

Encouraging “Outside-the-box” Thinking in Crowd Innovation Through Identifying Domains of Expertise

Lixiu Yu, Aniket Kittur, Robert E. Kraut
 Carnegie Mellon University
 5000 Forbes Avenue, Pittsburgh, PA 15213
 {lixuiyu, nkittur, robert.kraut}@cs.cmu.edu

ABSTRACT

People are more creative at solving difficult design problems when they use relevant examples from outside of the problem’s domain as inspirations. However, finding such “outside-the-box” inspirations is difficult, particularly in large idea repositories such as the web, because without guidance people select domains to search based on surface similarity to the problem’s domain. In this paper, we demonstrate an approach in which non-experts identify domains that have the potential to yield useful and non-obvious inspirations for solutions. We report an empirical study demonstrating how crowds can generate domains of expertise and that showing people an abstract representation rather than the original problem helps them identify more distant domains. Crowd workers drawing inspirations from the distant domains produced more creative solutions to the original problem than did those who sought inspiration on their own, or drew inspiration from domains closer to or not sharing structural correspondence with the original problem.

Author Keywords

Crowdsourcing; problem solving; design; idea generation

ACM Classification Keywords

H.5.3 Group and Organization Interfaces

INTRODUCTION

Important innovations and discoveries often come from drawing upon knowledge from domains outside of a target problem as an inspiration to solve the problem [8, 10]. For example, using origami-folding techniques, NASA’s space engineers designed a space array, which can be folded compactly and then deployed while in outer space. The new

design solved the 50-year-old space problem of transporting large objects in narrow rockets [2].

The emergence of large online idea repositories has the potential to radically accelerate innovation by increasing designers’ access to analogical ideas. There now exists an enormous selection of ideas that could spark creative ideas through analogy. For example, InnoCentive (innocentive.com) contains more than 40,000 business, social, policy, scientific, and technical problems and solutions in various domains, hundreds of new product ideas are submitted to Quirky (quirky.com) every day by a pool of over a million inventors, and OpenIDEO has collected hundreds of solutions for a variety of social problems since 2010 (openideo.com). More generally, information available online including scientific literature, patents, webpages, and images and videos represent a treasure trove of potential analogies in different domains. For example, a car mechanic adapted a method for extracting a lost cork from a wine bottle seen in a YouTube video to save a baby stuck in the birth canal, described as the most important innovation in birthing since forceps [17, 22].

However, our ability to process this deluge of information to find and use analogies is severely bottlenecked by individual cognitive limits, as the speed and capacity with which individuals can learn and explore new domains have not kept up with the rapid growth in online information from which analogies can be discovered. Even when they have the appropriate knowledge, people often become fixated on surface-level details that prevent them from retrieving useful analogs from memory or external repositories or applying them for problem solving [9].

Previous work has shown that using crowds to mine online idea repositories for such analogous inspirations can enable search at a scale beyond the individual and has identified ways of reducing fixation by having different workers see different representations of the problem [24, 25]. However, this work has looked at relatively small repositories of ideas, such as Quirky’s, where searchers were only looking through hundreds of ideas. In contrast, the web holds orders of magnitude more potentially useful analogous ideas, not only in explicit idea repositories, such as the US patent

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/10.1145/2818048.2820025>

database, but also in expert-generated content in nearly any domain.

In this paper, we describe and evaluate a two-stage process that enables crowds to search the web for useful and novel design ideas than alternative approaches. In the first stage, crowds identify domains of expertise remote from the initial problem but relevant enough to provide ideas to inspire useful and non-obvious solutions. They are best able to identify relevant but remote domains when the original concrete problem is represented as an abstract problem schema. In the second stage, crowds search in these domains to find inspirational examples they can adapt to solve the original problem. The key insight here is that a rich set of expert-generated ideas, solutions, and skills have already been documented on the web, and non-experts can find and appropriate these resources in the design process even if they did not possess the expertise to generate this knowledge in the first place.

Identifying “Outside-the-box” Domains

Although distant analogies can inspire innovation, systematically identifying them in distant domains is challenging. First, people must identify the domains where relevant inspirational sources might be found. The most useful inspirations often come from distant domains that have little surface similarity with the target problem [11, 18, 21]; for example, in the NASA example, the problem comes from the aerospace field while paper folding is from origami art. Identifying such domains can be challenging for human, because they often become fixated on the surface details of the problem domain. Outside-the-box domains are also challenging for automated search algorithms to identify, since the algorithms also generally search on surface features, leading to same-domain recommendations [3, 16, 20, 26].

Yu et al. [24] demonstrated a way for crowds to find distant domains more effectively by re-representing the problem. A critical step was abstracting surface details of the original problem to reduce people’s fixation on surface features, such as domain-specific vocabulary and objects [7, 15]. For example, in the previous NASA story, space and rockets are surface features while the abstract representation of the problem can be “how to contain more content within a limited space”. Using this abstract representation rather than the literal description of the problem frees people to identify remote areas of knowledge containing useful inspirations for the problem. Yu et al. then showed that crowd workers given the abstract representation of a problem (i.e., the schema) rather than the concrete representation were able to find analogies in more distant domains that were more useful for solving the problem creatively.

However, Yu et al. treated the search process as a black box, with no instructions on how to identify potentially fruitful domains to search within. This approach may have

been successful on the tens to hundreds of ideas easily available on Quirky.com, but may not scale well to the billions of web pages available on the Internet, a subset of which could have a potentially relevant solution to the target problem. The difficulty of winnowing distant domains for useful analogs has posed a crucial challenge to many design-by-analogy approaches, causing “this influential technique to be limited to little more than interesting examples with accompanying direction to simply “try to find analogies.” [14].

Exploring “Outside-the-box” Domains

A second problem is finding useful ideas within the identified domain. An expensive but powerful way to do so would be to have a panel of relevant experts on call from that domain. Indeed major design firms like IDEO are successful in part because they hire employees who are diverse in demographics, education and personal interests and can apply ideas from their areas of expertise to the problems clients bring them [10]. Although retaining experts may be necessary in research-investment intensive industries, like pharmaceuticals, for many more routine problem challenges in engineering and design, experts may have already documents relevant solutions in language accessible to laymen.

Crowdsourcing approaches have been increasingly used to expand the pool that innovation companies can draw upon to solve problems, ranging from product ideation (e.g., Quirky.com) to research and development (e.g., Innocentive.com) to societal challenges (OpenIDEO.com). Companies describe their problems and then use what is called “broadcast search” to invite anyone with a good idea to contribute solutions [6, 13]. Here instead of contributing original ideas, we aim to use non-expert crowd workers to mine the web for ideas in various domains that might help solve the target problem. Knowledge, methods, principles, skills and tools that experts use in their domains have often been amply documented on the Internet. For example, the origami folding techniques used in the previously mentioned NASA example can be found through a search for “origami patterns”. Thus, problem solvers who might not possess the desired expertise themselves might profitably search for outside-the-box knowledge online.

EXPERIMENT 1: IDENTIFYING OUTSIDE-THE-BOX DOMAINS

Our overall goal in this research is to elicit distant domains and useful inspirations from a large, open-ended idea repository such as the web. In particular, we aim to go beyond previous design-by-analogy work in providing a structured process for identifying fruitful domains for exploration, rather than simply asking individuals to find useful examples directly. We hypothesize that doing so requires generating and exploring domains that are distant from the original problem domain, yet contain structurally similar and relevant inspirations. In Experiment 1 we

explore a process for generating distant structurally similar domains, and in Experiment 2 we test whether the found domains yield relevant and useful inspirations.

In a pilot study we asked workers directly to identify domains that might be useful for solving the schematic representation of the problem. However, the resulting domains were generally too vague to be useful (e.g., “engineering”). To address this problem we turned to a mechanism that people use in real life when they are looking for knowledge in a different domain: they look for referrals. The intuition is that even if people might not have the desired knowledge a problem needs and are not accustomed to thinking about “domains” of knowledge, they may be able to identify a type of expert in a different domain who might deal with relevant problems. We thus reframed our elicitation from crowd workers, asking them to identify professions that could have useful perspectives. For example, if the goal were to design a new power strip, one might ask about professions that deal with problems of packing things into a limited space. Such professions could be as diverse as landscapers packing plants into a small yard; warehouse loaders packing products into crates; user interface designers packing information into interfaces; or even contortionists packing their bodies into small containers. Some of these domains might yield interesting solutions, such as a terraced landscaping solution inspiring the design of a power strip with different height levels to avoid plugs obstructing each other.

The above example also highlights the importance of abstraction in the domain elicitation process. Providing an abstract representation of a problem can increase the diversity of the resulting domains found [25]. For example, the abstract representation of the power strip’s problem for “packing things into a limited space” might elicit more diverse professions than a concrete representation such as “fitting different sized plugs into a power strip”.

In Experiment 1 we manipulated the level of abstraction of the problem representation (i.e., original, concrete problem or abstracted problem) and asked people to identify professions that might have interesting perspectives on the problem. Our goal was to find distant yet structurally relevant domains that might yield useful inspirations.

Participants

One hundred and twenty-two Amazon Mechanical Turk workers [12] participated in this experiment. Forty percent were women, and 93% were native English speakers. Their average age was 33 and ranged from 19 to 68.

Design and Procedure

We selected two design challenges from a crowd innovation website (Quirky.com) to use in the experiment (see Table 1). For each challenge we generated an abstract schematic representation by first abstracting the goal and the sub-goals, and replacing concrete objects with generic objects sharing essential attributes. For example, the original power strip

Original description	Schematic representation
(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?	<ol style="list-style-type: none"> 1. How can you fit objects of different sizes into a container? (goal) 2. Prevent objects blocking each other (sub-goal) 3. Fully make use of a container’s capacity (sub-goal)
(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?	<ol style="list-style-type: none"> 1. How can you dry multiple stackable objects? (goal) 2. Prevent multiple similar objects from taking up much space (sub-goal) 3. Separate multiple stackable objects easily (sub-goal)

Table 1. Original and schematic representations for the two design problems (power strip, cup) used in the experiments.

problem listed on the website talks about plugs and outlets, while the schematic version talks about objects fitting into a container or blocking each other. While in this research the experimenters created the schematic problem representation for convenience and standardizing the input to the experiment, previous research shows schematization of a concrete problem can be done by lay crowd workers following instructions [25].

After accepting the task, participants were randomly assigned to either the original or schematic representation of the power strip or cup problem and asked to recommend types of experts for a design problem. Specifically, they were asked, “*Please read the design problem below and suggest three types of experts who might provide useful or interesting perspectives in solving it and explain why*”.

Rating the Recommended Experts

We hypothesized that the schematic representation would return a more diverse set of experts. Table 2 shows examples of the recommended experts for the two problems in the two conditions. Some professions appeared in both problems (e.g., carpenter, construction worker) but for different reasons. For example, a rationale for a carpenter in the cup problem was that they could “*design a rack that separates the cups just enough so they can dry properly, but not take up too much space*” while for the power strip problem they “*Could design hidden compartments to run cords to other outlets that would also keep the cords hidden.*”

In the original concrete formulation of the cup problem, most recommendations appeared to invoke people who would have direct knowledge about kitchens and cups – e.g., housewives, chefs, bartenders, cup designers, and interior designers. Similarly, many participants who saw the original description of the power strip problem recommended electricians and electronic engineers.

In contrast, the schematic representations of the cup problem appeared to return a wider variety of experts, including bakers, forklift operators, and sculptors. Participants who saw the power strip problem’s schematic representation recommended a variety of experts who work with problems related to sizes, shapes, and blocking. For example, in the *schematic representation* of the power strip problem, a participant recommended a topology expert and explained, “*This person would know about the mathematical study of shapes and shapes in space which would help with deciding about objects of different sizes.*”

To analyze whether the conditions differed in generating domains two judges rated each recommendation on two dimensions: whether the expert was unique among all recommendations for the same problem and how distant the expert’s domain was from the original problem domain. To count unique experts, we combined similar recommendations into single one. For example, “mover” and “a moving expert-moving peoples’ possessions from one home to another” were both classified as “mover”. To judge distance, two judges blind to experimental condition rated each recommended expert on a 7-point Likert scale: “*How different is the above power strip problem from the problems that the recommended expert works with in his or*

Condition	Freq.	Percent of unique experts		Distance	
		Mean	S.E.	Mean	S.E.
Original	204	0.41	0.03	3.87	0.23
Schematic	162	0.60	0.04	5.74	0.16

Note on frequency: There were 68 participants in *Original representation* and 54 participants in *Schematic representation*. Each participant provided three recommendations.

Table 3. Experiment 1: Means and standard errors.

her domain?” By this metric, for example, a “mail sorter” was considered further from the cup problem than was a dishwasher. The judges achieved a good inter-rater reliability of 0.78, calculated using the Intraclass Correlation Coefficient (ICC) [5]. The final distance score was calculated by averaging the scores of the two judges.

Analysis and Results

Table 3 shows the means and standard errors of the percentages of unique experts and the distance scores. A proportion test showed participants identified a larger proportion of unique experts when shown schematic descriptions of the problem than when shown the more concrete, original problem descriptions ($z = 3.21, p < 0.001$). Participants shown the schematic problem descriptions also recommended domains more distant from problem domain than did those shown the original problem description ($z = -5.89, p < 0.001$, by a Mann-Whitney U test used because of the skewed distribution).

EXPERIMENT 2: GENERATING IDEAS

Although Experiment 1 suggests that that converting the original problem description into a schematic representation and eliciting domains in the form of professions led people to identify more distant structurally related domains. However, it fails to answer the key question of whether these distant domains yield more useful inspirations that lead to more creative solutions that the ones found in domains closer to the original problem domain. While more distant domains have the potential to yield inspirations that break fixation with the original problem and suggest interesting approaches from other domains, it is also possible that these domains might not be sufficiently relevant to the target problem to yield useful inspirations. For example, in Table 2 the schematic representation resulted in professions that at first glance might not seem relevant, such as “magician” for the power strip problem or “parking lot attendant” for the cups problem. However, on closer inspection workers provided rationales that make these potentially plausible sources, such as “*Can make objects appear to do many things you might think illogical*” for magician, perhaps suggesting that some trick mechanics might help fit things in places you wouldn’t expect; or

Cup problem		Power strip problem	
Original	Schematic	Original	Schematic
Chef, Housewife, Bartender, Carpenter, Counter-designer, Cup designer, Dishwasher, Gardener, Glass expert, Home-builder, Maid, HVAC-technician, Interior-designer, Professional-organizer, Rack-specialist, Waiter	Baker, Bodybuilder, Carpenter, Cashier, Construction – worker, Designer, Fireman, Forklift-operator, Landscape, Mail sorter, Mathematician specializing in geometry, Meteorologist, Parking-lot-attendant, Pianist, Potter, Sculptor	Electrician, Building-contractor, Computer-technician, Housewife, Network-engineer, Interior-designer, Artificial-intelligence-scientist, Building-maintainer, Cable installer, Carpenter, Construction-worker, Electrical-engineer	Contortionist, Geometry-expert, Graphic-designer, Landscape, Magician, Physicist, Sculptor, User-interface-designer, Architect, Warehouse-dock-loader, Artist, Expert-of-arithmetic, Expert-on-Japanese-aesthetics, Expert-on-topology

Table 2. Examples of the recommended expert domains from Experiment 1 in the Original and Schematic problem representation conditions.

“...needs to be wise about how they position cars so that they can get in and out whenever needed...would know the best way to position objects” for the parking lot attendant.

In this experiment we explore whether distant domains would yield useful inspirations by having participants generate new ideas using the inspirations found in the previous experiment. We are interested in whether participants can find more useful and interesting inspirational sources from an outside-the-box domain and solve the problem better than when searching on their own or when searching in a domain that shared surface similarities to the problem domain. Our hypothesis is that the benefits of far domains occur only when they are likely to contain ideas that have deep structural correspondences to the original problem. Thus, one purpose of this experiment is to insure that distance from the target domain by itself is not what leads to better problem solutions.

The experiment involved four conditions. In the *Problem-driven* condition, participants searched for inspiration in a domain recommended from the original problem description. In the *Schema-driven* condition, they searched for inspiration in a domain recommended from the schematic problem representation. The *Irrelevant schema-driven* control condition tested whether far domains that are not structurally similar to the original problem have the same benefits as far domains that were structurally similar. Specifically, participants searched for inspirations for the cup problem in a domain recommended from the schematic problem representation of the power strip problem and vice versa. Finally, we also included a control condition where participants were asked to search for inspiration on their own, called the *Self-selected* condition. In all these conditions, participants first searched for inspirations and then tried to solve the original problem using the inspirations they found.

Participants

Overall 130 Amazon Mechanical Turk workers participated in Experiment 2. Forty-six percent were women, and 96% were native English speakers. Their average age was 34 and ranged from 18 to 66. Participants from Experiment 1 were excluded from participating in Experiment 2.

Design and Procedure

After participants accepted the task, they were randomly

Independent variable	% creative	Odds Ratio	se	Z
Self-selected	0.26	5.44**	3.10	2.98
Problem-driven	0.12	13.53***	8.91	3.95
Schema-driven	0.65	NA	NA	NA
Irrelevant schema	0.25	5.45**	3.08	3.01

$p < .01 = **$ $p < .001 = ***$

Table 4. The effects of inspiration condition on design quality.

assigned to one of the four experimental conditions. They were asked to search for two inspirational examples for either the Cup or Power Strip problem and then generate a new product idea that solved the problem. They were either given no hint about where to search for inspiration (*Self-selected*) or were told to look for their inspirations from domains recommended in Experiment 1 from participants who saw the concrete version of an original problem (*Problem-driven*), the schematized version of the original problem (*Schema-driven*) or a schematized version of an irrelevant problem (*Irrelevant schema-driven*). Instructions for the conditions are described in more detail below:

1. *Self-selected*. Participants saw either the Cup or Power strip problem as described on the left column in Table 1 and then asked to search for useful ideas on the web that would help them design a solution to the problem. Specifically, they were told, “Please go to the Internet and find two 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 cup [power strip] problem. The useful ideas have to be about similar problems in a non-cup related domain.” After they found two useful ideas, they were asked to solve the problem. Specifically they were told, “Please generate a new product idea that solves the above design problem using the ideas you found.”

2. *Problem-driven*. After seeing one of the two design problems, participants were asked to search for useful ideas that an expert in the problem domain might have. For example, if they were assigned the housewife domain, their prompt would be, “Please go to the Internet and find two useful ideas a housewife might have that could inspire good




Inspiration	Expert domains and solutions
	An architect: Cup with studs on sides, so that they can be stacked while allowing airflow to dry easily.
	A carpenter: Maybe instead of a linear power strip it could be built as an arc or circle that allows more plugs to fit in basically the same amount of space. By turning a line into a circle it opens up more room around each plug or AC adaptor.
	An expert on Japanese aesthetics: If you made an outlet that looks like a Miyabi style building. You could go up levels and it would leave more room for the plugs to go around the outlet.

Table 5. Examples of inspirations and resulting solutions from participants.

solutions for the above problem. These ideas could be knowledge, skills, or methods a housewife uses to solve a similar problem in his or her own domain. Please don't search for ideas related to this cup problem. The useful ideas have to be about problems in the domain that a housewife deals with." An expert was randomly selected an expert without replacement from the domains recommended in the *Original representation* condition in Experiment 1, with a probability proportional to the number of times the expert was recommended in Experiment 1. After finding two useful ideas, participants were asked to solve the problem using those ideas as inspiration.

3. *Schema-driven*. This condition is identical to the *Problem-driven* condition except that participants were asked to find inspirations from domains identified in the *Schematic representation* condition in Experiment 1. After finding two inspirations, they then solved the problem.

4. *Irrelevant schema-driven*. This condition was identical to the *Schema-driven* condition except that participants were asked to find inspirations from the domains identified in the *Schematic representation* condition in Experiment 1 for an irrelevant problem. That is, if participants were to solve the cups problem, they were shown domains recommended for the power strip problem and vice versa. After finding two inspirations, they then solved the problem.

There were roughly 90 recommendations identified in each condition in Experiment 1. When selecting domains to give to participants, we coalesced domains to give related ones a common name and excluded generic recommendations such as "scientist" and "engineer".

Quality of the Design Solutions

Judging the creative quality of an idea can be difficult problem involving significant subjectivity. To judge the ideas we draw on previous research establishing methods for robustly rating creative idea quality, which considers an idea as being creative if it is both novel and practical [4]. Novelty was defined as an idea that was not obvious and differed from existing products on the market. Practicality was defined as how realistically an idea achieved its goal and could be designed and manufactured.

Two judges blind to experimental condition iterated on a rubric for rating product ideas, then used this rubric to independently rate each idea on two 7-point Likert scales measuring novelty and practicality. The judges achieved good ICC inter-rater reliabilities of 0.79 and 0.61 for novelty and practicality respectively. The final creativity score of each idea was computed as the mean of its novelty and practicality scores. Some limitations of this approach are discussed in the discussion section.

Below is an example of a solution to the power strip problem rated highly in terms of both novelty (6.0) and practicality (6.0).

Make a power strip that has some outlets raised about the others. That way a large plug can fit into the raised outlet and still leave room for another plug in the outlet that isn't raised.

In contrast, the design idea below was rated poorly on both dimensions of novelty (1.0) and practicality (2.5).

I would propose a specially made power strip that is longer, and wider than any power strip on the market. Would that has enough space between each socket to fit any plug, or power brick comfortably. Of course it would be bigger, but that is a fair trade off.

Following recommendations from [4], we then classified each design as creative if it was above a media-split threshold of 3.0 for both novelty and practicality or non-creative otherwise.

Analysis and Results

To determine whether the source of the domains participants used as inspiration influenced the quality of their designs, we conducted a logistic regression model predicting with whether an idea was creative (1) or not (0) from the experimental conditions, along with level of education and whether the participants' language was English as controls. Table 4 summarizes the results, with *Schema-driven domain* as the reference group. Results indicate that searching for inspiration in domains far from the original problem significantly increased the probability of generating a creative solution compared to searching in *Problem-driven* domains, in *Irrelevant schema-driven* domains or when given no direction about which domains to search. The odds of producing a creative idea when given recommendations to search in a *Schema-driven* domain were 13.53 times the odds when given a *Problem-driven* domain, 5.45 times the odds when given an *Irrelevant schema-driven* domain and 5.44 times of the odds when participants searched for inspiration with guidance about domain (*Self-selected*).

To further examine whether the domain of search is responsible for the difference in the quality of ideas, we added the *domain distance* score (judged in Experiment 1) as a mediator and re-run the logistic regression analysis. The results show that *domain distance* is significantly correlated with the odds of producing a creative idea: for every one-unit increase in the distance score, the odds of producing a creative idea increases by 2.07 times ($p < 0.01$). After adding *domain distance* as a mediator, the quality difference between the *problem-driven* condition and the *schema-driven* condition became non-significant ($b = -0.96, p = .23$). However, the difference of idea quality still exists between the *schema-driven* condition and the *irrelevant schema-driven* condition: even holding domain distance constant, the odds of producing a creative idea when given recommendations to search in a *schema-driven*

domain were 3.22 times of the odds when given an *irrelevant schema-driven* domain ($p < 0.05$).

The mediation analysis suggests that the distance of inspiration from the target domain created by the abstraction of the problem schemas completely explains the higher idea quality in the *schema-driven* condition compared to the *problem-driven* condition. That is, compared to the *problem-driven* condition, people in the *schema-driven* condition used examples from distant domains as inspiration, which in terms led them to generate creative ideas. However, the distance of inspiration from the target domain only partially explains why solutions inspired by *schema-driven* examples were more creative than ones inspired by *irrelevant schema-driven* ones. This failure of mediation shows that the benefits of far domains occur only when they are likely to contain ideas that have deep structural similarities to the original problem.

Search Mechanisms

In Experiment 2, we found that people were able to find relevant inspirational examples for a problem in remote domains and were more likely to generate creative ideas by using the inspirations that was found from remote domains.

These results, however, beg the question of how non-expert crowd workers were able to find useful ideas in unfamiliar domains. People found and used inspiration in the *schema-driven* domains using a variety of mechanisms. To illustrate, Table 5 shows the inspirations people found by searching in the domain of an architect, a carpenter, and an expert on Japanese aesthetics. Searching in an architect's domain, a participant found and adopted the energy-saving ventilation design of a building for the cup problem; searching in a carpenter's domain, a participant was inspired by the circle-shape of a drawer; and searching in the domain of Japanese aesthetics, a participant borrowed a design of the Japanese Miyabi. Other participants also utilized the mechanisms and principles of other domains. For example, a participant designed a power strip that has some outlets raised about the others by utilizing how warehouse dock loading systems work. Another participant designed a bendable power strip by searching in a contortionist's domain.

These forms of inspiration are much different from those found through domains recommended based on the original problem description or when participants searched for inspirations on their own. In these conditions, people found ideas related to organizing clutter, drying dishware, or rack design. When searching based on irrelevant domains, people often found information that appeared random and their solutions were not well connected to the inspirations they found.

The findings support the assumption that knowledge encoded on the web in various expert domains is useful in sparking creative solutions for a related problem in a different domain. However, the outside-the-box knowledge

had to be structurally alignable and relevant to the problem as well. The distance and relevance were created through the schematic representation in Experiment 1: the abstraction reduces the domain fixation leading to a distance between the original problem and the recommended domains, while goals and sub-goals keep the connection between them. The expert domains in Experiment 2 work as a mediation step to bridge the schematic representation and the useful ideas in the domains.

DISCUSSION

Previous research has shown that creative solutions to problems often emerge when experts from outside a problem domain apply their “foreign” skills, techniques and tools to the problem. This insight inspired our research in developing a systematic process to leverage crowd workers to harvest potentially applicable ideas from fields outside of a problem's domain. To do so, we had to get the workers to “think outside of the box” and overcome the challenge of functional fixedness [1], the cognitive bias that limits people to using objects only for their traditional purposes. We addressed this challenge by transforming the original, concrete problem statement to a more abstract representation. When combined with an instruction set framed in terms of “referrals”, this schematic representation of the problem allowed people to identify a more diverse set of domains to explore for potential solutions than did the original, concrete problem description.

Moreover, when non-expert crowd workers searched for inspirations for solutions in domains primed by the abstract representations, they returned a rich set of relevant examples. Their examples were often very different from the original problem in terms of surface features, but shared structural features. For example, at the surface level a Miyabi style building has little in common with an electrical power strip, but its three dimensional structure and arrangement of rooms might suggest ways to fit different plugs in a power strip. Crowd workers who collected inspirations from these remote domains produced more creative problem solutions than did those who searched for inspiration without guidance about where to look or who looked for inspiration in domains primed by the original concrete problem statement or by an abstract representation of a problem that did not share structural correspondence with the target problem.

We asked Amazon Mechanical Turk workers, who were not selected for engineering or design expertise, to solve relatively simple but real design challenges. The participants provided brief text descriptions of the approach they would take to a problem, such as building a power strip in two dimension (a curve) or three dimension (a tower); they were not required to provide details about implementation. Despite these limitations, the Cup and Power Strip problems were authentic design challenges taken from Quirky.com, the crowd innovation company

specializing in consumer products. In the past, members of the Quirky community designed detailed solutions for these challenges, and Quirky itself manufactured and sold consumer products to solve these problems. Interestingly, some of the solutions provided by crowd workers, such as a bendable power strip inspired by the contortionist domain, paralleled those from the Quirky community itself, such as the bendable PivotPower (one of Quirky's most popular products). Similarly, the architecture-inspired method for attaching studs to cups so that air could circulate between stacked cups to dry them is highly reminiscent of Totem Air Dry (another Quirky product). Workers did not indicate being influenced by Quirky products and none returned a Quirky product as an inspiration. When they submitted design ideas that paralleled Quirky products, their designs had clear links to examples retrieved from their assigned source domains, although we cannot guarantee that participants had never seen the relevant Quirky products. However, assuming no direct influence, parallels between crowd workers' solutions and those manufactured by Quirky suggest that the design ideas participants came up with in this controlled experiment could be commercially valuable in the world.

We consider our research, which challenged relatively unsophisticated workers to solve relatively simple design problems, only a first step. One area for profitable future research is investigating whether this process will be effective in crowd innovation sites like InnoCentive, TopCoder, or MathOverflow, which typically challenge sophisticated experts to solve complicated R&D problems. We believe the process of searching for inspirations outside of a problem's domain should be useful to overcome functional fixedness for sophisticated problems as well as simple ones and may actually be more useful for problems that require deep expertise. Abstracting away some of the non-essential features of the problem should help in identifying new domains in which to search for solutions. However, when the problem is complex or when the solution must be fully elaborated or reduced to practice, one may need to recruit true experts from the distant knowledge domains and not just novices reading what experts had previously written.

Another limitation in the current research is that one of the two judges in Experiment 2 was the first author. While both were blind to experimental condition and used rubrics to reduce bias and to standardize their ratings, future experiments could benefit from separate judges to reduce any biases resulting from trying to infer which experimental condition each idea came. Future research could also examine the robustness of using non-designers as judges in this context, though prior work shows high agreement between designers and non-designers in judging consumer products [e.g., 23].

The approach of using schemas as described here could also benefit from further research into the boundary conditions

for where schemas are useful. For example, it is not yet clear whether schemas would be as useful if the design process continued past ideation into the prototyping stage. Prototyping may force people to think more deeply about how to adapt specific mechanisms of inspirational examples to a solution, which may result in poor fits. For example, people may run into configurational or material use problem with designing a Japanese Miyabi-style power strip, which could limit the practicality of manufacturing it.

Representing Problems in Crowd Innovation

A core contribution of this work is that we use a problem's abstract representation as a cue to identify domains where a solution might be found, instead of directly searching for analogical ideas [25]. This opens up the search space from well-defined idea repositories (as in [25]) to the entire web. Starting a search for solutions in experts' domains can provide more concrete guidance about the type of knowledge needed to solve a problem than more abstract schematic representations of the problem, which can be ambiguous and difficult to interpret. Presenting different domains to different problem solvers opens up the search space and helps solvers find more varied inspirations even from the same schematic representation.

We argue that problem representation is an important research topic with the rise of crowd innovation. This new problem-solving model opens a rich research area on how to formulate problems so that a distributed group of people can solve them creatively. The schematic representation proposed in this paper is one approach constructing an alternative representation for a problem in order to encourage outside-the-box search for solutions. However, how to identify the most productive level of abstraction for a particular design challenge is still an open question. An interesting area for future study is whether more complex problems on sites such as InnoCentive could be re-represented and classified based on abstract goals and sub-problems, and thereby attract more actual outside-the-box experts to solve them.

ACKNOWLEDGEMENTS

This work was supported by NSF grants IIS-1526665, IIS-1149797, IIS-1217559, OCI-0943148, IIS-0968484, IIS-1111124, Bosch, Google, and Microsoft, Heinz College Center for the Future of Work, and Carnegie Mellon's Center for the Future of Work. We thank reviewers for useful feedback.

REFERENCES

1. Robert E. Adamson. 1952. Functional fixedness as related to problem solving: A repetition of three experiments. *Journal of experimental psychology* 44, 4: 288.
2. Meg Monk. 2013. The Digital Universe. <http://universe.byu.edu/2013/12/12/byu-engineers-use-origami-to-make-more-space-in-space/>

3. Maryam Fazel-Zarandi and Eric Yu. 2008. Ontology-based expertise finding. In *Practical Aspects of Knowledge Management*. Springer, 232–243.
4. Ronald A. Finke, Thomas B. Ward, and Steven M. Smith. 1992. *Creative cognition: Theory, research, and applications*. Bradford Book.
5. Ronald Aylmer Fisher. 1954. *Statistical methods for research workers*. Oliver and Boyd.
6. Karsten Frey, Christian Luthje, and Simon Haag. 2011. Whom Should Firms Attract to Open Innovation Platforms? The Role of Knowledge Diversity and Motivation. *Long Range Planning* 44, 5–6: 397–420.
7. Morten Friis-Olivarius. 2013. Overcoming design fixation through education and creativity methods. *The 19th International Conference on Engineering Design*. (ICED'13), 139–148.
8. Dedre Gentner, Sarah Brem, Ron Ferguson, Philip Wolff, Arthur B. Markaman, and Ken Forbus. 1997. *Analogy and creativity in the works of Johannes Kepler*. T. N. Ward, S. M. Smith, J. Vaid, eds. Creative Thought. Amer. Psych. Association, Washington, DC, 403–460.
9. Dedre Gentner, Russell Landers, and others. 1985. Analogical reminding: A good match is hard to find. In *Proceedings of the international conference on systems, man and cybernetics*, 607–613.
10. Andrew Hargadon and Robert I. Sutton. 1997. Technology brokering and innovation in a product development firm. *Administrative science quarterly*, 716–749.
11. Keith James Holyoak. 1996. *Mental leaps: Analogy in creative thought*. MIT press.
12. Aniket Kittur, Ed H. Chi, and Bongwon Suh. Crowdsourcing user studies with Mechanical Turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI'08), 453–456.
13. Karim R. Lakhani, Lars Bo Jeppesen, Peter Andreas Lohse, and Jill A. Panetta. 2007. *The value of openness in scientific problem solving*. Division of Research, Harvard Business School.
14. J.S. Linsey, A.B. Markman, and K.L. Wood. 2012. Design by analogy: A study of the WordTree method for problem re-representation. *Journal of Mechanical Design* 134, 4, 041009.
15. Tony McCaffrey. 2012. Innovation relies on the obscure a key to overcoming the classic problem of functional fixedness. *Psychological science*, 0956797611429580.
16. David W. McDonald and Mark S. Ackerman. 2000. Expertise recommender: a flexible recommendation system and architecture. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work* (CSCW '00), 231–240.
17. The New York Times. Car Mechanic Dreams Up a Tool to Ease Births. www.nytimes.com/2013/11/14/health/new-tool-to-ease-difficult-births-a-plastic-bag.html
18. Tony Poze. 1983. Analogical connections—The essence of creativity. *The Journal of creative behavior* 17, 4: 240–258.
19. Robert J. Sternberg. 1998. *Handbook of creativity*. Cambridge University Press.
20. Lynn Streeter and Karen E. Lochbaum. 1988. An expert/expert-locating system based on automatic representation of semantic structure. In *Proceedings of the Fourth Conference on Artificial Intelligence Applications* (IEEE '88), 345–350.
21. Thomas B. Ward. 1998. Analogical distance and purpose in creative thought: Mental leaps versus mental hops. *Advances in analogy research: Integration of theory and data from the cognitive, computational, and neural sciences*, 221–230.
22. Youtube, <http://www.youtube.com/watch?v=uL1ovAYtKuQ>
23. Lixiu Yu and Jeffrey V. Nickerson. 2011. Cooks or cobblers?: crowd creativity through combination. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (CHI'11), 1393–1402.
24. Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2014. Distributed analogical idea generation: inventing with crowds. *Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (CHI'14), 1245–1254. <http://dx.doi.org/10.1145/2556288.2557378>
25. Lixiu Yu, Aniket Kittur, and Robert E. Kraut. 2014. Searching for analogical ideas with crowds. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, (CHI '14), 1225–1234. <http://dx.doi.org/10.1145/2556288.2557378>
26. Jun Zhang and Mark S. Ackerman. 2005. Searching for expertise in social networks: a simulation of potential strategies. *Proceedings of the 2005 international ACM SIGGROUP conference on supporting group work*, 71–80