

Reports on the AAAI Spring Symposia

22–24 March 1999

Agents with Adjustable Autonomy

This symposium was motivated by the recognition that even as autonomous system technologies mature into practical applications, humans still refuse to disappear. Humans stay in the loop, so practical applications require that the autonomous software be understandable and adjustable.

Adjustable autonomy means dynamically adjusting the level of autonomy of an agent depending on the situation. For real-world teaming between humans and autonomous agents, the desired or optimal level of control can vary over time. Hence, effective autonomous agents will support adjustable autonomy, which contrasts with most work in autonomous systems, where the style of interaction between the human and the agent are fixed by design.

The adjustable autonomy concept includes the ability for humans to adjust the autonomy of agents, for agents to adjust their own autonomy, and for a group of agents to adjust the autonomy relationships within the group. Effective adjustable autonomy minimizes the necessity for human interaction but maximizes the capability for humans to interact at whatever level of control is most appropriate for any situation at any time.

A wide variety of papers were presented, with topics ranging from theoretical autonomy issues to various implemented autonomous systems and the practical autonomy issues faced during their design. In addition, three panels led discussions on the following topics:

First, Issues in Adjustable Autonomy discussed key technical, domain,

and social issues related to adjustable autonomy and its applications.

Second, Dimensions of Adjustable Autonomy discussed the varied dimensions and interpretations of adjustable autonomy, seeking a core definition and theme statement.

Third, Future Directions discussed the future research directions for adjustable autonomy systems, including challenge problems, practical applications, and metrics and evaluation.

The American Association for Artificial Intelligence, in cooperation, with Stanford University's Department of Computer Science, presented the 1999 Spring Symposium Series on 22 to 24 March 1999 at Stanford University. The titles of the seven symposia were (1) Agents with Adjustable Autonomy, (2) Artificial Intelligence and Computer Games, (3) Artificial Intelligence in Equipment Maintenance Service and Support, (4) Hybrid Systems and AI: Modeling, Analysis, and Control of Discrete + Continuous Systems, (5) Intelligent Agents in Cyberspace, (6) Predictive Toxicology of Chemicals: Experiences and Impact of AI Tools, and (7) Search Techniques for Problem Solving under Uncertainty and Incomplete Information.

This symposium was the first community event focused on the topic of adjustable autonomy and sparked lively conversation. Attendees included a particularly strong showing of industrial and National Aeronautics and Space Administration researchers, reflecting the application-motivated nature of the adjustable autonomy concept.

David Musliner
Honeywell Technology Center

Barney Pell
NASA Ames Research Center

Artificial Intelligence and Computer Games

We had a successful symposium this year, the first ever on this topic. We garnered nearly 50 participants, drawing fairly evenly between academia and the game industry. First, we heard from Ernest Adams and John Laird, both discussing what academia and the game industry had to share. We then had a session on the success and failure of AI, followed by a discussion of NPC design, which largely covered emotional aspects of AI. The day ended with demonstrations of research and new products. On Tuesday, we discussed NPC control, most specifically robotic approaches to AI, because as game environments get closer to simulation, the more they begin to look like robotics. Andrew Stern of PFMagic expertly moderated a panel on new directions in AI, looking beyond a "game" to an "interactive entertainment." We then had a long, acrimonious full-group discussion on the future of AI in video games, specifically trying to figure out what the "killer app" in AI gaming is going to be (analogous to DOOM's effect on three-dimensional games). Ian Davis of Activision bet a shiny new quarter that such an application would be out in 2024, but John Laird bet a silver dollar that the AI killer app would be in believable, real-world characters such as Furbys or Tamagotchis (Mike Van Lent made a bet that John Laird would appear shirtless in *PC Gamer* magazine. This prediction is the likeliest of the three). Our last day saw Henrik Lund of Denmark's Lego Lab discussing Lego Mindstorms and Ian Davis revealing the secrets of Civilization III and Dark Reign's AI. Last, we had a full-group discussion on building bridges between academia and the industry, with an emphasis on finding ways to pay for technology development and transfer.

Wolff Dobson
Northwestern University

AI in Equipment Maintenance Service and Support

This symposium was held under the premise that manufacturing compa-

nies offer their customers novel and aggressive service contracts in which the old parts and labor billing model is replaced by guaranteed up time. In turn, the motivation is placed on maintaining equipment in working order by the servicing company with a renewed emphasis on AI technologies. Symposium discussion topics presented here are modeling issues, prognostics, internet-supported diagnosis, and diagnostic information fusion.

With regard to modeling issues, there appears to be a dichotomy between diagnostic design and the diagnostic modeling approach. This dichotomy might be reflected, for example, in the differences in results when probabilities are elicited from field engineers and from design engineers (or getting information from experts versus cases). A common problem is brand new diagnostic designs where no case studies and no field data are available. Design knowledge might here be admissible to determine the bounds of legal values using known components.

To deal with the balance between granularity of diagnosis and cost, it might be practical to start with an abnormal condition detection scheme and then refine the classification by adding particular fault types. This scheme also circumvents the problem of unanticipated failure modes, provided the abnormal condition detector picks up pertinent features. As a next step, the information should be used for design for maintainability and, eventually, concurrent design.

With regard to prognosis, it appears the field is still undergoing the definition phase. One definition for prognosis was the prediction of future faults. A second definition included diagnosis that initializes prognosis by identifying incipient failures; prognosis is then strictly predicting the remaining useful life of a product. A third definition viewed prognosis as the prediction of continuous variables (life) using time-series analysis (without any diagnostics) as a black-box approach. As time passes, the period for prognostics overlaps with that of diagnostics. Similar to diagnosis, prognosis suffers sometimes from the unavailability of appropriate sensors for measuring

desired features indicative of remaining life.

Internet-supported diagnostics hold the potential to share the data of different customers through a common service provider. This service provider would supply the diagnostic algorithms to the shared client database. The use of agents could be leveraged for completing customer-specific tasks and carrying out service provider queries. Potential was also seen for using chip-based mini-web servers, which would allow extended remote diagnosis. IEEE Standards for Information Exchange for Diagnostics could be used as a common denominator for service providers. Bandwidth limitations are a bottleneck to sending actual sensor data, which might in the interim prompt service providers to move the diagnostic algorithms to the site and report only diagnostic results. Another advantage of internet-supported diagnostics is the supportability of products through multimedia support and advanced search and help features. Finally, internet-supported diagnosis could also promote rapid prototyping through the availability of early product versions on the web, which would foster fast feedback and evaluation.

Diagnostic information fusion is concerned with methods and tools that can be used to aggregate the information stemming from different diagnostic tools to arrive at a unified—and presumably better—diagnostic estimate about the state of the system. Any one diagnostic tool has shortcomings in dealing with all faults of interest at the desired level of accuracy. It seems plausible that a fusion scheme could be better than the best of any tool because there is a fair amount of redundant information available that should compensate deficiencies. If several tools agree on the diagnostic state, the task is straightforward, and the resulting output should be assigned higher confidence. However, if tools disagree, a decision has to be made about which tool to believe or to what degree. In addition, information is likely expressed in different domains, such as probabilistic information, fuzzy information, binary information, weights, and so on. The

fusion scheme has to map the different domains into a common one to be able to properly use the encoded data. The fusion scheme also has to deal with tools that are operated at different instantiations in time.

Kai Goebel
GE Corporate Research
and Development

Hybrid Systems and AI: Modeling, Analysis, and Control of Discrete + Continuous Systems

This symposium attracted 54 researchers both from the hybrid systems communities in electrical engineering and computer science and from the AI community.

The use of digital computers to control complex continuous, dynamic processes has contributed to the development of a new field of research that focuses on techniques for analyzing, synthesizing, and controlling dynamic systems whose behavior is modeled by hybrid (discrete + continuous) models. Hybrid system models include intervals of piecewise continuous behavior interleaved with discrete transitions. Each interval of continuous behavior represents a so-called mode of system operation; transitions between modes are discrete and can cause discontinuous changes in the system configuration and variables. Examples of hybrid systems include robots, air-traffic control systems, chemical plants, autonomous spacecraft control, smart buildings, and automated multivehicle highway systems.

The hybrid systems community is a cross-disciplinary community that combines modeling and analysis techniques from discrete-event systems, continuous dynamic systems, and control theory. The growing field of hybrid systems has seen a great deal of activity over the last few years, often focusing on synthesis, verification, and stability analysis of controllers for hybrid systems. Interestingly, a number of the problems addressed by this community are shared by AI researchers studying robotics, online time-critical computation, planning, simulation, verification, execution

monitoring, decision analysis, reasoning about action, diagnosis, modeling and analysis of physical systems, and perception. This symposium brought together these different communities to explore opportunities for exploiting AI representation and reasoning techniques for hybrid system modeling and analysis and integrating techniques from hybrid systems into current AI research.

To accommodate the diverse background of workshop participants, the symposium included four invited talks by researchers from the two communities. Alan Mackworth (who also graciously agreed to be our plenary speaker) presented "The Dynamics of Intelligence: Constraint-Satisfying Hybrid Systems for Perceptual Agents" in which he described the *CONSTRAINT-NET* model, a unitary framework for building hybrid intelligent systems as situated agents. His team has applied this framework to several applications, including soccer-playing robots. Shankar Sastry's talk entitled "Algorithms for the Design of Networks of Unmanned Aerial Vehicles" linked problems in nonlinear control to formal verification methods used in computer science and game theory. He discussed applications in intelligent vehicle highway systems, unmanned aerial vehicles, and air-traffic management systems. Tom Henzinger's talk entitled "Hybrid Games" presented a classification of verification problems based on varying models of hybrid automata. He extended this classification to two-player structures (plant versus controller) to similarly classify control problems. The presentation in particular provided a number of interesting results for polyhedral automata. Brian Williams's talk entitled "Model-Based Programming of Reactive Systems: The Journey of Deep Space One (DS1)" presented the concept of model-directed autonomous systems and his group's experiences in developing the *REMOTE AGENT* autonomous control system. *REMOTE AGENT* is soon to be demonstrated as a technology experiment on the Deep Space One mission.

There were five theme sessions coordinated by session chairs, each of whom provided an overview of his/her session area and facilitated discussion.

The first session of the symposium, Behavioral Programming (chair Michael Branicky) discussed a body of techniques for predictably composing lower-level behaviors into solutions that satisfy higher-level goals. Typically, the lower-level behaviors are given by sensorimotor loops or controllers, operating in a continuous domain, but the higher-level goals are encoded symbolically. Papers covered a range of topics related to the general problem of procedural learning in domains, including walking robots, artificial fish, and cooperating robots for manufacturing. The session entitled Formal Methods (chair Howard Wong-Toi) examined the use of logic to model and analyze hybrid systems. Several papers discussed expressive logic-based theories of action and how to extend them to represent and reason about hybrid systems. Another paper discussed situated multiagent architectures and the mapping of logic-based theories of action to these architectures. A final paper discussed the modal μ calculus, demonstrating that it and various extensions provide an expressively rich yet highly usable logical framework for formal analysis of hybrid systems. There were two sessions on synthesis and control (chairs Feng Zhao and Claire Tomlin) that discussed theory and tools for analyzing, synthesizing, and verifying multimodal hybrid systems. A number of papers went beyond analytic approaches and exploited geometric structure in the phase space to achieve computational efficiency. Other papers used game-theoretic approaches to verify and synthesize controller function. The last session, entitled Applications (chair Dan Clancy), included papers that discussed computational issues, such as the tracking of piecewise continuous behaviors and the enhancing of discrete-event simulation by including continuous system models. The symposium also included a poster session for all authors to discuss their work in detail.

The symposium was a success, and plans are under way for including sessions with AI themes at the next International Hybrid Systems Workshop (HS'00) to be held at Carnegie-Mellon University in March 2000. For further

information about this symposium, see www.ksl.stanford.edu/springsymp99; information on HS'00 is available at www.ece.cmu.edu/~hs00.

Gautam Biswas Vanderbilt
Sheila McIlraith
Stanford University

Predictive Toxicology of Chemicals: Experiences and Impact of AI Tools

This year presented at least two AI opportunities to discuss this topic: One was hosted by IJCAI in summer; the other was hosted by AAAI at the Stanford Spring Symposia. Why this interest in the topic, and what is the topic (and the challenge)? Because it would be impossible to review all the papers presented at our symposium, this article presents my personal impression of some of the major themes of the symposium, organizing them as questions and answers that emerged from the discussions.

Why and When Did the Problem of Toxicity Prediction Emerge?

The goal of toxicity prediction is to describe the relationship between chemical properties, on the one hand, and biological and toxicological processes, on the other. As "computational prediction," researchers intend a prediction based on theoretical values, so the aim is to study a compound without making experiments (either toxicological or physical) with it, possibly before synthesizing it. Why is this topic so important?

We are more and more aware of the need to understand and predict the consequences of chemicals on the health of human beings and wildlife, which is now done through ad hoc experiments that are incredibly expensive, are years long, and involve animal studies. The huge number of compounds to be studied makes this effort especially challenging. Furthermore, in many cases, a single chemical compound can generate many transformation products that are released in the environment over years; each of these transformation product requires, in principle, the same attention devoted

to the parent compound. As a consequence, the number of compounds to be studied becomes enormous.

Consider that 19 million different compounds have now been identified, and only for 10 percent of the industrially produced chemicals are some data on toxicity available. Moreover, a wider use of the newly introduced “combinatorial chemistry” that has been adopted by chemical companies will tremendously increase the number of compounds to be considered; an early example of combinatorial chemistry produced a library of more than 25 billion different compounds.

Several studies have been done from the basic idea that some of the activities of a compound might be related, and some simple toxicity tests have been proposed. It is known, for example, that mutagenic compounds can be also carcinogenic, and indeed, mutagenicity represents a warning for possible carcinogenicity. However, in many cases, no clear relationship can be drawn because the effects are complex and nonlinear.

Decades ago, chemists investigated the effects of particular groups in a family of molecules on particular properties of this family. A famous example is the influence of substituents on the dissociation constant (pKa) of benzoic acids, and it might be possible to predict it with good agreement. Again, this is a kind of knowledge that cannot explain all the toxic effects.

In the 1970s, the great development of ecotoxicology was initiated by the discovery that a certain amount of the toxicity on animals and plants might be explained on the basis of physicochemical properties of compounds and especially of partition coefficients of compounds between *n*-octanol and water ($\log P$). Also, this result cannot explain everything.

Other physicochemical descriptors of chemicals, such as molecular volume, dipole moment, electronegativity, molecular shape, or theoretical indexes, obtained by quantum-chemical calculations, have been proposed. All can be found in toxicological studies. It is now well recognized that no single descriptor can satisfy all the requirements in principle needed to

model highly variegated phenomena.

Indeed, thousands of studies on quantitative structure activity relationship (QSAR) use linear multivariate regressions models. The basic idea beyond QSAR studies is that the activity shown by a given molecule is encoded in its structure. Each chemical has its individual identity, but similar compounds exhibit similar activity. One of the important points is the definition of molecular similarity, which has been also tackled by a few presentations.

Why Toxicity Prediction Meets AI, and Machine Learning in Particular?

Is it possible to use the data and the knowledge available to predict the effects of chemicals? Is this a problem for the machine learning community, or something else?

The problem is how to connect molecular descriptors or physical properties or simple *in vivo* tests with the biological (toxic) activity. In most cases, deterministic or statistical approaches have been used to investigate this relationship within QSAR. There is often a correlation among the parameters, and thus, the interpretations can easily be misleading.

Without forgetting about statistics, it seems important to reason more about finding this relationship. AI techniques are, in principle, a good candidate; they can allow reasoning on the data, extracting knowledge, finding nonlinear function approximations, and building hypotheses. However, looking at the way machine-learning methods evolve, we can find that something is different here. In machine learning, data sets are used to build the model and eventually a separate set of test data. Every individual in the data set is characterized by a number of features. Data are considered good, and missing values can be dealt with by some methods. A large number of such data sets are available for the machine-learning community to try.

Why Is Toxicity Prediction of Chemicals Not There?

Some participants expressed the hope that data sets of toxicity prediction can be posted on the machine-learning


repository, but many difficulties and criticisms emerged.

The National Toxicology Program data set has been proposed by many participants, especially by Douglas Bristol, as a reasonable example. These data sets are based on the study of about 300 compounds to be used as a training set and the definition of small test sets (30 compounds). For all the molecules, the carcinogenicity is available, expressed as yes-no. This is the data set proposed for the AI challenge that was discussed at IJCAI. The reasons for some criticisms, mainly from the experts (that is, the toxicologists), are that the data set presents a mix of chemicals (300 organic and inorganic as representative of 19 million); some chemical classes are not represented; and some known biological mechanisms are not represented. Thus, how can we infer something significant?

Another criticism is of the choice of data sets and the distribution between genotoxic and nongenotoxic ones, which might be biased by the current need to explore some of them in greater detail.

Another, more consistent criticism emerged about the quality of the data available worldwide on biological effects. Every experimental datum has been produced by one or more institute in one or more experiment. What we know about a studied compound is clear and unambiguous. A comparative study presented by Emilio Benfenati on a data set extracted from different toxicology databases recognized as the principal ones has shown magnitude-of-order differences in the same datum, showing the questionable quality of the data, at least according to machine-learning standards. Because biological activities are largely variable, variability in the data should be dealt with, but how do we correctly use machine-learning paradigms?

This fact seems to strongly encourage relying more on the chemical structure of the compound, and on computed properties and computed indices, rather than on experimental values. In fact, a similar analysis on the variability of parameters computed with different programs gives better results. Results presented by Alan



Artificial Intelligence and Mobile Robots

Case Studies of Successful Robot Systems

Edited by David Kortenkamp, R. Peter Bonasso, and Robin Murphy


The mobile robot systems described in this book were selected from among the best available implementations by leading universities and research laboratories. These are robots that have left the lab and been tested in natural and unknown environments. They perform many different tasks, from giving tours to collecting trash. Many have distinguished themselves (usually with first- or second-place finishes) at various indoor and outdoor mobile robot competitions.

Each case study is self-contained and includes detailed descriptions of important algorithms, including pseudo-code. Thus this volume serves as a recipe book for the design of successful mobile robot applications. Common themes include navigation and mapping, computer vision, and architecture.

6 x 9, 400 pp. \$37.50, ISBN 0-262-61137-6

Prices higher outside the U.S. and subject to change without notice.

**To order, call 800-356-0343 (US and Canada) or (617) 625-8569.
Distributed by The MIT Press, 55 Hayward, Cambridge, MA 02142**



Katritzky shows promise for predicting physicochemical properties (to be used then for toxicity prediction).

Where Do We Look for Data Sets?

A huge effort should be made by toxicologists and is indeed in progress. Standardized protocols for experiments are used more and more, which will afford more reproducible data. Meanwhile, several databases are available, a few large and complex (more than 100,000 compounds), others smaller but more homogeneous. A list of such web sites is available on the symposium web site.

Machine Learning or Scientific Discovery?

A new emphasis on this question arose from the observation of John Frazier, who said, "Computer scientists, you are looking at the wrong problem, you will not go far trying to connect directly the structure with the activity, because we need to model all the process that brings the chemical through the living systems and indirectly produces the toxic effect. We know very little of this process." Indeed, no machine learning can discover knowledge not present in the data, and we are not sure that our data contain all the relevant knowledge.

We are sure that representing and understanding all the biological processes will be a greater challenge for AI, too, but it seems far off. We can

start now with machine learning because the problem of the regulatory commissions is great and important now, and we cannot wait. According to end users of predictive systems, as expressed by the regulatory commissions in America and Europe, the problem of assessing the risks of chemicals is huge and urgent and will offer a perfect-order integration of research and real-world applications.

Which AI Technique Is Better?

This symposium tried to highlight the potential of different AI approaches, either individually or combined, for computational toxicity prediction.

AI tools have yet to be fully evaluated in this domain. Which techniques are better for toxicity prediction, especially given our changing understanding of toxicology? Hybrid approaches combining ILP, argumentation, ANN, Bayesian networks, fuzzy logic, GA, and rough sets with mathematical and statistical methods such as discriminant analysis and PCA have been discussed by many. The report presents some results of predictive models done with some AI contributions that show promising success in areas where the classical statistical methods failed. Is this enough?

I would like to conclude this point with the words of Ann Richard: "Design models that speak the language of toxicologists and chemists, that provide a valid framework for biological and chemical conceptualization, that produce predictions that can be rationalized and justified, that can generate testable hypotheses concerning toxicity mechanisms, and that have scientific credibility."

Further information is posted at www.elet.polimi.it/AAAI-PT.

Giuseppina Gini

Politecnico di Milano, Italy

Search Techniques for Problem Solving under Uncertainty and Incomplete Information

Many real-world systems have to perform robustly in the presence of uncertainty. This symposium, chaired by Weixiong Zhang and Sven Koenig,

covered topics such as how to search different representations of uncertain information (including belief networks and Markov models), how to explore and map unknown environments, how to use learning and reasoning about uncertainty to improve search performance, and how to allocate resources under uncertainty. Application domains included manufacturing; linguistics; genetics; robotics; design; scheduling; health care; coding theory; constraint satisfaction; and games such as Chess, Poker, Go, and Tetris.

Eric Horvitz, Rina Dechter, and Murray Campbell (representing IBM's Deep Blue chess team) gave invited talks. Common themes of their presentations, as well as the presentations by the other participants, covered the need to use nonuniform search techniques to cope with the search complexity in the presence of uncertainty, including combining search methods that run in software and hardware, combining simulation and search, combining learning and search, interleaving search with action executions, and using multiple abstraction levels with different evaluation functions during the search. It was also discussed how to generalize the scalar evaluation functions often used by heuristic search methods to intervals (using either total or partial orderings), probability distributions, or tuples of values that can be used in conjunction with possibly nonadditive multiattribute utility functions. Other issues included nonmyopic search control, the selection of search methods based on domain properties, and the development of standard search architectures and tools. Despite the impressive progress, the general feeling was that we are just beginning to understand the hard issues involved in search control under uncertainty and incomplete information. These issues will remain at the core of AI research.

Sven Koenig

Georgia Institute of Technology

Shlomo Zilberstein

University of Massachusetts

Weixiong Zhang

University of Southern California/
Information Sciences Institute