

# UT Austin Villa: RoboCup 2017 3D Simulation League Competition and Technical Challenges Champions

Patrick MacAlpine and Peter Stone

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

**Abstract.** The UT Austin Villa team, from the University of Texas at Austin, won the 2017 RoboCup 3D Simulation League, winning all 23 games that the team played. During the course of the competition the team scored 171 goals without conceding any. Additionally, the team won the RoboCup 3D Simulation League technical challenge by winning each of a series of three league challenges: free, passing and scoring, and Gazebo running challenge. This paper describes the changes and improvements made to the team between 2016 and 2017 that allowed it to win both the main competition and each of the league technical challenges.

## 1 Introduction

UT Austin Villa won the 2017 RoboCup 3D Simulation League for the sixth time in the past seven years, having also won the competition in 2011 [1], 2012 [2], 2014 [3], 2015 [4], and 2016 [5] while finishing second in 2013. During the course of the competition the team scored 171 goals and conceded none along the way to winning all 23 games the team played. Many of the components of the 2017 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 2017 UT Austin Villa agent, the base foundation of which is the team’s 2011 championship agent fully described in a team technical report [6], but instead focuses on changes made in 2017 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: free, passing and scoring, and Gazebo running 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 the most important improvement to the 2017 UT Austin Villa team: fast walk kicks, while Section 4 analyzes the contribution of this improvement 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 [7, 8], a generic physical multiagent system simulator. SimSpark uses the Open Dynamics Engine (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 two 5 minute halves of soccer on a 30 X 20 meter field. The robot agents in the simulation are modeled after the Aldebaran Nao robot, 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/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 ( $120^\circ$ ). 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 main change for the 2017 RoboCup 3D Simulation League competition from previous years was the removal of crowding rules. Previously, too many players crowded around the ball caused players to be penalized and beamed to the sideline. Crowding rules were primarily enforced to decrease the number of collisions between robots as multiple collisions at the same time can slow down the simulator and potentially cause it to crash. With existing touching rules such that a player is beamed to the sideline if a group of three or more players are touching each other, and the addition in 2016 of charging fouls that penalize players for running into opponents, it was determined that crowding rules were no longer needed.

Figure 1 shows a visualization of the Nao robot and the soccer field during a game.



Fig. 1: A screenshot of the Nao humanoid robot (left), and a view of the soccer field during a 11 versus 11 game (right).

### 3 Fast Walk Kicks

Many components developed prior to 2017 contributed to the success of the UT Austin Villa team including dynamic role assignment [9], marking [10], and an optimization framework used to learn low level behaviors for walking and kicking via an overlapping layered learning approach [11]. This section discusses the development of a new and important component for 2017: fast walk kicks. Fast walk kicks refer to the ability of agents to approach the ball and quickly kick it without having to first come to a stop and enter a stable standing position. The amount of time it takes for agents to approach and kick the ball is an important consideration as kick attempts that take longer to perform give opponents a better chance to stop them from being executed.

The UT Austin Villa team specifies kicking motions through a periodic state machine with multiple key frames, where a key frame is a parameterized static pose of fixed joint positions. Figure 2 shows an example series of poses for a kicking motion. While some joint positions are specified by hand, a subset of values for joint positions are optimized using the CMA-ES [12] algorithm and overlapping layered learning [11] methodologies.

Prior to the 2017 competition all kicking motions performed by the UT Austin Villa team first required the agent to come to a stable standing position with both feet on the ground before kicking the ball. The team’s fastest kicks took about 0.5 seconds to execute but only traveled a little over 5 meters. Longer kicks, traveling as far as 20 meters, were slower and could take up to 2 seconds to execute.

The UT Austin Villa team has noticed a couple trends when optimizing parameter values for kicks: policies with more parameters allow for longer kicks, and policies with more parameters allow for kicking motions with shorter durations that are quicker to execute without the robot becoming unstable and falling over.




Fig. 2: Example of a fixed series of poses that make up a kicking motion.

As adding more parameters to a policy increases the space of policies that can be represented, it is not surprising that policies with more parameters have allowed for kicks that can travel farther and be executed faster. However, adding more parameters to a kick can make learning slower and more difficult, and there is likely an upper limit on the number of parameters that can effectively be learned as CMA-ES does not scale well to thousands of parameters [13].

Given a desire to develop a kick with good distance that is very fast to execute, we decided to learn kicking motion parameters for every joint over 12 simulation cycles (24 ms)—such a kicking motion is thus learned over the entire range of possible poses for any kick less than 0.25 seconds in duration. We optimized  $\approx 260$  parameters for this kick across 1000 generation of CMA-ES using a population size of 300—previously we have used a CMA-ES population size of 150 when optimizing kicks consisting of  $\approx 75$  parameters, however we decided to double the size of the population due to the larger number of parameters being optimized. Initial parameter values were seeded with joint angles taken from a subset of poses used by our longest kick: joint angle values across a 12 simulation cycle window of the kick that include when the ball is struck by the foot. During learning we used the following fitness function that rewards the agent for the distance the ball is kicked, encourages accuracy by giving a Gaussian penalty for the difference/offset between the desired and actual angles that the ball travels, and promotes stability via a negative value if the agent falls over during kicking:

$$fitness_{\text{kick}} = \begin{cases} -1 & : \text{Agent Fell} \\ \text{distBallTraveledForward} * e^{-\text{angleOffset}^2/360} & : \text{Otherwise} \end{cases}$$

The resulting kick learned from this optimization takes 0.24 seconds to execute, travels close to 20 meters in distance (nearly the same distance as our longest kicks that can take up to 2 seconds to execute), and provides a substantial increase in the teams performance—a performance analysis of using fast walk kicks is provided in Section 4.1.

As our learned kick takes less than 0.25 seconds to execute, the robot must begin the kicking motion starting from a walking position, and perform a “walk

kick” due to there not being enough time for the robot to first assume a standing position before striking the ball. During walk kicks it is important that a robot has its non-kicking support leg on the ground before initiating a kicking motion as otherwise the robot will likely fall over. When attempting walk kicks, the UT Austin Villa agent will wait until its support leg is on the ground—as determined by a large enough force measured by the force resistance perceptor on the sole of the agent’s support leg’s foot—before beginning a kick. The magmaOffenburg team, who also developed a walk kick for this year’s competition, similarly ensure that a robot’s support leg is on the ground before attempting a kick [14].

## 4 Main Competition Results and Analysis

In winning the 2017 RoboCup competition UT Austin Villa finished with a perfect record of 23 wins and no losses.<sup>1</sup> During the competition the team scored 171 goals while conceding none. 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 twelve 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 2017. This includes place (the rank a team achieved at the 2017 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)
magmaOffenburg	2	3.756 (0.057)	983-0-17	3778/22
FUT-K	3	4.793 (0.056)	995-0-5	4823/30
AIUT3D	5	5.946 (0.054)	1000-0-0	5981/35
BahiaRT	6	6.677 (0.055)	1000-0-0	6677/0
FCPortugal	7	6.753 (0.062)	1000-0-0	6818/65
Nexus3D	11	7.486 (0.035)	1000-0-0	7486/0
KgpKubs	8	7.510 (0.057)	1000-0-0	7510/0
RoboCanes	4	7.801 (0.066)	1000-0-0	7806/5
HfutEngine3D	12	7.952 (0.049)	1000-0-0	7957/5
Miracle3D	10	8.404 (0.056)	1000-0-0	8404/0
ITAndroids	9	11.169 (0.057)	1000-0-0	11169/0
RIC-AASTMT	13	11.466 (0.051)	1000-0-0	11466/0

UT Austin Villa finished with at least an average goal difference greater than 3.75 goals against every opponent. Additionally, UT Austin Villa won all but

<sup>1</sup> Full tournament results can be found at [https://www.robocup2017.org/file/results/0730/0730\\_soccer\\_simulation\\_3d.pdf](https://www.robocup2017.org/file/results/0730/0730_soccer_simulation_3d.pdf)

22 games that ended in ties—no losses—out of the 12,000 that were played in Table 1 with a win percentage greater than 98% against all teams. These results show that UT Austin Villa winning the 2017 competition was far from a chance occurrence. The following subsection analyzes the contribution of fast walk kicks (described in Section 3) to the team’s dominant performance.

#### 4.1 Analysis of Fast Walk Kicks

To analyze the contribution of fast walk kicks (Section 3) to the UT Austin Villa team’s performance, we played 1000 games between a version of the 2017 UT Austin Villa team with fast walk kicks turned off—and no other changes—against each of the RoboCup 2017 teams’ released binaries. Results comparing the performance of the UT Austin Villa team with and without using fast walk kicks are shown in Table 2.

Table 2: Average goal difference achieved by versions of the UT Austin Villa team with and without fast walk kicks, and the gain in average goal difference by using fast walk kicks, when playing 1000 games against all teams at RoboCup 2017.

Opponent	Fast Walk Kicks	No Fast Walk Kicks	Fast Walk Kicks Gain
UTAustinVilla	0 <sup>a</sup>	-0.557	0.557
magmaOffenburg	3.756	2.159	1.597
FUT-K	4.793	2.912	1.881
AIUT3D	5.946	4.824	1.122
BahiaRT	6.677	4.051	2.626
FCPortugal	6.753	4.701	2.052
Nexus3D	7.486	7.002	0.484
KgpKubs	7.510	5.532	1.978
RoboCanes	7.801	6.113	1.688
HfutEngine3D	7.952	6.827	1.125
Miracle3D	8.404	6.830	1.574
ITAndroids	11.169	10.294	0.875
RIC-AASTMT	11.466	10.526	0.940

<sup>a</sup> Games were not played, but assumed to be an average goal difference of 0 in expectation with self play.

Against all opponents the average goal difference was higher when using fast walk kicks, with the gain in average goal difference performance against each opponent averaging 1.423 goals. These results show that fast walk kicks provide a substantial improvement in game performance to the UT Austin Villa team.

#### 4.2 Additional Tournament Competition Analysis

To further analyze the tournament competition, Table 3 shows the average goal difference for each team at RoboCup 2017 when playing 1000 games against all other teams at RoboCup 2017.

Table 3: Average goal difference for each team at RoboCup 2017 (rows) when playing 1000 games against the released binaries of all other teams at RoboCup 2017 (columns). Teams are ordered from most to least dominant in terms of winning (positive goal difference) and losing (negative goal difference).

	UTA	mag	FUT	FCP	AIUT	Rob	Bah	Kgp	Hfut	Mir	Nex	ITA	RIC
UTAustinVilla	—	3.756	4.793	6.753	5.946	7.801	6.677	7.510	7.952	8.404	7.486	11.169	11.466
magmaOffenburg	-3.756	—	0.058	1.698	1.441	2.692	2.593	2.863	3.627	5.340	4.088	4.760	6.210
FUT-K	-4.793	-0.058	—	1.621	1.989	2.077	2.165	2.603	4.151	4.986	5.297	5.313	6.597
FCPortugal	-6.753	-1.698	-1.621	—	0.325	0.695	1.031	1.735	2.388	3.991	2.950	4.287	5.352
AIUT3D	-5.946	-1.441	-1.989	-0.325	—	0.506	0.449	0.738	2.077	2.721	3.223	2.609	5.536
RoboCanes	-7.801	-2.692	-2.077	-0.695	-0.506	—	1.068	1.852	2.279	4.175	2.612	3.309	4.791
BahiaRT	-6.677	-2.593	-2.165	-1.031	-0.449	-1.068	—	0.393	1.005	1.437	1.448	1.645	3.120
KgpKubs	-7.510	-2.863	-2.603	-1.735	-0.738	-1.852	-0.393	—	0.974	1.115	0.638	1.664	3.027
HfutEngine3D	-7.952	-3.627	-4.151	-2.388	-2.077	-2.279	-1.005	-0.974	—	0.436	0.081	0.180	1.426
Miracle3D	-8.404	-5.340	-4.986	-3.991	-2.721	-4.175	-1.437	-1.115	-0.436	—	0.562	0.525	2.693
Nexus3D	-7.486	-4.088	-5.297	-2.950	-3.223	-2.612	-1.448	-0.638	-0.081	-0.562	—	0.179	1.691
ITAndroids	-11.169	-4.760	-5.313	-4.287	-2.609	-3.309	-1.645	-1.664	-0.180	-0.525	-0.179	—	1.889
RIC-AASTMT	-11.466	-6.210	-6.597	-5.352	-5.536	-4.791	-3.120	-3.027	-1.426	-2.693	-1.691	-1.889	—

It is interesting to note that the ordering of teams in terms of winning (positive goal difference) and losing (negative goal difference) is strictly dominant—every opponent that a team wins against also loses to every opponent that defeats that same team. Relative goal difference does not have this same property, however, as a team that does better against one opponent relative to another team does not always do better against a second opponent relative to that same team. UT Austin Villa is dominant in terms of relative goal difference, however, as UT Austin Villa has a higher goal difference against each opponent than all other teams against the same opponent.

## 5 Technical Challenges

For the fourth straight year there was an overall technical challenge consisting of three different league challenges: free, passing and scoring, and Gazebo running 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 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.<sup>2</sup>

<sup>2</sup> Videos of the passing and scoring challenge and the Gazebo running challenge can be found at

<http://www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/#2017challenges>

Table 4: Overall ranking and points totals for each team participating in the RoboCup 2017 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		Free		Passing and Scoring		Gazebo Running	
	Rank	Points	Rank	Points	Rank	Points	Rank	Points
<b>UTAustinVilla</b>	<b>1</b>	<b>75.00</b>	<b>1</b>	<b>25.00</b>	<b>1</b>	<b>25.00</b>	<b>1</b>	<b>25.00</b>
magmaOffenburg	2	43.00	3	11.67	4	13.00	2	18.33
AIUT3D	3	30.33	4	5.00	3	17.00	3-4	8.33
FCPortugal	4	29.33	—	—	2	21.00	3-4	8.33
BahiaRT	5	23.33	2	18.33	6	5.00	—	—
HfutEngine3D	6	9.00	—	—	5	9.00	—	—

## 5.1 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 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 5 shows the results of the free challenge in which UT Austin Villa was awarded first place.

Table 5: Results of the free challenge.

Team	Votes
<b>UTAustinVilla</b>	<b>59</b>
BahiaRT	49
magmaOffenburg	46
AIUT3D	42

UT Austin Villa’s free challenge submission<sup>3</sup> presented the team’s fast walk kicks discussed in Section 3. Additionally, UT Austin Villa’s free challenge submission divulged preliminary work on representing the policy of a kicking motion as a neural network, and using deep learning [15] and the Trust Region Policy Optimization (TRPO) algorithm [16] to learn longer kicks. The BahiaRT team provided details about an optimization framework they created, the magmaOffenburg team talked about a 2D simulator they use for testing the strategy layer of their team, and the AIUT3D team introduced a motion editor for 3D Simulation League agents.<sup>4</sup>

<sup>3</sup> Free challenge entry description available at <http://www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/AustinVilla3DSimulationFiles/2017/files/UTAustinVillaFreeChallenge2017.pdf>

<sup>4</sup> <https://github.com/AIUT3D/aiut3d-motion-editor>



## 5.2 Passing and Scoring Challenge

In the course of the passing and scoring challenge,<sup>5</sup> a group of four agents on one team attempts to pass the ball between themselves—such that each agent touches the ball at least once—before scoring a goal in as little time as possible. At the beginning of the challenge the ball is placed at the center of the field and the agents must start with at least a three meter distance, along the X axis, from each other. If the initial position of the agents does not comply with the rules, the team is awarded a score of 85. The challenge ends when a goal is scored, the ball leaves the field, or 80 seconds have passed. For each distinct agent kicking the ball—judged as the ball traveling freely for at least 2.5 meters after being kicked, the score is reduced by one point. If a goal is scored, the score is reduced by one point. If the goal is scored after the ball has been kicked by all four players, the score is the time (in seconds) from the start of the trial until the scoring event. The objective of the challenge is to get as low a score as possible.

The starting position and strategy used by UT Austin Villa for the passing and scoring challenge is shown in Figure 3. Whichever agent is closest to the ball passes the ball to a position about a meter in front of the next farthest agent from the goal as shown by the yellow arrows in Figure 3. Once the ball has been sequentially passed forward between agents and the agent closest to the goal receives the ball, that agent kicks the ball in the goal as shown by the pink arrow in Figure 3. When agents are not the closest agent to the ball they just stand in place.

Table 6 shows the results of the passing and scoring challenge where teams were ranked by the average score of a team’s best (lowest) three out of four trials. UT Austin Villa won the challenge with an average score/time of less than 20. Each of UT Austin Villa’s passing and scoring challenge trial scores were better than all the scores of other teams’ trials.

Table 6: Scores for each of the teams competing in the passing and scoring challenge.

Team	Trial 1	Trial 2	Trial 3	Trial 4	Average of Best Three Trials
<b>UTAustinVilla</b>	<b>19.34</b>	<b>20.26</b>	<b>21.28</b>	<b>19.38</b>	<b>19.66</b>
FCPortugal	82.00	21.58	22.06	22.18	21.94
AIUT3D	31.26	53.66	33.98	35.98	33.74
magmaOffenburg	30.38	83.00	23.66	82.00	45.35
HfutEngine3D	35.58	84.00	82.00	83.00	66.86
BahiaRT	81.00	83.00	83.00	82.00	82.00

<sup>5</sup> Details and framework for the passing and scoring challenge at <https://github.com/magmaOffenburg/magmaChallenge#passing-challenge>

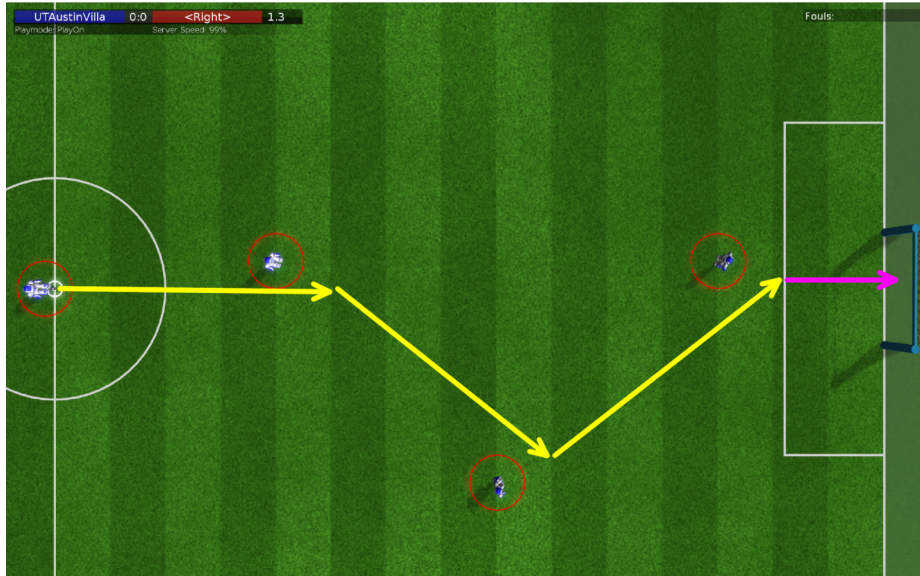


Fig. 3: Starting positions and strategy for the passing and scoring challenge. Yellow arrows represent passes between agents and the pink arrow represents a shot on goal.

### 5.3 Gazebo Running Challenge

Ongoing work within the RoboCup community is the development of a plugin<sup>6</sup> for the Gazebo [17] robotics simulator to support the RoboCup 3D Simulation League. As such, a challenge<sup>7</sup> was held where robots attempt to walk forward as fast as possible for 20 seconds in the Gazebo simulator without falling. In preparation for the challenge UT Austin Villa optimized fast walking parameters for the team’s omnidirectional walk engine [18] within the Gazebo simulator using the CMA-ES algorithm [12]. Walk engine parameters were optimized for 300 generations of CMA-ES with a population size of 150.

Results of the Gazebo running challenge are shown in Table 7. Each participating team performed four running attempts and were scored by the average forward walking speed across their three best attempts. UT Austin Villa won the challenge with all of the team’s runs having a speed of over 1.15 m/s. Each of UT Austin Villa’s running attempt speeds were greater than all other teams’ attempts. UT Austin Villa also won this same challenge at RoboCup 2016 [5].

<sup>6</sup> <https://bitbucket.org/osrf/robocup3ds>

<sup>7</sup> Framework for running the Gazebo running challenge at <https://github.com/magmaOffenburg/magmaChallenge>

Table 7: Speed in meters per second for each of the teams competing in the Gazebo running challenge.

Team	Run 1	Run 2	Run 3	Run 4	Average of Best Three Runs
<b>UTAustinVilla</b>	<b>1.176</b>	<b>1.210</b>	<b>1.159</b>	<b>1.179</b>	<b>1.188</b>
magmaOffenburg	0.354	0.365	0.417	0.256	0.379
AIUT3D	0.000	0.000	0.000	0.000	0.000
FCPortugal	0.000	0.000	0.000	0.000	0.000

## 6 Conclusion

UT Austin Villa won the 2017 RoboCup 3D Simulation League main competition as well as all technical league challenges.<sup>8</sup> Data taken using released binaries from the competition show that UT Austin Villa winning the competition was statistically significant. The 2017 UT Austin Villa team also improved dramatically from 2016 as it was able to beat the team’s 2016 champion binary by an average of 1.339 (+/- 0.039) goals across 1000 games.

In an effort to both make it easier for new teams to join the RoboCup 3D Simulation League, and also provide a resource that can be beneficial to existing teams, the UT Austin Villa team has released their base code [19].<sup>9</sup> This code release provides a fully functioning agent and good starting point for new teams to the RoboCup 3D Simulation League (it was used by six teams at the 2017 competition: AIUT3D, HfutEngine3D, KgpKubs, Miracle3D, Nexus3D, and RIC-AASTMT). Additionally the code release offers a foundational platform for conducting research in multiple areas including robotics, multiagent systems, and machine learning.

## Acknowledgments

Thanks to members of the BahiaRT and magmaOffenburg teams for helping put together the passing and scoring challenge and the Gazebo running challenge. This work has taken place in the Learning Agents Research Group (LARG) at the Artificial Intelligence Laboratory, The University of Texas at Austin. LARG research is supported in part by NSF (IIS-1637736, IIS-1651089, IIS-1724157), Intel, Raytheon, and Lockheed Martin. Peter Stone serves on the Board of Directors of Cogitai, Inc. The terms of this arrangement have been reviewed and approved by the University of Texas at Austin in accordance with its policy on objectivity in research.

## References

1. MacAlpine, P., Urieli, D., Barrett, S., Kalyanakrishnan, S., Barrera, F., Lopez-Mobilia, A., Ştiurcă, N., Vu, V., Stone, P.: UT Austin Villa 2011: A champion

<sup>8</sup> 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/#2017>

<sup>9</sup> Code release at <https://github.com/LARG/utaustinvilla3d>

- 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. LNAI. Springer (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. LNAI. Springer (2015)
  4. MacAlpine, P., Hanna, J., Liang, J., Stone, P.: UT Austin Villa: RoboCup 2015 3D simulation league competition and technical challenges champions. In: RoboCup-2015: Robot Soccer World Cup XIX. LNAI. Springer (2016)
  5. MacAlpine, P., Stone, P.: UT Austin Villa: RoboCup 2016 3D simulation league competition and technical challenges champions. In: RoboCup 2016: Robot Soccer World Cup XX. LNAI. Springer (2016)
  6. 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)
  7. Boedecker, J., Asada, M.: Simspark—concepts and application in the robocup 3d soccer simulation league. In: SIMPAR-2008 Workshop on the Universe of RoboCup Simulators. (2008) 174–181
  8. Xu, Y., Vatankeh, H.: Simspark: An open source robot simulator developed by the robocup community. In: RoboCup 2013: Robot World Cup XVII. Springer (2014) 632–639
  9. 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)
  10. MacAlpine, P., Stone, P.: Prioritized role assignment for marking. In: RoboCup 2016: Robot Soccer World Cup XX. LNAI. Springer (2016)
  11. MacAlpine, P., Stone, P.: Overlapping layered learning. *Artificial Intelligence* **254** (2018) 21–43
  12. Hansen, N.: The CMA Evolution Strategy: A Tutorial. (2009) <http://www.lri.fr/~hansen/cmatutorial.pdf>.
  13. Omidvar, M., Li, X.: A comparative study of cma-es on large scale global optimization. *AI 2010: Advances in Artificial Intelligence* (2011) 303–312
  14. Dorer, K.: Learning to use toes in a humanoid robot. In: RoboCup 2017: Robot Soccer World Cup XXI. LNAI. Springer (2017)
  15. Bengio, Y., Goodfellow, I.J., Courville, A.: Deep learning. *Nature* **521** (2015) 436–444
  16. Schulman, J., Levine, S., Abbeel, P., Jordan, M., Moritz, P.: Trust region policy optimization. In: Proceedings of the 32nd International Conference on Machine Learning (ICML-15). (2015) 1889–1897
  17. Koenig, N., Howard, A.: Design and use paradigms for gazebo, an open-source multi-robot simulator. In: Intelligent Robots and Systems (IROS). (2004)
  18. MacAlpine, P., Barrett, S., Urieli, D., Vu, V., Stone, P.: Design and optimization of an omnidirectional humanoid walk: A winning approach at the RoboCup 2011 3D simulation competition. In: Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence (AAAI-12). (2012)
  19. MacAlpine, P., Stone, P.: UT Austin Villa robocup 3D simulation base code release. In: RoboCup 2016: Robot Soccer World Cup XX. LNAI. Springer (2016)