

Computational Creativity: Coming of Age

*Simon Colton, Ramon López de Mántaras,
and Oliviero Stock*

■ This editorial provides an introduction to current AI research on computationally created artifacts as well as to the papers and topics covered by this special issue on computational creativity.

At its heart, computational creativity is the study of building software that exhibits behavior that would be deemed creative in humans. Such creative software can be used for autonomous creative tasks, such as inventing mathematical theories, writing poems, painting pictures, and composing music. However, computational creativity studies also enable us to understand human creativity and to produce programs for creative people to use, where the software acts as a creative collaborator rather than a mere tool. Historically, it's been difficult for society to come to terms with machines that purport to be intelligent and even more difficult to admit that they might be creative. For instance, in 1934, some professors at the University of Manchester in the United Kingdom built meccano models that were able to solve some mathematical equations. Groundbreaking for its time, this project was written up in a piece in *Meccano Magazine*. The article was titled "Are Thinking Machines Possible" and was very upbeat, but surprisingly ends by stating that "Truly creative thinking of course will always remain beyond the power of any machine."

Surely, though, this attitude has changed in light of the amazing advances in hardware and software technology that followed those meccano models? Sadly, no. Not much. Even within computer science, people are still skeptical about the creative potential of software. For instance, in *Non-Photorealistic Rendering*, a graphics textbook published in 2000, authors Thomas Stothotte and Stefan Schlechtweg boldly state that "simulating artistic techniques means also simulating human thinking and reasoning, especially creative thinking. This is impossible to do using algorithms or information processing systems." We could-

n't disagree more. As is hopefully evident from the articles in this volume, creativity is not some mystical gift that is beyond scientific study but rather something that can be investigated, simulated, and harnessed for the good of society. And while society might still be catching up, computational creativity as a discipline has come of age. This maturity is evident in the amount of activity related to computational creativity in recent years; in the sophistication of the creative software we are building; in the cultural value of the artifacts being produced by our software; and most importantly, in the consensus we are finding on general issues of computer creativity.

In the early days of AI research, many of the projects had grand goals: to build software able to compose sonatas, invent mathematical theories, and so on. In those days, the artifact-generation paradigm reigned supreme: the emphasis was on software that simulated entire reasoning and invention processes in order to build artifacts of real value. However, as further studies highlighted how difficult it is to implement artificially intelligent programs, AI researchers adopted the problem-solving paradigm. Here, an intelligent task is immediately broken down into manageable subparts and characterized as either a machine-learning problem, a planning problem, a theorem-proving problem, and so on. This kind of shoehorning enabled AI research to advance and—importantly—to demonstrate advances in concrete terms. Research in these subareas flourished, but it fragmented AI research: how many people regularly go to both machine-learning and theorem-proving conferences, for instance?

While problem solving certainly requires creativity, the majority of computational creativity researchers are more interested in the artifact-generation paradigm and are actively engaged in putting the pieces back together again. In this way, by combining AI systems, we are able to “climb the metamountain.” By this, we mean that individual techniques from AI and other areas of computing such as graphics are sufficiently powerful enough for us to enable software to take on tasks at increasing metalevels. For instance, automated painting has largely been associated with graphics research, where researchers have produced great tools that simulate some physical aspects of the painting process but not the cognitive aspects. Computational creativity researchers take their techniques, and ones from machine vision and other areas of AI, and simulate the way in which artists use these tools, for instance, by implementing scene construction methods. After this, we can look at the motivation for scene construction to express certain concepts and ideas, and so on. In this fashion, more of the cognitive aspects of creative processes are simulated, and artifacts of higher cultural val-

ue are produced by increasingly autonomous creative systems.

In the article by Amilcar Cardoso, Geraint Wiggins, and Tony Veale (“Converging on the Divergent: The History [and Future] of the International Joint Workshops in Computational Creativity”), more of the issues about which the community is coming to consensus are explored. This is done within the context of a survey of 10 years of computational creativity workshops. The authors start by separating the word *creativity* from the concept of “creativity,” and put both of these into historical context. In addition to historical case studies, they survey, compare, and contrast the philosophical and practical contributions made by key players in the field over the years and describe some logistics, such as the rise of the “show and tell” sessions at the annual meetings. They end by highlighting the prospects for research in this area and arguing that the field of computational creativity is very close to having firm enough foundations for rigorous scientific investigations in the future. This will be tested in coming years, as the series of computational creativity workshops evolves into the first computational creativity conference, to be held in Lisbon in January 2010.¹

Margaret Boden has been a key player in the field of computational creativity for a long time, having authored a number of influential books and papers. Her scholarly article here (“Computer Models of Creativity”) examines some of the thorny philosophical issues surrounding human and computer creativity. Unlike some authors working on creativity issues, she starts from the premise that creativity isn't a magical process that should be held in awe. From this starting point, Boden goes on to describe various ways in which computers can—and do—exhibit the same kinds of behaviors that creative humans do. With special emphasis on humor and visual arts as application domains, and looking at evolutionary processes as particularly valuable for computational creativity, Boden characterizes creative behaviors in terms of computer exploration, whether combinatorial and transformational. The article ends by posing a question about whether computers can be truly creative and by pointing out that through AI research, we now have an understanding of what sort of phenomenon creativity is.

The article by Gerhard Widmer, Sebastian Flossmann, and Maarten Grachten (“YQX Plays Chopin”) describes a computer program that learns to expressively perform classical piano music. The approach is data intensive and based on statistical learning. Performing music expressively certainly requires high levels of creativity, but the authors take a very pragmatic view to the question of whether their program can be said to

be creative or not and claim that “creativity is in the eye of the beholder.” In fact, the main goal of the authors is to investigate and better understand music performance as a creative human behavior by means of AI methods. AI has played a crucial role in the history of computer music almost since its beginning in the 1950s. However, until quite recently, most effort had been on compositional and improvisational systems and little efforts had been devoted to expressive performance.

The article by Pablo Gervás (“Computational Approaches to Storytelling and Creativity”) first addresses the issue of creativity from a general perspective, not only emphasizing the main features of interest in any creative process, such as novelty, unexpectedness, and usefulness (or goal satisfaction) of the output of the creative process, but also considering who is the intended audience and whether the output is expected to be novel with respect to previous outputs by the same system (or only with respect to an inspiring set). He then focuses on computational creativity and discusses a number of issues that have been identified as relevant to the computational modeling of creativity (whether it can be reduced to search or whether it should involve transforming a space, and to what extent it is influenced by the social context in which it takes place). Next he discusses the main features of interest in storytelling in terms of basic concepts from literary theory and then reviews representative storytelling systems and discusses them both in terms of how they deal with the features of interest in storytelling discussed in the first part of the paper and in terms of whether they introduce new additional features that need to be contemplated.

Chess has been a traditional challenge for AI. It took longer than expected, but now there is little controversy about the fact that the best programs outperform the best human players. At some recent matches between a series of outstanding champions and Bushinsky’s and Ban’s Deep Junior (of course won by the latter, still computer chess world champion), the most appreciated feature of the program that emerged in the after-match comments of the losers and of the independent specialized commentators was its creativity. The paper by Shay Bushinsky (“Deus Ex Machina—A Higher Creative Species in the Game of Chess”) talks specifically about that, with a number of very interesting examples and discussions about the prospects for the field. Even if it achieved remarkable results, Bushinsky points out that there are still several challenging developments ahead.

The article by Graeme Ritchie (“Can Computers Create Humor?”) deals with computational humor, that is, the field devoted to the automated generation of humor and to the appreciation of

humor. To some, the field may sound too ambitious: many people view humor as one of the most intriguing and distinguished forms of intelligence. In fact, if an artificial system were to produce jokes and forms of humor of all sorts, it would be based on the realization of most intelligent capabilities. The computational humor scope is currently much narrower, but some interesting forms of humor, like puns, funny ambiguous references, or for that matter ironical expressions, can be produced automatically. Ritchie provides an overview of some realizations, which in fact display a certain level of creativity, and discusses various aspects involved.

Our final article is a surprise, a kind of special guest for this issue of *AI Magazine*. It is a pleasure to include in this collection of written contributions something different, food for the brain in an artistic literary form, produced by EWI: an essay in the style of Douglas Hofstadter. Well, we cannot say more, you have to read it.

Computational creativity is a very lively subject area, with many issues still open to debate. For instance, many people still turn to the Turing test (and variants of it) to approximate the value of the artifacts produced by their software. That is, if a certain number of people cannot determine which artifacts were produced by computer and which were produced by a human, then the software is doing well. Other people believe that the Turing test is inappropriate for creative software. One has to ask the question, “Under full disclosure, would people value the artifacts produced by a computer as highly as they would the human produced ones?” In some domains, the answer could be yes: for instance, a joke is still funny whether or not it is produced by a computer. In other domains, such as the visual arts, however, the answer is very likely to be no. This highlights the fact that the production process is taken into account when assessing artworks. Hence, one could argue that such Turing-style tests are essentially setting the computers up for a fall. “Aha,” says the Turing tester, “these were produced by a human artist.... And these were produced by a convicted murderer!” We hope this makes the point without being too crass. Building creative software provides both a technical challenge and a societal one. To proceed further, we need to embrace the fact that computers aren’t human. We should be loud and proud about the artifacts being produced by our software, not ashamed that they weren’t produced by humans. We should celebrate the sophistication of the AI techniques we’ve employed to endow the software with creative behavior. And we should help the general public to appreciate the value of these computer creations by describing the methods employed by the software to create them.

New technologies, and in particular artificial intelligence, are drastically changing the nature of

creative processes. Computers are playing very significant roles in creative activities such as music, architecture, fine arts, and science. Indeed, the computer is already a canvas, a brush, a musical instrument, and so on. However, we believe that we must aim at more ambitious relations between computers and creativity. Rather than just seeing the computer as a tool to help human creators, we could see it as a creative entity in its own right. Creativity seems mysterious because when we have creative ideas it is very difficult to explain how we got them and we often talk about vague notions like “inspiration” and “intuition” when we try to explain creativity. The fact that we are not conscious of how a creative idea manifests itself does not necessarily imply that a scientific explanation cannot exist. As a matter of fact, we are not aware of how we perform other activities such as language understanding, pattern recognition, and so on, but we have better and better AI techniques able to replicate such activities. We agree with those that believe that creativity is an advanced form of our reasoning that involves memory, analogy, learning, and reasoning under constraints, among others, and is therefore possible to replicate by means of computers. This special issue addresses the question of the possibility of achieving artificial creativity through some examples of computer programs capable of replicating some aspects of creative behavior in several areas. We did not intend to cover the full range of AI approaches to computational creativity and we could not include all areas of application. The issue contains details of projects in the visual arts, music, storytelling, chess, and humor, in addition to providing an overview of the recent history of computational creativity and plenty of fuel for thought from a philosophical perspective.

To end, a final thought that helps to distinguish computational creativity research from other areas of artificial intelligence, and that might provide some ammunition against people who worry about our technological future. Good art makes you think. Whether it is a quick belly laugh at the punch line from a comedian or an in-depth reflection on the futility of war when viewing a painting like Picasso's *Guernica*, people who create artifacts do so in the hope that the artifact will make people engage their brains at an emotional or a cognitive level. This could be taken as our stated aim in computational creativity research: to build machines that can challenge us through artistic and scientific means. Of course, this will require the simulation of intelligence, but simulating thought processes is only one of our aims: we want to encourage thought processes in others, which is where we differ from other areas of AI. And when someone worries about the human race losing the ability to think for themselves because of smart

machines, please mention the automated painters, writers, composers, and scientists we're building specifically to encourage them to think more.

Note

1. See creative-systems.dei.uc.pt/icccx.

Simon Colton (www.doc.ic.ac.uk/~sgc) is a senior lecturer in computing at Imperial College, London. He leads the Computational Creativity Group, which investigates the application of AI and human-computer interaction techniques to creative applications in mathematics, graphic design, the visual arts, and video game design. In particular, Colton has developed the HR system, which has produced mathematical discoveries worthy of the mathematical literature; is currently working with video games companies on developing adaptive games; and has developed the Painting Fool, an automated artist that he hopes will be taken seriously as a creative painter in its own right—one day (www.thepaintingfool.com). He served on the AISB committee from 2001 to 2007, serves on the automated reasoning workshop committee and the Computational Creativity Conference steering group, and is one of the leaders of the AI and Games Academic/Industry research network. He is the author of more than 100 research papers and his research has won awards both nationally (BCS/CPHC distinguished dissertation award 2001) and internationally (AAAI best paper award 2000).

Ramon López de Mántaras (www.iiia.csic.es/~mantaras) is a research professor of the CSIC (Spanish National Research Council) and director of the Artificial Intelligence Research Institute of the CSIC. He holds an M.Sc. in computer science from the University of California Berkeley, a Ph.D. in physics from the University of Toulouse III, and a Ph.D. in computer science from the Technical University of Barcelona. His research focuses on case-based reasoning, autonomous robots, and AI and music. Former editor-in-chief of *Artificial Intelligence Communications*, he is currently associate editor of the *Artificial Intelligence Journal* and editorial board member of other international journals including the *Autonomous Robots Journal* and *AI Magazine*. He is an ECCAI Fellow and recipient of several awards including the City of Barcelona Research Prize. He was president of the Board of Trustees of IJCAI from 2007 to 2009.

Oliviero Stock (cit.fbk.eu/stock) is an Istitutoper la Ricerca Scientifica e Tecnologica (IRST) senior fellow. He has been at IRST since 1988 and was its director from 1997 to 2001. His activity is mainly in AI, natural language processing (including verbal humor), intelligent interfaces, and cognitive technologies. He is the author of 200 peer-reviewed papers and author or editor of 12 volumes. He has also been an invited speaker at over 60 conferences and participant or coordinator at over 20 panels. He has been chairman of the European Coordinating Committee for AI, president of the Association for Computational Linguistics and of the Italian AI Association, and is an ECCAI Fellow and an AAI Fellow.