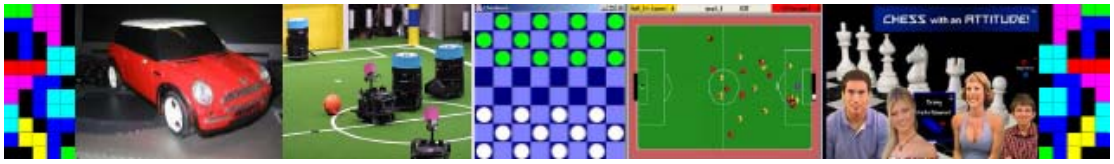

IEEE 2005 Symposium on Computational Intelligence and Games
(CIG'05)

IEEE CIG 2005

April 4-6 2005

Essex University, Colchester, Essex, UK

Programme

Computational intelligence Society

Home of Neural Networks, Fuzzy Systems and Evolutionary Computation

EPSRC

Engineering and Physical Sciences Research Council

Programme Overview

Sunday 3rd April 2005 : Tutorials

12:30 - 14:00	Particle Swarm Optimisation for Learning Game Strategies Andries P. Engelbrecht
14:00 - 14:30	Coffee
14:30 - 16:00	Coevolving game strategies: How to win and how to lose Evan J. Hughes
16:00 - 16:30	Coffee
16:30 - 18:00	Temporal Difference Learning for Game Strategy Acquisition Thomas P. Runarsson

Monday 4th April 2005

08:00 - 08:45:	Registration
08:45 - 09:00:	Opening Remarks
09:00 - 10:00:	**Plenary**: Jordan Pollack (*Is Progress Possible?*)
10:00 - 10:30:	Coffee
10:30 - 12:30:	Session 1: **Lucky Dip** (Chair: Evan Hughes)
12:30 - 14:00:	Lunch
14:00 - 15:30:	Session 2: **Board Games** (Chair: Thomas Runarsson)
15:30 - 16:00:	Coffee
16:00 - 18:00:	Session 3: **Creating Characters** (Chair: Andries Engelbrecht)
18:30 - 20:30:	Poster Session

Tuesday 5th April 2005

09:00 - 10:00:	**Plenary**: Risto Miikkulainen (*Creating Intelligent Agents through Neuroevolution*)
10:00 - 10:30:	Coffee
10:30 - 12:30:	Session 4: **Planning, Searching and Learning** (Chair: Jaap van den Herik)
12:30 - 14:00:	Lunch
14:00 - 15:30:	Session 5: **Real-Time Video Games** (Chair: Sushil Loius)
15:30 - 16:00:	Coffee
16:00 - 17:30:	Session 6: **Computer Go** (Chair: Martin Müller)
	Banquet

Wednesday 6th April 2005

09:00 - 10:00:	**Plenary**: Martin Müller (*Challenges in Computer Go*)
10:00 - 10:30:	Coffee
10:30 - 12:30:	Session 7: **(Co-)evolution 1** (Chair: Jordan Polloack)
12:30 - 14:00:	Lunch
14:00 - 15:30:	Session 8: **(Co-)evolution 2** (Chair: Risto Miikkulainen)
15:30 - 16:00:	Coffee
16:00 - 17:00:	**Plenary**: Jaap van den Herik (*Opponent Modelling and Commercial Games*)
17:00 - 17:15:	Closing remarks

Monday 4th April 2005

09:00 - 10:00: Plenary: Jordan Pollack (*Is Progress Possible?*)

10:30 - 12:30: Session 1: **Lucky Dip (Chair: Evan Hughes)**

- 1044: Utile Coordination: Learning interdependencies among cooperative agents, Jelle R. Kok, Pieter Jan 't Hoen, Bram Bakker, Nikos Vlassis
- 1056: Forcing neurocontrollers to exploit sensory symmetry through hard-wired modularity in the game of Cellz, Julian Togelius, Simon M. Lucas
- 1040: Designing and Implementing e-Market Games, Maria Fasli, Michael Michalakopoulos
- 1028: Dealing with parameterized actions in behavior testing of commercial computer games, J¨org Denzinger, Kevin Loose, Darryl Gates, John Buchanan

14:00 - 15:30: Session 2: Board Games (Chair: Thomas Runarsson)

- 1060: Board Evaluation For The Virus Game, Peter Cowling
- 1035: An Evolutionary Approach to Strategies for the Game of Monopoly®, Colin M. Frayn
- 1061: Further Evolution of a Self-Learning Chess Program, David B. Fogel, Timothy J. Hays, Sarah L. Hahn, James Quon

16:00 - 18:00: Session 3: Creating Characters (Chair: Andries Engelbrecht)

- 1011: Combining coaching and learning to create cooperative character behavior, J¨org Denzinger, Chris Winder
- 1017: Evolving Reactive NPCs for the Real-Time Simulation Game, JinHyuk Hong, Sung-Bae Cho
- 1005: A Generic Approach for Generating Interesting Interactive Pac-Man, Georgios N. Yannakakis and John Hallam
- 1008: Building Reactive Characters for Dynamic Gaming Environments, Peter Blackburn and Barry O'Sullivan

18:30 - 20:30: Poster Session

Tuesday 5th April 2005

09:00 - 10:00: Plenary: Professor Risto Miikkulainen (*Creating Intelligent Agents through Neuroevolution*)

10:30 - 12:30: Session 4: Planning, Searching and Learning (Chair: Jaap van den Herik)

- 1018: Adaptive Strategies of MTD MTD-f for Actual Games, Kazutomo SHIBAHARA, Nobuo INUI, Yoshiyuki KOTANI
- 1051: Monte Carlo Planning in RTS Games, Michael Chung, Michael Buro, and Jonathan Schaeffer
- 1039: Fringe Search: Beating A* at Pathfinding on Game Maps, Yngvi Bjornsson, Markus Enzenberger, Robert C. Holte and Jonathan Schaeffer
- 1046: Adapting Reinforcement Learning for Computer Games: Using Group Utility Functions, Jay Bradley, Gillian Hayes

14:00 - 15:30: Session 5: Real-Time Video Games (Chair: Sushil Loius)

- 1049: Academic AI and Video games: a case study of incorporating innovative academic research into a video game prototype, Aliza Gold
- 1062: Case-Injection Improves Response Time for a Real-Time Strategy Game, Chris Miles, Sushil J. Louis
- 1047: A Hybrid AI System for Agent Adaptation in a First Person Shooter, Michael Burkey, Abdennour El Rhalibi

16:00 - 17:30: Session 6: Computer Go (Chair: Martin Müller)

- 1030: Dynamic Decomposition Search: A Divide and Conquer Approach and its Application to the One-Eye Problem in Go, Akihiro Kishimoto, Martin Muller
- 1022: Combining Tactical Search and Monte-Carlo in the Game of Go, Tristan Cazenave, Bernard Helmstetter
- 1019: Bayesian generation and integration of K-nearest-neighbor patterns for 19x19 go, Bruno Bouzy, Guillaume Chaslot

Wednesday 6th April 2005

09:00 - 10:00: Plenary: Martin Müller (*Challenges in Computer Go*)

10:30 - 12:30: Session 7: (Co-)evolution 1 (Chair: Jordan Polloack)

1013: Evolving Neural Network Agents in the NERO Video Game, Kenneth O. Stanley, Bobby D. Bryant, Risto Miikkulainen

1053: Coevolution in Hierarchical AI for Strategy Games, Daniel Livingstone

1032: Coevolving Probabilistic Game Playing Agents Using Particle Swarm Optimization Algorithms, Evangelos Papacostantis, Andries P. Engelbrecht, Nelis Franken

1058: Evolving a Neural Network Location Evaluator to Play Ms. Pac-Man., Simon Lucas

14:00 - 15:30: Session 8: (Co-)evolution 2 (Chair: Risto Miikkulainen)

1025: Co-evolutionary Strategies for an Alternating-Offer Bargaining Problem, Nanlin Jin, Edward Tsang

1016: A New Framework to Analyze Evolutionary 2 × 2 Symmetric Games, Umberto Cerruti, Mario Giacobini, Ugo Merlone

1010: Synchronous and Asynchronous Network Evolution in a Population of Stubborn Prisoners, Leslie Luthi, Mario Giacobini, Marco Tomassini

16:00 - 17:00: Plenary: Professor Jaap van den Herik (*Opponent Modelling and Commercial Games*)

Paper Abstracts (sorted by paper number)

1003: Poster

A Study of Machine Learning Methods using the Game of Fox and Geese, Kenneth J. Chisholm & Donald Fleming

The game Fox and Geese is solved using retrograde analysis. A neural network trained using a co-evolutionary genetic algorithm with the help of the expert knowledge database was found to be a very capable Fox and Geese player after training, and quickly learned to beat training opponents.

1005: Oral

A Generic Approach for Generating Interesting Interactive Pac-Man, Georgios N. Yannakakis and John Hallam

This paper follows on from our previous work focused on formulating an efficient generic measure of user's satisfaction ('interest') when playing predator/ prey games. Viewing the game from the predators' (i.e. opponents') perspective, a robust on-line neuroevolution learning mechanism has been presented capable of increasing—independently of the initial behavior and playing strategy—the well known Pac-Man game's interest as well as keeping that interest at high levels while the game is being played. This mechanism has also demonstrated high adaptability to changing Pac-Man playing strategies in a relatively simple playing stage. In the work presented here, we attempt to test the on-line learning mechanism over more complex stages and to explore the relation between the interest measure and the topology of the stage. Results show that the interest measure proposed is independent of the stage's complexity and topology, which demonstrates the approach's generality for this game.

1008: Oral

Building Reactive Characters for Dynamic Gaming Environments, Peter Blackburn and Barry O'Sullivan

Interactive computer games are widely seen as a killer application domain for Artificial Intelligence (AI) [8]. Quite apart from the significant size of the games market in terms of revenue [3], computer games provide complex, dynamic, uncertain and competitive environments that are perfect for developing, testing and deploying AI technologies. While many researchers currently focus on enhancing games with sophisticated AI, most overlook the role that AI has to play in the development of the games themselves. In this paper we present an approach to building non-player characters (NPCs) for a well-known computer game, Unreal Tournament. Specifically, we use decision trees induced from human player strategies to define how an NPC in the game performs in a highly dynamic environment. The benefits of this approach are twofold. Firstly, it provides a basis for building competitive AI-based NPCs for interactive computer games. Secondly, this approach eases the development overhead of such characters. Our empirical evaluation demonstrates that the NPCs we create are very competitive against hand-crafted ones in a number of simulated gaming sessions.

1009: Poster

Teams of cognitive agents with leader: how to let them some autonomy, Damien Devigne, Philippe Mathieu, Jean-Christophe Routier

Most often situated multi-agent simulations, of which platform-games are an example, uses reactive agents. This approach has limitations as soon as complex behaviours are desired. For these reasons we propose an approach using cognitive agents. They have knowledge, objectives and are able to build plans in order to achieve their goals and then execute them. In this paper we particularly address the problem of teams of cognitive agents. We chose to build teams directed by a leader. One major problem is the building of the team plan and in particular one difficulty is to find the means in order to let autonomy to the team members. This can be done if the leader builds abstract plans. We present in this article a solution to this problem.

1010: Oral

Synchronous and Asynchronous Network Evolution in a Population of Stubborn Prisoners, Leslie Luthi, Mario Giacobini, Marco Tomassini

We study by computer simulation a population of individuals playing the prisoner's dilemma game. Each player has an invariable strategy (cooperate or defect) but the network of relationships between players is allowed to change over time following simple rules based on players' degree of satisfaction. The population almost always reaches a stable state and we observe that, in the long run, cooperators tend to cluster together in order to maintain a high average payoff and to protect themselves from exploiting defectors. Thus network topology plays an important role even though strategies are not allowed to evolve. We investigated both synchronous and asynchronous network dynamics, observing that asynchronous update, in addition of being more reasonable in a social setting, induces system stability more often than the synchronous one.

1011: Oral

Combining coaching and learning to create cooperative character behavior, J¨org Denzinger, Chris Winder

We present a concept for developing cooperative characters (agents) for computer games that combines coaching by a human with evolutionary learning. The basic idea is to use prototypical situation-action pairs and the nearest-neighbor rule as agent architecture and to let the human coach provide key situations and his/her wishes for an associated action for the different characters. This skeleton strategy for characters (and teams) is then fleshed out by the evolutionary learner to produce the desired behavior. Our experimental evaluation with variants of Pursuit Games shows that already a rather small skeleton –that alone is not a complete strategy– can help solve examples that learning alone has big problems with.

1012: Poster

A Survey on Multiagent Reinforcement Learning Towards Multi-Robot Systems, Erfu Yang, Dongbing Gu

Multiagent reinforcement learning for multirobot systems is a challenging issue in both robotics and artificial intelligence. With the ever increasing interests in theoretical research and practical applications, currently there have been a lot of efforts towards providing some solutions to this challenge. However, there are still many difficulties in scaling up multiagent reinforcement learning to multi-robot systems. The main objective of this paper is to provide a survey on multiagent reinforcement learning in multi-robot systems, based on the literature the authors collected. After reviewing some important advances in this field, some challenging problems are analyzed. A concluding remark is made from the perspectives of the authors.

1013: Oral

Evolving Neural Network Agents in the NERO Video Game, Kenneth O. Stanley, Bobby D. Bryant, Risto Miikkulainen

In most modern video games, character behavior is scripted; no matter how many times the player exploits a weakness, that weakness is never repaired. Yet if game characters could learn through interacting with the player, behavior could improve during gameplay, keeping it interesting. This paper introduces the real-time NeuroEvolution of Augmenting Topologies (rtNEAT) method for evolving increasingly complex artificial neural networks in real time, as a game is being played. The rtNEAT method allows agents to change and improve during the game. In fact, rtNEAT makes possible a new genre of video games in which the player teaches a team of agents through a series of customized training exercises. In order to demonstrate this concept in the NeuroEvolving Robotic Operatives (NERO) game, the player trains a team of robots for combat. This paper describes results from this novel application of machine learning, and also demonstrates how multiple agents can evolve and adapt in video games like NERO in real time using rtNEAT. In the future, rtNEAT may allow new kinds of educational and training applications that adapt online as the user gains new skills.

1016: Oral

A New Framework to Analyze Evolutionary 2 × 2 Symmetric Games, Umberto Cerruti, Mario Giacobini, Ugo Merlone

In this paper we present a new framework to analyze the behavior of evolutionary 2×2 symmetric games. The proposed approach enanbles us to predict the dynamics of the system using the parameters of the game matrix above, without dealing with the concepts of Nash equilibria and evolutionary stable strategies. The predictions are in complete accordance with those that can be made with these latter concepts. Simulations have been performed on populations with spatial structures, and show a good agreement with the model's predictions. We also analyze the dynamics of a particular system, showing how effectively the framework applies to it.

1017: Oral

Evolving Reactive NPCs for the Real-Time Simulation Game, JinHyuk Hong, Sung-Bae Cho

AI in computer games has been highlighted in recent, but manual works for designing the AI cost a great deal. An evolutionary algorithm has developed strategies without using features that are based on the developer. Since the real-time reactive selection of behaviors for NPCs is required for better playing, a reactive behavior system consisting neural networks is presented. Using only the raw information on games, the evolutionary algorithm optimizes the reactive behavior system based on a co-evolutionary method. For demonstration of the proposed method, we have developed a real-time simulation game called 'Build & Build'. As the results, we have obtained emergent and interesting behaviors that are adaptive to the environment, and confirmed the applicability of evolutionary approach to designing NPCs' behaviors without relying on human expertise.

1018: Oral

Adaptive Strategies of MTD MTD-f for Actual Games, Kazutomo SHIBAHARA, Nobuo INUI, Yoshiyuki KOTANI

MTD algorithm developed by Plaat is a variation of SSS* which uses the depth-first strategy to resolve the storage problem coming from the best-first strategy. Since MTD algorithm is based on the zero window search algorithm, the initial range of the searching windows plays an important role in the performance. In this paper, we show some basic experimental results of MTD algorithm. From considerations of the results, the performance of MTD algorithm is expected to be improved by the reduction of the number of times that zero window procedure is called. Then we propose two variations of MTD algorithm. Our purpose is the reduction of searching nodes by the setting of initial window ranges and the way of narrowing the range in the search. Experimental results show 6 percents improvement of searching nodes against MTD-f algorithm.

1019: Oral

Bayesian generation and integration of K-nearest-neighbor patterns for 19x19 go, Bruno Bouzy, Guillaume Chaslot

This paper describes the generation and utilization of a pattern database for 19x19 go with the Knearest-neighbor representation. Patterns are generated by browsing recorded games of professional players. Meanwhile, their matching and playing probabilities are estimated. The database created is then integrated into an existing go program, INDIGO, either as an opening book or as an enrichment of other pre-existing hand-crafted databases used by INDIGO move generator. The improvement brought about by the use of this pattern database is estimated at 15 points on average, which is significant on go standards.

1022: Oral

Combining Tactical Search and Monte-Carlo in the Game of Go, Tristan Cazenave, Bernard Helmstetter

We present a way to integrate search and Monte-Carlo methods in the game of Go. Our program uses search to find the status of tactical goals, builds groups, selects interesting goals, and computes statistics on the realization of tactical goals during the random games. The mean score of the random games where a selected tactical goal has been reached and the mean score of the random games where it has failed are computed. They are used to evaluate the selected goals. Experimental results attest that combining search and Monte-Carlo significantly improves the playing level.

1023: Poster

Pared-down Poker: Cutting to the Core of Command and Control, Kevin Burns

Poker poses cognitive challenges like those of warfare, business and other real world domains. This makes poker a good test bed for basic research on how people make Command and Control decisions and for applied research on how systems might help people make better decisions. In this paper, I compare the cognitive challenges of poker and warfare, and present a new suite of "Pared-down Poker" games that cut to the core of Command and Control. Compared to full-scale poker, Pared-down Poker is more tractable to normative analyses in the lab and more relevant to cognitive challenges in the world. The games have been programmed in Java along with various "animal archetypes" that simulate poker personalities. One game has been used to study the computational effectiveness of cognitive style against normative skill, and the findings from this study highlight questions for further research.

1024: Poster

On TRACS: Dealing with a Deck of Double-sided Cards, Kevin Burns

TRACS (Tool for Research on Adaptive Cognitive Strategies) is a new suite of card games played with a special deck, where the back of each card is a clue to the front of the card. This design simulates the clue/truth structure of real world domains like medicine and warfare, where truths (fronts of cards) must be diagnosed from clues (backs of cards) in order to make decisions (cards to choose, chips to bet, etc.). Here I present the cards and rules of TRACS. I also discuss how the games have been used for computational investigations of memory limits and Bayesian inference. The methods for these studies include human experiments and agent simulations, both of which are facilitated by the unique features of TRACS. The products of TRACS research include a computational model of memory limits and a decision support system for Bayesian inference.

1025: Oral

Co-evolutionary Strategies for an Alternating-Offer Bargaining Problem, Nanlin Jin, Edward Tsang

In this paper, we apply an Evolutionary Algorithm (EA) to solve the Rubinstein's Basic Alternating- Offer Bargaining Problem, and compare our experimental results with its analytic game-theoretic solution. The application of EA employs an alternative set of assumptions on the players' behaviors. Experimental outcomes suggest that the applied co-evolutionary algorithm, one of Evolutionary Algorithms, is able to generate convincing approximations of the theoretic solutions. The major advantages of EA over the game-theoretic analysis are its flexibility and ease of application to variants of Rubinstein Bargaining Problems and complicated bargaining situations for which theoretic solutions are unavailable.

1028: Oral

Dealing with parameterized actions in behavior testing of commercial computer games, J¨org Denzinger, Kevin Loose, Darryl Gates, John Buchanan

We present a method that enhances evolutionary behavior testing of commercial computer games, as introduced in [CD+04], to deal with parameterized actions. The basic idea is to use a layered approach. On one layer, we evolve good parameter value combinations for the parameters of a parameterized action. On the higher layer, we evolve at the same time good action sequences that make use of the value combinations and that try to bring the game in a wanted (or unwanted) state. We used this approach to test corner kicks in the FIFA-99 game. We were able to evolve many parameter-value-action sequence combinations that scored a goal or resulted in a penalty shot, some of which are very short or represent a rather unintelligent behavior of the players guided by the computer opponent.

1030: Oral

Dynamic Decomposition Search: A Divide and Conquer Approach and its Application to the One-Eye Problem in Go, Akihiro Kishimoto, Martin Muller

Decomposition search is a divide and conquer approach that splits a game position into sub-positions and computes the global outcome by combining results of local searches. This approach has been shown to be successful to play endgames in the game of Go. This paper introduces dynamic decomposition search as a way of splitting a problem dynamically during search. Our results in solving one-eye problems in the game of Go show the promise of this approach. Additionally, we propose relaxed decomposition, a more ambitious way of splitting positions.

1032: Oral

Coevolving Probabilistic Game Playing Agents Using Particle Swarm Optimization Algorithms, Evangelos Papacostantis, Andries P. Engelbrecht, Nelis Franken

Coevolutionary techniques in combination with particle swarm optimization algorithms and neural networks have shown to be very successful in finding strong game playing agents for a number of deterministic games. This paper investigates the applicability of a PSO coevolutionary approach to probabilistic games. For the purposes of this paper, a probabilistic variation of the tic-tac-toe game is used. Initially, the technique is applied to a simple deterministic game (tic-tac-toe), proving its effectiveness with such games. The technique is then applied to a probabilistic 4x4x4 tic-tac-toe game, illustrating scalability to more complex, probabilistic games. The performance of the probabilistic game agent is compared against agents that move randomly. To determine how these game agents compete against strong non-random game playing agents, coevolved solutions are also compared against agents that utilize a strong hand-crafted static evaluation function. Particle swarm optimization parameters/topologies and neural network architectures are experimentally optimized for the probabilistic tic-tac-toe game.

1033: Poster

Incrementally Learned Subjectivist Probabilities in Games, Colin Fyfe

In this paper, we show how our AI opponents learn internal representations of probabilities. We use a Bayesian interpretation of such subjectivist probabilities but do not implement full Bayesian methods of parameter estimation since we wish the AIs to be as human-like as possible. Thus the parameters of the subjectivist probabilities are learned incrementally.

1034: Poster

Similarity-based Opponent Modelling using Imperfect Domain Theories, Timo Steffens

This paper proposes a similarity-based approach for opponent modelling in multi-agent games. The classification accuracy is increased by adding derived attributes from imperfect domain theories to the similarity measure. The main contributions are to show how different forms of domain knowledge can be incorporated into similarity measures for opponent modelling, and to show that the situation space of the opponent modelling approach is not required to be the same as the situation space of the opponent players. Our approach has been implemented and evaluated in the domain of simulated soccer.

1035: Oral

An Evolutionary Approach to Strategies for the Game of Monopoly®, Colin M. Frayn

The game of Monopoly® is a turn-based game of chance with a substantial element of skill. Though much of the outcome of any single game is determined by the rolling of dice, an effective trading strategy can make all the difference between an early exit or an overflowing property portfolio. Here I apply the techniques of evolutionary computation in order to evolve the most efficient strategy for property valuation and portfolio management.

1037: Poster

How to Protect Peer-to-Peer Online Games from Cheats, Haruhiro Yoshimoto, Rie Shigetomi

Recently, P2P (peer-to-peer) online game systems have attracted a great deal of public attention. They work without central servers, thus, the maintenance and organization costs have been drastically reduced. However, in P2P systems, it is difficult for game creators to prevent cheats by malicious players, due to the lack of trusted servers. In order to solve the problem, we propose a practical and secure protocol based on public key cryptography, suitable for such P2P online game systems. Our scheme guarantees that players can immediately detect when cheating by other malicious players happens, and that honest players can prove their innocence, on condition that more than half of the participants are honest. We categorized cheating into four groups, Crack-The-Game software attack, hange-The-Input-After-Communication attack, Forge-The-Result attack and Be-Offline-When-Losing attack, and proved that our system is secure against each attack. Moreover, we showed that our scheme is general and directly applicable to almost all of existing online games including simulation games, fighting games, and MMORPG.

1038: Poster

Training an AI player to play Pong using a GTM, Gayle Leen, Colin Fyfe

We extend the work of [McGlinchey 2003], in which the author trained an AI player to play Pong from game observation data recorded from games played by humans. The data trained a Self Organising Map (SOM) , and it was found that the AI player played Pong with a human style of play. However one of the drawback of using the SOM was that the movement of the bat was jerky, due to quantisation of vectors. The author had to applying smoothing to the AI player's bat to make the movement more realistic. It was also found that the AI players were easy to beat. In this paper we train the AI player using Generative Topographic Mapping (GTM) and we show that using the mean of the conditional density to estimate the bat's position is better than using the mode.

1039: Oral

Fringe Search: Beating A* at Pathfinding on Game Maps, Yngvi Bjornsson, Markus Enzenberger, Robert C. Holte and Jonathan Schaeffer

The A* algorithm is the de facto standard used for pathfinding search. IDA* is a space-efficient version of A*, but it suffers from cycles in the search space (the price for using no storage), repeated visits to states (the overhead of iterative deepening), and a simplistic leftto- right traversal of the search tree. In this paper, the Fringe Search algorithm is introduced, a new algorithm inspired by the problem of eliminating the inefficiencies with IDA*. At one extreme, the Fringe Search algorithm expands frontier nodes in the exact same order as IDA*. At the other extreme, it can be made to expand them in the exact same order as A*. Experimental results show that Fringe Search runs roughly 10-40% faster than highly-optimized A* in our application domain of pathfinding on a grid.

1040: Oral

Designing and Implementing e-Market Games, Maria Fasli, Michael Michalakopoulos

Trading in electronic markets has been the focus of intense research over the last few years within areas of Artificial Intelligence and Economics. This paper discusses the need for tools to support the design and implementation of electronic market games. Such games simulating real life problems can be used in order to conduct research on mechanism design, markets and strategic behaviour. To this end e present the e-Game platform which was developed to support the design, implementation and execution of market game scenarios involving auctions. How game development is aided is demonstrated with a simple game.

1044: Oral

Utile Coordination: Learning interdependencies among cooperative agents, Jelle R. Kok, Pieter Jan ’t Hoen, Bram Bakker, Nikos Vlassis

We describe Utile Coordination, an algorithm that allows a multiagent system to learn where and how to coordinate. The method starts with uncoordinated learners and maintains statistics on expected returns. Coordination dependencies are dynamically added if the statistics indicate a statistically significant benefit. This results in a compact state representation because only necessary coordination is modeled. We apply our method within the framework of coordination graphs in which value rules represent the coordination dependencies betweenthe agents for a specific context. The algorithm is first applied on a small illustrative problem, and next on a large predator-prey problem in which two predators have to capture a single prey.

1046: Oral

Adapting Reinforcement Learning for Computer Games: Using Group Utility Functions, Jay Bradley, Gillian Hayes

Group utility functions are an extension of the common team utility function for providing multiple agents with a common reinforcement learning signal for learning cooperative behaviour. In this paper we describe what group utility functions are and suggest using them to provide non-player computer game character behaviours. As yet, reinforcement learning techniques have very rarely been used for computer game character specication. Here we show the results of using a group utility function to learn an equilibrium between two computer game characters and compare this against the performance of the two agents learning independently. We also explain how group utility functions could be applied to learn equilibria between groups of agents. We highlight some implementation issues arising from using a commercial computer game engine for multi-agent reinforcement learning experiments.

1047: Oral

A Hybrid AI System for Agent Adaptation in a First Person Shooter, Michael Burkey, Abdennour El Rhalibi

The aim of developing an agent that is able to adapt its actions in response to their effectiveness within the game provides the basis for the research presented in this paper. It investigates how adaptation can be applied through the use of a hybrid of AI technologies. The system developed uses the pre-defined behaviours of a finite state machine and fuzzy logic system combined with the learning capabilities of a neural network. The system adapts specific behaviours that are central to the performance of the bot in the game, with the main focus being on the weapon selection behaviour; selecting the best weapon for the current situation. As a development platform, the project makes use of the Quake 3 Arena engine, modifying the original bot AI to integrate the adaptive technologies.

1049: Oral

Academic AI and Video games: a case study of incorporating innovative academic research into a video game prototype, Aliza Gold

Artificial intelligence research and video games are a natural match, and academia is a fertile place to blend game production and academic research. Game development tools and processes are valuable for applied AI research projects, and university departments can create opportunities for student-led, team-based project work that draws on students' interest i n video games. The DigitalMedia Collaboratory at the University of Texas at Austin has developed a project in which academic AI research was incorporated into a video game production process that is repeatable in other universities. This process has yielded results that advance the field of machine learning as well as the state of the art in video games. This is a case study of the process and the project that originated it, outlining methods, results, and benefits in order to encourage the use of the model elsewhere.

1051: Oral

Monte Carlo Planning in RTS Games, Michael Chung, Michael Buro, and Jonathan Schaeffer

Monte Carlo simulations have been successfully used in classic turn–based games such as backgammon, bridge, poker, and Scrabble. In this paper, we apply the ideas to the problem of planning in games with imperfect information, stochasticity, and simultaneous moves. The domain we consider is real–time strategy games. We present a framework—MCPlan—for Monte Carlo planning, identify its performance parameters, and analyze the results of an implementation in a capture–the–flag game.

1053: Oral

Coevolution in Hierarchical AI for Strategy Games, Daniel Livingstone

Real-Time Strategy games present an interesting problem domain for Artificial Intelligence research. We review current approaches to developing AI systems for such games, noting the frequent decomposition into hierarchies similar to those found in real-world armies. We also note the rarity of any form of learning in this domain – and find limitations in the work that does use learning. Such work tends to enable learning at only one level of the AI hierarchy. We argue, using examples from real-world wars and from research on coevolution in evolutionary computation, that learning in AI hierarchies should occur concurrently at the different strategic and tactical levels present. We then present a framework for conducting research on coevolving the AI

1055: Poster

Nannon™: A Nano Backgammon for Machine Learning Research, Jordan B. Pollack

A newly designed game is introduced, which feels like Backgammon, but has a simplified rule set. Unlike earlier attempts at simplifying the game, Nannon maintains enough features and dynamics of the game to be a good model for studying why certain machine learning systems worked so well on Backgammon. As a model, it should illuminate the relationship between different methods of learning, both symbolic and numeric, including techniques such as inductive inference, neural networks, genetic programming, co-evolutionary learning, and reinforcement learning based on value function approximation. It is also fun to play.

1056: Oral

Forcing neurocontrollers to exploit sensory symmetry through hard-wired modularity in the game of Cellz, Julian Togelius, Simon M. Lucas

Several attempts have been made in the past to construct encoding schemes that allow modularity to emerge in evolving systems, but success is limited. We believe that in order to create successful and scalable encodings for emerging modularity, we first need to explore the benefits of different types of modularity by hard-wiring these into evolvable systems. In this paper we explore different ways of exploiting sensory symmetry inherent in the agent in the simple game Cellz by evolving symmetrically identical modules. It is concluded that significant increases in both speed of evolution and final fitness can be achieved relative to monolithic controllers. Furthermore, we show that simple function approximation task that exhibits sensory symmetry can be used as a quick approximate measure of the utility of an encoding scheme for the more complex game-playing task.

1058: Oral

Evolving a Neural Network Location Evaluator to Play Ms. Pac-Man., Simon Lucas

Ms. Pac-Man is a challenging, classic arcade game with a certain cult status. This paper reports attempts to evolve a Pac-Man player, where the control algorithm uses a neural network to evaluate the possible next moves. The evolved neural network takes a handcrafted feature vector based on a candidate maze location as input, and produces a score for that location as output. Results are reported on two simulated version of the game: deterministic and non-deterministic. The results show that useful behaviours can be evolved that are frequently capable of clearing the first level, but are still susceptible to making poor decisions. Currently, the best evolved players play at the level of a reasonable human novice.

1060: Oral

Board Evaluation For The Virus Game, Peter Cowling

The Virus Game (or simply Virus) is a turnbased two player perfect information game which is based on the growth and spread of competing viruses. This paper describes a CPU efficient and easy to use architecture for developing and testing AI for Virus and similar games and for running a tournament between AI players. We investigate move generation, board representation and tree search for the Virus Game and discuss a range of parameters for evaluating the likely winner from a given position. We describe the use of our architecture as a tool for teaching AI, and describe some of the AI players developed by students using the architecture. We discuss the relative performance of these players and an effective, generalisable scheme for ranking players based on similar ideas to the Google PageRank method.

1061: Oral

Further Evolution of a Self-Learning Chess Program, David B. Fogel, Timothy J. Hays, Sarah L. Hahn, James Quon

Previous research on the use of coevolution to improve a baseline chess program demonstrated a performance rating of 2550 against Pocket Fritz 2.0 (PF2). A series of 12 games (6 white, 6 black) was played against PF2 using the best chess program that resulted from 50 generations of variation and selection in self-play. The results yielded 9 wins, 2 losses, and 1 draw for the evolved program. This paper reports on further evolution of the best-evolved chess program, executed through 7462 generations. Results show that the outcome of this subsequent evolution was statistically significantly better than the prior player from 50 generations. A 16-game series against PF2, which plays with the rating of a high-level master, resulted in 13 wins, 0 losses, and 3 draws, yielding a performance rating of approximately 2650.

1062: Oral

Case-Injection Improves Response Time for a Real-Time Strategy Game, Chris Miles, Sushil J. Louis

We present a case-injected genetic algorithm player for Strike Ops, a real-time strategy game. Such strategy games are fundamentally resource allocation optimization problems and our previous work showed that genetic algorithms can play such games by solving the underlying resource allocation problem. This paper shows how we can learn to better respond to opponent actions (moves) by using case-injected genetic algorithms. Case-injected genetic algorithms were designed to learn to improve performance in solving sequences of similar problems and thus provide a good fit for responding to opponent actions in Strike Ops which result in a sequence of similar resource allocation problems. Our results show that a case-injected genetic algorithm player learns from previously encountered problems in the sequence to provide better quality solutions in less time for the underlying resource allocation problem thus improving response time by the genetic algorithm player. This improves the responsiveness of the game and the quality of the overall playing experience.