Cooks or Cobblers? Crowd Creativity through Combination

Lixiu Yu
Stevens Institute of Technology
Castle Point on Hudson, Hoboken NJ 07030 USA
lisayu1103@gmail.com

Jeffrey V. Nickerson
Stevens Institute of Technology
Castle Point on Hudson, Hoboken NJ 07030 USA
jnickerson@stevens.edu

ABSTRACT
A sketch combination system is introduced and tested: a crowd of 1047 participated in an iterative process of design, evaluation and combination. Specifically, participants in a crowdsourcing marketplace sketched chairs for children. One crowd created a first generation of chairs, and then successive crowds created new generations by combining the chairs made by previous crowds. Other participants evaluated the chairs. The crowd judged the chairs from the third generation more creative than those from the first generation. An analysis of the design evolution shows that participants inherited and modified presented features, and also added new features. These findings suggest that crowd based design processes may be effective, and point the way toward computer-human interactions that might further encourage crowd creativity.

Author Keywords
Crowdsourcing, creativity, conceptual combination, human based genetic algorithms, social computing, design sketches.

ACM Classification Keywords
H5.3. Group and Organization Interfaces

ACM General Terms
Design

INTRODUCTION

Can a crowd design well? On the one hand, diversity of thinking should lead to creativity, as expressed in a Chinese proverb: Three cobblers combined make a genius mind. On the other hand, conflicting styles may lead to chaos: Too many cooks spoil the broth.

Scientific research on social creativity reflects this division: for example, brainstorming has been shown by some to be a powerful technique for encouraging ideas and by others as a waste of time [21, 25, 27]. Technology has played a role in the debate: some of the problems encountered in brainstorming, such as production blocking, can be mitigated by the use of computer-based tools [26].

Indeed, the creativity tool landscape is diverse. Among these, automated design tools often make use of blind variation, followed by selection: several random choices are made, and the results evaluated [6]. Of particular note are genetic algorithms: computer simulations of evolution in which gene-like representations are altered and combined in a search for optimal solutions [7, 11, 15]. Such algorithms, however, don't work well if the candidate solutions cannot be formally represented and computationally evaluated. Because of this, *human based* genetic algorithms have been proposed; these algorithms allocate many or all of the combination and evaluation functions to human labor [19].

Until recently, large-scale experiments to test human-based algorithms were prohibitively complex. But web technologies, and the organizational structures provided by crowdsourcing marketplaces now make it feasible to coordinate large numbers of people: Participants can collaborate with each other by working on parts of an overall task, facilitated by crowdsourcing processes and technology [1, 4, 9, 16, 17, 22, 23]. Some of these techniques are designed so that each member of the crowd functions like a processing node in a computer system [20, 29].

These techniques may not work, because computers and humans are so different. Computers execute the instructions of an algorithm without complaint or caprice. In contrast, while human beings can bring a wealth of alternative perspectives to any task, they may also choose to ignore algorithmic instructions. Will humans participating in such systems be productive, like the cobblers, or destructive, like the cooks?

We answer this question by performing an experiment, based on a system we have built. This sketch combination system is a variant of a human-based genetic algorithm: one crowd creates a first generation of sketches, and then successive crowds create additional generations by combining the sketches made by the previous crowds. A total of 1047 participants played a variety of roles, leading to the creation of 80 third-generation design sketches.

The contributions of this research are two-fold. First, the effectiveness of the sketch combination system is demonstrated: the creativity of later design sketches increased significantly, compared to the initial sketches. This sketch system has the potential to be applied to design problems in many fields. Second, this system provides evidence for the role of combination in creativity. Specifically, designers inherit and modify existing features and also added new features when combining the given designs. In addition, the affordances of the drawing tool affect the emergence of features. The work has practical implications—the system can solve design problems—and theoretical implications—the crowd can function as a combination process.

In the following sections of the paper, the sketch combination system is described, along with an experiment illustrating its use. First, previous research on combination and creativity is discussed.

COMBINATION AND CREATIVITY

The combinatorial conjecture of creativity and related theories claim that new ideas come through combination [6, 35, 40]. Combination is also central to the success of genetic algorithms [15]. But such algorithms express combination in a very specific way.

Genetic algorithms start with a population of genes. A gene can be imagined as sequence of bits that represent features. Once a population is ranked for fitness, parents are picked to produce offspring. As in natural selection, the fit parents are more likely to produce offspring than the unfit parents. The children are created by crossover: some proportion of the bits will come from the first parent, and the rest from the other parent. Thus, new combinations are produced. Prior to this combination step, the population undergoes a mutation stage, in which some parents are altered by flipping their bits [7, 11, 15]. In practice, some genetic algorithms rely mainly on crossover, others mainly on mutation [37]. Whatever operator is being used, the search space is bounded. That is, the bits in a genetic algorithm represent a finite if large search space, and the crossover operation can only produce a child by combining the bits of the parents.

Human conceptual combination shares some similarities with machine combination. Simonton sees scientific creativity as a kind of natural selection: ideas are combined at random, and some of them lead to discoveries [35]. We don't have direct access to combinatory processes in the brain, so when a fortuitous combination occurs, we are not sure if it is the result of a directed search to find complementary ideas, or instead the result of many blind trails run below the level of consciousness [6]. Thagard claims that inference processes guide search [40]. In the design community, a variety of computer based and human assisted genetic algorithms have been used successfully [10].

But human combination processes also differ from machine processes in important ways. The search space of human processes is not neatly bounded: in combining ideas our minds have access to many past experiences and present perceptions, not just the representations at hand (c.f. [24]). Consequently, even if instructed to combine two ideas, a person may associate the presented ideas with any of a vast number of other ideas, remembered or constructed on the fly. These other ideas may or may not be as good as the presented ideas. Thus, we could argue either that human based genetic algorithms will work, because the human process is close enough to the machine process, or will not work, because there will be too much variability in the human based process. We predict that the human combination process will work much like a genetic algorithm – that even if there is more variability, it will often be to advantage. In the end, the designs evolved through a human based genetic algorithm will be judged more creative than the initial designs.

In testing this through experiment, we will ask the participants to sketch. Sketches play a strong role in design. On the one hand, they are used as a way of refining ideas, and work as stimuli to elicit creativity [12, 38]. On the other hand they can constrain ideas: that is, designers will fixate upon seeing a sketch, and likely fail to design something original. This fixation has been described as conformity to the sketches just seen [36]. If designers fully conform to the provided sketches, then the proposed combination system will result in a reshuffling of the initial feature set. In contrast, if the instructions are ignored, and ideas are generated through free association, then the algorithm will play little role in the final designs. To our knowledge there has been no empirical research on human based genetic algorithms that has addressed this issue. Because of human variability, we predict that there will be mix: there will be some cases of conformity to the features in parent ideas, and some cases of modification of these features, or introduction of new features outside the two given ideas, a process we call augmentation.

People may conform to or augment the features of the given designs during the combination process. Then how will features move through the generations of the combination process? Atypical features are attended to more, remembered better [28, 34], and selected more often than typical features [14]. Extending this idea to the context of a genetic algorithm, we predict that atypical features will occur more often in the third generation than in the first, and that typical features will occur less often in the third generation than in the first.

The drawing tool should exert an additional influence on combination and the eventual creativity of the sketches: participants may produce only what the tool makes easy to produce, or may associate one shape with another shape based on the options the tools provide.

In sum, we make several testable predictions. The sketch combination system will produce more creative designs in the last generation than the first. Atypical features will increase more often than typical features. Combination will include instances of conformity and augmentation.

METHOD

Problem Definition

Sketches provide a rich way to study design: sketches are the way conceptual designs are developed and shared with others [38, 41]. We looked for a design problem that both novices and experts could accomplish, and decided to use a problem that has been the focus of previous creativity research [12]: the design of chairs for children. Chairs are objects that are universally understood, but allow for a range of variations. Moreover, in the psychological literature, chairs are considered a basic level category [31], and thus have a recognizable prototypical shape, which might make it easier for participants to evaluate the originality of design alternatives. Children's chairs are not necessarily different from adult chairs in anything but size, but in practice are often more whimsically designed and utilize a wide range of geometries, which may encourage variation in the designs.

The Measurement of Creativity

Innovation is generally believed to be the successful implementation of creative ideas [3]: In the current study, we focus on the creativity of the designs, and will leave the study of their implementations to future research. A creative idea should be both novel and potentially useful [3]. Operationally, this leads to a binary measure of creativity: Creative designs are designs that exceed a certain threshold on both the scales of originality and practicality [8]. Thus in our study, the designs are evaluated on two scales: originality and practicality. The designs will then be classified as being creative or not based on the their scores on the two scales.

The Sketch Combination System

We brought together a crowdsourcing marketplace, design tools, and organizational processes to create a sketch combination system. We will look at each in turn. Briefly, the technology was an integration of the Amazon Mechanical Turk platform [2] and the Google Docs drawing platform [13]. The organizational process was based on the idea of a genetic algorithm [11, 15] implemented with human participants [19].

The crowdsourcing marketplace

Amazon Mechanical Turk was used by participants in this study. There are two types of users in Amazon Mechanical Turk: requesters and workers. Requesters post their tasks by creating HITs (Human Intelligence Tasks) and workers perform tasks in return for compensation. Previous studies have described the characteristics of these workers [18, 32].

The design software

Google Drawing was the tool used by participants to produce designs. The drawing tool provides options to perform freehand sketching, vector line drawing, text entry, and shape drawing. Participants were directed to a Google document page already opened as a drawing. Such a page, in later generations, presented the designs to be combined. Participants had access to all the features of that drawing tool. In the combination tasks, the presented sketches were rasterized, so that participants were required to generate the new features they saw in the images.

The organizational process

In order to clearly specify the sketch combination system, we first need to describe the way a genetic algorithm works. Such an algorithm starts with a first generation population, and performs a fitness ranking. Then, members of the population are selected to become parents of the next generation. Most often, tournament selection is used: two parents are selected at random, and the fitter chosen. Another two parents are chosen, and the fitter chosen. The two chosen parents then produce offspring through a combination procedure. These offspring serve as a new population, which is then ranked, and the process repeats. Because combination can sometimes take the worst features of highly ranked parents, and thereby degrade the available genetic pool, a set of the strongest parents are promoted into the next generation without change, an aspect of the algorithm referred to as *elitism* [7].

The sketch combination system was based on the above description, and designed to run for three generations, as summarized in Figure 1.

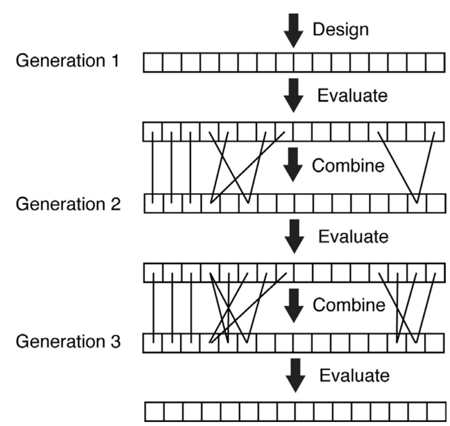

Figure 1. The generations of the experiment.

Below, we elaborate how the six steps in Figure 1 were implemented during the experimental test of the sketch combination system.

Generation 1 design

Eighty randomly selected participants were instructed to generate the population of the first generation:

Please design a chair for children.

Generation 1 evaluation

Two hundred different participants, randomly selected, then evaluated the designs generated by generation 1. Each participant rated 12 randomly selected designs on originality and practicality using 7-point Likert scales, as shown in Figure 2.

In response to a request to design a chair for children, your fellow workers created the following chairs for children. Please evaluate the originality and the practicality of each design based on the given scales.

Not original at all	○1	○2	○3	○4	○5	○6	○7	Extremely original
Not practical at all	○1	○2	○3	○4	○5	○6	○7	Extremely practical

Figure 2. The evaluation prompt.

Generation 2 combination

Tournament selection was used to select sixty pairs of generation-1 designs to serves as parents for generation 2. Participants combined the pairs:

The designs on the right are from your fellow workers. Please create a new chair for children by combining aspects of the two chairs shown.

The twenty highest-rated chairs from generation 1 were automatically promoted to generation 2.

Generation 2 evaluation

All 80 chairs were evaluated by a new set of 247 participants in the same way generation 1 was evaluated.

Generation 3 combination and evaluation

Again, tournament selection produced sixty pairs of generation-2 designs that a new crowd of sixty participants combined to create new chairs for Generation 3. In addition, the twenty top-rated chairs from Generation 2 were also promoted into generation 3. Then, 400 new participants rated the originality and practicality of all designs from all generations. Each participant rated 10 designs, including four from generation 1, three from generation 2, and three from generation 3.

RESULTS

Overall, there were 1047 participants involved in the experiment. Eighty participants took part in generation 1 design, 200 participants in generation 1 evaluation, 60 participants in generation 2 combination, 247 participants in generation 2 evaluation, 60 participants in generation 3 combination, and 400 participants in the overall cross-generation evaluation. Among all participants, 47% were female, 59% were native English speakers and 79% had earned college or graduate degrees. Ages ranged from 19 to 60 with a mean age of 28. Among the 200 participants who created or combined the chairs, 20% claimed 1-3 years of design experience.

Overall, 200 sketches of chairs were collected. They varied in many aspects, including shape and color. Participants made use of the menu choices of the Google Docs drawing tool to draw the chairs–some sketched freehand, and others started with shape stencils. Some sketches portray everyday chairs. Others depict unusual chairs, including child-attractive features, such as smiley faces, animal figures, and heart shapes. Several example chairs are shown in Figure 3.

Figure 3. Examples of the chairs.

The Generational Effect

Are designs from the last generation better than the first generation? Participants had been asked to rate the practicality and originality of both generation 1 and generation 3 designs. Figure 4 shows a scatter plot of generation 1 and generation 3 designs with respect to originality and practicality.

As discussed in the previous section, we used a binary measure of creativity that includes only designs that exceed a certain threshold on both the scales of originality and practicality [8]. We chose the approximate mean value of the ratings across all designs, 4.0, as that threshold. The dots in the upper right quadrant of Figure 4, with positions above and to the right of (4, 4), qualify as creative designs. From the figure, we can see that generation 3 designs tend to be shifted toward more creative side than generation 1 designs. The proportion of creative designs per generation is shown in Figure 5.

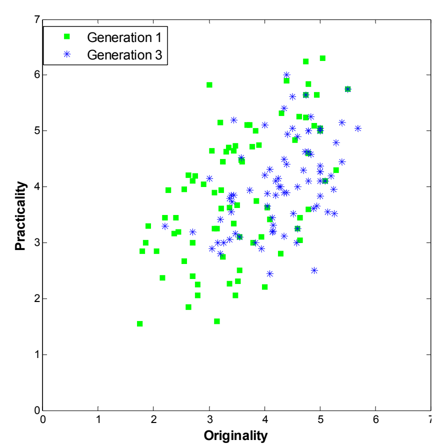

Figure 4. The originality and practicality of all designs from generation 1 and generation 3.

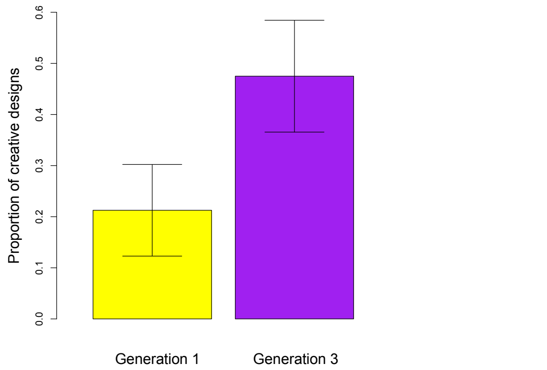

Figure 5. Proportion of creative designs. Error bars represent 95% confidence intervals.

The proportions are significantly different ($x^2(1, N=160)=11.08$, $P<0.01$). The difference is not an artifact of elitism: only five of the 20 elite chairs from generation 1 are still present in generation 3, and, if they are removed, the difference between the generations remains significant ($x^2(1, N=155)=6.55$, $P<0.01$).

This result is encouraging for the future of human based genetic algorithms. Any genetic algorithm needs to produce some children that are superior to their parents; otherwise the population regresses. Here not only do the children not get worse, the number of top children increases significantly. At the end of any design process one can expect implementation of only the best few designs, so a design process that generates a few high quality designs and many poor ones is preferable to a process that generates all medium quality designs.

We next compared the means of originality and practicality across all designs in the first and last generation. T-tests confirmed that originality increased significantly from generation 1 to generation 3 ($t(158)=5.67$, $P<0.01$) while practicality did not ($t(158)=1.04$, $P=0.15$) (see Figure 6).

Participants' Ratings

Are the participants performing well on the combination and rating tasks? The designs shown in Figure 7 are example chairs with low creativity. The left chair is rated relatively high in practicality but low in originality. The right chair, high in originality, low in practicality. The designs in Figure 8 are example of chairs rated as creative designs. Participants do appear to perform well on the combination and rating tasks: their ratings appear to capture the originality and practicality of the design, and at least some participants carefully combine features from one generation to the next.

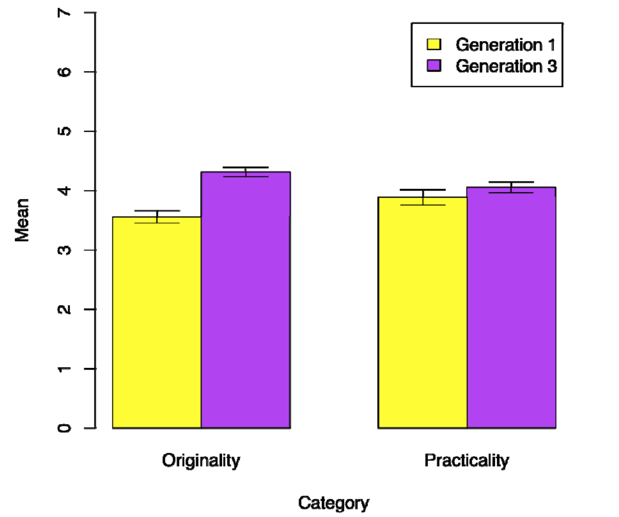

Figure 6. The originality and practicality of all chairs in generation1 and generation 3. Error bars represent standard error.

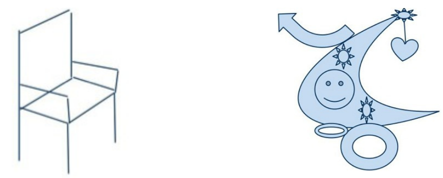

Figure 7. Two chairs rated with low creativity, from generations 1 and 3.

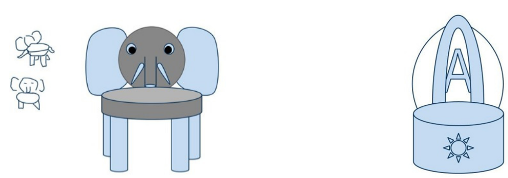

Figure 8. Two chairs rated with high creativity, from generation 1 and 3.

The demographics of the participants were examined to exclude the possibility that the designers' characteristics might have determined the difference in creativity between generation 1 and generation 3. No significant correlations between gender, education, and age were found with the originality or practicality of the sketches produced. The language spoken (classified as English vs. non-English) did correlate with creativity, but there was no difference in the proportion of participants with respect to this language distinction in generation 1 and 3.

How Participants Combine Designs

By looking at the way the designs were combined we might gain insight into the sources of increases in creativity. We may also learn how to improve the system in order to amplify this increase.

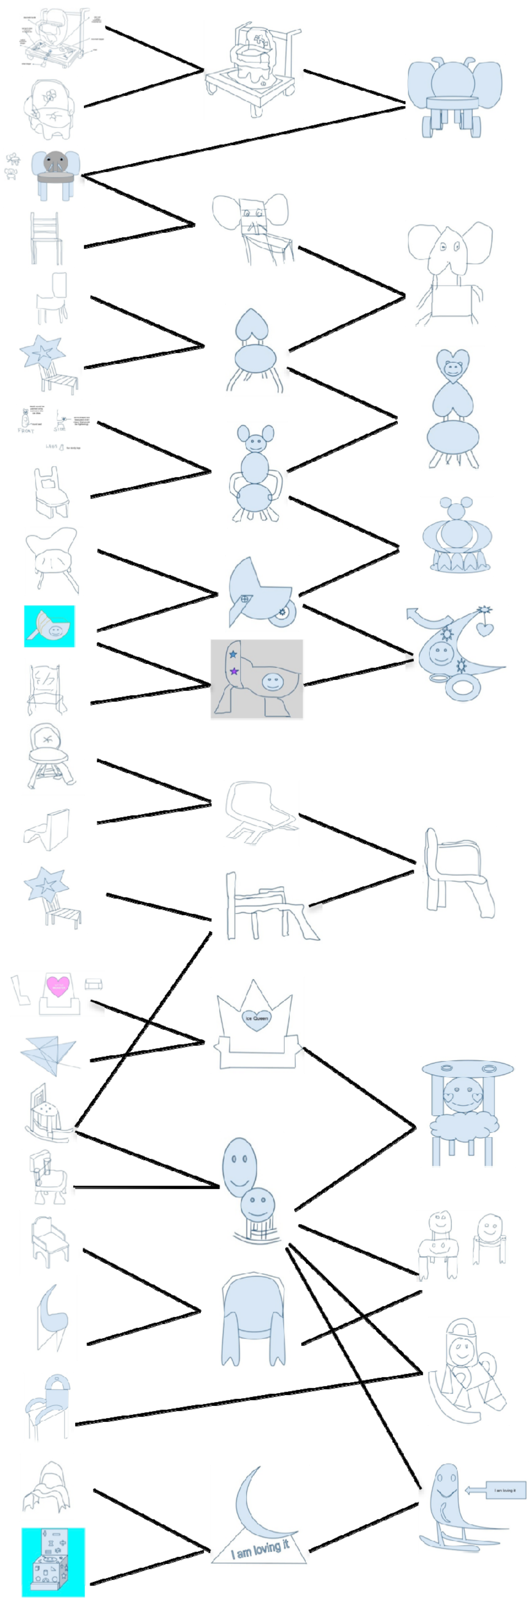

Figure 9. Evolution of several chairs.

Participants combined the given designs in many different ways, as shown in Figure 9. On the left are the chairs created in generation 1. In the middle are the chairs of generation 2, created through combining the chairs of generation 1. On the right are the generation 3 chairs, combined from generation 2. The lines indicate parent-child relationships.

As predicted, there appear to be two forces at work in the combination process. The first is conformity: the designers faithfully create a new chair by combining features of the provided chairs. The second is augmentation: the designers modify features inherited from the given chairs, or add features not present in the given chairs.

Figure 10 shows an example of conformity. The chair inherits features from parent chairs in a direct way. In addition, the designer explicitly describes the features that were combined:

> *The chair is more similar in design to the second than the first, with the straight armrests and the legs, but has the cushioning that the first seems to have. It also has some of the fun shapes that the first chair had.*

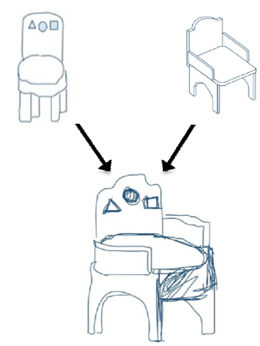

Figure 10. An example of conformity.

In contrast, the evolution tree shown in Figure 11 is an example of augmentation.

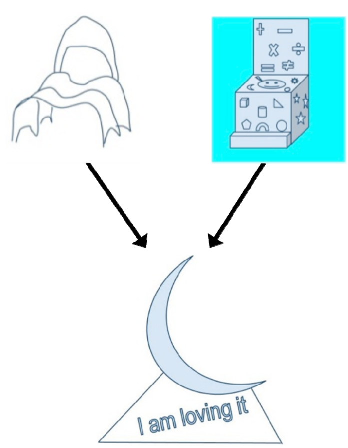

Figure 11. An example of augmentation.

			G1	G3
Typical legs	Four straight legs	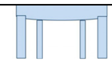	44	22
	Four straight legs with crossing bars	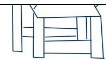	5	3
	Two straight legs	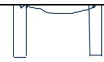	8	4
Atypical legs	Pointing legs		1	2
	Rocking chair legs	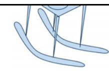	2	6
	Animal legs		0	5
	Flower legs	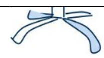	0	3
	Wheel legs	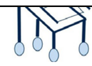	4	10
	Crossing legs	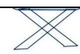	0	1
	Ball legs	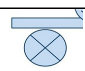	0	1
	Triangle legs	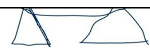	0	1
	Semi-circle legs		0	1
	Curly legs		0	2
	Goblet legs		0	2
	Base legs		1	1
	Three legs		0	1
Legless			15	15

Table 1. Chair legs.

One small feature of the right parent, a crescent moon, is amplified into a major feature. And text is added to the chair inside a new legless base. The explanation of the designer doesn't mention the parent chairs:

It is a round chair which will allow the child to relax on it. The legs are straight to give the chair a firm and stable bottom. The half moon like shape makes the chair attractive and fun for the child to use.

The Evolution of Features

The typical features of the chairs are those that occur in prototypical chairs [31]. For example, prototypical chairs have four straight legs, a typical feature. A non-prototypical chair might have legs that are designed to look like the legs of animals, an atypical feature. We predicted that such atypical features would be more likely to increase in subsequent generations than typical features. In order to test this, the legs of the chairs were examined. Coding of the legs was accomplished by two raters by first determining a set of distinct categories, and then classifying all chairs into those categories. These include 17 different categories that are summarized in Table 1. The chair legs were then coded into three major categories: typical, atypical, and legless by the two raters. Inter-rater reliability, measured by a Kappa score, is 0.82.

We contrast the proportions of atypical legs and typical legs in Figure 12. We found that atypical features increase significantly in frequency ($x^2(1, N=160)=19.59, P<0.01$), at the same time as typical features decrease significantly ($x^2(1, N=160)=13.30, P<0.01$) in generation 3.

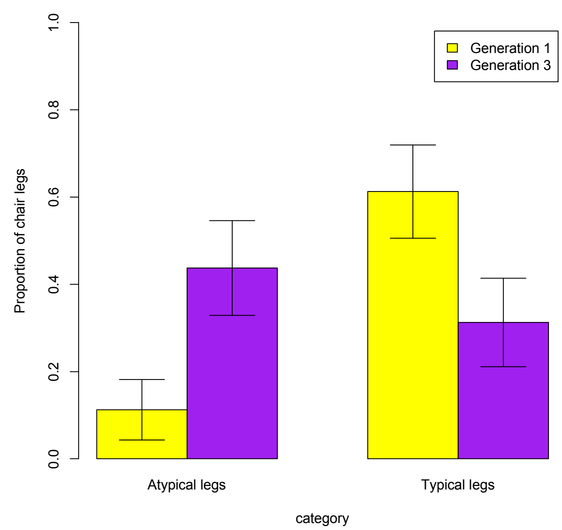

Figure 12. The evolution of chair legs from generation 1 to generation 3. Error bars represent 95% confidence intervals.

Affordances of the Tool

The Google Docs drawing tool provides several options for drawing sketches, including a pull-down shape palette, as shown in Figure 13. There are two ways the tool appears to be affecting the sketches.

Many participants made use of the shapes shown in the palette to design their chairs. In our experiment, we asked participants to use the scribble freehand drawing function when designing their chairs. However, many participants didn't comply, and used shapes instead. Specifically, we found that 17 sketches out of 80 in generation 1 used the shapes palette to design their chairs, while 40 sketches out of 80 in generation 3 used the shapes palette. This difference is significant according to a proportion test ($x^2(1, N=160)=13.19$, $P<0.01$). Many of the shapes used from the palette are atypical with respect to chair design, so one possible explanation for the increase is that these atypical features tend to be noticed, and are carried over into the next generation, just like the atypical chair legs.

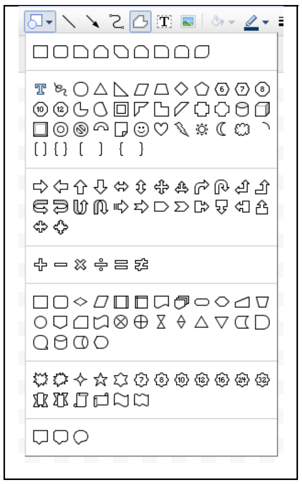

Figure 13. Pull-down shape palette in the Google Docs Drawing application.

Another explanation is that the palette serves as source of association [24]. For example, the number of hearts, smiley faces, crescent moons, circles, cloud and suns is high in the sketches. These all occur in the same section of the palette. Of the 57 sketches that use any shapes, 25 sketches have at least two of the mentioned shapes as decorations or as the shapes of the seat, back or arms. For example, Figure 14 shows a combined chair with multiple shapes as decorations, while the parent chairs only have the sun as decoration. Perhaps a participant pulled down the shape palette to add one shape, and decided to also add several neighboring shapes. This conjecture might be tested through experiments that provide different palettes and measure the frequency of shapes used in relation to their distance from each other on the palettes.

FUTURE RESEARCH

The described system affords future studies that might contrast different tool and processes. Future research might extend the reported experiment in three ways. First, the evaluation of the designs was done by a general population, recruited through Amazon Mechanical Turk. Other populations, and in particular expert or specialist populations, might have different perceptions of creativity. They may be especially skilled at assessing pragmatic issues such as manufacturing costs. Moreover, if indeed experts do have different perceptions, then it might be important to involve them in not just the evaluation but also the generation of the designs.

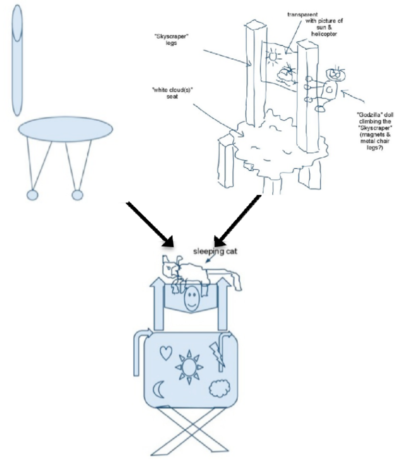

Figure 14. A combined chair with shapes.

Then, the design problems could contain more constraints that experts are more likely to be familiar with: time to market, materials, labor costs. Because it is difficult to assemble a crowd of experts, those that can be assembled might be given different tasks or their work assigned different weightings, as has been practiced in other domains [30]. Such experiments would need to be designed in domains where there are a sufficient number of experts to allow for multiple conditions and multiple trials.

Second, another important test of a creative design happens after it is implemented. That is, eventually the consumer or user evaluates the product for purchase and ongoing use. How could such tests be run? For many products, such as the chair example here, production tests are probably infeasible. But there exist sites that specialize in customized manufacturing and online sales of inexpensive objects (e.g. [5]), and so artifacts might be designed by the crowd and offered for sale in those environments.

Third, the system might extend its use of collaborative mechanisms: The coordination between participants provided in the current system is minimal. That is, no participant interacts directly with another. Instead, only some participants see the previous work of two others: all coordination is mediated through the designs. The independence of the participants makes possible large-scale parallel efforts, and guards against the many problems of tightly coupled collaborations, as described in studies of brainstorming [21, 25, 26]. But more interaction might be useful – for motivation, for clarification, and for deeper idea integration. Thus future studies might try a variety of ways of linking participants to each other: for example, through game-like interfaces that allow for two-way interaction (e.g. [1]), or through structured dyadic conversation (e.g. [42]), or through the communication patterns used in open source environments (e.g. [9, 33, 39]).

CONCLUSIONS

We asked if the crowd functions as a set of cooks or cobblers. We conclude that the systems' users were more like cobblers: they combined with each other to form a creative mind. That is, the generation 3 designs were significantly more creative than the generation 1 designs. The sketch combination system also provided a way to study the interrelations between creativity, combination, and tools. We showed that atypical features were favored in combinations, and that a drawing tool influences the evolution of these features.

That we can perform large-scale experiments related to creativity, and contemplate much larger ones, is perhaps the most important general message of this work. Collaboration necessarily takes place beyond the individual: in this case, it took place among over a thousand individuals, mediated entirely through the designs they produced. Such crowd-based systems are built in the hopes of catalyzing new kinds of collective minds.

ACKNOWLEDGMENTS

Funding for this research was provided by the National Science Foundation, awards IIS-0855995 and IIS-0968561. We are thankful for the suggestions of our colleagues Yasuaki Sakamoto and John Voiklis. This paper is taken in part from a dissertation to be submitted in partial fulfillment of the requirements for the Ph.D. degree at Stevens Institute of Technology.

REFERENCES

1. von Ahn, L. and Dabbish, L. Labeling images with a computer game. In *Proc. CHI 2004*, ACM Press (2004), 319-326.
2. Amazon Mechanical Turk. https://www.mturk.com/mturk/welcome.
3. Amabile, T. M., R. Conti, H. Coon, J. Lazenby, and M. Herron. Assessing the work environment for creativity. *Academy of Management Journal*, 39 (1996), 1154 – 1184.
4. Benkler, Y. *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press, New Haven, CT, USA, 2006.
5. Brabham, D. Moving the crowd at Threadless: Motivations for participation in a crowdsourcing application. *Information, Communication & Society,* 13 (2010), 1122 - 1145.
6. Campbell, D. T. Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review*, 67 (1960), 380 - 400.
7. Deb, K. *Multi-Objective Optimization using Evolutionary Algorithms*. Wiley, Chichester, UK, 2001.
8. Finke, R., Ward, T., and Smith, S. *Creative Cognition: Theory, Research, and Applications*. MIT Press Cambridge, MA, USA, 1996.
9. Fischer, G. Social Creativity: Turning Barriers into Opportunities for Collaborative Design. In *Proc. PDC2004*, 152-161.
10. Gero J. S., and Maher M. L. (eds) *Modeling creativity and knowledge-based creative design*. Lawrence Erlbaum, Hillsdale, NJ, USA, 1993.
11. Goldberg, D. E. *Genetic Algorithms in Search, Optimization and Machine Learning*. Kluwer Academic Publishers, Boston, MA, USA, 1989.
12. Goldschmidt, G, and Litan Sever A. From text to design solution: inspiring design ideas with texts. In *Proc. ICED 2009.*
13. Google Docs Drawing Application. http://docs0.google.com/demo/edit?id=scACRQaIm3t83kVISWPhWfrqx#drawing.
14. Heit, E. Influences of prior knowledge on selective weighting of category members. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24 (1998), 712-731.
15. Holland. J. H. *Adaptation in Natural and Artificial Systems.* University of Michigan Press, Ann Arbor, 1975.
16. Hu, C., Bederson, B.B., and Resnik, P. Translation by iterative collaboration between monolingual users. *In Proc. of Graphics Interface 2010*, Canadian Information Processing Society (2010), 39-46.
17. Kittur, A. Crowdsourcing, collaboration and creativity. *XRDS: Crossroads*, December (2010).
18. Kittur, A., Chi, E. H., and Suh, B. Crowdsourcing user studies with Mechanical Turk. In *Proc. CHI 2008*, ACM Press (2008), 453-456.
19. Kosorukoff, A. Human based genetic algorithm. In *Proc. IEEE Conference on Systems, Man, and Cybernetics,* 2001, 3464-3469.
20. Little, G., Chilton, L. B., Goldman, M., and Miller, R. C. Exploring iterative and parallel human computation processes. In *Proc. of the ACM SIGKDD Workshop on Human Computation* (2010), 68-76.
21. Lorge, I., Fox, D., Davitz, J., and Brenner, M. A. Survey of Studies Contrasting the Quality of Group Performance and Individual Performance, 1920-1957, *Psychological Bulletin* 55, 6 (1958), 337-372.
22. Maher, M. L. Design Creativity Research: From the Individual to the Crowd. *Design Creativity* 2010.
23. McDonald, D. W. Social Computing: A Research Topic in Search of Disciplinarity. *In the Research Directions for Social Computing Workshop at the ACM* International Conf. on Supporting Group Work (2007).

24. Mednick, S. A. The associative basis of the creative process. *Psychological Review* 69, 1962, 220-232.

25. Mullen, B., Johnson C. and Salas E. Productivity Loss in Brainstorming Groups: A Meta-Analytic Integration, *Basic and Applied Social Psychology* 72, 1 (1991), 3-23.

26. Nunamaker, J., Dennis, A. R., Valacich, J. S., Vogel, D. R. and George, J. F. Electronic Meeting Systems to Support Group Work, *Communications of the ACM* 34, 7 (1991), 40-61.

27. Osborn, A.F. *Applied imagination: Principles and procedures of creative problem solving (3rd Edition).* Scribners, New York, USA, 1963.

28. Palmeri, T. J., and Nosofsky, R. M. Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, & Cognition,* 21 (1995), 548-568.

29. Quinn, A. J., and Bederson, B. B. Human Computation: A Survey and Taxonomy of a Growing Field, *CHI 2011,* ACM Press (2011).

30. Raykar, V. C., Yu, S., Zhao, L. H., Valadez, G. H., Florin, C., Bogoni, L., and Moy, L. Learning from crowds. *Journal of Machine Learning Research*. 11, 7 (2010), 1297-1322.

31. Rosch, E. H. Principles of categorization. In E. Rosch and B. Lloyd, eds. *Cognition and Categorization.* Lawrence Erlbaum Associates, 1978, 27-48.

32. Ross, J., Irani, L., Silberman, M. S., Zaldivar, A., and Tomlinson, B. Who are the Crowdworkers? Shifting Demographics in Mechanical Turk. *Ext. Abstracts CHI 2010,* ACM Press (2010).

33. Sack, W., Détienne, F., Burkhardt, J-M., Barcellini, F., Ducheneaut, N., and Mahendran, D. A methodological framework for socio-cognitive analyses of collaborative design of Open Source Software. *Distributed Collective Practices Workshop at CSCW,* (2004).

34. Sakamoto, Y., and Love, B. C. Vancouver, Toronto, Montreal, Austin: Enhanced oddball memory through differentiation, not isolation. *Psychonomic Bulletin & Review,* 13 (2006), 474-479.

35. Simonton, D. K. Scientific creativity as constrained stochastic behavior: The integration of product, person, and process perspectives. *Psychological Bulletin,* 129 (2003), 475-494.

36. Smith, S. M., Ward, T. B., and Schumacher, J. S. Constraining effects of examples in a creative generation task. *Memory & Cognition,* 21 (1993), 837-845.

37. Spears, W. M. *Crossover or Mutation? In Foundations of Genetic Algorithms 2*. L.D. Whitley (ed). Morgan Kaufmann, CA, USA, 1993.

38. Suwa, M. and Tversky, B. What do architects and students perceive in their design sketches? A protocol analysis. *Design Studies* 18, 4 (1997), 385-403.

39. Tech Guru, *Open Source Cars*, http://www.techquark.com/2010/05/ open-source-cars.html, May 26, 2010, accessed January 10, 2011

40. Thagard, P. *Conceptual revolutions.* Princeton University Press, USA, 1992.

41. Tversky, B. and Chou, J. Creativity: depth and breadth. *ICDC 2010.*

42. Voiklis, J. *A Thing Is What We Say It Is: Referential Communication and Indirect Category Learning*, Ph.D. dissertation, Columbia University, New York, USA, 2008