2007

Generating Plausible Individual Agent Movements From Spatio-Temporal Occupancy Data

Ben Sunshine-Hill

University of Pennsylvania, bsunshin@seas.upenn.edu

Jan M. Allbeck

University of Pennsylvania, allbeck@seas.upenn.edu

Nuria Pelechano

University of Pennsylvania

Norman I. Badler

University of Pennsylvania, badler@seas.upenn.edu

Follow this and additional works at: <https://repository.upenn.edu/hms>

Part of the [Engineering Commons](#), and the [Graphics and Human Computer Interfaces Commons](#)

Recommended Citation

Sunshine-Hill, B., Allbeck, J. M., Pelechano, N., & Badler, N. I. (2007). Generating Plausible Individual Agent Movements From Spatio-Temporal Occupancy Data. *2007 Workshop on Massive Datasets*, 5-7.
<http://dx.doi.org/10.1145/1352922.1352924>

This conference took place on November 15, 2007, Nagoya, Japan.

This paper is posted at ScholarlyCommons. <https://repository.upenn.edu/hms/212>
For more information, please contact repository@pobox.upenn.edu.

Generating Plausible Individual Agent Movements From Spatio-Temporal Occupancy Data

Abstract

We introduce the Spatio-Temporal Agent Motion Model, a datadriven representation of the behavior and motion of individuals within a space over the course of a day. We explore different representations for this model, incorporating different modes of individual behavior, and describe how crowd simulations can use this model as source material for dynamic and realistic behaviors.

Keywords

crowd modeling, motion sensors, crowd simulation

Disciplines

Computer Sciences | Engineering | Graphics and Human Computer Interfaces

Comments

This conference took place on November 15, 2007, Nagoya, Japan.

Generating Plausible Individual Agent Movements from Spatio-Temporal Occupancy Data

Ben Sunshine-Hill¹, Jan Allbeck², Nuria Pelechano³, and Norm Badler⁴

Center for Human Modeling and Simulation

University of Pennsylvania

Philadelphia, PA 19104-6389

ABSTRACT

We introduce the Spatio-Temporal Agent Motion Model, a data-driven representation of the behavior and motion of individuals within a space over the course of a day. We explore different representations for this model, incorporating different modes of individual behavior, and describe how crowd simulations can use this model as source material for dynamic and realistic behaviors.

Categories and Subject Descriptors

I.3.7 [Computer Graphics]: Animation.

General Terms

Algorithms, Measurement, Human Factors.

Keywords

Crowd modeling, motion sensors, crowd simulation.

1. INTRODUCTION

Existing crowd simulation systems often consider crowds as homogeneous masses, full of agents with similar goals, strategies, and modes of behavior. Others assign each agent to one of a small set of goals, in order to process each member of a “goal group” using the same computations. Likewise, the time of day is not usually an explicit factor in a crowd simulation; a particular scenario represents a single mode of crowd behavior, accurate only during a particular time period.

1.1 Spatio-Temporal Agent Motion Models

We have processed the MERL motion sensor data to build a “Spatio-Temporal Agent Motion Model” (STAMM), a time-dependent probabilistic model of the motion of individuals within a building. Using this model, an arbitrary number of virtual agents can be simulated, their behavior mimicking that of the actual observed agents in all practical respects. Smoothing can be used to produce useful results even with relatively small datasets. The STAMM is trained from available tracklet data, using a random-walk scheme.

1.2 Related Work

1.2.1 Small-Scale Locomotion

Simulation of the locomotion of individuals within crowds has been undertaken using a variety of approaches. Reynolds used a weighted combination of flocking behaviors to coordinate individual motions[5], while Helbing et al. modeled crowd behavior using fluid dynamics-like laws[3]. Braun et al. extended the Helbing model to allow for different individuals to behave differently[2]. Pelechano et al. fused rule-based and social-forces-based models and incorporated psychological state into the simulation model[4]. These methods model small-scale motions rather than path-planning, and can be driven directly by the methods given in this paper.

1.2.2 Multilayer Crowd Simulation

By layering higher-level planning on top of small-scale locomotion systems, more intelligent and realistic crowd behavior is obtained. Sung et al. used probabilistic roadmaps to pre-plan locomotion sequences for agents[6], while Bayazit et al. layered high-level roadmaps on flocking behaviors in order to obtain larger-scale behaviors[1]. In contrast, Treuille et al. used dynamic potential fields in order to integrate large-scale and small-scale behaviors[8]. Shao et al. used a complex cognitive and behavior model for planning, but did not attempt realistic small-scale locomotion[6].

2. AGENT CREATION

The MERL motion sensor data does not include data for offices, so agents are considered to appear and disappear at certain points for the purposes of modeling. The creation of agents in a STAMM is based on two probability distributions. The first is $\Pr[n_{\text{creations}}|t]$, specifying the probability that a given number of agents will be created during a particular time period of the day; empirically we have found this distribution to be Gaussian with respect to $n_{\text{creations}}$ and multimodal Gaussian with respect to Time, confirming our assumptions. We then use a second distribution, $\Pr[X_{\text{start}}|t]$, to determine where each agent starts in the world. While it would of course be possible to combine these two distributions, separating them gives the advantage of allowing a finer time resolution for the former without overfitting the latter. If lunch starts at 12:15 sharp, a 60-second time resolution for $\Pr[n_{\text{creations}}|t]$ can precisely

¹ sunshine@seas.upenn.edu

² allbeck@seas.upenn.edu

³ npelecha@seas.upenn.edu

⁴ badler@seas.upenn.edu

capture that, while $\Pr[X_{\text{Start}}|t]$ can use 15-minute resolution in order to avoid artifacts from sparse data.

3. AGENT BEHAVIOR

In much of the MERL data, agents are observed to walk purposefully from point A to point B, making them amenable to a model that assumes only simple paths and near-constant speed in the absence of agent interactions. At the same time, however, in the data we have observed other agents lingering in hallways and making U-turns, behaviors that such a model could not easily represent. Therefore, we have explored both goal-directed and non-goal-directed representations for the STAMM.

The probability distributions for these behaviors are learned directly from random walks over the MERL data's tracklet graph, starting at tracklets with no predecessors and ending at tracklets with no successors. Although we expected this random walk sampling to produce problems when agents passed each other in hallways, reducing the efficiency of goal-directed paths, in practice we found no such problems.

3.1 Goal-Directed Behavior

A goal-directed STAMM is one in which each agent in the simulation has an associated goal position. For the sake of simplicity, we used the set of motion sensor locations as the set of possible goals; it is the goal of a goal-directed agent to move to this location and then go "off the radar". A goal-directed agent, therefore, has two phases of action: an initial phase, in which the goal is chosen, and an execution phase in which the agent tries to move towards the goal.

For the initial phase, the relevant probability distribution is $\Pr[X_{\text{Goal}}|X_{\text{Start}}, t]$, which describes the likelihood that an agent who emerges at a particular location during a particular time of day will have a given goal in mind. With the goal chosen, the agent moves to the execution phase. Here a Markov process based on the probability distribution $\Pr[X_{i+1}|X_i, X_{\text{Goal}}]$ is used to model the motion from one sensor zone to the next. We have chosen to make this distribution time-independent based on the assumption that agents driven by goal-directed behavior do not choose their path based on time of day, but by distance. Making the distribution time-independent avoids sparseness.

Once the agent arrives at X_{Goal} , it is removed from the system.

3.2 Non-Goal-Directed Behavior

For non-goal-directed behavior, a Markov process again evolves the location of the agent over time, but the probability distribution involves the start position instead of a goal position: $\Pr[X_{i+1}|X_i, X_{\text{Start}}, t]$. This distribution is time-dependent, and with a time resolution of 5 minutes smoothed with a 3-segment-wide box filter exhibited artifacts suggesting overfitting; using a 30-minute time resolution with the same filter removed these artifacts.

Since the agents have no set goal in mind, there is no fixed time at which they should disappear from the simulation. Therefore, path completion is integrated directly into the Markov process: for any X_i , X_{Start} , and t , one of the possible values for X_{i+1} is a dummy node representing the completion of the path.

For non-goal-directed behavior, the assumption that agents move at constant speed was removed. Instead, a probability distribution

on the activation duration within a motion sensor zone was given based on X_i and t , with the agent's desired velocity directly derived from the activation duration.

3.3 Comparison and Evaluation

Goal-directed behavior, as expected, produced completely reasonable and realistic paths, with few significantly inefficient paths observed. Non-goal-directed behavior was acceptable but produced some unrealistic behaviors, in particular oscillation between two sensor zones. We tested an order-2 Markov process based on the previous probability distribution and the additional artificially applied posterior that $\Pr[X_{i+1}=x|X_{i-1}=x]$ is low; this removed most of these behaviors (other than a few observed order-3 cycles), but at the cost of artificially suppressing U-turn behaviors which had been observed in the original data.

4. PERFORMANCE

Performance of the STAMM can be accomplished using any existing crowd simulation system that allows for individual behaviors. The number of agent creation events during a time segment is randomly sampled at the beginning of the time segment, and scheduled for times uniformly randomly sampled over the duration of that segment. As an agent passes into a motion sensor zone that had been chosen as its immediate destination, its associated Markov process chooses the next destination zone. In the case of non-goal-directed behavior, the agent's desired velocity is also set as described in section 3.2. Other than during zone transition times, the STAMM system is not used during the simulation.

5. CONCLUSIONS

Goal-directed behaviors learned from the MERL data were reasonable and realistic, with few artifacts, and playback based on the behaviors seemed to correlate well with the original data. Non-goal-directed behaviors were more problematic; while some realistic-looking loitering was observed, so were unrealistic cycles and overly frequent speed inflections. We believe that increasing the realism will require a more thorough analysis and classification of these behaviors, with agents given explicit behavioral intentions.

The time-dependent nature of the STAMM system was apparent. Agents were observed to flood to the elevators and stairwells during the lunch hour, and (with non-goal-directed behavior) to linger in the elevator lobby during high-traffic periods. This model goes well beyond previous work in accurately typifying and simulating the evolution of agent motion behavior based on aggregate spatio-temporal data over the course of a workday.

6. REFERENCES

- [1] Bayazit, O. B., Lien, J.-M., and Amato, N.M. 2002. Roadmap-based flocking for complex environments. In *Computer Graphics and Applications*, 2002. 104-113.
- [2] Braun, A., Musse, S. R., Oliveira, L. P., and Bodmann, B. E. 2003. Modeling Individual Behaviors in Crowd Simulation. In *Proceedings of the 16th international Conference on Computer Animation and Social Agents (CASA 2003)* (May 08 - 09, 2003). CASA. IEEE Computer Society, Washington, DC, 143.

- [3] Helbing, D., Buzna, L., Johansson, A., and Werner, T. 2005. Self-Organized Pedestrian Crowd Dynamics. In *Transportation Science* 39, 1 (Feb. 2005), 1-24.
- [4] Pelechano, N., Allbeck, J. M., and Badler, N. I. 2007. Controlling individual agents in high-density crowd simulation. In *Proceedings of the 2007 ACM Siggraph/Eurographics Symposium on Computer Animation*. Eurographics Association, Aire-la-Ville, Switzerland, 99-108.
- [5] Reynolds, Craig. 1999. Steering Behaviors for Autonomous Characters. *Game Developers Conference 1999*. <http://www.red3d.com/cwr/papers/1999/gdc99steer.html>.
- [6] Shao, W. and Terzopoulos, D. 2005. Autonomous pedestrians. In *Proceedings of the 2005 ACM Siggraph/Eurographics Symposium on Computer Animation*. ACM Press, New York, NY, 19-28.
- [7] Sung, M., Kovar, L., and Gleicher, M. 2005. Fast and accurate goal-directed motion synthesis for crowds. In *Proceedings of the 2005 ACM Siggraph/Eurographics Symposium on Computer Animation*. ACM Press, New York, NY, 291-300.
- [8] Treuille, A., Cooper, S., and Popović, Z. 2006. Continuum crowds. In *ACM SIGGRAPH 2006 Papers*. ACM Press, New York, NY, 1160-1168.