

Quo vadis, Artificial Intelligence?

Daniel Berrar, *University of Ulster, United Kingdom and Cancer Institute, Japan*

Naoyuki Sato, *Future University Hakodate, Japan*

Alfons Schuster, *University of Ulster, United Kingdom and RIKEN Brain Science Institute, Japan*

© 2010 Daniel Berrar et al. Article first published in "Advances in Artificial Intelligence", V. 1 (2010) as open access article, distributed under the terms of the Creative Commons Attribution License <http://www.hindawi.com/journals/aai/2010/629869.html>

ABSTRACT. Since its conception in the mid 1950s, artificial intelligence with its great ambition to understand and emulate intelligence in natural and artificial environments alike is now a truly multidisciplinary field that reaches out and is inspired by a great diversity of other fields. Rapid advances in research and technology in various fields have created environments into which artificial intelligence could embed itself naturally and comfortably. Neuroscience with its desire to understand nervous systems of biological organisms and systems biology with its longing to comprehend, holistically, the multitude of complex interactions in biological systems are two such fields. They target ideals artificial intelligence has dreamt about for a long time including the computer simulation of an entire biological brain or the creation of new life forms from manipulations of cellular and genetic information in the laboratory. The scope for artificial intelligence in neuroscience and systems biology is extremely wide. This article investigates the standing of artificial intelligence in relation to neuroscience and systems biology and provides an outlook at new and exciting challenges for artificial intelligence in these fields. These challenges include, but are not necessarily limited to, the ability to learn from other projects and to be inventive, to understand the potential and exploit novel computing paradigms and environments, to specify and adhere to stringent standards and robust statistical frameworks, to be integrative, and to embrace openness principles.

KEYWORDS: *Artificial intelligence, Computer simulation, Multidisciplinary field, Neuroscience, Systems biology*

Since its foundation, which is often associated with a conference held at Dartmouth College in New Hampshire in 1956, artificial intelligence (AI) can look back on a rather exciting and successful

(though not unproblematic) career. Per definition, artificial intelligence - which may be described as a field of study embracing concepts, methodologies, and techniques as part of a computer program that exhibits characteristics akin to intelligent behavior (Dubitzky, Azuaje, 2004) - is a sensitive and vulnerable field as the term *intelligence* itself lacks a generally (ideally universally) acknowledged definition. There is, of course, no shortage of characterizations for the term intelligence (e.g., Pfeifer, Scheier, 2000, pp. 6-21 or Negnevitsky, 2004, pp. 1-4), at large; however, the term remains a complicated, subtle, and malleable topic. In addition, although, arguably, many animals are intelligent to some degree, by intelligence people often assume the same scope of intelligence as can be seen in general purpose human action (e.g., Brighton and Selina, 2003 mentions the creation of artifacts that are capable of mimicking and expressing human intelligence, thought, consciousness, and emotion). Although the field has been exposed to some (often well-deserved) criticisms over the years, it may be fair to say that AI is a truly great scientific endeavor that reaches out and embraces a great variety of scientific disciplines ranging from areas that are mathematically and conceptually well defined (e.g., computer science, machine learning, biology, and neuroscience) to areas that are more difficult to quantify and deal with (e.g., philosophy of mind, cognitive science, and emotional intelligence). Of course, it is not possible here to acknowledge adequately the many contributions (including those coming from rapid advances in computer hardware and software) that shaped AI into the exciting discipline it is today, but it may be possible to categorize, crudely, the development of AI into a few key stages (e.g., see Negnevitsky, 2004, pp. 1-24).

Early AI projects had to come to terms with the many challenges related to the production of knowledge-based (expert) systems, including the acquisition of knowledge, its representation and evolution, among other challenges. The General Problem Solver, the early DENDRAL and MYCIN expert systems, or the ELIZA program may stand representatively for this period. The General Problem Solver is a general purpose program to simulate human decision-making, thinking, and reasoning; DENDRAL (a system for analyzing chemicals) and MYCIN (a system for the diagnosis of infectious blood diseases) are knowledge-based expert systems and focus on the problem solving in specific domains rather than on

a general problem-solving strategy that is applicable to many (all) domains; ELIZA is a program related to natural language processing (and the Turing test) where humans interacting with a machine in a question-answer script-based mode were led to believe that they were actually participating in a human-human interaction.

Artificial intelligence received a major boost in the mid 1980s from works on artificial neural networks (ANNs). It is possible to say that since the early days of the Perceptron the story of artificial neural networks is one of the most successful chapters in the voluminous book of AI. Since the emergence of several techniques, many of them united under the soft computing paradigm then provided AI with a capable and flexible repository for both theoretical research as well as hands-on problem-solving applications. Possibly, from an application point of view, this was one of the most exciting times for AI and the term *knowledge engineering*, which suggests that systems showing some degree of intelligent problem-solving ability could be assembled from a toolbox of available techniques, may capture this excitement quite well.

In more recent times, AI has made it into the limelight through the DARPA (Defense Advanced Research Projects Agency) Grand Challenges¹. In these challenges, vehicles had to navigate autonomously in increasingly challenging environments in an intelligent manner, avoiding obstacles and solving problems of increasing levels of complexity, without human intervention (e.g., lagnemma, Buehler, 2006). Without hesitation one needs to acknowledge the achievements in these tasks, paralleled perhaps in their prestige and popularity by the outstanding successes of IBM's Deep Blue supercomputer (the machine that recorded the first win in a game of chess against a reigning World Chess Champion, Garry Kasparov, in 1996) and the Deep Fritz chess engine (the commercial chess engine that triumphed 4-2 in December 2006 against the reigning champion Vladimir Kramnik in a six-game match)². It is beyond the scope of this paper to pay credit to the great variety of important and popular areas in which AI has made substantial contributions (often without being duly credited, taken for granted, or simply neglected) but it is helpful to emphasize that these areas include mainstream utilities such as the World Wide Web (e.g., the so-called *intelligent web* draws its power from algorithms that process information intelligently - Google's page ranking algorithm or algorithms discovering matches on social-networking sites are

1. <http://www.darpa.mil/grandchallenge/index.asp>

2. <http://www.research.ibm.com/deepblue/>

two examples, see Marmanis, Babenko, 2009), toys and gadgets (e.g., computer games, Millington, 2006, or the Lego Mindstorms Robotics Invention System, Kelly, Schuster, 2005), extremely helpful human-computer interfaces (e.g., user-friendly interfaces using natural language recognition and visualizations including brain-computer interfaces, in Duch, Mandziuk, 2004; Lebedev, Nicolelis, 2006), or humanoid robots functioning in various roles as partners for people in the immediate environment of human beings (e.g., AIBO, ASIMO)³. Artificial intelligence has even reached out beyond earthly confines and is heavily utilized for numerous tasks by several space agencies including the European Space Agency (ESA) (e.g., Girimonte, Izzo, 2007) and NASA (National Aeronautics and Space Administration)⁴.

The previous points may suggest that contemporary AI should find itself in a comfortable situation from which it should firmly and enthusiastically continue its cause - the study and creation of artificial entities that are capable of expressing human intelligence, thought, consciousness, and emotion. The paper argues that this may not necessarily be so and that AI may need to be cautious and perhaps alert to some degree about its standing because just around the corner there are extremely exciting and powerful fields that touch upon areas - and begin to outshine AI in areas - that are at the very heart of AI itself. These fields include, but are not necessarily limited to, neuroscience (Purves et al., 2004), synthetic biology, and systems biology (Kitano, 2002). These disciplines comprise several subfields with boundaries that are often blurred. They also exemplify many modern research endeavors that are characterized by their complexity and a truly interdisciplinary nature that draws upon the expertise of scientists from a wide range of academic backgrounds.

The forthcoming sections investigate how AI is situated in this extended environment. Initially, the following section takes a closer look at the interplay between AI, neuroscience, synthetic biology, and systems biology. Section *Hurdles and the road ahead* identifies several challenging hurdles in this interplay and includes suggestions for how these hurdles may be overcome in order to create a synergetic and productive environment that is beneficial and supportive to practitioners working in these fields. The final section provides concluding remarks and ends the paper with a summary.

3. <http://world.honda.com/ASIMO/>

4. <http://www-aig.jpl.nasa.gov/>

Quo vadis, Artificial Intelligence?

A starting point for an answer to this question may be one of the most successful scientific endeavors in the last century - the Human Genome Project. The Human Genome Project achieved its goal, the definition of the sequence of chemical base pairs which make up DNA, on 26 June 2000 with the publication of a first draft of the human genome - acknowledged in special issues in "Nature" (Lander et al., 2001) and "Science" (Venter et al., 2001). Simply speaking, the human genome consists of DNA (deoxyribonucleic acid). This DNA is assembled from a small number of basic nucleotides (adenine, guanine, cytosine, and thymine). These elementary building blocks bond together into DNA sequences. Genes, which are another important component in this context, are particular DNA sequences that play a fundamental role in the evolution and production of organisms (e.g., the development of a human being). At root, DNA is a code, and like any code DNA contains instructions and information (e.g., for the building of complex, 3-dimensional proteins). The genome contains the full set of instructions. Understanding (in its entirety) this code (the genome) is a major goal today. Crucially, the information in the human genome unfolds itself fully over time at different biological levels of structural organization and complexity, each level with its own important functions. From the bottom upwards, these levels may be molecules, genes, more complex cell formations, organs, and, on the highest level, complex higher organisms including human beings. Of course, the abundance of inspiring challenges in this environment invites a growing number of scientists from many fields, including biologists, AI practitioners, and other individuals whose fields have a long tradition in studying complex systems, data, codes, and information processing to pick and investigate hugely exciting problems in this overwhelmingly rich application area of study.

AI in synthetic biology and systems biology

A feature that unites systems biology and synthetic biology is the tremendous complexity that is inherent in both fields. Asked for a dividing line, one may argue that systems biology is motivated to create a complete and detailed understanding of existing biological systems whereas synthetic biology is driven by the vision to create novel biological entities from first principles.

Synthetic biology

Synthetic biology is defined as the design and construction of new biological parts, devices, and systems, and the redesign of existing, natural biological systems for useful purposes⁵. Hence, synthetic biology envisages the design and fabrication of biological components and systems that do not exist in the natural world from non-living, raw material and programming them with desired (novel) chemical functionality. The field also envisages the redesign of existing biological systems (Rasmussen et al., 2004). In addition, it is helpful to understand that the terms (artificial) synthetic life and artificial life (not to be confused with Alife, the field of study that examines systems related to life mainly through computer simulation applications) are related to synthetic biology and apply when the goal is to recreate life from non-living (abiotic) materials (e.g., Holmes, 2005). It is important to understand that artificial synthetic life and synthetic biology have the support of a strong lobby and that there are strong beliefs that artificial cells will eventually be created. Anyway, in the context of this paper it is necessary to underline that there are various concepts in these fields that should be of great interest for the AI community. For instance, the *minimal* genome (the smallest set of genes needed to support a simple living cell) is an important aim of artificial synthetic life (Ainsworth, 2003). From an AI perspective, this concept of a minimal genome is rather interesting and may map to the term *minimal intelligence*, which may be defined as the smallest amount of information needed to support intelligent behavior. The production of a minimal cell could intrinsically involve some intelligent processing that may be realized on the DNA level. It is possible to envisage a simulation of this processing involving some form of AI in an environment that exploits features of the popular DNA computing paradigmas models for various machine learning techniques (e.g., rough set analysis, Schuster, 2004) and other modeling approaches (e.g., Petri nets, Schuster, 2008) exist in the DNA computing world for quite some time. In this case, questions may arise such as: how much artificial intelligence or computing is needed for the construction or support of an artificial cell? Concepts such as *minimum cell computation* or *minimum cell information* may also emerge in this context. Overall, it is relatively easy to see that all these terms are highly relevant to AI and that synthetic biology certainly and quite naturally addresses several fundamental issues AI has had on its agenda since its very beginning.

5. <http://syntheticbiology.org/>

6. <http://www.systems-biology.org/>

Systems biology

The field of systems biology dedicates itself to the study of complex biological systems, their properties, interactions and dynamics (Kitano, 2002), and draws on a similarly wide pool of fields including several of the fields mentioned before. Of paramount importance is the holistic paradigm⁶, which is in stark contrast to the reductionist approach that was prevailing in molecular biology during the last decades. It should be noted here, however, that systems-level approaches to natural phenomena are not an invention of the 21st century. For example, Alan Turing, widely regarded as the grandfather of AI, adopted systems-level approaches in his studies on the chemical basis of morphogenesis (Turing, 1952), and could therefore rightfully be regarded as a pioneer in systems biology.

The characteristics of complex biological systems often emerge naturally via the interplay of individual system components. A typical example of such an emergent phenomenon is *robustness* (Kitano, 2007). Robust systems are inherently able to maintain their function despite internal or external perturbations, changes in the environment they operate (or live) in, or unreliable components. However, these individual components themselves may not be considered robust. Hence, robustness is an emergent property of the system as a whole and cannot be described or understood by studying the components alone. This preservation of function does not prevent complex systems from evolvability - in contrast, evolutionary principles appear to favor robust systems (Kitano, 2007). Although the principles behind robustness are not fully understood, there are specific cornerstones upon which robustness is believed to rely in various domains: (i) *modularity* the system is composed of elements that work together synergistically; (ii) *redundancy* some of the components share identical function, hence the depletion of one component can be compensated by others; (iii) *feedback control* the system is equipped with a sensor for the detection of changes in the environment and a controller for reacting to these changes, which allows dynamic system behavior. This feedback control also includes mechanisms for systems repair such as purging, the deliberate destruction of components that fail to operate properly (e.g., outside of their defined range of action); (iv) *spatial compartmentalization* the complex system has a physical or virtual embodiment that is subdivided into areas or compartments that may exchange information with each other;

(v) *distributed processing* the modular elements in the compartments collectively give rise to a higher, system-level function, phenotype, or morphology. It is hypothesized that for example cancer is a complex robust system (Kitano, 2004), relying on a functional redundancy of genetically heterogeneous cells, and on the ability to maintain homeostasis and functionality despite changes of the tumor microenvironment. Although these changes may affect an individual tumor cell to a larger extent than a healthy cell, the collective of tumor cells is robust due to functional redundancy caused by the genetic variability of its components. Hence, insights about the mechanisms leading to robustness could help us identify potential Achilles' heels in systems whose function we seek to destroy. Unraveling key mechanisms leading to robustness represents one of the grand challenges of systems biology. While methods from cybernetics and control theory are arguably at the forefront of modeling systems dynamics, there is a clear scope for AI approaches. For example, the multi-faceted appeal of robustness for AI has been addressed in a recent edited book (Schuster, 2008) that investigates robustness in a diverse range of areas related to AI including computer hardware and software, computer networks and protocols, brain-computer interfaces, biological networks and immune systems, humanoid robotics, image processing, artificial neural networks, genetic algorithms, chaos theory, and other soft computing techniques, as well as space system design and bioregenerative life support systems. This could be complemented by approaches from artificial life, which aims at understanding properties of life by abstracting its fundamental, evolutionary principles, and recreating and emulating them using computer programs. In fact, experiments with digital organisms have already revealed astonishing insights into the genetic basis of evolution (Adami, 2006). In any case, the aforementioned book agrees that robustness is not a trivial topic, that in truth robustness is a rather elusive and challenging concept, but the book also clearly emphasizes that features of robustness, as they appear in nature, have a great potential for being utilized as general design principles for AI systems (Schuster, 2008).

Further, the philosophy of systems biology is rooted in the desire to develop mathematically-founded, testable theories to *understand* whole biological organisms. Such theories, however, are only in their infancy. Mathematical models are still simplistic compared to

the daunting complexity of real biology and therefore often met with a justified skepticism (Voit, Schwacke, 2007). Yet, even the very notion of understanding in this context may be debatable - what do we mean by *understanding* a living system, actually? Does understanding imply that we grasp the *nuts and bolts* of the intricate living machinery? Or is it sufficient to make reliable predictions about the behavior of the system under normal conditions and under specific perturbations? Would it be possible to devise tests for systems biology models that are able to determine whether we have *understood* the system at hand? In AI, the Turing test is the ultimate (though not unchallenged) benchmark test for machine intelligence (Turing, 1950). What could represent an equivalent test for a computational model in systems biology that claims to *understand* the living system? Harel (2005) proposed several modifications to the Turing test in order to validate an artificial model of a living system, but the question of *understanding* is more fundamental in essence. It seems sensible to consider these delicate questions in order to have a yardstick against which we can measure our models (in systems biology, neuroscience and other disciplines); this would also preclude a moving of the goalpost after the models have been developed. Scientists from (philosophical) AI could provide valuable input to this debate.

AI and neuroscience

As mentioned before in section *Quo vadis, Artificial Intelligence?*, the information in the genome unfolds itself fully over time at different biological levels of organization and complexity (e.g., from molecules to genes, more complex cell formations and organs, up to complex higher organisms such as a human being). Neuroscience with its diverse range of subfields complements this chain in a natural way by investigating (among many other things) the cognitive functions inherent in such organisms. The phrase *from molecules to cognition* therefore is sometimes used to summarize the field.

Neuroscience is a tremendously complex field with many subfields (e.g., Purves et al., 2004) and offers a wide range of opportunities for AI. Neural signalling, for instance, investigates how neural systems acquire, coordinate, and disseminate information. Knowledge about these processes is fundamental to understanding brain pathologies, but also for the development of novel approaches to diagnosing and treating such problems. Artificial intelligence systems can benefit

from such an understanding in the field of artificial neural networks in particular and other application areas where networks play a fundamental role in general (e.g., pervasive/ubiquitous computing - Saha, Mukherjee, 2003 - and autonomic computing - Kephart, Chess, 2003 - aim for a new type of networked communication systems that are autonomously controlled, self-organized, radically distributed, technology independent, scale-free, and can manage themselves to various degrees given high level objectives from administrators).

Sensations and sensory signal processing, including data storage in memory, are other areas that are key to neuroscience (e.g., Sato, Yamaguchi, 2008). The signal processing in biological vision and olfactory systems, for instance, is heavily investigated and there is a direct link to pattern recognition, knowledge representation, and related AI areas. The perception of information from an environment (internal and external) and the accurate and meaningful digestion and interpretation of such information in real-time is a vitally important task for biological nervous systems in general, and, holistically, the field of robotics is an AI working area that encapsulates many of these issues.

The changing brain and complex brain functions are fundamentally important in neuroscience. The development of a brain throughout the lifespan of a biological organism, where the brain changes both its size (natural development) and its content (real-life experiences) is also highly relevant to AI. For example, it is relatively easy to draw parallels to research on autonomous intelligent agents roaming in an (intelligent) World Wide Web that undergoes changes in size and information content in perpetual motion (Marmanis, Babenko, 2009).

It is important to understand that the relationship between AI and neuroscience works vice versa. For example, the insight that many networks in nature are scale-free (Barábasi, Bonabeau, 2003) may be informative for neuroscience, it is, however, a finding that comes from a study researching the large-scale topology of (a portion) of the World Wide Web that had very little to do with neuroscience. It is also important to understand that some projects in neuroscience are similar in character and scope to the Human Genome Project, they are large-scale and have an extremely ambitious, clear-cut goal. The prime example is the well-known Blue Brain Project at École Polytechnique Fédérale de Lausanne⁷ in its attempt to

7. <http://bluebrain.epfl.ch/>

reverse-engineer the mammalian brain in order to understand brain function and dysfunction through detailed simulations, which is an incredible undertaking. However, this goal (the creation of artificial entities that may eventually be capable of intelligence, thought, consciousness, and emotion) is very much what AI practitioners have been dreaming about for a long time - but, per se, the Blue Brain Project is a neuroscience project and not a project in AI. Note that a similar argument holds for the previous sections on synthetic biology and systems biology. These fields address several key AI issues, but the term AI, in relation to its core business, the creation of artificial intelligence, appears relatively rarely in these areas. There are exceptions of course, but very often when AI is mentioned in these areas it is usually in a data analysis, data mining, or applied machine learning way, (e.g., Dubitzky, Azuaje, 2004). This does not mean at all that the Blue Brain Project, and other brain projects, explicitly avoid the mentioning of AI and its mission (Potter, 2007). It is just that AI, at the moment, simply does not play the role it could play (and perhaps due to its tradition should play) in such projects. It is as if AI has somehow disappeared from the Blue Brain Project and neuroscience (and synthetic/systems biology) at large.

On the basis of the previous sections - which identified a possible rich interplay between AI, neuroscience, synthetic biology, and systems biology - this paper feels that this is a somewhat surprising situation. The forthcoming sections therefore examine this situation in more detail and make suggestions that may help AI to raise its profile in neuroscience, synthetic and systems biology.

Hurdles and the road ahead

This section takes its motivation from the findings just mentioned and suggests several challenges AI may face in order to become more mainstream in neuroscience, synthetic and systems biology. These suggestions include, but are not necessarily limited to the following: (i) to be able to learn from other projects and to be inventive, (ii) utilization of novel computing paradigms and environments, (iii) development of standards, (iv) stringent computational and statistical frameworks, (v) adoption of *openness* principles, and (vi) to be integrative and creative.

Learning from other projects and being inventive

Although these first items are more like soft factors than hard factors they can be rather important for AI because success is as much the result of hard work as it is the result from being open and flexible enough to learn from others. It is not without reason that the capacity to learn stands at the heart of many AI systems, and it is clear that learning is an extremely powerful problem-solving strategy not only for AI systems but humans (and other organisms) in general. So what can AI learn from projects such as the Human Genome Project or the Blue Brain Project? Initially, it is necessary perhaps to highlight some of the outstanding features in such projects simply from a project management perspective. In general, these projects are: large-scale, team-oriented, inherently cross-disciplinary, collaborative, distributed, heterogeneous, and global. Crucially, on top of all this, these projects are garnished by a truly *grand vision!* (The history of mankind is full of examples for this. In more recent times, the early Apollo missions with their wonderfully inspiring goal of bringing men to the Moon - and the currently orbiting International Space Station research facility alike - perhaps hit the nail on the head). This does not necessarily imply, of course, that projects lacking these features are doomed to fail, but several positive developments in the former framework suggest that AI may benefit from adopting such an approach for future projects. The DARPA Grand Challenges are hugely attractive and successful but carry an ethical dimension with them that should not be underestimated. Some people may be reluctant to participate in these challenges or oppose them in principle because they are organized by the United States Department of Defense, which is responsible for the development of new technology for use by the military; the ethical dimension plays a similarly important role in neuroscience and synthetic biology, too. The DARPA Grand Challenges mentioned earlier or the concepts behind the organization of various X-Prizes⁸ may be trendsetting in this regard in many ways. Essentially, these events generate an extremely rich pool of expertise by drawing together elements of public interest, entrepreneurial spirit, and cross-disciplinary innovation. The dream is that this pool eventually creates breakthroughs that are beneficial to mankind at large.

This does not mean that there is a general lack of projects in AI

8. <http://www.xprize.org/>

9. <http://www.loebner.net/Prize/loebner-prize.html>

10. <http://www.robocup.org/>

11. <http://www.santafe.edu/>

12. <http://www.edge.org>

embracing this philosophy in principle. In many cases, however, such projects operate on a somewhat smaller scale for several reasons (financial support, popularity, etc.). The Loebner Prize in Artificial Intelligence⁹ (a competition in the format of a standard Turing test where a machine has to demonstrate human-like ability/intelligence), or the RoboCup¹⁰ (a competition where robots are tested in their ability to play soccer) may stand representatively for such projects. It is important now to understand that the scale of a project is a crucial factor, but also that scale alone may not be enough. There is another major obstacle that can deny even extremely exciting projects the support, attention, and success they may deserve: *individualism*. Take the case of humanoid robotics as an example. In the RoboCup challenges just mentioned, teams develop their applications individually and compete against each other again individually. The same accounts for industry driven projects in humanoid robotics. AIBO, for instance, is a Sony project, and ASIMO a humanoid robot created by Honda. If humanoid robotics is important enough for so many enthusiasts and such industry powerhouses then why is there no Humanoid Robot Grand Challenge with all resources open to all? Since robotics is only one of many examples, perhaps, what is required is a new way of thinking and collaboration in science at large. Actually, there are several rather interesting initiatives that point in the right direction. For example, the Santa Fe Institute¹¹ devotes itself to creating a new kind of scientific research community, and the Edge Foundation, Inc.¹² may be seen as a super sophisticated sphere where some of the greatest minds of our time engage in an intelligible and enthralling discourse on a diverse range of fundamental issues. Since section *Adoption of openness principles* is going to comment further on this attitude of *putting all cards on the table*, this section wants to add that being able to learn from other projects and to be open for change may get you some way but to make a truly large step forward may require another feature: the ability to be *inventive*.

Being inventive

To be able to be creative and inventive is a great thing to be and nature reminds us kindly and patiently day by day how important it is to be clever and original, equally, in the small and in the large. It is difficult, if not impossible, to specify a process for inventiveness, but one way for AI may be to first of all identify novel and exciting fields

that are relatively unexplored by AI or to get involved more deeply into current working areas with great potential for the future. From a potentially larger pool of fields, computer games, virtual reality, nanotechnology, quantum computing, and space exploration may be such fields.

Artificial intelligence has been in touch with the fields of computer games (Millington, 2006) and virtual reality (Luck, Aylett, 2000) almost from the point when these fields took off. Many projects, including those of commercial organizations, provide game engines and other artifacts for free but what is missing, and what may be desirable according to this paper, is the transition from small-scale investigations conducted by individuals or individual institutions to a coordinated effort in the spirit emphasized in the previous sections. The virtual world *Second Life*¹³ developed by Linden Lab and launched in 2003 demonstrates rather well the great potential such an approach may have for computer games, virtual reality, and AI. Second Life exploits free software utilization and the internet and provides users with an open environment in which they are free to express their creativity by designing artifacts with which they are free to interact and participate individually or in groups in a continuously changing and evolving virtual world.

In comparison to computer games and virtual reality, nanotechnology and quantum computing are two areas that are rather unexplored by AI. Nanotechnology is a rather diverse field again, and among other things the field researches new approaches to developing and processing new materials with dimensions on the nanoscale. Most people possibly are familiar with ideas such as micro-robots cleaning blood vessels, but then, where there is a robot there should be room for AI and indeed there is (Huw Arnall, 2003). Quantum computing is another field with great potential for AI and there are attempts for combining the two fields (e.g., Sgarbas, 2007; Kak, 2007). In many ways the weird world of quantum computing with its sometimes unsettling conundrums (parallel universes, entanglement, etc.) has much of the charm of the early AI (but arguably a higher degree of robustness and credibility when it comes to interpretations that could easily come from the science fiction camp).

The exciting domain of space exploration with its large scale dimension has a similar appeal for AI, and in fact there are many projects in which AI already faces the great challenge of space.

13. <http://secondlife.com/>

NASA and ESA, for example, pursue large space research programs in which AI plays a significant role; be it in the design of life-support systems (Rodriguez et al., 2007) or the development of autonomous systems for satellite path planning (Izzo, Pettazzi, 2007). Actually, a very attractive thing to do for AI in this spirit could be a grand challenge with the goal of sending a humanoid robot to the Moon. Perhaps, one of the tasks (not to be taken too seriously here) for this robot could be to play a round of golf on the Moon, mimicking the attempts demonstrated by NASA astronaut Alan Shepard who was the first human being hitting golf balls (with a six iron) on the Moon on the Apollo 14 mission in 1971. Overall, these examples indicate that there is great potential for AI in several cutting-edge areas. Whether AI can exploit these challenges to its advantage (e.g., by clearly defining a stimulating set of - truly - grand challenges, in the areas mentioned here or any other area for that matter) is a question open for speculation that only future AI enthusiasts may be able to answer.

Exploitation of novel computing areas and computing paradigms

The mentioning of the terms nanotechnology or quantum computing in the previous section provides an introduction to an issue that is rather important for AI: the exploitation of novel computing areas and computing paradigms. In the past, a standard computer was largely the platform for most AI systems (e.g., a desktop PC or an application embedded in a mobile robot). Our time, however, has seen the coming of new computing paradigms such as quantum computing or DNA computing. This development creates fundamental changes in terms of computer hardware and software (e.g., the difference between a standard PC and a DNA computer is sometimes expressed by saying that for a standard PC computing is computing with *bits*, whereas for a DNA computer it is computing with *molecules*). Similar arguments hold for work in synthetic biology (Section *Synthetic biology*) where the main working entity (processor) is the cell (e.g., the relatively young field of membrane computing is a new computational model inspired by the processing found in biological cells, Păun, 2006).

There are also dramatic developments happening in neuroscience at the moment. Historically, computational neuroscience has been a branch of AI (Arbib, 2003), with a focus on modeling and theorizing

functional aspects behind brain signals. Artificial, yet biologically plausible, neural networks like in the Blue Brain Project have long been playing a central role in the interpretation of brain signals and mechanisms of brain functions. The evaluation of artificial neural networks in this field often requires realistic, real-world environments, which are facilitated by robotics and computer vision. Methods that originated from AI play now an equally important role in the analysis of brain signals from both animals and humans. Sophisticated brain-machine interfaces (Lebedev, Nicolelis, 2006), for instance, exploit the computational power of a human brain, aiming at the translation of brain signals into motor commands for robotic device control (Santhanam et al., 2006). Pilot studies in this field have shown remarkable results with great promises for severely disabled people (e.g., for neuroprostheses Hochberg et al., 2006).

Along this line, *brain-like* computers are another exciting area in neuroscience. For instance, IBM has engaged into an activity researching brain-like computers and the newly coined term *cognitive computing* encapsulates the key idea to engineer *mind-like* intelligent machines by reverse engineering the structure, dynamics, function and behavior of the brain (Freidman, 2009). Crucially, this project targets an understanding of the *mind* (decade of the mind, Albus et al., 2007) based on the findings from the last decade of brain research.

The development of new neuroscience technologies for learning and memory in vitro is another relatively new and highly promising area in neuroscience. Researchers in this field grow mammalian brain cells in culture on multielectrode arrays in order to form a long-term, two-way interface between cultured networks and a computer. The generated cultured nets then can serve as brains of simulated so-called animats or robotic creatures, which opens the door widely for all sorts of AI research related to neuroscience driven robotics (e.g., Potter, 2007; Bakkum et al., 2007). Other interesting contributions, related to these studies, come from work investigating predictive behavior within microbial genetic networks where bacteria anticipate changing environments (Saigusa et al., 2008). The fact that bacteria have no brains or nervous system makes this work particularly interesting, and there are suggestions that these microbes experience and learn through evolutionary changes in their complex networks of interacting

genes and proteins (i.e., the problem-solving potential is part of the configuration of the system, Tagkopoulos et al., 2008). A key element of this system is *quorum sensing*, the regulation of gene expression in response to fluctuations in cell-population density (Miller, Bassler, 2001). Bacteria can exploit this naturally occurring, self-organizing principle as a means of communication with each other in order to collectively solve problems, or to give rise to collective phenomena such as bioluminescence. Interestingly, there is a parallel between quorum sensing and *swarm behavior*, the emergent property of a collective of individuals acting in concert. Key concepts of the collective behavior are bottom-up (instead of top-down) approaches, and the lack of a central command-and-control structure (i.e., decentralization), which have also been described as *swarm intelligence* - the collective is able to solve problems that are beyond the information processing abilities of the individual. Swarm behavior has been extensively studied in artificial life research and led to numerous algorithms, from ant colony algorithms to find optimal paths to crowd simulations in movie animation technology. There are certainly other examples worth mentioning in this section but for the sake of brevity this text moves on to other issues that may be equally important for AI in the wider context of this paper.

Development of standards

Standards facilitate many aspects of product quality such as correctness, reliability, reproducibility, robustness, understandability, interoperability, and maintainability. In a global setting, and in a heavy-demanding computing-permeated environment in particular, standards are an unavoidable necessity. There are several activities directed at the implementation of such standards for systems biology (Brazma et al., 2006), and the *Systems Biology Markup Language* (SBML)¹⁴ is just one example (Hucka et al., 2003). SBML provides a machine-readable, XML based format for describing models of biological processes. Projects in systems biology often involve genomic data from high-throughput experiments such as DNA microarrays (Schena et al., 1995). The *Minimum Information About a Microarray Experiment* (MIAME) (Brazma et al., 2001) is a standard that was developed for the exhaustive and unambiguous description of such experiments. Over 50 journals, including the "Nature" series, now require that any published study involving

14. <http://sbml.org/>

microarray data makes this data publicly available in a format compliant to the MIAME standard.

The development of standards could also address the terminological aspect. In multidisciplinary research environments, the same technical term can refer to very different things (Dubitzky et al., 2007), adding to the already existing language barrier. For example, the term *sample* may bear a completely different meaning to a statistician and a biologist. While such a misunderstanding can be easily detected and resolved, problems arising from other terms may be more difficult to spot: test set versus validation set, supervised versus unsupervised, and so forth, (Troester et al., 2009).

Standards for describing computational models do exist, for example, the Predictive Model Markup Language (PMML), an XML-based schema that allows the detailed description of statistical models. In metabolomics research, efforts for establishing minimum standards for reporting data analysis have been undertaken (Goodacre et al., 2007). However, in general, standards for computational and statistical models do not seem to be widely embraced yet, at least life science publications do not widely adhere to these standards, arguably because they are not enforced by journals and other publication sources - yet. There are many critical voices, though, calling for standards implementing *good statistical guidelines* (Troester et al., 2009). It can be speculated that in the near future journals will begin enforcing standardized reports of computational models, similarly to the MIAME standard for microarray data.

Adherence to stringent computational and statistical frameworks

Despite the pivotal role that computational and statistical models play in today's interdisciplinary research, their deployment (and the analysis workflow) are often only imprecisely and rudimentarily described. Problems can arise from uncertainties such as: How were the data preprocessed and normalized, precisely? Were outliers removed? If a data re-sampling strategy such as cross-validation was adopted, then how was this performed in detail? Which loss function was used and why? How were the model's parameters calibrated? Published reports of computational and statistical models generally leave many open questions, even in the context of data mining competitions such as the KDD cup¹⁵. Of crucial importance in this context is an audit trail that logs the detailed analysis steps,

15. <http://www.sigkdd.org/kddcup/>

allowing a reproducible study. This *bookkeeping*, albeit apparently simple, has been described as the *number one* problem in today's bioinformatics research (Savage, 2008).

Practitioners developing AI-based solutions for interdisciplinary life science research could perhaps play an active role in the specification of the aforementioned standards (even though it may be a bit of a stumbling block for some *old workhorses* in AI). For example, artificial neural networks have shown remarkable performance for classification and prediction tasks, and they are also an integral part of the toolbox for analyzing genomic data (Khan et al., 2001). However, their intrinsic black-box character may represent a major impediment to their widespread use in the life sciences. Practitioners in AI must be aware that applications in these fields should, ideally, also have explanatory power, and must be embedded in statistical frameworks that account for experimental artifacts and biases, issues due to multiple hypotheses testing, and the notorious curse of dimensionality (i.e., many more features than samples, which presents a huge challenge for clustering, classification and rule extraction), to name but a few.

Adoption of openness principles

The definition and adherence of rigorous standards and frameworks is closely linked to the principle of *openness*. For software development, for instance, the primary goal of open source is the production of reliable software, its main technique is code sharing, and the main tool for achieving this is the internet. There is a mountain of evidence from many projects (Linux, Ruby on Rails, Python, etc.) about the quality demonstrated by such systems and there is a great belief in many places that the open source model may be superior in many ways to traditional (commercial, individual) development approaches that are happening in closed environments. Openness should not be confused with decentralization or disorganization. Quite the opposite is true and an organization such as the World Wide Web Consortium (W3C), the organization which oversees the standards for the World Wide Web, may hold as the best example for this mode of operation. The W3C is pioneering in its organization and mode of operation and provides a blueprint for how large-scale projects may benefit from the adoption of open and streamlined processes and procedures.

From an AI perspective, openness relates very well to many issues

mentioned earlier in this section, and although it may be speculative to say so, there is a feeling that AI may benefit greatly if it decides to embrace some of the key openness concepts that are behind the successes demonstrated by organizations such as the W3C. The same is true for modern biology (note that although this section focuses somewhat on biology, the same is true for the field of neuroscience). Recent breakthroughs in modern biology would arguably not have been possible without the support of highly efficient and freely available, non-patented bioinformatics software like the BLAST algorithm (Quackenbush, Salzberg, 2006). There is a mountain of evidence from many systems for the impact of open source projects (e.g., the R language and environment for statistical computing, the CellML markup language for the description of biological models, or MathML, which describes mathematical notation, capturing both its structure and content).

Further testimony to the road towards openness comes from the rising number of highly cited articles published in, for instance, the Hindawi, the BioMed Central and the Public Library of Science series embracing the open access policy (Eysenbach, 2006). These journals adopt a licensing scheme of the Creative Commons¹⁶, which the Neurocommons project¹⁷ is based on, too. This project develops an open source knowledge management platform for biological research, with a focus on neuroscience. These are just some examples of efforts towards increased openness in science, which will enhance the reproducibility, transparency, dissemination and ultimately the quality of research results. But the trend towards openness is not limited to science. Google, for example, is spearheading an openness initiative by allowing users to extract their data from its proprietary products (e.g., Gmail), thereby opening the door for users who wish to switch to another provider¹⁸.

In addition, an increasing number of biomedical journals now require that the experimental data are deposited in open repositories (such as Gene Expression Omnibus) prior to publication. While this is a laudable practice, it is not sufficient to warrant truly reproducible results. Ideally, open source principles would be adopted not only for the data, but also for the source code and the implemented analysis workflow(s) in order to allow truly reproducible experiments, and to resolve arising scientific debates (Baggerly et al., 2008). While the need for publishing open source code has been pointed out for years, and many authors do provide it voluntarily in supplementary

16. <http://creativecommons.org/>

17. <http://neurocommons.org/>

18. <http://www.dataliberation.org/>

materials, most life science journals do not require this explicitly. It is tempting to speculate, however, that, in the foreseeable future, numerous life science journals will require the publication of open source code (which nonetheless can be copyrighted, of course), analogously to the publication of raw experimental data.

Integration of data, information, and knowledge

Modern life science disciplines generate experimental data at an enormous scale and complexity. These disciplines are also characterized by an enormous body of background knowledge in the form of domain expert knowledge, scientific articles, data warehouses, or graphical representations of, for instance, biological pathways. No comparable knowledge infrastructure exists in *classic* domains for AI applications such as data mining for marketing and retail or stock market prediction. The intelligent integration of this wealth of data, information and knowledge presents unprecedented analytical and technological challenges. Intelligent grid computing for a massive parallel processing of this deluge of geographically distributed data could represent a further interesting challenge for AI-inspired methods (Gil et al., 2004). Natural language processing and text mining of scientific articles may be the tool for digging out hidden gold nuggets of knowledge and for enriching life science data analysis (Natarajan et al., 2006). Still, truly intelligent approaches for leveraging this domain knowledge are only in their infancy. Intelligent human computer interfaces, aided by natural language processing and pattern recognition software offer further scope for approaches from AI in this context (Duch, Mandziuk, 2004).

Advances of modern life science also go hand in glove with the development of increasingly sophisticated technological instruments. Gene expression profiling based on microarray technology, for example, generates thousands of measurements for one single biological specimen, entailing what is commonly known as the *curse of dimensionality* or *small-n-large-p problem*. This problem is exacerbated by an inherently high level of technical and biological noise, systematic errors (such as experimenter bias), and missing data. Conventional statistical methods that originate from *classic* data mining domains (marketing, retail, etc.) are not tailored for the idiosyncrasies of life science data. Here, innovations from machine learning have made significant contributions. But paradoxically, current high-throughput technologies generate both

too many and too little data. Ideally, in order to *understand* system dynamics, we wish to measure multiple biological parameters for the same sample, at the same time and under the same experimental conditions. However, this is generally not feasible with the currently available, highly specialized instruments (e.g., common microarray platforms quantify only RNA abundance of a sample at a moment of time). Hence, time series experiments are necessary to assemble the *snapshots of time* into a full picture that allows the analysis of system dynamics. Still, biochemical experiments generally involve a *breaking-up* of the biological entity in order to measure the molecular components, whereas what we are really interested in are the dynamics of the complete, living entity in its natural environment. For example, primary tumors are composed of a multiple, genetically heterogeneous subpopulations of cells with different proclivity to metastasize, and the tumor microenvironment is a crucial determinant in the pathogenesis (Fidler, 2002).

Computational creativity

Creativity is a frequent element in the mythology, philosophy, or religion of many cultures and it is fair to say that it is one of these malleable concepts again that has fascinated mankind for centuries. The term creativity (e.g., Taylor, 1959) is similarly problematic in nature to the term intelligence. Among a manifold of definitions, (Kryssanov et al., 2001) defines creativity as a cognitive process to generate novel or unconventional solutions. This cognitive process relies on two essential mechanisms: (i) divergent thinking, which generates original, new ideas and (ii) convergent thinking, which logically evaluates a variety of possible solutions to find the optimal one. For instance, philosophers such as Pythagoras and the Pythagoreans contemplated beauty as an objective principle in beings which maintain harmony, order, and balance. In this view, beauty could be the harmony witnessed in the cosmos, but also an expression of normal human behavior. From beauty, however, it is only a small step to creativity - people admire the beauty of artifacts of various kinds but very often these artifacts are the product of a creative process undertaken by an artist. Creativity and beauty are not restricted to the liberal arts only. Many theories in science are considered to be the outcome of an equally creative process and people often mention the elegance or beauty of a theory. In the more recent mid-80s, for instance, science encountered a discourse

with beauty and the creative forces in nature through chaos theory, the inspirational field of science that captured, among many other things, the dynamic of natural systems in images (called *fractals*) of astonishing beauty (Mandelbrot, 1982).

Creativity has strong ties to computing for some time. The goal of web design, for instance, is not to add to the functionality of an application but to make an application aesthetically pleasing and accessible to its users (this does not mean that the production of code is a mundane task. On the contrary, many regard good coding as a highly creative activity). As another example, take the field of humanoid robotics where the physical appearance of a robot, its gestures or its tone of voice may have an impact on user acceptance (e.g., in health care). Computational creativity, which is a relatively young field, relates to many of these issues. The field, which by its very nature is a multidisciplinary scientific endeavor again, carries the vision to better understand human creativity and to construct (via computers and intelligent algorithms) artifacts demonstrating human-level creativity or tools that are able to support the creative processes of humans. In relation to AI it is necessary to mention that creativity has been on the AI agenda for some time (e.g., Boden, 1998) and that it is a challenging question to ask whether AI could assist this creative process, or perhaps even emulate the human process of generating creative solutions? The BISON project is a recent initiative aiming at the design and implementation of a comprehensive computational realization of a bisociative information discovery framework with applications in systems biology¹⁹ (*Bisociation* is a term (Koestler, 1964) coined to denote a creative process involving “... *the sudden interlocking of two previously unrelated skills, or matrices of thought*”). One aspect of BISON is that in large, complex, and heterogeneous domains association techniques fail to discover relevant information that is not related in obvious associative ways, in particular information that is related across different contexts. In reality, however, it is often the case that context/domain-crossing associations are needed in order to generate innovative domains. From this perspective it is an interesting idea to think about the threesome of AI, neuroscience and systems biology where novel information related to these areas is generated by some form of creative AI.

19. <http://www.bisonet.eu/>

Summary

The motivation in this paper was to investigate and position the standing of AI in a modern science context, in particular a modern life science context with a focus on systems biology and neuroscience. An important finding in the paper is the fact that systems biology and neuroscience offer a fertile ground for approaches from AI and that it is fair to say that these fields are united by a great overlap in their defining dreams and visions. Research in these areas (and in life science in general) is characterized by an ever-increasing, complex data and knowledge proliferation, which presents unprecedented challenges on multiple levels, ranging from the acquisition of data up to its high-level interpretation and utilization. With the desire to model and understand complex dynamic systems, these disciplines therefore share a common goal, and they could undoubtedly benefit from a synergetic mutual exchange of ideas and discussion of problems. The paper mentioned that several methods from machine learning and data mining that have their roots in AI are now the backbone of data analysis in systems biology and neuroscience, but also that there may be an atmosphere of individualism that may stand against the creation of a stimulating synergistic environment that can be beneficial for practitioners working in these fields. The paper highlighted that some of the main challenges in such hugely interdisciplinary research environments are not of a technical nature but rather of a cultural nature. In order to overcome these challenges the paper suggested several measures (soft and hard) including, but not necessarily limited to: the ability to learn from other projects and to be inventive, to exploit novel computing paradigms and environments, to develop and adhere to stringent standards and frameworks, to be integrative, and to embrace openness principles. Taking a bird-eye view on AI, the paper believes that at this point in time AI is still a very lively and fascinating field, but also that it may find itself at a crossroads where it has to set its course intelligently and wisely in order to sustain its privileged standing as an inspiring and visionary modern discipline.

Acknowledgments

The first author gratefully acknowledges the support of the Japan Society for the Promotion of Science (JSPS fellowship no. P08625) and of Dr. Hiroaki Kitano, The Systems Biology Institute, Tokyo, Japan. The third author gratefully acknowledges the support of the Japan Society for the Promotion of Science (JSPS fellowship no. S09168) and wishes to express his thankfulness to Dr. Yoko Yamaguchi, laboratory head of Laboratory for Dynamics of Emergent Intelligence, RIKEN Brain Science Institute, Wakoshi, Saitama 351-0198, Japan.

References

All URLs checked
December 2010

- Adami Christoph (2006), *Digital genetics: unraveling the genetic basis of evolution*, "Nature Reviews Genetics", V. 7, n. 2, pp. 109-118
- Ainsworth Claire (2003), *The facts of life*, "New Scientist", V. 178, n. 2397, pp. 28-31
- Albus John, Bekey George, Holland John, Kanwisher Nancy, Krichmar Jeffrey, Mishkin Mortimer, Modha Dharmendra, Raichle Marcus, Shepherd Gordon, Tononi Giulio (2007), *A proposal for a decade of the mind initiative*, "Science", V. 317, n. 5843, p. 1321
- Arbib Michael A. (2003), *The Handbook of Brain Theory and Neural Networks*. 2nd ed., Cambridge, USA, The MIT Press
- Baggerly Keith A., Coombes Kevin R., Neeley Shannon E. (2008), *Run batch effects potentially compromise the usefulness of genomic signatures for ovarian cancer*, "Journal of Clinical Oncology", V. 26, n. 7, pp. 1186-1187
- Bakkum Douglas, Gamblen Philip, Ben-Ary Guy, Chao Zenas, Potter Steve (2007), *MEART: the semi-living artist*, "Frontiers in Neurobotics", V. 1, n. 5
http://www.frontiersin.org/journal/FullText.aspx?s=768&name=neurobotics&ART_DOI=10.3389/neuro.12.005.2007
- Barábasi Albert László, Bonabeau Eric (2003), *Scale-free networks*, "Scientific American", V. 288, n. 5, pp. 60-69

Boden Margaret A. (1998), *Creativity and artificial intelligence*, "Artificial Intelligence", V. 103, n. 1-2, pp. 347-356

Brazma Alvis, Hingamp Pascal, Quackenbush John et al. (2001), *Minimum information about a microarray experiment (MIAME) – toward standards for microarray data*, "Nature Genetics", V. 29, n. 4, pp. 365-371

Brazma Alvis, Krestyaninova Maria, Sarkans Ugis (2006), *Standards for systems biology*, "Nature Reviews Genetics", V. 7, n. 8, pp. 593-605

Brighton Henry, Selina Howard (2003), *Introducing Artificial Intelligence*, USA, National Book Network Inc.

Dubitzky Werner, Azuaje Francisco (2004), *Artificial Intelligence Methods and Tools for Systems Biology*, Berlin, Springer

Dubitzky Werner, Granzow Martin, Berrar Daniel (2007), *Fundamentals of Data Mining in Genomics and Proteomics*, New York, Springer

Duch Wlodzislaw, Mandziuk Jacek (2004), *Quo vadis, computational intelligence?*, in Sinčák Peter, Vaščák Ján, Hirota Kaoru (eds.), *Machine Intelligence: Quo Vadis? (Advances in Fuzzy Systems - Applications and Theory)*, World Scientific, pp. 3-28

Eysenbach Gunther (2006), *Citation advantage of open access articles*, "PLoS Biology", V. 4, n. 5, article 157

Fidler Isaiah J. (2002), *The organ microenvironment and cancer metastasis*, "Differentiation", V. 70, n. 9-10, pp. 498-505

Freidman Roberta (2009), *Reverse engineering the brain*, "Biomedical Computation Review", V. 5, n. 2, pp. 10-17

Gil Yolanda, Deelman Ewa, Blythe Jim, Kesselman Carl, Tangmunarunkit Hongsuda (2004), *Artificial intelligence and grids: workflow planning and beyond*, "IEEE Intelligent Systems", V. 19, n. 1, pp. 26-33

Girimonte Daniela, Izzo Dario (2007), *Artificial intelligence for space applications*, in Schuster Alfons (ed.), *Intelligent Computing Everywhere*, London, Springer, pp. 235-243

Goodacre Royston, Broadhurst David, Smilde Age K. (2007), *Proposed minimum reporting standards for data analysis in metabolomics*, "Metabolomics", V. 3, n. 3, pp. 231-241

HarelDavid(2005), *A Turing-like test for biological modeling*, "NatureBiotechnology", V. 23, n. 4, pp. 495-496

Hochberg Leigh R., Serruya Mijail D., Friebs Gerhard M. (2006), *Neuronal ensemble control of prosthetic devices by a human with tetraplegia*, "Nature", V. 442, n. 7099, pp. 164-171

Holmes Bob (2005), *Alive!*, "New Scientist", V. 185, n. 2486, pp. 28-33

Hucka Michael, Finney Andrew, Sauro Herbert M. (2003), *The systems biology markup language (SBML): a medium for representation and exchange of biochemical network models*, "Bioinformatics", V. 19, n. 4, pp. 524-531

Huw Arnall Alexander (2003), *Future technologies, today's choices. Nanotechnology, artificial intelligence and robotics: a technical, political and institutional map of emerging technologies*, Technical Report, Department of Environmental Science and Technology, Environmental Policy and Management Group, Faculty of Life Sciences, Imperial College London, University of London, UK

Iagnemma Karl, Buehler Martin (2006), *Editorial for Journal of Field Robotics - Special Issue on the DARPA grand challenge*, "Journal of Field Robotics", V. 23, n. 8, pp. 461-462

Izzo Dario, Pettazzi Lorenzo (2007), *Autonomous and distributed motion planning for satellites*, "Journal of Guidance, Control, and Dynamics", V. 30, n. 2, pp. 449-459

Kak Subhash (2007), *Quantum mechanics and artificial intelligence*, in Schuster Alfons (ed.), *Intelligent Computing Everywhere*, London, Springer, pp. 81-101,

Kelly Stephen, Schuster Alfons (2005), *Application of a fuzzy controller on a lego mindstorms robot*. Proceedings of the 1st International Conference on Automation, Control and Instrumentation (IADAT-05), 2-4 February 2005, Bilbao, Spain, pp. 200-203

Kephart Jeffrey O., Chess David M. (2003), *The vision of autonomic computing*, "Computer", V. 36, n. 1, pp. 41-50

Khan Javed, Wei Jun S., Ringnér Marcus (2001), *Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks*, "Nature Medicine", V. 7, n. 6, pp. 673-679

Kitano Hiroaki (2007), *Towards a theory of biological robustness*, "Molecular Systems Biology", V. 3, article 137

Kitano Hiroaki (2004), *Cancer as a robust system: implications for anti-cancer therapy*, "Nature Reviews Cancer", V. 4, n. 3, pp. 227-235

Kitano Hiroaki (2002), *Systems biology: a brief overview*, "Science", V. 295, n. 5560, pp. 1662-1664

Koestler Arthur (1964), *The Act of Creation*, London, Hutchinson

Kryssanov Victor, Tamaki Hisashi, Kitamura Shinzo (2001), *Understanding design fundamentals: how synthesis and analysis drive creativity, resulting in emergence*, "Artificial Intelligence in Engineering", V. 15, n. 4, pp. 329-342

Lander Eric S., Linton Lauren M., Birren Bruce (2001), *Initial sequencing and analysis of the human genome*, "Nature", V. 409, n. 6822, pp. 860-921

Lebedev Mikhail, Nicoletis Miguel Angelo (2006), *Brain-machine interfaces: past, present and future*, "Trends in Neurosciences", V. 29, n. 9, pp. 536-546

Luck Michael, Aylett Ruth (2000), *Applying artificial intelligence to virtual reality: intelligent virtual environments*, "Applied Artificial Intelligence", V. 14, n. 1, pp. 3-32

Mandelbrot Benoit B. (1982), *The Fractal Geometry of Nature*, New York, W.H. Freeman and Co.

Marmanis Haralambos, Babenko Dmitry (2009), *Algorithms of the Intelligent Web*, Greenwich, Manning Publications Co.

Miller Melissa B., Bassler Bonnie L. (2001), *Quorum sensing in bacteria*, "Annual Review of Microbiology", V. 55, pp. 165-199

Millington Ian (2006), *Artificial Intelligence for Games*, San Francisco, USA, Morgan Kaufmann

Natarajan Jeyakumar, Berrar Daniel, Dubitzky Werner et al. (2006), *Text mining of full-text journal articles combined with gene expression analysis reveals a relationship between sphingosine-1-phosphate and invasiveness of a glioblastoma cell line*, "BMC Bioinformatics", V. 7, article 373

Negnevitsky Michael (2004), *Artificial Intelligence: A Guide to Intelligent Systems*. 2nd ed., Harlow, UK, Addison Wesley

Păun Gheorghe (2006), *Introduction to membrane computing*, in Ciobanu Gabriel, Pérez-Jiménez Mario J., Păun Gheorghe (eds.), *Applications of Membrane Computing*, Berlin, Springer, pp. 1-42

Pfeifer Rolf, Scheier Christian (2000), *Understanding Intelligence*, Cambridge, USA, The MIT Press

Potter Steve M. (2007), *What can artificial intelligence get from neuroscience?*, in Lungarella Max, Bongard Josh, Pfeifer Rolf (eds.), *Artificial Intelligence Festschrift: The Next 50 Years*, Berlin, Springer, pp. 174-185

Purves Dale, Augustine George J., Fitzpatrick David (2004), *Neuroscience*. 3rd ed., USA, Sinauer Associates

Quackenbush John, Salzberg Steven L. (2006), *It is time to end the patenting of software*, "Bioinformatics", V. 22, n. 12, pp. 1416-1417

Rasmussen Steen, Chen Liaohai, Deamer Davis (2004), *Transitions from non living to living matter*, "Science", V. 303, n. 5660, pp. 963-965

Rodriguez Luis, Jiang Haibei, Stark Kirsten, Bell Scott, Kortenkamp David (2007), *Validation of heuristic techniques for design of life support systems*. Proceedings of the Workshop on Artificial Intelligence for Space Applications (IJCAI-07), January 2007, Hyderabad, India

Saha Debashis, Mukherjee Amitava (2003), *Pervasive computing: a paradigm for the 21st century*, "Computer", V. 36, n. 3, pp. 25-31

Saigusa Tetsu, Tero Atsushi, Nakagaki Toshiyuki, Kuramoto Yoshiki (2008), *Amoebae anticipate periodic events*, "Physical Review Letters", V. 100, n. 1, Article ID 018101, 4 pages

Santhanam Gopal, Ryu Stephen I., Yu Byron M., Afshar Afsheen, Shenoy Krishna V. (2006), *A high-performance brain-computer interface*, "Nature", V. 442, n. 7099, pp. 195-198

Sato Naoyuki, Yamaguchi Yoko (2008), *Computational model-based prediction of human episodic memory performance based on eye movements*, "IEICE Transactions on Communications", V. 91b, n. 7, pp. 2142-2143

Savage Liz (2008), *Forensic bioinformatician aims to solve mysteries of biomarker studies*, "Journal of the National Cancer Institute", V. 100, n. 14, pp. 983-987

Schena Marc, Salon Dari, Davis Ronald W., Brown Patrick O. (1995), *Quantitative monitoring of gene expression patterns with a complementary DNA microarray*, "Science", V. 270, n. 5235, pp. 467-470

Schuster Alfons (2008), *Robust Intelligent Systems*, London, Springer

Schuster Alfons (2008), *Robustness in nature as a design principle for artificial intelligence*, in Schuster Alfons, *Robust Intelligent Systems*, London, Springer, pp. 173-197

Schuster Alfons (2008), *DNA algorithms for Petri net modeling*, "Informatica", V. 32, n. 4, pp. 421-427

Schuster Alfons (2004), *DNA algorithms for rough set analysis*. Proceedings of the 4th International Conference on Intelligent Data Engineering and Automated Learning (IDEAL-03), 21-23 March 2003, Hong Kong, pp. 498-513

Sgarbas Kyriakos N. (2007), *The road to quantum artificial intelligence*, in Papatheodorou Theodore, Christodoulakis Dimitris, Karanikolas Nikitas (eds.), *Current Trends in Informatics*, Athens, New Technologies Publications, pp. 469-477

Tagkopoulos Ilias, Liu Yir-Chung, Tavazoie Saeed (2008), *Predictive behavior within microbial genetic networks*, "Science", V. 320, n. 5881, pp. 1313-1317

Taylor Irving (1959), *The nature of the creative process*, in Smith Paul (ed.), *Creativity: An Examination of the Creative Process*, New York, Harcourt Brace, pp. 51-82

Troester Melissa A., Millikan Robert C., Perou Charles M. (2009), *Microarrays and epidemiology: ensuring the impact and accessibility of research findings*, "Cancer Epidemiology Biomarkers and Prevention", V. 18, n. 1, pp. 1-4

Turing Alan M. (1952), *The chemical basis of morphogenesis*, "Philosophical Transactions of the Royal Society", V. 237, n. 641, pp. 37-72

http://www.cecm.usp.br/~cewinter/aulas/artigos/2009/Turing_1952.pdf

Turing Alan M. (1950), *Computing machinery and intelligence*, "Mind", V. 59, pp. 433-460

Venter Craig J., Adams Mark D., Myers Eugene W. et al. (2001), *The sequence of the human genome*, "Science", V. 291, n. 5507, pp. 1304-1351

Voit Eberhard, Schwacke John (2007), *Understanding through modeling: ahistorical perspective and review of biochemical systems theory as a powerful tool for systems biology*, in Andrzej Konopka, *Systems Biology: Principles, Methods, and Concepts*, USA, CRC Press, pp. 27-82

Sintesi

Rispondere alla domanda "Quo vadis, intelligenza artificiale?" vuol dire indagare la natura dell'intelligenza artificiale e situarla correttamente nel più ampio contesto della scienza moderna. Emergono così immediatamente una fitta rete di scambi fecondi e di reciproche interrelazioni, significative intersezioni e sovrapposizioni di approcci e di contenuti, di temi e di prospettive, di chiavi interpretative e di obiettivi comuni, che collegano l'intelligenza artificiale ad alcune discipline d'avanguardia:

- *la biologia di sintesi, intesa come la progettazione e la costruzione di parti, mezzi e sistemi biologici affatto nuovi, ovvero la riprogettazione di sistemi biologici già esistenti in natura per scopi utili;*
- *la biologia dei sistemi, lo studio di sistemi biologici complessi, delle loro proprietà, interazioni e dinamiche, con il suo emergente paradigma olistico in contrasto con l'approccio riduzionistico prevalente in passato;*

- *il campo vastissimo della neuroscienza, riassumibile efficacemente nell'espressione "dalle molecole alla cognizione".*

Per rafforzare l'intelligenza artificiale come linea di sviluppo chiave in tutti questi ambiti di ricerca e potenziare la sua dimensione per eccellenza plurale e interdisciplinare, debbono essere però raccolte alcune sfide, non di natura tecnica, ma essenzialmente culturale.

Bisogna innanzitutto essere capaci di imparare da progetti differenti, anche assai distanti, ed essere aperti e disponibili al cambiamento, superando particolarismi, barriere e pregiudizi disciplinari, chiusure individualistiche.

È quindi indispensabile essere inventivi, sia spaziando in orizzonti inesplorati che approfondendo e riscoprendo con occhi nuovi ambiti già indagati: ad esempio, i giochi per il computer, la realtà virtuale, le nanotecnologie, l'elaborazione quantistica, l'esplorazione dello spazio.

Un terzo imperativo è il pieno sfruttamento di aree e paradigmi elaborativi nuovi: fra gli altri, si sono rivelati ricchi di sviluppi incoraggianti nell'ambito della ricerca sulla vita artificiale alcuni concetti basilari propri del comportamento collettivo, quali gli approcci bottom-up e il decentramento, ossia l'assenza di una struttura centrale insieme di controllo e di comando, cui corrispondono il comportamento e l'intelligenza "a sciame". Ancora, è una necessità imprescindibile lo sviluppo di standard, i quali possono garantire nell'orizzonte globale la qualità dei prodotti, soddisfacendo criteri come la correttezza, l'affidabilità, la riproducibilità, la robustezza, la comprensibilità, l'interoperabilità e la manutenibilità e chiamando altresì in causa prepotentemente l'aspetto terminologico.

Altre cruciali esigenze sono la conformità a strutture computazionali e statistiche rigorose e, strettamente legata ad essa, l'adozione di principi di apertura - da non confondersi con il decentramento e la disorganizzazione - e del modello open source in primis.

L'integrazione di dati, informazioni e conoscenze richiede poi un forte impegno: le moderne scienze della vita producono infatti dati sperimentali di enorme scala e complessità, che superano di gran lunga i limiti dei metodi statistici sorti dai dominî classici di data mining; l'integrazione intelligente di una simile mole di dati, informazioni e conoscenze pone pertanto problemi di carattere analitico e tecnologico senza precedenti, alla cui soluzione l'intelligenza artificiale può offrire evidentemente un contributo decisivo.

Risulta infine determinante la creatività computazionale: quantunque sia difficoltoso circoscrivere in maniera soddisfacente il concetto di "creatività", essa può essere utilmente definita come un processo cognitivo finalizzato a generare soluzioni nuove o non convenzionali, fondato su due meccanismi essenziali, il pensiero divergente, che produce idee nuove e originali, e il pensiero convergente, che valuta una molteplicità di soluzioni possibili fino a scegliere quella ottimale. La creatività è del resto al centro

del dibattito dell'intelligenza artificiale e la domanda principale è se questa possa davvero assistere il processo creativo o forse persino emulare il processo umano di generazione di soluzioni creative.

Ciò che risulta in ogni caso più promettente per il futuro dell'intelligenza artificiale - come per la maggior parte dei settori disciplinari - è la metodica transdisciplinarietà, l'associazione libera, fluida e dinamica e l'ibridazione fra contesti multipli e ambiti eterogenei.

