

TOWARDS OPTIMIZING ENTERTAINMENT IN COMPUTER GAMES

Georgios N. Yannakakis □ *Mærsk Mc-Kinney Møller Institute, University of Southern Denmark, Odense M, Denmark*

John Hallam □ *Mærsk Mc-Kinney Møller Institute, University of Southern Denmark, Odense M, Denmark*

□ *Mainly motivated by the current lack of a qualitative and quantitative entertainment formulation of computer games and the procedures to generate it, this article covers the following issues: It presents the features—extracted primarily from the opponent behavior—that make a predator/prey game appealing; provides the qualitative and quantitative means for measuring player entertainment in real time, and introduces a successful methodology for obtaining games of high satisfaction. This methodology is based on online (during play) learning opponents who demonstrate cooperative action. By testing the game against humans, we confirm our hypothesis that the proposed entertainment measure is consistent with the judgment of human players. As far as learning in real time against human players is concerned, results suggest that longer games are required for humans to notice some sort of change in their entertainment.*

Intelligent interactive opponents can provide more enjoyment to a vast gaming community of constant demand for more realistic, challenging, and meaningful entertainment (Fogel et al. 2004; Champanand 2004). However, given the current state-of-the-art in artificial intelligence (AI) in computer games, it is unclear which features of any game contribute to the satisfaction of its players, and thus it is also uncertain how to develop enjoyable games. Because of this lack of knowledge, most commercial and academic research in this area is fundamentally incomplete. The challenges we consider in this article are to provide qualitative and quantitative means for distinguishing a game's enjoyment value and to develop efficient AI tools to automatically generate entertainment for the player.

In our previous work (Yannakakis and Hallam 2004), we defined criteria that contribute to the satisfaction for the player, which map to

The authors would like to thank the anonymous reviewers for their valuable feedback and the players for their vital contribution to this work.

Address correspondence to G. N. Yannakakis, Mærsk Mc-Kinney Møller Institute, University of Southern Denmark, Odense M, Denmark. E-mail: georgios@mmmi.sdu.dk

characteristics of the opponent behavior in computer games. According to our hypothesis, the player-opponent interaction—rather than the audiovisual features, the context, or the genre of the game—is the property that primarily contributes the majority of the quality features of entertainment in a computer game. Based on this fundamental assumption, we introduced a metric for measuring the real-time entertainment value of predator/prey games.

According to our second hypothesis, entertainment is generated when adaptive learning procedures occur in real time. This allows for opponents to learn while playing against the player and adapt with regards to his/her strategy. When such a mechanism is built upon cooperative opponents, it is more likely that the game's interest value improves drastically. Using the Pac-Man game as a test-bed (Yannakakis and Hallam 2004a, 2005) and focusing on the non-player characters' (NPC's) cooperative behavior, a robust online (i.e., while the game is played) neuro-evolution (i.e., automatic shaping of artificial neural networks by the use of artificial evolution) learning mechanism was presented, which was capable of increasing the game's interest as well as keeping that interest at high levels while the game was being played. This mechanism demonstrated high robustness and adaptability to changing hand-crafted player strategies in a relatively simple playing stage. Additional experiments in a predator/prey game of a more abstract design called Dead End (Yannakakis and Hallam 2004b) displayed the effectiveness of the proposed methodology in dissimilar games of the same genre and expanded the applicability of the method.

In the work presented here, apart from experiments with computer-guided players, human players are used for testing the Pac-Man game in a more complex stage. The main objective of this work is to establish the interest measure proposed as an efficient generic predator/prey game metric, by the approval of human judgment. In other words, we attempt to cross-correlate the human notion of interest to the proposed interest metric in one of the most representative test-beds of this computer games domain. Furthermore, we examine the online learning algorithm's abilities against humans and attempt to discover whether it maintains its robustness and adaptability under real conditions, that is, against human players.

LEARNING IN GAMES

The majority of research on learning in games is built on board or card games. In the last decade, many researchers have been involved in the development of intelligent opponents in board or card games. Some of the attempts include evolutionary learning approaches applied, from tic-tac-toe (Fogel 1993), to checkers (Fogel [2002] among others) and Go (Richards et al. 1998). In Tesauro (2002), a temporal difference learning

mechanism generates computer opponents capable of beating even expert humans in backgammon. These games are board games simulated in computers and therefore sometimes people refer to them as “computer games.” However, when we refer to computer games, we refer to the category of commercial games played by NPCs in virtual worlds.

Based on the success of this mentioned research on board games, increasing computing power and the commercial possibilities of computer games, very recently, researchers have attempted to introduce AI into computer games and have discussed the theoretical perspective of learning in different categories of games. Laird (2002) surveys the state of research in using AI techniques in interactive computer games. He also provides a taxonomy of games and the importance of computer games as experimental environments for strong AI application. Furthermore, Isla and Blumberg (2002) suggest potential research directions in AI game development, emphasizing to the emotional state and the perceived information of the character. Taylor (2000) attempts to bridge the gap between game development and modern AI by proposing artificial life techniques for generating physically modeled characters.

Game AI researchers, in their majority, focus on the genre of first-person shooter (FPS) and real-time strategy (RTS) games, primarily because of their popularity and secondarily because of their open-source game engines. Alex J. Champandard (2004) is using an FPS game to propose and apply a plethora of forms of AI techniques (varying from simple scripting to adaptive learning) for specific tasks like movement, shooting, and weapon selection. Khoo (2002) developed an inexpensive AI technique based on the well-known Eliza program (Weizenbaum 1966), so that users get the impression of playing against humans instead of bots. In Cole et al. (2004), the parameters of the *Counter-Strike* built-in weapon selection rules are tuned by using artificial evolution. Furthermore, there have been attempts to mimic human behavior offline, from samples of human playing, in a specific virtual environment. Alternatively, dynamic scripting and evolutionary learning has been used in a real-time strategy (RTS) games (Ponsen and Spronck 2004). In Thureau et al. (2004), among others, human-like opponent behaviors are emerged through supervised-learning techniques in *Quake*. Even though complex opponent behaviors emerge, there is no further analysis of whether these behaviors contribute to the satisfaction of the player (i.e., interest of game). In other words, researchers hypothesize—by observing the vast number of multi-player online games played daily on the Web, for example—that by generating human-like opponents, they enable the player to gain 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. Such a

hypothesis is the core of Iida's work on board games. He proposed a general metric of entertainment for variants of chess games depending on average game length and possible moves (Iida et al. 2003).

Identifying and Augmenting Entertainment

There have been several psychological studies to identify what is "fun" in a game and what engages people playing computer games. Theoretical approaches include Malone's (1981) principles of intrinsic qualitative factors for engaging game play, namely, challenge, curiosity, and fantasy as well as the well-known concepts of the theory of flow (Csikszentmihalyi 1990) incorporated in computer games as a model for evaluating player enjoyment, namely, *Game Flow* (Sweetser and Wyeth 2005). A comprehensive review of the literature on qualitative approaches for modeling player enjoyment demonstrates a tendency of overlapping with Malone's and Csikszentmihalyi's foundational concepts. Many of these approaches are based on Lazzaro's "fun" clustering, which uses four entertainment factors based on facial expressions and data obtained from game surveys on players (Lazzaro 2004): hard fun, easy fun, altered states, and socialization. Koster's (2005) theory of fun, which is primarily inspired by Lazzaro's four factors, defines "fun" as the act of mastering the game mentally. An alternative approach to fun capture is presented in Read et al. (2002), where fun is composed of three dimensions: endurability, engagement, and expectations. Questionnaire tools and methodologies are proposed in order to empirically capture the level of fun for evaluating the usability of novel interfaces with children.

Work in the field of quantitative entertainment capture and augmentation is based on the hypothesis that the player-opponent interaction—rather than the audiovisual features, the context, or the genre of the game—is the property that contributes to the majority of the quality features of entertainment in a computer game (Yannakakis and Hallam 2004a). Based on this fundamental assumption, a metric for measuring the real-time entertainment value of predator/prey games was designed, and established as a generic interest metric for prey/predator games (Yannakakis and Hallam 2005a; 2005b). Further studies by Yannakakis and Hallam (2006) have shown that artificial neural networks (ANN) and fuzzy neural networks can extract a better estimator of player satisfaction than a human-designed one, given appropriate estimators of the challenge and curiosity of the game (Malone 1981) and data on human players' preferences. Similar work in adjusting a game's difficulty include endeavors through reinforcement learning (Andrade et al. 2005), genetic algorithms (Verma and McOwan 2005), probabilistic models (Hunicke and Chapman 2004), and dynamic scripting (Spronck et al. 2004). However, the aforementioned attempts

are based on the assumption that challenge is the only factor that contributes to enjoyable gaming experiences while results reported have not been cross-verified by human players.

A step further to entertainment capture is towards games of richer human-computer interaction and affect recognizers, which are able to identify correlations between physiological signals and the human notion of entertainment. Experiments by Yannakakis et al. (2006) have already shown a significant effect of the average heart rate of children's perceived entertainment in action games played in interactive physical playgrounds. Moreover, Rani et al. (2005) propose a methodology for detecting the anxiety level of the player and appropriately adjusting the level of challenge in the game of "Pong." Physiological state (heart-rate, galvanic skin response) prediction models have also been proposed for potential entertainment augmentation in computer games (McQuiggan et al. 2006).

We choose predator/prey games as the initial genre of our game research since, given our aims, they provide us with unique properties. In such games, we can deliberately abstract the environment and concentrate on the characters' behavior. The examined behavior is cooperative since cooperation is a prerequisite for effective hunting behaviors. Furthermore, we are able to easily control a learning process through online interaction. In other words, predator/prey games offer a well-suited arena for initial steps in studying cooperative behaviors generated by interactive online learning mechanisms. Even though other genres of games (e.g., FPS games) offer similar properties, researchers have not yet focused on the novel directions of human-verified entertainment capture and augmentation that are presented in this article.

PREDATOR/PREY GAMES

Predator/prey games have been a very popular category of computer games and among its best representatives is the classical Pac-Man game released by Namco (Japan) in 1980. Even though Pac-Man's basic concept—the player's (*PacMan's*) goal is to eat all the pellets appearing in a maze-shaped stage while avoiding being killed by four opponent characters named "*Ghosts*"—and graphics are very simple, the game still keeps players interested after so many years, and its basic ideas are still found in many newly released games.

Kaiser et al. (1998) attempted to analyze emotional episodes, facial expressions, and feelings of humans playing a predator/prey computer game similar to Pac-Man (Kaiser and Wehrle 1996). Other examples in the Pac-Man domain literature include researchers attempting to teach a controller to drive *Pac-Man* in order to acquire as many pellets as possible and to avoid being eaten by *Ghosts*. Koza (1992) considers the problem of

controlling an agent in a dynamic nondeterministic environment and, therefore, sees Pac-Man as an interesting multi-agent environment for applying offline learning techniques based on genetic programming. Other approaches, such as incremental learning (Gallagher and Ryan 2003) and neuro-evolution (Lucas 2005), have also been applied for producing effective Pac-Man playing strategies. The same Pac-Man application domain has been used for analyzing size and generality issues in genetic programming (Rosca 1996).

On the other hand, there are many researchers who use predator/prey domains in order to obtain efficient emergent teamwork of either homogeneous or heterogeneous groups of predators. For example, Luke and Spector (1996), among others, have designed an environment similar to the Pac-Man game (the Serengeti world), in order to examine different breeding strategies and coordination mechanisms for the predators. Finally, there are examples of work in which both the predators' and the prey's strategies are co-evolved in continuous or grid-based environments (Haynes and Sen 1995; Miller and Cliff 1994).

Similar to Luke and Spector (1996), we view the domain from the predators' perspective and we attempt to emerge effective hunting teamwork offline based on evolutionary computation techniques applied to homogeneous neural controlled (Yao 1999) predators. However, playing a predator/prey computer game like Pac-Man against optimal hunters cannot be interesting because the player is consistently and effectively killed.

Researchers have generally shown that online learning in computer games is feasible through careful design and effective learning methodologies. On that basis, Yannakakis et al. (2004a; 2004c) introduced a neuro-evolution mechanism that acts in real time that optimizes each opponent individually (heterogeneous game environment) for generating appealing games rather than high opponent performance (Demasi and Cruz 2002; Graepel et al. 2004).

INTERESTING BEHAVIOR

As noted, predator/prey games will be our test-bed genre for the investigation of enjoyable games. More specifically, in the games studied, the prey is controlled by the player and the predators are the computer-controlled opponents (nonplayer characters, or NPCs).

In the approach presented in this section, a quantitative metric of player satisfaction is designed based on general criteria of enjoyment. The first step towards generating enjoyable computer games is therefore to identify the criteria or features of games that collectively produce enjoyment (or else interest) in such games. Second, quantitative estimators for these criteria are defined and combined, in a suitable mathematical

formula, to give a single quantity correlated with player satisfaction (interest). Finally, this formulation player interest is tested against human players' judgment in real conditions using the Pac-Man test-bed.

Following the principles of Yannakakis and Hallam (2004a), we will ignore mechanics, audiovisual representation, control, and interface contributions to the enjoyment of the player and we will concentrate on the opponents' behaviors. A well-designed and popular game such as Pac-Man can fulfil all aspects of player satisfaction incorporated in the mentioned design game features. The player, however, may contribute to his/her entertainment through interaction with the opponents of the game and, therefore, it is implicitly included in the interest formulation presented here, see also Yannakakis and Maragoudakis (2005), for studies of the player's impact on his/her entertainment.

Criteria

By observing the opponents' behavior in various predator/prey games, we attempted to identify the key features that generate entertainment for the player. These features were experimentally cross-validated against various opponents of different strategies and redefined when appropriate. According to Yannakakis and Hallam (2004a) and (2005b) the criteria that collectively define interest on any predator/prey game are briefly as follows.

1. Appropriate level of challenge (*when the game is neither too hard nor too easy*).
2. Behavior diversity (*when there is diversity in opponents' behavior over the games*).
3. Spatial diversity (*when opponents' behavior is aggressive rather than static*). That is, predators that move constantly all over the game world and cover it uniformly. This behavior gives the player the impression of an intelligent strategic opponents' plan, which increases the game interest.

Metrics

In order to estimate and quantify each of the three aforementioned interest criteria, we let the examined group of opponents play the game N times (each game for a sufficiently large evaluation period of t_{\max} simulation steps) and we record the simulation steps t_k taken to kill the player, as well as the total number of the opponents' visits v_{ik} at each cell i of the grid game field for each game k . In the case where the game's motion is continuous, a discretization of the field's plane up to the character's size can serve this purpose.

Given these, the quantifications of the three interest criteria proposed above can be presented as follows.

1. *Appropriate Level of Challenge*: According to the first criterion, an estimate of how interesting the behavior is, is given by T in (1).

$$T = [1 - (E\{t_k\} / \max\{t_k\})]^{p_1}, \quad (1)$$

where $E\{t_k\}$ is the average number of simulation steps taken to kill the prey-player over the N games; $\max\{t_k\}$ is the maximum t_k over the N games; and p_1 is a weighting parameter.

p_1 is adjusted so as to control the impact of the bracketed term in the formula for T . By selecting values of $p_1 < 1$, we reward quite challenging opponents more than near-optimal killers, since we compress the T value toward 1. More details on the adjustment of the p_1 value for the Pac-Man game will follow.

The T estimate of interest demonstrates that the greater the difference between the average and the maximum number of steps taken to kill the player, the higher the interest of the game. Given (1), both easy killing (“*too easy*”) and near-optimal (“*too hard*”) behaviors receive low interest estimate values (i.e., $E\{t_k\} \simeq \max\{t_k\}$). This metric is also called “challenge.”

2. *Behavior Diversity*: The interest estimate for the second criterion is given by S in (2).

$$S = (\sigma / \sigma_{max})^{p_2}, \quad (2)$$

where

$$\sigma_{max} = \frac{1}{2} \sqrt{\frac{N}{(N-1)} (t_{max} - t_{min})} \quad (3)$$

and σ is the standard deviation of t_k over the N games; σ_{max} is an estimate of the maximum value of σ ; t_{min} is the minimum number of simulation steps required for predators to kill the prey obtained by playing against some “well”-behaved fixed strategy near-optimal predators ($t_{min} \leq t_k$); and p_2 is a weighting parameter.

The S increases proportionally with the standard deviation of the steps taken to kill the player over N games. Therefore, using S as defined here, we promote predators that produce high diversity in the time taken to kill the prey.

3. *Spatial Diversity*: A good measure for quantifying the third interest criterion is through entropy of the predators’ cell visits in a game, which

quantifies the completeness and uniformity with which the opponents cover the stage. Hence, for each game, the cell visit entropy is calculated and normalized into $[0, 1]$ via (4).

$$H_n = \left[-\frac{1}{\log V_n} \sum_i \frac{v_{in}}{V_n} \log \left(\frac{v_{in}}{V_n} \right) \right]^{p_3}, \quad (4)$$

where V_n is the total number of visits of all visited cells (i.e., $V_n = \sum_i v_{in}$) and p_3 is a weighting parameter. p_3 is adjusted in order to promote very high H_n values and furthermore to emphasize the distinction between high and low normalized entropy values. Appropriate p_3 parameter values, which serve this purpose, are those greater than one ($p_3 = 4$ in this article), since they stretch the value of H_n away from 1.

Given the normalized entropy values H_n for all N games, the interest estimate for the third criterion can be represented by their average value $E\{H_n\}$ over the N games. This implies that the higher the average entropy value, the more interesting the game is.

The three individual criterion metrics defined are combined linearly to produce a single metric of interest (Equation 5) whose properties match the qualitative considerations developed.

$$I = \frac{\gamma T + \delta S + \epsilon E\{H_n\}}{\gamma + \delta + \epsilon}, \quad (5)$$

where I is the interest value of the predator/prey game; γ , δ and ϵ are criterion weight parameters.

The interest metric introduced in (5) can be applied effectively to any predator/prey computer game, because it is based on generic features of this category of games (see Yannakakis and Hallam [2004b; 2005] for successful applications to dissimilar predator/prey games). These features include the time required to kill the prey and the predators' entropy throughout the game field. We therefore believe that (5)—or a similar measure of the same concepts—constitutes a generic interest approximation of predator/prey computer games. Moreover, given the two first interest criteria previously defined, the approach's generality is expandable to all computer games. Indeed, no player likes any computer game that is too difficult or too easy to play and, furthermore, any player would enjoy diversity throughout the play of any game. The third interest criterion is applicable to games where spatial diversity is important which, apart from predator/prey games, may also include action, strategy, and team sports games

according to the computer game genre classification of Laird and van Lent (2000).

The approach to entertainment modeling represented by equation (5) is both innovative and efficient. However, it should be clear that there are many possible formulae, such as equation (5), which would be consistent with the qualitative criteria proposed for predator/prey games. Other successful quantitative metrics for the appropriate level of challenge, the opponents' diversity and the opponents' spatial diversity may be designed and more qualitative criteria may be inserted in the interest formula. Alternative mathematical functions for combining and weighting the various criteria could be employed.

For example, other metrics for measuring the appropriate level of challenge could be used: One could come up with a T metric assuming that the appropriate level of challenge follows a Gaussian distribution over $E\{t_k\}$ and that the interest value of a given game varies, depending on how long it is—*very short* ($E\{t_k\} \approx t_{min}$) games tend to be frustrating and *long games* ($E\{t_k\} \approx \max\{t_k\}$) tend to be boring. (However, very short games are not frequent in the experiments presented here and, therefore, by varying the weight parameter p_1 in the proposed T metric [see (1)], we are able to obtain an adequate level of variation in measured challenge.)

The question remains, however, whether the number produced by such a formula really captures anything useful concerning a notion so potentially complex as human enjoyment. That question is addressed next.

THE PAC-MAN TEST-BED

The computer game test-bed studied is a modified version of the original Pac-Man computer game released by Namco. The player's (*Pac-Man's*) goal is to eat all the pellets appearing in a maze-shaped stage, while avoiding being killed by the four *Ghosts*. The game is over when either all pellets in the stage are eaten by *Pac-Man*, *Ghosts* manage to kill *Pac-Man*, or a pre-determined number of simulation steps is reached without any of these occurring. In that case, the game restarts from the same initial positions for all five characters.

Compared to commercial versions of the game, a number of features (e.g., power pills) are omitted for simplicity; these features do not qualitatively alter the nature of "interesting" in games of low interest. Cross-validation of this statement appears through the judgment and the beliefs of human players of both the original and this version of the game.

As stressed previously, the Pac-Man game is investigated from the viewpoint of *Ghosts* and more specifically how *Ghosts'* emergent adaptive behaviors can contribute to the interest of the game. Pac-Man—as a computer game domain for emerging adaptive behaviors—is a two-dimensional,




FIGURE 1 Snapshot of the Pac-Man game.

multi-agent, grid-motion, predator/prey game. The game field (i.e., stage) consists of corridors and walls. Both the stage's dimensions and its maze structure are predefined. For the experiments presented in this article, we use a 19×29 grid maze-stage where corridors are one grid-cell wide (see Figure 1).

The characters visualized in the Pac-Man game (as illustrated in Figure 1) are a white circle that represents *Pac-Man* and four ghost-like characters representing the *Ghosts*. Being *Ghosts*, one of their properties is permeability, i.e., two or more *Ghosts* can simultaneously occupy the same cell of the game grid. Additionally, there are black squares that represent the pellets and dark gray blocks of walls.

Pac-Man moves at double the *Ghosts'* speed and since there are no dead ends, it is impossible for a single *Ghost* to complete the task of killing it. Since *Pac-Man* moves faster than a *Ghost*, the only effective way to kill a well-performing *Pac-Man* is for a group of *Ghosts* to hunt cooperatively.

Pac-Man

Both the difficulty and, to a lesser degree, the interest of the game are directly affected by the intelligence of the *Pac-Man* player. In Yannakakis and Hallam (2004a; 2005), three fixed *Ghost*-avoidance and pellet-eating strategies for the *Pac-Man* player, differing in complexity and effectiveness, are presented. Each strategy is based on decision making applying a cost or probability approximation to the player's four neighboring cells (i.e., up, down, left, and right). We present them briefly in this article.

- Cost-Based (CB) *Pac-Man*: Moves towards a cost minimization path that produces effective *Ghost*-avoidance and (to a lesser degree) pellet-eating behaviors, but only in the local neighbor cell area.
- Rule-Based (RB) *Pac-Man*: This is a CB *Pac-Man*, plus an additional rule for more effective and global pellet-eating behavior.
- Advanced (ADV) *Pac-Man*: The ADV moving strategy generates a more global *Ghost*-avoidance behavior built upon the RB *Pac-Man*'s good pellet-eating strategy.

Ghosts

The arcade version of Pac-Man uses a handful of simple rules and scripted sequences of actions that control each *Ghost* individually (Wikipedia), combined with some random decision making to make the *Ghosts*' behavior less predictable. Even though this design yields quite complex *Ghost* behaviors, the player's interest decreases at the point where *Ghosts* are too fast to beat (Rabin 2002). In our Pac-Man version, we require *Ghosts* to keep learning and constantly adapting to the player's strategy, instead of being opponents with fixed strategies and furthermore maintain a constant real-time speed.

Neural networks are a suitable host for emergent adaptive behaviors in complex multi-agent environments (Ackley and Littman 1992) and have been successfully applied for adaptive learning in real time in computer games (Stanley et al. 2005). A multi-layered fully connected feedforward neural controller, where the sigmoid function is employed at each neuron, manages the *Ghosts*' motion. Using their sensors, *Ghosts* inspect the environment from their own point of view and decide their next action. Each *Ghost*'s perceived input consists of the relative coordinates of *Pac-Man* and the closest *Ghost*. We deliberately exclude from consideration any global sensing, e.g., information about the dispersion of the *Ghosts* as a whole, because we are interested specifically in the minimal sensing scenario. The neural network's output is a four-dimensional vector with respective values from 0 to 1 that represents the *Ghost*'s four movement options (up, down, left, and right, respectively). Each *Ghost* moves towards the available—unobstructed by walls—direction represented by the highest output value. Available movements include the *Ghost*'s previous cell position.

Fixed Strategy Ghosts

Apart from the neural controlled *Ghosts*, three additional non-evolving strategies have been tested for controlling the *Ghost*'s motion. These strategies are used as baseline behaviors for comparison with any neural-controller emerged behavior.

- Random (R): *Ghosts* that randomly decide their next available movement. Available movements have equal probabilities to be picked.
- Followers (F): *Ghosts* designed to follow *Pac-Man* constantly. Their strategy is based on moving so as to reduce the greatest of their relative coordinates from *Pac-Man*.
- Near-Optimal (O): A *Ghost* strategy designed to produce attractive forces between *Ghosts* and *Pac-Man*, as well as repulsive forces among the *Ghosts*.¹

ADJUSTING INTEREST PARAMETER VALUES FOR PAC-MAN

In this section, we present the procedures followed to obtain the appropriate parameter values of the interest estimate (5) for the Pac-Man game. In this article t_{\min} is 35 simulation steps, which is obtained as the minimum simulation time that *Pac-Man* survives when playing against the best-performing near-optimal *Ghosts*. In addition, t_{\max} is 320 simulation steps, which corresponds to the minimum simulation period required by the RB *Pac-Man* (best pellet-eater) to clear the stage of pellets.

In order to obtain values for the interest formula weighting parameters p_1 , p_2 , and p_3 , we select empirical values based on each interest criterion. For the first interest metric presented in (1), p_1 is adjusted so as to give T a greater impact or else a boost when even a slight difference between the maximum and the average lifetime of the player (i.e., challenge) is noted ($p_1 < 1$). This way we reward quite challengeable opponents more than near-optimal killers. For the third interest metric presented in (4), p_3 is adjusted in order to press for very high H_n values and furthermore to provide a clearer distinction between high and low normalized entropy values ($p_3 > 1$). Finally, p_2 is set so as σ has a linear effect on S . By taking this into consideration, $p_1 = 0.5$, $p_2 = 1$, and $p_3 = 4$, for the experiments presented in this article.

Moreover, values for the interest criteria weighting parameters γ , δ , and ϵ are also selected empirically based on the specific game. In Pac-Man, aggressive opponent behavior is of the greatest interest. The game loses any reliability when *Ghosts* stick in a corner instead of wandering around the stage. Thus, diversity in gameplay (S) and appropriate level of challenge (T) should come next in the importance list of interest criteria. Given the previously mentioned statements and by adjusting these three parameters so that the interest value escalates as the opponent behavior changes from Random to near-optimal, and then to follower, we come up with $\gamma = 1$, $\delta = 2$ and $\epsilon = 3$.

Since the interest value changes monotonically with respect to each of the three criterion values T , S , $E\{H_n\}$, sensitivity analysis is conducted on

the interest metric parameters, aiming to portray the relation between these parameters as well as their weighting degree in the interest formula. We therefore proceed by seeking opponent behaviors that generate ten different T , S , and $E\{H_n\}$ values, equally spread in the $[0,1]$ interval. Given these thirty values as input, p_1 , p_2 , p_3 , γ , δ , and ϵ parameters are systematically changed one at a time so that their percentage difference lies in the interval $[-50\%, 50\%]$. Each time a parameter change occurs, the absolute percentage difference of the game's interest is computed. The function between the absolute percentage differences of the interest value and the percentage differences of the interest weighting parameters is illustrated in Figure 2.


As seen in Figure 2, changes on the p_1 , p_2 , and p_3 parameters seem to influence the I value more than γ , δ , and ϵ . The observed difference in interest sensitivity is reasonable since the first three parameters represent powers, while the latter three correspond to product weights. More specifically, p_2 and p_3 reveal significant differences (i.e., greater than 5%) in I when decreased by 15% (i.e., $p_2 = 0.85$) and 9% (i.e., $p_3 = 3.64$) or increased by 20% (i.e., $p_2 = 1.2$) and 10% (i.e., $p_3 = 4.4$), respectively. For p_1 significant change in I is observed only when decreased by up to 35% (i.e., $p_1 = 0.325$). Accordingly, both ϵ and δ parameters reveal significant differences in I only when decreased by 40% and 45% respectively. Finally, for γ no significant change in I is observed even when changed by up to 50%.

Regardless of the sensitivity of the I value, as far as mainly the p_2 and p_3 parameters are concerned, we believe that the selected values project a robust I value considering that they constitute power parameters in the interest formula.


MEASURING PERFORMANCE

When a predator/prey game is investigated from the predator's viewpoint, optimality can be measured in the predators' ability to kill the prey. Thus, prey-killing ability of the *Ghosts* is the primary factor that determines how well-performed a behavior is in the Pac-Man game. Furthermore, preventing *Pac-Man* from eating pellets, which also implies fast-killing capabilities, constitutes an additional factor of the desired optimal behavior. Given these, a measure designed to give an approximation of a group of *Ghosts*' performance over a specific number N of games played, is

$$P = \frac{\alpha(k/N) + \beta \min\{1 + (e_{\min} - E\{e\})/e_{\max} - e_{\min}), 1\}}{\alpha + \beta}, \quad (6)$$



(a) Average



(b) Standard deviation

FIGURE 2 Absolute percentage differences of I over ten runs for each weighting parameter.

where P is the performance of a *Ghost* group behavior taking values from 0 to 1; k is the number of *Pac-Man* kills within N games; $E\{e\}$ is the average number of pellets eaten by *Pac-Man* over the N games; e_{\min} , e_{\max} are the lower and upper bound of the eaten pellets e , respectively (in this article $e_{\min} = 80$, $e_{\max} = 227$); α , β are weight parameters (in this article $\alpha = \beta = 1$).

OFFLINE LEARNING

We use an offline evolutionary learning approach in order to produce some “good” (i.e., in terms of performance) initial behaviors. An additional aim of the proposed algorithm is to emerge dissimilar behaviors of high fitness—varying from blocking to aggressive—offering diverse seeds for the online learning mechanism in its attempt to generate emergent *Ghost* behaviors that make the game interesting. The offline learning mechanism used is presented in Yannakakis and Hallam (2004a) and Yannakakis et al. (2004c).

ONLINE LEARNING (OLL)

As previously noted, games which can learn and adapt to new playing strategies offer a richer interaction to entertain the player. For that purpose, we use an evolutionary machine learning mechanism for the Pac-Man game, which is based on the idea of *Ghosts* that learn while they are playing against *Pac-Man*. Or else, *Ghosts* that are reactive to any player’s behavior learn from its strategy instead of being opponents with fixed strategies that exist in all versions of this game today. Furthermore, this approach’s additional objectives are to keep the game’s interest at high levels as long as it is being played and to achieve good real-time performance (i.e., low computational effort during gameplay). This approach is first introduced in Yannakakis et al. (2004a) for a prey-predator game called “Dead-End” and in (Yannakakis and Hallam (2004a) for the Pac-Man game.

Beginning from any initial group of OLT *Ghosts*, the OLL mechanism transforms them into a group of heterogeneous opponents that are conceptually more interesting to play against. An OLT *Ghost* is cloned four times and its clones are placed in the game field to play against a selected fixed *Pac-Man* type in a selected stage. Then, at each generation:

Step 1. Each *Ghost* is evaluated every t simulation steps via (7), while the game is played— $t = 50$ simulation steps in this article.

$$f' = \sum_{i=1}^{t/2} \{d_{P,2i} - d_{P,(2i-1)}\}, \quad (7)$$

where $d_{P,i}$ is the distance between the *Ghost* and *Pac-Man* at the i simulation step. This fitness function promotes *Ghosts* that move towards *Pac-Man* within an evaluation period of t simulation steps.

- Step 2. A pure elitism selection method is used where only the fittest solution is able to breed. The fittest parent clones an offspring with a probability p_c that is inversely proportional to the normalized cell visit entropy (i.e., $p_c = 1 - H_n$) given by (4). In other words, the higher the cell visit entropy of the *Ghosts*, the lower the probability of breeding new solutions. If there is no cloning, then go back to Step 1, else continue to Step 3.
- Step 3. Mutation occurs in each gene (connection weight) of the offspring's genome with a small probability p_m (e.g., 0.02). A Gaussian random distribution is used to define the mutated value of the connection weight. The mutated value is obtained from (8).

$$w_m = \mathcal{N}(w, 1 - H_n), \quad (8)$$

where w_m is the mutated connection weight value and w is the connection weight value to be mutated. The Gaussian mutation, presented in (8), suggests that the higher the normalized entropy of a group of *Ghosts*, the smaller the variance of the Gaussian distribution and therefore, the less disruptive the mutation process as well as the finer the precision of the GA.

- Step 4. The mutated offspring is evaluated briefly via (7) in offline mode, that is, by replacing the least-fit member of the population and playing an offline (i.e., no visualization of the actions) short game of t simulation steps. If there is a human playing Pac-Man, then the *Pac-Man's* motion trail of the last t simulation steps is recorded and opponents are evaluated against it in offline mode. The fitness values of the mutated offspring and the least-fit *Ghost* are compared and the better one is kept for the next generation. This pre-evaluation procedure for the mutated offspring attempts to minimize the probability of group behavior disruption by low-performance mutants. The fact that each mutant's behavior is not tested in a single-agent environment but within a group of heterogeneous *Ghosts*, helps more towards this direction. If the least-fit *Ghost* is replaced, then the mutated offspring takes its position in the game field as well.

The algorithm is terminated when a predetermined number of games has been played or a game of high interest (e.g., $I \geq 0.7$) is found.

We mainly use short simulation periods ($t = 50$) in order to evaluate *Ghosts* in OLL, aiming to the acceleration of the online evolutionary process. The same period is used for the evaluation of mutated offspring; this is based on two primary objectives: 1) to apply a fair comparison between the mutated offspring and the least-fit *Ghost* (i.e., same evaluation period) and

2) to avoid undesired high computational effort in online mode (i.e., while playing).

EXPERIMENTS AGAINST FIXED PLAYING STRATEGIES

Results obtained from experiments applied on the Pac-Man game against fixed playing strategies are presented in this section. These include offline training and online learning emergent behavior analysis as well as experiments for testing robustness and adaptability of the online learning approach proposed.

Offline Training

The procedure presented in this subsection is focused on generating well-behaved *Ghosts* in terms of the performance measure described previously. We train *Ghosts* against all three types of *Pac-Man* player through the neuro-evolution offline learning mechanism presented in previously.

In order to minimize the non-deterministic effect of the *Pac-Man's* strategy on the *Ghost's* performance and interest values as well as to draw a clear picture of these averages' distribution, we apply the following bootstrapping procedure. Using a uniform random distribution we pick 10 different 50-tuples out of the 100 previously mentioned games. These 10 samples of data (i.e., e, k, t_k, v_{ik}) from 50 games (i.e., $N = 50$) are used to determine the *Ghosts's* average performance and interest values and their respective confidence intervals. The outcome of this experiment is presented in Table 1.

TABLE 1 Performance (P) and Interest (I) Values (Average Values of 10 Samples of 50 Games Each) of Fixed Strategy (R, F, O) and OLT *Ghosts* (B, A, H) Playing against All Three *Pac-Man* Types (CB, RB, ADV). Average P and I Values ($E\{\}$) of All Six Strategies Appear in the Bottom Row. Experiment Parameters: Population Size is 80, $g = 1000$, $t = 320$ Simulation Steps. $N_i =$ Games, $p_m = 0.02$, 5-Hidden Neurons Controller

	Trained offline by playing against					
	CB		RB		ADV	
	P	I	P	I	P	I
R	0.423	0.547	0.363	0.586	0.356	0.523
F	0.754	0.771	0.701	0.772	0.621	0.771
O	0.891	0.729	0.897	0.749	0.964	0.686
B	0.734	0.576	0.689	0.412	0.869	0.442
A	0.661	0.654	0.606	0.652	0.662	0.555
H	0.348	0.190	0.310	0.250	0.467	0.423
$E\{\}$	0.635	0.578	0.592	0.570	0.656	0.566

Offline trained emergent solutions are the OLL mechanisms' initial points in the search for more interesting games. OLT obtained behaviors are classified into the following categories.

- Blocking (B): These are the OLT *Ghosts* that achieve the best performance against each *Pac-Man* type. Their behavior is characterized as 'Blocking' because they tend to wait for *Pac-Man* to enter a specific area that is easy for them to block and then kill. Their average normalized cell visit entropy value $E\{H_n\}$ lies between 0.55 and 0.65.
- Aggressive (A): These are OLT *Ghosts* that achieve lower performance in comparison to the blockers. Their behavior is characterized as "aggressive" because they tend to follow *Pac-Man* all over the stage in order to kill it. This motion feature generates the highest I value ($E\{H_n\} \geq 0.65$) among the interest values generated by the three emergent behaviors.
- Hybrid (H): These are suboptimal OLT *Ghosts* that achieve the lowest performance ($P \leq 0.55$) and low interest value in comparison to the aforementioned B and A *Ghosts* ($E\{H_n\} < 0.55$). Their behavior is characterized as "hybrid" because they tend to behave as a blocking-aggressive hybrid, which proves to be ineffective at killing *Pac-Man*.

As far as the interest value generated by the mentioned behaviors is concerned, confidence intervals (± 0.0647 maximum, ± 0.0313 on average) obtained by the bootstrapping procedure previously described indicate that B, A, and H are significantly different.

According to Table 1, near-optimal and blocking behavior *Ghosts* achieve high-performance values against all three *Pac-Man* types, whereas their interest value is not as high as their performance value. This reveals the compromise between optimality and interest it has to be made because, in a predator/prey computer game, optimal killing behaviors are almost never interesting behaviors. On the other hand, followers are likely to produce the most interesting behaviors (among the behaviors examined in Table 1) for the game.

Viewing results presented in Table 1 from the *Pac-Man* type perspective (i.e., the average values in the bottom row of the table), it looks as if the RB and ADV are, respectively, the hardest and easiest *Pac-Man* players to kill. Concerning the three *Pac-Man* types' generated interest, it seems that there is no significant difference amongst them.

Online Learning Experiments

As previously mentioned, the offline learning procedure is a mechanism that produces near-optimal solutions to the problem of killing

Pac-Man and minimizing the pellets eaten in a game. These solutions are the OLL mechanisms' initial points in the search for more interesting games. The OLL experiment is described as follows.

- Pick the nine emerged *Ghosts*' behaviors produced from the offline learning experiments presented previously (i.e., B, A and H behaviors emerged by playing against each *Pac-Man* type).
- Starting from each OLT behavior, apply the OLL mechanism by playing against the same type of *Pac-Man* as was used offline.
- Calculate the interest (bootstrapping procedure with $N = 50$ of the game every 500 games during each OLL attempt).

The evolution of interest over the OLL games of each one of the three OLT behaviors is presented in a subfigure of Figure 3. For each of the three subfigures, three lines are illustrated, representing the interest values and their respective confidence intervals of the OLL attempt playing against the three *Pac-Man* types. Figure 3(d) illustrates the overall picture of the OLL experiments by comparing the initial interest of the game against the best average interest value achieved from OLL. Clearly, the OLL approach constitutes a robust mechanism that, starting from near-optimal or suboptimal *Ghosts*, manages to emerge interesting games (i.e., interesting *Ghosts*) in the vast majority of cases (i.e., in 8 out of 9 cases $I > 0.7$). In 5 out of 9 OLL attempts, the best interest value is significantly greater or statistically equal to the respective follower's value (i.e., 0.771 against CB, 0.772 against RB and 0.771 against ADV).

As seen from Figure 3, OLL enhances the interest of the game independently of the initial OLT behavior or the *Pac-Man* player *Ghosts* play against. In all experiments presented here, the learning mechanism is capable of producing games of higher than the initial interest as well as keeping that high interest for a long period. There is obviously a slight probability of disruptive mutations (the higher the game's interest, through the cell visit entropy value, the less the probability of mutation) that can cause undesired drops in the game's interest. However, OLL is robust enough to recover from such disruptive phenomena (Figure 3).

Given an interesting initial behavior (e.g., aggressive behavior, $I > 0.6$), it takes some few thousands of games for the learning mechanism to produce games of high interest. On the other hand, it takes some several thousand games to transform an uninteresting near-optimal blocking behavior (see Figure 3(a) and Figure 3(c)) into an interesting one. That is because the OLL process requires an initial long period to disrupt the features of an uninteresting blocking behavior, in order to be able to boost the interest of the game.




FIGURE 3 Game interest over the number of OLL games. For reasons of computational effort, the OLL procedure continues for a number of games large enough to illustrate its behavior, after a game of high interest ($I \geq 0.7$) is found. Initial *Ghost* behaviors appear in (a), (b), and (c) subfigure caption, whereas (d) illustrates the overall picture of the experiment. Experiment parameters: $t = 50$ simulation steps, $p_m = 0.02$, 5-hidden neurons controller.

It is obvious that a number in the scale of 10^3 constitutes an unrealistic number of games for a human player to play. On that basis, it is very unlikely for a human to play so many games in order to notice the game's interest increasing. The reason for the OLL process being that slow is a matter of keeping the right balance between the process' speed and its "smoothness" (by smoothness we define the interest's magnitude of change over the games). A solution to this problem is to consider the initial long period of disruption as an offline learning procedure and start playing as soon as the game's interest is increased. Moreover, other online learning approaches like co-evolution (Demasi and Cruz 2002), dynamic scripting (Spronck et al. 2004), and reinforcement learning (Graepel et al. 2004; Andrade et al. 2005) could, in part, provide a solution for the long convergence times observed.

How effective will this mechanism be in a potential change from a fixed strategy to a dynamic human *Pac-Man* player? The next subsection provides evidence in order to support the answer.

Adaptability

In order to test the OLL approach's ability to adapt to a changing environment (i.e., change of *Pac-Man* strategy), the following experiment is proposed. Beginning from an initial behavior of high interest value I_{init} , we apply the OLL mechanism against a specific *Pac-Man* type. During the online process, we keep changing the type of player as soon as interesting




FIGURE 4 Online learning *Ghosts* playing against changing types of *Pac-Man*. Sub-figure captions indicate the playing *Pac-Man* sequence.

games (i.e., $I \geq I_{init}$) are produced. The process stops when all three types of players have played the game.

Since we have three types of players, the total number of such experiments is six (all different player type sequences). These experiments illustrate the overall picture of the approach's behavior against any sequence of *Pac-Man* types. As seen in Figure 4, OLL is able to quickly recover a sudden change in the player's strategy and boost the game's interest at high levels after sufficient games have been played (i.e., 100 to 500 games). The mechanism demonstrates a similar adaptive behavior for all six sequences of *Pac-Man* players, which illustrates its independence of the sequence of the changing *Pac-Man* type.

Results obtained from this experiment provide evidence for the approach's ability to adapt to new types of players as well as its efficiency in producing interesting games against humans with dynamic playing strategies.

EXPERIMENTS AGAINST HUMAN PLAYERS

Experiments against fixed playing strategies portrayed the OLL mechanism's ability to generate interesting *Pac-Man* games. Apart from being fairly robust, the proposed mechanism demonstrated high and fast adaptability to changing types of player (i.e., playing strategies). The next obvious step to take is to let humans judge whether generated games are realistically interesting, or not, and whether OLL indeed enhances the level of entertainment during play. For this, we conducted a survey, with human subjects as *Pac-Man* players, that primarily aimed to obtain answers for the following questions.

1. Does the interest value computed for a game correlate with human judgment of interest?
2. Does the online learning mechanism cause perceived interest to change? Do perceived changes match computed ones?

The experiment is described next. Then the statistical method used and the analysis of obtained results are presented, respectively.

EXPERIMENT DESCRIPTION

Answers to the target questions presented previously are based on statistical analysis of data acquired from a questionnaire applied for the *Pac-Man* game. The main prerequisite for a subject to participate in this experiment is to have played the original version (Namco) of the *Pac-Man* game at least once. Subject age covers a range between 17 and 51

years, where both sexes are almost equally represented (43.3% females, 56.7% males). In addition, all subjects speak English as a foreign language since their nationality is either Danish (90%) or Greek (10%). The questionnaire is divided into three parts (A, B, and C) and the steps that the subjects go through at each part are presented as follows.

Personal Data

Subjects are asked to define their interest in the Pac-Man game before they play it. Answers categorize participants into three types of Pac-Man player' represented as "Like," "Neutral," and "Don't like."

Subjects familiarize themselves with the game by playing 50 games against specific OLT opponents (i.e., opponent 4 presented in Table 2). Online learning is used during this testing period, which is not noticeable to the player. At the end of the testing period, each subject's opponents trained online are saved.

1st Objective

We pick opponents differing in interest measured against the ADV player (as the most advanced computer-guided *Pac-Man* player). We select five opponents whose interest values uniformly cover the $[0, 1]$ space. The selected opponents' number, which is used as an id-code, and their respective interest values are presented in Table 2.

By experimental design, each subject plays against three of the selected opponents in all permutations of pairs. In addition, we require equal participation of all three player types. For this experiment, we use 30 subjects divided into three equal subsets for each of the three player types (Like, Neutral, Don't Like), since $C_3^5 = 10$ subjects are required for each player type. Moreover, observed effects show that 30 subjects constitute a statistically significant sample.

TABLE 2 The Selected Opponents and Their Respective Interest— I and 95% Confidence Interval (I_u , I_l) Values

Opponent	Interest		
	I_u	I	I_l
1	0.2043	0.1793	0.1494
2	0.3673	0.3158	0.2670
3	0.5501	0.4943	0.4420
4	0.6706	0.6484	0.6267
5	0.8180	0.8023	0.7858

- As previously mentioned, each subject plays sets of games (five games in each set) against three of the selected opponents in all permutations of pairs and each time a pair of sets is completed, the player is asked whether the first set was more interesting than the second set of games. We use the 2-alternative forced choice (2-AFC) approach since it offers several advantages for a subjective interest capture. The 2-AFC comparative fun analysis (Read et al. 2002) minimizes the assumptions about people's different notions of entertainment and provides data for a fair comparison among answers of different people (Yannakakis et al. 2006)

The total number of sets of games that is played by each subject is 12 (all permutations of three pairs. Given thirty subjects, there are nine observed incidents for each pair of sets.

2nd Objective

Each subject plays 25 games against the initial training phase opponents (i.e., opponent 4—OP4) and 25 games against the online trained opponents that were saved (i.e., two sets of games). We let each subject play another two sets against these opponents in different order. Half of the subjects play these four sets of games in the sequence OLL-OP4, OP4-OLL, whereas the other half play them in the sequence OP4-OLL, OLL-OP4, since we require minimization of any potential ordering effect. Each time a pair of sets (two pairs here) is finished, the player is asked whether the first set was more interesting than the second set of games.

Subjects are asked to list the criteria they used for their assessment of which set of games was more interesting.

STATISTICAL METHOD

For this experiment, there are three null hypotheses formed.

- H_0 : The correlation between observed human judgment of interest and the computed interest value, as far as the different opponents are concerned, is a result of randomness.
- H_1 : Observed human judgment of interest does not correlate with the computed interest value, as far as the different opponents are concerned.
- H_2 : Observed human judgment of interest does not correlate with performance during play.

Given the interest metric (5) and two sets of games A and B , it can be determined that “game A is more (or less) interesting than game B .” In answer to the same question, a human subject can indicate that either

$I_A > I_B$ or $I_A < I_B$. In order to measure the degree of agreement between the human judgment of interest and the interest value given by (5), we calculate the correlation coefficients

$$c(\vec{z}) = \sum_{i=1}^N z_i / N, \quad (9)$$

where N is the number of incidents to correlate and

$$\vec{z} = \begin{cases} 1, & \text{if subject agrees with (5);} \\ -1, & \text{if subject disagrees with (5).} \end{cases} \quad (10)$$

The test statistic (9) is used to assess the truth of all three null hypotheses. However, for the null hypothesis H_2 , the correlation coefficients $c(\vec{z}')$ are computed where z' values are obtained from (11).

$$\vec{z}' = \begin{cases} 1, & \text{if subject chooses according to performance;} \\ -1, & \text{if subject does not choose according to performance.} \end{cases} \quad (11)$$

The distribution used for obtaining the correlation coefficient probabilities (p-values— $P(C \geq c)$) is the binomial. The observed effect is “highly significant” and “significant” if $P(C \geq c) < 1\%$ and $1\% < P(C \geq c) \leq 5\%$, respectively.

For the design of the subjects’ self reports we follow the principles of comparative fun analysis (Read et al. 2002; Yannakakis et al. 2006). The endurance and expectations for the majority of subjects that played Pac-Man were very high, indicating that the game design used was successful. More specifically, all subjects were excited to play Pac-Man as soon as they were informed about the rules of the game (derived through a *Funometer* tool application [Read et al. 2002]) and the majority of subjects stressed that they would like to play the game again (derived through an *Again-Again* table [Read et al. 2002]). As previously mentioned, we use the 2-alternative forced choice (2-AFC) approach since it offers several advantages for a subjective entertainment capture. The 2-AFC comparative fun analysis minimizes the assumptions about people’s different notions of entertainment and provides data for a fair comparison among answers of different people.

STATISTICAL ANALYSIS

As noted previously, this article concentrates on the characters’ behavioral aspect of interesting games. More specifically, it focuses on the opponent’s rather than the graphics’ or the sound’s impact on the player’s entertainment. Apart from the opponent, there are two additional factors that may affect the interest of a computer game, that are examined in this

section. These are the player-subject type (degree of *a priori* game liking) and the order of play.

Opponent

Each entity in Table 3 represents a subject’s answer to the question from the 1st objective, equivalent to “Is $I_i > I_j?$,” where i, j , are the row and column number, respectively. Given the interest values of the five opponents (see Table 2), “O and X” stand, respectively, for the subject’s agreement and disagreement with this ranking (in other words, O and X characters are selected for visual purposes to symbolize the respective z values—see (10). As stressed before, given 30 subjects, there are nine incidents for each pair of opponent which are represented in a 3×3 matrix. Rows within this matrix denote the type of subject that answered the specific question.

Table 4 presents the correlation coefficients and their respective $P(C \geq c)$ values for each one of the ten combinations of opponent pairs ($N = 18$) and in total ($N = 180$). There is an obvious disagreement between the interest metric and the human’s notion of interest in opponent pairs 1–2 and 3–4. Even though humans seem to agree with the interest metric in the pairs 1–3 and 1–4, the obtained p-values reveal statistically insignificant results. For the rest of the pairs, we experience statistically highly significant (i.e., 2–3, 2–5, 4–5) and significant (i.e., 1–5, 2–4, 3–5) matching to

TABLE 3 Agreement Between the Subject’s Judgment of Interest and the Interest Metric—O: $z = 1$, X: $z = -1$

Subject type		Is $I_{Row} > I_{Column}?$				
		1	2	3	4	5
1	Like		O O X	O O X	O O X	O O O
	Neutral		O O X	O O O	O O X	O O X
	Don’t Like		O O O	O O O	O O O	O X X
2	Like	X X X		O O O	O O X	O O O
	Neutral	X X X		O O X	O O O	O O X
	Don’t Like	O X X		O O O	O O X	O O X
3	Like	O X X	O O O		O X X	O O X
	Neutral	O O X	O O X		O X X	O O O
	Don’t Like	O X X	O O X		O O X	O O X
4	Like	O O X	O O O	X X X		O O X
	Neutral	O O X	O O X	X X X		O O X
	Don’t Like	O X X	O O X	O X X		O O O
5	Like	O O X	O O O	O O O	O O O	
	Neutral	O O O	O O X	O O X	O O O	
	Don’t Like	O O O	O O O	O O X	O O O	

TABLE 4 Interest Metric – Subject Judgment Correlation Coefficients c , $P(C \geq c)$ Values, Order of Play Test Statistic z' , and $P(Z \geq |z'|)$ Values for all Pairs of Opponents and in Total

Pair	c	$P(C \geq c)$	z'	$P(Z \geq z')$
1–2	–0.111	0.7596	0.2222	0.2403
1–3	0.3333	0.1189	0.3333	0.1189
1–4	0.3333	0.1189	–0.222	0.2403
1–5	0.5555	0.0154	–0.111	0.4072
2–3	0.6666	0.0037	0.1111	0.4072
2–4	0.5555	0.0154	0.1111	0.4072
2–5	0.7777	0.0006	0.0000	0.5927
3–4	–0.444	0.9845	–0.111	0.4072
3–5	0.5555	0.0154	–0.333	0.1189
4–5	0.6666	0.0037	0.2222	0.2403
Total	0.3888	$1.31 \cdot 10^{-7}$	0.0222	0.4818

observed human judgment. Finally, the total agreement correlation coefficient ($c = 0.3888$), as well as its p-value ($P(C \geq c) = 1.31 \cdot 10^{-7}$), demonstrate a statistically highly significant effect that rules out the null hypothesis H_1 . Thus, it appears that the observed human judgment of interest correlates with the computed interest value, as far as the different opponents are concerned. Moreover, the obtained p-values presented in Table 4 illustrate that the sample size of 30 subjects is adequate to produce statistically significant observed effects.

Opponent 1

Further investigation of the interest value generated by opponent 1 showed high dependence on the player type. More specifically, when opponent 1 plays against the CB *Pac-Man* and the RB *Pac-Man*, it generates interest, which is respectively statistically not different and significantly higher than the interest generated by opponent 2. Opponent 1 constitutes a particular case since no such change in the opponent ranking (i.e., ranked by interest) occurs for any other of the four remaining opponents.

Given the ranking instability of opponent 1, we recalculate the z values as if (1) $I_1 > I_2$ and (2) $I_1 = I_2$ and proceed. In the former case, the z values of the [1–2] pair swap their sign and the obtained p-values for this pair and in total are 0.4072 and $2.57 \cdot 10^{-8}$, respectively. For the latter case, the z values of the [1–2] pair are not taken into consideration and the two first (triplets of) rows and columns of Table 3 are merged into one by adding up their z values. The obtained p-values for the merged pairs and in total are presented in Table 5. For both cases, changes in the opponent 1 ranking increase the significance of the observed effects.

TABLE 5 Interest Metric: Subject Judgment Correlation Coefficients c and $P(C \geq c)$ Values When $I_1 = I_2$ Is Assumed

Pairs	c	$P(C \geq c)$
(1,2)-3	0.5000	0.0019
(1,2)-4	0.4444	0.0056
(1,2)-5	0.6667	$3.5 \cdot 10^{-5}$
Total	0.4444	$1.17 \cdot 10^{-8}$

Order of Play

In order to check whether the order of playing Pac-Man games affects the human judgment of interest, we hypothesize that there is no order effect and proceed as follows. For each pair of opponents, that a subject played in both orders, we count (a) the times K that the subject agrees with the interest value only in the first pair played and (b) the times J that the subject agrees with the interest value only in the latter pair played. In the case where the subject agrees or disagrees with the interest value in both pairs played, we take no action. To this end, we compute the z'' value (12) for each pair of opponents ($N = 9$) and in total ($N = 90$).

$$z''(K, J) = (K - J)/N \quad (12)$$

The greater the absolute value of $z''(K, J)$, the more the order of play tends to affect the subjects' judgment of interest. This value defines the test statistic used to assess the truth of the hypothesis that there is no order effect. The obtained z'' value is trinomially distributed.

As seen from Table 4, there are no statistically significant effects in any pair of opponents or in total. Therefore, the null hypothesis is not rejected and it seems that the order of play does not affect the human judgement of interest.

Opponent 1

The order of play is not affected by the particular behavior of opponent 1 either. That is, if $I_1 > I_2$, there is no difference in the obtained $P(Z \geq |z''|)$ values and, if $I_1 = I_2$, there is no statistically significant effects in any pair of opponents or in total (i.e., $P(Z \geq |z''|) = 0.5312$).

Subject Type

In this section, we present how the subject's type, which corresponds to the subject's "liking of the Pac-Man game," correlates with the subject's judgment of interest. To this end, we compute the correlation coefficients c and their respective probabilities $P(C \geq c)$ for each subject type (60 incidents for each type).

TABLE 6 Interest Metric – Subject Judgment Correlation Coefficients c , $P(C \geq c)$ Values of the Three Types of Subject and Correlation Variance (σ_c^2) Over the 10 Subjects of Each Type

Subject type	c	(σ_c^2)	$P(C \geq c)$
Like	0.4000	0.0691	0.0013
Neutral	0.6333	0.1234	0.0067
Don't like	0.4333	0.1493	0.0005
Total	0.3888	0.1079	$1.31 \cdot 10^{-7}$

As seen from Table 6, all three types of subject's observed judgment of interest collectively demonstrate a highly significant agreement ($P < 1\%$) with the interest metric. However, it appears that there is no significant difference between the three types and, therefore, no secure conclusions about the subject's type effect on its notion of interest can be arisen.

Opponent 1

By following the procedure described previously, for the particular case of opponent 1, we also come up with highly significant values for all three subject types and no significant difference between them for both cases of $I_1 > I_2$ and $I_1 = I_2$.

Online Learning

In this section, we analyze the observed effects from the on-line learning experiment (Part C) presented previously. In Part C, subjects play 2 sets of 50 games in total. Interest values calculated (bootstrapping procedure with $N = 25$) and presented in Table 7 show that, in 18 out of 30 cases, the human player managed to produce more interesting games by the use of the online learning procedure. However, it is not clear whether OLL used in humans cause the interest value to proliferate. Thus, it seems that 50 OLL games (testing period in Part A) are not adequate for the OLL mechanism to cause a significant difference in the interest value (see Table 7).

Choosing an online learning period (or else testing period) of 50 games is an empirical way of balancing efficiency and experimental time. The duration of the testing period lasted 20 minutes on average, whereas the whole experiment exceeded 65 minutes in many cases, which is a great amount of time for a human to be constantly concentrated. Fixed strategy *Pac-Man* player results showed that more online learning games are required for the interest value to change significantly, which appears to be the case for human players as well.

By calculating the correlation coefficient (9) between the computed interest values (presented in Table 7) and the human judgment of interest

TABLE 7 Interest I and Confidence Interval (I_u, I_l) Values Against All 30 Human Players Ranked by Subject Type. i.e. 1–10: Like, 11–20: Neutral, 21–30: Don't Like

Subject	OLL			No OLL		
	I_u	I_l	I	I_u	I_l	I
1	0.721	0.575	0.671	0.745	0.393	0.630
2	0.753	0.588	0.669	0.767	0.593	0.703
3	0.733	0.614	0.669	0.755	0.607	0.694
4	0.805	0.672	0.735	0.792	0.520	0.677
5	0.802	0.644	0.711	0.720	0.582	0.665
6	0.763	0.598	0.676	0.733	0.531	0.647
7	0.725	0.638	0.689	0.698	0.559	0.644
8	0.751	0.566	0.673	0.804	0.603	0.720
9	0.746	0.568	0.681	0.751	0.630	0.698
10	0.780	0.531	0.670	0.780	0.616	0.715
11	0.692	0.469	0.619	0.750	0.576	0.695
12	0.802	0.678	0.748	0.865	0.700	0.778
13	0.806	0.530	0.662	0.716	0.532	0.638
14	0.799	0.589	0.715	0.805	0.678	0.738
15	0.776	0.636	0.707	0.782	0.656	0.706
16	0.812	0.658	0.749	0.806	0.689	0.745
17	0.784	0.601	0.706	0.743	0.609	0.679
18	0.796	0.595	0.708	0.740	0.567	0.655
19	0.780	0.612	0.702	0.718	0.626	0.670
20	0.749	0.666	0.717	0.759	0.646	0.716
21	0.753	0.625	0.684	0.757	0.659	0.706
22	0.790	0.660	0.728	0.831	0.625	0.733
23	0.774	0.640	0.709	0.762	0.663	0.700
24	0.752	0.599	0.668	0.754	0.612	0.681
25	0.741	0.635	0.696	0.705	0.589	0.660
26	0.825	0.697	0.770	0.781	0.681	0.728
27	0.799	0.622	0.732	0.782	0.640	0.724
28	0.786	0.630	0.719	0.755	0.570	0.693
29	0.745	0.607	0.690	0.748	0.606	0.705
30	0.793	0.673	0.738	0.782	0.591	0.678
$E\{\}$	0.771	0.614	0.700	0.763	0.605	0.694

obtained the first question of our second objective, we get a value of $c = 0.1666$, with corresponding probability of $P(C \geq c) = 0.1225$ for $N = 60$. This does not constitute a statistically significant effect and suggests that humans were not able to tell the difference between opponent 4 and the opponents trained online at the end of the testing period.

Performance Factor

As noted before, each subject plays eight pairs of sets of games in total during this experiment (six in Part B and two in Part C), and each set is assigned a score that corresponds to the performance of the subject. More

specifically, the score is directly proportional to the number of pellets eaten by the player. Given the subjects' scores and the reported interest judgement, the z' values are computed as follows. If the subject chooses the set of games with the higher score obtained as being more interesting, then the z' value is 1. Accordingly, the z' value is -1 if the subject chooses the set of games with the lower score obtained as being more interesting. By computing (9) for all 30 subjects ($N = 8 \cdot 30 = 240$), we get $c(z') = -0.05$ and $P(C \geq -0.05) = 0.7994$, which constitutes the effect as statistically not significant. Therefore, the null hypothesis H_2 is not rejected and it seems that observed human judgment of interest does not correlate with performance during play.

However, before abandoning the hypothesis of the performance impact on human judgment totally, we attempt to draw the relation between the two from another perspective. Figure 5 illustrates a scatter plot of the correlation coefficients between the performance and the subject's judgement of interest against the correlation coefficients between the interest metric and the subject's judgement of interest for each subject. In addition, the line of the statistical correlation ($(f(x) = cor(c(\vec{z}), c(z')) \cdot x$, where $cor(c(\vec{z}), c(z')) = cov(c(\vec{z}), c(z')) / \sigma_{c(\vec{z})} \sigma_{c(z')} = -0.5864$) between the two samples of data is plotted. If we examine Figure 5 in detail as well as the reported interest criteria in the last question of the survey, there seems to be a classification of the subjects into three groups. These are:




FIGURE 5 Scatter plot of $c(z)$ and $c(z')$ values for each subject and their statistical correlation' line. The circular marker' \leq radius is increased in respect to the number of occurrences (i.e., 1, 2 or 3).

- Subjects that judge interest according to their performance ($c(z') \geq 0.5$), size: 6 out of 30 subjects. As far as their agreement with the interest metric is concerned, their observed judgment portrays a rather random behavior ($0.0 \leq c(z) \leq 0.25$). The reported interest criteria in the last question of the survey are explicit. Randomness and scoring performance are the major criteria in selecting the most interesting set between two.
- Subjects that do not judge interest according to their performance ($c(z') \leq 0.0$) and whose interest judgment correlates with the interest metric ($c(z) \geq 0.5$), size: 12 out of 30 subjects. This group's reported interest criteria are focused on the opponent's contribution to the player's satisfaction.
- Subjects that do not judge interest according to their performance ($-0.5 < c(z') < 0.5$) and whose interest judgment does not correlate with the interest metric ($c(z) < 0.5$), size: 12 out of 30 subjects. Subjects of this group concentrate on a variety of Pac-Man aspects different or implicitly syngenic to the *Ghosts'* behavior, as acquired from the reported interest criteria. These aspects include performance, game control ability, graphics, difficulty, and duration of game.

The computed statistical correlation value and Figure 5 provide evidence that human judgment of interest, that agrees with the interest metric, is not correlated with the human judgment of interest based on performance. In other words, it seems that subjects agreeing with the interest metric do not judge interest by their performance. Or else, subjects disagreeing with the interest metric seem to judge interest by their score and/or other criteria such as game controls and graphics.

Opponent 1

By assuming that $I_1 > I_2^2$, we reveal a slightly higher statistical correlation value $cor(c(\vec{z}), c(z')) = -0.5341$, but conceptually the same effects and subject classification groups as the above-mentioned.

CONCLUSIONS

Predator strategies in predator/prey computer games are still nowadays based on simple rules, which make the game pretty predictable and, therefore, somewhat uninteresting (by the time the player gains more experience and playing skills) (Woodcock 2001; Rabin 2002). A computer game becomes enjoyable primarily when there is a richer online interaction between the player and its opponents who demonstrate interesting behaviors (Yannakakis and Hallam 2004a; 2005c). Machine learning techniques applied online (Stanley et al. 2005) can generate behaviors that give the

illusion of intelligence, which is an important criterion for the human player's perceived entertainment. On top of that, playing against cooperative opponents makes the game more realistic and appealing to the human eye (Yannakakis and Hallam 2005b).

Given some objective criteria for defining interest in predator/prey games, we introduced a generic method for measuring interest in such games. The I metric presented in this article captures the concept of interest objectively and it is dependent on the player-game opponent interaction. We saw that by using the proposed online learning mechanism on the Pac-Man game, maximization of the individual simple distance measure (see (7)) coincides with maximization of the game's interest. Apart from being fairly robust, the proposed approach demonstrated high and fast adaptability to changing types of player (i.e., playing strategies). Results obtained against fixed strategy *Pac-Man* players showed that such a mechanism could be able to produce interesting interactive opponents (i.e., games) against dynamic human playing strategies.

By testing the game against humans, we managed to confirm our hypothesis that the interest value computed by (5) is consistent with the judgment of human players. In fact, human player's notion of interest of the Pac-Man game correlate highly with the captured interest value. However, there are instances where humans' reported notion of interest does not match the respective calculated I value of the game. Since the proposed interest metric was designed and evaluated on computer-controlled *Pac-Man* players, the reported mismatches confirm the fact that a human playing behavior differs from a computer-controlled designed player. In addition, it is revealed that both the subject type (i.e., experience/likeliness with the game) and the order of playing the game do not affect their judgment. Moreover, given each subject's score, it was demonstrated that humans agreeing with the interest metric do not judge interest by their performance. Or else, humans disagreeing with the interest metric judge interest by their score or based on other personal criteria like game control and graphics.

The main assumption drawn for the interest metric proposed is that players overall have a basic level of gaming skills for the test-bed game. In that sense, the computer-guided players used are models of some well-behaved, average skill players based on similar motion patterns that do not leave much space for significant differences in their best generated interest values. Human players that tested Pac-Man cross-validate this assumption since their generated interest values against the same opponent were not significantly different from each other.

As far as online learning against human players is concerned, results show that more online learning games are required for the interest value to change significantly and for humans to notice some sort of change in

the interest of the game. This is a function of the way that human-game interaction is used to train the opponents. Given more computing power, it may be possible to use the data provided by human-game interaction more efficiently and therefore achieve significant change in interest in fewer games. For instance, thousands of mutants could be evaluated in parallel over longer periods—which would provide better behavior estimates—and moreover the frequency of evolutionary iterations could be increased. Using this approach, we could accelerate the learning where appropriate and minimize the probability of unwanted, unrealistic, non-intelligent generated behaviors due to mutation. Clearly, a single unrealistic emerged AI behavior is sufficient to impair the “intelligent” image any adaptive approach is attempting to present and furthermore to diminish the satisfaction of the player (Champandard 2004; Funge 2004).

Discussion

As a novel direction in AI in computer games, cooperative opponents that learn in real time for optimizing the entertainment value of the game constitutes this work’s proposed step for future game development. Showing that such a learning mechanism maintains high levels of player satisfaction makes this approach appealing for application to the vast majority of game genres where online learning and opponent cooperation are, until nowadays, deliberately absent or optional. Here we discuss the potential of the methodology in other genres of multi-opponent games where the complication of the opponents’ tasks may differ. More specifically, we analyze the extensibility of the interest metric proposed, the online evolutionary learning mechanism and the neuro-controller used.

1. *Interest Metric* As already mentioned, the criteria of challenge and behavior’s diversity may be effectively applied for measuring the real-time entertainment value of any genre of games. Spatial diversity may in a sense also contribute to the interest value of specific genres (e.g., team-sport, real-time strategy, and first-person shooter games). As long as game developers can determine and extract the features (e.g., through online questionnaires) of the opponent behavior that generate excitement for the player, a mathematical formula can be designed in order to collectively represent them.
2. *Learning Methodology* The proposed online evolutionary learning method may also be successfully applied to any game during active real-time player-opponent interactions. Extracted features of this interaction may be used in order to estimate the fitness of the involved opponents according to their tasks. The replacement of the worst-fit opponent(s) method may be applied in frequent game periods to enhance the

group's fitness. See also Stanley et al. (2005) for a successful application of this method in the NERO game.

Artificial evolution can explore complex search spaces efficiently and, when combined with NNs, it can demonstrate fast adaptability to dynamic and changing environments. Therefore neuro-evolution is recommended for learning in real time. However, convergence time and unpredictability of the emergent behaviors constitute the disadvantages of the methodology, which can be dealt with by careful design of the learning mechanism. The tradeoff between opponent behavior stability and speed of learning online when using neuro-evolution was addressed in this article. Both the breeding of offspring and the variance of the Gaussian mutation used are inversely proportional to the cell visit entropy. This GA scheme minimizes unpredictability and allows for rapid genotype alterations when the interest value of the game is high and low, respectively.

In addition to a careful GA design, player modeling techniques are able to decrease convergence time down to realistic periods of time (i.e., tens of games) and furthermore proliferate the efficiency and justifiability of learning in real time (Yannakakis and Maragondiakis 2005).

3. *Controller* Artificial neural networks serve successfully the adaptability requirements for predator/prey reactive games in real-time. However, as the complexity of the opponents' tasks increases, there might be a need for more sophisticated structures of distributed representation. Memory of previous behaviors learned through the player-opponent interaction may very well be essential when a combination of various tasks is required. Recurrent NNs or augmented NN topologies with hidden states (Stanley and Miikulainen 2002) may be more appropriate when the opponents' tasks proliferate. Moreover, a hierarchy design of neuro-controllers that serve different opponent tasks could also provide the online learning mechanism with more flexibility and faster adaptability. Decision trees, adaptive scripts (Spronck et al. 2001), or classifier systems (Champanard 2004) could also host adaptive behaviors in real time successfully.

REFERENCES

- Ackley, D. H. and M. L. Littman. 1992. Interactions between learning and evolution. In: *Artificial Life II*, eds. C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, 478–507. Reading, MA: Addison-Wesley.
- Andrade, G., G. Ramalho, H. Santana, and V. Corruble. 2005. Extending reinforcement learning to provide dynamic game balancing. In: *Proceedings of the Workshop on Reasoning, Representation, and Learning in Computer Games, 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 7–12.
- Champanard, A. J. 2004. *AI Game Development*. New Riders Publishing.
- Cole, N., S. J. Louis, and C. Miles. 2004. Using a genetic algorithm to tune first-person shooter bots. In: *Proceedings of the 2004 Congress on Evolutionary Computation*, pages 139–145.

- Csikszentmihalyi, M. 1990. *Flow: The Psychology of Optimal Experience*. New York: Harper & Row.
- Demasi, P. and A. J. de O. Cruz. 2002. On-line coevolution for action games. In: *Proceedings of the 3rd International Conference on Intelligent Games and Simulation (GAME-ON)*, pages 113–120.
- Fogel, D. B. 1993. Using evolutionary programming to construct neural networks that are capable of playing tic-tac-toe. In: *Proceedings of the IEEE International Conference on Neural Networks*, page 875–880, San Francisco, CA, USA: IEEE Press.
- Fogel, D. B. 2002. *Blondie 24: Playing at the Edge of AI*. Morgan Kaufmann.
- Fogel, D. B., T. J. Hays, and D. R. Johnson, 2004. A platform for evolving characters in competitive games. In: *Proceedings of the Congress on Evolutionary Computation (CEC-04)*, pages 1420–1426, June 2004.
- Funge, J. D. 2004. *Artificial Intelligence for Computer Games*. Wellesley, MA: A. K. Peters Ltd.
- Gallagher, M. and A. Ryan. 2003. Learning to play Pac-man: An evolutionary, rule-based approach. In: *Proceedings of the Congress on Evolutionary Computation (CEC)*, pages 2462–2469.
- Graepel, T., R. Herbrich, and J. Gold. 2004. Learning to fight. In: *Proceedings of the International Conference on Computer Games: Artificial Intelligence, Design and Education*, pages 193–200, Reading, UK.
- Haynes, T. and S. Sen. 1995. Evolving behavioral strategies in predators and prey. In: *IJCAI-95 Workshop on Adaptation and Learning in Multiagent Systems*, pages 32–37, Montreal, Quebec, Canada: Morgan Kaufmann.
- Hunicke, R. and V. Chapman. 2004. AI for dynamic difficulty adjustment in games. In: *Proceedings of the Challenges in Game AI Workshop, 19th Nineteenth National Conference on Artificial Intelligence (AAAI'04)*, pages 91–96.
- Iida, H., N. Takeshita, and J. Yoshimura. 2003. A metric for entertainment of boardgames: Its implication for evolution of chess variants. In: *IWEC2002 Proceedings*, pages 65–72, Kluwer.
- Isla, D. and B. Blumberg. 2002. New challenges for character-based AI for games. In: *Proceedings of the AAAI Spring Symposium on AI and Interactive Entertainment*, pages 41–45, Stanford, CA, USA: AAAI Press.
- Johnson, S. 2004. *Adaptive AI*. Hingham, MA: Charles River Media.
- Kaiser, S. and T. Wehrle. 1996. Situated emotional problem solving in interactive computer games. In: *Proceedings of the VIXth Conference of the International Society for Research on Emotions*, pages 276–280, ISRE Publications.
- Kaiser, S., T. Wehrle, and S. Schmidt. 1998. Emotional episodes, facial expressions, and reported feelings in human computer interactions. In: *Proceedings of the Xth Conference of the International Society for Research on Emotions*, pages 82–86, ISRE Publications.
- Khoo, A. and R. Zubeck. 2002. Applying inexpensive techniques to computer games. *IEEE Intelligent Systems*, 17(4): 48–53.
- Koster, R. 2005. *A Theory of Fun for Game Design*. Scottsdale, AZ: Paraglyph Press.
- Koza, J. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press.
- Laird, J. E. 2002. Research in human-level AI using computer games. *Communications of the ACM* 3(8): 32–35.
- Laird, J. E. and M. van Lent. 2000. Human-level AI's killer application: Interactive computer games. In: *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI)*, pages 1171–1178, Saint Paul, MN.
- Lazzaro, N. 2004. *Why We Play Games: Four Keys to More Emotion Without Story*. XEO Design Inc., Technical Report.
- Lucas, S. 2005. Evolving a neural network location evaluator to play Ms. Pac-Man. In: *Proceedings of the IEEE Symposium on Computational Intelligence and Games*, pages 203–210, Cambridge, MA: Colchester, UK.
- Luke, S. and L. Spector. 1996. Evolving teamwork and coordination with genetic programming. In: *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 150–156, Palo Alto, California, USA: MIT Press.
- Malone, T. W. 1981. What makes computer games fun. *Byte* 6:258–277.
- McQuiggan, S., S. Lee, and J. Lester. 2006. Predicting user physiological response for interactive environments: An inductive approach. In: *Proceedings of the 2nd Artificial Intelligence for Interactive Digital Entertainment Conference*, pages 60–65.

- Miller, G. and D. Cliff. 1994. Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In: *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior (SAB-94)*, pages 411–420, MIT Press.
- Ponsen, M. and P. Spronck. 2004. “Improving adaptive game AI with evolutionary learning. In: *Proceedings of the International Conference on Computer Games: Artificial Intelligence, Design and Education*, pages 389–396, Reading, UK.
- Rabin, S. 2002. *AI Game Programming Wisdom*.: Charles River Media, Inc.
- Rani, P., N. Sarkar, and C. Liu. 2005. Maintaining optimal challenge in computer games through real-time physiological feedback. In *Proceedings of the 11th International Conference on Human Computer Interaction*, pages 184–192.
- Read, J., S. MacFarlane, and C. Cassey. 2002. Endurability, engagement and expectations. In *Proceedings of International Conference for Interaction Design and Children*, ACM Press.
- Richards, N., D. E. Moriarty, and R. Miikulainen. 1998. Evolving neural networks to play go. *Appl. Intell.* 8(1):85–96.
- Rosca, J. 1996. Generality versus size in genetic programming. In *Genetic Programming 1996: Proceedings of the First Annual Conference*, pages 381–387, Palo Alto, California, USA: MIT Press.
- Spronck, P., I. Sprinkhuizen-Kuyper, and E. Postma. 2004. Difficulty scaling of game AI. In: *Proceedings of the 5th International Conference on Intelligent Games and Simulation (GAME-ON 2004)*, pages 33–37.
- Stanley, K., Bryant, B., and R. Miikulainen. 2005. Real-time evolution in the NERO video game. In: *Proceedings of the IEEE Symposium on Computational Intelligence and Games*, pages 182–189, Colchester, UK.
- Stanley, K. and R. Miikulainen. 2002. Evolving neural networks through augmenting topologies: *Evolutionary Computation* 10(2):99–127.
- Sweeter, P. and P. Wyeth. 2005. GameFlow: A model for evaluating player enjoyment in games. *ACM Computers in Entertainment* 3(3).
- Taylor, T. 2000. “Artificial life techniques for generating controllers for physically modelled characters.” In: *Proceedings of the First International Conference on Intelligent Games and Simulation (GAME-ON 2000)*, ■.
- Tesauro, G. 2002. Programming backgammon using self-teaching neural nets. *Artificial Intelligence* 134:181–199.
- Thurau, 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)*, pages 315–323, Santa Monica, CA: The MIT Press.
- Verma, M. A. and P. W. McOwan. 2005. An adaptive methodology for synthesising mobile phone games using genetic algorithms. In: *Proceedings of the Congress on Evolutionary Computation (CEC-05)*, pages 528–535, Edinburgh, UK.
- Weizenbaum, J. 1966. ELIZA—a computer program for the study of natural language communications between men and machines. *Communications of the Association for Computing Machinery* 9:36–45.
- Wikipedia, the Free Encyclopedia. 2005. Pac-Man, [Online]. Available: <http://en.wikipedia.org/wiki/Pac-man>
- Woodcock, S. 2001. *Game AI: The State of the Industry 2000–2001: It’s not Just Art, It’s Engineering*.
- Yannakakis, G. N. 2005. *AI in Computer Games: Generating Interesting Interactive Opponents by the use of Evolutionary Computation*. Ph.D. thesis, University of Edinburgh.
- Yannakakis, G. N. and J. Hallam. 2004a. “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)*, pages 499–508, Santa Monica, CA, USA: The MIT Press.
- Yannakakis, G. N. and J. Hallam. 2004b. *Interactive opponents generate interesting games*. In: *Proceedings of the International Conference on Computer Games: Artificial Intelligence, Design and Education*, pages 240–247, Reading, UK.
- Yannakakis, G. N. and J. Hallam. 2005a. A generic approach for generating interesting interactive Pac-man opponents. In: *Proceedings of the IEEE Symposium on Computational Intelligence and Games*, pages 94–101, Colchester, UK.
- Yannakakis, G. N. and J. Hallam. 2005b. A generic approach for obtaining higher entertainment in predator/prey computer games. *Journal of Game Development* 1(3):23–50.

- Yannakakis, G. N. and J. Hallam. 2005c. A scheme for creating digital entertainment with substance. In: *Proceedings of the Workshop on Reasoning, Representation, and Learning in Computer Games, 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 119–124.
- Yannakakis, G. N. and J. Hallam. 2006. Towards capturing and enhancing entertainment in computer games. In: *Proceedings of the 4th Hellenic Conference on Artificial Intelligence, Lecture Notes in Artificial Intelligence*, vol. 3955, pages 432–442 Heraklion, Greece: Springer-Verlag.
- Yannakakis, G. N. and M. Maragoudakis. 2005. Player modeling impact on player's entertainment in computer games. In *Proceedings of the 10th International Conference on User Modeling; Lecture Notes in Computer Science*, vol. 3538, pages 74–78. Edinburgh: Springer-Verlag.
- Yannakakis, G. N., J. Hallam, and H. H. Lund. 2006. Capturing entertainment through heart-rate dynamics in the playware playground. In *Proceedings of the 5th International Conference on Entertainment Computing, Lecture Notes in Computer Science*, vol. 4161, pages 314–317 Cambridge, UK: Springer-Verlag.
- Yannakakis, G. N., J. Hallam, and H. H. Lund. 2006. Comparative fun analysis in the innovative playware game platform. In *Proceedings of the 1st World Conference for Fun'n Games*, pages 64–70, Preston, UK.
- 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)*, pages 986–993.
- Yao, X. 1999. Evolving artificial neural networks. *Proceedings of the IEEE* 87(9):1423–1447.

NOTES

1. Further details of this strategy are presented in Yannakakis and Hallam (2004).
2. The case of $I_1 = I_2$ is not investigated since the 1–2 pair is not taken into consideration and z' values cannot be computed for subjects that played that particular pair of sets.