How the *Obscure Features Hypothesis* Leads to *Innovation Assistant* Software

Tony McCaffrey
Cognitive Psychology
University of Massachusetts Amherst
Amherst, MA 01002 USA
amccaffr@psych.umass.edu

Lee Spector[1]
Cognitive Science
Hampshire College
Amherst, MA 01002 USA
lspector@hampshire.edu

Abstract

A new cognitive theory of innovation, the *Obscure Features Hypothesis* (OFH), states that almost all innovative solutions result from two steps: (1) noticing a rarely noticed or never-before noticed (i.e., obscure) feature of the problem's elements, and (2) building a solution based on that obscure feature (McCaffrey 2011). Structural properties of the human semantic network make it possible to locate useful obscure features with a high probability. *Innovation Assistant* (IA) software interactively guides human users to the most likely obscure features for solving the problem at hand.

Introduction

The OFH articulates a core principle for solving the problems requiring innovation used in psychology experiments (i.e., insight problems) as well as real world problems in engineering and design. As an example of an insight problem, consider the *Two Rings Problem* in which you have to fasten together two weighty steel rings using only a long candle, match, and a two-inch cube of steel. Melted wax is not strong enough to bond the rings together so the solution relies on noticing the obscure feature that a wick is a string that can be extricated from the candle by scraping away the wax on the cube of steel before tying the rings together.

Two lines of evidence suggest that all insight problems used in the psychology literature follow the OFH. First, the key features for solving known insight problems are not listed among that object's common associates in the Association Norms Database (Nelson et al. 1998) and are thus obscure (e.g., *string* is not listed as a close associate of either *candle* or *wick*). Second, Chrysikou (2006) had subjects list features from the key objects used in insight problems and found that the key feature from these objects was listed by only 9% of the subjects and are thus obscure.

The OFH opens up a research program to improve innovation based on two questions. What inhibits humans from noticing the obscure? What techniques can help overcome these sources of inhibition? In general, the normal processing of our perceptual, motor, and semantic systems leads us to notice typical features. Our everyday experience entrains us towards the typical, which is efficient for everyday life but an enemy of innovation. Moving humans from the typical to the obscure seems to be a promising way to improve innovation. Immediately, however, a major challenge presents itself. The number of features of an object (e.g., a candle) is potentially unlimited (e.g., my aunt collects candles, candles are generally smaller than a toaster) and perhaps infinite (e.g., a candle generally weighs less than 100 pounds, weighs less than 101 pounds, etc.: Murphy and Medin, 1985). What techniques can lead humans to obscure features that have a high probability of leading to innovation? Specifically, can a computer program steer human users through the potentially infinite search space to the most promising areas of the human semantic network?

McCaffrey (2011) analyzed the semantic networks of all insight problems used in the psychology literature plus a collection of engineering design problems and discovered several important structural properties relevant for innovation. In diagramming a semantic network graph for an object, a box is drawn around the object name and all its commonly noticed features. Analysis shows that the key obscure features for solving the examined problems are just outside the box. The *Just Outside the Box Hypothesis* (JOTB) posits that the key feature for most innovative solutions is one or two steps outside the box of commonly noticed features. Searching through a semantic network, a computer could help locate those concepts that are at the proper distance from the original concept according to the JOTB. This technique helps narrow the search a bit, but another principle related to the direction of the search from the original concept can narrow the search even more.

Although there are potentially an infinite number of features for an object, there is a finite taxonomy that characterizes the most prevalent types of features and relations for an object that are known to be useful for innovation. McCaffrey (2011) has developed such a taxonomy called the *Feature Type Taxonomy* (FTT) and uses it to characterize the most promising direction to search in from the initial concept. The most promising direction will vary depending on the kind of problem being addressed. For example, an examination of all known insight problems involving concrete objects reveals that the key obscure feature for 68% of them resides in noticing the material composition of the parts (e.g., examining the material make-up of a wick and wax for the *Two Rings Problem*) (McCaffrey 2011). So, given a concrete object insight problem, looking at the material category of the parts is a good idea. Other feature/relation types are useful for other innovative tasks. For example, a candle is often present during romantic events such as dinner dates. A comedian might take this information and create a funny situation by placing a candle in an unromantic event such as divorce court proceedings. In other words, the kind of innovative problem stipulates which feature types have the highest probability of being useful (e.g., Material or Events). For every type of problem articulated, we can create a probability distribution that measures the likely usefulness of any given feature type.

Together, the JOTB and the probability distributions formed from the FTT help specify the best distance and direction to search in from the initial concept for a given problem. The JOTB operationalizes the idea of innovation standing between order and randomness (Watts 1999). A concept inside the box is too strongly connected to the initial concept (i.e., too much order) while a concept many links away is too weakly connected (i.e., too much randomness). A concept just outside the box is on the boundary between order and randomness, which seems to be ideal for innovation.

Innovation Assistant Software

We created a software prototype, called an *Innovation Assistant* (IA), which guides human users to the most promising obscure features based on problem type. Our overall goal is to create a human-machine interaction that is more innovative than a human working in isolation. Whereas traditional AI programs attempt to solve problems completely on their own, the IA guides humans to notice the obscure features that we tend to overlook. The human user must then construct the solution based on the key obscure feature.

Currently, the prototype is instantiated with the probability distribution of the FTT for concrete object insight problems. Further, presently it only executes a technique to help users notice the material make-up of an object's parts; but other techniques for noticing the obscure members of other feature/relation types are ready for implementation.

Evidence: Lab and Real-World Problems

We tested the IA prototype on the eight insight problems used on human subjects in the experiments of McCaffrey (2011). Because the material composition of the parts is central to solving these problems, the IA asks the user to enter this information (e.g., "Can a wick be decomposed further into parts?" and "What material is a wick made of?"). While the IA successfully leads users to the key obscure features for these problems, on three of the problems the IA exceeded its goal and actually located the key feature and how it could be used to solve the problem. For example, for the *Two Rings Problem*, after the user enters the goal "fasten rings," enters the candle's parts as "wick" and "wax," and answers the IA's question about what material a wick is made of, the IA is able to generate how the string can solve the problem. The IA suggests: "a candle's wick is made of string, which might be able to tie the rings." Basically, the synonyms of the goal verb "fasten" are intersected with the verbs related to the potential actions of the known objects and parts. In this case, the verb "tie" is a synonym of the goal verb "fasten" so the IA can make the connection and solve the problem. In sum, the IA is designed to help users notice useful obscure features that are often overlooked. However, the IA can periodically go above this standard and actually solve the problem—if the user has not already solved it.

Most real-world problems require noticing features other than the material make-up of an object's parts, so we will present how several of the other techniques have helped with real-world problems.

First, a company challenged us to create an idea to detect roadside bombs. Most proposed solutions focus on either detecting the bomb itself through its metal or disrupting the electronics of the bomb. After shifting focus from the bomb's physical make-up to the events that the bomb is involved in (as suggested by the probability distribution for real-world engineering problems), a technique guided the user to articulate a hierarchy of these events and their sub-events (e.g., building the bomb, transporting the bomb, burying the bomb, and detonating the bomb). This hierarchy led to an idea for a mechanical method that could reliably detect the displacement of the dirt above a buried bomb rather than trying to detect the bomb itself. No further information will be disclosed at this time due to this problem being an ongoing open military problem. The technique performed its job by shifting attention to an obscure but promising feature type.

Second, another company gave us the unsolved problem of adhering a coating to Teflon. After switching focus to the goal, a technique helped users to reveal the common assumptions of the verb *adhere:* including (1) direct contact between (2) two surfaces using a (3) chemical process. Questioning these assumptions led to the novel and workable idea of sticking something "through" Teflon to a magnetic surface beneath the Teflon (i.e., a sandwich of three surfaces in which the coating indirectly sticks to the Teflon due to its attraction to the magnetic surface). Of course, the coating would need to possess the proper make-up in order to induce sticking. Again, the technique shifted our focus to an obscure way to adhere things together.

Third, at the beginning of the BP oil spill in late April 2010, we worked with the goal "contain oil" and a technique helped suggest a pipe wide enough to enclose the spill area. After making this "pipe" out of plastic, several engineers verified the workability of this plastic "curtain" that reached from the ocean floor to its surface and would guide the spilled oil to the top in a confined space where it could be dealt with. We submitted our idea to the BP disaster website as did other engineers.

Fourth, two candle companies verified the novelty of our candle designs that were created using a technique which helped us notice that candles are built out of one piece of wax and are motionless when they burn. We created novel candle designs consisting of multiple pieces of wax as well as designs in which the entire candle moves based on its own dynamics. These designs are under consideration by two candle companies so will not be disclosed.

Finally, the U. S. Post Office solicited the public for money-saving ideas. Working with the goal "deliver mail," a technique helped us to consider places to "pick up mail." The option of drive-through windows led to the idea of car commuters signing up to pick up their mail at the drive-through window where they get their daily coffee. Train commuters could pick up their mail at the subway stations they frequent every weekday. We added more logistics (e.g., banks of post office boxes at key pick-up locations), but estimates indicate that this method could eliminate up to one-third of daily home deliveries without delaying the reception of this mail.

The Next Phases

Given the promise of the IA prototype on insight problems and other techniques on real-world problems, we are moving forward in two directions. First, after implementing our other techniques, we will conduct controlled experiments with engineering and design students to further test the IA's effectiveness for improving innovation. Students will be given a design task. Students in a control group will complete the design problem on their own. Students in an experimental group will use the IA to help complete the design problem. Judges will rate the creativity and workability of the designs. We hypothesize that the IA users will produce designs deemed by the judges to be more creative than the designs produced by the control group.

Second, we will record all the runs of the interactions between human users and the IA program in order to try to detect patterns in the use of the program that led to successful innovation, as well as interactional patterns common to unsuccessful use. To all these run traces, we will apply a genetic programming (GP) system (Koza, 1994). Note that GP systems have been shown to be effective on a variety of difficult problems, including pattern detection (Koza, 1994). Specifically, when all the parameters to explore are known, GP systems are competitive with humans on design problems thought to require innovation (Spector, 2008). A GP system is therefore a good method for evolving and detecting patterns in the traces of the IA runs.

Further, the IA runs involve accessing the human semantic network as embodied by WordNet (Miller 1995). A GP system will be used to detect patterns in the semantic network searches that led to an innovative (i.e. novel and useful) connection. Using this method, we may discover other structural properties of the semantic network that are important for innovation. Each structural property we uncover will permit us to craft a new heuristic to exploit the semantic network for the benefit of innovation.

All our current research focuses on the first step of the OFH (i.e., helping humans notice obscure features) and ignores the second step (i.e., building a solution based on an obscure feature). Because solutions can become intricate causal networks, human cognitive limitations may inhibit us from piecing together the solution once we have noticed all the relevant features. Future research will address this possible limitation to human innovation.

Conclusion

Initial evidence from laboratory insight problems and a collection of real-world design problems points to the potential for the IA approach to improve human innovation. Testing of the IA in controlled experiments will lead to scientific evidence measuring the effectiveness of the IA as well as the truth of the OFH, its underlying principle. Discovery of new semantic network properties relevant to innovation will lead to new heuristics that improve innovation. In sum, the IA approach is a promising one for creating software that can improve human innovation based on the psychologically plausible theory called the OFH.

References

Chrysikou, E. (2006). When shoes become hammers: Goal-derived categorization training enhances problem-solving performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(4), 935-942.

Koza, J.R. (1994). *Genetic Programming II: Automatic Discovery of Reusable Programs*, MIT Press.

McCaffrey, A. J. (2011). The obscure features hypothesis for innovation: One key to improving human innovation. Unpublished doctoral dissertation, University of Massachusetts Amherst.

Miller, G. A. (1995). WordNet: A lexical database for English. *Communications of the ACM,* 38(11), 39-41.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.

Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (1998). The University of South Florida word association, rhyme, and word fragment norms.

Spector, L. (2008). Introduction to the special issue on Genetic Programming for human-competitive designs. In AI-EDAM: *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 22(3), 183-184.

Watts, D. J. (1999). *Small worlds: The dynamics of networks between order and randomness*, Princeton University Press.

---

[1] This material is based upon work supported by the NSF under Grant No. 1017817. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the NSF.