

AI's

10



to Watch


As we enter the second 50 years of AI's history, it's clear that many deep and challenging puzzles still remain before we can realize AI's full promise. Indeed, we don't really understand yet whether AI systems will ultimately be based on mathematics or biological models, learned or engineered. Thus, AI's future depends critically on its ability to attract the brightest new researchers—young people with broad knowledge and insight, great creativity, prodigious systems-building abilities, management skills, and the charisma to inspire generations of future graduate students. We envisioned the *IEEE Intelligent Systems 10 to Watch* award as a way to recognize new researchers' promise and early accomplishments, as well as to inspire those pursuing AI research in graduate school or still deciding whether to choose AI as their research specialty.

We didn't know what to expect in response to the call for nominations, but we entered into this hoping for a respectable response. As it turns out, we needn't have worried. We received more than 50 nominations, and they're a most impressive group. The selection committee was truly inspired by the quality, accomplishments, diversity, and depth of this remarkable group of young men and women. It wasn't easy to narrow the list to only 10, but we think that you too will be impressed by the winners. They represent many subspecialties, many approaches, and many countries. Our hearty congratulations to all!

On a final, sad note: one of our winners, Push Singh, died shortly after being selected. He is memorialized by Jim Hendler separately in this issue, on page 15.

—David L. Waltz, Center for Computational Learning Systems, Columbia University





Eyal Amir

University of Illinois at Urbana-Champaign

Eyal Amir is an assistant professor of computer science at the University of Illinois at Urbana-Champaign. Prior to joining the UI staff in January 2004, he was a postdoctoral researcher with Stuart Russell at the University of California, Berkeley. Amir's research focuses on building systems that reason, learn, and make decisions with logical and probabilistic knowledge. Applications of his research include the World Wide Web, adventure games, formal verification of circuits and programs, and controlling mobile robots. His research goal is human-level AI. Amir received the US National Science Foundation's Early Career Development award (2006) and Stanford University's Arthur L. Samuel Award for best computer science PhD thesis (2001–2002). His PhD was with John McCarthy and focused on logical AI. He received his BSc and MSc in mathematics and computer science from Bar-Ilan University, Israel. When he's away from his office, Amir likes to dance, practice Shotokan karate, and practice Olympic-style rifle shooting. The latter has won him multiple Israeli national medals. Contact him at eyal@cs.uiuc.edu.

Steps on the Way to Human-Level AI

Three challenges to achieving human-level AI are merging knowledge representation with machine learning, scaling up reasoning with logical and probabilistic knowledge to real-world-size problems, and developing a theory of human-level AI.

Addressing these challenges, my research focuses on the ability to represent, learn, and

reason about objects, relationships, and personal knowledge. These are central to AI—for example, natural-language processing, sequential decision making, and diagnosis problem solving. Such applications must consider myriad objects and relations, and real-world AI applications require representation and reasoning mechanisms that can scale to thousands of objects and relations and more.

With my coauthors (who have my deep gratitude), I discovered and am using two key mathematical tools. The first is a graph-based structure based on Craig's Interpolation Theorem in relational domains (equivalent to tree-width in propositional domains). I speed up inference in relational representations using this structure. The second tool is a family of logical representations that enables efficient tracking of knowledge and beliefs in dynamic domains. Such knowledge representations enable efficient learning in partially observable domains. I have used these tools to develop algorithms for factored planning;

behavior-based robot control architectures; and algorithms for reasoning, learning, and acting with explicit knowledge in large, partially observable domains.

In the near future, I expect AI to focus on its now-visible strengths: knowledge in partially observable domains and connecting the combinatorial (logical) and analytical (probabilistic) in solutions to problems. My work and my colleagues' shows how partitioning and relational inference methods apply to systems that connect logical, probabilistic, and behavioral knowledge.

In the more distant future, I expect a shift in AI toward massive data sets of distributed knowledge and resources. This will transform AI into a science that *uses* large scale instead of just grappling with it. For example, a program would sift through masses of knowledge and choose the pieces that best match its situation. Although it might not be provably correct, it would work well in practice. The advent of WWW search methodologies and

data mining's evolution from machine learning are examples of this transformation.

I plan to develop a methodology for inference in first-order logic and probabilistic representations that can ignore most interactions between objects, functions, and predicates and also be fast and correct (for an example, see www.cs.uiuc.edu/~eyal/compact-prop). My research will enable large, autonomous learning and expansion of knowledge by exploration. Autonomous agents would address unfamiliar situations by choosing knowledge from the vast amount available and would refine this knowledge by purposeful exploration.

Developing new foundations and a sound AI theory is the longer-term challenge in AI. A better theory of AI would help researchers build steps to human-level AI. In that vein, one of my goals in the next few years is to develop a theory of AI-completeness that would distinguish between humans and machines using an AI-complete problem. There, I wish to create a concept akin to NP-completeness that would serve to focus theoretical research in AI.

AI and Linguistics

My research focuses on robust and efficient computational modeling of complex linguistic phenomena in pragmatics, discourse, and lexical semantics. Linguists and computational scientists have studied these phenomena extensively for decades—they're at the core of language processing. To date, rule-based approaches have predominated

in modeling discourse and pragmatics. However, these models are hard to incorporate in modern systems: they're valid only for limited domains, with no guarantee of scalability or portability.

My research aims to develop models that combine the robustness of probabilistic techniques and the richness of representations proposed in linguistic theories. The underlying phenomena's complexity rules out using out-of-the-box approaches and requires novel model formulations to make them amenable to statistical analysis. To enable accurate and efficient learning and inference, I employ powerful algorithmic tools ranging from optimization approaches for modeling distributional properties of coherent texts to graph-theoretic methods for capturing context dependences in content selection.

My current work focuses on two fundamental, orthogonal dimensions of text—

content and coherence. Content models characterize text structure in terms of the topics addressed and the order in which such topics appear. These models can specify, for example, that articles about earthquakes typically contain information about quake strength, location, and casualties and that descriptions of casualties usually precede those of rescue efforts. But rather than manually determine a given domain's topics, we take a distributional view, learning them directly from unannotated texts via analysis of word-distribution patterns. Experiments show that automatically derived content models yield significant performance gains in text summarization and information ordering.

Although content models are domain dependent, coherence models aim to capture the properties that make well-written texts easier to read and understand than a sequence of randomly concatenated sen-

tences. These models' goal is to capture text relatedness at the level of sentence-to-sentence transitions. Our work's key premise is that distribution of entities in locally coherent texts exhibits certain regularities that can be induced automatically from raw texts. The coherence models operate over an automatically computed representation that reflects distributional, syntactic, and referential information about discourse entities. Our model learns mapping between transitional patterns and the degree of text coherence, thereby making it possible to automatically assess text quality. In fact, the rankings our coherence model produces show significant agreement with human coherence judgments elicited for automatically generated texts.

My long-term goal is to continue developing methods with sound probabilistic and linguistic foundations to obtain robust and powerful natural language processing systems.

AI's 10 to Watch

Regina Barzilay

Massachusetts Institute of Technology

Regina Barzilay is an assistant professor in the Department of Electrical Engineering and Computer Science and a member of the Computer Science and Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. She received her BSc in mathematics and MSc in computer science from the University of Ben Gurion in the Negev, Israel, and her PhD in computer science from Columbia University. Following her PhD, she was a postdoctorate associate at Cornell University. Her research interests include statistical text generation, discourse modeling, paraphrasing, and summarization. She serves on the editorial board of *Computational Linguistics* and is a member of the executive committee of the North American Association of Computational Linguistics (NAACL). Her honors include a US National Science Foundation Early Career Development award (2004) and a Human Language Technology Conference (HLT-NAACL 04) Best Paper award. In 2005, Barzilay was named one of *Technology Review's* TR35 for being a top young innovator of the 21st century. Contact her at regina@csail.mit.edu.





Jennifer Golbeck

**Joint Institute for Knowledge Discovery,
University of Maryland**

Jennifer Golbeck is a postdoctoral researcher and research coordinator of the newly founded Joint Institute for Knowledge Discovery at the University of Maryland, College Park, in the University of Maryland Institute for Advanced Computer Studies. She received her AB in economics and her SB and SM in computer science from the University of Chicago. Golbeck received her PhD in computer science from the University of Maryland in 2005. Her dissertation addressed computing and applying trust in Web-based social networks, and she's generally interested in how to use social-network analysis with online communities and communication networks to create intelligent applications. In 2005, Golbeck was selected as a DARPA IPTO Young Investigator and received an American Association for University Women Graduate Award. When she's away from her computer, she runs marathons and spends time with her two golden retrievers, π and K. Contact her at golbeck@cs.umd.edu.

AI and Social Networks

I research the dynamics of social networks found in online communities and email networks. I believe that we can analyze these networks to compute useful data about each user's social environment and that we can use the result to develop intelligent user interfaces and inform an understanding of communication patterns.

Tens of millions of people participate in Web-based social networks, and more than 140 Web sites have more than 200 million user accounts among them. Because the data is on the Web, it's publicly accessible—particularly with the growth of Semantic Web technologies for representing social networks. I'm also looking at networks built from the Enron Email Corpus, a public collection of mailboxes from 150 Enron executives comprising over 500,000 messages and 20,000 unique users.

With these data sets, I'm continuing work I began with my dissertation, investigating trust in social networks. As part of that project, I developed algorithms for computing personalized inferences of trust relationships between individuals in the network, based on trust values on the paths that connect them. I have used those algorithms to create systems that incorporate users' social information. An ongoing project,

FilmTrust, uses trust values to generate predictive movie ratings and determine the order in which movie reviews are presented (<http://trust.mindswap.org/FilmTrust>). I'm currently extending my analysis to develop a model of what factors people use in assigning trust. I'm also working on several other applications of trust, including filtering open source intelligence and prioritizing default rules for nonmonotonic reasoning.

I'm extending my work to understanding social networks' temporal dynamics. First, I'm interested in how real-world events affect the connection patterns visible in a network. Particularly, I'm interested in tracking whether identifiable clusters exhibit changes in activities at identifiable times. If we can recognize these clusters and correlate them with events or stages of projects, we can use their communication patterns predictively. For example, if a cluster is known to be

active in the early stages of certain project types, detecting familiar activity in real time could signal the beginning of a new project. If we can do this reliably, it's applicable in a number of spaces, including intelligence analysis.

I'm also looking at how network centrality correlates to temporal statistics for users. I'm testing how an individual's centrality changes over time as well as how an individual's centrality correlates to the person's time in the network or since his or her last activity. Certain centrality measures indicate influence in the network and thus suggest which members control communication or access to other members. Understanding how these factors are manifested might let us engineer social networks to optimize participation or access to information.

Computing for everyday users has largely become a social, Internet-based activity. I believe an important direction for AI is one that explores what it means to be intelligent with respect to users' social contexts.

AI and Human Cognition

Traditionally, AI research has drawn inspiration from human cognition, seeking to produce similarly intelligent behavior in artificial systems. In my research, I turn this relationship around, drawing inspiration from the successes of AI, machine learning, and statistics and using ideas from these disciplines to better understand how human

minds work. Because human beings are still the standard to which artificial systems are compared in many contexts, from learning language to scientific discovery, this approach also has the potential to provide new insights that can lead to advances in AI.

My main interest is in understanding how people make inductive inferences, reasoning to underdetermined conclusions from limited evidence. Despite the notorious difficulty of such inferences, people make them successfully every day, learning the meaning of words, recognizing new causal relationships, and making predictions about future events with ease. I try to identify the computational problems that underlie these inferences, think about these problems' optimal solutions, and examine how well the optimal solutions correspond to human behavior. Many inductive problems can be formulated as

problems of statistical inference, so I often end up evaluating the connection between statistical methods developed in AI and machine learning research and people's judgments. Thinking about tasks that people seem to be able to perform so effortlessly is a rich source of problems for which good formal analyses and algorithms have yet to be developed. Because of this, I spend a lot of time exploring new statistical methods in their own right.

At present, my research efforts focus on two inductive problems: learning causal relationships and learning language. Recent work in AI on causal graphical models provides a great foundation for asking questions about human causal learning and makes it possible to define models that make remarkably accurate quantitative predictions about people's judgments. Learning language is a perennial problem in both AI and cognitive science, and my

emphasis has been on developing statistical models that capture some of the structure of human languages. In both enterprises, methods from nonparametric Bayesian statistics have enabled me and my colleagues to define statistical models that increase in complexity as more data becomes available—something necessary to capture the scope and flexibility of human cognition.

My great hope for the future is to see a closer integration between the studies of artificial and natural intelligence. Psychologists were present at the birth of AI, and computer scientists were present at the birth of cognitive science. Over the last 50 years, the two disciplines have drifted in and out of contact with one another. I see the rigorous formal methods developed in AI research as intensely valuable to understanding human cognition and human cognition as a guide to solving some of AI's hardest problems. A closer integration between the two disciplines strikes me as the best way to develop a deeper understanding of intelligent systems.


AI's 10 to Watch

Tom Griffiths

University of California, Berkeley

Tom Griffiths is an assistant professor of psychology and cognitive science at the University of California, Berkeley. He received his BA in psychology from the University of Western Australia and his MS in statistics and PhD in psychology from Stanford University (the latter in 2005). Griffiths was a faculty member at Brown University's Department of Cognitive and Linguistic Sciences before moving to Berkeley. His research focuses on the development of computational models of human cognition with an emphasis on using ideas from Bayesian statistics and machine learning. In 2003, he won prizes for best student paper in both the natural-systems and synthetic-systems categories at the Neural Information Processing Systems Conference. Having been a competitive fencer for several years, he struggles to integrate his work life with his enthusiasm for the epee. Contact him at tom_griffiths@brown.edu.





Steven Gustafson

General Electric Global Research Center

Steven Gustafson is a computer scientist at the General Electric Global Research Center in Niskayuna, New York. As a member of the Computational Intelligence Lab, he develops and applies advanced AI and machine learning algorithms for complex problem solving. He received his PhD in computer science from the University of Nottingham, UK, where he was a research fellow in the Automated Scheduling, Optimisation and Planning Research Group. He received his BS and MS in computer science from Kansas State University, where he was a research assistant in the Knowledge Discovery in Databases Laboratory. His PhD dissertation, an analysis of a biologically inspired search algorithm in the space of computer programs, was nominated for the British Computer Society and the Conference of Professors and Heads of Computing Distinguished Dissertation award, which recognizes the top PhD thesis in the UK computer science community. In 2005 and 2006, he coauthored papers that won the Best Paper Award at the European Conference on Genetic Programming. Outside of work, Gustafson enjoys literature and traveling with his wife and infant son. Contact him at steven.gustafson@research.ge.com.

Creative Problem Solving with Genetic Programming

Genetic programming is a heuristic search method that uses a population of variable-length computer programs and a search strategy based on biological evolution. It represents an intuitive method for automatically evolving programs. (For examples of genetic programming, see the latest editions of IEEE publications.)

Several grand challenges exist in genetic programming, one of which helps illustrate an advantage of the method. By representing solutions with programs, practitioners specify the solution primitives in familiar algorithmic building blocks. These primitives are usually similar to, or the same as, those in most high-level programming languages. Actually constructing programs from these primitives (including the program's topological structure and content) is left to the transformation operators and evolutionary process in the genetic programming search strategy. However, a very complex problem arises: how do we transform solutions represented by programs? Although a high-level language such as Java is amenable to human programming, it's extremely sensitive to small changes in syntax and semantics. In genetic programming, solutions represented by programs are very sensitive to transformation operators. Thus, the appeal of specifying a solution as a program can be offset by the complexity of

defining representation-sensitive transformation operators. In genetic programming, the difficulty in designing well-behaving operators is typically avoided by using random-transformation operators, carrying out a computationally complex search process, and tolerating overly complex solutions.

My PhD research addressed the need to understand the dynamics of genetic programming that encourage efficient, effective search. I focused on a critical property of genetic programming search: the population. This relates to many aspects of the genetic programming algorithm. Diversity was used to describe and analyze populations and their effect on search. The research led to several informative measures of diversity, useful for controlling and predicting the outcome of search, and algorithm enhancements that help explain the dynamics of genetic programming search on new problems. My future work addresses genetic programming's potential to leverage knowledge from other

fields (such as programming languages and software engineering) to improve search (for example, to design better-behaving transformation operators and representations) and evolve better programs (that is, programs that are maintainable, extendable, and self-reconfigurable and that better use the vast array of existing software libraries).

My current work aims for innovation in science and algorithm research via real-world problem solving. At the GE Global Research Center, I research advanced learning and search algorithms for solving a wide range of data mining, optimization, and learning problems. Working in the scope of solving real-world problems requires balancing long-term research agendas with measurable short-term results. This creates the need to address often-overlooked research problems, such as the inefficiency of genetic programming transformation operators and solutions. Applying adventurous AI methods such as genetic programming to real-world problems has the potential to create new, innovative solutions as well as provide critical advancements to the method.

AI and Description Logics

My research is concerned with the logical foundations of computer science and AI. In AI, I have a particular interest in knowledge representation for two reasons: it's a central issue in almost every AI application and, accordingly, located at the very heart of AI; and it crucially depends on formally well-understood languages and

automated reasoning techniques. This indicates the high benefit of developing a logical underpinning of the field.

My research in knowledge representation emphasizes description logics, a popular family of logic-based knowledge representation formalisms. DLs' main applications are providing a formal description of an application domain's relevant notions and reasoning about these notions. For example, DLs have been successfully used to systematize medical terms in the SNOMED-CT terminology, which is widely used in US health care. Recently, DLs have experienced tremendous interest because of their use as ontology languages and the W3C's standardization of OWL-DL as the ontology language of the Web.

In description logic, the connection between logic and knowledge representation

has blossomed. In particular, there is a tight connection between theoretical research into DLs' logical foundations and the highly efficient DL reasoners many applications use today. Theoretical results map out DLs' landscape and interrelate expressive power and computational complexity, and implementation work provides important feedback about which DL features are useful and feasible in practice. This allows the achievement of the careful balance between expressive power and computational complexity that characterizes modern DLs.

In my own DLs research, I have, for example, contributed to the theoretical understanding of DLs that lets us refer to numerical data in terminologies, to finite-model reasoning in DLs, and to identifying DLs with tractable reasoning prob-

lems. The latter are currently being implemented in the reasoning system CEL, and the first evaluations regarding performance and usability are very promising. I've also worked on understanding the connection between DLs and other AI subfields, linking them to nonmonotonic logic and reasoning about action.


Concerning DLs' future, I believe that their use as an ontology language opens up many opportunities but also poses serious new research challenges. For example, integrating and interoperating multiple ontologies is crucial in many applications. I believe that if DLs are properly developed as ontology languages with an eye on both theory and practice, they might be able to come out of the AI experimentation lab to mainstream computer science.


AI's 10 to Watch

Carsten Lutz

Institute of Theoretical Computer Science

Carsten Lutz is a postdoctoral researcher at the Institute of Theoretical Computer Science of Dresden University of Technology in Germany. He received his diploma (comparable to MA) in computer science from the University of Hamburg and his PhD in computer science from RWTH Aachen University in 2002. His PhD thesis developed and analyzed a family of description logics that allow the representation of abstract conceptual knowledge and concrete numerical data in an integrated way. His current research interests revolve around modal and temporal logic's use in computer science and AI, including description logics for knowledge representation, temporal and dynamic logics for hardware and software systems verification, modal logics of space, and logics for multiagent systems. In 2006, he submitted his habilitation thesis to Dresden University of Technology. In his spare time, he can usually be found exploring the beautiful surroundings of Dresden. Contact him at lutz@tcs.inf.tu-dresden.de.





Pragnesh Jay Modi

Drexel University

Pragnesh Jay Modi is an assistant professor in Drexel University's Department of Computer Science. One of the first in his family to graduate from college, he received his PhD in computer science from the University of Southern California (2003) and his BS with honors in computer science and math from Carnegie Mellon University. His dissertation, "Distributed Constraint Optimization for Multiagent Systems," made several contributions to distributed constraint reasoning. Modi was chair of the Americas School on Agents and Multiagent Systems, chair of the Distributed Constraint Reasoning workshop, and a program committee member for the Autonomous Agents and Multiagent Systems Conference and the National Conference on Artificial Intelligence. He has reviewed for *Artificial Intelligence*, *Journal of Artificial Intelligence Research*, and *Journal of Autonomous Agents and Multiagent Systems*. He's interested in developing representations and reasoning algorithms for automated and semiautomated coordination of agent decisions and activities. Contact him at pmodi@cs.drexel.edu.

AI and Multiagent Coordination

Complex real-world tasks such as disaster rescue, security patrolling and policing, and large-event crowd control are extremely difficult in part because they require humans to coordinate an overwhelming number of decisions and actions in real time under stressful dynamic conditions. One potential application of my research is the development of intelligent

systems that can facilitate information sharing, conflict detection and resolution, collaborative planning, and other intelligent functions to substantially assist humans in performing complex coordination tasks.

From its inception, the field of AI has focused on how agents can make intelligent decisions, and I believe a key manifestation of such intelligence is the ability to coordinate effectively with others. While AI research sees several approaches to multiagent coordination (such as economic game theory and distributed Markov models), I'm particularly interested in developing and applying the Distributed Constraint Optimization Problem as a key paradigm for multiagent coordination. I believe DCOP is a promising approach because it leverages existing research in constraint-based representations that have a long history of proven success in AI.

In my thesis research, I developed the first algorithm for DCOP, named Adopt (Asynchronous Distributed Optimization), where

communication is completely asynchronous and guaranteed to obtain globally optimal solutions (under reliable communication conditions). Asynchronous communication lets agents make decisions in parallel, substantially decreasing the time it takes to find a globally optimal solution. Another interesting feature of Adopt is that it admits bounded-error approximation. Bounded-error approximation is useful when time and communication resources are limited and agents must quickly find a solution even if it's not optimal. A key feature of this approximation in Adopt is that provable quality guarantees are still available.

Looking forward, we can see new computing paradigms for intelligent agents on the horizon. Computing devices are becoming more portable and powerful (such as handhelds and cell phones) and network technologies are more effectively supporting wireless data, audio, and video traffic (such as ad hoc, Wi-Fi, and mobile networks). Unfortunately,

many of today's multiagent coordination algorithms weren't explicitly designed with such computing environments in mind, and significant challenges remain. Here are three open questions I consider crucial:

- How can we design distributed agent systems that solve difficult coordination problems with limited time or communication (for example, when communication is unreliable, expensive, or risky, or when privacy restrictions prevent information exchange)?
- How can we design distributed agent systems that are open (that is, design algorithms that can cope with fluidity), where agents come and go from the system?
- How can we design distributed agent systems that work in dynamic environments (that is, where communicated information can quickly become obsolete)?

These research challenges go to the heart of interaction between intelligent entities. If we want computers to help us with complex distributed real-world tasks, we must address these challenges as a central priority of AI research.

Harnessing AI Techniques for the Semantic Web

I believe that the best research brings together theory and practical needs. The Semantic Web context offers a perfect environment for applying theoretical methods developed in AI to the practical problems of the Web. My work has long been concerned with exploring knowledge acquisition and representation techniques for improving application

development and the user experience on the Web.

The practical motivation behind my thesis topic was the rapidly developing Web services area: as Web services have increased in number and become more critical to Web applications, the need to develop methods and tools that support (semi)automatic service discovery and composition has grown. Research on Semantic Web services approaches these problems by developing ways to produce formal service specifications that can be used to automate typical Web service tasks. Some major issues, which my thesis addresses, relate to how these descriptions should be structured and whether they can be (semi)automatically generated. By using Semantic Web services technology in several projects, I derived a set of concrete knowledge representation suggestions regarding Semantic Web service descriptions' content and structure. The OWL-S committee—the international standardization body in this

area—used these. Also, my research has shown that it is possible to automatically acquire a significant part of the domain knowledge needed to describe Web services from sources such as their textual descriptions or underlying tools' APIs. These algorithms' evaluation in the practical context of bioinformatics Web services has shown that, on average, half of the needed knowledge can be acquired automatically.

In recent years, the Semantic Web community has made a conscious effort to migrate its technologies into the real (that is, large, heterogeneous, distributed, dynamically changing, and open) Web environment, thus gradually dropping the simplifying assumptions inherent in much Semantic Web research. I believe that this tendency brings ample opportunities to further explore the synergy between AI and the Web.

For example, I'm already adapting my doctoral work to enhance online Web service repositories. For this, I primarily rely on knowledge

acquisition and presentation techniques developed by the Semantic Web community. These are currently poorly structured due to the difficulty of acquiring metadata about their large, dynamically changing service collections.


I'm also interested in investigating the challenges raised by extending Semantic Web tools. Most current tools rely on a single ontology and therefore function only within the domain this ontology defines. An ontology-based question-answering system, for example, is constrained to answer questions phrased using concepts defined by the ontology that it incorporates. Obviously, to fully benefit from the Semantic Web, the new generation of tools should dynamically select and combine appropriate semantic data directly from the Web (thus automatically adapting themselves to the user's context and information need). Therefore, a need exists for efficient, robust techniques to perform automatic ontology selection and evaluation, identify relevant ontology modules, and dynamically combine them according to the application's needs. These issues are fundamental to establishing a new, real Semantic Web.

AI's 10 to Watch

Marta Sabou

Knowledge Media Institute

Marta Sabou is a research fellow at the Knowledge Media Institute of the Open University, UK. She received a bachelor's degree in system engineering from the Technical University of Cluj-Napoca, Romania, and a master's degree in AI from the Vrije Universiteit, Amsterdam. In 2005, she completed her PhD at the Vrije Universiteit's Knowledge Representation Group on the topic of enhancing and (semi)automatically learning ontologies to specify the semantic description of Web services. Much of her doctoral work was carried out in the context of major European research projects, such as WonderWeb, KnowledgeWeb, and SWAP. Her research also involved significant cooperation with the UK-funded myGrid project and the OWL-S standardization committee and was published at the WWW and International Semantic Web Conferences and in the *Journal of Web Semantics*. Currently she focuses on developing ontology evaluation and selection techniques that can be applied in the new environment provided by the emerging Semantic Web. This topic is part of the research agenda of two European projects, Open Knowledge and NeOn. In her spare time, she enjoys swimming, working out, and learning foreign languages (currently Dutch and Italian). Contact her at r.m.sabou@open.ac.uk.





Richard A. Watson

University of Southampton

Richard A. Watson is a senior lecturer in the natural systems research group at the University of Southampton's School of Electronics and Computer Science. He received his BA in AI from the University of Sussex in 1990 and then worked in industry for about five years. Returning to academia, he chose Sussex again for an MSc in knowledge-based systems, where he was introduced to evolutionary modeling. His PhD in computer science at Brandeis University (2002) resulted in 22 publications and a dissertation addressing the algorithmic concepts underlying the major transitions in evolution. A postdoctoral position at Harvard University's Department of Organismic and Evolutionary Biology provided training to complement his computer science background. He now has over 35 journal and conference publications on topics spanning artificial life, robotics, evolutionary computation, and computational biology. At Southampton, he's establishing a new research group and leading preparation of a new MSc in complexity science. He is the author of *Compositional Evolution: The Impact of Sex, Symbiosis, and Modularity on the Gradualist Framework of Evolution* (MIT Press, 2006). Contact him at raw@ecs.soton.ac.uk.

AI and Algorithmic Biology

I've taken to calling my research field *algorithmic biology*—the use of algorithmics and complexity theory to understand the scientific principles that underlie both computational methods and biological systems. AI has always been interdisciplinary, with one foot in cognitive science and the other in computer science. Nouvelle AI, as some have called

it, has goals (among others) in animal behavior and robotics. And Artificial Life both asks deep questions about the nature of living things and perhaps promises new techniques of self-organization, self-repair, and self-reproduction in engineered systems. However, my favorite algorithmic process in biology isn't any one of these but rather the process that has driven things from one end of this spectrum to the other: from simple, self-replicating molecules, through numerous scales of organization, to the vast space of complex systems that includes intelligent animals like us.

Darwin provided the first algorithmic account of how biological complexity has arisen by showing how heritable variation and differential reproduction interact to produce evolution by natural selection. As computer scientists, we recognize this process as a means for biology to implement a form of population-based stochastic local search or hill-climbing. Hill-climbing is a simple, easy-

to-understand algorithm that has become deeply ingrained in evolutionary thinking in the last 150 years. But is such a trivial algorithm sufficient to derive all this complexity? As computer scientists, we know that hill-climbing has limitations.

Recently, there's been increasing awareness of some previously underresearched biological phenomena, including lateral gene transfer, endosymbiosis, and mechanisms involved in the major evolutionary transitions. Although such mechanisms act within natural selection, they present a problem for the hill-climbing model of adaptation. Exchange and fusion of genetic material between populations potentially allows innovation that neither ancestral population could achieve via random modifications within its own local neighborhood of genotypes. However, adherence to a simplified algorithmic paradigm has caused such phenomena to be treated as curios and squeezed into the hill-climbing framework as just another form of

variation. But the space of algorithmic processes includes more than just hill-climbing, and biology is pretty open-minded about using what works. In fact, we can better understand the biological phenomena mentioned earlier as a form of a bottom-up divide-and-conquer process, and this means that they can evolve systems that are impossible within the hill-climbing paradigm.

This observation's significance for both biology and evolutionary computation is scalability—in particular, the potential for adaptation to form assemblages of simpler units into aggregate units through several scales of hierarchical complexity. Personally, I think much is still missing in our understanding of the algorithmic processes in nature that couple microevolution and macroevolution. Likewise, there is a lot missing in our ability to make automatic design and optimization methods that scale up to provide really sophisticated solutions worthy of comparison with hand-designed systems, let alone biological complexity. But research in both these arenas can benefit greatly by learning from each other. ■

Push Singh (1972–2006)

James Hendler, *University of Maryland*

It was with great sadness that I learned of Push Singh's death on 28 February this year. This remembrance is in lieu of the essay we expected to publish for his "AI Ten to Watch" award. It would be hard to write even if I knew

him only through the paperwork submitted with his nomination. In Push's case, however, it's much harder, for I had recently begun talking with him about research ideas, and he had some exciting thoughts that I was eager to pursue during an upcoming sabbatical.

Push started at the Massachusetts Institute of Technology as an undergraduate in 1988, staying on as a graduate student and completing his thesis in MIT's Department of Electrical Engineering and Computer Science in 2005. His advisor was Marvin Minsky, one of AI's founders. Push was slated to begin a position as a faculty member in the MIT Media Laboratory in 2007 after, as he described it to me, a much-needed year off "to think."

Push's research was based partly on Marvin's society-of-minds approach, exploring what common sense was and how it could develop. His thesis, "EM-ONE: An Architecture for Reflective Commonsense Thinking," included the creation of Open Mind Common Sense (<http://csc.media.mit.edu>), a Web-based approach to acquire commonsense knowledge from the general public. He also developed a layered cognitive architecture aimed at exploring how reasoning about physical, social, and mental domains could work.

In remembering him, Marvin Minsky said, "Push had just been appointed to become a new professor at MIT, to pursue what we knew would be a brilliantly productive career. For several years we have been designing an ambitious project to develop [commonsense] theories, which

have slowly been gaining the interest of many researchers inside and outside the Media Lab. To us his loss is indescribable because of how we could communicate so much and so quickly in so very few words, as though we were parts of a single mind."

Henry Lieberman, who worked with Push on the Open Mind project, said, "Push and his students worked on collecting common sense and developing architectures and tools surrounding the knowledge. My students

and I concentrated on applying this knowledge to improve all kinds of interactive applications: browsers, editors, games, phones, etc. Push helped me teach my course, where we taught students how to understand and work with commonsense knowledge. It has now grown to the point that there are numerous projects involving common sense in many, if not most, of the other groups at the Media Lab, and also elsewhere. This collaboration has been one of the most productive and fruitful of my career, and I thus owe him a tremendous debt that I shall never have the opportunity to repay." (For these and many other tributes to Push,

see http://pedia.media.mit.edu/wiki/Push_Singh.)

In my own discussions with Push, we were exploring how to put his commonsense work, especially the Open Mind project, together with the Semantic Web work I do. From the first moment we talked, I could tell he was a brilliant young man with inspiring ideas. I looked forward to working with him, fully expecting to learn more than I could teach. We planned a visit for him to discuss his research and how we could proceed. While I didn't know him as well as Marvin, Henry, and the many students, faculty, and friends he had at MIT and beyond, I can't believe I'll never get a chance to pursue these ideas with this talented young scientist. He will be missed. ■



Photograph by Jim Nihart, PUSH Conference, 2005