

Design and Evaluation of an Extended Learning Classifier-based StarCraft Micro AI

Authors redacted for double blind review.

Abstract. Due to the manifold challenges that arise when developing an artificial intelligence that can compete with human players, the popular realtime-strategy game *Starcraft: Brood War* (BW) has received attention from the computational intelligence research community. It is an ideal testbed for methods for *self-adaption at runtime* designed to work in complex technical systems. In this work, we utilize the broadly-used Extended Classifier System (XCS) as a basis to develop different models of BW micro AIs: the Defender, the Attacker, the Explorer and the Strategist. We evaluate these AIs with a focus on their adaptive and co-evolutionary behaviors. To this end, we stage and analyze the outcomes of a tournament among the proposed AIs and we also test them against a non-adaptive player to provide a proper baseline for comparison and learning evolution. Of the proposed AIs, we found the Explorer to be the best performing design, but, also that the Strategist shows an interesting behavioral evolution.

1 Introduction

*Starcraft: Broodwar*¹ (BW, sometimes also referred to as only *Starcraft* or *Broodwar*) is one of the most famous instances of real-time strategy (RTS) games. It was released in 1998 for PCs and since then nearly 10 million copies have been sold. Founded on this number and on a huge number of players attracted to the game until today, it is seen as one of the most successful RTS game to date. RTS games can be characterized by three main tasks that the player has to fulfill: (i) collecting resources, (ii) creating buildings/units and (iii) controlling the units.

BW takes place in a science fiction setting, where three species compete for dominance in the galaxy. This are *Terrans*, a human-like species, *Protoss*, a species that is very advanced in technology and has psionic abilities, and *Zerg*, an insect swarm inspired species. The game has been extensively used for competitions, i.e., tournaments and leagues. These competitions usually consist of several 1-on-1 matches.

BW represents exactly the kind of training ground needed for testing and honing online learning methods and their capacity to function in complex real-world scenarios. BW challenges the learner through its great complexity, the arising dynamics, and the fact the fitness landscapes targeted by the learner are self-referential [8]. In BW, we face a set of entities (units and buildings) that interact with an environment (map and units of other players) in non-trivial ways. Furthermore, the environment is only partially observable and brings different types of uncertainty with it. Compared to other games that have been used as scientific testbeds, such as chess, go or poker, it creates a much bigger challenge. Another reason to chose BW as an application to test and hone online

¹ *Starcraft* and *Starcraft: Broodwar* are trademarks of Blizzard Entertainment

learning methods fit for real-world scenarios is the availability of an easy to use C++ library² that provides an interface to the game and therefore allows the development of artificial players as well as automated test runs of it.

Learning classifier systems, in particular variants of the extended learning classifier system (XCS), have been successfully deployed in various online learning tasks in real-world scenarios. In this work, we present an XCS-based model design for the artificial intelligence assuming the role of a player in *Starcraft: Brood War*. The remainder of this paper is structured as follows. In Section 2, we touch upon various related works in the context of RTS and corresponding machine learning approaches. We also introduce XCS as the learning system our approach is based on. In Section 3 we detail our model and the specific *Starcraft: Brood War* scenario it was developed for. Section 4 presents and discusses the results of our co-evolutionary learning experiments. Afterward, we conclude with a short summary and an outlook on potential future work.

2 Related Work

In this section, we first touch upon the numerous approaches to realising and deploying artificial intelligence techniques and machine learning in the Starcraft domain. Second, we present the Extended Classifier System (XCS) as the machine learning system used as the learning method for BW AIs in this work.

2.1 AI approaches in Starcraft

A recent survey covering bot architectures, i.e. the algorithmic architectures for automated players, is given in [10]. It identifies *learning and adaptation* as an open question in RTS game AI, which is addressed in this work. Numerous works in the field target prediction and handling uncertainty. In [16], for instance, a method is introduced to predict openings in RTS games. As another example, [19] presents an approach for estimating game states. In contrast, this work focuses on learning, but, the presented methods could be combined with the approach given here. Another direction of research is the exploration of methods for the engineering of bots. To this end, [15] proposes to follow the paradigm of agent-oriented programming, and [1] presents a method for automated testing of bots. Some works concentrate on providing data sets of BW games, e.g., [13] and [18]. There are also works about making and executing plans, such as [2] that proposes a method for the opening strategy optimization, or [4], where a method for the navigation of units is presented. Another category of works are the ones that innovate on mechanisms of strategy selection, e.g., [22], or of choosing tactical decisions [17]. Finally, there are several works on the control of units, e.g., [11], where a Bayesian network is utilized for the unit control, or [20], where Reinforcement Learning methods are applied to learn *kiting*, a hit and run technique for a special unit type. In this work, we propose the use of Learning Classifier Systems for providing an AI that both evolves new behaviours through evolutionary computation and hones and refines established ones through reinforcement learning. We provide four according AI designs which exhibit different focusses of the learning system's deployment.

² <https://github.com/bwapi/bwapi>

2.2 Extended Learning Classifier Systems

A *Learning Classifier System* (LCS) has originally been proposed in [5] by Holland. Later, he reworked the idea and proposes what today is considered a standard LCS in [6]. The most common extension of his work is the *Extended Classifier System* (XCS) of Wilson. It has been originally introduced in [21].

Since we adopted this variant for this work, the essence of the XCS is presented now. The basic architecture of an XCS is depicted in Figure 1. It represents a very elaborate learning system tailored towards real-world applications. Accordingly, in Figure 1, we see that the XCS gets a situation description of the environment through detectors. The situation is in the basic version of the XCS encoded as a bitstring. The population consists of classifiers, which hold several values:

- The **condition** is a string of 0, 1 and *don't cares* (often represented by X). The purpose of the condition is to determine, if the classifier matches the situation given by the detector. A match is given, if for every 0 in the situation there is a 0 or an X at same position in the condition.
- The **action** is also encoded as a bit string. The set of available actions is typically provided by the designer of the system and depends on the application.
- The **prediction** is a value that approximates the expected reward, given the action of this classifier is executed in the situations described by the condition. It is constantly adapted by taking new observations into account.
- The **prediction error** is a value that reflects how much the prediction deviated from the actual reward.
- The **fitness** expresses the accuracy of the prediction of the classifier.

The match set holds all classifiers that match the current situation. Most often, the match set holds classifiers that suggest different actions. The purpose of the prediction array is to decide which action is applied. To this end, it uses a fitness-weighted average of the predictions for each action that is mentioned in the match set. The classifiers proposing the chosen action are transferred to the action set and the action is applied through the effector.

In the next step, a reward is provided by the environment that values the current state. Based on the reward the prediction, error and fitness values of all the classifiers in the previous action set are adjusted according to a given update rule. In addition, a genetic algorithm is applied to the action set in order to create more appropriate rules. It selects parent classifiers for generating new ones based on their fitness values and can apply different crossover and mutation operators, which is up to the designer of the system.

3 Approach

In this section, we first explain the general scenario the AIs we have developed had to train and proof themselves. Based on this knowledge it is easier to follow the motivation for their individual designs which follows next.

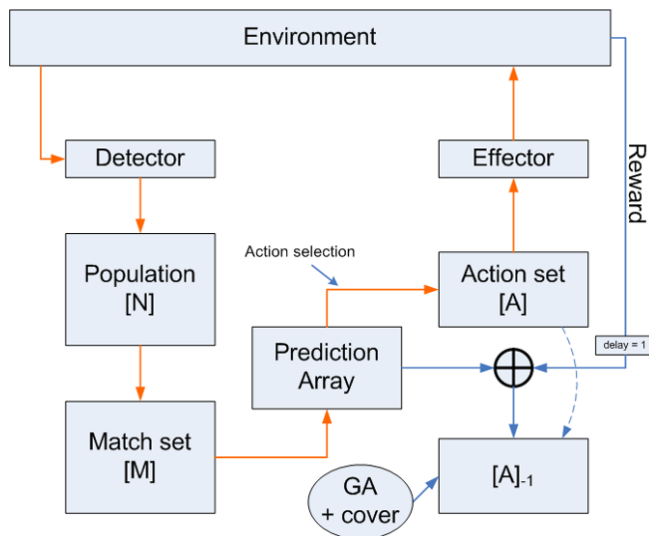


Fig. 1: The basic architecture of an extended learning classifier system or XCS.

3.1 Competition Scenario

We let our AIs compete and train in so-called *micro* matches, which implies that each AI was only able to control one group of units. In the given scenario, we did not consider the collection of resources and the production of units and buildings. A game is won, if all the enemy units or all the enemy buildings are destroyed. The match comes to a draw, if neither of the two competing AIs wins within a period of five minutes of simulated time. Each player starts out with the following heterogeneous set of predefined units.

Zergling Each player has control over 64 Zerglings at the start of the match. They are light units dealing little damage and they can only suffer little damage before they are destroyed. Zerglings are melee units, which means they can only attack, if they are close to enemy units.

Hydralisk Hydralisks can attack over distance but are not robust units, i.e., they should try to keep distance to the enemy units since they can be destroyed fast if they are attacked. Each player has control over 12 Hydralisks at the beginning of a match.

Ultralisk Each player only has two Ultralisks at their disposition. They are heavy units that are very robust, i.e. they can sustain a high number of hitpoints. Like Zerglings, Ultralisks are melee-only units.

Scourge Scourges are airborne units. Primarily, they attack other airborne units. At the loss of the Scourge unit itself, it can crash into other units to explode and damage the enemy. The player is provided with four Scourges at the beginning of a match.

Zerg Queen The Queen has no direct means of attack. Yet, it can slow down other units in a small quadratic area for 25 to 40 seconds, depending on the game speed. The enemies' movement velocity is halved, their rate of attack is reduced by between 10 to 33%, depending on the affected unit type. In addition, the queen can

hurl parasites at enemy units at a larger distances. The infested units' views extend directly add to the reconnaissance of the Queen's player.

The tournament map as well as the initial spatial arrangement of the given units is shown in Figure 2. It has been established by the SCMAI³ tournament. In Figure 2a, we see the used map. The starting points of the players are marked with (1). At the positions marked with (2), there are buildings that can attack units that are within their range. Two of these buildings belong to each player. If a player destroys one of the opponent's buildings, additional Zerglings appear in the center of the map as a reinforcement.




Fig. 2: (a) The map used for tournaments in our co-evolutionary experiment setup. (b) The initial spatial arrangement of the units made available to the AIs.

3.2 AI Components

In analogy to the components of an extended learning classifier system (Section 2.2), we considered the following basic building blocks for creating an effective Starcraft AI. (1) Behavioral rules to classify and react to a given situation, (2) a reinforcement component to adjust the rules' attributes in order to increase the achieved reward, (3) a covering mechanism to generate rules and fit newly encountered situations, and (4) a genetic algorithm for evolving the existing set of behavioral rules. In addition, we considered the ability to progress in battle formation based on individual *boid* steering urges [12], see Figure 3. Different from dynamically chosen but otherwise fixed formations, inferring the individual accelerations based on the units' neighborhoods results in emergent, adaptive formations [7]. In particular, the units sense their neighbors and (a) align their heading and speed with them, (b) tend towards their geometrical center, and (c) separate, if individual units get too close. As a result, battle formations such as the row formation in Figure 3(d) emerge.

3.3 XCS-based AIs

We combined the AI components outlined above in four different ways to trigger interesting competition scenarios and to trace and analyze the components' effectiveness. In

³ Starcraft Micro AI Tournament




Fig. 3: The augmented screenshots (a) to (c) depict the steering urges as defined by Reynolds’ flocking algorithm [12]. Based on these urges, formations emerge such as the one in (d).

particular, we implemented and co-evolved one defensive AI, one aggressive one, one that focuses on exploration and one where XCS takes global strategic decisions. Screen shots of their respective activities are seen in Figure 4.

The Defender Right at the beginning of the match, all units move to the upper of two buildings on the map and stay there for its defense until the end. Using the full functionality of an XCS, the Hydralisks’ as well as the Queen’s optimal behaviors are learned. For the Hydralisks, the condition part of the classifier rules considers the distance to the next visible enemy. Six actions are offered: (1) Approach and attack the closest ground or (2) airborne enemy, (3) move to a predefined point, (4) support a friendly unit, (5) protect the hatchery, or (6) burrow. For the Queen the proximity to the next enemy unit can trigger escape or *ensnare* airborne units or to hurl *parasites* at ground units. The XCS’ reinforcement component positively rewards any attacks, whereas the other actions are only remunerated, if the player is attacked itself or if the units have built up a great distance to their buildings.

The Attacker This is an offense-oriented AI. It will attack the next visible enemy. If none are in sight, the next enemy building is attacked. The units move in flocks, often they break into two clusters to attack both the enemy’s buildings simultaneously. In perilous situations (considering the numbers of enemies), however, an

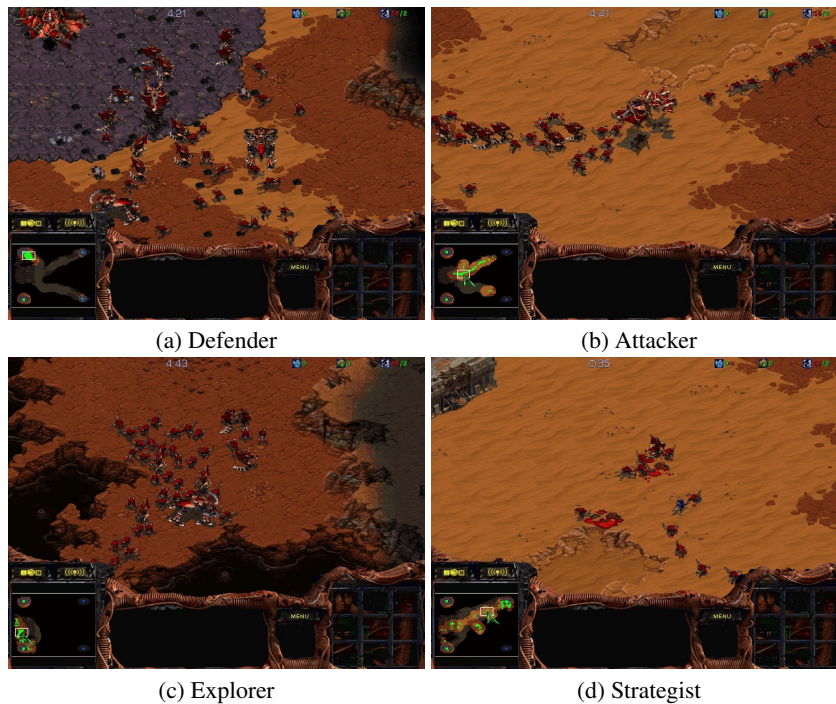


Fig. 4: Screen shots of representative behaviors of the four implemented AIs. The Defender assembles his troops to defend the buildings. The Attacker storms towards the enemy’s buildings to attack. The units of the Explorer swarm in different directions from the base. The Strategist decided on attacking enemy units.

XCS may reinforce the attack or instigate retreat. Any kind of attack (the XCS also decides the Queen’s mode of attack) is rewarded, suffering damage results in negative reinforcement.

The Explorer The units are divided into two clusters that head into randomly chosen directions to explore the environment. When an enemy is sighted, an XCS determines to attack or to escape. Successful attacks directly translate into positive rewards, whereas loss of health points implies negative reinforcement. First strike is additionally greatly rewarded, whereas suffering a surprise attack results in an equally great loss—adding or subtracting 100 reward points, respectively. Similarly, winning and losing a match results in adding/deducting the comparatively small reinforcement value of 10.

The Strategist Here, XCS is used to determine the overall strategy of the player. Based on the remaining time, the available and the opposing units, XCS determines whether to (1) attack enemy units or (2) buildings, whether to (3) defend one’s buildings or (4) to idle. The respective strategies imply according convoy movements, if necessary. Independently of the strategy, the next sighted enemy is always attacked. The

exhibited behavior is rewarded with the number of remaining units and buildings at the end of each match.

3.4 Learning Scenario

For the evaluation, we set up a learning scenario that addresses the issues of on-line learning, adaption and co-evolution. To allow this, we let the AIs compete with and learn from each other in several matches in a row. In particular, we first let the AIs train for 100 matches in a row with each of the other three AIs. In a second round, the previous adaptation is put to the test in the course of another 50 matches against each enemy AI. As all the AIs are designed to improve themselves by means of the combined reinforcement and evolutionary learning components of XCS, the tournament allows the AIs to co-evolve. Furthermore, for a better comparability of the approaches, we conducted additional experiments that include a non-learning AI. Competing against a non-learning AI ensures that any observed improvements do not emerge from co-evolutionary dynamics but are the result of the individual learners themselves. In addition, the non-learning AI also provides a clear baseline against which all the other AIs can be measured against. The non-learning AI mainly defends its position by splitting the given units in two groups which defend the two buildings. It has no intention of winning the game by moving on to attack the enemy.

4 Evaluation & Discussion

Considering 600 matches, the Explorer with 269 won matches clearly represents the best designed AI. The Attacker is second best with 105 wins, followed by the Defender (30 wins) and the Strategist (7 wins). For the complete stats, please refer to Table 1. These performances are both the product of the AIs basic strategies but also of the sequence of their co-evolutionary learning experiences. Therefore, it is important to analyze the relative progress each of the AIs has achieved. An according, high-level summary is depicted in Table 2. It shows the consecutive changes regarding the number of frames a simulation runs, the number of health points left of a player and the number of experienced winning situations. A decrease in frames signals an increase in clarity regarding the winner, as draws become less likely and as quicker solutions take over. An increase of left-over health points of an AI may be considered an improvement. However, an AI may also learn to sacrifice more health points in order to win a match in the end. Therefore, only an increase in wins statistically indicates an improved, i.e. learned, behavior.

AI	wins	draws	defeats	AI	wins	draws	defeats
<i>Defender</i>	30	547	33	<i>Strategist</i>	7	364	229
<i>Explorer</i>	269	323	8	<i>Attacker</i>	105	354	141

Table 1: The total numbers of wins, draws and defeats of the co-evolutionary runs of the AIs.

Defender				Strategist					
enemy	matches	frames	hp left	winner	enemy	matches	frames	hp left	winner
Explorer	100	+	+	0	Attacker	100	-	+	-
Attacker	100	-	-	+	Explorer	100	+	+	0
Strategist	100	+	-	-	Defender	100	+	+	0
Strategist	50	+	+	-	Explorer	50	-	0	-
Explorer	50	+	+	+	Attacker	50	-	+	0
Attacker	50	-	0	+	Defender	50	+	0	0

Explorer				Attacker					
enemy	matches	frames	hp left	winner	enemy	matches	frames	hp left	winner
Defender	100	+	+	-	Strategist	100	-	+	+
Strategist	100	+	-	-	Defender	100	-	-	+
Attacker	100	-	0	+	Explorer	100	-	0	-
Attacker	50	-	0	-	Explorer	50	+	+	-
Strategist	50	-	0	+	Strategist	50	-	+	+
Defender	50	+	+	+	Defender	50	-	-	0

Table 2: Each of the four AIs trains with and competes against all the other ones. This table depicts the consecutive changes in the number of frames the simulation ran for, the health points successfully maintained by the AIs, and the frequency of winning situations.

However, as pointed out above, in an attempt to objectively compare the learning successes of each AI, we reset the knowledge bases of all four of them and let them train and compare against a simple, non-learning AI for another 200 matches. The results in terms of averaged fitness evolution as well as in terms of averaged prediction error can be seen in Figure 5. Although they do not seem to improve much, the Defender and the Attacker AIs have rather high average fitness values to begin with. Their averaged prediction error does not change over the course of the evolutionary experiment either. The average fitness of the Strategist AI, however, rises continuously and converges quickly, despite the fact that its average prediction error rises sporadically as well. The most likely explanation for this discrepancy is that the prediction errors rise so uniformly across the whole population of classifiers that a greater error value would not impact the selection and thereby the whole interaction process. In analogy to the achieved changes as shown in Table 2 for the co-evolutionary experiments, the experiments against the non-learning AI yielded the following results. The Defender achieved an increase in health points left and stagnated regarding frames and wins. The Explorer improved to maintain greater numbers of health points, to reduce the numbers of frames needed in each match, and also to win more matches. The Strategist also cut down the number of frames for each match. Its health points stagnated and its number of wins increased. The evolution of the Attacker yielded longer matches (increase in computed frames), fewer health points left and an increase in won matches. To conclude these observations, like in the co-evolutionary experiments, where the Explorer showed the best performance based on wins, draws and defeats, the Explorer also exhibited the best performance in the baseline experiments against the non-learning AI. The increase in maintained health points, the reduction of frames per match and the increased number of wins are all desirable achievements unmet by its competitors.




Fig. 5: The four XCS-based AIs’ training progress when learning to compete against a simple non-learning AI.

4.1 Co-Evolution Qualitatively

The Attacker AI reinforces its aggressive behavior, especially if it does not suffer from inflicted damages. As a consequence, it learns to ruthlessly exploit the Defender AI’s feeble assaults by being even more fierce. When exposed to the other AIs, the Attacker AI quickly adapts to be slightly less aggressive. Similarly, the Strategist AI reinforces behaviors that minimize damage. As a result, when facing the Defender AI, the Strategist AI is not motivated to learn well-directedly. Instead, any behavior leads to success. When facing more aggressive opponents such as the Attacker or the Explorer, the Strategist AI receives smaller rewards, almost independently of the ingenuity of a selected strategy. The Defender AI adapts its Hydralisks to shy away from enemies as they get destroyed too quickly, otherwise. Towards simple AIs, such as the simple non-learning AI mentioned above, the Defender AI increases aggressive behaviors. The Defender’s Queen is mostly on the run, too, as otherwise the distance to the enemy’s units becomes too small. The Explorer as the overall best designed AI presented in this work is discussed more in depth in the next section.

4.2 Decentralized XCS-Concept

After the description of the evaluation and results, we want to provide further details about the Explorer, the most successful AI in the tournament. It utilizes a decentralized

XCS-concept, i.e., it adopts multiple XCS, that act in common. It uses one for each unit type (which are presented in 3.1). Each XCS decides whether the units of the respective type will engage or retreat when faced with an enemy. For the XCS the following configuration is used. Regarding the genetic algorithm, we empirically found the following parameters to be effective. Crossover is applied with a probability of 1% and mutation with a probability of 1.5% for each bit in a classifier. The parents are chosen by means of tournament selection with a tournament size of 5. The reinforcement component is configured with a learning rate of $\beta = 0.2$, i.e., the adaption of the prediction value is rather careful, and a discount factor of $\gamma = 0.71$ is used, i.e., future rewards are valued rather high. The reward is the difference between the damage the units have dealt and damage they have took, as it has been proposed in [20]. Additionally, there are some rewards that are considered in special situation. There is an extra reward of 10 if the game has been won and negative reward of -10 if the game is a draw or lost. The action selection is ϵ -greedy with $\epsilon = 0.02$, i.e., the action is selected randomly with a probability of 2%, otherwise the best action is selected. Additionally, a battle formation procedure based on Reynolds' flocking algorithm is utilized, where the parameters in the algorithm are optimized by a genetic algorithm. Through the optimization, the player evolves a very tight formation, where the Queen is positioned in the center.

Based on this architecture, the Explorer exhibits the following behavior. If there are no enemy units in sight, it creates two separate swarms from the set of available units. These two groups go in an individual, randomly chosen direction in order to explore the map. Utilizing a formation for movement can lead to a tactical advantage, if the opponent's forces are met. The Strategist develops different behavior against the different opponents. If facing the Defender, it will attack immediately, since it appears that this enemy will not be harmful for the strategist. Against the other two AIs, the Strategist is more reluctant since these more offensive AIs tend to deal more damage than the defender.

5 Summary & Future Work

Concluding the work, a motivation for *Starcraft: Brood War* as an ideal testbed for *self-adaption at runtime* is given in Section 1, followed by an overview over the state-of-the-art in the scientific developments in the BW domain and a short description of the XCS in Section 2. Next, in Section 3, four XCS-based AIs – the Defender, the Attacker, the Explorer and the Strategist – are presented. Furthermore, in Section 4, the evaluation scenario in a tournament scenario and against a non-learning AI with a special focus on the co-evolutionary behavior is presented and discussed.

Overall, we see two main results: The first is that the Explorer shows the best performance of all proposed XCS-based AIs. The second one is that, even though some players performed much worse than the Explorer, the *self-adaption at runtime* worked out for each player. This can be concluded since every player - after a period of adaption - shied away from the more aggressive AIs and, in turn, became more offensive against the more passive AIs.

For future work, we are planning in two directions. First, we want to refine the XCS-based approach for micro-management in BW. In particular, we want to provide

all the units' degrees of freedom available to individual, decentralized XCS learners and also feed them with pre-processed data that indicates general trends in the evolution of the game by means of correlation factors (e.g.[14]) or ascertainment of structural emergence (e.g. [3]). Second, we want to realise an AI that considers a broader managerial scope including micro-management and strategic group activities. To this end, we deem a multi-layered AI architecture taking on different responsibilities through the consideration of different time-scales and levels of abstraction a first important step [8].

References

1. Blackadar, M., Denzinger, J.: Behavior learning-based testing of starcraft competition entries. In: Bulitko, V., Riedl, M.O. (eds.) *AIIDE*. The AAAI Press (2011), <http://dblp.uni-trier.de/db/conf/aiide/aiide2011.html#BlackadarD11>
2. Churchill, D., Buro, M.: Build order optimization in starcraft. In: Bulitko, V., Riedl, M.O. (eds.) *AIIDE*. The AAAI Press (2011), <http://dblp.uni-trier.de/db/conf/aiide/aiide2011.html#ChurchillB11>
3. Fisch, D., Jänicke, M., Sick, B., Müller-Schloer, C.: Quantitative emergence – a refined approach based on divergence measures. In: *Self-Adaptive and Self-Organizing Systems (SASO)*, 2010 4th IEEE International Conference on. pp. 94–103 (Sept 2010)
4. Hagelback, J.: Potential-field based navigation in starcraft. In: *Computational Intelligence and Games (CIG)*, 2012 IEEE Conference on. pp. 388–393 (Sept 2012)
5. Holland, J.H.: *Adaptation**. In: ROSEN, R., SNELL, F.M. (eds.) *Progress in Theoretical Biology*, pp. 263 – 293. Academic Press (1976), <http://www.sciencedirect.com/science/article/pii/B9780125431040500123>
6. Holland, J.H., Reitman, J.S.: Cognitive systems based on adaptive algorithms. *SIGART Bull.* (63), 49–49 (Jun 1977), <http://doi.acm.org/10.1145/1045343.1045373>
7. Lin, C.S., Ting, C.K.: Emergent tactical formation using genetic algorithm in real-time strategy games. In: *Proceedings of the 2011 International Conference on Technologies and Applications of Artificial Intelligence*. pp. 325–330. TAAI '11, IEEE Computer Society, Washington, DC, USA (2011), <http://dx.doi.org/10.1109/TAAI.2011.63>
8. Müller-Schloer, C., Schmeck, H.: *Organic Computing - Quo Vadis?* In: Müller-Schloer, C., Schmeck, H., Ungerer, T. (eds.) *Organic Computing - A Paradigm Shift for Complex Systems*, chap. 6.2, pp. 615 – 625. Birkhäuser Verlag (2011)
9. Müller-Schloer, C., Schmeck, H., Ungerer, T. (eds.): *Organic Computing - A Paradigm Shift for Complex Systems*. Birkhäuser (2011)
10. Ontañón, S., Synnaeve, G., Uriarte, A., Richoux, F., Churchill, D., Preuss, M.: A survey of real-time strategy game AI research and competition in starcraft. *IEEE Trans. Comput. Intellig. and AI in Games* 5(4), 293–311 (2013), <http://dx.doi.org/10.1109/TCIAIG.2013.2286295>
11. Parra, R., Garrido, L.: Bayesian networks for micromanagement decision imitation in the rts game starcraft. In: Batyrshin, I., Mendoza, M. (eds.) *Advances in Computational Intelligence*, *Lecture Notes in Computer Science*, vol. 7630, pp. 433–443. Springer Berlin Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-37798-3_38
12. Reynolds, C.W.: Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics* 21(4), 25–34 (1987)
13. Robertson, G., Watson, I.: An improved dataset and extraction process for starcraft ai. In: *The Twenty-Seventh International Flairs Conference* (2014)

14. Rudolph, S., Tomforde, S., Sick, B., Hähner, J.: A Mutual Influence Detection Algorithm for Systems with Local Performance Measurement. In: Proceedings of the 9th IEEE International Conference on Self-adapting and Self-organising Systems (SASO15), held September 21st to September 25th in Boston, USA. pp. 144–150 (2015)
15. Shoham, Y.: Agent-oriented programming. *Artif. Intell.* 60(1), 51–92 (Mar 1993), [http://dx.doi.org/10.1016/0004-3702\(93\)90034-9](http://dx.doi.org/10.1016/0004-3702(93)90034-9)
16. Synnaeve, G., Bessière, P.: A bayesian model for opening prediction in rts games with application to starcraft. In: Cho, S.B., Lucas, S.M., Hingston, P. (eds.) CIG. pp. 281–288. IEEE (2011), <http://dblp.uni-trier.de/db/conf/cig/cig2011.html#SynnaeveB11a>
17. Synnaeve, G., Bessière, P.: A bayesian tactician. In: in "Proceedings of the Computer Games Workshop at the European Conference of Artificial Intelligence 2012. pp. 114–125 (2012)
18. Weber, B.G., Ontañón, S.: Using automated replay annotation for case-based planning in games. In: ICCBR Workshop on CBR for Computer Games (ICCB-Games) (2010)
19. Weber, B.G., Mateas, M., Jhala, A.: Applying goal-driven autonomy to starcraft. In: Youngblood, G.M., Bulitko, V. (eds.) AIIDE. The AAAI Press (2010), <http://dblp.uni-trier.de/db/conf/aiide/aiide2010.html#WeberMJ10>
20. Wender, S., Watson, I.D.: Applying reinforcement learning to small scale combat in the real-time strategy game starcraft: Broodwar. In: CIG. pp. 402–408. IEEE (2012), <http://dblp.uni-trier.de/db/conf/cig/cig2012.html#WenderW12>
21. Wilson, S.W.: Classifier fitness based on accuracy. *Evol. Comput.* 3(2), 149–175 (Jun 1995), <http://dx.doi.org/10.1162/evco.1995.3.2.149>
22. Yi, S.: Adaptive strategy decision mechanism for starcraft ai. In: Han, M.W., Lee, J. (eds.) EKC 2010, Springer Proceedings in Physics, vol. 138, pp. 47–57. Springer Berlin Heidelberg (2011), http://dx.doi.org/10.1007/978-3-642-17913-6_7