UT Austin Villa 2011: A Champion Agent in the RoboCup 3D Soccer Simulation Competition

Patrick MacAlpine, Daniel Urieli, Samuel Barrett, Shivaram Kalyanakrishnan, Francisco Barrera, Adrian Lopez-Mobilia, Nicolae Ştiurcă, Victor Vu, Peter Stone

Department of Computer Science, The University of Texas at Austin

June 7, 2012

RoboCup 3D Simulation Domain

- Teams of 9 vs 9 autonomous agents play soccer
- Realistic physics using Open Dynamics Engine (ODE)
- Agents modeled after Aldebaron Nao robot
- Agent receives noisy visual information about environment
- Agents can communicate with each other over limited bandwidth channel

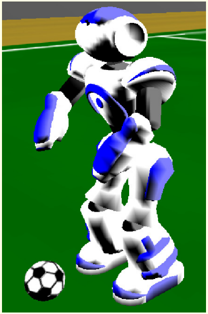

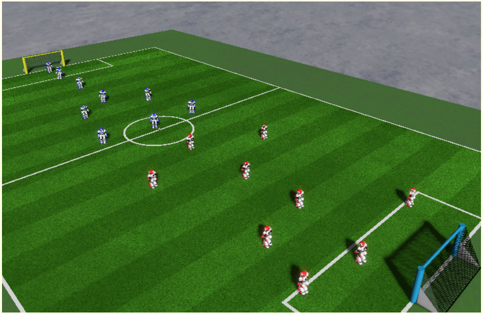

Competition Results

RoboCup	2010	2011
Goals For:	11	
Goals Against:	17	
Record (W-L-T):	4-5-1	
Place:	Outside Top-8	

Competition Results

RoboCup	2010	2011
Goals For:	11	136
Goals Against:	17	
Record (W-L-T):	4-5-1	
Place:	Outside Top-8	

Competition Results

RoboCup	2010	2011
Goals For:	11	136
Goals Against:	17	0
Record (W-L-T):	4-5-1	
Place:	Outside Top-8	

Competition Results

RoboCup	2010	2011
Goals For:	11	136
Goals Against:	17	0
Record (W-L-T):	4-5-1	24-0-0
Place:	Outside Top-8	

Competition Results

RoboCup	2010	2011
Goals For:	11	136
Goals Against:	17	0
Record (W-L-T):	4-5-1	24-0-0
Place:	Outside Top-8	1st

RoboCup	2010	2011
Goals For:	11	136
Goals Against:	17	0
Record (W-L-T):	4-5-1	24-0-0
Place:	Outside Top-8	1st

BIG IMPROVEMENT!

Key Components

- Omnidirectional Walk and Parameter Optimization
- Inverse Kinematics Based Kicking Engine
- Dynamic Role Assignment and Positioning System

Omnidirectional Walk Engine

- Double linear inverted pendulum model
- Based closely on that of walk engine by Graf et al
- Mostly open loop but not entirely
- Designed on actual Nao robot

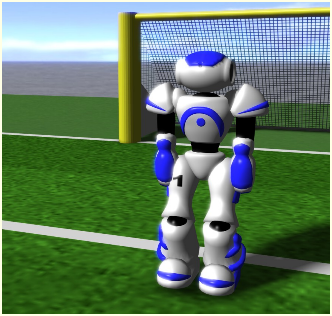

Initial Walk Parameters

- 40 different parameters
- Designed and hand-tuned to work on the actual Nao robot
- Provides a slow and stable walk

2011 Omnidirectional Walk Optimization

- Parameters (14) optimized through CMA-ES across a cluster
  - Population of 150 across 200 generations = 210,000 evaluations in less than a day
- Learn three different parameter sets for three different subtasks
  - Going to a target
  - Sprinting forward
  - Positioning around the ball when dribbling
- Parameters learned through a layered learning approach
  - Parameter sets learned sequentially
  - Each parameter set learned *in conjuction* with each other
  - Agent able to seamlessly transition between parameter sets

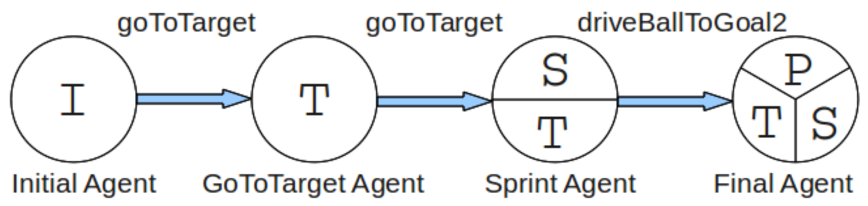

I = *initial*, T = *goToTarget*, S = *sprint*, P = *positioning*

Final Agent Video

Red 'T' = *gotoTarget* parameters, yellow 'S' = *sprint* parameters, cyan 'P' = *positioning* parameters

Optimized Omnidirectional Walk Performance

- Beat agent with initial hand-tuned walk parameters by average of 8.84 goals across 100 games
- Beat agent using non-omnidirectional walk used in 2010 by average of 6.32 goals across 100 games

Kick Engine Kinematics

- Define waypoints relative to ball for foot to reach
- Cubic Hermite splines used to compute path for foot to follow
- Inverse kinematics system determines if kick can be executed
- Optimize parameters of kick: waypoint values, speed, ball offset
- Learn kicks for multiple directions and orientations to the ball

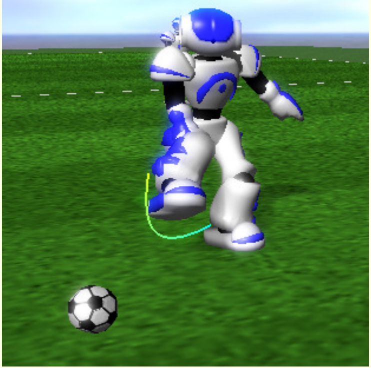

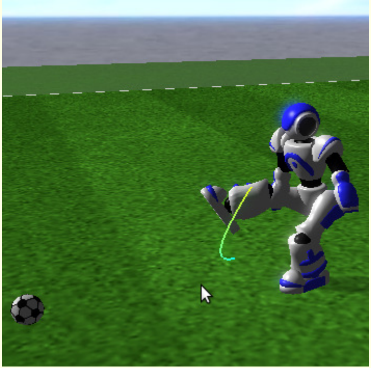

Kicking Video

Different directional kicks

Kick Performance

- Kicking agent loses by .15 goals on average to *dribble only* agent
- Strategy for best using kick not yet implemented (no passing yet)
- Shows improvement when used with an agent with a less effective walk (agent with *initial* walk parameters)
  - Kicking agent scored 8 goals while non-kicking agent failed to score when playing 100 games against each other

Role Assignment Mapping

- Every player assigned to a role (position) on the field
- Positions based on offsets from ball or endline
- *onBall* role assigned to the player closest to the ball
- One-to-one mapping of agents to positions
- Can be thought of as a role assignment *function*

Role Assignment Mapping

- Every player assigned to a role (position) on the field
- Positions based on offsets from ball or endline
- *onBall* role assigned to the player closest to the ball
- One-to-one mapping of agents to positions
- Can be thought of as a role assignment *function*

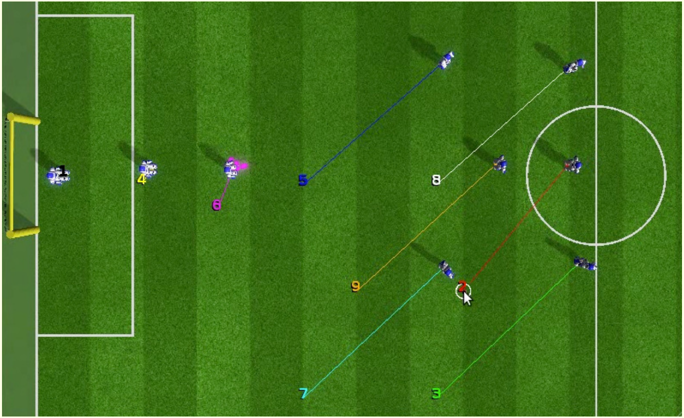

Role Assignment Function (f_v)

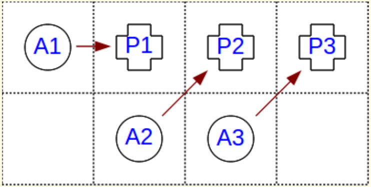

Lowest lexicographical cost (shown with arrows) to highest cost ordering of mappings from agents (A1,A2,A3) to role positions (P1,P2,P3). Each row represents the cost of a single mapping.

1:	$\sqrt{2}$ (A2$\rightarrow$P2),	$\sqrt{2}$ (A3$\rightarrow$P3),	1 (A1$\rightarrow$P1)
2:	2 (A1$\rightarrow$P2),	$\sqrt{2}$ (A3$\rightarrow$P3),	1 (A2$\rightarrow$P1)
3:	$\sqrt{5}$ (A2$\rightarrow$P3),	1 (A1$\rightarrow$P1),	1 (A3$\rightarrow$P2)
4:	$\sqrt{5}$ (A2$\rightarrow$P3),	2 (A1$\rightarrow$P2),	$\sqrt{2}$ (A3$\rightarrow$P1)
5:	3 (A1$\rightarrow$P3),	1 (A2$\rightarrow$P1),	1 (A3$\rightarrow$P2)
6:	3 (A1$\rightarrow$P3),	$\sqrt{2}$ (A2$\rightarrow$P2),	$\sqrt{2}$ (A3$\rightarrow$P1)

- Mapping cost = vector of distances sorted in decreasing order
- Optimal mapping = lexicorgraphically sorted lowest cost mapping
- Recursively minimizes longest distance any agent must travel
- Avoids collisions

Dynamic Programming Algorithm for Role Assignment

- Any subset of a lowest cost mapping is itself a lowest cost mapping
- Begin evaluating mappings of 1 agent and build up to n agents
- Only evaluate mappings built from subset mappings returned by f_v
- Evaluates $n2^{n-1}$ mappings, for $n = 8$ is 1024 (brute force = 40,320)

Positioning Video

Each position is shown as a color-coded number corresponding to the agent's uniform number assigned to that position. Agents update their role assignments and move to new positions as the ball or an agent is beamed (moved) to a new location.

Positioning System Evaluation

Team	Goal Difference
Static	.32 (.07)
AllBall	.43 (.09)

Static Each role is statically assigned to an agent
AllBall Every agent except goalie goes to the ball

Competition Analysis

Average goal difference across 100 games against other agents in the competition

Rank	Team	Goal Difference
3	apollo3d	1.45 (.11)
5-8	boldhearts	2.00 (0.11)
5-8	robocanes	2.40 (0.10)
2	cit3d	3.33 (0.12)
5-8	fcportugal3d	3.75 (0.11)
9-12	magmaoffenburg	4.77 (0.12)
9-12	oxblue	4.83 (0.10)
4	kylinsky	5.52 (0.14)
9-12	dreamwing3d	6.22 (0.13)
5-8	seuredsun	6.79 (0.13)
13-18	karachikoalas	6.79 (0.09)
9-12	beestanbul	7.12 (0.11)
13-18	nexus3d	7.35 (0.13)
13-18	hfutengine3d	7.37 (0.13)
13-18	futk3d	7.90 (0.10)
13-18	naoteamhumboldt	8.13 (0.12)
19-22	nomofc	10.14 (0.09)
13-18	kaveh/rail	10.25 (0.10)
19-22	bahia3d	11.01 (0.11)
19-22	l3msim	11.16 (0.11)
19-22	farzanegan	11.23 (0.12)

- Across 2100 games played won all but 21 games which ended in ties (no losses)

Performance Contributions

Performance Contributions

- Agent with 2010 non-omnidirectional walk would have finished in tenth place

Performance Contributions

- Agent with 2010 non-omnidirectional walk would have finished in tenth place

- Number of times goalie touched the ball during the 2011 competition = 0

Summary

Summary

- UT Austin Villa is a complete agent that won the 2011 RoboCup 3D simulation competition

Summary

- UT Austin Villa is a complete agent that won the 2011 RoboCup 3D simulation competition
- Key components of the agent are it's omnidirectional walk, kicking engine, and dynamic positioning system

Summary

- UT Austin Villa is a complete agent that won the 2011 RoboCup 3D simulation competition
- Key components of the agent are it's omnidirectional walk, kicking engine, and dynamic positioning system
- The omnidirectional walk proved to be the crucial component in winning the competition

Summary

- UT Austin Villa is a complete agent that won the 2011 RoboCup 3D simulation competition
- Key components of the agent are it's omnidirectional walk, kicking engine, and dynamic positioning system
- The omnidirectional walk proved to be the crucial component in winning the competition
- Optimizing parameters though machine learning is the underlying theme for the team's success

Future Work

- Attempt to apply learned walks in simulation to actual Nao robots
- Improve kicking strategy and add passing
- Attempt to learn better formations with machine learning

2012 Kickoff

More Information

UT Austin Villa 3D Simulation Team homepage:
www.cs.utexas.edu/~AustinVilla/sim/3dsimulation/

Email: patmac@cs.utexas.edu

This work has taken place in the Learning Agents Research Group (LARG) at UT Austin. LARG research is supported in part by NSF (IIS-0917122), ONR (N00014-09-1-0658), and the FHWA (DTFH61-07-H-00030).

Related Work

- N. Hansen. The CMA Evolution Strategy: A Tutorial, January 2009.
- C. Graf, A. Härtl, T. Röefer, and T. Laue. A robust closed-loop gait for the standard platform league humanoid.
- N. Shafii, L. P. Reis, and N. Lao. Biped walking using coronal and sagittal movements based on truncated Fourier series, January 2010.
- J. E. Pratt. Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, June 2000.
- N. Kohl and P. Stone. Machine learning for fast quadrupedal locomotion, 2004.
- D. Urieli, P. MacAlpine, S. Kalyanakrishnan, Y. Bentor, and P. Stone. On optimizing interdependent skills: A case study in simulated 3d humanoid robot soccer, 2011.
- P. MacAlpine, S. Barrett, D. Urieli, V. Vu, and P. Stone. Design and Optimization of an Omnidirectional Humanoid Walk: A Winning Approach at the RoboCup 2011 3D Simulation Competition, 2012.
- P. MacAlpine, F. Barrera, and P. Stone. Positioning to Win: A Dynamic Role Assignment and Formation Positioning System, 2013.
- P. Stone and M. Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork, 1999.

Related Work

- N. Hansen. The CMA Evolution Strategy: A Tutorial, January 2009.
- C. Graf, A. Härtl, T. Röefer, and T. Laue. A robust closed-loop gait for the standard platform league humanoid.
- N. Shafii, L. P. Reis, and N. Lao. Biped walking using coronal and sagittal movements based on truncated Fourier series, January 2010.
- J. E. Pratt. Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, June 2000.
- N. Kohl and P. Stone. Machine learning for fast quadrupedal locomotion, 2004.
- D. Urieli, P. MacAlpine, S. Kalyanakrishnan, Y. Bentor, and P. Stone. On optimizing interdependent skills: A case study in simulated 3d humanoid robot soccer, 2011.
- P. MacAlpine, S. Barrett, D. Urieli, V. Vu, and P. Stone. Design and Optimization of an Omnidirectional Humanoid Walk: A Winning Approach at the RoboCup 2011 3D Simulation Competition, 2012.
- P. MacAlpine, F. Barrera, and P. Stone. Positioning to Win: A Dynamic Role Assignment and Formation Positioning System, 2013.
- P. Stone and M. Veloso. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork, 1999.