors from parallel ads and 35 of 158 visitors from serial ads visited pages beyond the front page of *Ambidextrous*’ Web site ($\chi^2 = 1.18$, $p > 0.05$).

4.1.3 Independent expert ratings. The overall rating contained five 10-point rating scales: adherence to the client’s theme, creativity/originality, visual appeal, tastefulness, and adherence to graphic design principles. The average expert rating across all ads was 23.0 out of 50 (35.6 high and 15.0 low). The three top-rated ads were all from the parallel condition. An analysis of variances was performed with condition (parallel/serial), prior design experience (some/none), rater (seven independent raters), and rater type (client or professional) as factors and overall rating as the dependent variable. Parallel ads were rated higher ($\mu = 24.4$, $SD = 9.7$) than serial ads ($\mu = 21.7$, $SD = 8.8$) ($F(1,203) = 3.871$, $p < 0.05$) (see Figure 8). Experienced participants created higher-rated ads ($\mu = 25.7$, $SD = 9.6$) than novices ($\mu = 21.0$, $SD = 8.6$) ($F(1,203) = 20.98$,

⁵According to Google Analytics, the *Ambidextrous* Web site received 422 total visitors during the 15-day campaign, 79 less than the number of clicks reported by MySpace. One explanation for the disparity could be that users clicked the ad and then hit “back” before the browser loaded the client site. The 264 visitors from parallel ads are significantly more than the 158 visitors from serial ads when compared to impressions ($\chi^2 = 6.61$, $p < 0.05$).

$p < 0.05$) (see Figure 10). There was no interaction effect between condition and prior experience.

Some raters had higher average ratings than others ($F(5,203) = 18.88$, $p < 0.05$). There was no interaction between rater and condition; raters generally agreed that parallel ads outperformed serial ads. Analyses of variances were conducted separately for all five dimensions with condition and experience as factors. All dimensions skewed towards parallel ads, but only two—tastefulness ($F(1,227) = 7.527$, $p < 0.05$) and adherence to graphic design principles ($F(1,227) = 4.188$, $p < 0.05$)—were independently significant in favor of parallel ads. The ratings provided by the clients were higher on average ($\mu = 24.3$, $SD = 9.5$) than those provided by external ad professionals ($\mu = 22.0$, $SD = 9.1$) ($F(1,203) = 4.376$, $p < 0.05$). There was no interaction effect between rater type and condition.

Ads that performed well online generally also received high ratings by the clients and ad professionals. The ad with the best overall click-through rate received the sixth highest rating by the clients and ad professionals (see Figure 6, left). Likewise, the highest rated ad achieved the fourth highest click-through performance (see Figure 6, middle). There were anomalies, such as the top two ads in the serial condition. These two ads were ranked twenty-fifth and thirty-second (out of 33) by the expert raters, but received the third and fourth best overall click-through rates. The latter of those designs does not even mention the client (see Figure 6, right). Statistically speaking, online click performance was not a predictor of overall expert rating ($R^2 = 0.057$, $F(1,31) = 1.858$, $p > 0.05$, $b = 0.192$).

4.2 Parallel Ads Were Rated More Diverse Than Serial Ads

Workers on Amazon Mechanical Turk rated parallel ads as more divergent than serial ads. Raters performed pairwise similarity comparisons on a scale of 0 to 7 within each participant's set of six prototype ads. An analysis of variances was performed with condition (serial/parallel) and prior design experience as factors and pairwise similarity rating as the dependent variable. Serial ads were deemed significantly more similar ($\mu = 3.25$, $SD = 1.96$) than parallel ads ($\mu = 2.78$, $SD = 1.66$) ($F(1,14816) = 239.3$, $p < 0.05$). Parallel ads were rated more divergent.⁶

Similarity ratings were not predictive (or inversely predictive) of online click performance ($R^2 = 0.032$, $F(1,31) = 0.030$, $p > 0.05$, $b = 0.009$) or overall independent ratings ($R^2 = 0.030$, $F(1,31) = 1.999$, $p > 0.05$, $b = 0.246$).

⁶Similarity ratings changed depending on whether they were generated early or late in the process. Pairwise comparison of pairs 1-2, 2-3, and 1-3 were labeled “early” designs; pairs 4-5, 5-6, and 4-6 were labeled “late” designs. An analysis of variances was performed with condition (serial/parallel) and design-stage pairs (early/late/other) as factors and similarity rating as the dependent variable. Across conditions, ads created later were deemed more similar ($\mu = 3.41$, $SD = 2.03$) than early ads ($\mu = 2.97$, $SD = 1.77$) ($F(1,14814) = 107.835$, $p < 0.05$). The interaction between condition and design stage was marginally significant ($F(1,14814) = 2.460$, $p = 0.085$). Serial ads were rated more similar than parallel ads, both for early and late pairs but the similarity was greater for later ads.




Fig. 9. Participants in the parallel condition reported a greater increase in self-efficacy from pre- to post-design task.

4.3 Parallel Participants' Ad Design Self-Efficacy Increased

A self-efficacy assessment measured participants' belief in their ability to perform the design task. The difference between the pre- and post-task scores provides an indication of how participants' beliefs changed. Across all participants, self-efficacy rose from 10.85 to 12.12 (out of 20); a paired-samples T-test shows a significant difference ($t(32) = 2.355$, $p < 0.05$). Examining inter-question effects, each question independently resulted in a significant rise from pre- to post-task except for question four ("rate your ability to incorporate feedback into a design idea") ($t(32) = 0.154$, $p > 0.05$). This rise is consistent with prior findings that show individual self-efficacy beliefs increase with practice [Bandura 1997; Hall 2008].

An analysis of variances was performed with condition (serial/parallel) and prior task experience (experienced/novice). Participants in the parallel condition reported a significant increase in self-efficacy scores (see Figure 9), a net gain of 2.5 points ($F(1,29) = 4.210$, $p < 0.05$), while the serial condition essentially remained even (net gain $\mu = 0.4$).

4.4 Experienced Participants Outperformed Novices

4.4.1 Online click-through rates. The fourteen participants with prior experience in ad or graphic design significantly outperformed novices. Ads by participants with prior experience received 350 clicks on 752,424 impressions compared to 151 clicks on 427,896 impressions by novices ($\chi^2 = 8.10$, $p < 0.05$). There was no interaction effect between condition and prior participant experience.

4.4.2 Visitor behavior on client site. Visitors spent more time on the client's site after clicking ads created by experienced participants (38.0sec/visitor) compared to those created by novices (7.6sec/visitor) ($F(1,491) = 8.456$, $p < 0.05$). An interaction between condition and prior experience showed that having prior experience in the parallel condition led to more time on site than prior experience in the serial condition, 57.0 to 18.9 seconds/visitor ($F(1,491) = 4.045$, $p < 0.05$). Visitors from experienced ads were also more active navigators; 88 of 296 visitors from experienced ads and 12 of 126 visitors from novice ads




Fig. 10. Novice participants in the parallel condition reported an increase in self-efficacy from pre- to post-design task; self-efficacy for novices in serial decreased.

visited pages beyond the front page of *Ambidextrous*' Web site ($\chi^2 = 19.96$, $p < 0.05$).

4.4.3 Divergence ratings. Participants with prior experience created significantly more similar ads ($\mu = 3.15$, $SD = 1.86$) than novices ($\mu = 2.88$, $SD = 1.80$) ($F(1,14816) = 76.70$, $p < 0.05$). Ads created by novices were rated more divergent. There was also an interaction effect indicating that experienced participants in the serial condition created the most similar ads ($F(1,14816) = 36.45$, $p < 0.05$).

4.4.4 Self-Efficacy assessment. Participants with prior ad design experience reported a similar gain in self-efficacy ($\mu = 1.93$) as novices ($\mu = 0.79$) ($F(1,29) = 1.088$, $p > 0.05$). There was an interaction effect between condition and prior experience: novices reported a 2.9 increase in self-efficacy in parallel, but a 0.73 decrease in serial ($F(1,29) = 6.331$, $p < 0.05$) (see Figure 10). In short, parallel prototyping positively affected an individual's belief in their ad design ability, especially for novices.

5. ANALYSIS

A parallel approach led to ad creations with better performance by every measure: higher independent ratings, more impressions served up by MySpace, better click-through rates, more visitors to the client Web site, and more site interaction per visitor. Participants created the same number of prototypes and received equivalent feedback in the same time period. The only difference between conditions was a matter of when participants received a critique on their ideas, after each concept or after multiple creations.

Why did the process manipulation affect performance? This section offers three explanations for the differential results: comparison helped the parallel participants learn ad design principles, parallel participants better explored the design space, and serial participants perceived the critique as negative and thus gained no confidence in ad design.

Did parallel feedback impact how participants learned? Comparison processes can facilitate inductive reasoning on rival observations [Colhoun et al. 2008; Thompson et al. 2000]. Since parallel participants received feedback on multiple ideas simultaneously, they were more likely to read and analyze critique statements side-by-side. Direct comparison perhaps helped parallel participants better understand key design principles and led to more principled choices for subsequent prototypes. In serial prototyping, participants' ideas tended to follow directly from the feedback. This serial approach may implicitly encourage refinement at the expense of exploration. Performance likely improves in parallel because people exercise their comparative abilities to learn contextual constraints, malleable variables, and their interrelations.

Learning a parallel approach may change future behavior. When asked to describe their process for future design projects, 11 of 16 parallel participants said they would create more than one prototype and obtain copious feedback; only 5 of 17 serial participants made similar claims ($\chi^2 = 2.63$, $p > 0.05$). As one parallel participant said, "not spending too much time on any single prototype is useful because then you don't go into details as much."

Did a parallel process impact how participants explored concepts? The study showed parallel participants created significantly more divergent prototypes; serial participants tended to create more similar designs. The interviews revealed the role of the critique as one serial participant explained, "I think the feedback helped. I kept repeating the same mistakes, but maybe less and less each time . . . the feedback reiterated that." Another serial participant said:

"I would try to find a good idea, and then use that idea and keep improving it and getting feedback. So I pretty much stuck with the same idea."

This notion of sticking with an idea or using the feedback to decide where to go next did not surface in the parallel condition. As one parallel participant reported: "I didn't really try to copy off of the ads that I did before...I just made new ideas." Both the divergence measure and the qualitative interviews suggest the parallel structure supports more generative thinking and reduces fixation.

Parallel prototyping may encourage both a broad enumeration stage and a subsequent reflection stage. By contrast, immediate feedback from serial prototyping implicitly encourages refinement. On this view, the fact that parallel prototyping delays feedback is actually an advantage. From a behaviorist perspective, this can seem counterintuitive because immediate feedback highlights the connection between cause and effect. However, delay helps learners reflect: readily available, immediate feedback can be a crutch that discourages building a deep understanding [Anderson and Schooler 1990; Schmidt et al. 1989].

There are countless ways to combine text, images, and backgrounds in a 160×600 pixel ad design; some combinations perform better than others. To use an analogy, exploring design possibilities is like simulated annealing [Granville 1994]. Creative work often benefits from broadly exploring a design space with high entropy before optimizing in one direction. Perhaps serial participants hill-climbed to local, rather than global optima.

Experienced participants created a less diverse set of designs than novices; they also outperformed novices. In general, experts may know a priori which areas are promising and which to avoid. By contrast, novices have to learn what is effective through trial and error.

Did parallel participants gain more confidence in their ad-design ability? Parallel participants reported self-efficacy gains, while the serial participants reported no change. This effect was more pronounced for novices. Serial participants also perceived the expert feedback more negatively. In open-ended interviews, 13 of 16 parallel participants said the feedback was helpful or intuitive compared to 6 of 17 serial participants ($\chi^2 = 7.13$, $p < 0.05$). More notably, 8 of 17 of the serial participants reported the feedback as negative, compared to no such reports in the parallel condition ($\chi^2 = 9.94$, $p < 0.05$). One participant in the serial condition said:

“I received really negative comments saying (the clients) are looking for a creative and clever ad, which in other words is saying that this is stupid or ridiculous.”

Moreover, participants were asked to leave their email if they wanted to later volunteer for *Ambidextrous* magazine. Twelve out of sixteen parallel participants provided their email, while only five of seventeen serial participants did the same ($\chi^2 = 6.86$, $p < 0.05$), which suggests the parallel process may have helped motivate future action.

Perhaps having multiple alternative designs encourages investment in a creative process rather than a particular idea. Consequently, the parallel process encourages viewing the critique as an opportunity for improvement. In contrast, the fixation engendered by serial prototyping may cause people to take the critique as a catastrophic rebuke of their only option. With only one option there is no separation between designer and design. Parallel offers people distance between ego and object; serial conflates them.

6. FOLLOW-UP STUDIES

Two short follow-up experiments examined questions raised by the main study.

Did the experimenters (possibly subconsciously) provide better critique to the parallel participants? To assess bias, two ad professionals unfamiliar with the experimental manipulation provided blind-to-condition independent assessments of the critique statements. The expert judges performed a selection task resembling the task performed by the experimenter. After reading about the client’s advertising needs, the judge viewed an ad prototype and two triads of critique statements; one triad contained the three statements chosen during the actual experiment and the other triad was a random selection from the critique statement corpus. Judges were instructed to select a triad that “provides the most appropriate critique for the advertisement.”

An intra-class correlation (ICC) with a two-way mixed model [Shrout and Fleiss 1979] calculated the reliability between the experimenter’s choice of statements and each expert judge’s choice. The ICC(3,1) single measure correlation among raters on parallel ads is 0.434, and 0.496 for serial ads. There is no significant difference between these numbers and both represent a

moderate level of inter-rater reliability [Landis and Koch 1977]. In short, experts generally agreed with the feedback provided, and the level of agreement was comparable across conditions.

Did the critique statements help participants produce better ads? A follow-up study examined the value of the scripted ad critique statements in Appendix A. Thirty participants followed a serial process to create three prototypes and one final advertisement. The final ads were launched in an online ad campaign and rated by experts. Participants were randomly assigned to one of three conditions: *targeted*, *random*, and *none*. In the targeted condition, an experimenter selected three critique statements intended to help the participant improve their design. In the random condition, a random algorithm selected three critique statements. Participants in the none condition received no critique; rather, they viewed the client Web site during an equivalent critique interval.

In a 3-week campaign, ads that received a targeted critique had 49,414 impressions with 203 clicks, ads with no feedback had 49,663 impressions with 179 clicks, and ads that received randomly selected critique statements received 49,543 impressions with 157 clicks ($\chi^2 = 6.01$, $p < 0.05$). Moreover, twenty independent experts rated ads (on a 0 to 30 scale) with the targeted critique higher $\mu = 15.9$ (SD = 5.4) than ads with a random critique $\mu = 15.1$ (SD = 5.2) and ads with no critique $\mu = 14.4$ (SD = 6.2) ($F(2,597) = 3.287$, $p < 0.05$). The study found that the targeted critique helped participants learn basic principles of graphics design and produce better ads.

7. DISCUSSION ABOUT THE EXPERIMENTAL APPROACH

The experimental paradigm introduced in this article provides several important benefits for studying design. First, design results can be objectively measured through real-world analytics and subjectively assessed through crowdsourced and stakeholder ratings. Second, solutions demonstrate creative diversity and exhibit a broad range of performance. Third, it offers a mechanism for presenting feedback interactively and studying its effects. The advertising domain achieves these goals particularly well. Hopefully this paradigm will prove useful in additional domains.

Web analytics can be tremendously valuable for experimental work, but also present several challenges. Web hosts often show ads differentially based on performance. Poorly performing ads must have a large enough number of impressions to yield a robust measure of click-through rate. Additionally, click-through rate can vary over time. Fair comparison requires holding the number of impressions constant, analyzing data from a time interval with a roughly balanced impression rate, or using more sophisticated statistical analysis to factor out time effects.

8. CONCLUSIONS AND FUTURE WORK

This article found that when people create multiple alternatives in parallel, they produce higher-quality, more-diverse work and experience a greater increase in self-efficacy. Many excellent designers practice this approach already;

their successes inspired this research. Hopefully, these results will encourage many more practitioners and teachers to adopt a parallel approach. Integrating the parallel approach into design practicum can inculcate healthy prototyping habits and help foster a positive outlook toward critique. In the future, software tools and infrastructure providers could provide a powerful benefit by enabling people to rapidly design alternatives and experimentally compare them. More broadly, this research seeks to develop a theoretical understanding of creative work to help practitioners and students solve design problems more effectively. An important direction for future work is to study the impact of parallel design in other contexts, especially for longer time scales and for design teams.

APPENDIXES

APPENDIX A: EXPERT CRITIQUE STATEMENTS

1. Overall/Thematic

Ambidextrous seeks an ad with a single clear message that matches the theme of their journal.

Ambidextrous wants an ad that clarifies the product: a journal about design and design process.

Ambidextrous desires an ad that is simple, readable, consistent, and deliberate.

Ambidextrous does not want the ad to sound exclusive; they are open to anyone with interest.

Ambidextrous is looking for a creative and clever ad.

Ambidextrous is looking for a professional and tasteful ad.

Ambidextrous wants an exciting and visually appealing ad.

Ambidextrous wants an ad that matches the journal's style.

Ambidextrous wants an ad that reaches out to design practitioners, students, and researchers.

Use graphics/images that support the overall message. What message are you trying to convey?

Use colors/fonts that support the overall message. What message are you trying to convey?

Remember that the ad is a link; the URL does not *necessarily* have to be on the ad design.

2. Composition & Layout

Visual Flow and Balance

Try to create a balanced layout where the graphics don't tilt to one side or the other.

Try to create a visual flow for the viewer—what should the viewer see first?

Think about the proximity of different elements. How close together or far apart elements are placed suggests a relationship (or lack thereof) between otherwise disparate parts.

To help balance the ad, leave slightly more space at the bottom relative to the top of the ad.

Contrast the position of elements to draw the viewer's attention to the most important parts.

To create consistency for the viewer, create a consistent and balanced look using repetition.

Spacing and Alignment

Align text and graphics to create more interesting, dynamic, and appropriate layouts.

Use alignment to create a clean and organized look.

It's ok to break alignment only to draw the viewer's attention to important elements in the ad.

Use white around text and images to help frame the content.

Use space—the absence of text and graphics—to provide visual breathing room for the eye.

Try to balance the spacing around the border of the ad design.

These visual elements in the ad don't line up.

Consider playing around with different ways to justify the text (e.g., center, left, or right-justified).

Emphasis & Hierarchy

Be conscious of competing elements in the ad. Think about what should have emphasis.

Draw the viewer's attention to elements by contrasting size (scale).

Think about the visual hierarchy of the different elements (texts, images, colors, etc) of the ad. What is the most important?

Help the viewer recognize, identify and comprehend the most important information in the ad.

Use elements with visual intensity or color for emphasis.

3. Fonts, Colors, Images

Font Type

Try not to distort the font so that it becomes hard to read.

Use large, bold font/graphics to create focus or emphasis on the ad design.

If using text over an image, make the text bigger and darker than normal; make sure it is readable.

For text to stand out it has to be substantially different than other text.

Try not to mix serif and sans serif fonts.

Avoid using two different fonts that are too similar.

Try not to over emphasize text elements. (ex. a font does not need to be large, bold, and italic).

Images

Use large, bold graphics to create the focus of the ad design.

Consider using images for more visual impact.

Consider using fewer images.

Try not to over-rotate images, as it often distorts the content.

Color

Use color to create emphasis, to separate different elements, or to categorize content.

Avoid really light, bright colors.

Avoid colors together that look too similar (ex. brown & grey).

Try to use different colors that go well together.

Avoid complicated backgrounds.

Try to create a good visual separation between the text and the background

APPENDIX B: SELF-EFFICACY QUESTIONS (PRE & POST TASK)

On a scale from “Not confident at all” (1) to “Very confident” (7), how confident are you:

- (1) With your ability to design advertisements?
- (2) At understanding design problems?
- (3) With detecting problems in your design?
- (4) With incorporating expert feedback in your design?

APPENDIX C: POST INTERVIEW GUIDE

These questions provided guidance for the final interview; the exact order and phrasing varied.

- Please describe the final design you came up with.
- What do you think are the strengths and weaknesses of your ad design?
- Describe your design process and why you created each design.
- How did the feedback affect you? Was it helpful? What did you learn about graphic ad design?
- If you created another ad, how would you approach it? Describe your design process. Would you seek feedback? How many prototypes would you create?

APPENDIX D: ADVERTISING DESIGN BRIEF

Assignment

You have been hired to design a graphic advertisement for *Ambidextrous*, Stanford University's Journal of Design. You will learn a new graphic design tool, prototype a number of example ads, receive feedback from an ad design expert, and then create a final ad that will launch on MySpace.com.

Goals

Keep in mind the following goals as you create your ads:

- (a) Increase traffic to the *Ambidextrous* website: <http://ambidextrousmag.org/>
- (b) Reach out to the target audience: designers, researchers, practitioners, and students who are interested in stories about the process of design.
- (c) Impress the editors of *Ambidextrous*. The client wants an ad that fits their overall aesthetic and theme (see below).
- (d) Create ads with effective graphic design.

What is Ambidextrous?

Ambidextrous is Stanford University's Journal of Design. In its 3rd year of publication, *Ambidextrous* is a quarterly subscription-based print magazine that features people and processes involved in design and design thinking. *Ambidextrous* is a forum for the cross-disciplinary, cross-market community of people with an academic, professional and personal interest in design. Each issue focuses on one general topic. Previous topics have included Space, Secrecy, Food, The Future, Danger, Developing, Borders & Interfaces, etc. Articles are written by the community at large, and edited, illustrated, and photographed entirely by volunteers.

Theme and Aesthetic for the Ambidextrous Ad

The *Ambidextrous* editors would like an ad that embodies the theme and general aesthetic of the journal. The journal tells stories about people who do design and the process of getting there, not just final products. Readers of the journal are not an exclusive club—it's intended to be accessible to folks without formal design training. In general they are looking for an ad that is tasteful, creative, professional, visually appealing, and conveys a clear message about the product.

Rules/Requirements

- You may download and use graphics, images, text etc. as you see fit.
- You may not use another company's logo, copyrighted images, profanity, obscenity or nudity. Unacceptable ads will be rejected by the research team.
- Do not include the magazine's URL on the ad. Clicking the ad will direct the user to the site.

ACKNOWLEDGMENTS

We thank Ugochi Acholonu, Lera Boroditsky, Jodi Forlizzi, Jeffrey Heer, Erik Stolterman, Ewart Thomas, and Terry Winograd for valuable discussions about

this research. We thank Rachel Lopatin for analyzing the targeted critique data.

REFERENCES

- ALEXANDER, C., ISHIKAWA, S., AND SILVERSTEIN, M. 1977. *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
- ARKES, H. R. AND BLUMER, C. 1985. The psychology of sunk cost. *Organiz. Behav. Hum. Decision Proc.* 35, 124–140.
- BALL, L. J., EVANS, JStB, DENNIS, I. D., AND ORMEROD, T. C. 1997. Problem-solving strategies and expertise in engineering design. *Think. Reason.* 3, 247.
- BALL, L. J. AND ORMEROD, T. C. 1995. Structured and opportunistic processing in design: A critical discussion. *Int. J. Hum.-Comput. Stud.* 43, 131–151.
- BANDURA, A. 1997. *Self-Efficacy: The Exercise of Control* 1st Ed. Worth Publishers.
- DE BONO, E. 1999. *Six Thinking Hats* 2nd Ed. Back Bay Books.
- BORODITSKY, L. 2007. Comparison and the development of knowledge. *Cognition* 102, 118–128.
- BUXTON, B. 2007. *Sketching User Experiences: Getting the Design Right and the Right Design*. Morgan Kaufmann.
- CARSON, S. H., PETERSON, J. B., AND HIGGENS, D. M. 2005. Reliability, validity, and factor structure of the creative achievement questionnaire. *Creativity Resear. J.* 17, 37–50.
- COLHOUN, J., GENTNER, D., AND LOEWENSTEIN, J. 2008. Learning abstract principles through principle-case comparison. In *Proceedings of Cognitive Science Society*. 1659–1664.
- CROSS, N. 2006. *Designersly Ways of Knowing* 1st Ed. Springer.
- CROSS, N. 2004. Expertise in design: An overview. *Design Studies* 25, 427–441.
- CSIKSZENTMIHALYI, M. 1991. *Flow: The Psychology of Optimal Experience*. Harper Perennial.
- DODGSON, P. G. AND WOOD, J. V. 1998. Self-esteem and the cognitive accessibility of strengths and weaknesses after failure. *J. Personality Social Psy.* 75, :178–197.
- DORST, K. AND CROSS, N. 2001. Creativity in the design process: Co-evolution of problem-solution. *Design Studies* 22, 425–437.
- DOW, S. P., HEDDLESTON, K., AND KLEMMER, S. R. 2009. The efficacy of prototyping under time constraints. In *Proceeding of the 7th ACM Conference on Creativity and Cognition*. 165–174.
- DUNCKER, K. 1945. *On Problem-Solving*. Greenwood Press Reprint.
- DWECK, C. 2007. *Mindset: The New Psychology of Success*. Ballantine Books.
- DYM, C. L., AGOGINO, A. M., FREY, D. D., AND LEIFER, L. J. 2005. Engineering design thinking, teaching, and learning. *J. Engin. Educ.* 94, 103–120.
- FREDRICKSON, B. L. 2001. The role of positive emotions in positive psychology: The broaden-and-build theory of positive emotions. *American Psych.* 56, 218–226.
- GENTNER, D., LOEWENSTEIN, J., AND THOMPSON, L. 2003. Learning and transfer: A general role for analogical encoding. *J. Educ. Psych.* 95, 408, 393.
- GENTNER, D. AND MARKMAN, A. B. 1997. Structure mapping in analogy and similarity. *Amer. Psych.* 52, 45–56.
- GICK, M. L. AND KEITH J. H. 1983. Schema induction and analogical transfer. *Cognitive Psych.* 15, 1–38.
- GOEL, V. AND PIROLI, P. 1992. The structure of design problem spaces. *Cognitive Science* 16, 395–429.
- GOPNIK, A., MELTZOFF, A. N., AND KUHL, P. K. 2001. *The Scientist in the Crib: What Early Learning Tells Us About the Mind*. Harper Paperbacks.
- GRANVILLE, V., KRIVÁNEK, M., AND RASSON, J. P. 1994. Simulated annealing: A proof of convergence. *IEEE Trans. Patt. Anal. Mach. Intell.* 16, 652–656.
- HALL, T. S. 2008. Improving self-efficacy in problem solving: Learning from errors and feedback. PhD thesis, University of North Carolina, Greensboro.
- HARTMANN, B., KLEMMER, S. R., BERNSTEIN, M., ABDULLA, L., BURR, B., ROBINSON-MOSHER, AND GEE, J. 2006. ‘Reflective physical prototyping through integrated design, test, and analysis. In *Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*. ACM, 299–308.

- HERRING, S. R., CHANG, C.-C., KRANTZLER, J., AND BALLEY, B. P. 2009. Getting inspired!: Understanding how and why examples are used in creative design practice. In *Proceedings of the 27th International Conference on Human Factors in Computing Systems*. ACM, 87–96.
- JANIS, I. L. 1982. *Groupthink: Psychological Studies of Policy Decisions and Fiascoes* 2nd Ed. Wadsworth Publishing.
- JANSSON, D. AND SMITH, S. 1991. Design fixation. *Design Studies* 12, 3–11.
- KELLEY, T. 2002. *The Art of Innovation*. Broadway Business Books.
- KERSHAW, T. C. AND OHLSSON, S. 2004. Multiple causes of difficulty in insight: The case of the nine-dot problem. *J. Exper. Psych.; Learn., Memory, Cognition* 30, 3–13.
- KOLODNER, J. 1993. *Case-Based Reasoning*. Morgan Kaufmann.
- KOSARA, R. 2007. Visualization criticism—the missing link between information visualization and art. In *Proceedings of the Conference on Information Visualization*. IEEE Computer Society, 631–636.
- LANDIS, J. R. AND KOCH, G. G. 1977. The measurement of observer agreement for categorical data. *Biometrics* 33, 159–174.
- LASEAU, P. 1988. *Graphic Thinking for Architects and Designers* 2nd Ed. John Wiley & Sons Inc.
- LEE, B., SRIVASTAVA, S., KUMAR, R., BRAFMAN, R., AND KLEMMER, S. R. 2000. Designing with interactive example galleries. In *Proceedings of the ACM SIGCHI Conference on Human Factors in Computing Systems*.
- MELE, A. R. 2005. *Motivation and Agency*. Oxford University Press.
- NEWELL, A. 1972. *Human Problem Solving*. Prentice Hall.
- NIELSEN, J. AND FABER, J. M. 1996. Improving system usability through parallel design. *Computer* 29, 29–35.
- OSBORN, A. F. 1963. *Applied Imagination: Principles and Procedures of Creative Problem Solving*. Charles Scribner's Sons.
- PIAGET, J. 2001. *The Psychology of Intelligence*. Routledge.
- RITTEL, H. AND WEBBER, M. 1973. Dilemmas in a general theory of planning. *Policy Sci.* 4, 169, 155.
- SCHMIDT, R. A., YOUNG, D. E., SWINNEN, S., AND SHAPIRO, D. C. 1989. Summary knowledge of results for skill acquisition: Support for the guidance hypothesis. *J. Exper. Psych.: Learn., Memory, Cognition* 15, 352–359.
- SCHON, D. A. 1990. *Educating the Reflective Practitioner: Toward a New Design for Teaching and Learning in the Professions* 1st Ed. Jossey-Bass.
- SCHON, D. A. 1995. *The Reflective Practitioner: How Professionals Think in Action*. Ashgate Publishing.
- SCHOOLER, L. J., AND ANDERSON, J. R. 1990. The disruptive potential of immediate feedback. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society*.
- SCHRAGE, M. 1999. *Serious Play: How the World's Best Companies Simulate to Innovate*. Harvard Business School Press.
- SHROUT, P. E. AND FLEISS, J. L. 1979. Intraclass correlations: Uses in assessing rater reliability. *Psych. Bull.* 86, 420–428.
- SIMON, H. A. 1996. *The Sciences of the Artificial* 3rd Ed. The MIT Press.
- SMITH, S., KOHN, N., AND SHAH, J. 2008. What you see is what you get: Effects of provocative stimuli on creative invention. In *Proceedings of the NSF International Workshop on Studying Design Creativity*.
- STERNE, J. 2002. *Web Metrics: Proven Methods for Measuring Web Site Success* 1st Ed. Wiley.
- THOMPSON, L., GENTNER, D., AND LOEWENSTEIN, J. 2000. Avoiding missed opportunities in managerial life: Analogical training more powerful than individual case training. *Organiz. Behav. Hum. Decision Process.* 82, 60–75.
- TOHIDI, M., BUXTON, W., BAECKER, R., AND SELLEN, A. 2006. Getting the right design and the design right. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1243–1252.
- TORRANCE, E. P. 1974. *Torrance Tests of Creative Thinking*. Personnel Press, Ginn and Co., Xerox Education Co.

Received May 2010; revised August 2010; accepted August 2010