ATTUnited-2001: Using Heterogeneous Players

Peter Stone

AT&T Labs — Research
180 Park Ave.,Florham Park, NJ 07932, USA
pstone@research.att.com
http://www.research.att.com/ pstone

1 Introduction

ATTUnited-2001 is a successor of the CMUnited teams: CMUnited-97, CMUnited-98, CMUnited-99, and ATT-CMUnited-2000. It is built mainly upon CMUnited-99 [4]. It also incorporates the team action architecture from ATT-CMUnited-2000 [3], but not the dynamic planning of set-plays capability [2]. The main research focus of ATTUnited-2001 is the affect of heterogeneous players in the RoboCup simulator [1].

Heterogeneous players were introduced to the RoboCup simulator for the first time this year, in version 7.0 of the simulator [1]. In particular, in any given game, each team is able to select from identical pools of players, including the default player type from years past and six randomly generated players. At startup, teams are configured with all default players. However, the autonomous on-line coach may substitute in the randomly generated players for any player other than the goalie. The only restriction is that each random player type can be assigned to at most 3 teammates.

The random players are generated by the simulator at startup time by adjusting five parameters, each representing a trade-off in player abilities:

1. Maximum speed vs. stamina recovery
2. Speed vs. turning ability
3. Acceleration vs. size
4. Leg length vs. kick accuracy
5. Stamina vs. maximum acceleration

For the remainder of this paper, the parameters will be referred to by their associated numbers above.

At the outset, it was not known whether using heterogeneous players would be advantageous. The research reported here was designed to generate a prediction of the advantage of including a player with a specific player type on a team, and as a secondary goal, characterizing the utility of playing a given player in a given position (location on the field).

2 Setting the Stage

To experiment with heterogeneous players, the ATTUnited-2001 players had first to be made aware of their player type and made to alter their play according to their type. The players' decisions that are affected are include kicking skills, stamina management, and prediction of player locations, among others [4]. As all of these decisions were already parameterized based on the players' physical ability, this change was a straightforward modification to the existing code.

3 Establishing a Potential Effect

3.1 Isolating the Parameters

The initial question to be examined was whether there was any potential whatsoever for heterogeneous players to improve team performance. In order to examine this question, we endeavored first to determine the parameters' effects in isolation. That is, we matched a team with non-default players (the "tested" team) against a team comprised of all default players. The players on the tested team were able to adjust their skills to their player type, but were otherwise identical to the players on the opposing team.

All the players on the team being tested were created with one of the 5 player parameters set to its maximum value. The rest of the parameters, as well as those for the opponent players, were kept at the default values. We assume that the effect of each parameter is linear within the range of possible values, thus requiring only that we test the endpoints of each range. Since keeping the tested parameter at its minimum value corresponds to having all default players, there is no expected advantage (or disadvantage) in that case.

The following results were observed. In all cases, the parameter number being varied is listed first, followed by the tested team's average score and then that of the default team. The number of games run is listed in parentheses.

1. 0 – 4.3 (19 games)
2. 2.1 – .5 (10)
3. 1.7 – .8 (12)
4. 2.3 – .8 (16)
5. .2 – 6.8 (5)

These results indicate that, considered independently and assuming monotonic functions of parameters to advantage, players are better off with parameters 1 and 5 set to their default (low) values and parameters 2–4 set to their maximum (high) values.

3.2 Bounding the Advantage

The experiments reported in Section 3.1 suggest that there is a potential for heterogeneous players, when selected optimally from the within the parameter space, to improve a team's performance. To bound the potential advantage, we tested a team comprised of players with their most beneficial parameter settings (as indicated in Section 3.1) against a default team.

We ran the test under two conditions. First, the tested team's players were fully aware of their player types and reacted accordingly (knowledge). Second, while the players were still given the advantageous capabilities, they acted as if they were default players (no knowledge). As before, the tested team's score is shown first and the number of games run is indicated in parentheses.

knowledge: 2.9 – .35 (17)
no knowledge: .9 – 1 (11)

These results suggest that heterogeneous players can improve a team's performance by *at most* 3 goals per game. However, the advantage is only observed if they are aware of their types and adjust their skills accordingly.

4 Distributing the Players

The bound discovered in Section 3.2 was obtained assuming that *all* players on the team are given the *optimal* player parameters. But in practice, some player types are better than others, and only 3 players are allowed to assume any given type. In this section we address the question of whether it would be more advantageous to have the strongest players on defense or on offense.

Since the team's default formation uses 4 defenders and 6 midfielder/attackers ("attackers"), we began by testing a team composed of 4 "optimal" defenders (i.e. players with the best parameters from Section 3.1) and 6 default attackers against a team composed of 4 default defenders and 6 optimal attackers[1]. In this case, the first team out-scored the second by an average of 1.4 – .9 over 19 games, suggesting that it it is preferable to put the strongest players on defense.

To verify this result, we tested a team with optimal defenders and default attackers, and then, conversely, a team with default defenders and optimal midfielder/attackers, both against an all-default team. In these experiments, we observed the following results.

Optimal defenders: 2.2 – .6 (32 games)
Good attackers: 3.3 – 1.6 (12)

In this case, although the goal difference per game was roughly equivalent in the two cases, the ratio of goals scored to goals suffered is much more favorable for the team with strong defenders. Therefore, we conclude that, all else being equal, the strongest players should be placed in defensive positions.

5 Establishing a Real Effect

In a real game, the 6 player types from which to select are generated randomly. To this point, it was still not clear whether selecting from among these player types could give a team any expected advantage over using all default players.

In order to establish a real effect, we assumed that the 5 player type parameters were linearly independent, and weighted them based on the goal differences in the experiments reported in Section 3.1. Doing so led to the weights for parameters (1,2,3,4,5) of (-4.5,1.5,1.0,1.5,-5.0).

[1] For the purpose of the experiments reported in this paper, the players are not allowed to exchange positions with teammates.

Multiplying the weight vector by the vector of percentage values of the parameters (i.e. $\frac{val - val_{min}}{(val_{max} - val_{min}}$), we then assigned the best player types to the players 3 at a time, starting with the defenders. The resulting heterogeneous team out-scored the default team by an average of 2.2 – .8 over 10 games, and in 9 of the 10 games, the heterogeneous team won.

We therefore conclude that using heterogeneous players can improve a team's performance by *at least* 1.4 goals per game.

6 Weighting the Parameters

Finally, to more precisely weight the parameters, we conducted a series of experiments in which all players on the test team were set to the same random type and played against the default team. Still assuming a linear model and independence among the parameters, we then found a least-squares fit from the parameters to the average score difference. Doing so led to a parameter weight vector of (-2.8, 4.2, 0.8, 1.4, -1.7) with a bias term of -0.6.

7 The Competition

In RoboCup-2001, ATTUnited-2001 used the weight vector determined in Section 6 to rank the player types, and assigned them 3 players at a time starting with the best player types as defenders (provided the predicted score difference for the player type is positive.). Overall, the team won 6 games, tied 2, and lost 1, failing to qualify for the quarterfinals based only on an unfavorable goal difference. In the coach competition, the coach was able to improve the performance of another team by 2 goals (against a constant opponent) simply by assigning the player types as indicated above.

Acknowledgements

ATTUnited-2001 was based closely on CMUnited-99, which was co-developed with Patrick Riley and Manuela Veloso, as well as parts of ATT-CMUnited-2000 co-developed with David McAllester. Thanks to Sanjoy Dasgupta for consultation regarding the best fit of the parameters to the data in Section 6.

References

1. M. Chen, E. Foroughi, F. Heintz, S. Kapetanakis, K. Kostiadis, J. Kummeneje, I. Noda, O. Obst, P. Riley, T. Steffens, Y. Wang, and X. Yin. Users manual: Robocup soccer server manual for soccer server version 7.07 and later. Available at http://sourceforge.net/projects/sserver/.
2. P. Riley and M. Veloso. Adaptive team coaching using opponent model selection. In *Fifth International Conference on Autonomous Agents*, 2001.
3. P. Stone and D. McAllester. An architecture for action selection in robotic soccer. In *Proceedings of the Fifth International Conference on Autonomous Agents*, 2001.
4. P. Stone, P. Riley, and M. Veloso. The CMUnited-99 champion simulator team. In M. Veloso, E. Pagello, and H. Kitano, editors, *RoboCup-99: Robot Soccer World Cup III*, pages 35–48. Springer Verlag, Berlin, 2000.