

Improving Crowd Innovation with Expert Facilitation

Joel Chan Steven Dang Steven P. Dow

Human-Computer Interaction Institute
Carnegie Mellon University, Pittsburgh, PA USA
{joelchuc, stevenda, spdow}@cs.cmu.edu

ABSTRACT

Online crowds are a promising source of new innovations. However, crowd innovation quality does not always match its quantity. In this paper, we explore how to improve crowd innovation with real-time expert guidance. One approach would for experts to provide personalized feedback, but this scales poorly, and may lead to premature convergence during creative work. Drawing on strategies for facilitating face-to-face brainstorming, we introduce a crowd ideation system where experts monitor incoming ideas through a dashboard and offer high-level "inspirations" to guide ideation. A series of controlled experiments show that experienced facilitators increased the quantity and creativity of workers' ideas compared to unfacilitated workers, while Novice facilitators reduced workers' creativity. Analyses of inspiration strategies suggest these opposing results stem from differential use of successful inspiration strategies (e.g., provoking mental simulations). The results show that expert facilitation can significantly improve crowd innovation, but inexperienced facilitators may need scaffolding to be successful.

Author Keywords

Creativity; crowdsourcing; brainstorming; expert facilitation

ACM Classification Keywords

H.5.3. Information interfaces and presentation (e.g., HCI): Group and Organization Interfaces: Computer-supported cooperative work.

INTRODUCTION

From complex R&D problems [28], to product design [4], to social innovation challenges [29], organizations increasingly turn to online crowds to obtain fresh perspectives on challenging problems. Theoretically, the scale and diversity of crowds offer increased chances of obtaining exceptional solutions. In practice, crowds often excel at generating *many* ideas, but often fail to reliably generate many *creative* ideas, i.e., ideas that are both novel and valuable [18]. For

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CSCW '16, February 27-March 02, 2016, San Francisco, CA, USA

© 2016 ACM. ISBN 978-1-4503-3592-8/16/02...\$15.00

DOI: <http://dx.doi.org/10.1145/2818048.2820023>

example, Dell's IdeaStorm platform has implemented approximately 550 product ideas gathered from the crowd, but these are laboriously culled from more than twenty thousand idea submissions, many of which are duplicate ideas or too vague/impractical to add value as a new product. Crowd workers may lack the ability to identify and productively build on promising solutions, whether due to lack of expertise [23] or overreliance on signals such as community upvotes, which may simply reflect the popularity of ideas, as opposed to their creativity [5,24].

Prior research has explored strategies for integrating experts into crowd innovation processes, from establishing creative goals [22], to leading coordination efforts [33,43], to providing timely, task-specific feedback [14]. These strategies improve creative outcomes, but they can be difficult to perform at crowd scale. Further, while expert guidance can help crowds focus their efforts and converge on high-value solutions, it might prevent divergent thinking. For example, gold standard examples [38] (e.g., showing workers exemplary solutions) could lead to premature convergence during creative tasks, since people often have a hard time breaking away from solutions known to be successful in the past [4,32]. Strict assessment can also lead to evaluation apprehension [13], causing people to be reluctant to explore "wild" ideas, an important strategy for finding exceptional (not just "good") ideas [53].

This paper explores how we might improve crowd innovation by adapting expert facilitation strategies from face-to-face brainstorming [20,41,52]. Expert facilitators guide ideation by pointing out promising solution approaches to inspire further ideation. Importantly, skilled facilitators do not simply highlight particular ideas: they often highlight key high-level characteristics or schemas exemplified by the idea, and provoke reconsideration of implicit assumptions about the design problem [52]. For example, a common facilitation strategy is to say, "*X is an interesting idea. How else might we <leverage feature Y of idea X>?*"

In this paper, we adapt strategies for expert facilitation into a system for real-time crowd ideation and evaluate its potential with a series of online controlled experiments. IdeaGens provides a dashboard for a skilled facilitator to monitor the evolving solution space and to offer inspirations (i.e., ideas, questions, provocations) for crowd ideators. Crowd workers can request these inspirations from a queue to inspire their thinking on a problem.

We evaluated IdeaGens with two controlled experiments. In Experiment 1, crowd workers ($N=87$) on Amazon Mechanical Turk (MTurk [35]) ideated solutions for a common social predicament (forgetting an acquaintance's name). Participants either brainstormed independently with no facilitation or received high-level guidance (i.e., inspirations) from two facilitators with prior experience managing brainstorming sessions. Results show that facilitated participants generated more ideas of higher creativity (as rated by blind-to-condition judges) than unfacilitated participants. As measured by Latent Semantic Analysis [30], facilitated participants had higher convergence (i.e., higher occurrence of highly similar idea pairs, an index of design iteration) than and equal divergence (semantic diversity of ideas) as unfacilitated participants.

In Experiment 2, we recruited three facilitators with little to no prior experience leading brainstorming sessions to use IdeaGens to guide the crowd. Using an identical study design as Experiment 1, 85 crowd participants either received facilitation or not. In contrast to Experiment 1, facilitated participants generated less creative ideas than unfacilitated participants. Content analyses of inspirations generated across both experiments suggest that these opposing results could be explained by differences in inspiration strategies employed by facilitators. Experienced facilitators used more open-ended questions and provoked more mental simulation (which was significantly correlated with higher creativity ratings), while inexperienced facilitators relied heavily on simply highlighting and distributing examples.

This paper makes three contributions:

- 1) We designed a crowd ideation system inspired by a successful strategy from face-to-face group brainstorming and addressed key challenges related to the crowd context (e.g., monitoring a large, evolving solution space, providing flexible and timely guidance at scale)
- 2) We ran experiments that demonstrate the value of expert facilitation for online crowds, and
- 3) We conducted a content analysis to illuminate successful strategies for facilitating crowd ideation and discuss design considerations for more effective facilitation.

RELATED WORK

Ensuring Quality Crowd Work

There are a range of strategies for dealing with low quality crowd work, such as comparing worker output to known high quality “gold standard” answers to screen workers [38], using behavioral traces to identify good workers [47], and weeding out “bad answers” with aggregation techniques like majority voting [19]. However, many of these techniques do not apply straightforwardly to crowd innovation. For example, innovators rarely know in advance what the best solutions will be, ruling out the possibility of using gold standard items. Aggregation techniques like majority

voting might miss fresh perspectives that deviate from the consensus.

Our research builds on techniques designed to enhance quality on more open-ended crowd tasks by leveraging experts for real-time input [14,22]. For example, *Shepherd* [14] enables requesters to provide timely expert feedback on crowd workers’ product reviews. *Ensemble* enables an expert lead-author to provide creative direction for ensembles of crowd workers who generate content for small pieces of a larger short story [22]. Lead authors guide work by defining “story problem” prompts for specific scenes in the story (e.g., “How can Character X meet Character Y?”), and provide feedback on contributions through comments.

One key challenge of applying real-time expert guidance to crowd innovation problems is scaling it to tens or potentially even hundreds to many thousands of participants (vs. less than 10 per team in *Ensemble*). In this research, we consider key design parameters that might influence how guided facilitation scales, such as the granularity (e.g., feedback on individual ideas vs. validation of general solution approaches) and input source (experts vs. novice peers). Drawing inspiration from systems that surface real-time information on crowd work [47], IdeaGens features a “dashboard” that shows submitted ideas and visualizes semantic information to highlight the evolution of the solution space.

Facilitating Effective Idea Generation

The goal of idea generation is to discover exceptional ideas that can provide a solid foundation for later stages of the creative process (e.g., prototyping), ultimately culminating in a creative product, i.e., one that is both novel and valuable [18]. The literature on creative ideation emphasizes two aspects of ideation that must be simultaneously optimized to achieve this goal. On the one hand, the search for solutions in the design space must be sufficiently *divergent* in order to not miss promising solution approaches. Divergence involves exploring many ideas [1,44,53] and searching broadly in the solution space to encompass a variety of distinct solution approaches [9,48,54] (e.g., many ideas are semantically distant from each other). On the other hand, *convergent* search is needed to combine and refine shallow or half-baked ideas into more creative ones [36,39,46]. Convergence involves elaborating ideas with more detail [17] and exploring variations on themes [10] (e.g., at least a few ideas are semantically close in the solution space).

There are known strategies for promoting divergence with crowd ideation. For example, innovators can increase the number and/or diversity of individuals recruited [3,15]. In contrast, innovators lack reliable strategies for promoting convergence. Signals of idea value based on community upvotes or comments are frequently employed, but are often unreliable: these signals are often driven more by popularity [5] or “rich get richer” effects [24], rather than ideas’ actual innovative potential. Other innovators manage the convergence process themselves, spending valuable in-house time (on the order of many weeks) to consolidate

promising solutions based on the crowds’ ideas [5]. However, this strategy scales poorly, leading to failures to actually benefit from the crowd’s contributions in a timely manner. For example, the change.gov online ideation website was shut down prematurely because the staff were not able to meaningfully process the huge volume of contributions. Similarly, in their 10 to the 100th crowdsourced innovation project, Google had to recruit 3,000 employees to prune the 150,000 ideas received, resulting in a nine-month delay in their project timeline.

One successful strategy for simultaneously improving divergence and convergence (in face-to-face group brainstorming) is to employ a skilled *facilitator* [20,41,52]. Prior studies show that face-to-face groups with a dedicated facilitator outperform groups with no facilitation in terms of both divergent and convergent performance [20,25,42]. The literature on facilitation distinguishes between two overarching categories of actions that facilitators can take to improve ideation. Process facilitation focuses on the group’s process or relationships, e.g., by ensuring equitable opportunities for member contributions and managing/mediating group conflict [12]. Content facilitation directly influences the substance/content of the group’s work [12], e.g., by providing inspiring images or prototypes [41] and calling attention to emergent themes and unique ideas [52]. This paper focuses more on enabling content facilitation of crowd ideation through inspirations that stimulate ideation along promising solution paths.

IDEAGENS

Drawing on principles and strategies for improving crowd work quality and facilitating effective ideation, we designed IdeaGens with the following guidelines in mind:

- *Responsiveness*: Enable facilitators to monitor and responsively guide ideation as it unfolds over time
- *Flexibility*: Support a range of inspiration strategies that

apply to diverse types of innovation problems

- *Scalability*: Allow one or a few people to manage a large crowd of workers

IdeaGens is built in MeteorJS, a full-stack Javascript web application framework built on Node.js. The system includes an *ideator interface* where crowd workers can generate ideas in parallel, and a *facilitation dashboard* that enables real-time monitoring and guiding of the crowd’s ideation. The core of IdeaGens is an *inspiration system* that links the dashboard and individual ideator interfaces. The dashboard enables facilitators to create inspirations (as open-ended text-based messages) that call out interesting themes or frame the problem in new ways.

One key design consideration is how to distribute inspirations across ideators. In typical face-to-face brainstorming, facilitators typically “push” guidance, gently interrupting the discussion at an appropriate time (e.g., during lulls in the discussion) with prompts or questions that are tailored to the group’s discussion. However, we felt that this “push” model would not scale to facilitating many tens to potentially hundreds of ideators working in parallel. Indeed, in pilot testing with earlier iterations of the tool, we found that facilitators were not able to effectively and efficiently decide when and to whom to distribute inspirations, even with as few as 8-10 ideators. Therefore, we implemented a “pull” mechanism for inspiration distribution. The system collects inspirations in a queue, which ideators can “pull” from on-demand in a simple first-in-first-out algorithm (i.e., older inspirations pulled first). The system keeps a tally of the number of ideators and ensures that there are always enough “copies” of each inspiration for all workers to access if they choose. This “pull” approach supports greater scalability and was motivated by prior work showing that ideators benefit most from inspirations when delivered “on demand” versus pushed or on a regular interval [49].

Ideation Interface

The ideation interface enables entry of new ideas for the brainstorming prompt, either for the general prompt (left column of the interface), or related to particular inspirations (right column). At any time they wish, ideators can press the “Inspire Me” button (located below the brainstorming prompt) to pull new inspirations from the inspiration queue (see Fig. 1). Each button press yields a single new inspiration, which appears directly below the button. Each inspiration includes its own text entry box, which ideators can use to enter ideas inspired by that particular inspiration. As ideators request additional inspirations, older inspirations move to the right and users can scroll left and right to review their inspirations as well as any relevant ideas. In this version of the system, we only allow ideators to see their own ideas. We do this because there is a lack of principled guidance from the literature on how to best enable ideators to benefit from large numbers of ideas from other members of the crowd. Research suggests that ideators benefit most from other contributions when they are




Figure 1. Ideator interface allows ideators to receive inspirations on-demand by clicking on the “Inspire Me” button. To provide feedback to facilitators, ideators are encouraged to enter ideas sparked by an inspiration into the inspiration-specific entry box.




Figure 2. Dashboard enables facilitators to monitor the evolving solution space, as well as guide crowd ideation through the creation of inspirations. Facilitators also receive feedback on their inspirations by inspecting ideas that were inspired by each inspiration.

able to give careful attention to those contributions [21,51]; but at the scale of crowd ideation, the volume of ideas may overwhelm ideators' limited cognitive resources. Some research also suggests that, when given too many ideas as potential inspiration, people may stop attending to them, or only build on them in superficial ways [21,51].

Facilitation Dashboard

The facilitation dashboard provides two primary systems for comprehension of the evolving solution space. First, the dashboard's left panel includes a live updated list of all ideas submitted by the crowd (Fig. 2, left panel). Facilitators can explore ideas using keyword search, sorting by various attributes of the ideas (e.g, submission time, alphabetical order), and bookmarking notable ideas using the "thumbs-up" feature. Second, in the center of the interface, a word cloud derived from the ideas submitted (removing common stopwords and words present in the brainstorming prompt) provides keyword insights into the idea pool, allowing facilitators to search for high-level trends as well as surprising submissions. Keywords are sized by frequency, and clicking any word triggers a search for that word in the idea list. Some basic summary statistics are provided on the top to give a high-level sense of activity levels, but the system emphasizes semantics to maintain focus on the primary task of guiding exploration of the solution space.

On the right side of the interface is the inspiration panel. Facilitators enter new inspirations into the textbox as freeform text messages. The freeform text format is designed to encourage reflection on higher-level themes, rather than simply curating and distributing examples. The inspiration panel also allows facilitators to monitor the effects of their inspirations. The panel lists each inspiration and includes counts for how many ideas have been submitted as well as a list of all submitted ideas for each inspiration (accessible through an expand/collapse feature).

Crowd Recruitment

For this research, IdeaGens also includes an interface with MTurk through the open-source LegionTools framework [31]. LegionTools enables requesters to queue crowd workers into a retainer and then simultaneously launch the crowd into a task. Workers are paid for waiting in the retainer, though they can perform other tasks during the waiting time, and they are given a bonus for completing the assigned task. Using the interface to LegionTools, innovators can assemble crowds of varying sizes to work on a brainstorming problem at the same time.

EXPERIMENT 1: EXPERIENCED FACILITATORS

To evaluate IdeaGens, we conducted a quantitative controlled experiment with in-person facilitators and online crowd workers. In this initial evaluation, we sought experienced facilitators (i.e., individuals with significant expertise in facilitating creative idea generation). We hypothesize that, relative to unfacilitated workers, facilitated workers will produce ideas with both greater divergence (higher fluency and broader search) and improved convergence (deeper search and more creative ideas overall).

Method

Participants

This study leverages two populations of participants to serve the two roles in the study: 1) *facilitators* and 2) *ideators*. Two experienced facilitators were recruited, each participating in a separate experimental trial (with a different crowd). SF (male, 28 years old) is experienced with leading idea generation sessions, and an expert in the literature on effective creative idea generation. JC (female, 37 years old) is a game designer with 10 years of experience leading group brainstorms. We recruited two facilitators to reduce the probability that positive effects would be due to idiosyncrasies of a single facilitator.

113 workers (44% female, mean age = 33.23 [$SD = 11.76$]) from MTurk were recruited to participate as crowd ideators for the two trials. We restricted recruitment to US workers with at least 80% approval rate. Ideators were paid \$1.50 (\$0.50 for waiting, \$1.00 for participation in the study).

Study Design

The evaluation was conducted as a single-factor between-subjects experiment. In each trial, recruited workers were randomly assigned to one of two conditions:

- 1) In the **Facilitated** condition, ideators generated ideas using the ideator interface as depicted in Figure 1 (i.e., with facilitation through IdeaGens)
- 2) In the **Unfacilitated** condition, ideators generated ideas in a modified ideator interface that removed the "Inspire Me" button. These ideators received no facilitation; their ideas were also not fed to the facilitator dashboard (so that their ideas would not provide an unfair advantage to the **Facilitated** ideators).

Running Facilitated and Unfacilitated ideators in the same trial (rather than collecting data from Unfacilitated ideators

separately) helps reduce the risk of confounds (e.g., differences in time of day, etc.). Eighty seven of 113 recruited ideators completed all study procedures and generated at least 2 ideas. These 87 ideators constitute the sample for this study. There were no differences in attrition between the Facilitated ($N=46$, 27% attrition) and Unfacilitated ($N=41$, 19% attrition) conditions, Z test for difference in proportions = -0.83 , $p = 0.41$.

Brainstorming Prompt

Participants generated ideas for the following problem: *“Imagine you are in a social setting and you’ve forgotten the name of someone you know. How might you recall their name without directly asking them? Be as specific as possible in your descriptions.”*

Two key properties of the problem make it a suitable choice for this study. First, because the problem is a common social predicament, both facilitators and MTurk workers likely have sufficient expertise and interest to generate interesting ideas, maximizing the probability that we would be able to observe authentic creative phenomena in our experiments. Second, unlike many classic brainstorming problems, which also have low requirements for prior knowledge (e.g., alternative uses for a brick), this problem has articulable dimensions of both novelty and value. That is, there are obvious (e.g., “ask someone else”) and highly unusual solutions (e.g., “ask them what their full tombstone inscription would say”) to the problem, but also solution approaches that are more (e.g., “ask them for their phone number and give them your phone to put it in”) or less likely (e.g., “Talk about random names and then maybe you’ll guess their name”) to succeed in getting someone’s name. This allows us to distinctly observe divergence (e.g., added novelty) and convergence (e.g., elaborating, increasing value) and how the manipulation affects each characteristic separately. This is especially important because our system is designed to help balance divergence and convergence. Further, many classic brainstorming prompts have been compiled into online collections, so we wanted to ensure people were not simply yielding long lists of answers from an Internet search.

Procedure

The facilitators went through the following procedure. After obtaining informed consent, the experimenter explained the overall task to the facilitator, noting that their main goal was to help a group of ideators come up with the most creative ideas possible for a brainstorming problem. The experimenter further explained that the primary mechanism for achieving this main goal would be to create inspirations (i.e., thought-provoking questions, insights, or themes drawn from brainstormers’ ideas or their own thoughts). The facilitator then completed a brief (10 min) tutorial of IdeaGens before facilitating the crowd brainstorm (20 min). Finally, the experimenter conducted a semi-structured interview with the facilitator, focusing on understanding the

facilitators’ rationale and strategies for creating various inspirations.

Ideators went through a different procedure in parallel with the facilitators. Once launched into IdeaGens from Legion-Tools, ideators provided informed consent and were randomly assigned to condition. The ideators then completed a self-paced tutorial. Integrated into the tutorial was a baseline fluency task, where ideators were given 1 minute to produce as many alternative uses of a bowling pin as possible (more details below). After completing the tutorial, ideators were automatically launched into the brainstorm. After 10 minutes, ideators were automatically directed to a short survey with demographics information and questions about their experiences during the brainstorm.

Measures

We measure key aspects of participants’ *divergence* (fluency and breadth of search), *convergence* (depth of search), and creative *outcomes* (rated creativity of ideas).

Fluency: Number of Ideas

We removed ideas that were either incomplete (and therefore unintelligible; e.g., “ask how to”) or in clear violation of the stated constraints of the problem (e.g., proposing to ask the person directly: “Just ask the person again”).

Creativity: Combination of Novelty and Value

Creativity was operationalized as the product of novelty and value scores for each idea. Taking the product rather than the sum of the two scales places higher weight on ideas that are high on *both* novelty and value, and captures the theoretical intuition that highly novel, useless ideas, and highly obvious, valuable ideas are not creative [50]. Novelty is the degree to which an idea surprised a judge. Value is the estimated likelihood that the idea would work (i.e., recover the person’s name), given that it was actually implemented as stated.

Two trained judges (both graduate researchers) exhaustively evaluated all non-redundant ideas for novelty and value, providing ratings on a 1 (worst) to 7 (best) Likert scale. The judges had appropriate domain knowledge for this task given that it addressed a common social predicament. Judges were blind to experiment condition during the rating task. To ensure internal consistency, judges sorted ideas by rating and time of rating after completing each set of 100 ratings, adjusting earlier ratings if necessary. Inter-rater reliability was acceptable for both scales, at $r = .72$ for novelty, and $r = .65$ for value. All disagreements greater than 2 points on the scale were resolved through discussion. All ideas’ final novelty and value scores were then computed by averaging the ratings from both judges; creativity scores were the product of the novelty and value scores.

For brevity, we only report analyses of the creativity scores; however, the pattern of effects for novelty and value separately are substantially similar (with slightly stronger trends for novelty). Our results are also robust to an additive formulation of the creativity combination function (i.e., sum




Figure 3. Ideators facilitated by experienced facilitators generated more ideas (left panel) and had higher max creativity scores (right panel). Error bars are ± 1 standard error.

of novelty and value). In our analyses we consider both *mean* (how creative are participant's ideas, on average) and *max* (what is the highest creativity score attained by each participant) creativity.

An example of a *low* creativity idea is “*think really hard*” (Novelty = 1, Value = 2.5, Creativity = 2.5); an example of a *high* creativity idea is “*ask for their best impression of their mom yelling their first, middle, and last name as a kid*” (Novelty = 6.5, Value = 6, Creativity = 39).

Breadth and Depth of Search of Solution Space

To characterize the structure of the solution space, we trained a semantic model of the set of ideas using Latent Semantic Analysis (LSA [30]), which estimates a high-dimensional semantic space representation of a corpus of documents based on word co-occurrence patterns. LSA is widely used in creativity and design research to characterize the semantics of ideation, particularly diversity of ideas [2,16,45]. Experiments 1 and 2 in this paper yielded 2,425 ideas. Since the accuracy of vector space models can be significantly improved by increasing the size of the training corpus, we combined the 2,425 ideas from our study with 2,307 raw ideas on the same problem, collected in a separate study [26], yielding a training corpus with 4,732 ideas. In [26], 59 MTurk workers (all U.S. residents) generated ideas for the same problem under similar conditions to our unfacilitated participants (i.e., with no external stimulation). The main difference was that, rather than setting a time limit, participants were given varying instructions regarding the target number of ideas (e.g., 5, 10, 25, etc.). This prior study aimed to quantitatively model the temporal and semantic dynamics (e.g., semantic clustering of temporally adjacent ideas) of unconstrained brainstorming. The authors of the study shared their raw data with us.

Depth was operationalized as the maximum pairwise similarity between a given participant's ideas. Higher maximum similarity indicates a higher probability that at least one of the participant's ideas is a close variation/iteration of another of his/her own ideas. Pairwise similarity between ideas was the cosine between their

semantic vectors in the LSA space, yielding scores between 0 (semantically very different) to 1 (semantically identical).

Breadth was operationalized as the mean pairwise distance between a given participant's ideas. Higher mean pairwise distance indicates that participants' ideas are sampled from very diverse regions of the solution space. Distances were calculated by subtracting pairwise cosines from 1, yielding distance scores between 0 (semantically identical) and 1 (semantically very different).

Control Measure: Baseline Fluency

Our primary control measure is participants' performance on the baseline fluency task (i.e., number of bowling pin alternative uses generated). The task is meant to measure participants' base level of creative fluency (as a proxy for individual creativity), but also likely reflects familiarity with the interface and motivation, among other factors. All of these attributes are expected to influence work quality; therefore, we account for them in our statistical analyses by including baseline fluency as a covariate predictor in our statistical models.

Results

The two facilitators generated 30 inspirations in total. Each ideator received an average of 6.3 inspirations ($SD=3.5$). One example of an inspiration was “*How might you involve technology?*” which sparked ideas like “*Give them your phone and ask them to put their number in*”, and “*get their email address*”. Ideators across conditions generated 1,144 valid ideas (52 invalid ideas removed).

To statistically evaluate the effects of facilitation on the dependent measures, we estimate analysis of covariance (ANCOVA) models for each dependent measure, predicting performance on that measure for each participant as a function of baseline fluency, experimental trial, and experimental condition. Reported means and standard errors are model-adjusted (i.e., controlling for baseline fluency and averaged across experimental trials).

Facilitated Ideators Generated More Ideas

Facilitated ideators generated significantly more ideas ($M=12.9$ ideas/person, $SE=0.7$) than Unfacilitated ideators ($M=10.2$, $SE=0.8$), $F(1,83) = 6.4$, $p = .01$ (see Fig. 3, left).

Facilitated Ideators Generated More Creative Ideas

There were no differences between conditions on mean creativity, $F(1,83)=0.32$, $p=0.57$. However, Facilitated ideators did have significantly higher max creativity scores ($M=25.7$, $SE=1.1$), compared to Unfacilitated ideators ($M=22.1$, $SE=1.2$), $F(1,83)=4.8$, $p=0.03$ (see Fig. 3, right).

Facilitated Ideators Searched More Deeply and Equally Broadly in the Solution Space

For facilitated ideators, the maximum LSA-estimated pairwise similarity between ideas was marginally significantly higher ($M=0.74$, $SE=0.0$) than Unfacilitated ideators ($M=0.67$, $SE=0.0$), $F(1,83)=3.2$, $p=0.08$, suggesting that ideators were more likely to produce variations/iterations of




Figure 4. Ideators facilitated by experienced facilitators searched more deeply (left panel) and equally broadly (right panel) in the solution space.

ideas (i.e., increased depth of search) when provided with inspirations (see Fig. 4, left).

The diversity of ideas was equivalent across conditions, $F(1,83)=0.6$, $p=0.45$, indicating that the enhanced depth of search did not preclude breadth of search (see Fig. 4, right).

Discussion

Overall, Experiment 1 demonstrated IdeaGens' effectiveness at improving crowd ideation work quality. Facilitation with IdeaGens increased ideators' quantity, novelty, and creativity of ideas compared to unfacilitated ideators. The increased depth of search with facilitation aligns with our intuitions about how expert guidance might improve work quality; indeed, depth of search was significantly correlated with both mean creativity ($r=0.3$, $p=0.01$), and max creativity ($r=0.3$, $p=0.01$). Importantly, boosts in depth of search did not come at the expense of diversity of ideation.

Evaluating IdeaGens with experienced facilitators constitutes a fair test of its value, since it is designed to enable expert facilitation. However, this methodological choice leaves open the question of precisely what value is provided by an experienced facilitator. What facilitation strategies do experts employ and how do they affect ideators? Could less experienced facilitators provide comparable benefits? To gain insight into these questions, we ran a second controlled experiment with novice facilitators.

EXPERIMENT 2: NOVICE FACILITATORS

Method

Three members of the university community were recruited as facilitators. All facilitators have at least some prior experience with brainstorming, but no significant expertise with leading brainstorms. MB (Female, 58 years old) works in education programming for a local museum. MM (male, 40 years old) is a freelance mascot performer who collaborates with a local music band. NM (female, 20 years old) is a biology student at a private research university.

We recruited 137 MTurk workers (53% female, mean age = 33.23 [$SD=11.8$]) to participate as crowd ideators for the three different trials. Recruitment restrictions were the same

as in Experiment 1, except that we also added a restriction to bar repeat participation from workers who had participated in Experiment 1. We obtained valid data (i.e., completed all study procedures and generated at least 2 ideas) from 85 of the 137 recruited ideators. There were no differences in attrition between the Facilitated ($N=41$, 40% attrition) and Unfacilitated ($N=44$, 36% attrition) conditions, z test for difference in proportions = 0.33, $p=0.74$.

Apart from the participants, all methods (i.e., design, task, procedure, and measures) were identical to Experiment 1.

Results

Facilitators generated 35 inspirations in total. Each ideator received an average of 5.8 inspirations ($SD=5.0$). One example of an inspiration was "What other personal questions might lead to their name?", which sparked ideas like "ask if they were named after anyone", and "ask if they have met anyone with their name." Ideators across conditions generated 1,166 valid ideas (58 invalid ideas removed).

Novice Facilitators Did Not Increase Ideators' Fluency

Facilitated ideators did not generate significantly more ideas ($M=11.5$, $SE=0.7$) than Unfacilitated ideators ($M=10.9$, $SE=0.7$), $F(1,80)=0.34$, $p=0.56$.

Novice Facilitators Reduced Ideators' Creative Output

Facilitated ideators had significantly lower mean creativity scores ($M=10.8$, $SE=0.4$) compared to Unfacilitated ideators ($M=12.4$, $SE=0.4$), $F(1,80)=7.9$, $p=0.01$. Similarly, Facilitated ideators had marginally significantly lower max creativity scores ($M=18.8$, $SE=1.0$), compared to Unfacilitated ideators ($M=21.4$, $SE=1.0$), $F(1,80)=3.8$, $p=0.05$.

Novice Facilitators Did Not Influence Ideators' Breadth and Depth of Search

There were no significant differences between conditions on either depth (maximum pairwise similarity, $F(1,80)=0.5$, $p=0.49$), or breadth of search (diversity of ideas, $F(1,80)=1.0$, $p=0.32$).

Discussion

In summary, in contrast to Experiment 1, ideators did not benefit when guided by inexperienced facilitators; rather, facilitated ideators generated less valuable and creative ideas than unfacilitated ideators.

These results strengthen inferences from Experiment 1. For example, one might conclude from Experiment 1 that workers simply tried harder because they knew someone was paying attention. However, the opposing pattern of results between Experiment 1 and Experiment 2 helps to rule out this explanation. The results also suggest that experienced facilitators might provide value beyond simply curating examples: if mere exposure to external input was sufficient to benefit ideation, we would expect to have observed, at worst, a muted benefit of facilitation in Experiment 2. What were the experienced facilitators doing that worked, and how did novice facilitators manage to harm (not just fail to improve) overall creativity?




Figure 5. Inspirations from experienced facilitators yielded more ideas per ideator (left panel) and more creative ideas (right panel) than inspirations from novice facilitators.

ANALYSIS OF INSPIRATIONS

Experienced facilitators generated slightly more inspirations ($M=15.0$) than novice facilitators ($M=11.7$), respectively), and ideators received comparable numbers of inspirations each across the two experiments (6.3 vs 5.8). Thus, it seems unlikely that raw numbers of inspirations would explain the difference in effects. To gain more insight into the differences in effects between the experienced and novice facilitators, we conducted an exploratory analysis of the inspirations that were created by both expert and novice facilitators.

Ideators Were More Likely to Rate Experienced Facilitators' Inspirations as Helpful

We first consider ideators' reflections on the inspirations received (as measured in their post-brainstorm survey responses). Examining ideators' response to the question "Did you find any of the suggested inspirations helpful?", a significantly higher proportion of ideators facilitated by experienced facilitators said that they found inspirations to be helpful (91%), compared to ideators facilitated by novice facilitators (70%), $z=2.4$, $p=0.01$.

One interesting theme in ideators' open-ended comments was that helpful inspirations provoked new frames of thinking about the problem, e.g., "It helped me realize an angle for problem solving that I had not considered". Approximately 4% of comments for novice facilitators and ~18% of comments for experienced facilitators had this theme. One ideator's comment suggested that novice facilitators' inspirations could be causing fixation, stating "I kept seeing the inspiration and then all I could think of is what was already suggested."

Inspirations from Experienced Facilitators Led to Higher Fluency and Creativity

The ideators' comments suggest that experienced facilitators might be more likely to invoke new problem solving angles (not just spark individual ideas) and that novice facilitators might fixate ideators. But does this reflect actual differences in inspiration outcomes between experienced and novice facilitators? To explore this, we analyzed the

ideas yielded by each inspiration. Since ideators were instructed to enter ideas sparked by an inspiration directly into the inspiration's idea entry box, we associated all such ideas with their inspiration "parent" for this set of analyses. In line with our core findings concerning number of ideas and max creativity, we considered two key inspiration outcomes: 1) *yield*, i.e., number of ideas yielded per ideator who saw that inspiration, and 2) *max creativity* of ideas.

The results mirror the ideator-level outcomes (see Fig. 5). Inspirations from experienced facilitators had significantly higher yield ($M=0.9$, $SE=0.1$) than novice facilitators ($M=0.5$, $SE=0.1$), $F(1,63)=15.2$, $p=0.00$. Experienced facilitators' inspirations also yielded higher max creativity of ideas ($M=20.6$, $SE=1.2$) than novice facilitators ($M=12.4$, $SE=1.1$), $F(1,62)=24.3$, $p=0.00$.

Inspiration Strategies

Why were experienced facilitators' inspirations more successful with ideators? Experienced and novice facilitators could have employed different strategies to create inspirations. For example, one simple strategy would be to curate examples for distribution. More advanced strategies are also possible, e.g., highlighting high-level solution themes or provoking rich mental simulations of new scenarios.

We developed a coding scheme for inspiration strategies through an open-coding approach, iteratively abstracting the most common emerging themes from the inspirations. Table 1 shows the final strategy-centered coding scheme.

Inspirations were coded for the presence/absence of each strategy. Note that these strategies are not mutually exclusive: inspirations could combine multiple strategies. For example, "People's names are often on objects (e.g., driver's license). What other objects might their names be on? How might you exploit this fact?" employs both an example (i.e., provides the idea to "look for name on a driver's license") and an inquiry (i.e., "how might you...").

Two researchers independently coded 10% of the inspirations to estimate reliability of the coding scheme. Inter-rater agreement was acceptable to high for all strategies: Cohen's kappa=0.61 for examples, 0.74 for simulation, and 1.0 for inquiry. One author coded the remainder of the inspirations.

Simulations Led to More Creative Ideas

To test the effects of inspiration strategy, we estimated separate linear regressions of the two ideation outcomes (number and creativity of ideas sparked by the inspiration) on each of the strategies, controlling for facilitator experience. Table 1 shows how different inspiration strategies influenced ideation. The **examples** strategy had no effect on either yield, $F(1,61)=1.0$, $p=0.32$, or max creativity, $F(1,60)=0.4$, $p=0.54$. Likewise, **inquiries** had no significant effect on either yield, $F(1,61)=1.2$, $p=0.28$, or max creativity, $F(1,60)=0.5$, $p=0.49$.

In contrast, **simulations** were associated with higher max creativity of ideas ($M=19.5$, $SE=3.5$) compared to non-

Strategy	Description	Sample Inspiration with Strategy	Yield	Max creativity
Examples	Directly provide an idea	“Ask them to put their contact info in your phone”	+0.2	+1.8
Simulations	Invite ideators to generate ideas from a different perspective (e.g., from a different “persona” or specific situation/setting).	“Imagine if you had a different persona (e.g., a politician collecting signatures). What strategies might be available to you?”	+0.3	+8.2 **
Inquiries	Provoke open-ended reflection	“Where might their name be written?”	+0.2	-2.3

^m $p < .10$ * $p < .05$ ** $p < .01$,

Table 1. Strategies observed in facilitator inspirations. Simulations and inquiry led to higher yield. Simulations also led to higher creativity. Cell values are model estimates for mean difference vs absence of strategy, controlling for facilitator experience.

simulations ($M=15.3$, $SE=0.8$), $F(1,60)=8.1$, $p=0.01$.. However, simulations had no significant effect on yield, $F(1,61)=2.6$, $p=0.11$.

Experienced Facilitators Used Simulations More Often

Table 2 shows the distribution of inspiration strategies across experienced and novice facilitators. Experienced facilitators used significantly more inquiries ($z=4.1$, $p=0.00$) and simulations ($z=2.5$, $p=0.01$) than novice facilitators. Novice facilitators relied heavily on providing examples, significantly more often than experienced facilitators ($z=4.4$, $p=0.00$). These overall differences were mirrored at the ideator level. Ideators received fewer examples from experienced ($M=5.3$, $SD=3.5$) vs. novice facilitators ($M=8.5$, $SD=6.6$). In contrast, ideators received more inquiries from experienced ($M=6.3$, $SD=6.0$) vs. novice facilitators ($M=1.3$, $SD=2.0$), and more simulations from experienced ($M=3.0$, $SD=3.0$) vs. novice facilitators ($M=0.4$, $SD=1.0$). The large discrepancy in example and inquiry use between experienced and novice facilitators makes it difficult to interpret the null findings for example and inquiry ideation outcomes: our analyses might have been statistically underpowered because there were so few “non-examples” and inquiries with novices. However, the difference in simulations suggests that the opposing effects of IdeaGens for experienced vs. novice facilitators could be at least partially explained by differential use of the simulation inspiration strategy.

Strategy	Proportion of inspirations ^a	
	Experienced	Novice
Examples	0.37	0.89 ***
Simulations	0.23 **	0.03
Inquiries	0.63 ***	0.14

^aProportions do not sum to 1 because strategies are not mutually exclusive; ** $p < .01$, *** $p < .001$

Table 2. Experienced facilitators used more inquiry and provoked simulations more often than novice facilitators. Novice facilitators relied heavily on providing examples, more often than experienced facilitators. Significant comparisons are marked for each row.

Experienced Facilitators Were More Intentional About Tailoring Inspirations to the Ideators

Facilitators’ open-ended reflections (and observations of their behaviors during the brainstorm) suggest additional insights into successful facilitation. One key theme was that experienced facilitators appeared to be more intentional about tailoring and responding to ideators’ ideas. For example, JC reflected on her inspiration-making mid-session: “I feel like at this point I had hit the categories I could think of, the larger questions I could think of from what I was seeing, so I was trying to figure out what would be a more targeted question, but I didn’t want to be too targeted.” This revealed attention to the evolving structure of the solution space (at levels of abstraction), and a focus on tailoring the content and framing of the question to the ideators. Experienced facilitators were also observed more often checking what ideas came of their inspirations in order to gauge their effectiveness.

In contrast, 2 of the 3 novice facilitators created more than half of their inspirations before paying attention to any of the ideators’ ideas. In their reflections, they noted that they were mostly trying to think of ways that one could recall forgotten names, essentially serving as additional ideators, not facilitators. Novice facilitators would often verbally note that an idea was “interesting”, but fail to create an inspiration from it. For example, NM bookmarked an idea “start a game naming famous people who have the same name as yourself”, but did not convert it into an inspiration, noting, “I thought it was good for a specific setting...but I thought, if I was meeting someone on the street, and I’m trying to remember their name, it’s not probably the best idea.” This suggests novice facilitators perhaps overly focused on evaluating ideas, rather than extracting solution themes (e.g., games) and responding to the emerging solution space as a whole.

GENERAL DISCUSSION

Summary and Interpretation of Findings

This research adapted and evaluated a strategy for improving crowd innovation using real-time facilitation (in the form of inspirations). In our experiments, experienced facilitators increased the fluency, creativity, and convergence of crowd ideators (compared to unfacilitated idea-

tors) without sacrificing divergence; in contrast, novice facilitators negatively impacted the creativity of ideas and failed to help ideators converge.

The contrast between experienced and novice facilitators helps rule out plausible alternative explanations. First, the benefits of facilitation cannot be explained in terms of social facilitation [6], i.e., increased effort due to “being watched”: otherwise, novice facilitation would also have an advantage over no facilitation. Similarly, facilitation benefits cannot stem solely from mere exposure to additional stimulation (ideation prompts); if this were the case, novice-facilitated ideators would also have an advantage over non-facilitated ideators, and there would be no meaningful differences in the nature and impact of experienced vs. novice facilitators’ inspirations. Indeed, follow-up analyses of the inspirations revealed that experienced facilitators’ inspirations sparked more creative ideas, in part because they more frequently promoted mental simulations (a strategy that was correlated with more creative ideas).

This research provides evidence that real-time facilitation can positively influence crowd ideation, and (perhaps more importantly) uncovers evidence of *how* to best facilitate crowd ideation (e.g., encouraging more advanced inspiration strategies, such as provoking mental simulations).

Limitations

There are potential concerns about generalizability of our findings due to the chosen problem for our experiments. Brainstorming for a common social predicament is significantly simpler than brainstorming for high-impact design problems that are typically the focus of crowd innovation efforts (e.g., addressing climate change, inventing products). Some theoretical analyses suggest that, for extremely complex problems (e.g., highly challenging R&D problems addressed by Innocentive), it may be counterproductive to introduce inspiration-guided convergence [7,15]. In such problem spaces, theorists argue that the optimal strategy is to have as many ideators as possible explore the solution space independently. This minimizes the probability of the crowd’s search getting stuck in local optima and missing more high-impact globally optimal solutions. Future research should empirically test how the value of expert facilitation might vary as a function of problem complexity.

Regarding our analyses of inspiration strategies, there are reasons to expect that the benefits of the simulation strategy observed in our data are not idiosyncratic for our chosen brainstorming task. Human-centered designers employ a closely related strategy—persona-based ideation—to stimulate more creative thought during initial concept generation [37]. The inquiry strategy identified in our data might also prove useful in other brainstorming settings. While we did not observe direct evidence of their impact on number or creativity of ideas (in part due to lack of statistical power), we did observe that experienced facilitators used inquiry more often than novice facilitators. Further, the inquiry strategy is also reminiscent of “*How can we...*”

questions used often in design practice. On balance, we recommend that the present study be regarded as a proof of concept for the value of real-time expert facilitation. Future studies are needed to verify whether and how our findings generalize to other ideation settings.

Separately, some might be concerned that we found benefits for *max* creativity but not *mean* creativity. However, arguably, innovators care more about increasing the number of exceptional ideas, rather than simply raising the average creativity of ideas. If facilitation helps ideators function at their maximum creative potential (a natural interpretation of the max creativity results), we would expect to see a higher number of exceptionally creative ideas in the facilitated condition. Indeed, in Experiment 1, the facilitated crowd yielded almost twice as many exceptional ideas (91 ideas with creativity rating greater than 1 standard deviation above the mean, or a creativity score of 19 or greater) compared to the unfacilitated crowd (56 exceptional ideas).

FUTURE WORK

Understanding The Pitfalls of Novice Facilitation

While our inspiration analyses partially explain the negative impact of novice facilitators (e.g., less use of simulations), more research is necessary to fully explore the reasons for this negative impact. For example, perhaps ideators were frustrated at being promised “inspirations” but not finding inspiration. Alternatively, perhaps novice facilitators simply focus on suboptimal parts of the solution space. Uncovering these reasons might point to further ways to ensure effective facilitation.

Improved Tools for Facilitation

Opportunities exist for exploring how to improve support for facilitation. For example, peer-review systems have successfully employed scaffolding techniques, such as pre-authored templates to improve feedback at scale [27,34]. Similarly, research on group decision-making has also developed templates for novice facilitators [8]. We believe these techniques could be useful for helping novice facilitators achieve greater success with crowd innovation.

Additionally, we made a key design decision to avoid burdening the facilitator with the task of manually distributing inspirations. Our simple queue-based mechanism seemed to work well enough, but ideators did comment relatively frequently that they received inspirations describing ideas they had already contributed. Future work might explore how to enable more personalized inspirations (e.g., inspirations that relate to previous ideas so as to expand their thinking, but not so unrelated as to cause process losses due to task switching [40]).

We also envision improvements to the monitoring aspect of the system. Qualitatively, the facilitators appeared to benefit from the word cloud, using it extensively to explore high-level themes as well as surprising ideas. Future work could explore the potential value of more sophisticated representations of the solution space. For example, perhaps

natural-language processing techniques (e.g., TF-IDF) could be used to identify semantically unique ideas so that facilitators could promote greater breadth of search. Also, one facilitator noted that it would be useful to see different “threads” of thought sparked by a particular inspiration (both to get feedback on the inspiration as well as create new inspirations). Such semantics might be best obtained through parallel human-powered analysis of the ideas [11].

Expanding the Facilitator’s Toolkit

In this research, we focused on enabling crowd facilitation through inspirations. It would be interesting to explore a broader range of facilitator interventions that might improve work quality in a crowd context. For example, facilitators could employ “meta-inspirations” that embody domain-specific ideation heuristics (e.g., in product design, “create modularity” [55]), or domain-general creativity techniques like reversing assumptions. It might also be useful to explore how process facilitation might work with online crowds: for example, ideators might benefit from encouragement to engage in extended effort.

We also noticed some hints of differing approaches between the experienced facilitators. When examining novelty and value separately, SF’s facilitated ideators had higher value but only slightly higher novelty of ideas; by contrast, JC’s facilitated ideators had higher novelty but approximately equivalent value of ideas. Both, however, produced more creative ideas. These distinct patterns point to two alternative approaches to improving creativity: not only promoting convergence in search of higher value ideas, but also promoting novelty through iteration (an oft-neglected approach discussed in prior work [39]). It would be valuable to analyze in more detail what facilitation strategies are more beneficial for either approach.

Integrating Expert Facilitation With Other Strategies and Models of Crowd Innovation

Finally, it would be fruitful to explore how expert facilitation might integrate with other strategies and models of effective crowd innovation. For example, in this research crowd workers brainstormed synchronously, but we believe our model of facilitation through inspirations could work equally well in asynchronous settings (which are common in crowd innovation). In mature innovation communities, such as OpenIDEO.com and Climate CoLab, senior members of the community could also serve as expert facilitators, increasing the scalability of the method.

CONCLUSION

This research explores how expert facilitation might be applied to improve work quality in crowd innovation. We embodied the strategy of expert facilitation through inspirations in IdeaGens, and demonstrated its ability to influence work quality across two quantitative controlled experiments: expert facilitation successfully increased the crowd’s convergence (more iteration on ideas, higher creativity of ideas) without sacrificing divergence (higher quantity of ideas, equivalent diversity of ideas). Content analyses of

facilitators’ inspirations underscore that IdeaGens’ benefits stem from the value added by the expert facilitator, and help to define a road map for effective facilitation of crowd ideation.

ACKNOWLEDGEMENTS

This research was supported by National Science Foundation grants #1208382, #1217096, and #1122206. We thank Peter Kremer, Lucy Guo, Michael Richardson, Ishan Vashishtha, Sejal Popat, and Angela Liu for assisting with the development of the system, and Aniket Kittur and Lixiu Yu for helpful feedback on early drafts of the paper. We would also like to acknowledge all the MTurk workers who participated in our study.

REFERENCES

1. Alfredo Muñoz Adánez. 2005. Does quantity generate quality? Testing the fundamental principle of brainstorming. *Spanish journal of psychology* 8, 2: 215–220. <http://doi.org/10.1017/S1138741600005096>
2. Alice M. Agogino, Shuang Song, and Jonathan Hey. 2006. Triangulation of Indicators of Successful Student Design Teams. *International Journal of Engineering Education* 22, 3: 617–625.
3. Ricardo Matsumura Araujo. 2013. 99designs: An Analysis of Creative Competition in Crowdsourced Design. *First AAAI Conference on Human Computation and Crowdsourcing*.
4. Barry L. Bayus. 2013. Crowdsourcing new product ideas over time: An analysis of Dell’s Ideastorm community. *Management Science* 59, 1: 226–244. <http://doi.org/10.1287/mnsc.1120.1599>
5. Osvald M. Bjelland and Robert Chapman Wood. 2008. An Inside View of IBM’s Innovation Jam. *MIT Sloan management review* 50, 1: 32–40.
6. Charles F. Bond and Linda J. Titus. 1983. Social facilitation: A meta-analysis of 241 studies. *Psychological Bulletin* 94, 2: 265–292. <http://doi.org/10.1037/0033-2909.94.2.265>
7. Kevin J. Boudreau and Karim R. Lakhani. 2015. Open disclosure of innovations, incentives and follow-on reuse: Theory on processes of cumulative innovation and a field experiment in computational biology. *Research Policy* 44, 1: 4–19. <http://doi.org/10.1016/j.respol.2014.08.001>
8. Robert O. Briggs, Gert-Jan J. De Vreede, and Jay Nunamaker Jr. 2003. Collaboration engineering with ThinkLets to pursue sustained success with group support systems. *J. of Management Information Systems* 19, 4: 31–64.
9. Joel Chan, Katherine Fu, Christian D. Schunn, Jonathan Cagan, Kristin L. Wood, and Kenneth Kotovsky. 2011. On the benefits and pitfalls of analogies for innovative design: Ideation performance based on analogical distance, commonness, and modality of examples. *Journal of Mechanical Design* 133: 081004. <http://doi.org/10.1115/1.4004396>

10. Joel Chan and Christian Schunn. 2015. The impact of analogies on creative concept generation: Lessons from an in vivo study in engineering design. *Cognitive Science* 39, 1: 126–155. <http://doi.org/10.1111/cogs.12127>
11. Lydia B. Chilton, Juho Kim, Paul André, et al. 2014. Frenzy: Collaborative Data Organization for Creating Conference Sessions. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM, 1255–1264. <http://doi.org/10.1145/2556288.2557375>
12. Victoria K. Clawson and Robert P. Bostrom. 1996. Research-driven facilitation training for computer-supported environments. *Group Decision and Negotiation* 5, 1: 7–29. <http://doi.org/10.1007/BF02404174>
13. William H. Cooper, R. Brent Gallupe, Sandra Pollard, and Jana Cadsby. 1998. Some Liberating Effects of Anonymous Electronic Brainstorming. *Small Group Research* 29, 2: 147–178.
14. Steven Dow, Anand Kulkarni, Scott Klemmer, and Björn Hartmann. 2012. Shepherding the crowd yields better work. *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, ACM, 1013–1022. <http://doi.org/10.1145/2145204.2145355>
15. Sanjiv Erat and Vish Krishnan. 2011. Managing Delegated Search Over Design Spaces. *Management Science* 58, 3: 606–623. <http://doi.org/10.1287/mnsc.1110.1418>
16. Adam E. Green, David J. M. Kraemer, Jonathan A. Fugelsang, Jeremy R. Gray, and Kevin N. Dunbar. 2010. Connecting Long Distance: Semantic Distance in Analogical Reasoning Modulates Frontopolar Cortex Activity. *Cerebral Cortex* 20, 1: 70–76. <http://doi.org/10.1093/cercor/bhp081>
17. Raymonde Guindon. 1990. Knowledge exploited by experts during software system design. *International Journal of Man-Machine Studies* 33, 3: 279 – 304. [http://doi.org/10.1016/S0020-7373\(05\)80120-8](http://doi.org/10.1016/S0020-7373(05)80120-8)
18. Beth A. Hennessey and Teresa M. Amabile. 2010. Creativity. *Annual Review of Psychology* 61: 569–98. <http://doi.org/10.1146/annurev.psych.093008.100416>
19. CJ Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Eighth International AAAI Conference on Weblogs and Social Media*.
20. Scott G. Isaksen and John P. Gaulin. 2005. A Reexamination of Brainstorming Research: Implications for Research and Practice. *Gifted Child Quarterly* 49, 4: 315–329. <http://doi.org/10.1177/001698620504900405>
21. Elahe Javadi and Wai-Tat -. T. Fu. 2011. Idea Visibility, Information Diversity, and Idea Integration in Electronic Brainstorming. *Proceedings of the 6th International Conference on Foundations of Augmented Cognition: Directing the Future of Adaptive Systems*, Springer-Verlag, 517–524.
22. Joy Kim, Justin Cheng, and Michael S. Bernstein. 2014. Ensemble: Exploring Complementary Strengths of Leaders and Crowds in Creative Collaboration. *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ACM, 745–755. <http://doi.org/10.1145/2531602.2531638>
23. Joy Kim, Mira Dontcheva, Wilmot Li, Michael S. Bernstein, and Daniela Steinsapir. 2015. Motif: Supporting Novice Creativity Through Expert Patterns. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 1211–1220. <http://doi.org/10.1145/2702123.2702507>
24. Vassilis Kostakos. 2009. Is the Crowd’s Wisdom Biased? A Quantitative Analysis of Three Online Communities. *IEEE Xplore* 4: 251–255. <http://doi.org/10.1109/CSE.2009.491>
25. Thomas J. Kramer, Gerard P. Fleming, and Scott M. Mannis. 2001. Improving Face-To-Face Brainstorming Through Modeling and Facilitation. *Small Group Research* 32, 5: 533–557. <http://doi.org/10.1177/104649640103200502>
26. Filip Krynicki. 2014. Methods and models for quantitative analysis of crowd brainstorming.
27. Chinmay Kulkarni, Koh Pang Wei, Huy Le, et al. 2013. Peer and self assessment in massive online classes. *ACM Transactions on Computer-Human Interaction (TOCHI)* 20, 6: 33.
28. Karim R. Lakhani. 2008. InnoCentive.com. *Harvard Business School Case*, 608–170.
29. Karim R. Lakhani, Anne-Laure Fayard, Natalia Levina, and Stephanie Healy Pokrywa. 2012. *OpenIDEO*. Social Science Research Network, Rochester, NY. Retrieved May 22, 2015 from <http://papers.ssrn.com/abstract=2053435>
30. Thomas K. Landauer, Peter W. Foltz, and Darrell Laham. 1998. An introduction to latent semantic analysis. *Discourse Processes* 25, 2: 259–284.
31. Walter S. Lasecki, Kyle I. Murray, Samuel White, Robert C. Miller, and Jeffrey P. Bigham. 2011. Real-time Crowd Control of Existing Interfaces. *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology*, ACM, 23–32. <http://doi.org/10.1145/2047196.2047200>
32. Abraham S. Luchins. 1942. Mechanization in problem solving: The effect of Einstellung. *Psychological Monographs* 54, 6: i–95. <http://doi.org/10.1037/h0093502>
33. Kurt Luther, Kelly Caine, Kevin Ziegler, and Amy Bruckman. 2010. Why It Works (when It Works): Success Factors in Online Creative Collaboration. *Proceedings of the 16th ACM International Conference on Supporting Group Work*, ACM, 1–10. <http://doi.org/10.1145/1880071.1880073>
34. Kurt Luther, Jari-Lee Tolentino, Wei Wu, et al. 2015. Structuring, Aggregating, and Evaluating Crowdsourced Design Critique. *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, ACM, 473–485. <http://doi.org/10.1145/2675133.2675283>

35. Winter Mason and Siddharth Suri. 2012. Conducting behavioral research on Amazons Mechanical Turk. *Behavior Research Methods* 44, 1: 1–23. <http://doi.org/10.3758/s13428-011-0124-6>
36. Jensen T. Mecca and Michael D. Mumford. 2013. Imitation and Creativity: Beneficial Effects of Propulsion Strategies and Specificity. *The Journal of Creative Behavior*. <http://doi.org/10.1002/jocb.49>
37. Tomasz Miaskiewicz and Kenneth A. Kozar. 2011. Personas and user-centered design: How can personas benefit product design processes? *Design Studies* 32, 5: 417–430. <http://doi.org/10.1016/j.destud.2011.03.003>
38. Tanushree Mitra, C.J. Hutto, and Eric Gilbert. 2015. Comparing Person- and Process-centric Strategies for Obtaining Quality Data on Amazon Mechanical Turk. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, ACM, 1345–1354. <http://doi.org/10.1145/2702123.2702553>
39. Bernard A. Nijstad, Carsten K. W. De Dreu, Eric F. Rietzschel, and Matthijs Baas. 2010. The dual pathway to creativity model: Creative ideation as a function of flexibility and persistence. *European Review of Social Psychology* 21: 34–77. <http://doi.org/10.1080/10463281003765323>
40. Bernard A. Nijstad and Wolfgang Stroebe. 2006. How the group affects the mind: a cognitive model of idea generation in groups. *Personality and Social Psychology Review* 10, 3: 186–213. http://doi.org/10.1207/s15327957pspr1003_1
41. Alex F. Osborn. 1963. *Applied Imagination: Principles and Procedures of Creative Problem Solving*. Charles Scribner's Sons, New York, NY.
42. Nicole L. Oxley, Mary T. Dzindolet, and Paul B. Paulus. 1996. The effects of facilitators on the performance of brainstorming groups. *Journal of Social Behavior & Personality* 11, 4: 633–646.
43. Cheong Ha Park, KyoungHee Son, Joon Hyub Lee, and Seok-Hyung - H. Bae. 2013. Crowd vs. Crowd: Large-scale Cooperative Design Through Open Team Competition. *Proceedings of the 2013 Conference on Computer Supported Cooperative Work*, ACM, 1275–1284. <http://doi.org/10.1145/2441776.2441920>
44. Sidney J. Parnes and Arnold Meadow. 1959. Effects of “brainstorming” instructions on creative problem solving by trained and untrained subjects. *Journal of Educational Psychology* 50, 4: 171–176.
45. Ranjani Prabhakaran, Adam E. Green, and Jeremy R. Gray. 2013. Thin slices of creativity: Using single-word utterances to assess creative cognition. *Behav Res Methods*. <http://doi.org/10.3758/s13428-013-0401-7>
46. Eric F. Rietzschel, Bernard A. Nijstad, and Wolfgang Stroebe. 2007. Relative accessibility of domain knowledge and creativity: The effects of knowledge activation on the quantity and originality of generated ideas. *Journal of Experimental Social Psychology* 43, 6: 933–946. <http://doi.org/10.1016/j.jesp.2006.10.014>
47. Jeffrey Rzeszotarski and Aniket Kittur. 2012. CrowdScape: Interactively Visualizing User Behavior and Output. *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology*, ACM, 55–62. <http://doi.org/10.1145/2380116.2380125>
48. Jami J. Shah, Roger E. Millsap, Jay Woodward, and S. M. Smith. 2012. Applied Tests of Design Skills Part 1: Divergent Thinking. *Journal of Mechanical Design* 134, 2: 021005–021005–10. <http://doi.org/10.1115/1.4005594>
49. Pao Siangliulue, Joel Chan, Kzryzstof Gajos, and Steven P. Dow. 2015. Providing timely examples improves the quantity and quality of generated ideas. *Proceedings of the ACM Conference on Creativity and Cognition*.
50. Dean K. Simonton. 2012. Combinatorial creativity and sightedness: Monte Carlo simulations using three-criterion definitions. *International Journal of Creativity & Problem Solving* 22, 2: 5–17.
51. Ut Na Sio, Kenneth Kotovsky, and Jonathan Cagan. 2015. Fixation or inspiration? A meta-analytic review of the role of examples on design processes. *Design Studies* 39: 70–99. <http://doi.org/10.1016/j.destud.2015.04.004>
52. Robert I. Sutton and A. Hargadon. 1996. Brainstorming groups in context: Effectiveness in a product design firm. *Administrative Science Quarterly* 41: 685–718.
53. Christian Terwiesch and Karl T. Ulrich. 2009. *Innovation tournaments: Creating and selecting exceptional opportunities*. Harvard Business Press, Boston, MA.
54. E. Paul Torrance. 1988. The nature of creativity as manifest in its testing. In *The nature of creativity: Contemporary psychological perspectives*, Robert J. Sternberg (ed.). Cambridge University Press, New York, NY, 43–75.
55. Seda Yilmaz and Colleen M. Seifert. 2011. Creativity through design heuristics: A case study of expert product design. *Design Studies* 32, 4: 384 – 415. <http://doi.org/10.1016/j.destud.2011.01.003>