

INTERACTIVE OPPONENTS GENERATE INTERESTING GAMES

Georgios N. Yannakakis and John Hallam
The Maersk Institute for Production Technology
University of Southern Denmark
Campusvej 55, 5230 Odense M, Denmark
E-mail: {georgios;john}@mip.sdu.dk

KEYWORDS

On-line Learning in Computer Games, Prey/Predator, Interactive Opponents, Interest Emergence.

ABSTRACT

In this paper we present experiments on neuro-evolution mechanisms applied to predator/prey multi-character computer games. Our test-bed is a computer game where the prey (i.e. player) has to avoid its predators by escaping through an exit without getting killed. By viewing the game from the predators' (i.e. opponents') perspective, we attempt off-line to evolve neural-controlled opponents, whose communication is based on partial implicit information, capable of playing effectively against computer-guided fixed strategy players. However, emergent near-optimal behaviors make the game less interesting to play. We therefore discuss the criteria that make a game interesting and, furthermore, we introduce a generic measure of this category of (i.e. predator/prey) computer games' interest (i.e. player's satisfaction from the game). Given this measure, we present an evolutionary mechanism for opponents that keep learning from a player while playing against it (i.e. on-line) and we demonstrate its efficiency and robustness in increasing and maintaining the game's interest. Computer game opponents following this on-line learning approach show high adaptability to changing player strategies which provides evidence for the approach's effectiveness against human players.

INTRODUCTION

In (Yannakakis et al. 2004), we introduced a predator/prey computer game named 'Dead End' for emerging complex and cooperative behaviors among agents through evolutionary procedures. In

this game the prey (i.e. player) has to avoid its eight predators (i.e. *Dogs*) by escaping through an exit without getting killed. Since there are eight *Dogs* on the game field, they are designed to be slower than the Player so that the game is fairer to play. This game's fundamental concepts are inspired from previous work of Yannakakis et al. (2003) where efficient cooperative behaviors, supported only by partial implicit communication, emerge amongst the agents of a complex multi-agent environment.

Similar to Luke's and Spector's (1996) work on the *Serengeti* world, we view Dead End from the predators' perspective. Our first aim is to emerge effective complex teamwork behaviors by the use of an off-line training approach, based on evolutionary computation techniques, applied to homogeneous neural controlled agents (Yao 1999). *Dogs* have to demonstrate good cooperative strategies in order to kill the Player and/or to defend the Exit. Such behaviors can be aggressive, defensive, or a hybrid of the two. Given the specific game, we believe that 8 predators are enough for cooperative behaviors to emerge.

However, playing a computer game like Dead End against well-playing opponents with fixed hunting behaviors cannot be regarded as interesting. The first stage of experiments on this test-bed, given an implicitly defined notion of interest, is presented in (Yannakakis et al. 2004). We believe that the interest of any computer game is directly related to the interest generated by the opponents' behavior rather than to the graphics or even the player's behavior. Thus, when 'interesting game' is mentioned we mainly refer to interesting opponents to play against. In (Yannakakis and Hallam 2004), we argue that the interest measure proposed (for the well-known Pac-Man game) defines a generic measure of any predator/prey game. Results obtained from Dead End and

presented here give evidence for this interest measure's generality, which defines one of the goals of this work.

We present a robust on-line neuro-evolution learning mechanism capable of increasing the game's interest (starting from well performing behaviors trained off-line) as well as maintaining that interest at high levels as long as the game is being played. In our Dead End predator/prey computer game we require *Dogs* to keep learning and constantly adapting to the player's strategy instead of being opponents with fixed strategies. In addition, we explore learning procedures that achieve good real-time performance (i.e. low computational effort while playing).

Recently, there have been attempts to mimic human behavior off-line, from samples of human playing, in specific virtual environments. In (Thureau et al. 2004) among others, human-like opponent behaviors are emerged through supervised learning techniques in a first person shooter console game. Even though complex opponent behaviors are emerged, there is no further analysis on whether these behaviors contribute to the satisfaction of the player (i.e. interest of game). In other words, researchers hypothesize --- by looking at the vast number of multi-player on-line games played daily on the web --- that by generating human-like opponents the player gains more satisfaction from the game. This hypothesis might be true up to a point; however, since there is no explicit notion of interest defined, there is no evidence that a specific opponent behavior generates more or less interesting games.

DEAD-END GAME

Dead End is a two-dimensional, multi-agent, grid-motion, predator/prey game. The game field (i.e. stage) is a two-dimensional square world that contains a white rectangular area named "Exit" (see Fig. 1) at the top. For the experiments presented in this paper we use the 16 X 16 cm stage presented in Fig. 1, which is divided into grid squares (of length 0.5 mm). The characters visualized in the Dead End game (as illustrated in Fig. 1) are a dark grey circle of radius 0.75 cm representing the Player and 8 light grey square (of

dimension 1.5 cm) characters representing the *Dogs*.

The relationship between the *Dogs* and the Player is mutually highly competitive. The aim of a Player is to reach the Exit, avoiding the *Dogs*. On the other hand, the aims of the *Dogs* are to defend the Exit and/or to catch the Player. In Dead End, if a Player succeeds in arriving at the Exit, this event is described as a *win*. Additionally, if a *Dog* manages to catch a Player, this event defines a *kill*. If there is neither a Player win nor a kill for a predetermined large period of time, then the outcome of the game is a *win* again. After either a win or a kill, a new game starts.

The Player moves at four thirds the *Dogs*' maximum speed and since there are no dead ends, it is impossible for a single *Dog* to complete the task of killing it. Since the Player moves faster than a *Dog*, the only effective way to kill the Player is for a group of *Dogs* to hunt cooperatively.

The simulation procedure of the Dead End game is as follows. Player and *Dogs* are placed in the game field (initial positions) so that there is a suitably

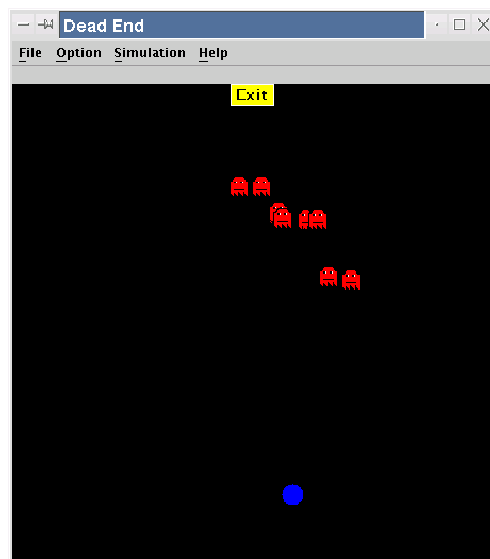


Fig. 1. Snapshot of the Dead End game

large distance between them. Then, the following occur at each simulation step. **(a)** Both *Dogs* and the Player gather information from their environment and take an individual movement decision, up, down, left or right. **(b)** If the game is

over (i.e. Player escapes through the Exit, Player is killed, or the simulation step is greater than a predetermined large number), then a new game starts from the same initial positions for the *Dogs* but from a different, randomly chosen, position at the bottom of the stage for the Player.

The Player

The difficulty of the Dead End game is directly affected by the intelligence of the Player. Its nature is significant because *Dogs*' emergent behavior is strongly related to their competitive relationship against it. To develop more diverse agents' behaviors, different playing strategies are required. We therefore chose three fixed *Dog*-avoidance and/or Exit-achieving strategies for the Player, differing in complexity and effectiveness. The non-deterministic initial position of the player is devised to provide *Dogs* with diverse examples of playing behaviors to learn from.

Randomly-moving (RM) Player

A Randomly-moving Player takes a movement decision by selecting a uniformly distributed random picked direction at each simulation step of the game.

Exit-achieving (EA) Player

An Exit-achieving Player moves directly towards the Exit. Its strategy is based on moving so as to reduce the greatest of its relative distances from the Exit.

Cost-based path planning (CB) Player

A cost-based path planning Player constitutes the most efficient *Dog*-avoiding and Exit-achieving strategy of the three different fixed-strategy types of Player. A discrete Artificial Potential Field (APF) (Khatib 1986), specially designed for the Dead End game, controls the CB Player's motion. The overall APF causes a force to act on the Player which guides it along a *Dog*-avoidance Exit-achievement path. For a more detailed presentation of the CB player, see (Yannakakis et al. 2004).

Any motion strategy that guides a Player to arrive quickly at the Exit, avoiding any *Dogs* and keeping to the straightest and fastest possible trajectory, is definitely a "good" strategy in terms of the Dead End game. Hence, the CB Player presents a "good"

behavior in this game and furthermore a reference case to compare to human playing behavior.

Neural Controlled Dogs

Artificial neural networks (ANNs) are a suitable host for emergent adaptive behaviors in complex multi-agent environments (Ackley and Littman 1992). A feedforward neural controller is employed to manage the *Dogs*' motion and is described in this subsection.

Using their sensors, *Dogs* inspect the environment from their own point of view and decide their next action. Each *Dog* receives input information from its environment expressed in the ANN's input array of dimension 6. The input array consists of the relative coordinates of (a) the Player, (b) the closest *Dog* and (c) the Exit. A *Dog*'s input includes information for only one neighbor *Dog* as this constitutes the minimal information for emerging teamwork cooperative behaviors. We deliberately exclude from consideration any global sensing, e.g. information about the dispersion of the *Dogs* as a whole, because we are interested specifically in the minimal sensing scenario.

As previously mentioned, a multi-layered fully connected feedforward ANN has been used for the experiments presented here. The hyperbolic tangent sigmoid function is employed at each neuron. The ANN's output is a two dimensional vector which represents the *Dog*'s chosen motion in X, Y coordinates.

Fixed strategy Dogs

Apart from the neural controlled *Dogs*, an additional fixed non-evolving strategy has been tested for controlling the *Dogs*' motion. *Dogs* of this strategy are called 'Followers' and they are designed to follow the Player constantly by moving at half the Player's speed (i.e. 1.0 cm/simulation step). This strategy is used as a baseline behavior for comparison with any emergent neural controller behavior.

INTERESTING OPPONENTS

In order to find, as objective as possible, a measure of interest in the Dead End computer game we first

need to define the criteria that make a game interesting. Then, second, we need to quantify and combine all these criteria in a mathematical formula. The game should then be tested by human players and have this formulation of interest cross validated against the interest the game produces in real conditions. This last part of our investigation constitutes a crucial phase of future work.

In order to simplify this procedure we will ignore the graphics' as well as the player's contribution to the interest of the game and we will concentrate on the *Dogs*' behavior that effects the game's interest. That is because, we believe, the computer-guided opponent character contributes the vast majority of features that make a computer game interesting.

By being as objective and generic as possible, we believe that the criteria that collectively define the interest of the Dead End game are as follows (see also (Yannakakis and Hallam 2004) for interest criteria definitions for the Pac-Man game).

- *When the game is neither too hard nor too easy.* In other words, the game is interesting when *Dogs* manage to kill the player sometimes but not always. In that sense, optimal behaviors are not interesting behaviors and *vice versa*.
- *When there is diversity in Dogs' behavior over the games.* That is, when *Dogs* are able to find different ways of hunting and killing the player in each game so that their strategy is less predictable.
- *When Dogs' behavior is aggressive rather than static.* That is, *Dogs* that move towards killing the player but meanwhile, move constantly all over the game field instead of simply following it. This behavior gives player the impression of an intelligent strategic *Dogs*' plan which increases the game interest.

In order to estimate and quantify each of the aforementioned criteria of the game's interest, we follow the same procedure introduced in (Yannakakis and Hallam 2004). Thus, the metrics for the three criteria are given by T (difference between maximum and average player's lifetime

over N games --- N is 50 in this paper), S (standard deviation of player's lifetime over N games) and $E\{H_n\}$ (stage grid-cell visit average entropy of the *Dogs* over N games) respectively. All three metrics are combined linearly (1)

$$I = \frac{\gamma T + \delta S + \varepsilon E\{H_n\}}{\gamma + \delta + \varepsilon} \quad (1)$$

where I is the interest value of the Dead End game; γ, δ and ε are criterion weight parameters (for the experiments presented here $\gamma = 1, \delta = 2, \varepsilon = 1$).

The measure of the Dead End game's interest introduced in (1) can be effectively applied to any predator/prey computer game (e.g. see (Yannakakis and Hallam 2004)) for a successful application on the Pac-Man game) because it is based on generic quantitative features of this category of games. These features include the time required to kill the prey as well as the predators' entropy throughout the game field. We therefore believe that (1) --- or a similar measure of the same concepts --- constitutes a generic interest approximation of predator/prey computer games. In fact, the two first criteria correspond to any computer game whereas the third criterion corresponds only to predator/prey games.

OFF-LINE LEARNING

We use an off-line evolutionary learning approach in order to produce some 'good' (i.e. in terms of performance) initial behaviors for the on-line learning mechanism. The ANNs that determine the behavior of the *Dogs* are themselves evolved (evolutionary process is limited to the connection weights of the ANN).

The evolutionary procedure is as follows. Each *Dog* has a genome that encodes the connection weights of its ANN. A population of 40 (we keep this number low because of the computational cost) ANNs (*Dogs*) is initialized randomly with initial uniformly distributed random connection weights that lie within $[-5, 5]$. Then, at each generation: **(a)** Each *Dog* in the population is cloned 8 times. These 8 clones are placed in the Dead End game field and play the game against a selected Player type for an evaluation period T

(e.g. 125 simulation steps). The outcome of this game is to ascertain the total number of wins (W) and kills (K). **(b)** Each *Dog* is evaluated via (2)

$$f = \alpha K - \beta W \quad (2)$$

where K and W are the total numbers of kills and wins respectively; α is the reward rate of a kill; β is the penalty rate of a win. **(c)** A pure elitism selection method is used where only the 20% fittest solutions are able to breed and, therefore, determine the members of the intermediate population. **(d)** Each parent clones an equal number of offspring in order to replace the non-picked solutions from elitism. **(e)** Mutation occurs in each gene (connection weight) of each offspring's genome with a small probability p_m (e.g. 0.01). A uniform random distribution is used again to define the mutated value of the connection weight.

The algorithm is terminated when a predetermined number of generations g is completed (e.g. $g=300$) and the fittest *Dog*'s connection weights are saved.

ON-LINE LEARNING

This evolutionary learning approach is based on the idea of *Dogs* that learn while they are playing against the Player. In other words, *Dogs* that are reactive to any player's behavior and learn from its strategy instead of being predictable and, therefore, uninteresting characters for game playing. Furthermore, this approach's additional objective is to keep the game's interest at high levels as long as it is being played.

Beginning from any initial off-line trained (OLT) group of homogeneous *Dogs*, the on-line learning (OLL) mechanism attempts to transform them into a group of heterogeneous *Dogs* that are interesting to play against. The OLL procedure is as follows. An OLT *Dog* is cloned 8 times and its clones are placed in the Dead End game field to play against a selected Player type. Then, at each generation:

(a) Each *Dog* is evaluated every T (T is 25 here) simulation steps via (3), while the game is played (where (x_d^k, y_d^k) and (x_p^k, y_p^k) are the cartesian

coordinates of the Player's and the *Dog*'s center respectively at simulation step k).

$$f' = \frac{1}{1 + \sum_{k=1}^T \{|x_d^k - x_p^k| + |y_d^k - y_p^k|\}} \quad (3)$$

By using (3), we individually promote each *Dog* that attempts to stay as close as possible to the Player during an evaluation period. **(b)** If the average fitness of the population is greater than a fixed threshold value then, go to **(a)** else, continue. **(c)** A pure elitism selection method is used where only the fittest solution is able to breed. The fittest parent clones an offspring that replaces the worst-fit member of the population. This offspring takes the worst-fit member's position in the game field. **(d)** Mutation occurs in each gene (connection weight) of the offspring's genome exactly as in the off-line learning algorithm.

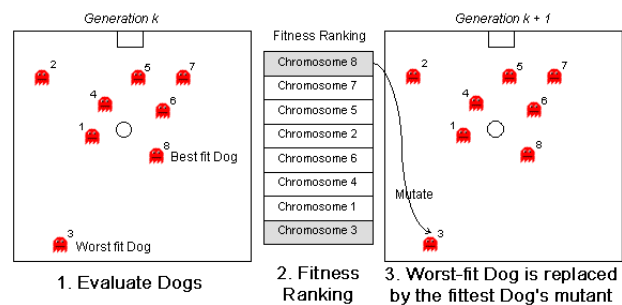


Fig. 2. The on-line learning mechanism

The algorithm is terminated when a predetermined number of generations g is completed (e.g. $g=5000$) and all 8 *Dogs*' connection weights are saved. Fig. (2) illustrates the main steps of the OLL algorithm.

We mainly use small simulation periods (i.e. $T=25$) to evaluate *Dogs* during OLL. The aim of this high frequency of evaluations is to accelerate the on-line evolutionary process. However, the evaluation function (3) constitutes an approximation of the examined *Dog*'s overall performance for large simulation periods. Keeping the right balance between computational effort and performance approximation is one of the key features of this approach. We therefore use

minimal evaluation periods capable of achieving good estimation of the *Dogs*' performance.

RESULTS

Results obtained from experiments applied on the Dead End game are presented in this section. These include, off-line and on-line learning emergent behavior analysis as well as experiments for testing robustness and adaptability of the OLL mechanism proposed.

Performance Measurement

In order to evaluate the performance of a team of *Dogs*, we record the total number of both kills K and wins W of the examined team, against a specific Player, by placing these agents in Dead End and letting them play the game for $12.5 \cdot 10^3$ simulation steps. We believe that this is a long enough period for testing a playing-behavior of a team of *Dogs* in an efficient way. This evaluation is called a *trial*. We then calculate the value $P = 100[K/(K+W)]$. This performance measurement (P) quantifies the Player-killing (K) percentage over the total number of games played ($K+W$).

Off-line Learning Experiments

The experiment presented in this subsection is focused on producing well-behaved *Dogs* in terms of the performance measure previously described. We train *Dogs* against all three fixed-strategy types of Player through the off-line learning mechanism. In this experiment we select $\alpha = \beta = 1$ in fitness function (2) --- providing equal opportunities for promoting both Player-hunting and Exit-defensive behaviors. The off-line learning experiment is described as follows.

(a) Apply the off-line learning mechanism by playing against each type of Player separately. Repeat the learning attempt (run) 10 times --- we believe that this number is adequate to illustrate a clear picture of the emergent behavior --- with different initial conditions. (b) Evaluate each one of the 10 teams of OLT *Dogs* against all three types of Player. Their performance and interest measurement are given by the average values obtained over the 10 trials. (c) Evaluate non-evolving randomly generated (i.e. untrained) as

well as Player-follower *Dogs* (i.e. Followers) against every Player type (run 10 trials and calculate their average performance and interest). The outcome of this experiment is presented in Table I.

Table I. The effect of off-line training on the *Dogs*' average performance ($E\{P\}$) and interest ($E\{I\}$) over 10 learning attempts

	Playing against					
	RM		EA		CB	
	$E\{P\}$	$E\{I\}$	$E\{P\}$	$E\{I\}$	$E\{P\}$	$E\{I\}$
OLT/RM	91.27	0.728	24.36	0.682	3.82	0.243
OLT/EA	62.55	0.555	96.01	0.661	51.27	0.486
OLT/CB	93.09	0.628	55.09	0.681	72.98	0.425
Followers	98.54	0.466	78.94	0.763	71.51	0.709
Untrained	75.58	0.401	62.46	0.498	17.77	0.425

As can be seen from Table I, there is a large performance improvement of the OLT *Dogs* in comparison to the untrained or even the Follower *Dogs* against all three types of Player. However, in most cases, OLT *Dogs* against a specific Player seem to get lower average performance values when playing against a Player other than the Player they have been off-line trained against. *Dogs* trained off-line against CB Players showed good overall performance against all types of Players. Therefore, among the three fixed-strategy Players, the CB Player provides the best off-line training for the opponent agents. This suggests that when *Dogs* learn from more complex and effective types of Players, they tend to generalize better.

An increased interest value when *Dogs* are trained off-line is also noticeable in all cases (see Table I). However, these emergent behaviors fail to compete the interest generated by the Followers in the majority of cases (mainly against the EA and CB Players).

The most typical emergent behaviors are pure Exit-defensive or pure Player-hunting behaviors but hybrids also occur frequently. The off-line learning mechanism, in the majority of cases, produces *Dogs* that defend the Exit and/or hunt the Player in a cooperative fashion. As stressed before, opponents in this game have to learn to cooperate in order to be successful (achieve a high performance value) against any playing strategy.

On-line Learning Experiments

The off-line learning procedure is a mechanism that attempts to produce near-optimal solutions to the problem of killing the Player and defending the Exit. These solutions will be the OLL mechanisms' initial points in the search for more interesting games. The OLL experiment is described as follows.

(a) Apply the OLL mechanism to all teams of OLT *Dogs* (see Off-line Learning Experiments section) playing against each type of Player separately. (b) Evaluate performance and interest values of each OLL attempt against each Player type. The outcome of this experiment is presented in Table II and Fig. 3.

As seen from Table I and Table II, the OLL mechanism manages to find ways of increasing the interest of the game regardless of the initial OLT behavior or the player. Due to space considerations we present only 3 out of the 9 OLL experiments in detail here. Fig. 3 demonstrates the learning mechanism's ability of producing games of higher than the initial interest as well as keeping that high interest for a long period. The mechanism demonstrated a similar adaptive behavior for all 9 different OLL experiments. This suggests that the evolutionary approach proposed shows a behavior of high robustness which furthermore manages to generate opponents' behaviors of much higher interest values.

The OLL mechanism tends to be a highly disruptive procedure (via the mutation operation) for high-interest group behaviors towards individual rewards. Such disruptive mutations can cause undesired drops in the game's interest generated by a team of *Dogs*. However, experiments show that *Dogs* trained by individual rewards (while playing) manage to maintain and even increase the game's interest.

Another important feature of the mechanism is its ability to quickly emerge interesting opponents to play against. It takes, in the worst case experienced, fewer than 500 OLL games for the mechanism to generate games of higher interest.

Table II. Best average interest values achieved by applying on-line learning on *Dogs* trained off-line. The respective average performance values are also presented

	Playing against – On-line learning					
	RM		EA		CB	
	$E\{P\}$	$E\{I\}$	$E\{P\}$	$E\{I\}$	$E\{P\}$	$E\{I\}$
OLT/RM	86.73	0.758	36.45	0.762	43.09	0.721
OLT/EA	95.64	0.707	84.18	0.701	20.91	0.617
OLT/CB	97.09	0.685	53.64	0.745	60.92	0.610




Fig. 3. Interest (averaging over 10 trials) evolution over the number of games played. Initial behavior: OLT/RM (initial and best interest values are presented in the first row of Table I and Table II respectively).

On the other hand (see Table I and Table II), in almost half cases, there is a decrease of the *Dogs*' average performance values. In general, *Dogs* that achieve high-performance values do not generate interesting games. This illustrates the tradeoff between optimality and interest in any computer game. In Dead End, optimal killing behaviors cannot produce interesting games.

CONCLUSIONS

The Dead End predator/prey computer game is devised as an interesting test-bed for studying the emergence of multi-agent cooperative behaviors supported by partial implicit communication through evolutionary learning mechanisms. We introduced an off-line learning mechanism, from which effective cooperative predator behaviors have rapidly emerged.

Predator strategies in predator/prey computer games are still nowadays based on simple rules which make the game quite predictable and, therefore, uninteresting --- by the time the player

gains more experience and playing skills. A computer game becomes interesting primarily when there is an on-line interaction between the player and his opponents who demonstrate interesting behaviors.

Given some objective criteria for defining interest in predator/prey games presented by Yannakakis and Hallam (2004) we introduced a method for explicitly measuring interest in the Dead End game. We saw that by using the proposed on-line learning mechanism (see also (Yannakakis et al. 2004)), maximization of the individual simple distance measure (see (3)) coincides with maximization of the game's interest. Apart from being robust, the proposed mechanism demonstrates fast adaptability to new types of player (i.e. playing strategies). Therefore, we believe that such a mechanism will be able to produce interesting interactive opponents (i.e. games) against even the most complex human playing strategy.

We believe that the methods used need to be tested on more complex Dead-End stages (i.e less *Dogs*) in order to provide more evidence for their generality, and the interest measure proposed needs to be cross-validated against human players. In addition, investigation of the heterogeneity's contribution on these results constitutes an important step for future work.

REFERENCES

Ackley D.H. and M.L. Littman. 1992. "Interactions between learning and Evolution." In *Artificial Life II*, C.G. Langton, C.Taylor, J.D. Farmer, and S. Rasmussen, eds. Sante Fe Institute Studies in the Sciences and Complexity. Reading, MA: Addison-Wesley, 478-507.

Khatib O. 1986. "Real-time obstacle avoidance for manipulators and mobile robots." *International Journal for Robotics Research*, vol. 5, no. 1, 90-99.

Luke S. and L. Spector. 1996. "Evolving teamwork and coordination with genetic programming." In *Genetic Programming 1996: Proceedings of the First Annual Conference*, J.R. Koza, D.E. Goldberg, D.B. Fogel, and R.L. Riolo, eds. Stanford University, CA, USA: MIT Press, 150-156.

Thureau C., C. Bauckhage and G. Sagerer. 2004. "Learning human-like Movement Behavior for Computer Games" In *From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior (SAB-04)*, Los Angeles, USA, July 13-16, 315-323.

Yannakakis G.N. and J. Hallam. 2004. "Evolving Opponents for Interesting Interactive Computer Games." In *From Animals to Animats 8: Proceedings of the 8th International Conference on Simulation of Adaptive Behavior (SAB-04)*, Los Angeles, USA, July 13-16, 499-508.

Yannakakis G.N., J. Levine, and J. Hallam. 2004. "An Evolutionary Approach for Interactive Computer Games." In *Proceedings of the Congress on Evolutionary Computation (CEC-04)*, June, 986-993.

Yannakakis G.N., J. Levine, J. Hallam, and M. Papageorgiou. 2003. "Performance, Robustness and Effort Cost Comparison of Machine Learning Mechanisms in *FlatLand*." In *Proceedings of the 11th Mediterranean Conference on Control and Automation MED'03*. IEEE, June.

Yao X. 1999. "Evolving artificial neural networks." In *Proceedings of the IEEE*, vol. 87, no. 9, 1423-1447.

BIOGRAPHY

Georgios N. Yannakakis (Student Member, IEEE) received both the 5-year Diploma (1999) and the M.Sc. (2001) degree from the Technical University of Crete, Chania, Greece. He is a Ph.D. candidate at the School of Informatics, Centre of Intelligent Systems and their Applications, University of Edinburgh. G. Yannakakis is currently a visiting researcher at the Maersk Institute for Production Technology, University of Southern Denmark, Odense. He has published several research papers in the areas of cooperative multi-agent systems, evolutionary computation and AI in computer games.