New Perspectives on Human Problem Solving

Robert L. Goldstone[1] and Zygmunt Pizlo[2]

In November 2008 at Purdue University, the 2nd Workshop on Human Problem Solving was held. This workshop, which was a natural continuation of the first workshop devoted almost exclusively to optimization problems, addressed a wider range of topics that reflect the scope of the Journal of Problem Solving. The workshop was attended by 35 researchers from around the world. The program of the workshop is available at: http://spiderman.psych.purdue.edu/problem_solving/workshop08/. There were two full days of talks and posters, covering insight, education, decision making and causality, combinatorial optimization and applications. The level of excitement was high, and the lively and informative discussions confirmed that, despite the fragmentation of the field, there is something special and unique about problem solving that is likely to help define the field, almost exactly 100 years after the Gestalt Psychologists established it for the first time. Defining the field is obviously important, but doing solid research and disseminating the results is absolutely critical. This special issue provides a nice sample of the current work on human problem solving illustrating a variety of experimental tools and conceptual frameworks.

Why has problem solving remained one of the most important topics in human cognition? Problem solving is the activity in which people show their uniquely human gift for flexible cognition and adaptation. We are not always shackled victims of habit. We can think flexibly about situations that we have never come across before, and arrive at apt and fresh responses to them. We can change how we represent problems. We can come up with novel strategies for solving problems. In other words, we can program ourselves to solve problems.

These capabilities still distinguish human problem solving from machine intelligence efforts. Long after humans have been replaced by machines in tasks where automatic algorithms can be deployed to effectively accommodate known situations, genuine problem solving requiring the creation of thinking tools, representations, strategies, and meta-cognitive awareness will still be a privileged domain of humans. Over time, we are getting a better understanding of what tasks require truly flexible problem solving

[1]Indiana University; [2]Purdue University

capacity, and what tasks can be solved using existing pre-determined strategies and representations. Surprisingly, world-class chess playing has turned out to be a domain in which fundamentally new board representations need never be created. All that is required is an extremely rapid search over cleverly pruned problem spaces generated by pre-established representations (still, this brute force approach is likely to fail if chess is played on a larger board, say 16x16, rather than the conventional 8x8 board. Larger boards will increase the search space, which may prevent the machine from playing in "real time"). By contrast, the easier problem (for humans) of discovering a rule that distinguishes the Category A items from the Category B items in each of the four problems shown in Figure 1 is beyond the capacity of any existing computer program, and is likely to be for some time to come. These Bongard problems are difficult for both humans and machines, but decidedly more difficult for machines, apparently because they require the problem solver to come up with novel, initially unforeseen descriptions of the objects. As machines get increasingly adept at search, planning, pruning, and optimization, it will become increasingly important to understand the kinds of problems that cannot be cast in these terms, and how people are nonetheless able to solve them. The article in this issue by Ivan Ash, Patrick Cushen, and Jennifer Wiley provides a critical examination of empirical evidence for the kind of restructuring and description-building implicated by Bongard problems and other insight problems. The authors closely examine claims for restructuring, and develop a useful skeptical stance toward claims like the one we just made above – that restructuring cannot be cast in terms of well-understood problem solving operations like search, planning, and pruning. Their article provides a perspicacious synthesis of the kinds of cognitive processes that can be imputed given different methodological paradigms.

Historically, research strands in problem solving and causality have remained relatively separate within cognitive psychology. However, there is a good argument for connecting them, because determining what causes what in a complex environment is one of the most challenging problems there is. Furthermore, making inferences about causation entails many of the same processes that feature prominently in other forms of problem solving. One such process, investigated by Marc Buehner and Jon May in their article, is searching for and comparing alternatives. They report evidence that is consistent with people comparing alternative possible causes of an event, and inferring which one is, in fact, the cause of the event on the basis of contingency, background knowledge, and temporal delay. Their results are less compatible with purely association-driven accounts. In this respect, their results provide a good justification for linking causal reasoning and problem solving. Both employ a critical process of assessing multiple options in the context of one another. In fact, the strong possibility exists that it is the same comparison and evaluation cognitive engine that determines causal relations and solves many problems for us.

Another area which is closely related to problem solving is decision making. In fact, it is not clear how the two differ. Perhaps difference related to the role of representation.
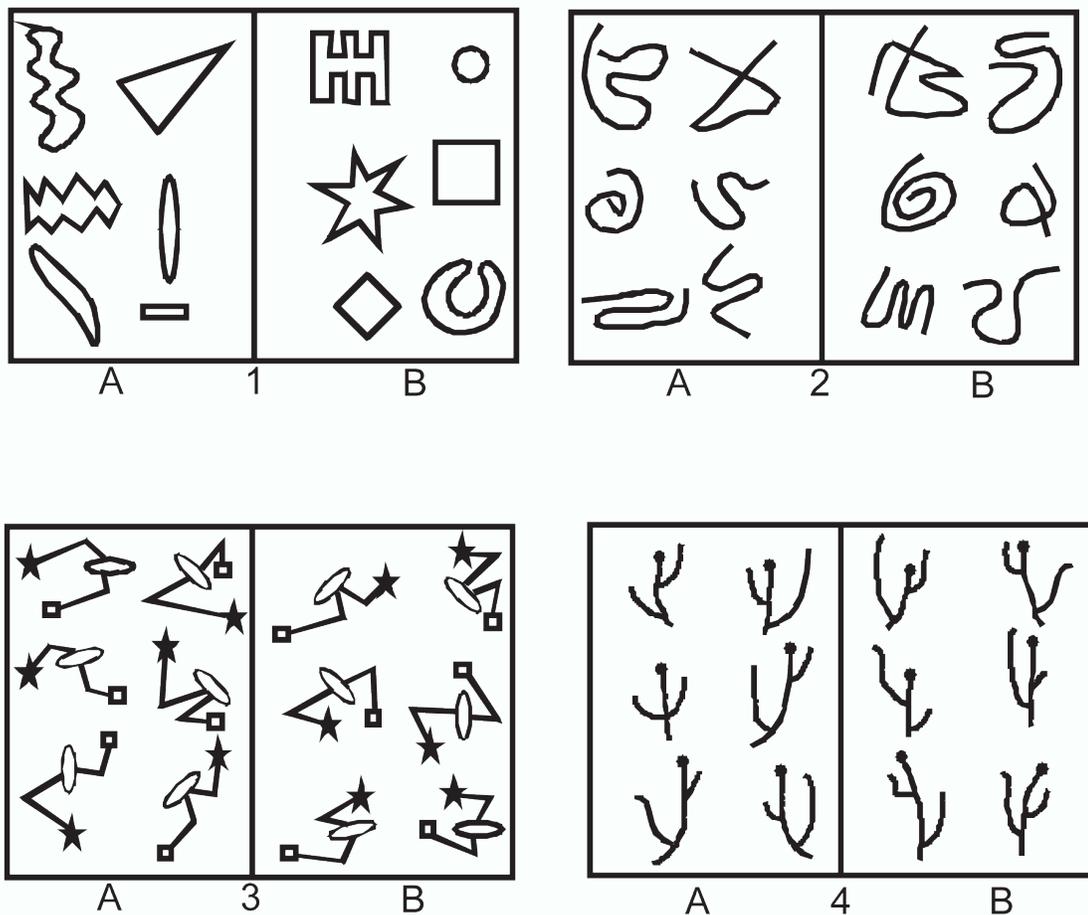
Figure 1. For each of these four problems, try to find the rule that distinguishes the 6 objects belonging in Category A (on the left) from the 6 objects belonging in Category B (on the right). These problems are loosely adapted from some of the Bongard problems presented on Harry Foundalis' web site: http://www.foundalis.com/res/bps/bpidx.htm.

Indeed, theories of decision making rarely explore the role of representation of the decision making situation. However, this is likely to be a matter of emphasis, rather than of some fundamental aspects of the underlying mechanisms. Both, decision making and problem solving often involve a sequence of "moves" that are aimed at bringing the agent to a goal. The sequential decision making process discussed by Michael Yi, Mark Steyvers, and Michael Lee is a good example. They examined the usefulness of particle filter models that have been used recently to build the visual representation of 3D scenes by robots. The fact that the same computational tool is used in decision making and in building representations is not a coincidence. There is something fundamentally similar in problem solving, perception and decision making.

We are encouraged by a renewed trend to connect problem solving with education. Education is very possibly THE "killer application" for problem solving research and cogni-

tive science more generally. It is generally acknowledged that our educational system could be a lot better, and part of what will make it better is understanding how people learn, and apply what they have learned. Accordingly, in this current special issue of Journal of Problem Solving, and in a future issue, we have invited several authors to present work on problem solving in mathematics, science, and engineering educational contexts. In this issue, the article by Daniel Belenky and Timothy Nokes, and another article by Katherine Gunion, Todd Milford, and Ulrike Stege, are excellent exemplars of the fusion of problem solving and education.

Belenky and Nokes study two important factors that have been shown to affect students' learning and transfer of knowledge: concreteness of learning materials and metacognitive stance. There are reasons for thinking that teaching students about the mathematics of combinatorics with concrete materials (animals) could produce either better or worse learning than when more abstract materials (blocks) are used. On the one hand, students may benefit from the familiarity, discriminability, and groundedness of the animals. On the other hand, the mathematics underlying combinatorics is domain-general and does not depend upon whether a problem is instantiated by animals, appliances, or automobiles, and so perhaps more idealized materials allow students to better appreciate the general mathematics. In fact, the authors find that concrete materials yield better transferable knowledge, but only when students were prompted to reflect on their problem-solving processes. One intriguing possibility is that this reflection allowed the students to transcend the specific details of the animals when given the concrete materials. The students could then simultaneously benefit from the grounding provided by the animals' familiarity but not be unduly tied to their specific forms. The authors' research points to the exciting possibility that when concrete materials are combined with efforts to reflect on what is being learned, it may be possible to achieve understandings that are both well-grounded *and* high-reaching.

The specific educational topic tackled by Gunion, Milford, and Stege is the difficult concept from computer science of recursion. Like Belenky and Nokes, these authors are grappling with the issue of how to have students learn what would normally be considered to be a highly abstract concept. Also like Belenky and Nokes, the authors gravitate toward a grounded, experience-based method of instruction (even Belenky and Nokes' "abstract" materials are physical manipulable objects, namely blocks). They achieve positive gains in even young students' understanding of recursion by engaging them in participatory simulations, active programming tasks, games, and identification of visual patterns. One interpretation of Gunion et al.'s successful pedagogy is that perception and action provide seeds from which can grow abstract problem solving capabilities. A more radical alternative interpretation to which we are partial is that problem solving never ceases to be based on perception and action. By this account, abstractions must be cashed out as the deployment of trained perceptual processes and actions. If this account has validity,

then we should not be seeking ways to trump or defeat potentially misleading perception so as to attain pure scientific or mathematical understanding. Instead, we should be trying to *train* our perceptual processes so that they provide us with interpretations of a situation that are congruent with formal understandings. While it might be thought that perception is too constrained to serve this function, an interesting question to ask yourself is whether and how the shapes in Figure 1 *look* different once you have learned the rules that cleanly separate the As from the Bs.

In the 6th paper of this issue, James MacGregor and John Cunningham use rebus puzzles to study insight problem solving. Insight problem solving has always occupied a special place in problem solving literature. The main reason is that the underlying mechanisms are so mysterious. The sudden change from not being able to "see" the solution to discovering it, makes it difficult to study the process that leads to insight. Rebus puzzles seem to offer a solution to this methodological problem because they allow the experimenter to systematically vary the difficulty of the problem. As a result, one is now able to study the nature of the process by examining what happens in a set of problems, rather than in a single one. If the researcher's understanding of the building blocks of rebus puzzles is correct, the results will shed light on the critical aspects that distinguish difficult from easy rebus puzzles. This will surely provide insight into insight itself.