

Kent Academic Repository

Full text document (pdf)

Citation for published version

Rizzi, Caroline and Johnson, Colin G. and Vargas, Patricia (2018) Fear Learning for Flexible Decision Making in RoboCup: A Discussion. In: Akiyama, Hidehisa and Obst, Oliver and Sammut, Claude and Tonidandel, Flavio, eds. Lecture Notes in Computer Science. RoboCup 2017: Robot World Cup XXI. Lecture Notes in Artificial Intelligence, 11175. Springer, Germany pp. 59-70.

DOI

https://doi.org/10.1007/978-3-030-00308-1_5

Link to record in KAR

<https://kar.kent.ac.uk/69611/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

Fear Learning for Flexible Decision Making in RoboCup: A Discussion

Caroline Rizzi¹, Colin G. Johnson¹, and Patricia A. Vargas²

¹ School of Computing, University of Kent, Canterbury, UK,
cr519@kent.ac.uk*, C.G.Johnson@kent.ac.uk

² Robotics Laboratory, School of Mathematical and Computer Sciences,
Heriot-Watt University, Edinburgh, UK
p.a.vargas@hw.ac.uk

Abstract. In this paper, we address the stagnation of RoboCup competitions in the fields of contextual perception, real-time adaptation and flexible decision-making, mainly in regards to the Standard Platform League (SPL). We argue that our Situation-Aware FEAR Learning (SAFEL) model has the necessary tools to leverage the SPL competition in these fields of research, by allowing robot players to learn the behaviour profile of the opponent team at runtime. Later, players can use this knowledge to predict when an undesirable outcome is imminent, thus having the chance to act towards preventing it. We discuss specific scenarios where SAFEL’s associative learning could help to increase the positive outcomes of a team during a soccer match by means of contextual adaptation.

Keywords: RoboCup, Cognitive Learning, Contextual Fear Conditioning, Brain Emotional Model, Affective Computing

1 Introduction

RoboCup is an important international scientific initiative with the goal to advance the state of the art of artificial intelligence for autonomous robots by proposing an ambitious challenge. The official challenge of the RoboCup initiative, established in 1997, states that “by the middle of the 21st century, a team of fully autonomous humanoid robot soccer players shall win a soccer game, complying with the official rules of FIFA, against the winner of the most recent World Cup.” [4].

The relevance of the RoboCup competition is not on the challenge itself, but on the intrinsic gains from the journey to accomplish such a goal. RoboCup’s initiative poses a challenge of high complexity that requires a significant body of research in the areas of artificial intelligence, sensor fusion, real-time planning and navigation, cooperation in multiagent robotics, context recognition, image processing, motor control, among others [7]. The RoboCup initiative has hosted annual competitions for more than 20 years now, a period over which significant advancements have been achieved towards autonomous robotics.

Despite RoboCup’s many achievements in a number of research fields related to autonomous robotics, the development of contextual perception and flexible decision making has made modest progress. In this paper, we address RoboCup’s stagnation in these areas, discussing the existing approaches and their limitations (Section 2). We propose in Section 3, as potential approach to tackle these limitations, our artificial emotional model named SAFEL (Situation-Aware FEAR Learning) [13, 14, 17], which consists of a hybrid architecture based on the fear-learning mechanisms of the brain. Finally, in Section 4, we suggest RoboCup scenarios where SAFEL could be useful to improve flexible decision making at both individual and multi-agent levels, followed by the conclusions in Section 5.

2 Intelligent Behaviour in RoboCup

Due to the inherent teamwork nature of soccer, most RoboCup-related works addressing intelligent behaviour and decision making tend to investigate techniques to optimise collaborative behaviour and pre-coordination [6, 10, 19]. These approaches are commonly based on pre-determined coordination strategies, such as predefined behaviour rules and pre-trained machine learning and/or evolutionary algorithms, which provide the robot players with a basis to decide when it is better to kick, dribble, pass, change roles, etc.

Although useful for training a team of robots to play collaboratively in most common soccer situations, these approaches lead to limited coordination strategies that are immutable and restricted to tactics stipulated prior to the actual soccer match. As a consequence, the same strategy is delivered against all opponent teams in the competition. However, different teams may use different tactics, and a specific pre-trained approach may fail against a particular opponent while being successful against another opponent team.

For example, suppose that strikers of a particular team may be more aggressive and negligent than normal, consequently causing more collisions and fouls, while another team’s striker may be excessively cautious, which could slow it down. Therefore, the goalkeeper may need to take risky actions against the former team in order to avoid collisions and goals, which may not be necessary against the latter team.

If we consider the many aspects involved in a robot’s action (e.g., the kick strength, the collision prevention strategy, the walking speed, the range and timing of vision perception, etc.), the final behaviour of a robot can be completely different from team to team, even when considering the same in-game situation (e.g., striker shooting to goal). Taking advantage of the behavioural variations among teams can be critical to winning a soccer match.

For this reason, a mechanism that allows the individuals of a team to learn and adapt at runtime to the playing behaviour of each opponent team is essential. The need for such a mechanism is accentuated by the ultimate goal of the RoboCup competition: a soccer match between a team of robots and a team of humans.

Usually, for a team of robots, the difference in behaviour is only meaningful when comparing two distinct teams. In other words, two different teams may have different approaches to how they implement their players under a particular role (e.g., striker, defender, goalkeeper, etc). But it is uncommon for a team in the RoboCup competition to have two or more completely different implementations of the same role.

On the other hand, the behavioural difference goes further than just group level for a team of humans. This is because, not only different teams have different tactics, but also different individuals in a team may behave completely different under the same situation, even when assigned to the same role. Therefore, a team of human players entails even more complex behavioural differences if compared to a team of robots.

The need for real-time adaptation capabilities has been previously addressed using case-based reasoning [1, 18]. In these works, case-based reasoning approaches are used for post-coordination as a mean to optimise players' positioning during the match. These works represent a great contribution towards post-coordination, flexible decision making and real-time adaptation.

Nevertheless, to the best of our knowledge, these works usually fall in one of the following limitations: (1) temporal information is not considered in the problem solving, (2) the approach is domain specific or (3) predictions' applicability is limited to the optimisation of players positioning. In addition, most of these approaches have been tested and applied only to RoboCup leagues based on simpler robots or on simulations. RoboCup leagues based on more complex robots, such as the Standard Platform League (SPL), still lack more robust real-time adaptation mechanisms. In the next section, we present a situation-aware fear-learning computational model, which is a real-time adaptation mechanism capable of overcoming the above-mentioned limitations.

3 SAFEL

SAFEL stands for *Situation-Aware FEAR Learning*. It is a situation-aware computational system capable of providing robots with fear-learning skills in order to predict threatening situations to their own well-being or to their goals. SAFEL's model has been first proposed by us in [17], partially implemented and tested in [13] and improved in [14]. In this section, we briefly introduce SAFEL's biological inspiration and design. For a deeper understanding of SAFEL's model and further details on its implementation and performance analysis, we refer the reader to our previous publications [13, 14, 17].

SAFEL is a hybrid computational architecture inspired by the LeDoux's fear-learning model of the human brain [9, 8]. According to LeDoux, fear learning greatly relies on two brain regions known as the *amygdala* and the *hippocampus*, as well as on a cognitive function known as the *working memory*.

Considerable evidence indicates the amygdala as an essential brain region for fear learning and memory [9, 8]. It is responsible for processing the emotional significance of sensed stimuli by creating associations between neutral and aver-

sive stimuli. On the other hand, the hippocampus is believed to be the main brain region involved in context processing [8]. In the hippocampus, sensory information is put together in order to form a unitary representation of the current state of affairs. Unlike information processed in the amygdala, representations formed in the hippocampus are not just visual, auditory or olfactory, but all of these at once, and includes the way these sensations relate to each other both in intensity and temporal order. Finally, the working memory creates associations between the contextual memory formed in the hippocampus with the emotional memory formed in the amygdala, giving emotional meaning to the contextual information acquired in past experiences.

SAFEL’s architecture is based on the task division proposed by LeDoux. Therefore, analogous to the LeDoux model, SAFEL is divided into three modules that work in an integrated and parallel manner: the amygdala, the hippocampus and the working memory modules. Fig. 1 depicts the SAFEL model, illustrating how the three modules of the architecture are interconnected.

Environmental stimuli detected by the robot (e.g., by means of sensors’ input or direct user input) must first be normalized and categorised into aversive and neutral stimuli by the robot’s controller before being delivered to the amygdala and hippocampal modules. The amygdala module is responsible for detecting threats by analysing the current values of aversive stimuli and associating them to simultaneously occurring neutral stimuli. This learning process is induced by means of a procedure analogous to *classical fear conditioning* [11].

In classical fear conditioning, associative learning is induced by pairing a neutral stimulus (i.e., a stimulus that initially elicits no specific response from the individual) with an aversive stimulus (i.e., a stimulus that naturally elicits fear or discomfort, such as pain, hunger, etc.). Eventually, the previously neutral stimulus acquires emotional meaning and becomes able to elicit the state of fear by itself, even in the absence of the aversive stimulus. When this happens, we say that the neutral stimulus is now a *conditioned stimulus*, which elicits fear as a *conditioned emotional response*.

In SAFEL’s model, the amygdala module is also responsible for providing emotional feedback to the hippocampus module, which in parallel generates complex contextual representations of the sensed environmental stimuli. In the hippocampus, the amygdala emotional feedback and the generated contextual information are associated.

Finally, pieces of contextual information and their emotional significance are memorised in the working memory. Later, any previously experienced pattern of contextual information will trigger the retrieval of that stored memory and its emotional meaning. Consequently, if a particular situation preceded the occurrence of an aversive stimulus in a past experience, the working memory will retrieve the same state of fear triggered by that situation in the past, warning the individual that an undesirable situation is likely to happen in the near future.

SAFEL’s amygdala module is based on a modified *artificial neural network* (ANN) proposed by us in [16], which allows robots to associate environmental stimuli at runtime based on the Pavlovian classical conditioning procedure [11].

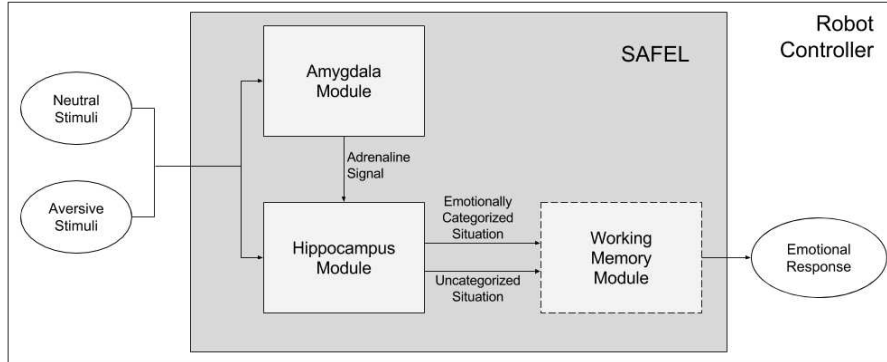


Fig. 1. The SAFEL model. Solid-border boxes represent areas of the brain whereas dotted-border boxes represent cognitive functions of the brain. The model receives neutral and aversive stimuli as input from the robot controller and outputs the corresponding emotional response back to the robot controller.

In the amygdala module, this modified ANN is used to associate neutral and aversive stimuli at runtime. The ANN is pre-trained to generate a high output value whenever any aversive input is also high, and a low output otherwise, regardless of the value of neutral inputs. Associative learning takes place by adjusting the first-layer weights of the ANN according to the coincidence of input values. In other words, the association takes place whenever a high neutral stimulus input and a high aversive stimulus input co-occur. Eventually, the neutral stimulus is turned into a conditioned stimulus, becoming able to trigger by itself the same ANN output that the aversive stimulus would, even in its absence. The output of the ANN is said to be the *adrenaline signal*, which is a real value in the range $[0,1]$ representing the current fear level of the system.

The hippocampus module is based on Dey’s [3] conceptualization of situation awareness for expert systems. It is responsible for collecting, understanding and managing the states of the robot over time. To accomplish that, we have modelled and implemented the hippocampus module using SCENE [12, 15], which is a powerful situation management platform that extends the JBoss Drools rule engine and its CEP (Complex Event Processing) platform [2].

The hippocampus module receives two inputs: *events*, which are sets of environmental stimuli at a given point in time, and the adrenaline signal relayed by the amygdala module. This module is responsible for assembling these events into pieces of information known as *situations*, which depict the robot’s state of affair during a particular period of time. Situations are stored in the hippocampus module as matrices $S_{m \times n}$, where m is the number of time steps encompassed in a particular situation’s duration and n is the number of stimuli being sensed by the robot.

Situations are later categorised in regards to their emotional meaning according to the subsequent emotional feedback from the amygdala. Situations are

categorized as *aversive situations* if preceding high adrenaline signals, and *safe situations* otherwise. Ongoing situations are left uncategorised because their true emotional meaning can only be determined sometime after their conclusion.

Finally, the working memory is the module of SAFEL where the association between context and “fear” takes place. In the working memory, the temporal patterns of situations are memorised and associated with their respective labels (safe or aversive). Here, two processes take place. First, a feature extraction is performed in order to generate compacted versions of situational information containing only the most relevant characteristics of the situations’ temporal patterns. This process transforms situation information $S_{m \times n}$ into $S'_{1 \times 3n}$, by extracting temporal information about each stimuli such as their average value over time, number of local maxima and skewness. These compacted situations are then delivered to a binary classification tree for learning and prediction.

The tree associates the emotional meaning of a situation with its temporal pattern. Then, whenever an emotionally uncategorised situation arrives, the tree attempts to predict its emotional meaning by comparing the temporal properties of this situation and of those previously learned. If the tree finds a match for that situation pattern, then it returns the emotional category linked to that pattern, which will be either safe or aversive. Ultimately, SAFEL’s final output is the emotional category retrieved by the classification tree and indicates whether something aversive is likely to happen in the near future.

4 RoboCup Use Cases

Robot soccer poses a great challenge to robotics. It is real-time and takes place in a highly dynamic environment, perceived by means of the robots’ sensors, which are in many cases susceptible to reading failures and noises. In such a challenging environment, advanced techniques and algorithms for flexible decision making, adaptation and fast reaction are essential. In this section, we discuss possible scenarios in the RoboCup SPL competition where a mechanism such as SAFEL, which provides flexible decision-making capabilities and real-time adaptation, would be desirable, if not essential.

4.1 Individuals’ Adaptation

As mentioned in previous Sections, most research related to robot soccer focuses on improving teamwork and cooperative behaviour. Teamwork is undoubtedly crucial in soccer, but its effectiveness is limited to the level of skill of the team members. Good cooperation strategy is of little aid if the members of the team are unqualified. In addition, despite the intrinsic team-work nature of soccer, there are many situations in which individual players find themselves isolated from the rest of their team. In these cases, they have no choice but to rely on their own skills and decision-making capabilities. For this reason, we argue that mechanisms to improve the adaptation skills and flexible decision-making of




Fig. 2. Scenario example. Red squares represent players from team A, blue circles represent players from team B and the black circle represents the ball. The circle and square with thicker border line represent the respective team’s goalkeeper.

individual players are also essential for robot soccer. This is, however, a neglected area of study in RoboCup.

We propose the scenario exemplified in Fig. 2: suppose a match between team A and team B, where team A is currently attacking. Now suppose that a defender from team B manages to take possession of the ball and switch fields (i.e., pass the ball from one side of the field to the other in one shot). Because team A was fully engaged in the attack, all members of team A are in team B’s side of the field, except the goalkeeper. Finally, also consider that, second by the goalkeeper from team A, the striker from team B is the closest player to the ball at this moment. The ball stops closer to the goalkeeper, but far enough so that the goalkeeper would have to leave the goal area vulnerable in order to reach it.

This hypothetical scenario is a clear example of a situation in which the goalkeeper is isolated from the rest of its team and is forced to rely only on its own judgement and skills. The striker from team B will certainly reach the ball before the other players of team A unless the goalkeeper intervenes in some way. The decision to be taken by the goalkeeper is, therefore, whether to intervene or not. Intervening would imply in leaving the goal area unattended and, consequently, vulnerable. On the other hand, not intervening would give an obvious advantage to the opponent striker for a clear shot to goal.

The answer to this question is not straightforward because it depends on the profile of the opponent team. If the opponent striker has weak shot, for instance, then it is likely to need more than one kick to attempt a goal, giving team A

time to retreat and aid in the defence. Also, if the striker's first kick is not strong enough, then the ball will consequently be closer to the goalkeeper so that it does not need to leave the goal area unattended in order to reach the ball. Thus, in this case, staying in place, protecting the goal and waiting for help to arrive is more advantageous to the goalkeeper.

On the other hand, if the striker has a strong shot and a good aim, it may be worth to risk leaving the goal area and trying to reach the ball first, since staying in the goal area would likely result in a scoring opportunity for team B. At first, it may seem a trivial problem that depends only on whether the striker has strong kick or not. However, there are many factors involved, including the distances and angles between the elements of interest in this situation (i.e., the striker, the goalkeeper, the goal and the ball).

The kick that is weak from a particular position in the field, may be enough to score a goal from another position if we consider the angle and distance between the ball, the goal and the goalkeeper. The complexity of the problem can be further increased by the number of undesirable outcomes that we intend to avoid. For instance, in the above example, we defined that goals are the only outcome to be avoided. However, if we add collisions as another undesirable outcome, the behaviour profile of opponent teams will diverge even more, and the sequence of events leading to the undesired outcomes (goals and collisions) will become even more complex.

It is also worth mentioning that the sequence of events and their outcomes can only be analysed and associated after they occur at least once. Because it is not part of the pre-existing set of knowledge of team A, a learning process must take place at runtime. This means that the situation of interest must be experienced by the individuals first in order for them to acquire new knowledge about the world. Only then an association between environmental stimuli can be induced, leading to the necessary adaptations in behaviour.

We are currently working on a case study based on the scenario discussed above. In this case study, the goalkeeper is placed under the same isolated situation above described and tested against four different team behaviours, where both goals and collisions are considered to be undesirable outcomes. By using SAFEL's mechanism of fear learning, we expect the goalkeeper to learn when it is advantageous to leave the goal area vulnerable and go for the ball, considering the sequence of events and their outcome for each particular team behaviour and the positioning of the elements of interest in the field.

4.2 Cooperation and Team Work

In real-life soccer, human players commonly use both pre- and post- coordination strategies in conjunction. Soccer tactics usually involve the training of an agreed formation and strategy prior to the match, which is the pre-coordination phase. Nevertheless, unforeseen events may occur during the match, forcing teammates to communicate and adapt the team's strategy, which can be seen as a post-coordination phase. While pre-trained coordination is well developed and studied

in the RoboCup competition (as discussed in Section 2), the development of post-coordination strategies is still overlooked.

Although SAFEL’s adaptive learning is intended mainly for robots as individuals, we claim that it can be successfully applied to multi-robot tasks by improving collaborative behaviour and post-coordination. Among the varied methods with which post-coordination can be accomplished using SAFEL, we highlight two approaches, which are discussed below.

Anticipated Help Request In many soccer situations, an undesirable outcome may be unavoidable, regardless of a player success in correctly predicting it and taking the appropriate actions. For instance, consider the example given in Section 4.1. The best way to avoid a goal when dealing with a particular opponent behaviour may be to leave the goal area and try to reach the ball first. However, this may also increase the chances of collision, since both striker and goalkeeper will attempt to reach the ball at the same time.

It may seem a dead-end situation, but using SAFEL’s predictions to send an anticipated help request to the rest of the team is a valid and advantageous action in this case. By requesting help before it is actually needed, the goalkeeper allows its teammates to act with antecedence and, perhaps, aid in situations where help would be impracticable without the opportunity to anticipate their actions.

The goalkeeper, knowing that a particular sequence of events recurrently leads to an undesirable outcome, could message its teammates and warn them of its current situation. The teammates, in turn, could use SAFEL for learning to predict the goalkeeper’s warning messages, by associating these messages with the sequence of events that recurrently precede them. Ideally, the teammates would become capable of predicting when the goalkeeper will be in trouble before the goalkeeper itself and have more options towards preventing undesirable outcomes. This cascade of predictions would not only improve the team’s collaborative behaviour, but also potentially enhance the understanding among teammates to the point where they may be aware of each other’s situation before any message is exchanged.

In the above discussion, we have instantiated the example of the goalkeeper, but the same approach could be used to any other playing role in a variety of situations. This is because SAFEL’s approach for fear-learning is domain independent, thus applicable to any scenario where predictions based on temporal information are necessary for environmental adaptation, in and outside the RoboCup domain.

The Coordinator-Robot Approach Similarly to SAFEL’s proposal, the work of Ros et al. [18] is based on a fully distributed design, where each robot has an independent reasoning and perception of the world. Communication between teammate robots is allowed, but there is no global perception or control.

In order to solve coordination and collaborative behaviour using their decentralised model, Ros et al. propose an approach in which one single robot is selected as the “coordinator”. The coordinator is responsible for reasoning over

the current problem, given the state of the world, and messaging to the remaining teammates the sequence of actions that should be executed by the group in order to solve that problem.

This approach provides a simple solution to the post-coordination problem while preventing decision-making conflicts that could arise from a distributed system. Ros et al. also argue that this approach is useful for heterogeneous teams, where the role of coordinator could be assigned to the group of robots with higher computing power, while the remaining robots would only execute the actions calculated by the coordinators.

Like the work of Ros et al., SAFEL’s adaptive fear learning is also based on decentralised perception and reasoning. For this reason, we argue that the approach of Ros et al. for coordination and collaboration could also be used with SAFEL. The inputs for SAFEL could be, for instance, the global state of the world (e.g., ball position, teammates positions, opponents positions, collisions, goals, etc.), estimated from the local state of affairs perceived by each teammate. Then, SAFEL’s emotional response would indicate whether there is an imminent threat that the team should prioritise.

For instance, suppose that team A has used a machine learning algorithm to pre-train cooperative behaviour for defence purposes, which is based on the attacking behaviour of the other teams in the previous years of the RoboCup competition. Let’s also suppose that team B has created a novel attack strategy, which is considerably distinct from those used by the other teams. Naturally, the pre-trained defence strategy of team A would be ineffective against team B.

In this scenario, SAFEL could be used with one or more coordinator robots. If there is a pattern in the attacking formation of team B that recurrently precedes goals from team B, then SAFEL would associate both and warn the coordinator robot whenever that formation pattern occurs again. This would give the coordinator the chance to act before another goal is scored, by adapting the defence strategy of the team. For example, instead of using the default pre-trained sequence of actions, the coordinator could change it to a more risky or aggressive defence, such as sending all teammates to defend the goal area, which should be used only in extreme situations. In this example, the adaptation capabilities provided by SAFEL go further than individuals’ level, by adapting the behaviour of the team as a whole.

4.3 Drop-in Competition

The drop-in competition, introduced in 2013 [5], encourages the creation of agents capable of coordinating and co-operating with other teammates in an ad-hoc manner. In this competition, robots of different RoboCup teams collaborate as a single team towards a common goal: win the match with the highest goal difference possible.

The biggest challenge of the drop-in competition reside in the the lack of pre-coordination, which affects the players’ capability to properly communicate. Because of the limited and possibly misleading communication in the drop-in competition, many RoboCup teams do not completely trust their teammates’

messages. Doing so could mislead their robots towards engaging in a disadvantageous or non-intelligent behaviour, which would negatively affect their score in the competition. According to Genter et al. [5], some RoboCup teams have developed strategies to determine the reliability of their teammates' messages and skills. Others, in turn, have solved the issue by simply not accepting any incoming messages, or even by not communicating at all with teammates. This has led RoboCup teams to converge to a single strategy in the drop-in competition that Genter et al. describe as "Play the ball if it is close and/or no other robot wants to play the ball. Take a supporting position otherwise."

This simplistic strategy clearly diverges from the main goal of the drop-in competition. Genter et al. argue that strategies for ad-hoc teamwork can be improved in many ways, among which they mention the creation of mechanisms to: (1) evaluate the reliability of each teammate's communication; (2) evaluate the relative skill of each teammate in particular roles and positions; and (3) identify when it is better to make or receive a pass. We argue that the fear-learning mechanism of SAFEL fulfils the exact needs indicated by Genter et al., since all the three above-listed real-time evaluations can be performed in terms of associations between sequences of events and their undesirable outcomes.

5 Conclusion

In this paper, we have discussed the current gaps in the areas of contextual perception, real-time adaptation and flexible decision-making within the RoboCup competitions, mainly in respect to the Standard Platform League (SPL). Later, we proposed SAFEL as a potential approach for tackling these gaps. SAFEL is an emotional model for artificial fear-learning, which is inspired by well-known neuroscience findings on the brain's mechanisms of fear learning. Ultimately, we discussed scenarios within the RoboCup SPL where real-time adaptation is essential at both individual and multi-agent levels and suggested approaches for addressing these situations using SAFEL's associative learning. Future work involves using SAFEL to implement the scenario discussed in Section 4.1.

Acknowledgments. The first author is financially supported by CAPES, a Brazilian research-support agency (process number 0648/13-2).

References

1. Ahmadi, M., Keighobadi-Lamjiri, A., Nevisi, M.M., Habibi, J., Badie, K.: Using a Two-Layered Case-Based Reasoning for Prediction In Soccer Coach. In: Proceedings of the International Conference of Machine Learning; Models, Technologies and Applications (MLMTA), pp. 181–185 (2003)
2. Bali, M.: Drools JBoss Rules 5.X Developer's Guide. Packt Publishing (2013)
3. Dey, A.K.: Understanding and Using Context. *Personal Ubiquitous Comput.* **5**(1), 4–7 (2001). DOI 10.1007/s007790170019
4. Ferrein, A., Steinbauer, G.: 20 Years of RoboCup. *Künstliche Intelligenz* **30**(3-4), 225–232 (2016). DOI 10.1007/s13218-016-0449-5

5. Genter, K., Laue, T., Stone, P.: Benchmarking Robot Cooperation without Pre-Coordination in the RoboCup Standard Platform League Drop-In Player Competition. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 3415–3420 (2015). DOI 10.1109/IROS.2015.7353853
6. Genter, K., MacAlpine, P., Menashe, J., Hannah, J., Liebman, E., Narvekar, S., Zhang, R., Stone, P.: UT Austin Villa: Project-Driven Research in AI and Robotics. *IEEE Intelligent Systems* **31**(2), 94–101 (2016). DOI 10.1109/MIS.2016.35
7. Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawai, E., Matsubara, H.: *RoboCup: A challenge problem for AI and robotics*, pp. 1–19. Springer Berlin Heidelberg (1998). DOI 10.1007/3-540-64473-3_46
8. LeDoux, J.: *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. Phoenix, London (1999)
9. LeDoux, J.: The Emotional Brain, Fear, and the Amygdala. *Cellular and Molecular Neurobiology* **23**(4–5), 727–738 (2003). DOI 10.1023/a:1025048802629
10. Nitschke, G.: Emergent Cooperation in RoboCup: A Review. In: *RoboCup 2005: Robot Soccer World Cup IX*, pp. 512–520. Springer Berlin Heidelberg (2005). DOI: 10.1007/11780519_48
11. Pavlov, I.P.: *Conditioned Reflexes: An Investigation of the Physiological Activity of the Cerebral Cortex*. Oxford University Press (1927)
12. Pereira, I., Costa, P., Almeida, J.: A Rule-Based Platform for Situation Management. In: 2013 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), pp. 83–90 (2013). DOI 10.1109/CogSIMA.2013.6523827
13. Rizzi, C., Johnson, C.G., Fabris, F., Vargas, P.A.: A Situation-Aware Fear Learning (SAFEL) model for robots. *Neurocomputing* **221**, 32–47 (2017). DOI 10.1016/j.neucom.2016.09.035
14. Rizzi, C., Johnson, C.G., Vargas, P.A.: Improving the predictive performance of SAFEL: A Situation-Aware FEAr Learning model. In: 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 736–742 (2016). DOI 10.1109/ROMAN.2016.7745201
15. Rizzi Raymundo, C., Dockhorn Costa, P., Almeida, J., Pereira, I.: An Infrastructure for Distributed Rule-Based Situation Management. In: 2014 IEEE International Inter-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), pp. 202–208 (2014). DOI 10.1109/CogSIMA.2014.6816563
16. Rizzi Raymundo, C., Johnson, C.G.: An Artificial Synaptic Plasticity Mechanism for Classical Conditioning with Neural Networks. In: Z. Zeng, Y. Li, I. King (eds.) *Advances in Neural Networks ISNN 2014, Lecture Notes in Computer Science*, pp. 213–221. Springer International Publishing (2014). DOI 10.1007/978-3-319-12436-0_24
17. Rizzi Raymundo, C., Johnson, C.G., Vargas, P.A.: An architecture for emotional and context-aware associative learning for robot companions. In: 24th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 31–36 (2015). DOI 10.1109/ROMAN.2015.7333699
18. Ros, R., Arcos, J.L., Lopez de Mantaras, R., Veloso, M.: A case-based approach for coordinated action selection in robot soccer. *Artificial Intelligence* **173**(910), 1014–1039 (2009). DOI 10.1016/j.artint.2009.02.004
19. Whiteson, S., Kohl, N., Miikkulainen, R., Stone, P.: Evolving Keepaway Soccer Players through Task Decomposition, pp. 356–368. Springer Berlin Heidelberg (2003). DOI 10.1007/3-540-45105-6_41