A Mathematical Theory of Creativity (A Grand Challenge)

Simon Colton

Department of Computing, Imperial College, London

'Creativity' is one of those words that some computer scientists avoid at all costs. Part of the reason for this is that detractors of research into Artificial Intelligence have used it against the discipline: "computers will never be intelligent because they can't be creative". Part of the reason is that it is such a mis-used, poorly understood, loaded word, that it seems bound to cause unease whenever used in conjunction with machines. Finally, until recently, it's been difficult to project the word creative onto any existing programs, for various reasons. However, creativity clearly plays a part in human intelligence and cannot be ignored if we are to write programs which exhibit intelligence. I suggest that a mathematical theory of creativity would provide the basis for the building of creative programs, just as logic supplies the basis for deductive programs.

My grand challenge is to devise such a theory. By stating *mathematical* rather than a more formal *computational* theory of creativity, I am acknowledging the fact that certain aspects of the theory may not be well suited to implementation in current computers (and this limitation itself may advance computer science in non-standard directions such as quantum and DNA computing). I believe that two important aspects of a mathematical theory of creativity will be the amalgamation of many forms of reasoning into a theory of creative reasoning and the separation of creative tasks into problem solving, discovery and synthesis. This latter aspect will require a paradigm change whereby Artificial Intelligence is no longer seen as just a tool for solving problems.

Few people these days say that: "computers will never be intelligent because they can't reason". This is a sign of the maturity of automated theorem provers (ATPs), which reason deductively, and machine learning (ML) programs which reason inductively. However, generally speaking, ML can discover, but not explain, while ATP can explain but not discover. The combination of ATP and ML has an exciting potential: much more creative approaches to the problems solved by these types of programs. Theorems would be proved by the invention of new concepts, and the finding of intermediate hypotheses empirically. Concept learning would be driven as much by theoretical exploration as empirical pattern discovery, and advantageous properties of the concepts learned would be explained rather than demonstrated. There are many technical problems to overcome in order to combine ATP and ML programs. Part of the problem is inflexibility in both types of programs. In particular, machine learning programs would have to discover axioms of the domain which are not necessarily related to the task they are being employed for. Automated reasoning programs would have to work with less well formed axioms and possibly produce unsound explanations.

The combination of inductive and deductive reasoning -- which is beginning to feature in AI research -- is just the tip of the iceberg, however. There are many forms of reasoning identified in the literature on creativity which will need to be firstly understood, and secondly incorporated into the theory. Other forms of reasoning include: abduction, invention, reflection, analogy, metaphor, serendipity, abstraction, experimentation, observation, evaluation, reparation, justification, exploitation, imagination, innovation and interpretation. I suggest that the problem of producing a mathematical theory of creative reasoning requires amalgamation, rather than combination, of established (and less well established) reasoning techniques. The notion of ways of reasoning which don't fit into categories like 'inductive' and 'deductive', etc., but somehow encapsulate the scope of many different reasoning techniques, is beyond current thinking. If we add to this the requirement that the theory takes into account creative teamwork, then we can see that devising such a theory is indeed a grand challenge, as we will also have to take into account theories of cooperation and competition.

As a research program, one way of breaking down the challenge of producing a theory of creativity would be in terms of the tasks being covered. I would suggest three levels of creative activity: problem solving (find me a solution), discovery (find me something I wasn't expecting) and synthesis (build me something special). Problem solving -- such as theorem proving -- can often be accomplished using one or two forms of reasoning; discovery -- such as the generation of scientific hypotheses -- involves a combination of different reasoning techniques, and requires the solution of many problems along the way; synthesis -- such as the invention of the vacuum cleaner -- involves prolonged processes of creative reasoning, and requires many discoveries along the way.

Artificial Intelligence is so often portrayed as just a set of techniques for solving problems. For instance, Luger and Stubblefield's well known AI textbook is subtitled: 'Structures and Strategies for Complex Problem Solving'. I believe it is naive to think that all intelligent activity can be characterised in terms of stating and solving a problem. In many cases, an exploration to find some artifacts is undertaken without a clear goal in mind, but in the knowledge that there are plenty of problems available that may end up being solved at some stage during the exploration. In other cases, the requirements for the solution to a problem evolve during the problem solving attempt. At the extreme -- in the case of serendipitous reasoning -- problems are invented to fit solutions. For example, the invention of Post-It notes is a classic case of prolonged serendipity: an attempt to solve one problem (finding a very sticky glue) failed, but eventually a problem which the artifact resulting from the original study (a re-usable glue) solved was invented. One could say that serendipity is just solving the problem of finding a problem given a solution, but I believe that this kind of shoe-horning is holding AI research back and we should aim to relax the grip that goal-directed methods have on AI currently. Exploration (curiosity) for its own sake surely has a role to play in intelligence. Necessity may be the mother of invention, but what about the other relatives?

This challenge arises from scientific curiosity about the limits of formalising human reasoning and the nature of creativity. Both of these enquiries were formulated long ago, and both still stand. A paradigm shift may occur if we stop thinking about AI as a set of reliable but uninspiring problem solving techniques and start to write programs which explore rather than solve, make mistakes along the way and evolve past their programming. This aim will not be met purely by commercially motivated evolutionary advance, and I think that the amalgamation of reasoning techniques (rather than combination) is beyond what is initially possible.

A useful decomposition of this challenge is into theories able to complete different levels of creative tasks: problem solving, discovery and synthesis, as described above. Of course, problem solving is certainly underway and has brought enormous benefit to science, and even if we fail to achieve the second and third aspects, it seems likely that their study will bring additional scientific benefit. It may not be entirely obvious to assess how far the challenge has been met, but we could set down certain markers. For example, if the same creativity theory could plausibly explain how musicians compose a sonata and how chemists discover a drug, then we could claim that the challenge has been met. Less ambitious goals in terms of problem solving and discovery tasks could be specified along the way.

It is impossible to imagine this grand challenge being met by anything other than an international consortium of researchers into Artificial Intelligence, Cognitive Science, Philosophy, Psychology and other fields. Research into machine creativity is already an international effort, but it is certainly held back by a lack of dedicated scientists of the right calibre and speciality. This is possibly because of the fact that creativity research does not (yet) have enthusiastic support from the entire research community, for the reasons given above, and others. By stating this grand challenge, I hope to have shown that it is time to reclaim the word creativity and admit that a thorough understanding of creative reasoning is an undeniable requirement for the next generation(s) of AI researchers.