

Virtual player design using self-learning via competitive coevolutionary algorithms

Mariela Nogueira Collazo · Carlos Cotta ·
Antonio J. Fernández-Leiva

Published online: 31 January 2014
© Springer Science+Business Media Dordrecht 2014

Abstract The Google Artificial Intelligence (AI) Challenge is an international contest the objective of which is to program the AI in a two-player real time strategy (RTS) game. This AI is an autonomous computer program that governs the actions that one of the two players executes during the game according to the state of play. The entries are evaluated via a competition mechanism consisting of two-player rounds where each entry is tested against others. This paper describes the use of competitive coevolutionary (CC) algorithms for the automatic generation of winning game strategies in Planet Wars, the RTS game associated with the 2010 contest. Three different versions of a prime algorithm have been tested. Their common nexus is not only the use of a Hall-of-Fame (HoF) to keep note of the winners of past coevolutions but also the employment of an archive of experienced players, termed the hall-of-celebrities (HoC), that puts pressure on the optimization process and guides the search to increase the strength of the solutions; their differences come from the periodical updating of the HoF on the basis of quality and diversity metrics. The goal is to optimize the AI by means of a self-learning process guided by coevolutionary search and competitive evaluation. An empirical study on the performance of a number of variants of the proposed algorithms is described and a statistical analysis of the results is conducted. In

addition to the attainment of competitive bots we also conclude that the incorporation of the HoC inside the primary algorithm helps to reduce the effects of cycling caused by the use of HoF in CC algorithms.

Keywords Coevolution · Competition · Self-learning · RTS game · Virtual player

1 Introduction

The videogame sector represents the largest of the entertainment industries and is an area that is constantly evolving. The main purpose of videogames is to provide entertainment to the player(s) but how to infact do this, is an open question that has yet to be fully answered. In the past, game developers primarily concentrated on having more realistic games and worked on implementing games with high quality graphics (i.e., having high resolution textures, a good measure of frames-per-second, etc.). Whilst, until recently, this policy guaranteed reasonable profits for the development company, in the last decade it has not been enough to ensure the success of a game, as players now demand other features for their games. Many of these required features are associated with the resolution of problems that demand knowledge from a wide set of research domains such as art, psychology, narrative, or music to name but a few.

In this heterogeneous context, one paradigm that is present in all (or in most) videogames is artificial intelligence (AI). Currently, it is applied to most aspects of game development and design such as learning, human player imitation, procedural content generation (PCG), intelligent camera control, automatic game testing, player/opponent modeling, and computational narrative, among others

M. Nogueira Collazo
University of Computer Science, Havana, Cuba
e-mail: mnogueira@uci.cu

C. Cotta (✉) · A. J. Fernández-Leiva
University of Málaga, Málaga, Spain
e-mail: ccottap@lcc.uma.es

A. J. Fernández-Leiva
e-mail: afdez@lcc.uma.es

(Lucas 2009). The application of AI, and computational intelligence (as the AI representative of the nature-inspired computational techniques for learning and optimization such as evolutionary algorithms, artificial neural networks, swarm intelligence and ant-colony optimization methods, to give a few examples) in games is thus an active research field that poses significant challenges to the game development community (Lucas et al. 2012). Nonetheless, the most traditional application of AI in games is to govern the behavior of the non-player-characters (NPCs) in the game with the aim of endowing intelligence on the enemy and as a consequence increasing the satisfaction of the player who demands an opponent who exhibits intelligent behavior. However, generating an AI to control NPCs is a hard task—traditionally tackled by hand—that also requires a quite a large dose of patience. For instance, in triple-A videogames the cost of developing the AI is very high due to the huge number of possible situations in which the NPC might be at a given instant of the game and to the fact that in many commercial games the AI that controls an NPC is still made up of very specialized scripts (previously programmed based on the experience of the designers/programmers and via some traditional AI technique—e.g., finite state machines, of fuzzy logic). These are usually at the root of very well-known problems associated with “artificial stupidity” (Lidén 2004; e.g. loss of reality, or the predictability and the existence of holes/bugs in the NPC behavior).

It is well-known that a NPC behavior which is too stupid is not desirable as this always leads to a very easy victory for the player thus decreasing his/her satisfaction; an analogous argument can be applied to a very specialized game AI as the player would then most probably suffer a heavy defeat with the same consequence. As an intermediate solution, commercial videogames offer different difficulty levels to the player so that in the most complex levels the player faces high-quality (again pre-programmed) difficult-to-beat opponents. Once the player is able to beat all the opponents in each level, they lose interest. In this context, the generation of opponents whose behavior evolves in accordance with the player’s increasing abilities is an appealing feature that makes the game more attractive. In the literature can be found many proposals for generating NPCs whose behaviors self-adapt to player skills (Szita et al. 2009; Gutiérrez et al. 2011).

One of the most interesting game research problems is that of developing AI for real-time strategy (RTS) games. In general, in an RTS game the players make decisions in a team (or army) consisting of a number of (not necessarily equal) units (e.g., soldiers, workers,...) and each player has two main objectives: (a) to destroy the opponent’s assets and (b) to create additional structures with some specific goal (e.g., construct buildings to protect their units or

defend a specific position in the map). The objectives have to be achieved with an initial limited number of resources because during the game it is possible to obtain more resources. So an RTS game can be viewed as a resource gathering game. Here, game AI means to define the behavior of a virtual player that controls one of the teams in the game. These features and the complexity of the search space (states describe large playing scenarios with hundreds of units simultaneously acting, and the environments consist of many thousands of possible positions for each of the hundreds, possibly thousands, of units) make the design of game AI a very interesting challenge. However this is at the same time, a very hard to handle task due to the specific problems caused by the large search spaces and the parallel nature of the problem—unlike traditional games, any number of moves may be made simultaneously (Corruble et al. 2002).

RTS games provide a range of challenging problems for AI design, e.g., planning in an uncertain world with incomplete information, learning, opponent modeling, and spatial and temporal reasoning (Buro 2004) just to name a few. One of the most significant challenges is to provide non-cheating and human-like virtual players. To quote Lucas et al. (2012), this basically involves “restricting the information available to the AI player to the information a human player may be able to gather in the game, and restricting the actions of the AI player to human player actions (executed in time and space”. Additionally, the virtual player should pose a challenge to the human players independently of their skills and of the strategies they adopt to play.

This paper deals with the automatic generation of self-adaptive AI for NPCs in RTS games. This can be catalogued as a form of PCG and combines the advantages of PCG and self-adaptation as for instance level adjustment to player skills and the possibility of producing endless games among others (Togelius et al. 2011). In particular, we consider coevolution, a biologically inspired technique based on the interaction between different species which in our opinion represents one of the most interesting approaches to be exploited in the evolutionary programming area.

Coevolutionary systems are usually based on two kinds of interactions, akin to symbiotic and predator/prey relationships in nature. The former is a cooperative approach in which an individual is typically decomposed in to different components that evolve simultaneously and the fitness of which depends on the interaction between these components; the latter is a competition-based approach in which an individual competes with other individuals for the fitness value and, if appropriate, will increase its fitness at the expense of its counterparts, whose fitnesses decrease. As in predator/prey relationships in nature, this second approach

is prone to trigger “army races” in which the improvement of some individuals stimulates the improvement in the opponents, and vice versa.

This technique has been shown to be successful in a number of applications but it also has a number of drawbacks (Ficici and Bucci 2007). Many approaches, e.g. Ebner et al. (2010), center on the analysis of the dynamics of the coevolutionary process with the aim of identifying both its weaknesses and strengths and to produce more solid techniques for coevolution support. Miconi (2009) also underlines the importance of the terms superiority and progress, to avoid fails in the coevolutionary search process. The first remedy to the inherent pathologies of coevolution consisted of proposing some forms of evaluating individuals (Rosin and Belew 1995), and the memorization of a number of successful solutions to guide the search. Following this idea, Rosin and Belew (1997) proposed the use of a *Hall-of-Fame* (HoF) based mechanism as an archive method. It acts as a long-term memory mechanism in competitive coevolutionary (CC) algorithms for managing the historical set of champions during the individuals’ evaluation.

The work presented here, builds on our previous work (Nogueira et al. 2012, 2013) on testing the use of HoF-based CC (HoFCC) algorithms for finding winning strategies in RTS games. In our previous papers we analyzed how the diversity and the growth of the HoF can influence the quality of the solutions obtained by HoFCC algorithms. This was done in the context of the RTS game *RobotWars*,¹ a self-developed game the simplicity and inherent limitations of which handicapped the scope of our experimentation; the CC algorithms that imposed certain diversity and tried to maintain a manageable size of the HoF (by removing those champions not contributing to the optimization process) demonstrated a better performance than the others that did not work on the two metrics considered but they still suffered from the appearance of cycling. This is a well-known problem of CC algorithms that work with ‘archive’ structures such as the HoF. Now we considerably extend our previous work by considering a new RTS game—namely *Planet Wars*, the *Google AI challenge in 2010*—that allows a deeper experimental analysis and therefore provides more consistent conclusions. We also propose a different evaluation mechanism to exploit the potential offered by archive methods to maintain transitivity between the solutions. Moreover we add novel strength indicators which are independent from the fitness function with the objective of avoiding the appearance of cycling. The novelty of this last aspect consists of incorporating into our prime CC algorithm, an additional archive (termed hall-of-celebrities, HoC) that contains a

team of experienced virtual players that are used to evaluate how strong a candidate is. The combined use of both halls (HoF and HoC) with the (possibly combined) employment of diversity and quality metrics helps the optimization to obtain competitive bots that self-adapt to beat their (co)evolved enemies.

2 Background

This section discusses a number of approaches that are related to competitive coevolution applied in games (Sect. 2.1), and describes the RTS game used in our experimental section (Sect. 2.2).

2.1 (Competitive) coevolution in games

Due to the intrinsically competitive nature of the games, many researchers have opted for the application of a competitive coevolution approach to solve searching and self-learning problems in games, e.g., the study described in Angeline and Pollack (1993) on competitive fitness functions in the Tic Tac Toe game, the application of simple competitive models for evolving strategies in a pursuit-evasion game (Reynolds 1994), or the evolution of both morphology and behaviors of artificial creatures through competition in a predator–prey environment (Sims 1994). Also Ashlock et al. (2012) analyzed the level of difficulty involved in finding solutions with a basic coevolutionary algorithm for zero-sum games as well as non-zero-sum games; a fitness metric with the score obtained by an agent when faced with another one belonging to the same generation was employed; the authors detected that co-evolving good strategies for zero-sum games is more difficult than for non-zero-sum games.

Competitive coevolution has been used heavily in complex scenarios such as those that emerge in strategy games; so, Smith et al. (2010) developed coevolved artificial intelligent opponents with the objective of training human players in the context of a capture-the-flag game. Also, Avery and Louis (2010) analyzed the use of coevolution for creating a tactical controller for small groups of game entities in a real-time capture-the-flag game; a representation for generating adaptive tactics using coevolved Influence Maps was proposed, and the result was the attainment of an autonomous entity that plays in coordination with the rest of the team to achieve the team objectives. Related with this line of research, Miles and Louis (2006) presented a spatial decision making system within the context of a 3D computer RTS game, that used a basic implementation of co-evolution in which players were firstly evolved against static hand-coded opponents and later against another population of co-evolving players.

¹ <http://wp.me/p2cObl-60>.

More recently, Dziuk and Miikkulainen (2011) has explored several methods for automatically shaping the coevolutionary process, and this is done by modifying the fitness function as well as the environment during evolution. Another interesting perspective was presented in Avery and Michalewicz (2007) where the authors, using the game of Tempo as a test space, facilitated the selection of optimal strategies by clustering the solutions in the population of a coevolutionary system through the concept of similarity. This cluster system integrated a long-term memory that valued the changes produced in the environment to trigger appropriate coevolution. The game of Tempo has also been used with the aim of improving the creation of smart agents in Johnson et al. (2004) and Avery et al. (2008).

More closely related to our work, we can cite several articles that address the use of archive methods as an alternative to avoid the effect of pathologies, and that present interesting proposals for optimizing their performance. For example, Nerome et al. (1998) propose a competitive co-evolution model that introduces the concept of package as a set of good strategies and the best package is the one that contains the smallest number of strategies providing, at the same time, the highest number of victories. Different forms of archives, like the layered Pareto-coevolution archive (de Jong 2004) and the coordinate system archive (Jaskowski and Krawiec 2010) have also been proposed. However, one can find other simpler structures like for instance the use of one simple champion memory that can be improved via different proposals. Lichocki (2008) describes three useful extensions of the HoF that include uniqueness, manual teachers, and competitive fitness sharing, CFS (Bosin and Belew 1997); the results showed that HoF works better than single elimination tournament (Angeline and Pollack 1993) but this method was not able to prevent the lack of diversity in the population. In Livingstone (2005) the author presents a system in which different levels of a hierarchical AI coevolve in a simple RTS game environment; archive methods were applied to avoid some of the pathologies of coevolution using fixed, non-evolving opponents during the evaluation process. More recently Samothrakis et al. (2013) presents a set of measurements to identify cycling in a population, and proposes an algorithm that minimizes the effect of cycling in a coevolutionary system for the game Othello; in the experiments the authors used different algorithms for generating artificial players, including one that uses a simple archive method with fixed size, which obtained good results.

According to Avery et al. (2008), the question of how to actually use the memory in the coevolution tends to be split into two areas: inserting individuals from memory into the coevolution, or evaluating individuals from the populations against the memory. Our investigation falls comes under

the latter, and we have implemented different variants of HoF for controlling evaluation process, in a CC algorithm.


However, although the contribution of the “archive methods” to reduce the number of drawbacks (e.g., missing features and cycling) of coevolutionary systems is undeniable, their implementation entails a number of issues that have to be taken into account such as the strength of the archive members and the control of the archive diversity. So for instance, a high risk of generating cycles can appear if the strength of a member is valued by pitting this member against a fixed set of opponents, and this is precisely what we did in our previous work (Nogueira et al. 2012, 2013). A key point to consider is the metric that defines the strength of an individual, which should indicate when an individual is qualitatively superior to the rest but without losing sight of the fact that sometimes quantitative superiority does not correspond with a high strength.

This paper addresses precisely this issue and proposes the consideration of an archive of experts in the evaluation of individuals, in addition to the classic archive of champions. This expert archive will include a fixed set of high quality opponents that will be used to provide an additional evaluation (to that already provided by the use of the champion memory, i.e., the HoF) to measure the quality of the individual; for our particular RTS game and this paper, this expert set is conformed by the collection of built-in bots provided by the *Google AI Challenge 2010* competition plus an optimized genetic algorithm (GA)-based bot from Fernández-Ares et al. (2011). The idea of adding an expert archive for the assessment of candidate fitness is to reduce the likelihood of cycling during coevolution. The fitness of an individual will be calculated by pitting it against its ancestral opponents, but individuals with better fitness will be up against one of the experts’ bot (chosen at random) in various scenarios (i.e., different maps); only in case that this individual beats them three times will it be included in the champions’ memory. In doing this, we are trying to exploit the potential offered by archive methods to maintain transitivity between the solutions, but adding strength indicators which are independent from the fitness function with the idea of preventing the appearance of cycling.

Another important aspect to consider is the control of diversity to avoid the redundancy in the memory of champions (which may also combat the over-specialization); we also consider this issue here and propose an algorithm (HoFCC-diversity) that updates the archive from a diversity metric and eliminates those solutions identified as “less diverse”. The paper also proposes a multi-objective HoFCC algorithm based on both metrics: strength and diversity, different variants from this algorithm are compared empirically.

Note that this section does not mention other forms of coevolution that have also been applied in games because

Fig. 1 Screenshot of Planet Wars: both planets and fleets are identified by a *number* which shows the number of ships that they have. Their *colors* also identify the player who owns them. (Color figure online)



they are not directly related with competition (e.g., Cook et al. 2012 focused on cooperation and not on competition).

2.2 Planet Wars

Planet Wars was the game used in Google AI Challenge 2010,² this contest was run by the University of Waterloo Computer Science Club and supported by Google. The game takes place on a map which contains several planets, each of which has a specific number of ships (*NS*) on it. The planets may belong to the player, the opponent, or just be neutral. Each planet owned by a player (not those that are neutral) will increase the forces there according to the “growth rate” of the planet. This rate indicates the ratio of growth to the *NS* on the planet. Figure 1 shows a screenshot of this game, in which the red planets belongs to the player, the green to the enemy, and the rest are neutral.

At the start of each turn, the player receives the current status of the game (i.e., information about the planets and fleets) and can only do one type of action: send fleets of ships from any planet the player owns (i.e., those planets where the *NS* is higher than 0 and the ships belong to the player) to any other planet on the map. Players may send as many fleets as they wish in a single turn as long as they own enough ships. After sending fleets, each planet owned by a player will increase the *NS* staying there proportionally to the planet growth rate. The fleets that were sent in a previous order take a number of turns to reach their destination according to the distance between the planet of origin and the destination one; upon the arrival of the ships if both the planet and the ships belong to the same player, then the *NS* increases by adding up the current *NS* on the planet and the number of newly arrived ships. Otherwise, if the arrival planet is neutral, then it has a fixed *NS* and the player must be sent at least $NS + 1$ ships to own the planet (i.e. for

reaching the neutral planet in order to conquer it); and if the player sends fleets to an enemy planet (i.e. the player attacks an enemy planet) a fight is initiated for the control of the planet (ships from both sides destroy each other until the player with the highest *NS* owns the planet, in this case the destination planet does not have a fixed *NS*, because he/she increments his/her force according his/her growth rate).

Fleets cannot be redirected once they start their journey. Players may continue to send more fleets in later turns even while older fleets are in transit. Although the players issue their orders on a turn-by-turn basis, they carry out these orders at the same time, so we can treat this game as a real-time one.

The player with the most ships at the end of the game wins. The game may also end earlier if one of the players loses all his/her ships and in this case the player that has ships remaining wins instantly; also if one player exceeds the time limit without completing his/her orders she forfeits the game. If both players have the same *NS* when the game ends this is considered a draw.

3 Competitive coevolution for self-learning in Planet Wars

This section describes our CC approaches. Section 3.1 centers on codification issues for the representation of a game strategy, and Sect. 3.2 details our primary CC algorithm that employs the HoF as a long-term memory mechanism, and the HoC as the expert archive for evaluating the strength of individuals; three variants of this algorithm are also explained.

3.1 Representation issues

Our goal is to find winning strategies, governing a bot (i.e., the virtual player) in Planet Wars by optimizing the rules

² <http://planetwars.aichallenge.org>.

that guide its decision making so that in each turn the bot should select the best action according to its current state in the game. Here the state at a specific instance of the game is determined by the advantage (with respect to its opponent) of the bot in terms of ships and growth rate. Both advantage metrics have three possible values ('high', 'null' or 'low') which indicate the level of advantage in each case. Lets say that, for player p (resp. opponent o), GR_p (resp. GR_o) is the total "growth rate", according to the "growth rate" of the planets owned by player p , and let $\Delta GR_{po} = GR_p - GR_o$; then if $\Delta GR_{op} > 0$, we say that player p has a 'high' advantage over opponent o in terms of the "growth rate"; 'null' in the case of a draw (i.e., $\Delta GR_{po} = 0$), and 'low' if $\Delta GR_{po} < 0$. The calculation of the advantage in terms of ships is similar but considering NS_p instead of GR_p (resp. NS_o instead of GR_o) as the total NS owned by player p (resp. opponent o) at the current instant of the game; in other words, $\Delta NS_{po} = NS_p - NS_o$. However, we considered different thresholds for the advantage values so that if $\Delta NS_{po} > 10$, then player p has a 'high' advantage, 'null' if $0 \leq \Delta NS_{po} \leq 10$, and 'low' otherwise (note that we consider that a difference of 10 ships is not significant enough to distinguish between the two players; this value was arrived at through our experience of playing).

So, a virtual player strategy is coded as a bidimensional matrix (see Table 1), where the first dimension symbolizes the player's advantage over his/her opponent in terms of the NS (i.e. the ΔNS_{po}), and the second dimension represents the advantage in terms of the total "growth rate" (i.e. the ΔGR_{op}), and as we explained above each axis has three possible values ('high', 'null' or 'low'). Each cell in the matrix acts as a gene and stores one of the next possible actions:

- (1) attack the strongest enemy planet (AS) (i.e., the enemy planet that owns the highest NS),
- (2) attack the weakest enemy planet (AW),
- (3) attack the closest enemy planet (AC),
- (4) conquer the strongest neutral planet (CS; again in terms of NS),
- (5) conquer the weakest neutral planet (CW),
- (6) conquer the closest neutral planet (CC),
- (7) follow the enemy (FE; in this action the order is to send fleets of ships to the planet to which the enemy is now sending his/her own fleets).

Table 1 A matrix of actions which was generated by our coevolutionary process

	$\Delta NS_{po} > 10$	$0 \leq \Delta NS_{po} \leq 10$	$\Delta NS_{po} < 0$
$\Delta GR_{op} > 0$	AS	AC	AW
$\Delta GR_{op} = 0$	AS	AC	CW
$\Delta GR_{op} < 0$	FE	AS	AW

This way, the whole matrix represents a strategy that controls, deterministically, the behavior of a bot during the game by executing the action associated with a specific instance of the game. For a virtual player there are nine possible different states (i.e. 3×3 , all the possible value combinations considering the two dimensions of the matrix). And basically, in a specific turn of the game the player will execute the action stored in the state in which he perceives that he is. See Table 1 for an example of the actions' matrix which shows the distribution of the actions for each possible state, e.g. when the player has a high advantage with respect to his/her opponent in terms of "growth rate" and "ships" ($\Delta GR_{op} > 0$ and $\Delta NS_{op} > 10$) the selected action will be "attack the strongest enemy planet (AS)". Note that the search space is $7^9 = 40,353,607 \in [2^{25}, 2^{26}]$, which cannot be exhaustively assessed due to the cost of the evaluation that requires a game simulation and thus metaheuristic techniques are used.

3.2 HoFCC algorithm and variants

Using Planet Wars we propose, in this paper, a modified and improved version of a CC algorithm that we already described in Nogueira et al. (2012, 2013) for a simplistic capture-the-flag game. The latter used the HoF as a long-term memory mechanism to keep the winning strategies found in each coevolutionary step and all of them were also used in the evaluation process (in the basic algorithm). In fact, the best individual from each coevolutionary iteration is retained for future testing, and so we obtain an historical set of champions which is used in the evaluation of the individual. Each population maintains its own HoF, in which its own winning strategies (with respect to the set of winning strategies of its opponent) found in each coevolutionary step will be saved. So, in the coevolutionary step n each possible solution of army A again fights with each solution in $\{B_1, B_2, \dots, B_{n-1}\}$, where B_i is the champion found by army B in the i th (for $1 \leq i \leq n - 1$) coevolutionary step.

Regarding the use and implementation of HoF some aspects must be defined. The first is the criteria for inserting a new member in the memory. In this paper we have changed our previous proposal (Nogueira et al. 2013) and we have set that an individual must defeat all the HoF members (i.e., all champions of the opponent population), and has to simultaneously beat an expert bot belonging to an archive of efficient (possibly hand-coded) bots; we call this archive HoC.

We have then considered different policies for maintaining the champions in the set. For this issue, we take into account the contribution of the individual (i.e., the champion) to the search process as, for instance, it might be the case that some opponents that belong to very old generations do not show a good performance in comparison with opponents generated in recent generations and thus they

might be easily beaten; it is therefore crucial to remove those champions that do not contribute to the solution which, in other words, represents a mechanism to control the size of the champions' memory, and is an important element because taking into consideration all the champions (during the evaluation process) might produce more consistent solutions at the expense of a very high computational cost (note that many simulations of the match must be executed for each champion involved in the evaluation; we will provide more details on this further on).

In what follows we present our new proposal for a prime HoFCC algorithm and three variants of it that differ precisely in the policy of establishing the aspects mentioned previously. From a general point of view, basically the variants differ from each other in the way that they periodically update the solution set on the basis of quality and diversity metrics.

3.3 Basic HoFCC

Algorithm 1 shows the schema of our primary algorithm HoFCC (a revised and improved version of that presented in Nogueira et al. 2013). The algorithm has two parameters: α that indicate (in percentage) the portion of the HoF to be removed with the aim of maintaining only those worthwhile champions that might contribute to the solution according to some specific metric (see details below) and λ the frequency of executing this updating of the *HoF*.

A specific strategy is considered 'winning' if it achieves a certain score when it deals with the strategies belonging

to the set of winning strategies of its opponent (i.e., the enemy HoF), and it defeats an expert bot belonging to the HoC. The initial objective is to find a winning strategy of player 1 with respect to player 2 (i.e. the initial opponent) so that the HoF of player 2 is initially loaded with some strategies (randomly or manually initialized: line 2). The HoC is also loaded with a set of (other demonstrated) efficient virtual players that will be used to evaluate the strength of the solutions (line 2).

The HoF of the player being evolved is updated (i.e., only those robust champions are kept) according to some specific criteria every λ coevolutionary step (lines 5–7). Then a standard evolutionary process tries to find a strategy for player 1 that can be considered victorious (lines 12–17). A strategy is considered winning if its fitness value is above a certain threshold value ϕ (line 18; which indicates that this strategy has defeated the members of the opponent HoF) and at the same time it has been able to defeat one of the (randomly chosen) celebrities (i.e., 'experts') in three consecutive battles on different maps (i.e., scenarios) of the game; if they are indeed successful, this strategy is added to the HoF of player 1 (line 20) and the process is initiated again but with the player roles reversed (line 21); otherwise (i.e. no winning strategy is found) the search process is restarted. If after a number of coevolutionary steps no winning strategy is found the search is considered to have stagnated and the coevolution ends (see the while condition in line 4). At the end of the entire process we obtain as a result two sets of winning strategies associated accordingly with each of the populations.

Algorithm 1: HoFCC(α, λ)

```

1   $nCoev \leftarrow 0$ ;  $A \leftarrow player_1$ ;  $B \leftarrow player_2$ ;  $\phi \leftarrow thresholdvalue$ ;
2   $HoF_A \leftarrow \emptyset$ ;  $HoF_B \leftarrow INITIALOPPONENT()$ ;  $HoC \leftarrow EXPERTBOTSET()$ ;
3   $pop \leftarrow EVALUATE(HoF_B)$ ; // Evaluate initial population
4  while  $nCoev < MaxFailCoevolutions \wedge NOT(timeout)$  do
5      if  $nCoev \bmod \lambda = 0$  then //HoF updating every  $\lambda$  Coevolutions
6          |  $PURGE_\alpha(HoF_A)$ ; // HoF updating
7      end if
8       $pop \leftarrow RANDOMSOLUTIONS()$ ; // pop randomly initialized
9       $i \leftarrow 0$ ;
10      $e \leftarrow RANDELEMENTFROM(HoC)$ ;
11     while  $(i < MaxGen) \wedge (fitness(best(pop)) < \phi) \wedge (beatExpert(best(pop), e))$  do
12         |  $parents \leftarrow SELECT(pop)$ ;
13         |  $childs \leftarrow RECOMBINE(parents, p_X)$ ;
14         |  $childs \leftarrow MUTATE(childs, p_M)$ ;
15         |  $pop \leftarrow REPLACE(childs)$ ;
16         |  $pop \leftarrow EVALUATE(HoF_B)$ ;
17         | if  $(fitness(best(pop)) \geq \phi) \wedge (beatExpert(best(pop), e))$  then //winner found!
18             |  $nCoev \leftarrow 0$ ; // start new search
19             |  $HoF_A \leftarrow HoF_A \cup \{best(pop)\}$ 
20             |  $temp \leftarrow A$ ;  $A \leftarrow B$ ;  $B \leftarrow temp$ ; // interchange player' roles
21         | else
22             |  $nCoev \leftarrow nCoev + 1$ ; // continue search
23         | end if
24         |  $e \leftarrow RANDELEMENTFROM(HoC)$ ;
25         |  $i \leftarrow i+1$ ;
26     end while
27 end while

```

Regarding the evaluation of candidates for a specific player p (where $p \in \{\text{player 1, player 2}\}$), the fitness of a specific strategy is computed by pitting it against a selected subset of the (winning) strategies in the HoF of its opponent player (that we call the “selected opponent set”) and evaluating his quality in a direct confrontation against any of the expert bots in the HoC. Given a specific strategy s its fitness is computed as follows:

$$\text{fitness}(s) = \frac{\sum_{j=1}^k (r_j^s - n\text{Turn}_s(j))}{k} + \text{extras}_s(e), \quad (1)$$

where $k \in \mathbb{N}$ is the cardinality of the selected opponent set, $r_j^s \in \mathbb{R}$ returns ϕ points if strategy s beats strategy h_j belonging to the selected opponent set (i.e. victorious case), and 0 if h_j wins over strategy s ; $n\text{Turn}_{s,j} \in \mathbb{N}$ is the number of turns spent on the game (this value is 0 in case of defeat); and $\text{extras}_s(e) \in \mathbb{N}$ is a bonus that individual s receives if it defeats the “expert bot” e . This fitness definition was formulated based on our game experience, and it values the victory above any other result.

In the next paragraph we describe in detail the three variants of our HoFCC algorithm which will be tested in the experimental section.

3.4 HoFCC-diversity

In this proposal the HoF acts as a long-term memory mechanism, but the content of the HoF is updated by removing those members that provide less diversity. The value of *diversity* that an individual in the HoF provides is calculated by the genotypic distance as follows: we manage the memory of champions as a matrix in which each row represents a solution and each column a gene (i.e., an action in the strategy). Then, we compute the entropy value for a specific column j as follows:

$$H_j = - \sum_{i=1}^k (p_{ij} \log p_{ij}), \quad (2)$$

where p_{ij} is the probability of action i in column j , and k is the memory length. Finally the entropy of the whole set is defined by:

$$H = \sum_{j=1}^n H_j. \quad (3)$$

The higher the value of H the greater the diversity of the set. For determining the diversity’s contribution to a specific solution, we calculate the value of entropy with this solution inside the set, and the corresponding value with this solution outside the set, and finally, the difference of these two values represents the contribution of diversity.

The number of individuals to be deleted from the memory should be set by the programmer as a percentage

value (α) representing the portion of the HoF to be removed; in other words, the HoF (with cardinality $\#HoF$) is ordered according to the diversity value in a decreased order and the last $\lceil \frac{\#HoF \times \alpha}{\lambda} \rceil$ individuals in this ordered sequence are removed. The frequency of updating (λ) is also a parameter of this version (i.e., the HoF is updated every λ coevolutions).

The motivation of this proposal is to maintain certain diversity among the members of the HoF, and at the same time reduce (or maintain an acceptable value for) the size of the memory. With this in mind, we assume that the deleted individuals will not affect the quality of the solutions found. Here, the cardinality of the selected opponent set k in the evaluation phase—see Eq. (1)—is the cardinality of the opponent HoF after executing the updating of the memory (i.e., after removing the individuals).

3.5 HoFCC-quality

In this version, we follow a similar approach to that applied in HoFCC-diversity but now the HoF is ordered with respect to a measure of quality that is defined as the number of defeats that an individual obtained in the previous coevolutionary step; in other words, a simple counter variable associated with each member of the HoF stores the number of defeats that were computed for the corresponding member during the evaluation process of the opponent army in the previous coevolutionary turn.

Based on our game experience, we assume that this metric is representative of the strength of a solution, and the aim is to keep only the robust individuals in the champions’ memory by removing the weak strategies. As in the HoFCC-diversity, the parameters α and λ have to be set, and the cardinality of the selected opponent set k is exactly the same.

3.6 HoFCC-U

This variant of HoFCC follows the idea of optimizing the memory of the champions, but in this case we propose a multiobjective approach where each solution has a diversity value and also a quality value, as previously described, associated with it. Then, a percentage value (α) from the set of dominated solutions according to the multiobjective values is removed; if the set of dominated solutions is empty then HoF is ordered according to the measure of quality and the solutions with worst quality will be removed. As in the previous algorithms (HoFCC-quality and HoFCC-diversity) the frequency of updating the HoF is an important parameter that must be defined.

This proposal uses a fitness function different to that shown in Eq. (1) the definition of which was inspired by

CFS (Bosin and Belew 1997). The main idea is that a defeat against opponent X has more importance if there are other individuals that defeated X .

So, a penalization value N for each individual i (for $1 \leq i \leq k$) in the population is then calculated as follows:

$$N_i = 1 - \frac{1}{k} \sum_{j=1}^k \frac{v_{ij}}{V(j)}, \quad (4)$$

where $v_{ij} = 1$ if the i th individual of the population defeats the j th strategy (or champion) in the HoF (whose cardinality is k) and 0 otherwise; and

$$V(j) = \sum_{i=1}^n v_{ij}, \quad (5)$$

is the number of individuals in the population which defeat the j th opponent of the HoF. As a consequence, $N_i \approx 0$ if the i th candidate defeats all opponents of HoF and the i th solution itself is one of the few candidates to do so; $N_i = 1$ if doesn't defeats any opponent; and $0 < N_i < 1$ depending on how many times it wins and how common it is to beat certain opponents. The fitness of a candidate i is then computed as follows:

$$F_i = P_i - \omega N_i, \quad (6)$$

where P_i is the result obtained in the battles by Eq. (1), and $\omega \in \mathbb{N}$ is a coefficient that scales N_i in order to make it meaningful with respect to the value P .

4 Experiments and analysis

This section describes the experimental analysis conducted on the three variants of the HoFCC described in the previous section. All experiments were executed using the Planet Wars engine version 1.2. Five maps from the collection of maps designed for the competition (specifically, maps 1, 10, 20, 30, 50) were considered, so that when an individual faces another one (i.e., its opponent) it must perform three consecutive battles on any of these maps having a maximum of 500 turns for each game. If the individual under evaluation defeats the opponent in all battles then it will be considered the winner.

The initial enemy strategy was defined as 'random' (i.e., we create a random action matrix), and the HoC (i.e., the selection of the "expert bots") comprised the set of bots that were originally provided as "example bots" in the competition (i.e., BullyBot, RandomBot, DualBot, ProspectorBot, RageBot), and specialized hand-coded (and later optimized via evolutionary algorithms) bot (GeneBot) that obtained a ranking position in the top-20 % of the Google AI competition 2010 and that was developed by a team from the University of Granada; we selected this bot

Table 2 Parameters of the coevolutionary cycle

Parameters	Values
<i>maxFailCoevolutions</i>	5
<i>maxCoevolutions</i>	15
<i>maxGen</i>	20
<i>popSize</i>	15
<i>maxEvaluations</i>	10,000
<i>crossProbability</i>	$p_X = 0.75$
<i>mutProbability</i>	$p_M = 1/nb$
ϕ	1,500
<i>expertBonus</i>	500
ω	50
λ	3

because it competed efficiently, the source code was available, and it has been extensively described in the scientific literature (Fernández-Ares et al. 14; Mora et al. 2012).

Next, we detail the configuration of this experiments, and later will discuss the obtained results.

4.1 Configuration of the experiments

Nine instances of our algorithms were used: three for each of the HoFCC-diversity, HoFCC-quality, and HoFCC-U varying according to the values of $\alpha \in \{10, 30, 50\ \%\}$. The notation *HoFX- α* —where $X \in \{Div, Qua, U\}$ —is used to denote each of these nine variants. In all cases, we set $\lambda = 3$ and perform 10 runs per algorithm instance, using a steady-state GA—note that this corresponds to lines 11–16 in Algorithm 1—with the aim of finding a winning strategy with respect to the set of strategies (all strategies stored in the HoF of the opponent) that were considered winning in previous stages of the coevolutionary algorithm; this GA employed binary tournament for selection, uniform crossover, bit-flip mutation, elitist replacement. Table 2 shows the parametrization values, where *maxFailCoevolutions* represents the limit of continuous coevolutions that one of the players can consume without finding a champion solution; the timeout condition (line 4 in Algorithm 1) is associated with *maxCoevolutions*, which indicates the maximum numbers of total coevolutions that can achieve an algorithm; and *maxEvaluations* is the limit of evaluations (i.e., $timeout = nCoev > maxCoevolutions \vee numEvaluations > maxEvaluations$, where *numEvaluations* represents the number of evaluations consumed at that instant by the algorithm). Mutation probability depends on *nb* which is the number of genes of the individual; ϕ is the threshold value for a winning solution.

Our analysis has been guided by the following indicators which are applied for all runs of each algorithm: *Best*

fitness shows the fitness of the best champion strategy found by the search process; *Average fitness* shows the average fitness value reached during the coevolutionary cycles; *Number of evaluations* indicates the total number of battles which are executed during the evaluation process; *Number of victories* indicates the total number of victories obtained in a *All versus All* fighting among the best solutions found by each version of algorithm.

Next we analyze the results obtained in ten independent executions for the nine versions of HoFCC, and focus on the indicators mentioned; we have used a non-parametric statistical test to compare ranks namely Kruskal–Wallis test (1952) with a significance of 95 %. When this test detects significant differences in the distributions, we have performed multiple tests using the Dunn–Sidak method (Sokal and Rohlf 1995) in order to determine which pairs of means are significantly different, and which are not.

4.2 Analysis of the results

Figure 2 shows the results of the fitness of the best champion strategy found in each independent execution of the algorithm instances. Bear in mind that in our experiments we use three versions of the HoFCC which optimize the use of HoF (in terms of diversity, quality, or both). The Kruskal–Wallis test shows that there are no significant differences between values (see the first row in Table 3). Note however the existence of three outliers (with fitness value 0) in the instances HoFU50, HoFU30 and HoFDiv30; this indicates that in their associated executions no winning




Fig. 2 Distribution of best fitnesses achieved by each algorithm. As usual, each *box* comprises the second and third quartiles of the distribution, the median is marked with a *vertical line*, and outliers are indicated with a *plus sign*

Table 3 Results of Kruskal–Wallis test for all the indicators

Indicators	<i>p</i> value
Best fitness	0.4605
Average fitness	1.356e–007
Number of evaluations	0.0266
Numbers of victories	0.0093




Fig. 3 Distribution of average fitnesses obtained in each algorithm

strategy was found after completing five consecutive coevolutions.

Figure 3 shows the behavior of the average fitness for each algorithm instance. The Kruskal–Wallis test confirms that the differences between values are statistically significant (see second row in Table 3). According to the graphic, the algorithms working on ‘diversity’ obtain the best results. Note also that HoFDiv50 again demonstrates the best performance and that the hybrid versions of the algorithm are not competitive. This last assertion might be made because it is more difficult to obtain a high fitness value due to the penalty that is applied to the score obtained by the individuals according to Eq. (4). A direct consequence is a decrease in the values of the objective function, which on the other hand should not affect the quality of the individual (this will be verified later in another test). In the results of multiple tests for the value of average fitness, the extreme values of the distribution are those that mark the difference: HoFDiv50 has significant differences respect to HoFDiv30, HoFU10, HoFU50, HoFU30; and HoFU30 is significantly different from HoFQua10, HoFQua50, HoFDiv10, and HoFDiv50.

In the case of the number of evaluations, there are significant differences (see third row in Table 3) between HoFDiv30 and HoFQua30. We noted however that the number of evaluations is directly proportional to the length




Fig. 4 Distribution of the number of evaluations used by each algorithm

of the coevolutionary cycle; remember that for these experiments we set a maximum quota of five consecutive coevolutions without success, and when this limit is reached the algorithm stops. We noted that the lengths of the coevolutionary cycles were not too long, which may be a direct consequence of the specialization of the solutions obtained during the search what makes more difficult the attainment of better solutions (Fig. 4).

Ten champions (the best individuals stored in the HoF at the end of the algorithm execution) for each algorithm instance (i.e., one for each execution) were obtained during the experiments and these champions fought in an *All versus All* tournament (note that in the cases of the outliers with fitness values 0 that were commented previously—regarding best fitness—the undefeated opponent of the HoF was taken as the “winning strategy” in this tournament). This means that every (best) individual competes against the ten champions of the opponent algorithm, and in each confrontation three battles (each of them in a distinct scenario) are executed. Figure 5 shows the results of this tournament. Note how each family of algorithms is clearly distinguished from the others by its results (this could not be perceived in our previous work on the RobotWars game Nogueira et al. 2013). According to the Kruskal–Wallis test there are significant differences between values (see fourth row in Table 3). The main differences are between HoFDiv and HoFU. Surprisingly all the instances of HoFU version obtain the best results, what means that the algorithms based on the multiobjective approach showed a more efficient behavior, whilst those instances based on the diversity metric exhibited a poor performance in direct combats. Certainly, this seems to contradict the results shown in Figs. 2 and 3 (see discussion on this in the next subsection) although in fact it




Fig. 5 Distribution of the number of victories obtained by each algorithm in the *All versus All* fighting tournament

is an expected result that confirms that the strongest and more consistent virtual players are those generated from the combined use of the two metrics considered in this paper. We also noted that the quality metric (as expected as well) marks a clear distinction with the diversity measure, although the combination of these is the most productive.

4.3 Summary of results

We observe that in the *All versus All* test, which attempts to measure the strength of the ‘champions’, the results of the algorithms are not related with the other indicators which focus on the analysis of the fitness values. This may be a sign that the coevolutionary process is affected by cycling and a solution with a high fitness might be theoretically considered as stronger although in practice it might not be; in fact the fitness score is not directly related to the concept of strength (or efficiency in combat). As already mentioned, the strongest virtual players (according to the test *All versus All*) were obtained by the family of algorithms that use a multi-objective approach based on diversity and quality values as metric to guide the search whilst the algorithms based on just one metric (specially those considering the diversity values) did not exhibit good results. The lowest values associated with the bots generated from the multi-objective instances might be caused by the penalty that is applied to the score in the fitness function according to Eq. (4). To relate all the algorithm instances, in the future we might think to include some of the champions obtained by each of the algorithms in the HoC of the rest of algorithms so that more pressure will be given to the attainment of champions and the fitness value of the individual might be more consistent with their performance.

The results obtained in this paper provide conclusive evidence in favor of the hypotheses set forth in previous work on the RobotWars game (Nogueira et al. 2013). Related to this, note that we include more battles against the members of the rival HoF in the candidate evaluation, and second (and most importantly) we believe that the application of the concept of CFS (Rosin and Belew 1997) to the fitness function, and the addition of pressure to the search via the inclusion of an expert archive caused a reduction in the appearance of cycling. In this regard, in our previous proposals we assumed that there would be transitivity between solutions and at the last solution added to the archive had to be the strongest, so the solutions were evaluated only by the score that they obtained against the members of the opponent HoF which favored the appearance of cycling and disengaged the search targets. The idea of including in the fitness functions combats against a group of experts (with which they have no evolutionary relationship) helped to break relationships between the champions found in each generation, and this surely eases the cycle breaking.

Regarding the percentage of updating the HoF (i.e., 10, 30, or 50 %), no significant difference was observed because their behaviors were indistinguishable. In this case, we believe that it is due to the fact that the HoF does not reach a significant size because in the experiments we set a limit of total coevolutions, so the HoF cannot grow significantly and it does not allow the capabilities of the updating process to be exploited.

5 Conclusions

Finding algorithms that can reduce the effect of the inherent pathologies of the coevolutionary models is an open research line. In previous work we explored the use of an archive method, in particular the use of a HoF, to keep the champions obtained in each coevolutionary step and that was used in the evaluation process to guide the search. An empirical study conducted on a RTS game—with intrinsic underlying limitations due to its simplicity—showed that the well-known problem of cycling continued appearing due to the assumption of a transitivity among the champions. Now, in this paper, we have extended that work significantly by introducing changes in the coevolutionary process, trying to promote those solutions that are truly strong. The proposed approach maintains the use of the HoF but also incorporates an additional memory (i.e., another archive termed HoC) to contain other efficient bots (i.e., ‘experienced’ or optimized virtual players) that are used here to assess the strength of the solution candidates. This concept of HoC allows pressure to be put on the algorithm search as the expert bots can (and should) be

implemented independently in other contexts. This means that these optimized virtual players do not necessarily maintain an evolutionary relationship with the evolutionary population. Moreover, in the experimental analysis we have considered a new RTS game that enables a deeper experimentation and thus the attainment of more consistent results to draw general conclusions. In addition, we have considered two quality metrics and, based on these, we have suggested a number of variants (i.e., families) of a primary CC algorithm that differ in the mechanism of updating the HoF; this updating is executed with the aim of removing those champions not contributing to the optimization. A new fitness function, inspired by the CFS (Rosin and Belew 1997), was also used in a multi-objective version of the primary CC algorithm.

The results obtained in the experiments clearly distinguish the performance of each algorithm family in the search for strong solutions. In this regard, our multi-objective CC version shows a consistent performance. This multi-objective variant is guided by the quality of the solution candidate (with respect to both HoF and HoC) and the diversity of the HoF. Taking into account these two metrics independently, the algorithmic versions led by the quality measure provides better (in the sense of “more competitive”) bots than those variants guided by the diversity metric.

An experimental analysis has also shown that our HoF/HoC-based proposal helps to reduce the occurrence of cycling in the coevolutionary process. Moreover, this research has also highlighted the self-adjustment capabilities of the CC algorithms described here by generating winning strategies with respect to both the (co)evolved enemies and other optimized enemies that form part of an experienced line-up. This opens up multiple application in the arena of videogames such as their use on adaptive games (Szita et al. 2009) among others.

Future work will involve analyzing the performance of new coevolutionary models and the application of those described here on other RTS games, incorporating new metrics for the HoF updating process, trying to design new evaluation mechanisms with the aim of reducing the effects of coevolutionary pathologies, and exploiting the potential of archive methods when they are combined with other approaches which help to optimize the performance of the solutions.

Acknowledgments This work is partially supported by Spanish MICINN under Project ANYSELF (TIN2011-28627-C04-01),³ by Junta de Andalucía under Project P10-TIC-6083 (DNEMESIS)⁴ and by Universidad de Málaga, Campus de Excelencia Internacional Andalucía Tech.

³ <http://anyself.wordpress.com/>.

⁴ <http://dnemesis.lcc.uma.es/wordpress/>.

References

- Angeline PJ, Pollack JB (1993) Competitive environments evolve better solutions for complex tasks. In: Forrest S (ed) 5th International conference on genetic algorithms (ICGA93). Morgan Kaufmann, San Mateo, pp 264–270
- Ashlock D, Ashlock W, Samothrakis S, Lucas SM, Lee C (2012) From competition to cooperation: co-evolution in a rewards continuum. In: IEEE conference on computational Intelligence and games (CIG 2012). IEEE, Granada, pp 33–40
- Avery P, Louis SJ (2010) Coevolving team tactics for a real-time strategy game. In: IEEE congress on evolutionary computation (CEC 2010). IEEE, Barcelona, pp 1–8
- Avery P et al (2008) Coevolving a computer player for resource allocation games: using the game of Tempo as a test space. PhD Thesis, School of Computer Science, University of Adelaide
- Avery PM, Greenwood GW, Michalewicz Z (2008) Coevolving strategic intelligence. In: IEEE congress on evolutionary computation (CEC 2008). IEEE, Hong Kong, pp 3523–3530
- Avery PM, Michalewicz Z (2007) Static experts and dynamic enemies in coevolutionary games. In: IEEE congress on evolutionary computation (CEC 2007). IEEE, pp 4035–4042
- Buro M (2004) Call for AI research in RTS games. In: Fu D, Orkin J (eds) AAAI workshop on challenges in game AI, San Jose, pp 139–141
- Cook M, Colton S, Gow J (2012) Initial results from co-operative co-evolution for automated platformer design. In: Chio CD, Agapitos A, Cagnoni S, Cotta C, de Vega FF, Caro GAD, Drechsler R, Ekárt A, Esparcia-Alcázar AI, Farooq M, Langdon WB, Guervós JJM, Preuss M, Richter H, Silva S, Simões A, Squillero G, Tarantino E, Tettamanzi A, Togelius J, Urquhart N, Uyar S, Yannakakis GN (eds) Applications of evolutionary computation—EvoApplications (EvoGAMES). Lecture notes in computer science, vol 7248. Springer, Berlin, pp 194–203
- Corruble V, Madeira CAG, Ramalho G (2002) Steps toward building of a good AI for complex Wargame-Type simulation games. In: Mehdi QH, Gough NE (eds) 3rd International conference on intelligent games and simulation (GAME-ON 2002), London, pp 155–159
- de Jong E (2004) Towards a bounded Pareto-Coevolution archive. In: Congress on evolutionary computation (CEC2004), vol 2. IEEE, Portland, pp 2341–2348
- Dziuk A, Mikkulainen R (2011) Creating intelligent agents through shaping of coevolution. In: IEEE congress on evolutionary computation (CEC 2011), IEEE, New Orleans, pp 1077–1083
- Ebner M, Watson RA, Alexander J (2010) Coevolutionary dynamics of interacting species. In: Chio CD, Cagnoni S, Cotta C, Ebner M, Ekárt A, Esparcia-Alcázar A, Goh CK, Guervós JJM, Neri F, Preuss M, Togelius J, Yannakakis GN (eds) Applications of evolutionary computation, EvoApplications 2010 (EvoApplications (1)). Lecture notes in computer science, vol 6024. Springer, Berlin, pp 1–10
- Fernández-Ares A, Mora AM, Guervós JJM, García-Sánchez P, Fernandes C (2011) Optimizing player behavior in a real-time strategy game using evolutionary algorithms. In: IEEE congress on evolutionary computation (CEC 2011), IEEE, New Orleans, pp 2017–2024
- Ficici SG, Bucci A (2007) Advanced tutorial on coevolution. In: 2007 GECCO conference companion on genetic and evolutionary computation. ACM, New York, pp 3172–3204
- Gutiérrez JAG, Cotta C, Fernández Leiva AJ (2011) Design of emergent and adaptive virtual players in a war RTS game. In: Ferrández JM, Sánchez JRÁ, de la Paz F, Toledo FJ (eds) 4th International work-conference on the interplay between natural and artificial computation (IWINAC 2011). Lecture notes in computer science, vol 6686. Springer, La Palma, pp 372–382
- Jaskowski W, Krawiec K (2010) Coordinate system archive for coevolution. In: IEEE congress on evolutionary computation (CEC 2010), IEEE, Barcelona, pp 1–10
- Johnson RW, Melich ME, Michalewicz Z, Schmidt M (2004) Coevolutionary Tempo game. In: Congress on evolutionary computation (CEC 2004), vol 2, IEEE, Portland, pp 1610–1617
- Kruskal W, Wallis W (1952) Use of ranks in one-criterion variance analysis. *J Am Stat Assoc* 47(260):583–621
- Lichocki P (2008) Evolving players for a real-time strategy game using gene expression programming. Master Thesis, Poznan University of Technology
- Lidén L (2004) Artificial stupidity: the art of intentional mistakes. In: AI game programming wisdom 2. Charles River Media, Inc., Hingham, pp 41–48
- Livingstone D (2005) Coevolution in hierarchical ai for strategy games. In: IEEE conference on computational intelligence and games (CIG 2005). IEEE, Colchester, Essex
- Lucas SM (2009) Computational intelligence and AI in games: a new IEEE transactions. *IEEE Trans Comput Intell AI Games* 1(1):1–3
- Lucas SM, Mateas M, Preuss M, Spronck P, Togelius J (2012) Artificial and computational intelligence in games (Dagstuhl seminar 12191). *Dagstuhl Rep* 2(5):43–70
- Miconi T (2009) Why coevolution doesn't "Work": superiority and progress in coevolution. In: Vanneschi L, Gustafson S, Moraglio A, Falco ID, Ebner M (eds) Genetic programming, 12th European conference (EuroGP 2009). Lecture notes in computer science, vol 5481. Springer, Berlin, pp 49–60
- Miles C, Louis SJ (2006) Co-evolving real-time strategy game playing influence map trees with genetic algorithms. In: International congress on evolutionary computation (CEC 2006). IEEE, Vancouver
- Mora AM, Fernández-Ares A, Guervós JJM, García-Sánchez P, Fernandes CM (2012) Effect of noisy fitness in real-time strategy games player behaviour optimisation using evolutionary algorithms. *J Comput Sci Technol* 27(5):1007–1023
- Nerome M, Yamada K, Endo S, Miyagi H (1998) Competitive co-evolution based game-strategy acquisition with the packaging. In: Jain LC, Jain RK (eds) Knowledge-based intelligent electronic systems, 2nd international conference (KES (3)). IEEE, Adelaide, pp 184–189
- Nogueira M, Gálvez J, Cotta C, Fernández-Leiva A (2012) Hall of Fame based competitive coevolutionary algorithms for optimizing opponent strategies in a new RTS game. In: Fernández-Leiva A et al A (ed) 13th Annual European conference on simulation and AI in computer games (GAME-ON 2012). Eurosis, Málaga, pp 71–78
- Nogueira M, Gálvez J, Cotta C, Fernández-Leiva A (2013) An analysis of hall-of-fame strategies in competitive coevolutionary algorithms for self-learning in RTS games. In: International conference on learning and intelligent optimization (LION7), LNCS. Springer, Catania
- Reynolds C (1994) Competition, coevolution and the game of tag. In: Brooks R, Maes P (eds) Proceedings of artificial life IV. MIT Press, Cambridge, pp 59–69
- Rosin C, Belew R (1997) New methods for competitive coevolution. *Evol Comput* 5(1):1–29
- Rosin CD, Belew RK (1995) Methods for competitive co-evolution: finding opponents worth beating. In: Eshelman LJ (ed) 6th International conference on genetic algorithms (ICGA). Morgan Kaufmann, San Mateo, pp 373–381
- Samothrakis S, Lucas SM, Runarsson TP, Robles D (2013) Coevolving game-playing agents: measuring performance and intransitivities. *IEEE Trans Evol Comput* 17(2):213–226

- Sims K (1994) Evolving 3D morphology and behavior by competition. *Artif Life* 1(4):353–372
- Smith G, Avery P, Housmanfar R, Louis SJ (2010) Using co-evolved RTS opponents to teach spatial tactics. In: Yannakakis GN, Togelius J (eds) *IEEE conference on computational intelligence and games (CIG 2010)*. IEEE, Copenhagen, pp 146–153
- Sokal Robert R, Rohlf James F (1995) *Biometry: the principles and practice of statistics in biological research*. W.H. Freeman and Company, New York
- Szita I, Ponsen MJV, Spronck P (2009) Effective and diverse adaptive game AI. *IEEE Trans Comput Intell AI Games* 1(1):16–27
- Togelius J, Yannakakis GN, Stanley KO, Browne C (2011) Search-based procedural content generation: a taxonomy and survey. *IEEE Trans Comput Intell AI Games* 3(3):172–186