

FRA-UNited — Team Description 2019

Thomas Gabel, Fabian Sommer, Steffen Breuer, Eicke Godehardt

Faculty of Computer Science and Engineering

Frankfurt University of Applied Sciences

60318 Frankfurt am Main, Germany

{tgabel|godehardt}@fb2.fra-uas.de, {fsommer|sbreuer}@stud.fra-uas.de

Abstract. The main focus of FRA-UNited’s effort in the RoboCup soccer simulation 2D domain is to develop and to apply machine learning techniques in complex domains. In particular, we are interested in applying reinforcement learning methods, where the training signal is only given in terms of success or failure. In this paper, we review some of our recent efforts taken during the past year.

1 Introduction

The soccer simulation 2D team FRA-UNited is a continuation of the former Brainstormers project which has ceased to be active in 2010. The ancestor Brainstormers project was established in 1998 by Martin Riedmiller, starting off with a 2D team which had been led by the first author of this team description paper since 2005. Over the years, a number of sister teams emerged (e.g. the Tribots, Twobots, or Icebots) participating in real robot leagues. Our efforts in the RoboCup domain have been accompanied by the achievement of several successes such as multiple world champion and world vice champion titles as well as victories at numerous local tournaments throughout the first decade of the new millennium.

While the real robot teams mentioned were closed down entirely, the 2D team has been in suspended mode since 2010 and was re-established in 2015 at the first author’s new affiliation, Frankfurt University of Applied Sciences, reflecting this relocation with the team’s new name FRA-UNited.

As a continuation of our efforts in the ancestor project, the underlying and encouraging research goal of FRA-UNited is to exploit artificial intelligence and machine learning techniques wherever possible. Particularly, the successful employment of reinforcement learning (RL, [6]) methods for various elements of FRA-UNited’s decision making modules — and their integration into the competition team — has been and is our main focus. Moreover, the extended use of the FRA-UNited framework in the context of university teaching has moved into our special focus. So, we aim at employing the 2D soccer simulation domain as a fundament for teaching agent-based programming, foundations of multi-agents systems as well as applied machine learning algorithms.

In this team description paper, we refrain from presenting approaches and ideas we already explained in team description papers of the previous years.

Instead, we focus on recent changes and extensions to the team as well as on reporting partial results of work currently in progress. We start this team description paper, however, with a short general overview of the FRA-UNited framework. Note that, to this end, there is some overlap with our older team description papers including those written in the context of our ancestor project (Brainstormers 2D, 2005–2010) which is why the interested reader is also referred to those publications.

1.1 Design Principles

FRA-UNited relies on the following basic principles:

- There are two main modules: the world module and decision making
- Input to the decision module is the approximate, complete world state as provided by the soccer simulation environment.
- The soccer environment is modeled as a Markovian Decision Process (MDP).
- Decision making is organized in complex and less complex behaviors where the more complex ones can easily utilize the less complex ones.
- A large part of the behaviors is learned by reinforcement learning methods.
- Modern AI methods are applied wherever possible and useful (e.g. particle filters are used for improved self localization).




Fig. 1. The Behavior Architecture

1.2 The FRA-UNited Agent

The decision-making process of the FRA-UNited agent is inspired by behavior-based robot architectures. A set of more or less complex behaviors realize the agents' decision making as sketched in Figure 1. To a certain degree this architecture can be characterized as hierarchical, differing from more complex behaviors, such as “no ball behavior”, to very basic, skill-like ones, e.g. “pass behavior”. Nevertheless, there is no strict hierarchical sub-divisioning. Consequently, it is also possible for a low-level behavior to call a more abstract one. For instance, the behavior responsible for intercepting the ball may, under certain circumstances, decide that it is better to not intercept the ball, but to focus on more defensive tasks and, in doing so, call the “defensive behavior” and delegating responsibility for action choice to it.

2 On the Use of Wiretapping Opponent Agent Communication

In last year’s team description paper, we have put the main emphasis on a deep learning approach for intercepting and decoding opponent agent communication using deep convolutional neural networks [3]. In so doing, we developed and implemented a neural network-based model that tries to predict whether an opponent message contains one of several considered chunks of communication. (e.g. a pass announcement or ball/player positions). If, for example, we find that the message contains a pass announcement, we try to get more information out of the message such as the movement vector of the announced pass by separate convolutional networks. For more details on this approach we refer to [5].

Besides, in [3] we presented a software engineering-focused approach that allowed us to use the established TensorFlow framework [1] into our team. Essentially, we engineered a suitable workflow for creating a shared library (C++) that can conveniently be linked to our (also, C++-based) FRA-UNited team binary and, hence, be utilized during test or competition matches. While these research questions and software engineering aspects could be solved successfully, a thorough analysis of the possible impact of the entire approach on the average performance of our team was still outstanding.

2.1 Empirical Analysis of Impact on Team Performance

Our experiments focused on four versions of our team (labelled A) and the 2016 and 2017 versions of Helios (labelled B [2]).

- (a) *A_Baseline17*: Version of FRA-UNited used at RoboCup 2017
- (b) *A_NoComm*: While no change to the agents’ playing behavior has been made, the agents of FRA-UNited no longer use *any* form of inter-agent communication. So, this version should perform worse than the baseline.
- (c) *A_NoOwnPasses*: No change in agent playing behavior, but passes will no longer be announced or communicated among the players of FRA-UNited. All other aspects of team-internal communication (like communicated player information or ball data) remain untouched.
- (d) *A_OppPassExploit*: This communication behavior variant was in the center of our interest since it represents the baseline version enhanced by the implementation of the approach to eavesdrop and understand opponent pass announcements, whose basic ideas are mentioned above.

The actual exploitation of an intercepted opponent pass announcement in (d) has intentionally been implemented in a straightforward manner in order to facilitate a fair assessment of the impact of the approach. Each decoded opponent pass is treated in exactly the same manner as a pass announcement received from a teammate (like in (a) or (b)). Accordingly, no additional logic for opponent passes were programmed, which allows us to assess the benefits of having wiretapped the opponent team as accurately as possible. Hence, from the point of view of an *A_OppPassExploit* agent, there are just “a few more” pass announcements compared to the baseline version of the agents.

2.2 Results

All four versions of FRA-UNITed were matched against both versions of Helios. For each combination, 5000 matches were played in order to form score averages and, hence, to account for the stochastic nature of the Soccer Server.

Figure 2 shows the variability in playing strength of the four FRA-UNITed version (Team A) considered. Here, *A_Baseline17* is considered as the baseline (100%) against which the three other variants are compared with respect to the relative amount of goals scored and conceded against both versions of Helios (Team B, baseline version from 2017 and predecessor version from 2016). Consistently, an increased utilization and exploitation of communicated information brings about an increased overall performance. Interestingly, when playing against *B_Baseline17*, the activation of our wiretapping and opponent pass announcement decoding approach yields an increase in the number of goals shot by 8.8% and a simultaneous decline in the number of goals conceded by 4.4%. Besides this and as a very general note, we can draw a general conclusion on the importance of communication as a whole in soccer simulation since, apparently, the playing performance (in terms of numbers of goals) drops/increases by more than 50% when switching off/on the entire team-internal communication.




Fig. 2. Taking FRA-UNITed version *A_Baseline17* as the ground truth, these charts visualize the relative changes in the average numbers of goals scored and goals conceded when the other three considered versions are matched against *Team B* (left: against Helios 2016, right: against Helios 2017).

2.3 Discussion

Here, we have presented only a brief excerpt of all the empirical experiments we conducted for analyzing the impact of communication and of wiretapping and decoding opponent pass announcements. Much more details on further results and conclusions can be found in a separate paper [4]. As a very basic note, we can draw a general conclusion on the importance of communication as a whole in soccer simulation since, apparently, the playing performance (in terms of numbers of goals) drops/increases by more than 50% when switching off/on the entire team-internal communication. With our experiments, we could show that the playing performance of a team that wiretaps an opponent and exploits

intercepted information (in our case pass announcements) can be boosted significantly. Clearly, when switching to another opponent a separate decoding model in the form of a deep convolutional neural network would have to be trained beforehand. Also, we should emphasize the fact that the advantages reported can only be exploited in real tournament games under the assumption that the opponent does not change or encrypt its communication. In the experimental evaluation at hand, our selected opponent (Helios) was the current world champion. Thus, all conclusions refer to this opponent in the first place. However, as pointed out by [5], similar or even better communication learning performance can be expected, when playing against other 2D top teams.

3 Goal Keeper Renewal

The implementation of FRA-UNited’s goal keeper dates back to 2003 with Daniel Withopf as its main author. Over the years, this goalie has witnessed a number of successes, including three world champion titles, three vice world champion titles, and six consecutive European-level titles¹. Interestingly, almost no changes were made to the strategic and tactical part of the goalie agent program throughout the years, except for minimal modifications that were required due to changes in the physics and rules of the soccer simulator.

For RoboCup 2018 in Montreal, we started to implement a new goal keeper behavior with the aim to get a more effective and maintainable goal keeper. We allowed a prototype to play at RoboCup 2018, in order to collect some first practical experience in the context of competitive games. Because of some clear goalie mistakes, we finally decided to revert to the original goal keeper during the ladder games². Nevertheless, the new goal keeper performed very well and we could use the experience in the tournament to improve the behavior.

The main intention behind the development of the new goalkeeper, was to close the critical gaps of the old FRA-UNited goal keeper. For example, it used to have difficulties in catching balls, which are kicked from a sharp angle onto the short post. Also, it has a passive and defensive behavior in terms of thwarting aggressive plays from the opposite team, like intercepting long passes behind the own defense line to stop the attack. Because maintaining and updating the old FRA-UNited goal keeper would require significant reverse engineering effort, we decided to start from the scratch. First, we developed a new movement behavior on the goal line. The goalkeeper now moves on a trapeze, so that it can cover a large area of the goal and in addition protect the short post more reliably, if an opponent attacks from a sharp angle.

¹ Having played at international competitions for 15 years, the duration of the international career of this goalie is comparable to Diego Maradona (1979-94), Zlatan Ibrahimovic (2001-16), or – matching perfectly the time frame – Wayne Rooney (2003-18).

² As ill luck would have it, we lost the game for rank 4 due to a known “short angle weakness” of the old goalie. Both goals that we conceded in that game exploited a flaw in the old goalie’s implementation.

Also new features are implemented, for example the new goal keeper is now able to distinguish between a goal shot from an opponent and a back pass. This is done by checking, which player has the ball in his control area and whether that player has recently performed a kick or tackle action. When a back pass is detected, the goal keeper tries to save the ball with a kick out of the dangerous zone; only if it finds no possibility to clear the ball, it will catch it.

Basically, the renewed goal keeper is aligned more offensively. It plays more aggressively in intercepting passes played behind our defense line as it is sure to be the first player at the ball. Therefore, it positions itself in a more offensive spot compared to the old goalkeeper. Current analysis of the current version, which is still work in progress, indicate a reduction of the number of goals conceded by 4% while gaining a far more maintainable implementation.

4 Conclusion

In this team description paper we have outlined the characteristics of the FRA-UNITed team participating in RoboCup's 2D Soccer Simulation League. We have stressed that this team is a continuation of the former Brainstormers project, pursuing similar and extended goals in research, development as well as for teaching purposes. Specifically, we have put emphasis on our most recent research activities and practical implementation of our results.

References

1. Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems (2015), <http://tensorflow.org/>, software available from tensorflow.org
2. Akiyama, H., Nakashima, T., Fukushima, T.: HELIOS2017: Team Description (2017), www.robocup2017.org/file/symposium/soccer_sim_2D/TDP_HELIOS2017.pdf, Supplementary to RoboCup 2017: Robot Soccer World Cup XXI
3. Gabel, T., Klppner, P., Godehardt, E., Tharwat, A.: FRA-UNITed - Team Description 2018 (2017), Supplementary material to RoboCup 2018: Robot Soccer World Cup XXII
4. Gabel, T., Klppner, P., Godehardt, E., Tharwat, A.: Communication in Soccer Simulation: On the Use of Wiretapping Opponent Teams. In: Proceedings of RoboCup 2018: Robot Soccer World Cup XXII. Springer, Montreal, Canada (2018)
5. Gabel, T., Tharwat, A., Godehardt, E.: Eavesdropping Opponent Agent Communication Using Deep Learning. In: Proceedings of Multi-Agent System Technologies (MATES 2017). pp. 205–222. Springer, Leipzig (2017)
6. Sutton, R.S., Barto, A.G.: Reinforcement Learning. An Introduction. MIT Press/A Bradford Book, Cambridge, USA (1998)