Bringing the Wisdom of the Crowd to an Individual by Having the Individual Assume Different Roles

Jaime Teevan
Microsoft Research
Redmond, WA USA
teevan@microsoft.com

Lisa Yu
Bosch Research
Pittsburgh, PA USA
lisa.yu@us.bosch.com

ABSTRACT
Groups of people tend to generate more diverse ideas than individuals because each group member brings a different perspective to the table. But while someone working alone can suffer from fixation and have difficulty of thinking outside-the-box, in this paper we show that it is possible to help them think more like a group by asking them to approach a problem from different perspectives. We present a study of 54 crowd workers in which some individual workers were asked to assume the role of various relevant experts while solving a problem. We find that participants who were asked to assume different roles came up with more creative ideas than those who were not. These findings suggest there is an opportunity for problem solving tools to bring the wisdom of the crowd to individuals.

Author Keywords
Creativity; role playing; crowdsourcing; selfsourcing.

ACM Classification Keywords
H.5.m. Info. interfaces and presentation (e.g., HCI): Misc.

INTRODUCTION AND RELATED WORK

Crowdsourcing researchers are familiar with the idea of the *wisdom of the crowd*, whereby information from multiple people is pooled to produce a better outcome than each individual could produce alone [18]. However, working with other people comes at a cost; the crowd, for example, costs money; explaining a task to someone else costs effort; waiting for a reply costs time. Tasks may also contain private information that we do not want to share with others [10], or we may not want to give up ownership of the output produced by a task [13]. As a result, we tend to do most of our tasks ourselves. The challenge is: we are just one person, and tend to fixate on the knowledge and skills we have [6].

Existing idea generation and creativity literature has explored a variety of approaches to help people think diversely [16], ranging from changing the aspects of looking at problems, to manipulating the media and tools used during problem solving, to controlling exposure to information and experience sources. For example, the *six hats* method asks people to wear different metaphorical hats representing different thinking perspective [5, 17]. Switching attentions between ambiguous sketches has been found to increase divergent thinking [20], and recent researchers harness the online crowds as sources of novel ideas [4, 22].

This paper builds on the wisdom of the crowd to ask: *Can we help an individual think like a crowd*? Instead of crowdsourcing a task to multiple people for diverse solutions, it may be possible to selfsource [19] the task to the same person in different contexts. There is some evidence this may be possible for concrete tasks [1, 7, 20]. For example, participants in one study were asked to answer twice to, "What percentage of the world's airports are in the United States?" [21]. When both answers were averaged it produced a more accurate response than either individually, but the accuracy was still much worse than when the two answers were drawn from different participants. This is consistent with work by Ariely et al. [1] that found that averaging multiple guesses from an individual helps somewhat, but not as much as average multiple guesses from different people. In this paper we extend these findings from concrete estimation tasks to more complex, creative tasks. Additionally, instead of simply eliciting multiple answers from an individual, we investigate approaches to bring out a crowd out of an individual by having the same person approach a problem multiple times as if they were a different person with a different mindset each time.

The wisdom of the crowd arises when we aggregate people's diverse, independent perspectives [1]. To create a similar phenomenon within an individual we must work to artificially create diverse, independent experiences for that individual. For example, when Vul and Pashler [21] made participants wait three weeks between providing answers, they found that "the benefit of re-asking yourself the same question rises to 1/3 the value of a second opinion." Asking people to consider alternative [9] or opposite [12] outcomes can help people overcome anchoring effects and produce less biased judgments, and Herzog and Hertwig [8] found that they could help individuals estimate the dates of historical events if they asked them to give a second answer from the opposite perspective of their first. Crowdsourcing research has shown a person's frame of reference impacts how they do a task, with, for example, the previous task influencing performance on the next task [3, 14].

C&C '17, June 27-30, 2017, Singapore, Singapore
 ISBN 978-1-4503-4403-6/17/06…$15.00
DOI: http://dx.doi.org/10.1145/3059454.3059467

We combine crowdsourcing with selfsourcing to help people approach creative problems from multiple perspectives by asking them to imagine assuming the role of various relevant experts. Role playing has been shown to help people experience new perspectives. Boess [2] found it helped students in their design work, and Matthews et al. [11] found it beneficial for therapeutic systems. Recent research by Yu et al. [22] used role playing to support the crowdsourcing of creativity, and found that crowd workers generate more creative ideas when asked to consider a problem from the perspective of an expert in an unexpectedly related domain. We build on these findings to look at how assuming different roles generated by crowd workers can help an individual experience different perspectives. We present a crowd-based study that asks a single person to assume multiple roles while generating creative ideas. By comparing solutions generated in this way with those from a single perspective, we find evidence that individuals produce more creative ideas when primed with roles.

METHODOLOGY

We conducted a crowd-based study to explore whether people generate better ideas by searching for inspiration from the perspectives of multiple out-of-the-box domains than just searching for inspirations on their own.

Idea Collection

To collect ideas, individuals were given a difficult design problem that required a creative solution, and asked to come up with three potential solutions. To assist with this, we asked participants to search the Internet for inspiration. We required that each inspiration be drawn from different domains so that it could not serve as solutions on its own. Instead, participants applied the inspiration directly or indirectly to the problem to generate their own solution.

Conditions

We designed a control condition, which we refer to as the *no-role* condition, where participants searched for inspiration on their own and then used the inspiration to solve the problem by following the instruction below:

> *Please go to the Internet and find three useful ideas that could inspire good solutions for the above problem. These ideas could be knowledge, skills, or methods other people use to solve a similar problem in their own domain. Please don't search for ideas related to this [task] problem. The useful ideas have to be about similar problems in a non-[task] related domain*.

Participants were provided with a form with three sections representing the three useful ideas they were asked to collect. For each idea they were asked to paste a link to the site they found for inspiration and then use that inspiration to generate text describing a solution for the given problem.

We designed the treatment condition to ask individuals to search for inspiration from the perspective of three different roles (the method of generating these roles will be explained in the later section). We refer to this condition as the *multiple-roles* condition. The instructions for read as follows:

> *Please go to the Internet and find three useful ideas a [role 1], [role 2], and [role 3] might have that could inspire good solutions for the above problem. These ideas could be knowledge, skills, or methods a [role 1], [role 2], or [role 3] uses to solve a similar problem in his or her own domain. Please don't search for ideas related to this [task] problem. The useful ideas have to be about problems in the domain that a [role 1], [role 2], or [role 3] deals with.*

As in the *no-role* condition, participants were provided with a form with three sections, each asking for a link and text. In this condition, however, sections were also labeled with one of the three roles they were to be using for inspiration.

Participants

We recruited 54 participants via Mechanical Turk. Seventeen were male, the rest female. Ages ranged from 21 to 61, with a mean of 34. Each participant was randomly assigned to one of the two conditions, with 33 falling in the *no-role* condition, and 21 falling in the *multiple-roles* condition. The discrepancy across conditions occurred naturally. It may be random, as both tasks were presented in similar ways, or more people may have quit mid-task in the *multiple-roles* condition. Each participant was paid $2 to complete the task. Using Mechanical Turk qualifications, we only allowed unique workers who have more than 95% task approval rate to take the task, and we auto-approved all of the submitted ideas.

Creativity Tasks

For the creativity tasks we selected two problems from the set of problems described by Yu et al. [22], one related to redesigning a power strip, and the other related to dry cups. We assigned one of the problems to a participant randomly.

> Power strip problem: *Have a look at the power strip under your desk. How many of its outlets are being used? How many of them would you like to use, but you can't, because a giant power brick (transformer) in the adjacent outlet is blocking it? How could you fit all the different plugs in all the outlets?*

> Cup problem: *When we finish washing cups and glasses, we have to either spread them out individually, but then they take up all the counter space. Alternatively, we can stack them, but then the cups never dry completely and it is hard to separate them from each other later. How can you dry many cups quickly so that they don't take up too much space and moisture doesn't get trapped between them?*

Roles

We likewise used the roles Yu et al. [22] developed for these tasks to represent unexpected potential experts in the problem domain. Each role was identified by crowd workers based on the abstract structures of the task. For example, an abstract structure for the power strip problem is, "How can you fit objects of different sizes into a container?" and this led crowd workers to identify "a warehouse dock loader" as a potential expert. Yu et al. found that when workers used

Condition	Idea 1	Idea 2	Idea 3
No-role	2.50	2.27	2.00
Multiple-roles	3.40	3.40	2.57

Table 1. The average creativity score across ideas as rated by an independent judge. Creativity is the average of the idea's novelty and practicality (based on a 7-point Likert scale).

these roles to generate creative ideas they found better inspiration because the roles led to analogies in different domains.

Given we want to provoke individuals to think outside-the-box, we borrowed the roles from their study. However, instead of asking multiple workers to assume the roles, we asked individual participants to assume three different roles. The order of the roles was randomized. We predicted that by taking on new roles, people would explore different ways of thinking and exceed their existing creative capability. The roles we used in the study included: a warehouse dock loader, a landscaper, a sculptor, a magician, an artist for the power strip problem, and a contortionist for the cup problem.

Idea Evaluation

After each idea was generated we asked participants to self-rate how hard it was for them to come up with the idea ("*How difficult was it for you to generate this idea?*"), and to report how creative they thought the idea was ("*How creative is this idea?*"). Responses were collected using a 7-point Likert scale. At the end of the entire task, participants were also asked to name their favorite idea of the three.

In addition to the self-reported data, we also had an independent judge (a HCI graduate student) evaluate ideas blind to condition. For this we drew on previous research for robustly rating creative idea quality, which considers an idea as being creative if it is both *novel* and *practical* [4, 22]. Novelty was defined for our judge as an idea that was not obvious and differed from existing available solutions. Practicality was defined as how realistically an idea achieved its goal and whether it could be designed and manufactured. The two measurements were provided using 7-point Likert scales and averaged to produce an overall creativity score.

RESULTS

Using the data we collected, we analyze the impact of assuming multiple roles on the creativity of the ideas generated by our participants, as well as on their ability to generate multiple ideas over time. Two example ideas generated for the power-strip problem are shown in Figure 1, the first by a worker in the *no-role* condition, and the second by a worker in the *multiple-roles* condition.

No-role:

> Make the power brick (transformer) smaller or located at a different spot along the power cord so that it is not in the way or interfering with the power outlets.

Multiple-roles (role: an expert on topology):

> For this design I would have it in the shape of the tree, and have a spot for the transformer at the bottom of the tree, and have all the other plugs at the top of the tree.

Figure 1. Example ideas generated by workers in the *no-role* and *multiple-roles* conditions for the power-strip problem.

Condition	Measure	Idea 1	Idea 2	Idea 3
No-role	Creativity	5.30	5.18	4.97
	Difficulty	3.24	3.88	4.21
	Favorite	11	8	14
Multiple-roles	Creativity	5.33	4.81	5.29
	Difficulty	3.81	3.71	3.81
	Favorite	6	5	10

Table 2. The average self-reported creativity and difficulty level across ideas (based on a 7-point Likert scale), and the number of times each idea was listed as favorite.

Measures of Creativity

As described in the previous section, we collected two measures of creativity for the ideas our participants generated: one that relied on the assessment of an independent judge (Table 1), and another based on participants' own self-reported assessment of their ideas (Table 2). In general, we find evidence to suggest that asking people to assume different roles encouraged creative thinking. As reported by one participant in the *multiple-roles* condition, "I tend not to be very creative, but this got my brain working."

According to the ratings provided by an independent judge, the ideas generated in the *multiple-roles* condition were generally more creative than the ideas generated in the *no-role* condition, as can be seen in Table 1. In fact, when considering only ideas with above average creativity scores, we find that 47% of the *multiple-roles* ideas are above-average, as compared with only 34% of the *no-role* ideas. We conducted a T.test analysis on the judged creativity, and found that creativity score averaged across the ideas from the *multiple-roles* condition (M=3.13, SD=1.40) was higher than that of *no-role* condition (M=2.26, SD=1.72), $t(52)=2.02$, $p<.05$. This result suggests that participants were able to generate more creative ideas when they changed their thinking by assuming the roles of other experts.

These differences do not, however, appear when we analyze the self-reported data (Table 2). Here we see no significant difference in the self-reported creativity of the ideas generated from a *no-role* (M=5.15, SD=1.51) as compared with *multiple-roles* (M=5.14, SD=1.54). This could mean that there was no difference. However, it could also be because our participants had a hard time judging the creativity of their ideas without the context of the ideas that other people had generated for the problem. Further, workers were probably biased towards viewing ideas that they came up with as creative, as it allowed them to view their work in a positive light. They may also have thought a high creativity score would make their work more likely to be accepted.

Generating Multiple Ideas

When we look an individual participant's creative problem solving ability over time, we observe that there was a general trend for participants to generate less creative ideas as they were asked to continue to come up with new ones. Idea 3

(M=2.22, SD=1.51), for example, was generally judged by our independent judge to be less creative than Idea 2 (M=2.71, SD=2.05), $t(53)=1.67$, $p<.05$. As might be predicted given what we know about the wisdom of the crowd, this suggests the best way to collect three creative ideas would be to ask three different people to each generate one idea, which is the equivalent of taking the first – and most creative – idea from three of our participants.

When we look at the favorite idea selected by each participants, however, we find they tended to prefer their most recent idea; 24 participants said they liked their third idea best, while only 17 and 13 liked the first and second ideas best, respectively. This could be a result of a recency effect, with participants reflecting a cognitive bias towards the most recent idea they generated. But it suggests that in the case of self-assessed creativity, the wisdom-of-an-individual-assuming-roles may produce a comparable outcome to taking the first idea from three different people assuming the role, as is done in a standard wisdom-of-the-crowd approach.

We expected the roles would make it easier for people to generate multiple ideas, as it can be progressively harder for an individual to come up new ideas while brainstorming alone. We hypothesized that the different roles would help keep participants from getting stuck viewing the problem from a single perspective. We see that there is a trend in the difficulty levels participants reported in generating ideas, as shown in Table 2. We observed that the difficulty of coming up with a new idea seemed to increase over time in the *no-role* condition, while the difficulty seemed staying the same in the *multiple-roles* condition. Further study is needed, however, to confirm whether this trend is true, as a repeated ANOVA analyses did not yield a significant difference. This could be an artifact of the study design, as participants were asked to generate three solutions at one time. Future study could separate the process into three isolated steps and only show each participant one step at a time.

Can We Help an Individual Think Like a Crowd?

Taken collectively, our findings suggest it may be possible to use the crowd to help an individual think like a crowd by having them assume different roles. Individuals appear to be more creative when they assume the roles of multiple other experts during idea generation, and may also have an easier time generating ideas over time. Additionally, when we compare the creativity of the ideas an individual was able to achieve by assuming different roles with the wisdom of the crowd, where three different individuals to solve the problem independently, it appears the approach may begin to help an individual approximate the aggregated output of a crowd.

DISCUSSION

Individuals often develop particular ways of thinking. Research consistently shows that people tend to fixate on a particular perspective and thus have difficulty addressing problems creatively [6]. With the refined disciplinary divisions in modern education, people are becoming increasingly specialized. While deep knowledge and skills in a domain has many advantages, it can also increase fixation. To break fixation, people often choose to solicit opinions from other people or work in groups. Authors, for example, send papers out for peer review, companies run focus groups, and friends ask each other for advice. However, external resources are not always available and, when they are, they come at a cost. In addition, to protect private information or control the intellectual property, people might be reluctant to have others solve their problems for them. In this paper we focused on how to bring multiple perspectives to an individual to help them perform tasks effectively and creatively. We were interested in methods and tools that break fixation and routine.

We explored the use of roles as a way to encourage multiple perspectives through a combination of crowdsourcing and selfsourcing. Our findings suggest that simply by asking people to assume new roles and explore new domains, people can become more creative than would be if they relied solely on their own way of thinking. We found evidence to suggest that different roles supported an individual's creative idea generation and that considering different roles may make it easier to continue generating ideas. Looking forward, we believe there may be many ways to help individuals solve their real-world problems in real-world contexts by thoughtfully combing crowdsourcing and selfsourcing [19].

The work presented in this paper should be interpreted within the context of understanding the limitations to the exploratory study we conducted. Additional raters could provide a more valid external representation of creativity, and a more structured presentation of information to participants during the study could help more tightly control the priming effects. Further, asking participants to think aloud as they generated their ideas or conducting interviews after the experiment could provide insight into how the roles actually changed participants' mindset during problem solving.

However despite the fact that our findings were drawn from an exploratory study, we think these preliminary findings are promising and suggest interesting areas for further research. For example, it may be possible to encourage people to think differently by actively changing their location, the time when information is requested, or their mode of thinking. Instead of asking participants to generate three ideas at once, it might be effective to solicit new ideas every hour.

CONCLUSION

In this paper we proposed supporting creativity by creating internal crowds within an individual by asking people to imagine themselves assuming different roles. We conducted a crowd-based study that demonstrated the potential of having a single person address creative problems by assuming the role of various crowd-generated experts. Our findings suggest that this approach could be used to help bring the wisdom of the crowd to an individual. Our findings can inform the design of tools intended to support selfsourcing personal creativity. Looking forward, we are interested in exploring different ways to change contexts, including location, time, and mode of thinking.

REFERENCES

1. Dan Ariely, Wint Tung Au, Randall H. Bender, David V. Budescu, Christiane B. Dietz, Honging Gu, Thomas S. Allsten, and Gal Zauberman. 2000. The effects of averaging subjective probability estimated between and within judges. *Journal of Experimental Psychology: Applied*. 6 (2): 130-147.
2. Stella U. Boess. 2008. First steps in role playing. In *CHI '08 Extended Abstracts on Human Factors in Computing Systems*, 2017-2024. DOI: 10.1145/1358628.1358632
3. Carrie Cai, Shamsi Iqbal, and Jaime Teevan. 2016. Chain reactions: The impact of order on microtask chains. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2016)*, 3143-3154. DOI: 10.1145/2858036.2858237
4. Joel Chan, Steven Dang, and Steven P. Dow. 2016. Improving crowd innovation with expert facilitation. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW 2016)*, 1223–1235. DOI: 10.1145/2818048.2820023
5. Edward De Bono and Marcela Pandolfo. 1999. *Six thinking hats*. Back Bay Books New York. Retrieved March 20, 2017 from http://capitalquality.org/wp-content/uploads/2012/11/CQI_Six_Hats_LL_Lite1.pdf
6. Kees Dorst and Nigel Cross. 2001. Creativity in the design process: Co-evolution of problem-solution. *Design Studies.* 22(5): 425–437. DOI: 10.1016/S0142-694X(01)00009-6
7. Sharad Goel, Daniel M. Reeves, Duncan J. Watts and David M. Pennock. 2010. Prediction without markets. *Electronic Commerce 2010*.
8. Stefan M. Herzog and Ralph Hertwig. The wisdom of many in one mind: Improving individual judgments with dialectical bootstrapping. *Psychological Science*, 20 (2), 231-237.
9. Edward R. Hirt and Keith D. Markman. 1995. Multiple explanation: A consider-an-alternative strategy for debiasing judgments. *Journal of Personality and Social Psychology,* 69(6): 1069-1086.
10. Walter Lasecki, Jaime Teevan, and Ece Kamar. 2014. Information extraction and manipulation threats in crowd-powered systems. In *Proceedings of the 17th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW 2014)*, 248-256. DOI: 10.1145/2531602.2531733
11. Mark Matthews, Geri Gay, and Gavin Doherty. 2014. Taking part: Role-play in the design of therapeutic systems. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2014)*, 643-652. DOI: 10.1145/2556288.2557103
12. Thomas Mussweiler, Fritz Strack, and Tim Pfeiffer. 2000. Overcoming the inevitable anchoring effect: Considering the opposite compensates for selective accessibility. *Personality and Social Psychology Bulletin*, 26 (9): 1142-1150.
13. Michael Nebeling, Alexandra To, Anhong Guo, Adrian A. de Freitas, Jaime Teevan, Steven Dow, and Jeffrey P. Bigham. WearWrite: Crowd-assisted writing from smartwatches. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2016)*, 3834-3846. DOI: 10.1145/2858036.2858169.
14. Edward Newell and Derek Ruths. 2016. How one microtask affects another. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2016)*, 3155-3166. DOI: 10.1145/2858036.2858490
15. Peter Organisciak, Jaime Teevan, Susan T. Dumais, Robert C. Miller and Adam Kalai. 2014. A crowd of your own: Crowdsourcing for on-demand personalization. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing (HCOMP 2014)*.
16. Mark A. Runco. 2010. Divergent thinking, creativity, and ideation. *The Cambridge handbook of creativity*: 413–446.
17. Robert J. Sternberg and Todd I. Lubart. 1999. The concept of creativity: Prospects and paradigms. *Handbook of creativity* 1: 3–15.
18. James Surowiecki, 2004. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations.* New York: Doubleday, 2004.
19. Jaime Teevan, Daniel Libeling, and Walter Lasecki. 2014. Selfsourcing personal tasks. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI 2014)*, 2527-2532. DOI: 10.1145/2559206.2581181
20. Barbara Tversky and Juliet Y. Chou. 2011. Creativity: Depth and breadth. In *Design Creativity 2010*. Springer, 209–214. Retrieved December 30, 2016 from http://link.springer.com/chapter/10.1007/978-0-85729-224-7_27
21. Edward Vul and Harold Pashler. 2008. Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science*, 19 (7).
22. Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2016. Encouraging "outside-the-box" thinking in crowd innovation through identifying domains of expertise. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW 2016)*, 1214-1222. DOI: 10.1145/2818048.2820025.