Brainstormers 2D — Team Description 2010

Thomas Gabel and Martin Riedmiller

Machine Learning Lab
Institute of Computer Science
University of Freiburg, 79110 Freiburg, Germany
{tgabel|martin.riedmiller}@informatik.uni-freiburg.de

Abstract. The main focus of the Brainstormers' 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. Our final goal is a learning system, where we only plug in "win the match" – and our agents learn to generate the appropriate behavior. Unfortunately, even from very optimistic complexity estimations it becomes obvious, that in the soccer simulation domain, both conventional solution methods and also advanced today's reinforcement learning techniques come to their limit – there are more than $(108 \times 50)^{23}$ different states and more than $(1000)^{300}$ different policies per agent per half time. This paper outlines the architecture of the Brainstormers team and also adopts a retrospective persepctive by investigating the progress made throughout the recent years.

1 Introduction

The Brainstormers project was established in 1998, starting off with a 2D team. Ever since we have been participating in RoboCup's annual soccer simulation tournaments. Over the years, the Brainstormers Tribots (competing in RoboCup's MidSize league since 2002), the Brainstormers 3D (soccer 3D simulation, 2004–2006), as well as the Brainstormers Twobots (Humanoid League, 2008) expanded the Brainstormers team. Our work has been accompanied by the achievement of several successes such as multiple World Vice Champion titles and the World Champion titles at RoboCup 2005 in Osaka (2D), RoboCup 2006 in Bremen (MidSize), RoboCup 2007 in Atlanta (2D + MidSize), and RoboCup 2008 in Suzhou (2D).

The team description paper at hand focuses on the Brainstormers 2D, our team competing in soccer simulation's 2D league. The underlying and encouraging research goal of the Brainstormers has always been to exploit AI and machine learning techniques wherever possible. Particularly, the successful employment of reinforcement learning (RL) methods for diverse elements of the Brainstormers' decision making modules – and their integration into the competition team – has been our main focus throughout the years.

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

Our recent focus has been on a restrospective and critical investigation of the general progress that has been made in the soccer simulation 2D league during the time of its existence. Therefore, we highlight some of our findings regarding that topic in this paper (Section 2). Additionally, we devote the major part on a retrospection of the 12 years the Brainstormers project has been running pointing to some achievements and contributions made. Besides, we start off by briefly describing the overall architecture and basic design principles of our team in remainder of this section.

1.1 Design Principles

The Brainstormers 2D rely on the following basic principles:

- There are two main modules: the world module and the decision making module.
- Input to the decision module is the approximate, complete world state as provided by the soccer simulation environment.
- The soccer environment is modelled as a Markovian Decision Process (MDP).
- Decision making is organized in complex and less complex behaviors.
- 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).

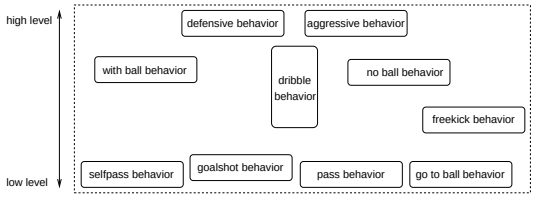

Fig. 1. The Behavior Architecture

1.2 The Brainstormers Agent

The decision making process the Brainstormers Agent is based upon 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 so doing, call the "defensive behavior" delegating responsibility for action choice to it. Our team's source code has been made publicly available in 2005 and can be retrieved from our team web site.

2 An Attempt to Retrospect

In human soccer playing, it is impossible to make sound statements about the development of the playing performance level over years. Although many experts will certainly agree that contemporary soccer exhibits more speed, dynamism, and superior athleticism than a few decades ago, an empirical proof on its superior performance is inconceivable. So, it remains unascertained whether the first winner of the European Champions Clubs' Cup (Real Madrid, 1956) would be defeated by the team of Real Madrid that won the Champions League in 2000 and, moreover, how these two teams would compare against the current "version" of FC Barcelona, Champions League winner 2009. For very human reasons, such as maturing and aging, comparisons and improvement analysis do not apply for longer periods of time.

2.1 Progress in RoboCup

With respect to long-term comparisons, one might argue that the situation is better in robot soccer. Robots are patient, could be stored away after a competition, and be unpacked after a few years in order to compete against a contemporary team. Though this is a nice thought experiment, practice in RoboCup has taught that such an approach is infeasible: The performance of soccer-playing robots strongly depends on the environment, ground, and lighting and is, thus, hard to reproduce. Besides, even unused hardware is aging, too, and is in general too expensive to be just locked away. Hence, a formal analysis of the progress of soccer-playing robots made throughout the years is difficult to establish. As a consequence, assessments concerning the exact year-to-year progress remain qualitative in nature.

RoboCup's simulation leagues adopt a special role. No hardware development and maintenance is required and software agents do not age. However, the performance of soccer-playing software agents strongly relies on the simulation of the physical world they are placed in. If the characteristics of that simulation change, a meaningful evaluation of the progress made over years is rendered impossible. The simulation league has experienced a highly dynamic history, where the simulators used have undergone tremendous changes over the years. Thus, statements about the general level of playing performance frequently remained vague and on a qualitative level, without empirical verification. For example, already in 2002 it was claimed that the "overall playing strength of the teams in the tournament was quite impressive" [1]. Moreover, "the playing level of the tournament showed increased and consistent improvement as compared to last year's tournament". Similarly, "the 2003 tournament again showed a big advance in the performance of the teams" [2] and the "level of play of the last twelve teams this year was very mature and close to each other" [3].

While we subscribe to these assessments, we need to stress that no empirical proof for their correctness exists and, in fact, cannot exist: From 2001 to 2002 and 2002 to 2003 several changes were introduced into the 2D soccer simulation

which is why a meaningful evaluation of the objective progress made in those days is infeasible.

By contrast, our focus has recently been on the subsequent five-year period of stability (2003-2007) where no changes were introduced into the Soccer Server. In those days, the goal was to retain the simulated soccer environment stable in order to facilitate measuring scientific progress from year to year. However, contrary to the original intention, year-to-year progress was not measured systematically in the time window in which the 2D soccer simulation represented a stable multi-agent testbed. We therefore tackled the following questions recently:

- What are the exact conditions that are required to allow for a *meaningful* retrospection?
- If such conditions can be identified, has there been any *provable* progress in performance in soccer simulation 2D during the recent years?
- If so, how close are teams and how reliable are the results of a single game with respect to the noise in the simulation? Has there been something like convergence in the teams' further development and has perhaps a certain saturation level in their performance been reached?

We conducted a large-scale empirical analysis of the progress made during the time interval mentioned. The exact results of this study are to appear in a yet to be published paper submitted recently (in the next section we present some of the results relating more strongly to our team). In brief, our findings indicate that the progress made during the stable period (2003-2007) is astonishing. For example, while admired for their sophisticated play in 2003 and 2004 [2], world champions of those times played at the level of a low-class team a few year later, sometimes losing in the double digits against top teams from 2006 or 2007. Moreover, given the fact that the competitive character of RoboCup competitions is nearly identical in all league, we may infer that the progress to be observed in the simulation league can serve as a showcase for other RoboCup leagues, too, even if their characteristics disallow for a quantitative comparison.

2.2 Brainstormers 2D: The Retrospective View

Over the years, the Brainstormers project has grown continuously in size and maturity. Undergraduate and graduate students as well as the team founder Martin Riedmiller have contributed piece by piece all components of the team. Figure 2 is a tribute to all the people who have actively contributed to the project.

When having a look at the left part of Figure 3, we can see that in the beginning phase of the project (1998) the number of lines of code grew quickly. This phase was characterized by attempts to realize significant parts of the agent behavior by reinforcement learning approaches and included first implementation of reinforcement learning libraries: Tools had to be developed that were capable of arranging the training process, including pattern generators, statistics modules, and of course the learning algorithms themselves.

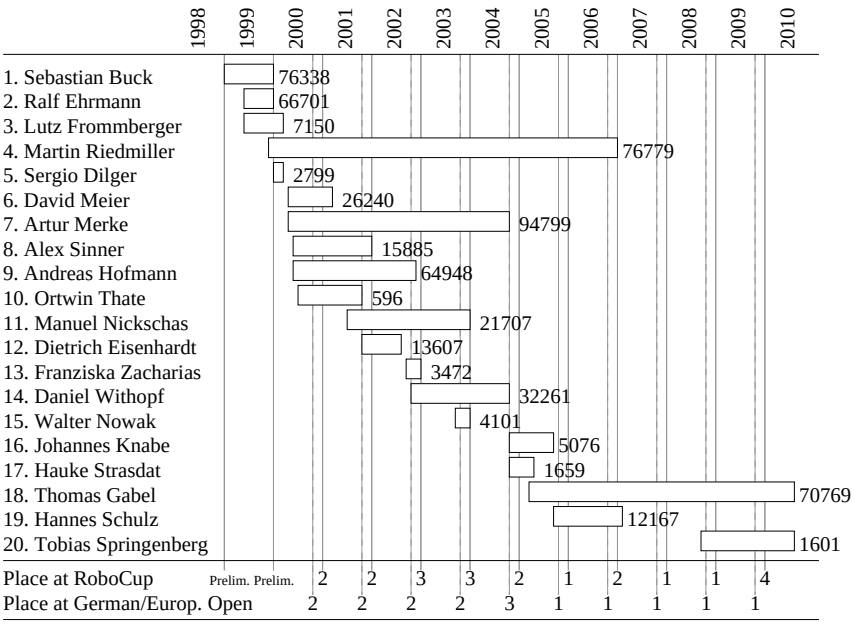

Fig. 2. Over the years, many people have contributed to the Brainstormers project. This listing shows the time intervals during which the different team members checked in source code into the Brainstormers' repository. The numbers indicate the total number of lines of code added by the respective author.

After having failed in the 1998 and 1999 tournaments, vast parts of the software were discarded and redeveloped, including various learning approaches. Most prominently, much more effort was then put into developing an accurate world model which is substantial to the successful employment of learning approaches. At the same time the Brainstormers had to keep pace with recent changes introduced into the soccer simulation environment. So, for example, the jump in the LOC curve (left part of Figure 3) at the beginning of 2002 is due to the introduction of a coach program. This period was accompanied by an astonishing number of second and third places: Our team became runner-up or third place in all competitions we participated in (till 2004, cf. Figure 2).

While in 2003 a major redesign of our agents' software architecture was undertaken (cf. Section 1.2) and the code was cleaned up heavily (see the jump in the LOC chart in Figure 3), in the subsequent period from 2004 to 2007, our team could greatly benefit from facing a nearly stable simulation environment. This allowed us to concurrently (a) redesign vast parts of our team play and (b) enhance several of the machine learning approaches we employ in such a manner that the resulting behaviors are highly competitive. As a result of this and thanks to the efforts made by all team members throughout the years, we finally achieved three world champion titles and five European level titles.

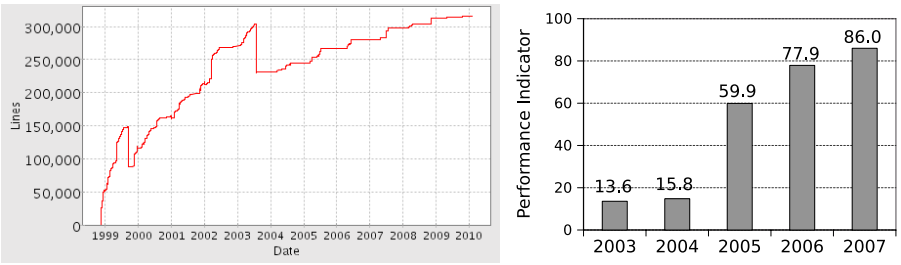

Fig. 3. Left: Number of lines of code the Brainstormers project comprises. See the text for more explanations. Right: Relative playing strength of the Brainstormers in their versions from 2003 till 2007.

As part of the empirical evaluation mentioned in Section 2.2, we also measured the performance level of our team during the stable period. Interestingly, the Brainstormers' playing strength has approximately increased sixfold from 2003 to 2007: The numbers reported in the right part of Figure 3 indicate the percentage share of points the different Brainstormers versions achieve on average, when playing repeatedly against the teams that placed among the top four during RoboCup world championship tournaments from 2003 to 2007. The corresponding average scores (averaged over 270 matches each) developed from 0.59:4.01 in 2003, to 0.60:3.65 in 2004, to 2.06:0.86 in 2005, to 2.70:0.29 in 2006, and finally to 2.87:0.21 in 2007.

3 Summary

In this team description paper we have outlined the characteristics of the Brainstormers team participating in RoboCup's 2D Soccer Simulation League. We have stressed that our main research focus lies on the development of reinforcement learning techniques and their integration into our team. We have also taken a look back to the 12 years-lasting history of the Brainstormers team as well as to the question whether verifiable progress has been made by the soccer simulation 2D league in general.

References

1. Obst, O.: Simulation League – League Summary. In: G. Kaminka, P. Lima, R. Rojas (Eds.): RoboCup 2002: Robot Soccer World Cup VI. Springer, Melbourne, Australia (2003) 443–452
2. Pagello, E., Menegatti, E., Bredenfeld, A., Costa, P., Christaller, T., Jacoff, A., Polani, D., Riedmiller, M., Safiotti, A., Sklar, E., Tomoichi, T.: RoboCup-2003. New Scientific and Technical Advances. AI Magazine **25** (2004) 81–98
3. Pagello, E., Menegatti, E., Bredenfeld, A., Costa, P., Christaller, T., Jacoff, A., Johnson, J., Riedmiller, M., Safiotti, A., Tomoichi, T.: Overview of RoboCup 2003 Competition and Conferences. In: D. Polani et al. (Eds.): RoboCup 2003: Robot Soccer World Cup VII. Springer, Padova, Italy (2004) 1–14