

UT Austin Villa: RoboCup 2015 3D Simulation League Competition and Technical Challenges Champions

Patrick MacAlpine, Josiah Hanna, Jason Liang, and Peter Stone

Department of Computer Science, The University of Texas at Austin, USA
{patmac, jphanna, jliang, pstone}@cs.utexas.edu

Abstract. The UT Austin Villa team, from the University of Texas at Austin, won the 2015 RoboCup 3D Simulation League, winning all 19 games that the team played. During the course of the competition the team scored 87 goals and conceded only 1. Additionally the team won the RoboCup 3D Simulation League technical challenge by winning each of a series of three league challenges: drop-in player, kick accuracy, and free challenge. This paper describes the changes and improvements made to the team between 2014 and 2015 that allowed it to win both the main competition and each of the league technical challenges.

1 Introduction

UT Austin Villa won the 2015 RoboCup 3D Simulation League for the fourth time in the past five years, having also won the competition in 2011 [1], 2012 [2], and 2014 [3] while finishing second in 2013. During the course of the competition the team scored 87 goals and only conceded 1 along the way to winning all 19 games the team played. Many of the components of the 2015 UT Austin Villa agent were reused from the team's successful previous years' entries in the competition. This paper is not an attempt at a complete description of the 2015 UT Austin Villa agent, the base foundation of which is the team's 2011 championship agent fully described in a team technical report [4], but instead focuses on changes made in 2015 that helped the team repeat as champions.

In addition to winning the main RoboCup 3D Simulation League competition, UT Austin Villa also won the RoboCup 3D Simulation League technical challenge by winning each of the three league challenges: drop-in player, kick accuracy, and free challenge. This paper also serves to document these challenges and the approaches used by UT Austin Villa when competing in the challenges.

The remainder of the paper is organized as follows. In Section 2 a description of the 3D simulation domain is given. Section 3 details changes and improvements to the 2015 UT Austin Villa team including those for variable distance kicks, set plays, and a kick decision classifier for deciding when to kick or dribble the ball, while Section 4 analyzes the contributions of these changes in addition to the overall performance of the team at the competition. Section 5 describes and analyzes the league challenges that were used to determine the winner of the technical challenge, and Section 6 concludes.

2 Domain Description

The RoboCup 3D simulation environment is based on SimSpark,¹ a generic physical multiagent system simulator. SimSpark uses the Open Dynamics Engine² (ODE) library for its realistic simulation of rigid body dynamics with collision detection and friction. ODE also provides support for the modeling of advanced motorized hinge joints used in the humanoid agents.

Games consist of 11 versus 11 agents playing on a 30 meters in length by 20 meters in width field. The robot agents in the simulation are modeled after the Aldebaran Nao robot,³ which has a height of about 57 cm, and a mass of 4.5 kg. Each robot has 22 degrees of freedom: six in each leg, four in each arm, and two in the neck. In order to monitor and control its hinge joints, an agent is equipped with joint perceptors and effectors. Joint perceptors provide the agent with noise-free angular measurements every simulation cycle (20 ms), while joint effectors allow the agent to specify the speed and direction in which to move a joint.

Visual information about the environment is given to an agent every third simulation cycle (60 ms) through noisy measurements of the distance and angle to objects within a restricted vision cone (120°). Agents are also outfitted with noisy accelerometer and gyroscope perceptors, as well as force resistance perceptors on the sole of each foot. Additionally, agents can communicate with each other every other simulation cycle (40 ms) by sending 20 byte messages.

In addition to the standard Nao robot model, four additional variations of the standard model, known as heterogeneous types, are available for use. These variations from the standard model include changes in leg and arm length, hip width, and also the addition of toes to the robot's foot. Teams must use at least three different robot types, no more than seven agents of any one robot type, and no more than nine agents of any two robot types.

The 2015 RoboCup 3D Simulation League competition included a couple key changes from the previous year's competition. The first of these was a rule change requiring that the ball must either touch an opponent, or touch a teammate outside the center circle, before a team taking a kickoff can score. This rule was put in place to prevent teams from attempting to score directly off a kickoff [3] so as to prevent the competition from potentially devolving into a kickoff taking contest. The second change was to add noise to the beam command used by agents to place themselves at specific positions on the field before a kickoff occurs. Adding noise to the beam command requires agents to use perception when kicking the ball—the added noise prevents agents from beaming to an exact position behind the ball and then blindly executing a kick.

¹ <http://simspark.sourceforge.net/>

² <http://www.ode.org/>

³ <http://www.aldebaran-robotics.com/eng/>

3 Changes for 2015

While many components contributed to the success of the UT Austin Villa team, including dynamic role assignment [5] and an optimization framework used to learn low level behaviors for getting up, walking, and kicking via an overlapping layered learning approach [6], the following subsections focus only on those that are new for 2015. Analysis of the performance of these components is provided in Section 4.1.

3.1 Variable Distance Kicks

In 2014 the UT Austin Villa team only had four kicks for its robots to choose from: a fast short kick that goes about 5 meters, both a low and a high kick that travel around 15 meters each, and a long kick that can propel the ball up to 20 meters. This coarse granularity in kick distances limits the set of locations that the ball can be kicked to. Ideally we would like our robots to be able to kick the ball to any precise position on the field such as professional human soccer players are capable of doing. Having the ability to kick the ball to any location opens up possibilities for better passing and teamwork.

For the 2015 competition the UT Austin Villa team added a set of 13 new kicks to its agents with each of the kicks optimized to travel a fixed distance of 3 to 15 meters in 1 meter increments. This series of new variable distance kicks allow a robot to kick the ball within half a meter of any target 2.5 to 15.5 meters away. The kicks, represented as a series of parameterized joint angle poses [7], were optimized using the CMA-ES algorithm [8] and the team’s optimization framework incorporating overlapping layered learning [6]. During learning of a d meter kick the robot attempts to kick the ball to a target position d meters directly in front of the robot, and a kick attempt is awarded a negative fitness value equal to the euclidean distance of the ball relative to the target position. Each kick was optimized for 400 generations of CMA-ES with a population size of 150. After optimization of each kick the top 300 highest fitness kick parameter sets were evaluated again over 300 kick attempts each to check for consistency. Finally, a parameter set with both high accuracy and low variance for the target distance was identified from collected data and chosen as the kick to use. This learning process was performed for each kick distance, and run across all five heterogeneous agent types, resulting in a total of $13 \times 5 = 65$ kicks learned.

Variable distance kicks allow for a richer set of passing options as robots can select from many potential targets to kick the ball to as shown in Figure 1. Each potential kick location is given a score according to Equation 1, and the location with the highest score is chosen as the location to kick the ball to. Equation 1 rewards kicks for moving the ball toward the opponent’s goal, penalizes kicks that have the ball end up near opponents, and also rewards kicks for landing near a teammate. All distances in Equation 1 are measured in meters.

$$\begin{aligned} \text{score}(target) = & -\|opponentGoal - target\| \\ & \forall opp \in Opponents, -\max(25 - \|opp - target\|^2, 0) \\ & + \max(10 - \|closestTeammateToTarget - target\|, 0) \end{aligned} \quad (1)$$




Fig. 1: Potential kick target locations with lighter circles having a higher score. The highest score location is highlighted in red.

Having many available targets to kick the ball to for passing was very important for implementing a keepaway task for the free challenge discussed in Section 5.3. Having accurate variable distance kicks was also imperative for doing well during the kick accuracy challenge described in Section 5.2.

In addition to allowing for more precise passing, variable distance kicks are also useful for taking shots on goal. Generally speaking, the greater the distance a kick travels the longer and higher the ball may travel in the air, and possibly over the goal, when shooting. To prevent accidentally shooting the ball over the goal, which happened quite frequently during the 2014 RoboCup competition, we limit kicks for shooting on goal to be no more than 7 meters in distance beyond the goal line. Having the ability to kick the ball with just the right amount of power such that the ball flies into the goal—but not over it—is a valuable skill during games.

3.2 Set Plays

During the 2014 RoboCup competition the UT Austin Villa team used a multi-robot behavior to score goals immediately off an indirect kickoff. This behavior consisted of having one robot lightly touch the ball before a second robot kicked the ball into the opponent’s goal [3]. As rules were changed for the 2015 competition, and now a teammate is required to touch the ball outside of the center circle before a goal can be scored, this kickoff tactic is no longer allowed. Instead the team created legal set plays for kickoffs to try and quickly score.

The first kickoff set play, shown in the left image of Figure 2, has the player taking the kickoff kick the ball slightly forward and to the left or right side of the field to a waiting teammate ready to run forward and take a shot. The player taking the kickoff chooses which side target to kick the ball to based on which target is furthest from any opponent. If there are opponents near both side targets then the player taking the kickoff instead chooses the kickoff set play shown in the right image of Figure 2. In this set play the ball is first kicked backwards and to the side to a waiting teammate. The player who receives this

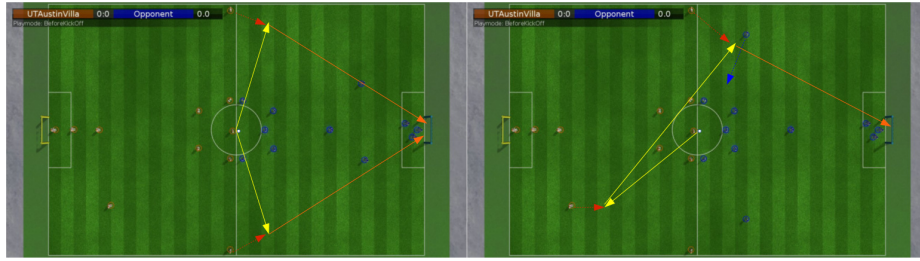


Fig.2: Kickoff set play to the sides (left image) and pass backwards (right image). Yellow lines represent passes and orange lines represent shots. Dashed lines represent agent movement (red for teammates and blue for opponents).

backwards pass then kicks the ball forward and across to the other side of the field where a teammate is waiting for a pass. It is expected that the player who receives the second pass will be in a good position to take a shot on goal as opponent agents will have been drawn to the other side of the field after the initial backwards pass off the kickoff.

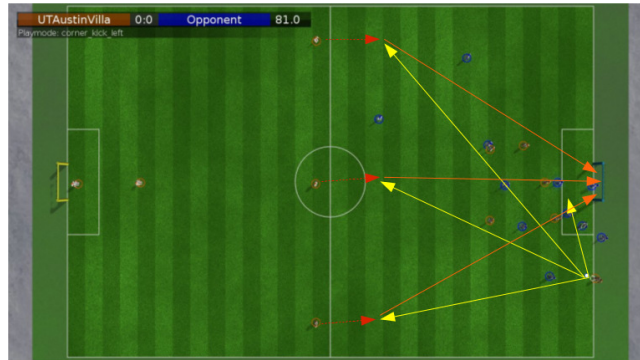


Fig.3: Corner kick set plays. Yellow lines represent passes and orange lines represent shots. Dashed red lines represent teammate movement. In the example shown the ball would be passed to the teammate waiting for the ball near the bottom of the image as that teammate is most open.

In addition to kickoff set plays the UT Austin Villa team also created set plays for offensive corner kicks. These set plays, shown in Figure 3, consist of having three teammates move to positions on the midline at the center and both sides of the field. The player taking the corner kick chooses to kick the ball to whichever of these three players is most open. If none of these players are open then the player taking the corner kick just chooses the default option of kicking the ball to a position in front of the goal where several teammates are waiting.

All set plays require passing the ball to specific locations on the field though the use of learned variable distance kicks discussed in Section 3.1. Approaching and kicking the ball must be quick as a team has only 15 seconds to kick the ball once a set play starts.

3.3 Kick Decision Classifier

Before deciding where to kick the ball, first a decision must be made as to whether to kick or dribble the ball. The 2014 UT Austin Villa team chose to always dribble if an opponent is within two meters of the ball—it was assumed that an agent might not have enough time to complete a kick if an opponent is less than two meters from the ball.

Rather than using a hand-picked value to determine if there is enough time to kick the ball, the 2015 UT Austin Villa team decided to train a logistic regression classifier to predict the probability of a kick being successful given the current state of the world. To do so, the team played many games against a common opponent in which agents were instructed to always try and kick the ball. During the course of kick attempts the following state features were recorded and then labeled as positive or negative kick examples based on whether kick attempts were successful.

1. Difference between angle of ball and the orientation of agent
2. Difference between angle of kick target and orientation of agent
3. Angle difference between closest opponent to ball (OPP*) and ball from agent's point of view
4. Difference between angle of ball (from OPP*'s point of view) and the orientation of OPP*
5. Is OPP* fallen or not
6. Magnitude of OPP* velocity
7. Angle between OPP* velocity and ball velocity
8. Distance from agent to ball / OPP* distance to ball
9. Distance from agent to ball / OPP* distance to agent
10. OPP* distance to ball / OPP* distance to agent
11. Distance from agent to ball - OPP* distance to ball
12. Distance from agent to ball - OPP* distance to agent
13. OPP* distance to ball - OPP* distance to agent
- 14-24. Same features as 3-13 except OPP* is the second closest opponent to ball
- 25-35. Same features as 3-13 except OPP* is the third closest opponent to ball

The output from a trained classifier is a probability of a kick attempt being successful. A threshold value for this probability, for which kicks are attempted when the probability of a successful kick exceeds this value, was chosen after experimenting with different threshold values while playing 100s of games against multiple opponents. Metrics monitored during these games were average goal differential, number of kicks performed, goals against, and the probability of a tie or loss.

During the competition the UT Austin Villa team used two different kick classifier models. One was trained against the Apollo3D team which was thought to be one of the fastest teams, and the other was trained against the BahiaRT team which was the opponent UT Austin Villa had the most trouble scoring against at the 2014 competition [3]. Ultimately it was decided to use the more conservative model trained against Apollo3D whenever the team was on defense, so as to be less likely to lose the ball on defense, and then half the time (playing as the left team) switch to the BahiaRT trained model that chose to kick more frequently when on offense and in range of being able to take a shot on goal.

4 Main Competition Results and Analysis

In winning the 2015 RoboCup competition UT Austin Villa finished with a perfect record of 19 wins and no losses.⁴ During the competition the team scored 87

⁴ Full tournament results can be found at http://wiki.robocup.org/wiki/Soccer_Simulation_League/RoboCup2015#3D

goals while only conceding 1. Despite finishing with a perfect record, the relatively few number of games played at the competition, coupled with the complex and stochastic environment of the RoboCup 3D simulator, make it difficult to determine UT Austin Villa being better than other teams by a statistically significant margin. At the end of the competition, however, all teams were required to release their binaries used during the competition. Results of UT Austin Villa playing 1000 games against each of the other 11 teams’ released binaries from the competition are shown in Table 1.

Table 1: UT Austin Villa’s released binary’s performance when playing 1000 games against the released binaries of all other teams at RoboCup 2015. This includes place (the rank a team achieved at the competition), average goal difference (values in parentheses are the standard error), win-loss-tie record, and goals for/against.

Opponent	Place	Avg. Goal Diff.	Record (W-L-T)	Goals (F/A)
FUT-K	2	2.082 (0.036)	927-2-71	2178/96
FCPortugal	3	2.399 (0.040)	945-4-51	2624/225
BahiaRT	4	2.496 (0.044)	944-1-55	2501/5
Apollo3D	5	3.803 (0.046)	995-0-5	3805/2
magmaOffenburg	6	4.167 (0.051)	999-0-1	4171/4
RoboCanes	7	4.187 (0.049)	998-0-2	4235/48
Nexus3D	8	5.571 (0.044)	1000-0-0	5573/2
CIT3D	9	6.321 (0.050)	1000-0-0	6321/0
ITAndroids	11	10.125 (0.041)	1000-0-0	10125/0
Miracle3D	12	10.521 (0.056)	1000-0-0	10521/0
HfutEngine3D	10	11.897 (0.068)	1000-0-0	11897/0

UT Austin Villa finished with at least an average goal difference greater than two goals against every opponent. Additionally UT Austin Villa only lost 7 games out of the 11,000 that were played in Table 1 with a win percentage greater than 92% against all teams. This shows that UT Austin Villa winning the 2015 competition was far from a chance occurrence. The following subsection analyzes some of the components described in Section 3 that contributed to the team’s dominant performance.

4.1 Analysis of Components

Table 2: Average goal difference (standard error shown in parentheses) achieved by different versions of the UT Austin Villa team when playing 1000 games against the top four teams at RoboCup 2015.

Opponent	UTAustinVilla	NoVarDistKicks	NoSetPlays	NoKickClassifier
UTAustinVilla	—	-0.042 (0.028)	-1.524 (0.048)	0.060 (0.029)
FUT-K	2.082 (0.036)	1.929 (0.036)	1.467 (0.035)	2.107 (0.036)
FCPortugal	2.399 (0.040)	2.536 (0.044)	2.478 (0.042)	2.181 (0.040)
BahiaRT	2.496 (0.044)	2.407 (0.044)	2.283 (0.040)	2.173 (0.039)

Table 2 shows the average goal difference achieved by the following different versions of the UT Austin Villa team when playing 1000 games against the top four teams at RoboCup 2015.

UTAustinVilla Released binary with all features.

NoVarDistKicks No variable distance kicks (except for those used during set plays).

NoSetPlays No set plays.

NoKickClassifier No kick decision classifier.

Table 3: Scoring percentage of kickoffs and corner kicks achieved by versions of the UT Austin Villa team with and without using set plays while playing 1000 games against the top four teams at RoboCup 2015.

Opponent	Kickoff Scoring %		Corner Kick Scoring %	
	SetPlays	NoSetPlays	SetPlays	NoSetPlays
UTAustinVilla	47.11	0.18	21.32	8.25
FUT-K	54.20	2.65	28.50	12.83
FCPortugal	10.45	6.34	26.03	16.83
BahiaRT	1.25	1.79	33.66	11.79

Using variable distance kicks slightly improves performance against most teams except for when playing against FCPortugal. When watching games against FCPortugal it was noticed that the UT Austin Villa team often makes short passes to players who have opponents running toward them and are no longer open by the time the ball is received. This suggests that Equation 1 in Section 3.1 for scoring kick target locations should be improved to take into account the velocity of opponents. Using opponents’ velocities, and also performing machine learning to train a function for computing the value of a location to kick the ball to, are future work.

Using set plays really improves performance against UTAustinVilla and FUT-K. Set plays are also beneficial when playing against BahiaRT. Table 3 shows the scoring percentages for kickoffs (measured as having scored within 30 seconds of a kickoff) and corner kicks (measured as having scored within 15 seconds of the kick being taken) when both using and not using set plays while playing against the top four teams at RoboCup 2015. Using set plays improves the scoring percentage against all opponents on corner kicks. Kickoffs are not very successful against BahiaRT and FCPortugal with set plays as both teams use formations during kickoffs that are spread out and cover the field well. These spread out formations make it difficult for opponents to find an area with enough free space to receive a pass. Using set plays actually lowers performance against FCPortugal which we attribute to the low kickoff scoring success rate—we find it beneficial to just kick the ball deep into the opponent’s side on kickoffs when playing against teams whose formations interfere with our kickoff set plays.

Using the kick decision classifier improves performance against BahiaRT and FCPortugal. This is not surprising as one of the classifiers used was trained against BahiaRT which was built on top of a version of FCPortugal’s code base. The kick decision classifier does not help against UTAustinVilla and FUT-K, which are likely the two quickest teams in the league, and are both faster at getting to the ball than either of the teams the kick decision models were trained against. Future work remains to train kick decision classifier models for playing against UTAustinVilla and FUT-K in order to verify if doing so can improve performance against these teams.

5 Technical Challenges

For the second straight year there was an overall technical challenge consisting of three different league challenges: drop-in player, kicking accuracy, and free

challenge. For each league challenge a team participated in points were awarded toward the overall technical challenge based on the following equation:

$$\text{points}(\text{rank}) = 25 - 20 * (\text{rank} - 1) / (\text{numberOfParticipants} - 1)$$

Table 4: Overall ranking and points totals for each team participating in the RoboCup 2015 3D Simulation League technical challenge as well as ranks and points awarded for each of the individual league challenges that make up the technical challenge.

Team	Overall		Drop-in Player		Kick Accuracy		Free	
	Rank	Points	Rank	Points	Rank	Points	Rank	Points
UTAustinVilla	1	75	1	25	1	25	1	25
FCPortugal	2	61.44	3-6	19.44	2	21	2	21
BahiaRT	3	41.44	3-6	19.44	6	5	3	17
FUT-K	4	37.22	9	7.22	3	17	4	13
magmaOffenberg	5	32.44	3-6	19.44	4	13	—	—
CIT3D	6	31.78	2	22.78	—	—	5	9
HFutEngine3D	7	19.44	3-6	19.44	—	—	—	—
Apollo3D	8	19	10	5	5	9	6	5
Nexus3D	9-10	10.56	7-8	10.56	—	—	—	—
RoboCanes	9-10	10.56	7-8	10.56	—	—	—	—

Table 4 shows the ranking and cumulative team point totals for the technical challenge as well as for each individual league challenge. UT Austin Villa earned the most points and won the technical challenge by taking first in each of the league challenges. The following subsections detail UT Austin Villa’s participation in each league challenge.

5.1 Drop-in Player Challenge

The drop-in player challenge,⁵ also known as an ad hoc teams challenge, is where agent teams consisting of different players randomly chosen from participants in the competition play against each other. Each participating team contributes two agents to one drop-in player team where drop-in player games are 10 vs 10 with no goalies. An important aspect of the challenge is for an agent to be able to adapt to the behaviors of its teammate. During the challenge agents are scored on their average goal differential across all games played.

Table 5 shows the results of the drop-in player challenge at RoboCup under the heading “At RoboCup 2015”. The challenge was played across 8 games such that every agent played at least one game against every other agent participating in the challenge. UT Austin Villa used the same strategy employed in the 2013 and 2014 drop-in player challenge [9], and in doing so was able to win this year’s drop-in player challenge.

Drop-in player games are inherently very noisy and it is hard to get statistically significant results when only playing 8 games. In order to get a better idea of each agents’ true drop-in player performance we replayed the challenge with released binaries across all $\binom{10}{5} * \binom{5}{5} / 2 = 126$ possible team combinations of drop-in player games ten times each. Results in Table 5 of replaying the

⁵ Details of the drop-in player challenge at http://www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/2015_dropin_challenge/3D_DropInPlayerChallenge.pdf

Table 5: Average goal differences for each team in the drop-in player challenge when playing all possible pairings of drop-in player games ten times (1260 games in total) and at RoboCup 2015.

Team	Avg. Goal Diff.	At RoboCup 2015	
		Rank	Avg. Goal Diff.
UTAustinVilla	1.823 (0.045)	1	1.625
FCPortugal	0.340 (0.067)	3-6	-0.125
BahiaRT	0.182 (0.068)	3-6	-0.125
magmaOffenburg	-0.039 (0.068)	3-6	-0.125
FUT-K	-0.052 (0.068)	9	-0.625
RoboCanes	-0.180 (0.068)	7-8	-0.375
CIT3D	-0.361 (0.067)	2	1.125
HfutEngine3D	-0.501 (0.067)	3-6	-0.125
Apollo3D	-0.593 (0.066)	10	-0.875
Nexus3D	-0.620 (0.066)	7-8	-0.375

competition over many games show that UT Austin Villa has an average goal difference more than five times higher than any other team, thus validating UT Austin Villa winning the drop-in player challenge.

5.2 Kick Accuracy Challenge

For the kick accuracy challenge⁶ robots are asked to kick a ball to the center point of the field from ten different starting ball positions ranging in distances to the center of the field of 3-12 meters in 1 meter increments. Having already optimized variable distance kicks for each of the integer distances in the 3-12 meter range as described in Section 3.1, the UT Austin Villa agent participating in the challenge simply chose to execute the appropriate kick for the distance the ball was from the field center during each kick attempt.

Table 6: Average kick distance error in meters for each of the participating teams in the kick accuracy challenge.

Team	Error
UTAustinVilla	0.200
FCPortugal	0.282
FUT-K	0.963
magmaOffenberg	1.302
Apollo3D	2.714
BahiaRT	4.652

Results of the kick accuracy challenge are shown in Table 6. UTAustinVilla won the challenge by having the lowest average kick distance error due to its very accurate learned kicks. It is worth noting that the FCPortugal team, who also had a very low average error and took second place in the challenge, used a different strategy of learning a single general kicking skill for different distances using contextual policy search [10].

5.3 Free Challenge

During the free challenge, teams give a five minute presentation on a research topic related to their team. Each team in the league then ranks the top five

⁶ Details and framework for the kick accuracy challenge at <https://github.com/magmaOffenburg/magmaChallenge#kick-challenge>

presentations with the best receiving 5 votes and the 5th best receiving 1 vote. Additionally several respected research members of the RoboCup community outside the league vote, with their votes being counted double. The winner of the free challenge is the team that receives the most votes. Table 7 shows the results of the free challenge in which UT Austin Villa got first place.

Table 7: Results of the free challenge.

Team	Votes
UTAustinVilla	79
FCPortugal	67
BahiaRT	49
FUT-K	32
CIT3D	29
Apollo3D	13

UT Austin Villa’s free challenge submission⁷ focused on describing the team’s approach, incorporating machine learning, for kicking and passing the ball to teammates. This included the following topics: how to approach and kick the ball to different targets (learning skills for walking up to and kicking the ball discussed in [6] as well as learning variable distance kicks presented in Section 3.1), where to kick the ball (using a kick location scoring function detailed in Section 3.1), when to kick the ball (by querying the kick decision classifier described in Section 3.3), and how to have teammates move to receive a pass (using kick anticipation explained in [3]).

UT Austin Villa’s free challenge presentation culminated in the demonstration of a keepaway task in which one team attempts to maintain possession of the ball and keep it away from another team for as long as possible. During the demonstration a team was shown to be able to maintain possession and keep the ball away from the 2014 RoboCup champion UT Austin Villa team for over two minutes.⁸

6 Conclusion

UT Austin Villa won the 2015 RoboCup 3D Simulation League main competition as well as all technical league challenges.⁹ Data taken using released binaries from the competition show that UT Austin Villa winning the competition was statistically significant. The 2015 UT Austin Villa team improved dramatically from 2014 as it was able to beat a version of the team’s 2014 champion binary (the NoScoreKO agent in [3] that does not attempt the now illegal behavior of scoring on a kickoff) by an average of 1.838 (+/-0.047) goals across 1000 games.

⁷ Free challenge entry description at <http://www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/AustinVilla3DSimulationFiles/2015/files/UTAustinVillaFreeChallenge.pdf>

⁸ A video of the keepaway task can be found at <http://www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/#2015challenges>

⁹ More information about the UT Austin Villa team, as well as video highlights from the competition, can be found at the team’s website: <http://www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/#2015>

A large factor in UT Austin Villa's success in 2015 was due to improvements in kicking and the coordination of set plays. In order to remain competitive, and challenge for the 2016 RoboCup championship, the team will likely need to improve multiagent team behaviors such as passing. Additionally, as other teams in the league advance their own passing capabilities, UT Austin Villa will look to implement marking strategies to account for opponents' offensive strategies.

Acknowledgments

Thanks to Klaus Dorer for putting together the kicking accuracy challenge. This work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (CNS-1330072, CNS-1305287), ONR (21C184-01), AFRL (FA8750-14-1-0070), and AFOSR (FA9550-14-1-0087).

References

1. MacAlpine, P., Urieli, D., Barrett, S., Kalyanakrishnan, S., Barrera, F., Lopez-Mobilia, A., Știurcă, N., Vu, V., Stone, P.: UT Austin Villa 2011: A champion agent in the RoboCup 3D soccer simulation competition. In: Proc. of 11th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2012). (2012)
2. MacAlpine, P., Collins, N., Lopez-Mobilia, A., Stone, P.: UT Austin Villa: RoboCup 2012 3D simulation league champion. In: RoboCup-2012: Robot Soccer World Cup XVI. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin (2013)
3. MacAlpine, P., Depinet, M., Liang, J., Stone, P.: UT Austin Villa: RoboCup 2014 3D simulation league competition and technical challenge champions. In: RoboCup-2014: Robot Soccer World Cup XVIII. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin (2015)
4. MacAlpine, P., Urieli, D., Barrett, S., Kalyanakrishnan, S., Barrera, F., Lopez-Mobilia, A., Știurcă, N., Vu, V., Stone, P.: UT Austin Villa 2011 3D Simulation Team report. Technical Report AI11-10, The Univ. of Texas at Austin, Dept. of Computer Science, AI Laboratory (2011)
5. MacAlpine, P., Price, E., Stone, P.: SCRAM: Scalable collision-avoiding role assignment with minimal-makespan for formational positioning. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15). (2015)
6. MacAlpine, P., Depinet, M., Stone, P.: UT Austin Villa 2014: RoboCup 3D simulation league champion via overlapping layered learning. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15). (2015)
7. Depinet, M., MacAlpine, P., Stone, P.: Keyframe sampling, optimization, and behavior integration: Towards long-distance kicking in the robocup 3d simulation league. In: RoboCup-2014: Robot Soccer World Cup XVIII. Lecture Notes in Artificial Intelligence. Springer Verlag, Berlin (2015)
8. Hansen, N.: The CMA Evolution Strategy: A Tutorial. (2009) <http://www.lri.fr/~hansen/cmatutorial.pdf>.
9. MacAlpine, P., Genter, K., Barrett, S., Stone, P.: The RoboCup 2013 drop-in player challenges: Experiments in ad hoc teamwork. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS). (2014)
10. Kupcsik, A.G., Deisenroth, M.P., Peters, J., Neumann, G.: Data-efficient generalization of robot skills with contextual policy search. In: AAAI. (2013)