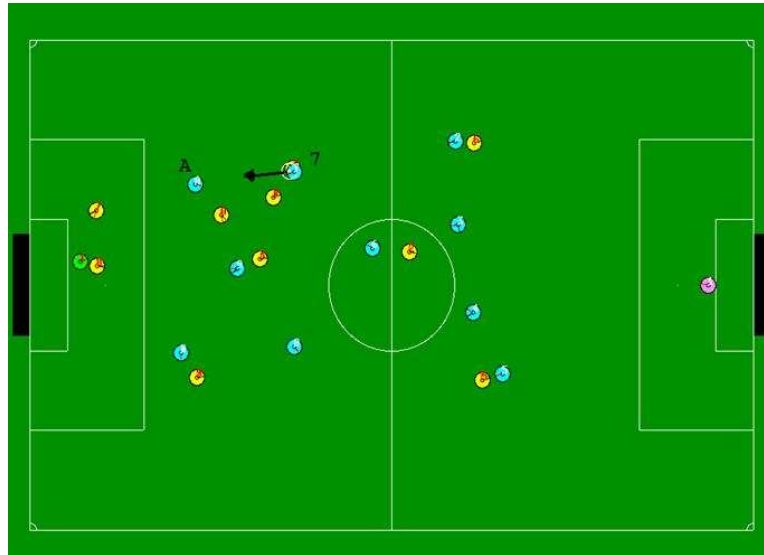


# Opponent modeling in the RoboCup Simulator




Gregory Kuhlmann and Peter Stone

# Outline

- The Coach Competition
- The UT Austin Villa Coach
- Changes for 2005
- Something completely different
  - General Game Playing

# RoboCup Simulator

- **Distributed**: each player a separate client
- Server models dynamics and kinematics
- Clients receive **sensations**, send **actions**



- Parametric actions: **dash, turn, kick, say**
- **Abstract, noisy** sensors, hidden state
  - **Hear** sounds from limited distance
  - **See** relative distance, angle to objects ahead
- $> 10^{9^{23}}$  states
- **Limited resources** : stamina
- Play occurs in **real time** ( $\approx$  human parameters)

# Motivation for Coaching

- MAMSIG
  - Aim: encourage research in opponent modeling
  - Challenge: create a simulated coach
    - \* autonomous agent that gives advice
    - \* improves performance of a team against a fixed opponent
- Power of a coach:
  - More a priori knowledge
  - Better view of world
  - More computational resources
- Prerequisites:
  - coachable players (programmed by others)
  - standardized coaching language

# RoboCup Coach Competition

- Sub-league of RoboCup Simulator League
- Coaching scenario:
  - Access to log files (“game films”) of fixed opponent
  - Noise-free, omniscient view of field
  - Limited communication (once every 300 cycles, 50 cycle delay)
    - can’t micromanage. No centralized control.
  - Advice sent in standardized coach language
  - Players to follow advice *most of the time*
  - Performance measured by goal difference
- Good test of opponent modeling?

# RoboCup Coach Competition (contd.)

- 4 International Competitions (*plus regional events*)
  - Early years - best result worse than no advice
    - \* teams already coherent and competent
    - \* probably stuck in local maximum
  - 2003 - coaching helped
    - \* team of players from several institutions (UT, CMU, USTC)
    - \* little or no default strategy.
  - 2004 - some rule changes
    - \* standardized communication language
    - \* new scoring metric
    - \* limited time to review logfiles
  - 2005 - ?


# CLang

- Standardized Coach Language
  - independent of coachable player's behavior representation
- If-then rules:  
 $\{\textit{condition}\} \rightarrow \{\textit{action}\}$
- Example:  
If **our player 7 has the ball**, then **he should pass to player 8 or player 9**

```
(definerule pass789 direc
  ((bowner our {7})
   (do our {7} (pass {8 9}))))
```

## Example: UT Austin Villa Coachable Player

- Candidate actions are assigned values using a heuristic
  - Based on probability and value of success
- Before advice:




- Action with highest value is chosen



## Example: UT Austin Villa Coachable Player (contd.)

- Advice bumps values up (or down)
- When rule `pass789` becomes active:



- generally takes best advised action
- possible to override advice

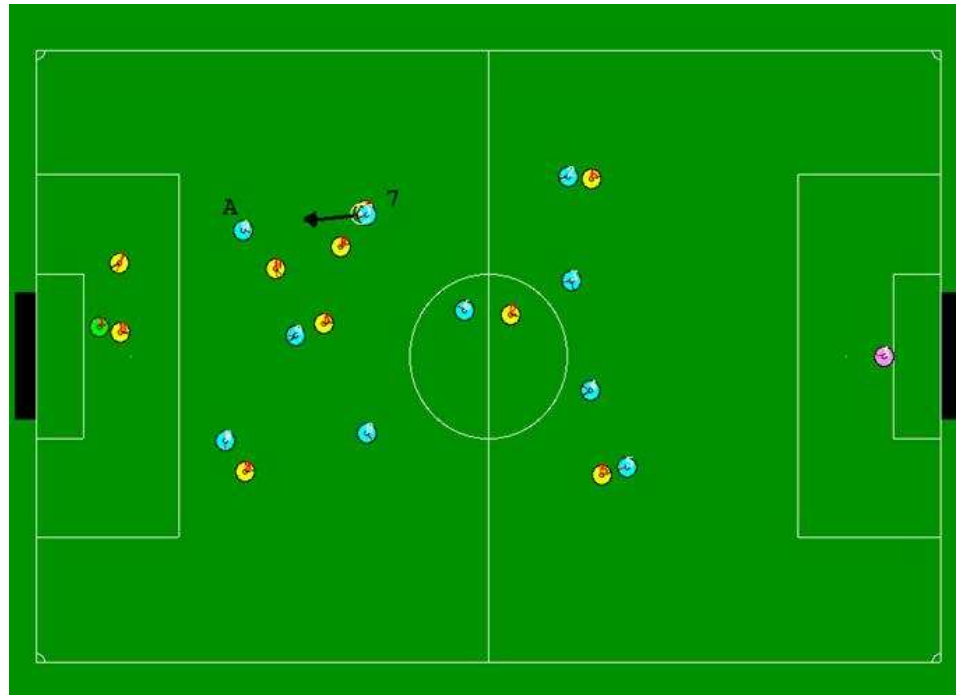
# The UT Austin Villa Coach

- Opponent-specific advice
  - Learned **defensive positioning** advice
    - \* predict opponent passes
    - \* advise player to block pass
  - Learned **offensive action selection**
    - \* mimic successful team's passing and shooting
  - Learned **formations**
    - \* mimic successful team's positioning
    - \* average position + ball attraction
- Handcoded rules
  - encode general soccer strategy

## The UT Austin Villa Coach (contd.)


- Game analysis
  - Given **x** and **y** coordinates
  - Detect high-level events: **play-by-play**
- Offline learning
  - Learn from **logfiles**
  - **Online** learning possible but **difficult**
  - All advice sent at **start** of game

# Predicting Agent Behavior



- **Inputs:** features of current world state
  - Player locations, distances to ball and goal, current score, etc.
- **Classification:**  $\text{PassFrom}_k$ 
  - **Example:**  $\text{PassFrom}_7$  stored in opponent 10's training set

# Model: Decision Trees



- Compile training instances
- Train decision tree for **each modeled player**
  - J48 algorithm (*weka*)
  - very much like C4.5

# Generating Advice

- Generate advice for each leaf node in tree
  - Action to counter predicted opponent action
  - Example:
    - \* If opponent 10's y-coordinate is greater than 17.7, then position our player 4 between opponent 10 and opponent 7

```
(definerule def4rule1 direc
  ((ppos opp {10} (rec (pt -52.5 34) (pt 52.5 17.7))))
  (do our {4} (pos (((pt opp 10) * (pt .7 .7)) +
                    (pt opp 7) * (pt .3 .3))))))
```

## Incorporating Advice



- Thanks to the advice, defender 4 is ready to intercept a pass from opponent 7 to 10.

## Competition Results

| Team            | 1st Round |     | 2nd Round |     | 3rd Round |     |
|-----------------|-----------|-----|-----------|-----|-----------|-----|
| UT Austin Villa | 0:19      | 7th | 0:2       | 1st | 8:2       | 1st |
| FC Portugal     | 1:21      | 8th | 0:8       | 4th | 7:3       | 2nd |
| Iranians        | 0:14      | 4th | 0:5       | 3rd | 3:2       | 3rd |
| Helli-Amistres  | 1:12      | 2nd | 0:3       | 2nd | 7:7       | 4th |

- 1st place in 2003 RoboCup Coach Competition
- Only one other team used learning
- Further experiments statistical tie with second place



# Experimental Results

| Opponent   | w/ HC | None | Formation | Offensive | Defensive | Full |
|------------|-------|------|-----------|-----------|-----------|------|
| BoldHearts | N     | -8.8 | -3.3      | -2.9      | -2.9      | -2.7 |
|            | Y     | -6.8 | -0.5      | -1.4      | -5.7      | -6.5 |
| Sirim      | N     | -4.1 | 2.6       | 1.2       | 0.9       | 1.7  |
|            | Y     | -5.4 | -1.6      | -0.3      | 0.8       | -0.4 |
| EKA-PWr    | N     | -0.6 | 2.8       | 2.9       | 3.4       | 2.7  |
|            | Y     | 1.0  | 3.62      | 2         | 2.12      | 2.43 |

- Formation learning helps
- Handcoded sometimes hurts
- Offensive and defensive advice mixed
- Why?

# Changes for 2005

- What happened in the 2004 competition?
  - Online learning (k-armed bandit).
  - Two of top three never saw opponent.
- Make opponent modeling necessary
  - Test prediction, not exploitation
  - Make defects more obvious
  - Take the human out of the loop

# 2005 Competition

- Detect **patterns**: simple exploitable behaviors
- **Offline phase**
  - Given one log file of **base** strategy
  - One log file for each **base+pattern** (labeled)
- **Online phase**
  - Play full match using standard coachable players
    - \* send **advice** to facilitate detection
  - Opponent with **two or more** patterns activated
    - \* base strategy may be different
  - Report active patterns
    - \* more points for reporting **sooner**
    - \* penalty for **incorrect** detection

# General Game Playing

- **Challenge**: play a game you've **never played before**
- Perfect information, deterministic
- Single or multi-player
- Simultaneous decision or turn-taking
- **Yes**: 8-puzzle, Tic-tac-toe, Go, Chess, Roshambo, Repeated Prisoner's dilemma
- **No**: Yahtzee, Backgammon, Battleship, Poker
- Given a game description in **first-order logic**
  - initial state
  - state transition function
  - legal moves
  - terminal and goal states

## General Game Playing (contd.)

- **Competition** at AAI 05
  - Description of **unseen** game sent to agents
  - 30 seconds to **think** between moves
  - Illegal moves punished
  - Best score wins
- **Agent modeling** in GGP
  - Have perfect model of all but other agents
  - Only have raw features
  - Must figure out **cooperative/competitive**
- Winner gets **\$10,000**