Why creative design is important.

Enhancing Creative Design via Software Tools

An agent (individual human, group, organization, or software agent) has a problem if there is a mismatch between the current state and some desired state and no preexisting algorithm is available for transforming the current state to the desired state. Problems can be categorized according to whether starting state, ending state, and allowable transformations are well-defined or ill-defined. Proving a Euclidean Theorem is an example of a completely well-defined problem while designing a house, a useful software system, a new business process, and writing a teaching story are examples of completely ill-defined problems. There are many software tools for helping people solve various types of well-defined problems; however, relatively few tools exist for helping people solve ill-defined problems such as design problems. Yet, design problems are extremely high leverage. For instance, design errors, whether in software, drugs, business processes, or automobiles are extremely costly, compared, for instance, with coding or manufacturing errors. Conversely, effective and innovative designs are extremely lucrative. Although creativity is certainly demonstrated in well-defined problems such as theorem proving and chess, ill-defined problems such as design require creativity quintessentially.

The world is changing rapidly but the ability of people to design creatively has not improved in any noticeable way. As a result, there is a widening gap between the degree of flexibility and creativity needed in order for individuals and organizations to adapt and their capacity to do so [4]. Failure to innovate is not random, but can generally be ascribed to three main difficulties:

- Individuals or groups do not engage in effective and efficient processes of innovative design;
- The necessary skills, talents, and *knowledge sources* are not brought to bear on the problem; and
- The appropriate level, type, and directionality of motivation are not brought to bear.

Laboratory [3, 10] as well as field research [2, 5] over the last several decades has established the major process difficulties of individuals and groups are mainly due to a limited number of errors and these errors can be avoided or ameliorated by providing appropriate structure.

The appropriate overall structure for innovation has several sub-steps and structure is necessary both to help facilitate the progress through these steps and to help guide the separate sub-steps; distinct guidelines are appropriate for each of these sub-steps [8, 9]. As an example of a common failure in the overall control structure, people typically fail to spend sufficient time in the early stages of design: problem finding and problem formulation [7]. As an example of a common failure during a specific stage of innovative design, people often bring critical judgment into play too early in the idea generation phase of problem solving. As another example, empirical evidence shows that people's behavior is path-dependent and they are often unwilling to take what appears to be a step that undoes a previous action even if that step is actually necessary for a solution [10].

Regarding the second issue (bringing to bear necessary skills, talents and knowledge sources), while software tools cannot fully substitute for human experts,

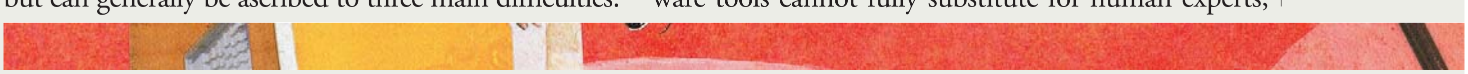

By John C. Thomas, Alison Lee, and Catalina Danis

evidence suggests individuals have a large amount of relevant implicit knowledge they often will not bring to bear on a problem and that providing appropriate strategies [10], knowledge sources [11] or representations [3] can significantly improve an individual's effectiveness in problem solving and innovation. In controlled laboratory experiments, subjects did significantly better in a temporal design task when they used a spatial representation; yet, very few subjects spontaneously adopted such a representation [3]. The impact of felicitous representations, however, is not confined to laboratory demonstrations. Speech research advancements accelerated greatly when waveforms were largely replaced with speech spectrograms and Feynman diagrams allowed similar breakthroughs in atomic physics. Regarding motivation, it may be often assumed that increased motivation will result in better performance. However, the reality is considerably more complex. The effect of increased motivation interacts with personality, complexity, novelty, and whether the motivation is intrinsic or extrinsic [1].

Some Proposed Creativity Tools

We believe several factors have converged that will now allow us to use the power of the computer, combined with a knowledge of how to impact creativity, to provide increases in creativity during the design process. Software tools can help address all the preceding issues by organizing the creative process (both overall and sub-steps), in providing people with a rich array of appropriate strategies, knowledge sources, and representations, and in providing the right level and type of motivation for the task at hand.

The ability to cross boundaries between fields—to apply an analogous solution to a new problem—is one of the hallmarks of creative problem solving. One interesting research question is how to augment or support this kind of analogical reminding and reasoning using technology. Recent work has helped outline several possible software applications for analogy-based creativity support, providing the tools necessary to explore this question in real-world systems by collecting and analyzing abstract planning strategies from a wide variety of domains [6]. We believe it likely that people within a given field tend to use a number of strategies common to that field but are likely to be unaware of potentially applicable strategies from entirely different fields. For instance, a medical doctor or a researcher in a pharmaceutical company may typically think of various strategies for "fighting" disease with a treatment. Suppose instead that the doctor examines strategies from sales. Then, the question becomes "How can you convince the prospect (the disease) that it is better to go elsewhere or do something else?" A strategy in sales might be to find out what the customer likes in the current situation as well as finding out what they see as unmet needs.

Another sales strategy might be to "ease the pain of transition," for example, by paying in installments. Applied to a medical situation, this might lead to thinking about trying to understand the conditions most favorable to this disease (perhaps a particular chemical environment) and then rather than trying to "kill" the disease by providing a totally hostile environment, instead provide an environment the disease "prefers" outside the patient along with a gradient pathway that allows a graceful exit toward the preferred environment. Alternatively, one might conceptualize the problem as a transportation problem. A disease pathogen is not an evil enemy to be destroyed. Rather, the pathogen simply consists of a large number of tiny objects that need to be transported from region A (the site of infection) to region B (somewhere outside the patient's body). Mechanical engineers, civil engineers, and operations research experts have various methods of solving transportation problems. How might these be applied to disease? Can a "hook" be attached to each pathogen making it easy to pick them up as on some assembly lines?

The suggestion is that a book of strategies from a wide variety of fields be indexed according to an abstract general description of the problem someone is trying to solve. To address a specific problem, a person can find its abstract description and then find numerous strategies that might be applied to this general type of problem but seldom have been.

Providing people engaged in design problem solving with a wider array of potential strategies is just one

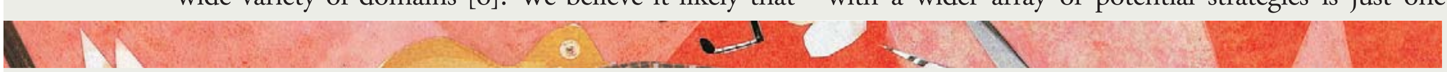

avenue for the machine augmentation of creative cognition by applying additional fields of knowledge. Other methods include teaching stories, and providing users with a wide range of potentially useful representations. We've developed the Narrative Insight Method for bringing together small groups of people with various levels of expertise in a certain area, inducing them to share stories, recording these stories, and turning them into teaching materials.

Another approach builds on the notion that people know much more than they know they know. Creativity can be enhanced by helping people "find" relevant knowledge. An example prototype Board of Directors can be used both to help generate ideas and to help people critically examine ideas from various perspectives (see www.research.ibm.com/knowsoc/prototypes_directors.html). Yet another tool we have envisioned is the Personal Knowledge and Experience Inventory. Here, the person is led, by a series of evocative questions, to make various maps of their potential sources of personal knowledge based on their experiences, the places they have been, the people they have known, the courses they have taken, the projects they have worked on, skills they have mastered and so forth. When faced with a novel task, the person now scans the map for potentially relevant sources. We believe such a map may serve a useful motivational purpose as well by reminding people of past successes. The construction of such a map necessarily relies on the memory of the user and requires considerable manual input. In the near future, such a map could be largely constructed automatically from various online records and transactions.

Another approach we are exploring is to develop Pattern Languages as tools of thought; specifically, we are developing patterns in the sociotechnical arena and believe such patterns may help people find, formulate, and solve complex problems requiring creativity. A Pattern is the named solution to a recurring problem. Typically, a Pattern includes an evocative picture, an abstract, a statement that describes the context in which this problem is likely to occur, a statement of the problem, an analysis of the problem in terms of countervailing forces, a solution to the problem, and a diagram that summarizes the solution. It may also include references, linkages to related patterns, a rating for how important the Pattern is, a section describing the resulting context, and known uses.

Our Patterns will also include, in some cases, pointers to software components that are meant to instantiate or support that Pattern. A Pattern Language is a lattice of Patterns that together claims substantive coverage for a coherent field. Pattern Languages have been applied to architecture, object-oriented programming, business processes, and human-computer interaction. Unlike a formula or an algorithm, a Pattern is meant to stimulate thought without providing an inflexible, possibly inappropriate structure. We are using the Pattern Language we are developing to provide guidance for the design of a Web site to foster a community of practice concerned with sociotechnical issues.

Any system of representation is a kind of double-edged sword. It makes certain things easier to see and manipulate and it makes other things more difficult. We believe that using the Pattern Language and associated tools will result, on average, in more creative and more competent sociotechnical systems than would otherwise be the case. However, some designers might use Patterns in an overzealous or overly rigid fashion and in so doing miss potentially creative solutions. The style of language used in the Patterns, the philosophy of how to use Patterns, and tools that guide the designer can help minimize this danger. Pattern Languages are also meant to evolve over time and be thought-provoking even if inapplicable. Moreover, not only may different community members share their various experiences by suggesting different individual Patterns. We also want to provide a means for community members to use their individual talents collectively to build single Patterns. Some individuals may be particularly good at finding recurring problems; others at organizing into forces; others at suggesting solutions; others at finding evocative images; others at editing; still others at recommending excellent names. We feel that on balance, Pattern Languages promise increases, not decreases, in creativity, both for individuals and for communities.

Helping to Improve Creativity Over Time

Practice in a complex field does not necessarily lead to better performance. Many golfers' handicaps, for instance, do not improve with time beyond the first two years. The behavior of the golf ball is influenced by many factors and correct attribution to the cause(s) of error is difficult. For instance, a putt may go awry because the golfer has misread the slope of the green or the grain. The golfer may have read the putt correctly but hit the ball poorly. Even a perfect putt may go astray because of irregularities in the green or changes in the wind. Unfortunately, the most salient feedback experienced by the golfer is simply the path the ball takes. But the golfer may easily make a wrong attribution about what error or combination of errors led to that path.

A similar, but even more challenging problem arises when a person attempts to improve their creativity. The evaluation of a product of creativity such as a novel, scientific finding, work of music, or painting may vary quite markedly over time. In particular, something may be first viewed as ridiculous, idiotic, outrageous and not be valued as "good" until years later; indeed, it may not be valued positively until after the creator's death. Under these circumstances, a person clearly cannot improve their creativity on the basis of extrinsic feedback! But even in the case where a creative work is seen as creative immediately, attribution is still complex and subject to error. What made this particular work better than others? Was it the original conception or the execution? It seems that highly creative people, like many star athletes, actually end up with what seems to the outsider fairly particular and superstitious behavior patterns.

In the case of golf, significant improvements in putting can come about via the use of devices that decouple feedback about the various components of putting. For example, you can practice your putting stroke against a straight surface to ensure the swing path is straight. Can we construct similar methods of feedback decoupling for something as complex as creativity? We believe it may be possible. The same normative process model that can be useful in providing overall guidance can also be used as a diagnostic aid. This can be applied in two modes. In one mode, the person or team is using the normative process model to help design a real system, artifact, or process. In the course of using the model, information is collected and this information can be analyzed to give clues about the effectiveness of various stages and transitions.

We are not supposing that an effective and correct analysis can be conducted automatically, but a combination of data mining, question asking, and reflection will produce a more disentangled account of what worked well and what worked poorly in a complex and lengthy process. For instance, suppose a team has designed and developed a system to support collaboration. As the team progresses through each stage of a defined process, statistics can be collected and displayed. These can serve as the basis for reflection. During the idea generation phase, for example, the team may choose to engage in a metaphor-enriched brainstorming session to generate ideas. How many ideas were generated? How were they distributed over time? Did they end the process at a point where very few new ideas were being generated or considerably before that point? How does this pattern compare with that of the last project? The notion is that the team can use the display of such data as a focus for a self-reflective discussion.

In a second possible mode of use, a person or team could run quick standardized tests of creative components under various conditions to determine what works. It would be both impractical and meaningless to write the same novel twice under different conditions—once by working in the morning and once working in the evening—in order to determine what works better. But one could take an adaptive word fluency test in the morning and in the evening. The results would not scientifically prove which time of day would work better for creative writing, but it might provide convincing guidance for some. It may be more important to converge on some consistent method for writing quickly than it would be to try to discover the optimum method, particularly in light of the virtual impossibility of the attribution (credit assignment) task.

A major challenge of all these approaches is how to provide aid that resonates with the natural processes in which individuals and groups are engaged. Intervening in the design process at the wrong time could be both detrimental to the creative process and decrease the probability that designers will want to use such tools. Empirical research will be necessary to determine the most appropriate synergistic interaction paradigms; however, our hunch says to provide a rich palette of tools for individuals and groups to use when they deem appropriate. Further empirical research (laboratory and field) on the spontaneous use of these tools may enable us to develop reasonable algorithms that are more activists and interventionists. C

References

1. Amabile, T.M. *The Social Psychology of Creativity*. Springer-Verlag, New York, 1983.
2. Carroll, J., Thomas, J., and Malhotra, A. A clinical-experimental analysis of design problem solving. *Design Studies 1*, 2 (1979), 84–92.
3. Carroll, J., Thomas, J., and Malhotra. Presentation and representation in design problem solving. *British Journal of Psychology 71*, 1 (1980), 143–155.
4. Drucker, P. *Managing in a Time of Great Change.* Truman Talley Books, New York, 1995.
5. Gordon, A. Playing chess with Machiavelli: Improving interactive entertainment with explicit strategies. *AAAI Spring Symposium*, Stanford, 2001.
6 Olson, M. and Bly, S. The Portland experience: A report on a distributed research group. *International Journal of Man-Machine Studies 34* (1991), 211–228.
7. Sobel, D. *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of his Time.* Penguin, New York, 1995.
8. Stein, M. *Stimulating Creativity*. Academic Press, New York, 1974.
9. Thomas, J. An analysis of behavior in the hobbits-orcs problem. *Cognitive Psychology 6* (1974), 257–269.
10. Thomas, J., Lyon, D., and Miller, L. *Aids for Problem Solving.* IBM Research Report, RC-6468, 1977.
11. Thomas, J. Problem solving by human-machine interaction. In K. Gilhooly, Ed. *Human and Machine Problem Solving*. Plenum, London, 1989.

John C. Thomas (jcthomas@us.ibm.com) is a member of the research staff at the IBM T.J. Watson Research Center in Hawthorne, New York.
Alison Lee (alee@us.ibm.com) is member of the research staff at the IBM T.J. Watson Research Center in Hawthorne, New York.
Catalina Danis (danis@us.ibm.com) is member of the research staff at the IBM T.J. Watson Research Center in Hawthorne, New York.