

A Layered Approach to Learning Client Behaviors in the RoboCup Soccer Server

- Peter Stone
 - Ph.D. thesis at CMU, 1998
 - Book from MIT Press
- Manuela Veloso
 - Advisor
- <http://www.cs.utexas.edu/~pstone/thesis/>


Multiagent Learning in Robocup Soccer Simulation

- An incremental approach to learning is used
- Difficult (impossible?) to learn complex behaviors based solely on the primitives provided by the server
- Layered Learning:
 - building increasingly complex behaviors on top of one another;
 - learning low-level behaviors before high-level ones;
 - Higher-level behaviors utilize lower-level ones as components

The Objective

- Teach the server clients two learned behaviors: one low-level and one high-level
- Low-level: intercepting a moving ball
- High-level: passing
- The high-level skill is dependent on the mastery of the low-level one.

The Robotic Competition Architecture



The Simulator

- Noda's Soccer Server
 - Players have limited vision
 - Players communicate via a blackboard visible to all players
 - All players controlled by separate processes
 - Actions and sensors are noisy
 - Real time play

Low-level skill: Ball interception

- Intercepting a moving ball a prerequisite for any kicking action
- Learning through supervised empirical learning: Neural Networks (NNs).
- Difficulties:
 - Ball movement unpredictable due to noise in the system
 - Players have limited vision

Low-level skill: Ball interception (cont.)

- Defender learns turn angle (TA—the angle to turn after facing the ball), based on the ball's distance and ball angle.
- NN training
 - randomized situations and defender actions
 - Results recorded
 - SAVE: ball is intercepted
 - GOAL: ball is not intercepted, ball goes into the goal
 - MISS: ball is not intercepted, no goal made

Low-level skill: Ball interception (cont.)

- NNs trained on 500+ examples had a SAVE rate of roughly 83-86%.
- NN results compared against a lookup table mapping ball angle to the average NN output for that ball angle—results were almost as good as with the NN itself.
- NN results also compared with an analytical predictor. Again, results were almost as good as the NN results.

High-level skill: passing

- Passing involves two team members: a passer and a receiver.
- The receiver's action in passing is identical to the defender's action in ball interception. The same NN is used.
- Decision Trees (DTs) were used to decide if the ball should be passed to a particular teammate.

High-level skill: passing (cont.)

- Passer uses the receivers' views of the field in addition to its own when making the decision whether to pass.
- During training, passer chose a random receiver in each trial.
- In addition to the intended receiver, 4 defenders also attempted to intercept the ball.
- Results:
 - SUCCESS: the intended receiver intercepts the ball
 - FAILURE: a defender intercepts the ball
 - MISS: no one intercepted the ball

High-level skill: passing (cont.)

- 174 attributes were recorded for each trial, to be used to predict whether a pass will be a success, a failure, or a miss.
 - Half the data points were from the passer's perspective, half from the receiver's
- Decision Tree software produced a pruned tree, in which insignificant data were left out.
 - Many more passer attributes were actually used than receiver attributes.
- 5000 training examples
 - 51% successes
 - 42% failures
 - 7% misses

High-level skill: passing (cont.)

- DT returns a confidence measure for each receiver candidate
 - If more than one will be successful, choose the one with the highest confidence level
 - If all are predicted to fail, choose the one with the lowest confidence level
- 5000 trials
 - Passer must pass, even if DT predicts failure for all receivers.
 - 79% success rate

Team-level strategies

13

- A set play was run several times to show that the higher-level learning (passing decision) could be incorporated into a game situation.
- Goal: keep layering new higher-level learned behaviors on top of old ones, resulting in a high-level functioning team that is robust and reliable in game situations.

Summary

14

- Learning strategies developed
- Successful in a large application
- Peter Stone is now at UT Austin