Automatically Improving Team Cooperation by Applying Coordination Models

Jan Wendler[*,†]
Department of Computer Science
Humboldt University Berlin
D-10099 Berlin, Germany
email: wendler@informatik.hu-berlin.de

Gal A. Kaminka
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
email: galk@cs.cmu.edu

Manuela Veloso
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
email: veloso@cs.cmu.edu

Abstract

In this paper we present general theoretical models to improve the coordination between team members. We apply these models for passing, a coordination problem in soccer, and demonstrate the usefulness of these models. For the improvement of the team a trainer has been used to extract good passing behavior out of a set of games. The trainer improves this behavior and gives advice for the agents to improve its coordination.

Introduction

Coordination is most critical for the overall success of a team. Teams are formed by individuals or other teams to pursuit a common goal. The better they want to perform their goal the more they need make use of their team capabilities. To optimize the use of their team capabilities they need to improve their coordination.

The process of improving the coordination between agent team members is still not well understood. Coordination is often addressed only in the context of overall task performance, and thus improvements are time-consuming and difficult. First the agent behavior has to be analyzed which is time consuming. After that some improvements has to be selected out of the wide range of possible improvements. Finally the improvements needs to be implemented and evaluated. This makes it difficult to improve the coordination of an agent team.

An autonomous team trainer can help to automate this process. For that he needs to know what good coordination is about. We propose to use models of coordination based on coordination theory (Malone & Crowston 1994). From this paper we apply the coordination paradigm of transferring a resource from a producer to a consumer. Such a physical transportation of a resource can be considered perfect if the right resource is delivered at the right time, in the right way and to the right consumer.

In the application scenario we use, the producer and the consumer are soccer players. The resource ball is transfered from one player to another by passing. Therefore the process of passing is a coordination problem for at least two players. We present and apply general theoretical models to improve the coordination among the soccer players during passes. We show that the models we developed are empirically predictive, in that they are able to successfully predict the pass success rate of different soccer teams which we did not develop ourselves.

Applying the theoretical models to find the optimal solution directly is very complex. For this reason it would take too much time to compute the solution by the agents itself in the game. Instead a team trainer analyze all passes of different games and applies the coordination models to compute perfect pass alternatives according to the models. For this a case based approach is used, all computed passes are collected in a case base and are used by the agents in further games to perform better coordinated passing. Whereas the improved coordination show an increased pass success rate it does not show any improved results in the score.

This paper is organized as follows: the next section explains the difficulty of the used multi agent environment by a short introduction into the simulated soccer game. The section after that will motivate, present and evaluate the theoretical coordination models. The next section shows how these models has been applied to improve the team coordination during passing. Finally we discuss related work and conclude.

The Simulated Soccer MAS

Soccer simulation is an example of a complex multi agent system (MAS) where coordination problems occur naturally, which is why we are using it for our investigation.

The *RoboCup Federation* provides a soccer simulation program for the *Simulation League* of the Robot World Cup Initiative (RoboCup). RoboCup (Kitano *et al.* 1997) is an attempt to foster AI and intelligent robotics research by providing a standardized problem and by holding yearly workshops and world championships.

Just like its human paragon, an artificial soccer player needs perception, rational and social reasoning, action execution and many more. In the simulated soccer system *SoccerServer* (Noda 1995) the perception and action execution of the agents is done via socket communication. The soccer agents observe other players, the ball and landmarks like goals, lines, and flags, if they are in their field of view. They

[*] This research has been supported by the German Research Society, Berlin-Brandenburg Graduate School in Distributed Information Systems (DFG grant no. GRK 316)

[†] Most of this work was done during a research stay at Carnegie Mellon University.

respond with primitive actions like *kick*, *dash* and *turn*. All information the agent perceives or executes is relative to the player and noisy.

Whereas the game chess maybe compared to solving a mathematical problem, artificial soccer is more like acting in real life because decision have to be make rapidly in a dynamic environment with limited local information, the perception and action execution is noisy and not reliable, communication between players is restricted, the player's stamina is restricted, the environment is continuous and requires real-time decisions.

Usually the agents use the observed information about the landmarks to estimate their position and build a world-model using the observed data about the other players and the ball. Based on this and earlier sensor information they reason what to do and decide for a behavior to execute. For this they use at least implicit models about their teammates, e.g. if an agent decides to pass into the direction of a teammate, it assumes that this teammate tries to intercept the ball. Furthermore in deciding this pass, the agent most likely has used models of the nearby opponents to avoid a pass failure.

Theoretical Coordination Models

Before we present the theoretical coordination models we motivate the need for a team trainer and the use of coordination models. Finally we evaluate the coordination models.

The Need for a Trainer

The development of a soccer team is a difficult and tedious process: the analysis of team behavior is time consuming, there are many degrees of freedom for possible improvements, and agents interactions are difficult to analyze manually. For instance the teams *AT Humboldt* (e.g. (Burkhard, Hannebauer, & Wendler 1998)) has been in development over the last 4 years. Further improvements are time consuming and frustrating, lots of data has to be analyzed, experiments has to be run repeatedly and it is not clear how to best improve a team.

We want to develop tools that will automatically upgrade team performance. An automated team trainer should observe the team-members behavior, analyze it to find weaknesses and use the analysis to improve the team's performance.

Since coordination is critical to the success of a team, one promising approach is for the trainer to try to improve the coordination among team-members. But coordination is poorly understood, and therefore difficult to analyze.

To address this difficulty, we are proposing to exploit coordination theory (Malone & Crowston 1994). We use theoretical models of coordination, to understand a team's coordinated behavior, to analyze the coordination and discover weaknesses, and to suggest improvements to team behavior.

For instance, passing in soccer is an example of a coordination problem for at least two agents. The ball can be considered as a resource who is transfered by passing from one agent to another. The agents are the producer and the consumer in the coordination paradigm of transferring a resource. The coordination between the agents can be considered perfect if the ball is passed to the right teammate at the right time and at the right way.

Doing it in the right way can really help a team which is validated by a correlation between the pass success rate and the score difference (see figure 1). 17 teams played against two fixed opponents. For every team its pass success rate was determined and related to its goal difference to the fixed opponents.

We assume that the decision to whom and at which time to pass has been already made by some higher level decision routine.

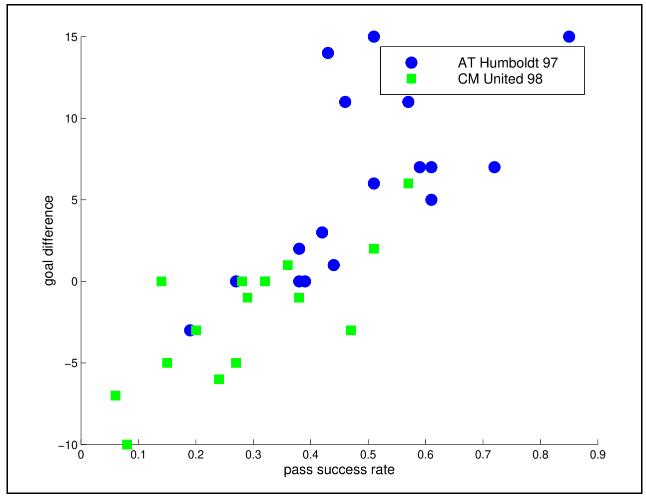

Figure 1: Correlation between pass success rate and the goal difference in games against the teams *AT Humboldt'97* and *CM United'98*

Passing the right way means for us that the receiver does not have to change its intended behavior more than necessary: the receiver has to spend the least amount of effort (stamina) in moving from the intended location to the interception point of the pass. As more stamina the receiver has to spend for this distance, as worse is the pass coordination. So the most important question is: how do we know what the receiver is intended to do?

Models of Intended Behavior

The intended location of a pass varies greatly depending on the situation of the game. It can vary between preferring to receive the ball at the current position of the receiver, or between receiving it at some target location to which the receiver is heading. Both strategies are widely used in soccer games, the first for a kind of structured game play and the second mostly for attacks. There may exist a wide range of other strategies of intended behavior so it may be that players intend to reach a fixed position on the field.

Based on the observations we propose two models of intended behavior to determine the intended location of the receiver:

Inertia model: In this model the receiver intends to slow down naturally, its intended location is determined by its inertia movement only. So this model focus on passing to the receiver's position directly.

Maintain model: In this model the receiver intends to keep the same speed and the same direction as at the pass start time, so its intended location is determined by its velocity vector and the time between the pass and the control of the ball by the teammate. This model simulates the passing to a position to whom the receiver is already heading.

In a given passing situation the kicker can choose its passing speed and angle. In figure 2 a pass situation is presented. For the kicker six different pass possibilities (three different angles and two different speeds) are shown. According to the inertia model the direct passes p_3 and p_4 are optimal because the receiver does not need to spend any effort, it slows down naturally to the position marked with a dotted circle where it awaits the ball. According to the maintain model the pass p_6 is optimal because for that pass the receiver can keep its current speed to reach the interception area, which is marked by a dashed circle, in time. The pass p_5 is too slow for the receiver to keep the same speed as at the pass start time.

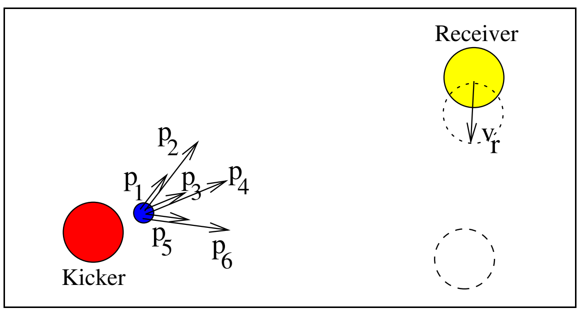

Figure 2: Kicker with six pass possibilities.

Because there are usually opponents on the soccer field the two proposed models also take care that all opponents need longer to control the ball than the target receiver. To determine the best pass according to these models the trainer assumes that both receiver and opponents have nearly perfect ball intercepting capabilities. This is a conservative constraint because most teams do not have perfect intercepting capabilities. Instead the intercepting capabilities are team dependent and therefore difficult to extract. So it is better to be on the safe side by assuming the opponent behaves best.

Just as the models account for the expected capabilities of opponents, they must also account for the capabilities of the producer of the resource, i.e. the kicker of the ball. The critical capability of the kicker for the purposes of evaluating a pass is the speed at which it can kick the ball towards the receiver. We rely on two models here: one in which the kickers capability is optimal, and one in which it is the same as its current level.

The latter model limits the pass possibilities to all passes which have same or less speed than the observed pass p_p. So if the speed of the observed pass for the situation shown in figure 2 is less than the speed of the pass possibilities p_2, p_4 and p_6 these possibilities are not considered to determine the optimal pass. According to the inertia model the pass p_3 is then still optimal. According to the maintain model the pass p_5 is now the best, but the receiver needs some effort because its intended location at the time when the ball reaches the dashed circle is far away from this area.

Combining the inertia and maintain model with the two models about the ball kicking capabilities we get the four models: **inertia-best**, **inertia-same-speed**, **maintain-best** and **maintain-same-speed**. Applying these coordination models to an observed pass p_p with a given passing speed and direction, the trainer can evaluate for a given situation how much this pass deviates from the optimal pass p_o according to these models. These deviation is measured by the stamina points the agent would need to spend by running from the intercept position of the observed pass to the intended intercept position of the optimal pass. The needed stamina for a straight route is directly related to its length and therefore directly related to the effort made by a player when running along this route.

To summarize, we propose 4 models of coordination in passing, which differ in their prediction of the intended behavior of the receiver, and in the passing capabilities of the kicker.

Evaluation of the Theoretical Models

To motivate these models practically we analyzed the evaluation session data of the RoboCup 99 world championship (Kaminka 1998). During this session, 17 teams played against two fixed opponent teams: *CM United'98* (Stone, Veloso, & Riley 1999) and *AT Humboldt'97* (Burkhard, Hannebauer, & Wendler 1998) [1]. We recognized that teams that behaved good according to the theoretical models are good in passing. We identified a negative correlation between the deviation from an optimal pass according to the models and the pass success rate.

Figure 3 shows this correlation for the 17 games which was played against *AT Humboldt'97* for the models inertia-best and inertia-same-speed. The average deviation from the optimal pass and the pass success rate were computed across all observed passes of each team. The x-axis shows the deviation in stamina points and the y-axis the success rate between 0 and 1. The teams with less deviation from the optimal pass effort coordinate better according to our models. A value of 0 in the deviation would indicate an optimal passing team.

The correlation values support our hypothesis that better-passing teams (in terms of pass-success rate) will behave as predicted by the models we introduced. Correlation values go from -1 to 1, and in our case a correlation of -1 would mean perfect correlation since negative correlation between pass-success rate and effort deviation means that the better the success rate, the less deviation from our models exist.

The correlation between the two models inertia-best resp. inertia-same-speed (cp. figure 3) and the pass success rate is -0.46 resp. -0.38. The correlation between the

[1] *CM United'98* was winner of RoboCup 98, *AT Humboldt'97* winner of RoboCup 97

maintain-best resp. maintain-same-speed and the pass success rate is -0.44 resp. -0.38.

These correlation values are quite good considering that the 17 evaluated teams are very different in play style and strength. Without the three worst outlier we get a correlation of -0.80 resp. -0.71 for the inertia models and a correlation of -0.86 resp. -0.77 for the maintain models, which is very high. For the inertia models the three outlier are presented in the figure by the unfilled circles and squares. For two of the outlier only a very few successful passes could be recognized so the statistics for these two teams may even be incorrect. The correlation values between all models and the pass success rate for the games against *AT Humboldt'97* are shown in the following table.

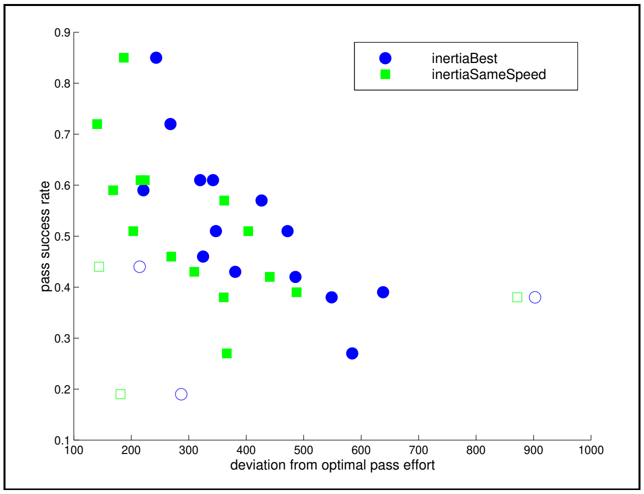

Figure 3: Correlation between deviation from optimal pass effort according the inertia models and the pass success rate in games against *AT Humboldt'97*

coordination model	correlation for all games	correlation w/o outlier
inertia-best	-0.46	-0.80
inertia-same-speed	-0.38	-0.71
maintain-best	-0.44	-0.86
maintain-same-speed	-0.38	-0.77

In the games against *CM United'98* the correlation values are a little bit worse but show the same tendency. Without the three worst outlier the correlation is less than -0.5 for all models.

Without the three worst outlier we have a strong negative correlation between the deviation from the optimal pass effort and the pass success rate. This provides empirical evidence that the models we introduce, successfully predict actual coordination by multi-agent systems not of our own design. With the previously demonstrated correlation between coordination (pass success-rate) and overall performance (goal difference), we now turn to using the coordination models to improve the performance of a team by modifying its passing behavior to conform to the models.

Applying the Models to Improve Coordination

This section starts with a description of how optimal coordination can be computed in general. Then our implementation for this computation is presented. The section finished with the results of the implementation and its discussion.

Calculating Optimal Coordination

In the last section we proposed and defined the optimal coordination solution according to the proposed models. To determine the optimal coordination out of all coordination solutions one has to search through the space of all solutions and compare them. For this we test a huge amount of transfer possibilities by simulation and evaluate and compare the possibilities to find the optimal transfer solution. This is computational intensive.

Furthermore for each transfer possibility the movement of the transfered object, the intended receiver and interacting adversaries who hinder the transfer has to be simulated. Whereas the movement of the transfered object is mostly well-known and the movement of the intended receiver can often be predicted easily, for the movement of the adversarial agents some models about them need to be applied. A safe strategy for modeling the adversarials is to expect that they behave optimally. During the simulation two time points can be determined, the earliest and the eventual time point for a successful transfer. The latter time point is given by the time the best adversarial needs to interrupt the transfer. For all time points between the earliest and the eventual time point the effort spent by the receiver for a successful transfer is determined, and its minimum gives the effort of the receiver for the simulated transfer. The best transfer is given by the transfer who minimizes the effort of all possible transfers.

With a time-consuming calculation we simulate a wide range of passes and test them for optimality. Such a calculation cannot be done online by the soccer agents during the game. So we decided to perform this calculation off-line by a trainer. The trainer saves the solutions in a case base which is used by the agents.

We choose a case based approach because of these two valuable properties:

- A case base is incremental. It is very easy to add new cases with new solutions even at runtime. The agents can communicate their cases among each other.
- The cases are readable by a human observer, he can evaluate the usefulness of particular cases and can even add, remove or modify the solutions of the cases.

Implementation

To acquire enough passes the trainer analyzed 30 games. For every detected successful pass it computed the passes who fulfilled our four models best. For this a simulation of 900 different passes was done, starting from the current pass direction, 30 different directions in a sector of 90 degree were simulated for 30 different passing speeds. For every simulated pass the ball, the receiver and all opponents were simulated to determine how much time the receiver and the best positioned opponent would need to intercept the passed ball (part 1). After that the minimal needed effort by the receiver was calculated for this pass (part 2). The overall minimal needed effort among all 900 passes results in the optimal pass and was saved as a case if it existed (part 3). Algorithm 1 shows the pseudo code for the computation of the optimal pass. If no optimal pass existed for a given pass situation, which means that it is impossible to pass the ball to a teammate if the opponent behaves optimal, no case was generated.

Algorithm 1 Computation of the optimal pass

```
for every pass (out of 900) do {
   // part 1
   until an opponent controls the ball {
      simulate ball by one step
      simulate intercepting receiver
      simulate intercepting opponents (all)
   } returns {
      receiverInterceptionTime
      bestOpponentInterceptionTime
   }
   // part 2
   for all times between receiverInterceptionTime and
         bestOpponentInterceptionTime do {
      compute minimalEffortByReceiverForTime
   } returns {
      minimalEffortByReceiverForPass
   }
   // part 3
   if( minimalEffortByReceiverForPass < minEffort ) {
      minEffort = minimalEffortByReceiverForPass
      optimalPass = pass
   }
}
```

Out of the 30 different games 2031 passes (cases) were extracted. Such an extracted case consist of a pass description and the best solution for six different passing strategies. Four of the strategies are just given by following the theoretical models, the other two strategies: **current** and **fastest** were added to compare the four models independently from the chosen implementation. For the strategy **current** the values of the successful pass are used as solution, the strategy **fastest** saves the same solution but only if the pass had been even successful against optimal ball intercepting opponents.

The description of a case consists of:

- the vector between passer and receiver,
- the absolute value of the receiver's speed,
- the body direction of the receiver,
- the vector between the passer and the opponent who can intercept the ball fastest, and
- the vector between the passer and the opponent goal.

The solution of a case consist of:

- the vector between the passer and the ball interception point for the receiver, and
- the ball passing speed.

To use these cases and its solution strategies by the agent team *AT Humboldt'2000* (Burkhard *et al.* 2001) we exchanged the pass behavior of the team by the new case based passing behavior. A case matches the given situation if at least the following conditions are true:

- It exist a possible receiver and the vector to the receiver does not differ significantly from the "passer to receiver vector" of the case.
- There exist no opponent which is nearer to the interception point of the case than the opponent with the best position in the case.

There are other conditions which also has to be true, e.g. the interception point has to be inside the field. If more than one case fit the conditions the cases are compared on how they advance the position of the ball. The case which is most promising in this sense is chosen.

Results

In the first experiments with our implementation we didn't differentiate between forward and backward passes. In these games we recognized a high increase in passing and in the pass success rate as well. The agent tended to perform much more backward passes than before, because they was nearly almost possible and safe. Unfortunately the score was much worse in these games. We adapted the selection of backward passes in the way that backward passes are only performed if the unmodified team would perform them. Thus the first experiments show that increasing the pass success rate does not necessary increase the overall task performance – scoring. Instead improving the coordination has to be performed in the context of increasing the overall task performance.

For the evaluation of the current implementation for every strategy 30 games has been played against the unmodified team. The table shows the results of these games:

strategy	wins:ties:losses	goal diff.
current	7:17:6	+2
fastest	4:22:4	+2
inertia-best	3:20:7	-5
inertia-same-speed	4:18:8	-3
maintain-best	3:19:8	-5
maintain-same-speed	4:20:6	-2

The results for the strategies current and fastest show that the case base consist of enough cases to mimic the pass behavior of the unmodified agents as well as it shows that the implementation of the case selection out of the case base is right. Unfortunately the results for the other strategies does not show an improvement in the score. Therefore we looked at the pass success rate and compared it to the pass success rate of the unmodified team and discovered a **small improvement of the pass success rate**.

Discussion of the Results Based on the disappointing results we looked into what went wrong in the case-based approach. The first problem was that we only tried to enhance the coordination among passer and receiver by adapting the behavior of the passer. For the receiver we assumed a fixed perfect intercepting model as we did for the opponents as well. Enhancing the coordination from only one side may work but assuming a perfect intercept model for the receiver is too restrictive. In fact we realized during the experiments that the ball intercepting behavior of the *AT Humboldt* team is anything but near optimality. The trainer should always generate cases for the team which are conform with its capabilities. The results of the strategies current and fastest which assume no perfect intercept capabilities of the receiver demonstrate this. So it seems promising to adapt the theoretical models to assume proper ball intercepting capabilities.

Related Work

In (Goldberg & Matarić 1997) the authors use interference – the amount of time which robots spend avoiding each other – as a measure of coordination in foraging tasks. Interference can be viewed as a measure of coordination failures, in contrast with our work, which attempts to measure coordination success.

The next two papers we discuss are earlier attempts at building automated trainers for RoboCup teams, but both approaches rely on human programmer to interpret the results. The authors of (Driessens *et al.* 1999) presents an approach to verify and validate the behavior of agents based on inductive reasoning. With predictive induction they generate classification trees and with descriptive induction they generate theories for soccer playing agents. The derived trees and theories are then validated by a human programmer if they differ from his expectations. With this approach the human observer is supported by a tool with whom he can judge if its expectations of coordination among teammates are fulfilled. Unfortunately he gets no support for his decision if the coordination can be improved nor he gets help in improving the coordination. Furthermore the examples presented in the paper are not discussed in case of coordination issues.

Another tool that helps humans to analyze, evaluate and understand agent and even team behavior is presented in (Raines, Tambe, & Marsella 2000). An automated analyst agent called ISAAC analyzes soccer games off-line after its end using data from the agents observable behavior traces. An impressive wide range of behaviors of the individual agent, of agent interactions and of team success or failure are analyzed. However compared to our approach the agent interactions are only sketchily analyzed. Data traces are matched against generic interaction pattern only to figure out the success or failure of the interaction behavior. This information is statistically processed and presented to the human observer. The agent interactions are not investigated for the reasons of success or failure nor the coordination among the agents is evaluated. As the other discussed related work the authors do not approach to use the analysis to give concrete hints for improvements nor they approach to automatically improve the team behavior. The paper by Raines et. al. and the paper by Driessens et. al. focus on using machine learning for their investigations to detect soccer-specific behavior, while we are attempting to rely on analytical methods of coordination to improve the agents' performance.

Conclusion and Future Work

After we motivated that coordination is critical for the success of a team and that improving coordination is difficult to perform manually, we suggested to use an autonomous trainer to support the improvement process. Based on coordination theory for transferring a resource we proposed theoretical models to improve the coordination between team members. We did evaluate these models empirically, and found that their predictions are correlated with higher pass-success rates among 17 soccer teams that are not of our own design. Because the computation of the best transferring coordination is very time consuming the team trainer computes for the proposed models the optimal solutions for given transfer situations. The situations and its solutions are collected in a case base which is used by the agents to improve their coordination. The approach has been applied to soccer playing agents for the coordination problem passing. With the improved passing the coordination success (as measured by the pass success rate) improved but the overall task performance was not improved.

An alternation we started to experiment with, was to use the cases also for the positioning of the possible receivers. This means the possible receivers are attracted by positions which would lead to good passes.

References

Burkhard, H.-D.; Bach, J.; Schröter, K.; Wendler, J.; Gollin, M.; Meinert, T.; and Sander, G. 2001. AT Humboldt 2000 (Team Description). In Stone, P.; Balch, T.; and Kraetzschmar, G., eds., *RoboCup-2000: Robot Soccer World Cup IV*, volume 2019 of *LNAI*, 405–408. Berlin, Germany: Springer Verlag.

Burkhard, H.-D.; Hannebauer, M.; and Wendler, J. 1998. AT Humboldt - Development, Practice and Theory. In Kitano, H., ed., *RoboCup-97: Robot Soccer World Cup I*, volume 1395 of *LNAI*, 357–372. Springer Verlag.

Driessens, K.; Jacobs, N.; Cossement, N.; Monsieurs, P.; and Raedt, L. D. 1999. Inductive Verification and Validation of the KULRoT RoboCup Team. In Asada, M., and Kitano, H., eds., *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *LNAI*, 193–206. Springer Verlag.

Goldberg, D., and Matarić, M. J. 1997. Interference as a Tool for Designing and Evaluating Multi-Robot Controllers. In *AAAI-97 Proceedings*.

Kaminka, G. A. 1998. The multi-agent systems evaluation repository. http://www.cs.cmu.edu/~galk/Eval/.

Kitano, H.; Tambe, M.; Stone, P.; Veloso, M.; Coradeschi, S.; Osawa, E.; Matsubara, H.; Noda, I.; and Asada, M. 1997. The robocup synthetic agent challenge 97. In Pollack, M. E., ed., *Proceedings of the IJCAI-97*, 24–29. Morgan Kaufmann.

Malone, T. W., and Crowston, K. 1994. The Interdisciplinary Study of Coordination. *ACM Computing Surveys* 26 (1):87–119.

Noda, I. 1995. Soccer Server: a simulator of RoboCup. In *Proceedings of AI symposium'95*, 29–34. Japanese Society

for Artificial Intelligence.
see http://ci.etl.go.jp/˜noda/soccer/server/index.html.

Raines, T.; Tambe, M.; and Marsella, S. 2000. Automated Assistants to Aid Humans in Understanding Team Behaviors. In Veloso, M.; Pagello, E.; and Kitano, H., eds., *RoboCup-99: Robot Soccer World Cup III*, volume 1856 of *LNAI*, 85–102. Berlin, Germany: Springer Verlag.

Stone, P.; Veloso, M.; and Riley, P. 1999. The CMUnited-98 Champion Simulator Team. In Asada, M., and Kitano, H., eds., *RoboCup-98: Robot Soccer World Cup II*, volume 1604 of *LNAI*, 61–76. Springer Verlag.