Paper:

The Omnipresent Computing Menace to Information Society

Alfons Schuster* **and Daniel Berrar****

*Laboratory for Dynamics of Emergent Intelligence, RIKEN Brain Science Institute
Wako-shi, Saitama 351-0198, Japan
**Interdisciplinary Graduate School of Science and Engineering, Tokyo Institute of Technology
G3-45, 4259 Nagatsuta, Midori-ku, Yokohama 226-8502, Japan
E-mail: aschuster@brain.riken.jp, berrar.d.aa@m.titech.ac.jp
[Received February 21, 2011; accepted May 27, 2011]

Computers have evolved from mere number crunchers to systems demonstrating an astonishing degree of sophistication, decision-making ability, and autonomy. Silicon is no longer the only substrate facilitating information processing. Despite these progresses, machine intelligence is still far from rivaling human intelligence. Nonetheless, we might be all too ready to rely on inferior agents for decision making, to give away sensitive information without fully understanding the consequences involved, or to tinker with genetic code to program carbon-based machines without fully appreciating the risks. This article explores the potentials and risks that information societies may face in the wake of current and emerging intelligent computing paradigms.

Keywords: intelligent computing, carbon-based computing, information society, information ethics

1. Introduction

On September 5, 1977, NASA launched the Voyager 1 spacecraft with the mission to study the outer solar system and potentially interstellar space. Voyager 1 has now become the farthest human-made object from Earth.[1] Voyager 1 is not only a demonstration of human ingenuity and modern technology, Voyager 1 is also an ambassador of mankind and an authoritative example that mankind has entered the *information age* – that mankind has successfully made the historic transition from an industrial society to an *information society*. This claim manifests itself in particular through the famous gold-plated audiovisual disc that Voyager 1 carries on its frame. This disc includes photos of the Earth and some of its life forms, as well as music and the sounds of whales. The disc and its mission therefore create an image of Voyager 1 of being both a tireless explorer as well as a fingerprint of mankind in the vastness of space. But is it really unproblematic to give away crucial information about mankind in such a lighthearted way [1]? Imagine what could happen to mankind if the probe is ever picked up by an intelligent but hostile alien life form.[2]

Admittedly, such mind games are mere science fiction. But as far-fetched as the Voyager 1 example may have been, it holds as a blueprint for a wide range of problem scenarios that an information society may encounter in general and the modern computer-permeated world in particular.[3] For example, a computer-based information exchange may happen on levels of various degrees of sensitivity, privacy, sophistication, and trust [2]. It may involve a range of requirements in terms of safety, confidentiality, reliability, verifiability, or risk, and it may not be clear whether such an information exchange might turn out to be harmful to any of the entities involved in it. As an example, think about the problem of security in online banking or any kind of online transaction for that matter. There simply is no 100% guarantee for any such business processes. Some of these critical issues may be rooted in the human camp, whereas others may originate from fundamental restrictions on computer hardware or software, or other process or system limitations.

This paper explores several of these issues from an *intelligent computing* perspective. This perspective considers intelligent computing in its widest definition, including systems that are able to learn and improve their performance in order to accomplish specific tasks. It also considers intelligent computing encompassing unconventional computing paradigms and applications that require ingenuity or creativity, new substrates, and novel ways of human-computer interaction. In addition, although intelligent systems are permeating and shaping information society in many ways, we need to acknowledge the vastness of the field of intelligent systems, which ranges from intelligent task scheduling to computation with slime molds. Because of this wide range, this article cannot do justice to all relevant issues. Therefore, the article focuses on the issues arising from the interaction between humans and intelligent systems (Sections 2–4), and discusses the impact these issues have on information society (Sections 5–7).

1. http://voyager.jpl.nasa.gov/

2. Star Trek fans will of course remember that the first Star Trek feature film revolved around a plot similar to the content provided in this introductory section.
3. Conceptually, the study of information societies belongs to the field of information studies, which is a multidisciplinary field investigating a wide range of information systems and information issues.

2. Information is Not Knowledge

A first area of computer-based impact on information society comes from a relatively unexpected direction, namely from the large amounts of data generated in highly complex domains, and the utilization of powerful modern computer systems for analyzing this data deluge. A key problem lies in the understandability and verifiability of computer-generated output. For instance, the Tevatron particle accelerator at the Fermi National Accelerator Laboratory (Fermilab) in Batavia, Illinois, USA, produces millions of collisions every second from which only a fraction of some thousand collisions per second (data that still amounts to over 10 petabytes) are kept for further analysis. The time frame to analyze the raw data is in the order of years. It is clear that such an amount of data can no longer be analyzed *manually*, i.e., by a human expert taking full control of the process, and that new ideas and sharper tools will be needed to extract meaningful information from the data. If the data is automatically analyzed (e.g., by tools using artificial intelligence or machine learning algorithms), however, then it is possible that the generated results may be difficult to interpret and explain. For example, Artificial Neural Networks (ANNs) are an extremely versatile tool for finding various types of relationships in large data volumes, but ANNs are also *black boxes*, i.e., they generally do not provide any means for explaining how the results were obtained.

Work by Schmidt and Lipson addresses the same problem [3]. Their work, which is related to automate (i.e., computerize) the process of finding natural laws and describing them by mathematical equations, introduces a rather general problem solving procedure that is based on evolutionary computing. In the evolutionary approach, roughly, a problem-solver (e.g., an algorithm) generates randomly a large number of potential solutions (e.g., equations) for a system under study. Imagine finding equations describing the dynamic behavior of the stock market. Each potential solution is evaluated in how well it describes or approximates the behavior of the target system. The most promising solutions are selected and some of their features combined or updated. This generates a new pool of, on average, improved potential new solutions. The process of selecting, breeding, and changing is repeated until the evolutionary algorithm generates an acceptable (ideally an optimal) solution for the system (e.g., a set of equations describing the dynamic of the stock market). Under certain assumptions, the approach works astonishingly well. But even in these cases, one of the problems with the approach can be as follows. The algorithm may be fed with data and generate an equation that describes the data well, but it may be extremely difficult to interpret the *meaning* of the generated output. In the end, the user has a tool that works well, but nobody knows why, because the tool has no explanation facilities. It is possible, of course, to turn a blind eye to this situation and to acknowledge that the machine is doing a good job, despite its black box character, but perhaps this is an unsatisfying outcome for the human quest for knowledge.

Related examples can be found in the life sciences, predominantly in post-genomic biology, where high-throughput technologies have caused a paradigm shift from a traditionally *hypothesis-driven* to a *data-driven* analysis [4]. Whereas in the past, it was common to investigate one gene or one protein at a time, it is nowadays common to screen thousands of traits for a biological specimen in a single experiment. Data storage capabilities are constantly going up while the costs for generating data are going down. The real challenge is now how to make intelligent use of this generated data. The machine learning and data mining community has contributed immensely to the analysis of the high-throughput data deluge, which in turn has motivated the development of new learning algorithms. The toolbox of sophisticated analytical methods is huge, and sharper tools are constantly added. But having such an arsenal of tools at one's disposal may be dangerous, too, as it can lead researchers to believe that the data can always be meaningfully analyzed – which is a fallacy because there are some biases that cannot be controlled for [5]. Exploratory analyses without a predefined hypothesis have also been criticized and compared to *fishing expeditions* – after all, once the data set is sufficiently large, you are bound to find something apparently interesting, like finding your date of birth in the infinity of π. Human beings are excellent at detecting patterns, which can be explained by evolution: being quick at detecting friend and foe brought a survival advantage. But our craving for patterns can also lead us astray into seeing patterns that simply do not exist. Some data simply are random, but "*[n]othing is so alien to the human mind as the idea of randomness*" [6].

On the other hand, the very nature of science is exploratory. Research often takes twists and turns before something interesting is found, and, indeed, the literature is replete with examples of serendipitous discoveries (e.g., that of penicillin). Despite this encouraging feature of the scientific endeavor, sophisticated analytical tools still often leave us with no more than information extracted from data. But we aim for knowledge, and in the words of Albert Einstein, however, "*information is not knowledge*." What, then, is knowledge?

3. Knowledge is Action

A fundamental problem related to this question is due to a current practice of scientific analysis that prevails in the life sciences such as biology, psychology, medicine, and epidemiology, but also in computer science and some engineering-related fields. This practice is dominated by Fisherian statistics, p-values, and the quest for *significant* results. At the root of the problem is the *null hypothesis testing*: assuming that the null hypothesis is true, what is the probability of observing data as extreme as or more extreme than the actually observed data? This probability is the p-value. Alas, this is rarely what the scientist is (or rather, should be) actually interested in. What the researcher wishes to find out is the probability that the hy-

pothesis is true, given the observations. Interpreting the p-value as this probability is known as the fallacy of the transposed conditional. Without assumptions about the prior probability of the hypothesis, hence, without looking at the problem through Bayesian lenses, it is impossible to calculate the probability of the hypothesis. The problems of the p-value and null hypothesis testing have been written about for decades (e.g., see [7] for an account of the most common misconceptions). Yet, the problem is not only related to the human interpretation of p-values. The current practice – pitting one hypothesis (the null hypothesis H_0) against an alternative hypothesis (H_1) and using p-values as decision criterion – results from a misalliance between inherently incompatible concepts: the Neyman-Pearson hypothesis test and Fisher's p-value [8]. It may not be widely appreciated that this misalliance is at the root of many problems and paradoxes in statistical analysis [8]. In addition, when Fisher's p-value is disentangled from the Neyman-Pearson hypothesis test, it can be seen that these great scientists did not exactly agree on what the scientific method should be about. Thus, probing deeper into this issue invariably leads us to a fundamental philosophical discussion on the nature of the scientific method itself.

Deficiencies of statistical hypothesis testing have been described in various fields of research, including epidemiology [8] and medicine [9], so one has to wonder why hypothesis testing is still so deeply entrenched in research practice. Johnson [10] offers several possible explanations, but essentially concludes that the reasons are psychological rather than rational. McClosky and Ziliak [11] see reasons to be worried about the vital consequences of the current practice of statistical significance testing because it may even kill people, which is a clear indicator that scientific insight influences decisions in public health. Poole [12], on the other hand, cautions that science and decision-making are not the same, as in science, "*... we seek to learn, to explain and to understand*," whereas "*in decision-making, we seek reasons to act or refrain from acting*." It is possible to agree with Poole in that understanding and deciding to act are indeed different thought processes. But should understanding not always precede deciding (to act), at least in an ideal world? These problems deserve to be taken very seriously, as statistical principles rather frequently provide fundamental underpinnings of intelligent systems such as decision support or knowledge-based systems [13].

4. Trust in Intelligent Systems

4.1. Privacy and Security

One area where autonomous intelligent agents infringe the privacy of individuals or that of other entities is that of covert robotics. The great ambition in the field is to design autonomous robots that can operate, or frankly speaking spy, around their targets without being detected [14]. These robots are typically equipped with sensors that allow them to create a map of their operation surroundings as well as software that enables them to avoid or evade attention. Another area closely related to this is the application of robotics in the field of deception. Wagner's work [15] mentions a definition for deception as a "*process by which actions are chosen to manipulate beliefs so as to take advantage of the erroneous inferences*." The paper also mentions that the ability to deceive may be an indicator of theory of mind and of social intelligence, which is a meaningful assumption that, however, should not be too surprising to those familiar with the so-called Turing Test (a benchmark test for machine intelligence of historical dimension).

An area where security plays a main role is the intelligent (smart) home domain. Here, integrated computer-based devices with various degrees of decision-making autonomy are blended into environments of everyday life until they are indistinguishable from it in order to provide increased quality of life for residents. A particularly caring (and arguably equally lucrative marked) in this context are fully automated homes and healthcare environments for elderly or disabled people. It is clear that not only vulnerable people such as disabled, or elderly people, or children, but anyone needs to be assured that his or her personal health, security, and safety must be guaranteed at all times in such a highly computerized environment.

Remaining in the health and care sector, we find that molecular biology has arguably revolutionized biomedical research in the late 20th century with the advent of high-throughput screening technologies. Technologies for gene expression and sequence analysis are paving the way for a personalized medicine, i.e., a medicine that is tailored for the genomic profile of an individual person. Genomic patient data are already publicly available in large databases, and their volume increases nearly exponentially. Intelligent systems are necessary to *mine* these data for the molecular characterization of diseases, the identification of patients at higher or lower risk for developing certain types of diseases, and for the prediction of a patient's drug response. However, the potential of misuse also exists. Imagine an organization that excludes an individual from access to some goods or services on the basis of the genetic profile of that individual. For example, an individual may be refused access to a health insurance or a particular job because his or her genetic makeup matches a high-risk profile. Heeney [16] correctly mentions that this process, known as *redlining*, may become a new form of serious discrimination. In addition, in the ideal case, in publicly available genomic databases, no personal patient data are disclosed. However, that does not mean that all data are guaranteed to remain anonymous. In fact, if it is known that an individual has participated in a genome-wide association study, then it is relatively straightforward to link that individual to his or her medical data [17].

4.2. Safety and Reliability

Now that the keyword safety has been mentioned, it is easy to quickly identify a few more scenarios where

this important issue comes to the fore. The field of telemedicine where the safety and perhaps the survival of patients depends on the need to communicate critical expertise and operations reliably in real-time over the Internet is one such example. Other examples can be a space mission where control tasks (e.g., flight control or monitoring the well-being of crew members in a space station or, imagining a more futuristic setting, in a fully automated life support system on Mars) are taken over by an automated (autonomous) computer system. Such a system may use artificial intelligence software or models and concepts from brain research and neuroscience (e.g., as in pervasive or ubiquitous computing, which uses the autonomic nervous system as a model for building systems that can act autonomously and demonstrate features such as self-configuration, self-healing, self-optimization, and self-protection, among several other useful features). In all these cases it is clear that the people involved (patients, astronauts, etc.) find themselves greatly exposed to various types of threats, and that there are huge demands on the systems involved to guarantee well-being and safety to the highest degree possible.

Unfortunately, there is more bad news. A considerable problem permeates all the areas mentioned before. This problem applies to the production of all large and complex systems in general and to the production of large and complex software systems in particular. It is the problem of requirements analysis and specification. Analyzing and specifying *all* requirements for a large and complex systems is a notoriously difficult task that usually penetrates all phases of a product life cycle (essentially, analysis and specification, design, implementation, testing, and maintenance). Indeed, the problem is of a scale that, in practice, perfect, error-free software remains an illusive goal at best – despite the application of a huge body of best practices, techniques, and experiences [18]. For instance, studies have shown that defect rates in typical commercial software have been estimated at 10 to 17 per 1000 lines of code (with systems produced under the open-source movement showing lower failure rates). Certainly, it is possible to achieve lower defect rates via rigorous specification procedures involving formal mathematics-based methods (e.g., some commercially successful software houses that use the best available practices regularly achieve defect rates of 0.03 to 0.05 per 1000 lines of code), but even in these cases the problem of defective software remains and creates the following dilemma. If large and complex software cannot be produced without errors, then it is fair to ask for who is accountable if things go really wrong?

4.3. Responsibility and Accountability

From a potentially larger pool of examples, consider applications where a large number of intelligent, autonomous, decentralized software agents operate as sophisticated and independent decision-makers in complex environments. One such complex application may involve organizing and regulating the supply of electricity in a country. In the ideal case, the load profile in an electricity grid is flat. In reality, however, the profile is volatile and depends on sudden surges of demand. In order to smother volatility and to manage electricity grids more efficiently, researchers investigate the ability of so-called smart agents and smart grids [19]. One task that smart agents aim to tackle is to provide alternatives to the common demand-supply pattern in the form of a more dynamic grid in which supply and demand are in continual feedback. To some extent, this may be achieved by recruiting energy from temporary storage units (e.g., batteries) at peak-times when demand is high and to draw energy from the grid at peak-off times when demand is low.

As good as it may sound, there is a general downside to the smart agent approach. One of the main problems lies in the term smart agent and the degree of smartness (intelligence) given to these entities.[4] The previous section already indicated that it is wishful thinking to expect error-free software in large application areas. In addition, and this is one of the more painful lessons from many years of artificial intelligence research, it is substantially more difficult to write error-free intelligent software for intelligent systems. Of course, on top of all this rests the complexity dimension of large systems (e.g., that of the electricity grid) makes it simply impossible to judge, and make accurate statements about the overall state of affairs and dynamic behavior of a system at any moment of time. So, if the smart (intelligent) agents are truly autonomous and software is almost guaranteed to be error prone, and nobody can really say at all times what is going on in detail in the grid, then who is answerable, responsible, or accountable when things go really wrong? Actually, the answer to this question is not so simple, though the field of information ethics may shed some light on it.

5. Information Ethics

Information ethics is a wide field and concerns most of the issues mentioned earlier in this text but also issues (e.g., ownership, copyright protection, and intellectual freedom) that have been neglected in this article so far. From this range of topics, this selection adjusts its focus onto a discussion that revolves around moral topics. The topics of violence in computer games, the usage of brain machine interfaces in military, as well as the topics of robot ethics and transhumanism appear under such an adjustment quite naturally.

Data analysts are mining the data mountains that are generated by millions of users in a diversity of applications (ranging from fantasy games to so-called serious games where, for example, rescue workers prepare and train for action in a factory blaze) for gaining insights into

4. The term *intelligence* is often avoided and bypassed by terms such a *smart*. One reason for this may involve the opinion that the term intelligence should apply exclusively to humans. Another reason may be that various visionary (sometimes premature) promises in artificial intelligence may have given the term a negative connotation that better should be avoided in some computer application domains.

a wider spectrum of human behaviors [20]. Some of the motives in the industry are the creation of personalized games and environments that adapt to each user's abilities and interests, or the generation of higher levels of user involvement or attractiveness by moving games from requirements in ability and skill (as in shooter games) to a more intellectual and emotional level.

In terms of user involvement, an interesting technology that has found its way into the military domain are brain-computer interfaces, i.e., devices that acquire and transform neural signals into actions intended by their user [21]. According to this research, brain-computer interfaces may be able to provide humans with rudimentary control over computer systems and robotic devices. If these systems and devices can live up to their expectations, then it is possible to see them in domains ranging from therapeutic applications to restore function after injury, to human performance enhancers, and to devices that can be used in military operations involving various degrees of risk, severity, and impact. Here, we may also mention the young field of robot ethics, a field that is fueled by the assumption that in the future, humanity is going to coexist with robots, and that the ethical implications in such a partnership need to be firmly investigated [22].

Another ethical dilemma arises when people leave (intentionally or unintentionally) a trace of data behind them (e.g., on the Internet, in databases of various types, or on private or other computer-based systems). In case one wishes to make this data disappear, it turns out rather quickly that this is not an easy task and the terms *data cemetery* and *digital soul* may describe the problem well. Both terms provide an interesting link to another instance of human-computer interaction that requires a considerable degree of sensitivity: computer applications in real cemeteries. Such applications use near-field communication technology for an information exchange between, say, an electronic device embedded in the gravestone and a mobile phone. A visitor using such an application can then be provided with information about a deceased person. Clearly, the presentation of data and information has changed dramatically in recent years. But even traditional information processing institutions do not remain untouched by this development. Newspapers are now available online; it is possible to download and read entire encyclopedias on e-books; to access a myriad of information on small hand-held devices; and, of course, there is YouTube, Facebook, and a variety of other social network services. And all of this is possible through modern technology and its ongoing advancement. Although this advancement seems to be unstoppable, there are critical issues indicating that it is necessary to keep a watchful eye on this development. Arguments about online game addiction and its effects of social isolation, up to cases where people collapse or even die from exhaustion because they spent too much time in front of a computer, are known.

In terms of the integrity and dynamic of social networks, it is relatively easy for an impostor to put up a fake profile, either a profile for a nonexistent person, or a profile for a real person but without that person's consent. It is even possible that the entity behind a user account is a so-called socialbot, a piece of software designed to mimic social interaction with humans or other agents. In a recent socialbots competition, socialbots have been able to attract hundreds of so-called followers on Twitter [23]. These socialbots have demonstrated the ability to heavily shape the structure of social network groups. Socialbots could therefore manipulate social networks on a large scale – for good or ill. From purely software-based entities with social abilities it is only a small step, however, to envisage entities synthesizing artificial and biological modes of interaction.

6. From Silicon- to Carbon-Based Computing

The relatively young field of synthetic biology, for instance, focuses on the design and engineering of new biological parts and biomolecular computers. Pioneering work in this direction has been done by Benenson et al. who designed a molecular Turing machine [24] and a stochastic molecular automaton [25]. Synthetic biology also focuses on synthetically creating gene regulatory networks [26] and even assembling new artificial life forms [27]. It has even become possible to reprogram entire organisms with specific functions. For example, Tamsir et al. [26] implemented all possible two-input gates including the XOR and equal function in a colony of *Escherichia coli* bacteria. Levskaya et al. [28] reprogrammed these blind bacteria to create a two-dimensional chemical image.

The scope for such genetically engineered machines is large, and of particular interest are of course applications in the biomedical field. Independent of the application domain, however, there should be an alertness about the risks that the (intentional or accidental) release of synthetic organisms may harbor. As in genetically modified crops, horizontal gene transfer[5] in synthetic organisms is a serious concern. It may be possible that synthetically modified bacteria transfer parts of their genetic program to other, naturally occurring bacteria and thereby altering their functions, with perhaps detrimental consequences. In principle, synthetically created artificial life forms are also subject to Darwinian processes; hence, they could evolve and adapt despite built-in functions triggering their self-destruction outside of the lab. It may therefore also be possible that a natural species is suddenly confronted with a competitor having an "unfair" evolutionary advantage. In both scenarios, the ramifications are unpredictable for natural ecosystems in general or human beings in particular. Note, that the human body contains at least ten times more bacterial cells than human cells [29]. Bacterial cells play a pivotal role in physiological and immunological processes in the human body. But despite groundbreaking research [30], very little is actually truly

5. Horizontal gene transfer is the transfer of genetical material from one organism A to another organism B that is not an offspring of A.

understood about the intricate cell-to-cell communication systems and the orchestrated interplays of gene regulatory networks. And that holds for one species – systems such as the human body that host interacting sub-systems are even orders of magnitude more complex. The potential of synthetic biology in biomedical and environmental applications is huge [31], but summoning a "living servant" may not be without risk, as Johann Wolfgang von Goethe already showed in the 1797 poem *The Sorcerer's Apprentice*.

Alan Turing, great grand-father of artificial intelligence, believed that, at some stage, it would be possible to construct machines that would simulate the behavior of the human mind [32]. Turing adopted systems-level approaches in his studies on the chemical basis of morphogenesis [33], and could therefore be rightfully regarded as a pioneer in systems biology, too. But even Turing might arguably not have foreseen that about five decades later, scientists would be able to reprogram entire organisms. When Turing wrote in 1951 [32] that "*[the machines] would be able to converse with each other to sharpen their wits. At some stage therefore we should have to expect the machines to take control ...*," Turing was talking about intelligent machines that emulate the human mind. In [32], *intelligence* refers to human intelligence, reasoning, and insight. However, by broadening the ill-defined term *intelligence* [34], we may regard collectives of engineered bacteria, programmed to perform a specific task, as an intelligent system. Bacteria communicate – or converse – with each other via *quorum sensing*, a chemical protocol that involves the production, release, and detection of signal molecules [30]. Using this *molecular language*, bacteria can indeed coordinate their collective behavior. Recent research has also revealed evidence for interspecies quorum sensing between different strains of bacteria and even between bacteria and eukaryotes[6] [35]. These examples indicate that synthetically engineered machines may indeed "sharpen their wits" through evolutionary principles, which is reminiscent of Turing's vision, albeit with a more dystopian twist.

Clearly, *wet* computing is still at its infancy, but some parallels to the development of modern computers are all too visible. In the not-too-distant future, it may indeed be possible to manipulate a cell or collectives of cells with an instructing language, similar to the programming languages that are currently in use to instruct silicon-based hardware. The molecular ingredients in synthetic biology's kitchen are relatively cheap and do not require highly specialized facilities. In addition, the basic equipment of university labs is sufficient and the costs are further falling [31]. Now, almost a decade has passed since the creation of a synthetic infectious virus [36], and synthetic biology progresses with giant leaps. In the early days of silicon computing, few may have foreseen how easy it would become to write and propagate malicious software, i.e., computer viruses. In a distant future, will we witness analogies in synthetic biology, only with real viruses involved?

6. Eukaryotes are a family of more complex organisms that have a membrane-bound nucleus. In contrast, bacteria belong to the family of prokaryotes, organisms that lack a true cell nucleus.

7. Conclusion

The great potential intelligent systems demonstrate necessarily involving risks. Some of these risks are well-known (e.g., the risks of stolen passwords) and in such situations a user can take precautionary steps. Other situations, such as trails of personal information in online social networks, or the manipulative potential of socialbots, for instance, involve less well-known risks. Arguably, however, these risks may not necessarily be detrimental to the well-being of human beings. Unfortunately, looking at the wider scope of intelligent systems (e.g., including their growing number of utilizations and application domains) also reveals more intimidating and threatening, and more unpredictable risks. This paper specifically addressed emerging intelligent systems that are based on molecular machines belonging to this latter category. The paper also reported on risks that may not be generally appreciated, although they may be very fundamental for intelligent systems. For example, the scientific method is increasingly data-driven, which poses new problems. Many (but not all) risks can be decreased by raising awareness and by implementing regulatory procedures and standards (e.g., the US Food and Drug Administration and the Environmental Protection Agency already provide regulatory frameworks for synthetic biology). We paid attention to the ethical dimension related to the pervasiveness of intelligent systems and the various impacts these systems may have on modern information-centered society. This somewhat *softer* dimension (as opposed to rigorous mathematical definitions and clearcut technical processes and standards) is highly relevant for the well-being of human beings in a modern information-driven society. Overall, it is with intelligent systems as it is with many discoveries and inventions – their glory as well as their perfidy lie in the integrity of the human beings behind them.

References:

[1] A. Kent, "Too damned quiet?" Available at: http://arxiv.org/abs/1104.0624, accessed May 16, 2011.
[2] A. Schuster and Y. Yamaguchi, "From foundational issues in artificial intelligence to intelligent memristive nano-devices," Int. J. Mach. Learn. & Cyber., DOI 10.1007/s13042-011-0016-1, 2011.
[3] M. Schmidt and H. Lipson, "Distilling free-form natural laws from experimental data," Science, Vol.324, No.5923, pp. 81-85, 2009.
[4] D. Berrar, M. Granzow, and W. Dubitzky, "Introduction to genomic and proteomic data analysis," In W. Dubitzky, M. Granzow, and D. Berrar (Eds.), Fundamentals of Data Mining in Genomics and Proteomics, Springer, pp. 1-37, 2007.
[5] D. F. Ransohoff, "Bias as a threat to the validity of cancer molecular-marker research," Nature Reviews Cancer, Vol.5, No.2, pp. 142-149, 2005.
[6] J. Cohen, "Chance, skill, and luck: The psychology of guessing and gambling," Baltimore, MD: Penguin Books, 1960.
[7] S. Goodman, "A dirty dozen: Twelve *p*-value misconceptions," Seminars in Hematology, Vol.45, No.3, pp. 135-140, 2008.
[8] S. Goodman, "*p* values, hypothesis tests, and likelihood: Implications for epidemiology of a neglected historical debate," American J. of Epidemiology, Vol.137, No.5, pp. 485-496, 1993.
[9] J. P. A. Ioannidis, "Why most published research findings are false," PLoS Medicine, Vol.2, No.8, p. e124, 2005.

[10] D. H. Johnson, "The insignificance of statistical significance testing," J. of Wildlife Management, Vol.63, No.3, pp. 763-772, 1999.
[11] D. N. McClosky and S. T. Ziliak, "The unreasonable ineffectiveness of Fisherian "tests" in biology, and especially in medicine," Biological Theory, Vol.4, No.1, pp. 44-53, 2009.
[12] C. Poole, "Beyond the confidence interval," American J. of Public Health, Vol.77, No.2, pp. 195-199, 1987.
[13] A. Isaksson, M. Wallman, H. Gransson, and M. G. Gustafsson, "Cross-validation and bootstrapping are unreliable in small sample classification," Pattern Recognition Letters, Vol.29, No.14, pp. 1960-1965, 2008.
[14] D. Hambling, "Surveillance robots know when to hide," New Scientist, Vol.209, No.2804, p. 25, March 2011.
[15] A. R. Wagner and R. C. Arkin, "Acting deceptively: providing robots with the capacity for deception," Int. J. of Social Robotics, Vol.3, No.1, pp. 5-26, 2011.
[16] C. Heeney, N. Hawkins, J. de Vries, P. Boddington, and J. Kayeb, "Assessing the privacy risks of data sharing in genomics," Public Health Genomics, Vol.14, No.1, pp. 17-25, 2010.
[17] J. Du and M. Gerstein, "Genomic anonymity: Have we already lost it?," The American J. of Bioethics, Vol.8, No.10, pp. 71-81, 2010.
[18] A. Schuster (Ed.), "Robust intelligent systems," Springer-Verlag, London, 2008.
[19] P. Vytelingum, T. D. Voice, S. D. Ramchurn, A. Rogers, and N. R. Jennings, "Agent-based micro-storage management for the smart grid," Proc. 9th Int. Conf. on Autonomous Agents and Multiagent Systems, pp. 39-46, 2010.
[20] J. Bohannon, "News Focus: IEEE International Conference On Computational Intelligence And Games, Game-miners grapple with massive data," Science, Vol.330, No.6000, pp. 30-31, 2010.
[21] I. S. Kotchetkov, B. Y. Hwang, G. Appelboom, C. P. Kellner, E. S. Connolly, Jr., "Brain-computer interfaces: Military, neurosurgical, and ethical perspective," Neurosurg Focus, Vol.28, No.5, 2010.
[22] P. Lin, K. Abney, and G. Bekey, "Robot ethics: Mapping the issues for a mechanized world," Artificial Intelligence, doi:10.1016/j.artint.2010.11.026, 2011.
[23] J. Giles, "Fake tweets by socialbot fool hundreds of followers," New Scientist, Vol.209, No.2804, p. 28, March 2011.
[24] Y. Benenson, T. Paz-Elizur, R. Adar, E. Keinan, Z. Livneh, and E. Shapiro, "Programmable and autonomous computing machine made of biomolecules," Nature, Vol.414, No.6862, pp. 430-434, 2001.
[25] Y. Benenson, B. Gil, U. Ben-Dor, R. Adar, and E. Shapiro, "An autonomous molecular computer for logical control of gene expression," Nature, Vol.429, No.6990, pp. 423-429, 2004.
[26] A. Tamsir, J. J. Tabor, and C. A. Voigt "Robust multicellular computing using genetically encoded NOR gates and chemical 'wires'," Nature, Vol.469, pp. 212-215, 2010.
[27] D. G. Gibson, J. I. Glass, C. Lartigue et al., "Creation of a bacterial cell controlled by a chemically synthesized genome," Science, Vol.329, No.5987, pp. 52-56, 2010.
[28] A. Levskaya, A. A. Chevalier, J. J. Tabor et al., "Synthetic biology: engineering Escherichia coli to see light," Nature, Vol.438, No.7067, pp. 441-442, 2005.
[29] R. D. Berg, "The indigenous gastrointestinal microflora," Trends in Microbiology, Vol.4, No.11, pp. 430-435, 1996.
[30] M. B. Miller and B. L. Bassler, "Quorum sensing in bacteria," Annual Review of Microbiology, Vol.55, pp. 165-199, 2001.
[31] H. Breithaupt, "The engineer's approach to biology," EMBO Rep., Vol.7, No.1, pp. 21-23, 2006.
[32] A. M. Turing, "Intelligent machinery, a heretical theory," Philosophia Mathematica, Vol.4, No.3, pp. 256-260, 1996.
[33] A. M. Turing, "The chemical basis of morphogenesis," Philosophical Trans. of the Royal Society of London B, Vol.237, No.641, pp. 37-72, 1952.
[34] D. Berrar, N. Sato, and A. Schuster, "Quo vadis, artificial intelligence?," Advances in Artificial Intelligence, doi:10.1155/2010/629869, 2010.
[35] C. A. Lowery, T. J. Dickerson, and K. D. Janda, "Interspecies and interkingdom communication mediated by bacterial quorum sensing," Chem Soc Rev. Vol.37, No.7, pp. 1337-1346, 2008.
[36] J. Cello, A. V. Paul, and E. Wimmer, "Chemical synthesis of poliovirus cDNA: generation of infectious virus in the absence of natural template," Science, Vol.297, pp. 1016-1018, 2002.

Name:
Alfons Schuster

Affiliation:
Laboratory for Dynamics of Emergent Intelligence, RIKEN Brain Science Institute

Address:
Wako-shi, Saitama 351-0198, Japan

Brief Biographical History:
1990 Bachelor of Science (BSc) in Applied Physics
1999 Doctor of Philosophy in Computer Science

Main Works:
• A. Schuster (Ed.), "Robust Intelligent Systems," Springer Verlag, London, 2008.
• A. Schuster (Ed.), "Intelligent Computing Everywhere," Springer Verlag, London, 2007.
• artificial intelligence, computing, and information study

Name:
Daniel Berrar

Affiliation:
Interdisciplinary Graduate School of Science and Engineering, Tokyo Institute of Technology

Address:
G3-45, 4259 Nagatsuta, Midori-ku, Yokohama 226-8502, Japan

Brief Biographical History:
1999 Master of Science in Medical Informatics
2004 Doctor of Philosophy (machine learning, bioinformatics)

Main Works:
• W. Dubitzky, M. Granzow, and D. Berrar, "Fundamentals of Data Mining in Genomics and Proteomics," Springer, Heidelberg, 2007.
• D. Berrar and P. Flach, "Caveats and pitfalls of ROC analysis in clinical microarray analysis (and how to avoid them)," Briefings in Bioinformatics, doi: 10.1093/bib/bbr008, 2011.
• artificial intelligence and machine learning