Data-Driven Generation of Simulated Soccer Behaviors

David C. Brogan
University of Virginia
dbrogan@cs.virginia.edu

Yannick Loitière
University of Virginia
ycl2r@cs.virginia.edu

ABSTRACT

The dynamics of how groups move through space to accomplish common goals must be understood to create realistic synthetic environments. One potential method for creating such multiagent behaviors is to replay prerecorded examples of group movements. While these data-driven methods effectively capture the original performance for a particular instance, the success of these methods for interactive, multiagent applications is limited by the large number of potential agent movements that must be prerecorded. To mitigate the scaling effects of data-driven multiagent behavior algorithms, we propose a behavior model that reduces the dimensionality of prerecorded data and decreases the amount of data required by effectively using available data. We have chosen to investigate the sport of simulated soccer and have developed behaviors for simulated soccer players from the data acquired from recent RoboCup games.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning

General Terms

Algorithms

Keywords

Coordinating multiagents and activities

1. INTRODUCTION

Multiagent behaviors play a vital role in coordinated robotics, scientific simulations, and surveillance systems. In the case of mobile agents, the behaviors must generate movements for each individual based on its state relative to other agents and the world. Our system automatically develops behaviors for mobile agents from the observation of the movements of multiagent groups. With these behaviors, we demonstrate a multiagent system that anticipates the future actions of the observed system and retargets behaviors to different groups in novel circumstances.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
AAMAS'02, July 15-19, 2002, Bologna, Italy.
Copyright 2002 ACM 1-58113-480-0/02/0007 ...$5.00.

We wish to predict and generate the behaviors of the simulated soccer players used in the popular simulated soccer testbed, RoboCup [3]. The game logs of RoboCup soccer games are widely available and provide our system with the observational data needed to build behavioral models of player actions and team strategies. Although the game logs provide extensive details of the game state, we only use the player positions in our algorithms and thus these algorithms would be compatible with the data generated by future image-based systems for human soccer games.

The machine learning and computer graphics communities have investigated learning by example [1] and motion capture [2] technologies, but all data-driven methods like these are confronted with the need to use models to fill gaps in the database, blend incongruous observations, and extrapolate to new circumstances. To be effective and realistic, these behavioral models must possess sufficient degrees of freedom to discriminate between similar, but distinct, observations while not becoming so rigidly defined that the model size and complexity scale uncontrollably due to the myriad observed interactions between the agents and the dynamic world. Our contribution is to map the high-dimensional player data observed during a simulated soccer game to a simplified representation that preserves the ability to synthesize and predict player movements.

2. METHODOLOGY

Simulated soccer data consists of the 22 player positions and the ball position sampled at 6,000 moments throughout the game, forming a sparse data set with high dimensionality. This data property impedes the ability of data-driven modeling techniques to automatically capture relationships between one state and another. Furthermore, the labeling of the data fields will cause comparison metrics to fail to indicate the similarity between two samples of game state when player-occupied field locations are unchanged, but the location of two players are exchanged. We argue that in these simulated soccer games, all players on a team can be treated equally and thus they need not be uniquely identified in the game state.

We convert each simulated soccer game log into a sequence of presence density maps (PDMs, see figure 1), with one data channel for each class of entity we are interested in, namely the players from each team and the ball: $\mathcal{T} = \{\tau_{team_A}, \tau_{team_B}, \text{and} \tau_{ball}\}$. Each channel of the PDM is defined as a 2-D function of the constituent members of its associated agent class; formed by summing the set of 2-D Gaussians centered over each entity in the class then sampled in a grid covering the entire playing field. The motivation for this transposition comes from fluid mechanics methodologies where it is common to switch focus from individual fluid particles (Lagrangian) to the space through which they move (Eulerian). This mapping permits our system to predict the image captured by an overhead camera at a soccer game rather than the positions of each individual player. This data representation enables image processing techniques analog to those developed by Schödl et al. [4], wherein finite-length image sequences are parsed and recombined into arbitrarily long continuous animations.

To predict how the state of a current game will change during δ frames, we developed a comparison metric, a predictor, and a corrector. The PDM sequences from multiple soccer games are placed in a database, $\mathcal{D}$, with ordering preserved to maintain time-dependent correlations. The system predicts future PDMs based on the PDM, g, of the current game state by comparing it to all the PDMs, $d \in \mathcal{D}$, using the comparison metric, $\mu(g, d)$:

$$\mu(g,d) = 1 - \exp(-\sum_{\tau \in \mathcal{T}} {\psi_\tau^+}^2(g,d)) \tag{1}$$

The cross correlation $\psi(x, y)$ is used to evaluate the similarity between each channel of $\mathcal{T}$ in the PDMs:

$$\psi(x,y) = \frac{1}{N} \cdot \sum_{i \in [1,N]} \frac{(x_i - \bar{x})(y_i - \bar{y})}{\sigma_x \cdot \sigma_y} \tag{2}$$

$$\text{and } \psi^+ = \begin{cases} 0 & \text{if } \psi \leq 0 \\ \psi & \text{otherwise} \end{cases} \tag{3}$$

where N is the number of samples, $\bar{x}$ is the mean, and σ is the variance.

The resulting comparisons are then used to generate the predicted PDM δ frames in the future, $\tilde{g}_\delta$. For each state in the database, d_i, the system accesses the successor state that is δ frames ahead, $d_{i+\delta}$. The function $\tilde{g}_\delta$ is a weighted average of the set of successor frames, where the weighting is derived from the comparison metric and from a reinforcement learning correction term, λ.

$$\tilde{g}_\delta(t) = \frac{\sum_i d_{i+\delta} \cdot \mu(g, d_i) \cdot \mu(d_i, d_{i+\delta}) \cdot \lambda_{d_i}}{\sum_i \mu(g, d_i) \cdot \mu(d_i, d_{i+\delta}) \cdot \lambda_{d_i}} \tag{4}$$

The correction term λ_d is associated with each PDM, $d \in \mathcal{D}$. Once δ frames have elapsed in the current game, the correction terms of all PDMs get updated so those that accurately predicted future states are rewarded:

$$\forall d \in \mathcal{D}, \lambda_d = (1 - \epsilon) \cdot \lambda_d + \epsilon \cdot \mu(\tilde{g}_\delta(t - \delta), d) \tag{5}$$

where ϵ is a tunable parameter.

3. EXPERIMENTAL RESULTS

We have completed two series of tests to evaluate the performance of our data-driven predictor. We measured how the accuracy of the prediction was affected by the lookahead term, δ. As a reference technique, we constructed a kinematic extrapolator that computes a future PDM based on the current velocities and positions of the ball and players, while our method only uses position information.

Figure 1: *presence density map* for a team defending its goal. Dark regions indicate player locations.

Games were played between the 2000 champion team, FC Portugal and Wahoo Wunderkind, comparing the performance of the predictor with and without the benefit of corrector feedback. Although the kinematic predictor performs better for small δ values, its performance degrades rapidly as δ increases and the data-driven predictor outperforms it for values of δ greater than 17 frames with corrector feedback, or 29 without (see figure 2). These results indicate the value of using velocity data to complement our behavior model, which we will investigate in future work.

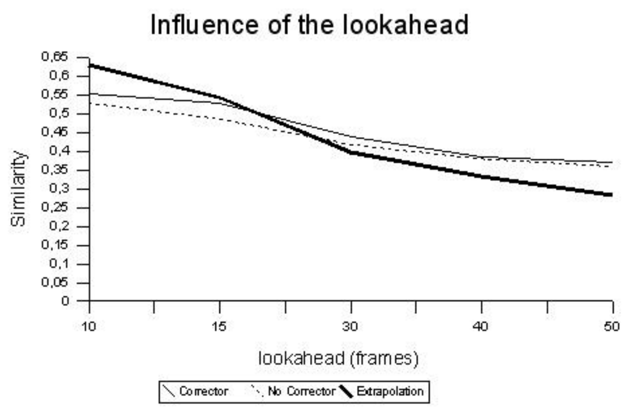

Figure 2: This graph plots the matching performance of the data-driven predictor and the kinematic extrapolator. A perfect score of 1.0 indicates an exact match between the predicted PDM and the actual future PDM from the test data.

4. REFERENCES

[1] C. Atkeson and S. Schaal. Robot learning from demonstration. In *Proceedings of the Fourteenth International Conference (ICML '97)*, pages 12–20, 1997.

[2] M. Gleicher. Rertagetting motion to new characters. In *Proceedings of SIGGRAPH 98*, pages 33–42, 1998.

[3] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, and E. Osawa. RoboCup: The robot world cup initiative. In *Agents '97*, pages 340–347, 1997.

[4] A. Schödl, R. Szeliski, D. Salesin, and I. Essa. Video textures. In *Proceedings of SIGGRAPH 2000*, pages 489–498, 2000.