CMLoki and CMOwl: Team and Coach Description for Qualification

Patrick Riley, Yang Gu, John Davin, and Manuela Veloso

Carnegie Mellon University, Computer Science Department, Pittsburgh, PA 15232
pfr@cs.cmu.edu, guyang@cs.cmu.edu, jdavin@cmu.edu, mmv@cs.cmu.edu

CMLoki is the latest simulation team instantiation from Carnegie Mellon. CMOwl continues as our entrant to the coach competition.

CMLoki is built on TsinghuAeolus 2002 and 2003. We hope to explore three major themes with this team: opponent modeling, communication, and coachability.

Opponent modeling has been a part of our simulation teams for many years [1–4]. Most recently we have experimented with the strategy that a single attacker uses to take on a goalie. We took an attacker from our team against three different goalies (TsinghuAeolus 2002, Everest 2002, and FCPortugal 2002). The variation of two parameters controlling how the attacker tries to dodge the goalie greatly affect the scoring performance. Our player can use feedback from past experience to choose a strategy with good performance. We hope to extend this type of modeling and adaptation to multi-agent strategies such as the coordinate strategy of multiple defenders.

We recently proposed a communication standard for player agents which has been accepted as a requirement for the coachable agents for RoboCup2004 [5]. Using this standard, we have run a series of experiments demonstrating the positive effect of communication on the agents' world models [6]. These experiments dealt with fixed strategies (based on the world state) for deciding when and what to communicate. We would like to further explore the range of possible strategies and evaluate their effect on agent world knowledge and performance. Ideally, the communication strategy would adapt to the current game conditions based on analysis of communication performance during the game.

Coachability is the third major theme of our team. The creation of a coach competition in 2001 has led to many interesting research problems, including how to incorporate advice into an agent's behavior. We are interested in creating coachable agents both for our participation in the team competition and as players for the coach competition. Participation in the team competition will allow us to explore coaching and coachability in ways not allowed in the coach competition.

We have previously worked on architectures for incorporating advice in both the ChaMeleons [7] and Wyverns teams. Fairly strict criteria for following advice have been used, leaving open the broader question of how advice can be most effectively used to improve performance.

While our design of architectures and algorithms for incorporating advice provides a good starting point, more remains to be done. In the past we have used estimates of the probability of success of actions as criteria for filtering and selecting between actions advice by the coach. More robust and varied criteria based on the state of the world and past performance provide interesting avenues for exploration.

Coaching is an active area of research for us and we have been involved with the initial design and further evolution of the coach competition. We consider coaching in a larger context than robot soccer [8, 9]. Our past work in coaching has consisted largely of creating models of various parts of the performance of teams in order to predict performance or to imitate the team [4, 3]. With these models we have shown improved performance for teams using our coach.

Recently we have examined the construction and use of an abstract MDP to describe the game [10]. An MDP can be constructed to capture agent behavior and improve performance of a team. Currently, the MDP is not adapted to the team being coached. We are currently exploring how the model provided by the MDP can be used to identify the particular abilities of the agents and opponents and adapt the advice accordingly. Ideally, this can be done both by training before the match and while a game is taking place.

CMLoki and CMOwl continue our long research tradition in opponent modeling and coordinated strategy in the simulated robot soccer framework.

References

1. Riley, P., Veloso, M.: On behavior classification in adversarial environments. In Parker, L.E., Bekey, G., Barhen, J., eds.: Distributed Autonomous Robotic Systems 4. Springer-Verlag (2000) 371–380
2. Stone, P., Riley, P., Veloso, M.: Defining and using ideal teammate and opponent models. In: IAAI-2000. (2000) 1040–1045
3. Riley, P., Veloso, M., Kaminka, G.: An empirical study of coaching. In Asama, H., Arai, T., Fukuda, T., Hasegawa, T., eds.: Distributed Autonomous Robotic Systems 5. Springer-Verlag (2002) 215–224
4. Riley, P., Veloso, M.: Planning for distributed execution through use of probabilistic opponent models. In: AIPS-2002. (2002) 72–81 *Best Paper Award.*
5. Davin, J., Riley, P., Gu, Y.: Coachable agents communication standard. `http://www.cs.cmu.edu/~robosoccer/simulator/comm_standard.html` (2004)
6. Davin, J., Riley, P., Veloso, M.: Commlang: Communication for coachable agents. (under review for the RoboCup Symposium 2004) (2004)
7. Carpenter, P., Riley, P., Veloso, M., Kaminka, G.: Integration of advice in an action-selection architecture. In Kaminka, G.A., Lima, P.U., Rojas, R., eds.: RoboCup-2002: Robot Soccer World Cup VI. Number 2752 in Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin (2003) 195–205
8. Riley, P., Veloso, M., Kaminka, G.: Towards any-team coaching in adversarial domains. In: AAMAS-02. (2002) 1145–1146
9. Riley, P., Veloso, M.: Coaching advice and adaptation. In Polani, D., Bonarini, A., Browning, B., Yoshida, K., eds.: RoboCup-2003: The Sixth RoboCup Competitions and Conferences. Springer Verlag, Berlin (2004) (to appear).
10. Riley, P., Veloso, M.: Advice generation from observed execution: Abstract markov decision process learning. (under review for AAAI-04) (2004)