CANADA RESEARCH CHAIR

Reprinted courtesy Government of Canada, Canada Research Chairs Program

CHAIRHOLDER PROFILES

Jonathan Schaeffer
Canada Research Chair in Artificial Intelligence
The University of Alberta
Tier 1 - January 1, 2002

Achievements: Inventor of Chinook, the world checkers champion. Chinook is the first computer program to win a human world championship, a feat recognized by the Guinness Book of World Records. Author of four books and more than one hundred articles and papers on artificial intelligence, parallel computing, and bioinformatics; holder of Informatics Circle of Research Excellence (iCORE) Chair; American Institute of Artificial Intelligence Fellow; recipient of E.W.R. Steacie Award from NSERC; co-founder of BioTools Inc.; consultant to Electronic Arts Canada and BioWare.

Research Involves: Development of high-performance, real-time artificial intelligence applications

Research Relevance: Development of high-performance, real-time artificial intelligence applications

GAMES THAT IMITATE LIFE

Dr Jonathan Schaeffer is one adult who knows that computer games represent much more than just a way of killing time. One of the world's leading authorities on artificial intelligence, Dr Schaeffer believes that games are ideal domains for exploring the capabilities of computational intelligence. Because, unlike life, the rules are fixed, the scope of the problem is constrained and the interactions of the players are well defined; games can act as perfect "control" situations. Games can also be a microcosm of the real world, and successfully achieving high computer performance in a non-trivial game can be a stepping stone toward solving more challenging real-world problems.

Dr Schaeffer's program as Canada Research Chair in Artificial Intelligence will use games as experimental test beds for artificial intelligence research. His intent is to achieve a better understanding of what it takes to build high-performance systems that operate in real time, where high-performance is defined as achieving a performance level comparable to, or better than, that of the best humans. Rather than seeking small incremental advances that are in isolation of the domain in which the ideas could be applied, Dr Schaeffer plans to take an application and solve all the problems necessary to achieve high performance. His research is driven by the requirements of the application, with the intent that, if interesting applications are selected, the research results will have wide applicability. His past research is a prime example, having led to the formation of a bioinformatics company whose software is being used by research laboratories around the world.

Specifically, his research project will tackle three priorities: single-agent search games (puzzles), with emphasis on path finding algorithms, DNA sequence alignment algorithms, and development of a generic planning system; algorithms for computer-based poker, with attendant problems of dealing with imperfect information and opponent modelling; and development of "realistic" characters for sports and role-playing games.

HIGH-PERFORMANCE ARTIFICIAL INTELLIGENCE SYSTEMS

iCORE Chair
Computing Science
University of Alberta

Dr Jonathan Schaeffer is iCORE Chair of High Performance Artificial Intelligence Systems at the University of Alberta. iCORE has committed $500,000 per year for five years for a total of $2.5 million dollars to develop this research group.

EXECUTIVE SUMMARY

This report represents a summary of the second year of the iCORE Chair in High-Performance Artificial Intelligence Systems. The group now consists of three professors, four affiliated professors, one post-doctoral fellow, nine PhD students, and 16 Masters students. In addition, there are five programmer/analysts (two part-time), and a half-time secretary.

The High-Performance Artificial Intelligence Systems research group specializes in artificial intelligence research, investigating new technologies for creating "intelligent" behaviour in a computer. Although the research spans many areas of artificial intelligence, including search, machine learning, and heuristic knowledge, historically the group has used games to demonstrate the ideas. Fundamental problems in artificial intelligence are being investigated in the context of computer programs that play chess, checkers, Go, and poker. Many of the game-playing programs have achieved a high level of performance and have challenged the best human players in the world.

Although the group's reputation was initially made by applying research work to classic board and card games, since 1999 the team has been moving more towards addressing the challenges of the commercial games industry. Commercial games (or, more precisely, interactive entertainment) is a maturing industry that had $15 billion in sales in North America last year, with an impressive 15 percent growth in the market. In the past year the team strengthened ties with Electronic Arts of Vancouver (the largest games company in the world) and BioWare of Edmonton (the world leader in role-playing games). The new technology has been well received by both companies, with good prospects for integration into commercial products.

Another thrust of this project is the development of parallel programming environments. For over 15 years, the team has been building tools to simplify the task of parallel programming. The third generation tool, CO_2P_3S, is now available on the web and actively promoted at major parallel computing conferences.

RESEARCH GOALS AND OBJECTIVES

The project is progressing extremely well along the lines outlined in the original iCORE proposal. The group has built an international reputation based on artificial intelligence (AI) research, using games as an experimental test-bed for this work. However, the research challenges from the classic board and card games are limited (the games of poker and Go being notable exceptions). Since 1999, we have been slowly moving our research efforts towards addressing the challenges of the commercial games industry. This represents a golden opportunity for us, since artificial intelligence research in this industry is still in its infancy. At this point in time, over 40 percent of our graduate students are working in this area, and this number will only increase. More details can be found at www.cs.ualberta.ca/~games.

RESEARCH PROJECTS

In the past year we have made major strides forward in engaging the commercial games industry and making significant progress in doing industry-based research. We have become one of the largest research groups in this area. However, as we are learning, there is a large gap between academic research and industry expectations. The commercial games industry in particular is heavily performance oriented. They need real-time solutions that use little CPU and memory. Few AI efforts address real-time constraints – an area which is one of our research group's strengths. We are not developing industrial strength solutions for our partners, but we are building proof-of-concept demonstrations that show that our technology can meet the stringent industry demands.

Our group continues to build on its past success in artificial intelligence. Most notable is the poker project, which is addressing the hard AI problems of reasoning with imperfect and incomplete information. Our progress in the past year has been excellent, developing new technology that has resulted in a quantum improvement in the state of the art, Our poker-playing program became the first such program to be competitive with a top human player (January 2003). In the upcoming year, we hope to challenge the best players in the world.

WE HAVE ONE OF THE STRONGEST GROUPS IN THE WORLD WORKING ON DEVELOPING HIGH-PERFORMANCE SEARCH ALGORITHMS. UNLIKE MOST RESEARCH GROUPS, WE BUILD COMPLETE AI SYSTEMS, ADDRESSING ALL THE ISSUES NEEDED TO ACHIEVE HIGH PERFORMANCE.

The long-term objective of our work is to enhance our understanding of search, knowledge and their interactions. We have one of the strongest groups in the world working on developing high-performance search algorithms. Unlike most research groups, we build complete AI systems, addressing all the issues needed to achieve high performance. It always starts with search (well defined and understood), integrating application-dependent knowledge (not yet well understood) only on an as-needed basis. Discovering new ways to lessen dependence on knowledge is critical to AI success; human knowledge is fraught with error and difficult to obtain. The goal is to automate this process as much as possible.

Part of the project funding supports research into parallel computing, which was not discussed in the original iCORE proposal. For over 15 years we have been developing new parallel algorithms and tools to simplify the difficult task of writing a correct parallel application. These activities have always hovered around 30 percent of the chair's research time. While this research area is not artificial intelligence, it is considered "high performance."

Of interest is that the technology we built to develop parallel applications (our CO_2P_3S parallel programming environment) is directly applicable to our artificial

S C H A E F F E R

intelligence scripting project. CO_2P_3S builds on the (sequential) software idea of design patterns – exploiting commonly occurring software designs. CO_2P_3S uses parallel design patterns. We have copied this technology for AI scripting. Character behaviour also follows patterns. If one describes a character as a "guard" then that conveys a lot of information about that character's behaviour. The guard notion becomes a behavioural pattern that can be customized to give the specific behaviour that is desired. It is interesting that the technology we developed for parallel computing is relevant in artificial intelligence.

Commercial Games Research

In the past, computer graphics were the major technological differentiators between competing games products. The realism of the graphics has increased consumer demand for realism in the game characters. The commercial games industry now recognizes that artificial intelligence has become a major consumer consideration in assessing the quality of a product. Unfortunately, the games industry has few AI experts researching new technologies, giving universities an opportunity to have a major impact in new technology development. In academia, the University of Alberta has the world's largest research group working in this area.

1.) The first major thrust is in AI scripting. Character behaviours in games are usually defined using scripts. However, the result is complex software that is hard to maintain. Further, the resulting performance of the characters is disappointing because the characters will only do precisely what has been scripted, and typically this is a very small (usually one) set of behaviours. We have been developing a tool that allows for the rapid construction of complex character behaviours. The tool, called ScriptEase, is based on having a rich set of pre-defined behaviours (for characters, speech, situations, and plot) that the user can select and then customize to their needs. This work is novel and, because our extensive experience with patterns (see the CO_2P_3S section below) gives us a competitive edge for developing the next generation of scripting technology. Our prototype tool has been used to build complex stories in a very short time. The work has been demonstrated to BioWare and been very well received. Creating realistic characters has many industrial applications, including training programs, web interfaces, and other forms of interactive entertainment.

2.) The second major thrust is pathfinding. For many computer games, the "simple" task of having a character find a path from their current position to a goal is a time-critical, CPU-intensive function. This is an instance of a problem domain called single-agent search, but in this case is restricted to a two-dimensional grid (with the intent of moving to three dimensions). We developed new algorthms for grid-based pathfinding, yielding some surprising results that run counter to conventional wisdom. BioWare has implemented some of our ideas in their next product and report that they resulted in improved performance. The same technology is applicable to a wider domain of applications, including robot planning.

3.) The third major thrust is applying machine learning to games. Game companies are reluctant to ship games that learn in response to the user's interactions. The reason for this is that it is difficult to control the learning, and a player can contrive to have a program learn poor behaviour. Also, conventional learning algorithms are either too slow, or learn too slowly. For example, in Electronic Arts successful FIFA soccer game, the computerized soccer players are incapable of adjusting their play to match that of their human opponents. We have developed new technology that allows computer soccer players to dynamically modify their behaviour in a controlled way, allowing the program to recognize when it has made a mistake and adjust its play so that the mistake is not repeated. This technology has been enthusiastically endorsed by Electronic Arts.

A major highlight of this year was Jack van Rijswijck's paper on the machine learning algorithms that he developed for FIFA soccer. This work was accepted for presentation at the annual Game Developer's Conference. This is the premier conference in the industry, with a heavy emphasis on new developments that can impact game-program development. Very few academic papers have ever been accepted for this conference.

Another commercial games-related research initiative is Michael Buro's work on real-time strategy games. He has developed a test-bed for exploring issues in real-time strategy games including client-server architectures, managing limited CPU resources, and complex group behaviours. Buro is working with Relic, a Vancouver-based games company.

Classic Games

Traditional games research has concentrated on two-player games of perfect information (the opponents are not hiding anything). Poker is very challenging because of hidden information (you do not know the opponent's cards), multiple players (typically 10 in a game), and deception (bluffing is critical to successful play). These dimensions significantly complicate the problem domain, making it an application domain that better represents the complexities of intelligence in real life. For example, poker is a model for economic game theory as well as business negotiations and Internet auctions.

For almost a decade we have been developing new technologies for dealing with imperfect information. We have applied the notion of Nash equilibriums to build a pseudo-optimal two-player poker program (an optimal program is too computationally expensive to build right now). This program achieved international success by narrowly losing a match to a world-class player in January 2003. Plans are in place for a real-money match against one of the best players in the world in 2003.

Other efforts in classic games include:

1. Martin Müller has built up a team of six people working on computer Go. Unless games like chess, search is ineffective here. Success in the game depends on using complex interacting knowledge.
2. For almost a decade we have been working on solving the game of checkers. It has a search space of $O(10^{20})$ – a daunting number. We believe it likely that we can solve the game in the next few years. That is we will have a program that will never lose (assuming checkers is a draw with perfect play, as seems likely). Although the final result – solving checkers – is not particularly exciting from the scientific point of view, the technology and tools developed to solve such a large computational problem are relevant to a wide audience.

We continue to improve our world-championship programs for the games of Lines of Action and shogi (Japanese chess). In addition, we were the first team to build a perfect program for the game of 10 x 10 domineering.

Other Artificial Intelligence Initiatives

Planning: Many of our search-based research contributions are applicable to the field of artificial intelligence planning systems. For the past year we have been building a hierarchical planning system. It takes a planning problem domain (e.g. a robot having to plan how to restock inventory) and decomposes it into a global problem (what has to be done) and a series of local problems (stocking individual items). The result is a system that can come up with workable plans considerably faster than conventional approaches. We are working on generalizing the technology to handle a wider set of application domains.

Optimal multiple sequence alignment: A cornerstone for understanding the human genome is the computational problem of sequence alignment-determining the (dis)similarity of DNA protein strands. We have developed new technology for performing an optimal alignment of multiple (long) DNA/protein strands that is roughly four times faster than existing approaches. In the past year we have worked with biologists to assess and improve the quality of our alignment results.

CO_2P_3S

The CO_2P_3S project attempts to use modern software technology to simplify the complexities of developing parallel applications. CO_2P_3S stands for Correct Object-Oriented Pattern-based Parallel Programming System. As the name suggests, the package uses objected-oriented technologies, design patterns and frameworks to facilitate code development. A user selects a parallel design pattern that best matches their application needs, selects some options to customize it to their application, and then fills in CO_2P_3S-generated sequential code stubs with application-dependent code. The result is a complete, functional parallel application. The software is available for download (www.cs.ualberta.ca/-systems/cops).

The state of the art in parallel programming tools remains primitive, and we face a difficult task to demonstrate the value of our tool set. Despite being well received in academia (for example, reflected by a best paper prize), we have not yet been able to build up a strong user community.

CISS

We initiated CISS, the Canadian Internetworked Scientific Supercomputer (www.cs.ualberta.ca/-ciss). We have worked on developing grid-like software that can be installed at the user level, without need for system administrator support. The software allows an otherwise idle computer to "pull" in computational tasks to be executed from a remote site. In effect, one can create a virtual supercomputer. The scalability and portability of the software was demonstrated on a national scale on November 4, 2002. On that date, we were able to harness 1,376 computers spanning 20 administrative domains at 18 different sites. In a single day, we were able to do three and a half years of computing to help solve an interesting computational chemistry property.

The purpose of CISS was threefold. First, it demonstrated the functionality of the software used. Second, it furthered chemistry research. Third, it helped build the social infrastructure for sharing high-performance computing resources in Canada.

WestGrid

Although this is not a direct research contribution, in many ways the WestGrid project may have the most long-term impact. WestGrid is a partnership of eight Alberta and British Columbia institutions to bring world-class high-performance computing resources to western Canada. The partners are the University of Alberta, University of British Columbia, University of Calgary, University of Lethbridge, Simon Fraser University, TRIUMF, Banff Centre, and NewMIC. This project was successful at achieving roughly $50 million of funding from the Canada Foundation for Innovation, the province of Alberta, the province of British Columbia, computer vendors, and the member institutions. The five co-principal investigators for the project are Jonathan Borwein (Simon Fraser University), Gren Patey (University of British Columbia), Jonathan Schaeffer (University of Alberta), Brian Unger (University of Calgary), and Mike Vetterli (SFU/TRIUMF). Although the AI research will benefit in only small ways from this infrastructure, the impact on the research productivity of Alberta and British Columbia researchers will be immense. There will be major benefits to researchers in areas diverse as biology, chemistry, physics, engineering, medicine, and the social sciences.

RESEARCH TEAM

TEAM LEADER	AWARDS
Jonathan Schaeffer	Canada Research Chair in Artificial Intelligence Fellow, AAAI NSERC E.W.R. Steacie Fellowship

TEAM MEMBERS	TITLE
Michael Buro	Associate Professor
Martin Müller	Associate Professor

OTHER TEAM MEMBERS	TITLE
Russ Greiner	Professor
Rob Holte	Professor
Paul Lu	Assistant Professor
Duane Szafron	Professor

POSTDOCTORAL FELLOWS	TOPIC	AWARDS
Yngvi Bjornsson	Learning Search Control	Gold medal-Computer Olympiad

PHD CANDIDATES	TOPIC	AWARDS
Darse Billings*	Computer Poker	+
Adi Botea*	Planning Systems	+
Markian Hlynka*	Learning Search Control	+
Akihiro Kishimoto	Computer Go	World Computer Shogi Champion
David O'Connell*	Single-agent Search	PGS-B
Ehud Sharlin*	Tangible User Interfaces	+
Brian Sheppard	Computer Scrabble	
Jack van Rijswijck*	AI Architectures for Sports Games	+
Peter Yap*	Pathfinding on a Grid	
Ling Zhao	High-level Planning	Alberta Ingenuity

MSC CANDIDATES	TOPIC	AWARDS
Michael Chung*	Real-time Strategy Games	Alberta Ingenuity
Patrick Earl*	Meta Parallel Programming	
Mark Goldenberg*	Parallel Job Scheduling	PGS-A
Dave Gomboc	Tuning Evaluation Functions	
Zhuang Guo*	Web-server Patterns	
Thomas Hauk*	Probabilistic Two-player Search	
Alex Kovarksy	Machine Learning in RTS	+
Jonathan Newton	Learning Mistakes	
Xiaochen Niu	Heuristic Knowledge and Search	
Dominique Parker*	Pattern-based AI Scripting	
James Redford*	Pattern-based AI Scripting	
Terry Schauenberg*	Opponent Modelling	PGS-A
Xiaomeng Wu	Bayesian Learning	
Jonathan Yip	Scripting in RTS	PGS-A
Haizhi Zhang*	Search Algorithms	
Jianjun Zhou	Incremental Search Algorithms	

Students that are supervised or co-supervised by Jonathan Schaeffer are indicated by a *. Faculty involved in (co-)supervising these students include Michael Buro, Rob Holte, Paul Lu, Martin Müller, Duane Szafron, Jaap van den Herik (University of Maastricht) and Ben Watson (Northwestern University). Students who are current or past holders of a major scholarship are indicated by a +.

SUPPORT	POSITION
Neil Burch	Programmer/Analyst
Aaron Davidson	Programmer/Analyst
Marcus Enzenberger	Programmer/Analyst
Amanda Hansen	Administration
Matthew McNaughton	Programmer/Analyst
Kai Tan	Programmer/Analyst

COLLABORATIONS

The group is actively working with several partners:

1. Electronic Arts (commercial games research) has historically provided cash and graduate student internships. In the past year, they made a (small) software donation to the group.

2. BioWare (commercial games research) sponsors the research with $10,000 per year.

3. Relic (commercial games research) is negotiating a project.

4. Joerg Denzinger, University of Calgary works on a joint research project, supported by Intelligent Robotics and Intelligent Systems (IRIS) NCE funding.

5. Strong research ties with IKAT at the University of Maastricht (The Netherlands) and the Computer Games Laboratory at Shizouka University (Japan) include annual visits and graduate student exchanges.

6. WestGrid is a multi-institutional initiative (University of Alberta, University of British Columbia, University of Calgary, University of Lethbridge, Simon Fraser University, TRIUMF, Banff Centre, and NewMIC) and multi-disciplinary initiative. The industrial partners include Hewlett Packard, IBM, and Silicon Graphics.

7. Alberta Ingenuity Center for Machine Learning (AICML). This research center was formed in the past year, with Jonathan Schaeffer one of the co-principal investigators. AICML is starting to work with a number of industrial partners.

FUNDING

In addition to the iCORE grant of $500,000 per year, Russ Greiner, Rob Holte, Randy Goebel, and Jonathan Schaeffer attracted $7 million over five years from Alberta Ingenuity's Centre of Excellence program. NSERC provides yearly operating grants of $20,000 and $46,200 to Drs Müller and Schaeffer, and part of the $600,000 NSERC MFA grant (Pollard et al. with Schaeffer as a coapplicant) provides infrastructure. Dr Schaeffer's Tier 1 CRC provides $200,000 per year for salaries and overhead. The National Centres of Excellence funding provides $155,000 (IRIS-Schaeffer) and $125,000 (PENCE-Szafron et al.) per year.

The $11,990,000 Canada Foundation for Innovation WestGrid grant is to build high-performance computing facilities in Western Canada. The PIs are Jonathan Borwein (Simon Fraser University), Gren Patey (University of British Columbia), Jonathan Schaeffer (University of Alberta), Brian Unger (University of Calgary), and Mike Vetterli (TRIUMF). Schaeffer and Unger leveraged the CFI funds to get $6 million in provincial matching funds. Combined with vendor contributions and operating funds from CFI, the total project budget is roughly $50 million.

INTELLECTUAL PROPERTY

Schaeffer is the co-founder of BioTools Inc. (www.biotools.com), a bioinformatics company. BioTools has three successful commercial products: PEPTOOL (protein analysis), GENETOOL (DNA analysis), and CHROMATOOL (DNA/protein assembly). These products are used in over 1,000 research laboratories around the world. Success with these products led to the opportunity to do contract work with some of the biggest players in the human genome efforts. Currently most of BioTool's work is on a contractual basis.

Chenomx is a spinoff from BioTools (www.chenomx.com). Chenomx has developed revolutionary software technology to do fluid analysis. From a spectrogram produced by a NMR machine, its programs can analyze the data to a level of detail not easily possible in a laboratory. The first application is to analyze urine. Conventional urine analysis (as prescribed by a doctor) returns the analysis of six (of over 250) compounds in the urine. Chenomx's software accurately returns an analysis of over 100 compounds, faster and at less cost. The company has partnered with Varian and Breuker, the two largest NMR manufacturers in the world. Its product, ECLIPSE, is currently under evaluation by a major pharmaceutical company.

BioTools and Chenomx are successes, but both have been hampered by a lack of venture capital. Together they employ over 20 people and have combined revenues of roughly $1 million.

Additional Activities

Our group has had several other notable events happen:

1. Hosted the third biennial Computers and Games conference in Edmonton (July 2002). Over 110 people attended from around the world.
2. Martin Müller organized the 21st Century Cup Computer Go championship in Edmonton (July 2002). Twelve programs from around the world competed. The two University of Alberta entries finished in the middle of the pack.
3. Schaeffer is the co-author of the official FIDE rules for man-machine chess matches (FIDE is the body that governs international chess). Schaeffer was one of the match officials for the Garry Kasparov-DEEP JUNIOR chess match in New York (January/February, 2003).
4. Schaeffer is on the executive committee of C3.ca, the national voice for high-performance computing in Canada. C3.ca is producing a long-range strategy for funding high-performance computing in Canada. The team is headed by Kerry Rowe (Vice-President Research, Queen's University) and Schaeffer is one of the seven co-authors.
5. Competed in the 2002 World RoboCup Championships (small-sized league). Matt McNaughton and his team won two games and lost two games, placing third in their division.

PUBLICATIONS

Refereed journal papers

1. R. Hayward and J. van Rijswijck, "Hex and mathematics," *Discrete Mathematics*, 2003, to appear.
2. S. MacDonald, J. Anvik, S. Bromling, D. Szafron, J. Schaeffer, and K. Tan, "From patterns to frameworks to parallel programs," *Parallel Computing*, vol. 8, no.12, 2002, pp. 1663-1683.
3. J. Zhou and M. Müller, "Depth-first discovery algorithm for incremental topological sorting of directed acyclic graphs," *Information Processing Letters*, 2003, to appear.
4. M. Müller, "Conditional combinatorial games and their application to analyzing capturing races in Go," *Information Sciences*, 2003, to appear.
5. M. Müller, "Counting the score: Position evaluation in computer Go," *Journal of the International Computer Games Association*, vol. 25, no.4, 2002, pp. 219-228.
6. M. Buro, "Report on the IWEC-2002 man-machine Othello match," *Journal of the International Computer Games Association*, vol. 25, no.2, 2002, pp. 113-114.
7. N. Bullock, "Domineering: Solving large combinatorial search spaces," International Journal of the Computer Games Association, vol. 25, no.2, 2002, pp. 67-84.
8. Y. Bjornsson and T. A. Marsland, "Learning extension parameters in game-tree search," Information Sciences Journal, 2003, to appear.

Refereed conference papers

1. D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron, "Approximating game-theoretic optimal strategies for full-scale poker," in International Joint Conference on Artificial Intelligence (IJCAI), 2003, pp. 661-668.
2. M. Goldenberg, A. Kovarksy, X. Wu, and J. Schaeffer, "Multiple agents moving target search," in International Joint Conference on Artificial Intelligence (JJCAI), 2003, pp. 1511-1512.
3. Y. Bjornsson, M. Enzenberger, R. Holte, J. Schaeffer, and P. Yap, "Comparison of different abstractions for pathfinding on maps," in International Joint Conference on Artificial Intelligence (IJCAI), 2003, pp. 1536-1538.
4. M. McNaughton, J. Redford, J. Schaeffer, and D. Szafron, "Pattern-based AI scripting using ScriptEase," in AI'2003: The Sixteenth Canadian Conference on Artificial Intelligence, 2003, pp. 35-49.
5. K. Tan, D. Szafron, J. Schaeffer, J. Anvik, and S. MacDonald, "Using generative design patterns to generate parallel code for a distributed memory environment," in ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), 2003, pp. 203-215.
6. A. Botea, M. Müller, and J. Schaeffer, "Extending PDDL for hierarchical planning and topological abstraction," in iCAPS workshop on PDDL, 2003, pp. 25-32.
7. C. Pinchak, P. Lu, J. Schaeffer, and M. Goldenberg, "The Canadian internetworked scientific supercomputer," in High Performance Computing Systems and Applications, D. Senechal, Ed., 2003, pp. 193-199.
8. A. Kishimoto and J. Schaeffer, "Distributed game-tree search using transposition table driven work scheduling," in International Conference on Parallel Processing (ICPP), 2002, pp. 323-330.
9. S. Bromling, S. MacDonald, J. Anvik, J. Schaeffer, D. Szafron, and K. Tan, "Pattern-based parallel programming," in International Conference on Parallel Processing (ICPP), 2002, pp. 257-265.
10. A. Kishimoto and J. Schaeffer, "Transposition table driven work scheduling in distributed game-tree search," in AI'02: Advances in Artificial Intelligence, ser. Lecture Notes in Artificial Intelligence, R. Cohen and B. Spencer, Eds. Springer Verlag, 2002, pp. 56-68.
11. S. MacDonald, D. Szafron, J. Schaeffer, J. Anvik, S. Bromling, and K. Tan, "Generative design patterns," in 17th IEEE International Conference on Automated Software Engineering (ASE), 2002, pp. 23-34.
12. M. McNaughton, P. Lu, J. Schaeffer, and D. Szafron, "Improving single-agent search using memory-efficient heuristics," in American Association for Artificial Intelligence (AAAI) National Conference, 2002, pp. 737- 743.

13. M. Müller, "Multicriteria evaluation in computer game-playing and its relation to AI planning," in AIPS-2002 Workshop on Planning and Scheduling with Multiple Criteria, 2002, pp. 1-6.
14. M. Müller, "A generalized framework for analyzing capturing races in Go," in Sixth Joint Conference on Information Sciences, 2002, pp. 469-472.
15. M. Buro, "ORTS: A hack-free RTS game environment," in Computers and Games, ser. Lecture Notes in Artificial Intelligence. Springer Verlag, 2003, to appear.
16. M. Buro, "ORTS: A hack-free environment for real-time strategy games," in International Joint Conference on Artificial Intelligence (IJCAI), 2003, to appear, pp. 1534-1535.
17. Y. Bjornsson and T. A. Marsland, "Learning control of search extensions," in Joint Conference on Information Sciences, 2002, pp. 446-449.
18. Y. Bjornsson and M. Winands, "YL wins Lines of Action tournament," Journal of the International Computer Games Association, vol. 25, no.3, 2002, pp. 185-186.
19. A. Botea, "Using abstraction for heuristic search and planning," in 5th International Symposium on Abstraction, Reformulation, and Approximation, ser. Lecture Notes in Artificial Intelligence, S. Koenig and R. Holte, Eds., vol. 2371. Springer Verlag, 2002, pp. 326-327.
20. L. Zhao, "Tackling Post's correspondence problem," in Computers and Games, ser. Lecture Notes in Artificial Intelligence. Springer Verlag, 2002, to appear.
21. E. Sharlin, Y. Itoh, B. Watson, Y. Kitamura, L. Liu, and S. Sutphen, "Cognitive cubes: A tangible user interface for cognitive assessment," in ACM Computer-Human Interfaces (CHI), 2002, pp. 347-354.
22. J. van Rijswijck, "Learning goals in sports games," Game Developers Conference, 2003, to appear
23. C. Pinchak, P. Lu, J. Schaeffer, and M. Goldenberg, "The Canadian Internetworked Scientific Supercomputer," in *High Performance Computing Systems and Applications (HPCS)*, 2003, pp. 193-199.
24. A. Botea, M. Müller, and J. Schaeffer, "Using abstraction for planning in sokoban," *in Computers and Games*, ser. Lecture Notes in Artificial Intelligence. Springer Verlag, to appear.
25. M. Müller, "Proof-set search," in *Computers and Games*, ser. Lecture Notes in Artificial Intelligence. Springer Verlag, 2003, to appear.
26. M. Buro, "The evolution of strong Othello programs," in *Entertainment Computing-Technology and Applications*, R. Nakatsu and J. Hoshino, Eds. Kluwer, 2003, pp. 81-88.

Magazine papers

1. J. Schaeffer, "Tangled up in blue," New Scientist, vol. 91, no. 3, May-June 2003, pp. 276-278.

Conference posters

1. Y. Xu, A. Huckauf, W. Jager, P. Lu, J. Schaeffer, and C. Pinchak, "CISS-l experiment: Ab initio study of chiral interactions," in 39th International Union of Pure and Applied Chemistry (IUPAC) Congress and 86th Conference of the Canadian Society for Chemistry, 2003, poster abstract, to appear.
2. B. Li, S. Buckingham, J. Schaeffer, A. Spencer, and W. Gallin, "Computational analysis of voltage-gated potassium channels," in Intelligent Systems for Molecular Biology (ISMB), 2003, p. 131, poster abstract.

Abstract

1. J. Schaeffer, "Solving the games people play," in Seventh Computer Olympiad. Technical Report CS 02-03, Department of Computer Science, University of Maastricht, 2002, abstract for keynote talk at the Seventh Computer Olympiad.

MACHINE LEARNING

Alberta is a leading centre for machine learning

- iCORE Chairs: Jonathan Schaeffer, Rich Sutton (starting 2003), Ian Witten (Visiting Professor)
- CRC Chair: Dale Schuurmans
- Alberta Ingenuity Centre for Machine Learning: Rob Holte, Russ Greiner, Randy Goebel, Jonathan Schaeffer, Michael Bowling, Rich Sutton, Dale Schuurmans
- U of A Computing Science:

Data mining	Joerg Sander, Mario Nascimento, Osmar Zaiane, Davood Rafei
Pattern recognition	Terry Caelli, Walter Bischof, Vadim Bulitko
Natural language	Andrew Lin, Greg Kondrak, Dekang Lin
Bioinformatics	Duane Szafron and Russ Greiner
Applications	Jonathan Schaeffer, Michael Buro, Martin Mueller, Michael Bowling, Duane Szafron
Robotics	Michael Bowling
Agents	Renee Elio

DATA OVERLOAD PROBLEM

- **In business:** Data volume doubles or triples every year, in decision making contexts (*Globe & Mail*, "How firms can cope with grip of data fear," 21 November 2002)
- **In biotechnology:** "Automated sequencing technology accelerates the pace of input to database, from the current rate of doubling every 20 months. High throughput from cDNA sequencing is expected to double the size of databases in less than one year." National Library of Medicine Centre for Biotechnology Information (NCBI)
- **Of online information:** 2.1 billion publicly accessible pages, 7.3 million adder per day (Cyveillance, July 2000)

WHAT IS MACHINE LEARNING?

Developing efficient and robust algorithms for finding useful patterns in data.
Machine learning refines raw data into useful information.

- Patterns in medical data diagnose diseases from early symptoms predict effectiveness of alternative therapies
- Patterns in manufacturing process data improve process control
- Patterns in atomic probe microscope data improve understanding of molecular structure
- Patterns in human web use improve effectiveness of e-business, improve accuracy of navigation

Computers have learned:

• to do accurate credit card approval	• Humans analysts are 50% accurate; Machine learning was more than 70% accurate
• to dispatch telephone technicians	• BellAtlantic saved $10 million per year
• to optimize parameter settings for separating oil from gas in an oil refinery	• in 10 minutes (human experts require more than one day)
• to catalogue celestial objects (Fayyad et al. 1993)	• automatically discovered 22 new quasars with more than 92% accuracy
• to identify genes (Delcher et al. 1995)	• automatically identifies more than 97% of genes

SCHAEFFER