

# Populating Reconstructed Archaeological Sites with Autonomous Virtual Humans

Wei Shao<sup>1</sup> and Demetri Terzopoulos<sup>2</sup>

<sup>1</sup> Media Research Lab, New York University, New York NY, USA

<sup>2</sup> Computer Science Department, University of California, Los Angeles CA, USA

**Abstract.** Significant multidisciplinary efforts combining archaeology and computer science have yielded virtual reconstructions of archaeological sites for visualization. Yet comparatively little attention has been paid to the difficult problem of populating these models, not only to enhance the quality of the visualization, but also to arrive at quantitative computer simulations of the human inhabitants that can help test hypotheses about the possible uses of these sites in ancient times. We introduce an artificial life approach to populating large-scale reconstructions of archaeological sites with virtual humans. Unlike conventional “crowd” models, our comprehensive, detailed models of individual autonomous pedestrians span several modeling levels, including appearance, locomotion, perception, behavior, and cognition. We review our human simulation system and its application to a “modern archaeological” recreation of activity in New York City’s original Pennsylvania Station. We also describe an extension of our system and present its novel application to the visualization of possible human activity in a reconstruction of the Great Temple of ancient Petra in Jordan.

## 1 Introduction

Significant multidisciplinary efforts combining archaeology and computer science have yielded virtual reconstructions of archaeological sites for visualization. State-of-the-art computer graphics enables architectural reconstructions that exhibit impressive geometric and photometric detail. Yet comparatively little attention has been paid to the difficult problem of populating reconstructed models with virtual humans, not only to enhance the quality of the visualization, but also to arrive at quantitative computer simulations of the human inhabitants that can help test hypotheses about the possible uses of these sites in ancient times.

Human simulation has become an increasingly “hot” topic in advanced computer animation. Its application to creating amazing background crowd scenes in feature films has proved its commercial value and this continuously pushes the technology forward. To date, however, the individual characters in such simulated crowds do not have much intelligence and can only serve as background “extras”. By contrast, our artificial life approach [1] to modeling humans spans the modeling of human appearance, locomotion, perception, behavior, and cognition. While our work is innovative in the context of behavioral animation, it

is very different from conventional crowd animation, where one character algorithmically follows another in a manner that is relatively easy to program with a few simple behavioral rules. We are uninterested in crowds *per se*. Rather, the goal of our work has been to contribute a comprehensive, self-animated model of individual human beings that incorporates nontrivial human-like abilities. In particular, we have paid serious attention to deliberative human activities over and above the underlying reactive behavior level.

In particular, we have been developing a decentralized, comprehensive model of pedestrians as autonomous *individuals* capable of a broad variety of activities in large-scale synthetic urban spaces. The pedestrian simulation system that we have implemented enables us to create lengthy animations of numerous pedestrians in large urban environments without manual intervention. We review our pedestrian simulation system in the context of a “modern archaeological” recreation of activity in New York City’s original Pennsylvania Station, which was tragically demolished in 1963. The contribution of this paper, however, is to develop our pedestrian simulator further through the introduction of new behavioral routines, and to report on its novel application to the visualization of possible human activity in a reconstruction of the Great Temple in the ancient Nabataean city of Petra.

The remainder of the paper is organized as follows: We briefly discuss related work in Section 2. Section 3 reviews our simulation system, which comprises an environment model and an autonomous pedestrian model. Section 4 describes a novel application of our simulator in the domain of archaeology. Finally, Section 5 presents our conclusions.

## 2 Related Work

State-of-the-art computer graphics modeling and rendering techniques enable artists, architects, and archaeologists to create or recreate static ancient scenes with high fidelity. There are many practitioners in this field, now called “Virtual Heritage,” which was pioneered in the early 1990s (e.g., [2,3,4]). As these reconstructed virtual sites become increasingly realistic, the lack of mobile human occupants becomes increasingly conspicuous and detracts from the overall realism. Only recently have researchers begun to populate various reconstructed models, especially ancient archaeological sites, with reactive human agents [5,6].<sup>1</sup> Although these characters do not have much intelligence and hardly locomote within their virtual environment, they offer new opportunities for scientific research. Recently, researchers have started to apply computer simulation

---

<sup>1</sup> Note, however, that there are several reasons why virtual humans have not commonly populated virtual models in archaeology, and they do not necessarily have to do with technical issues; for example, archaeologists are generally reluctant to show what they do not have sufficient evidence for, such as humans in action or their clothing (which does not generally survive), and archaeologists generally know very little about spatial behavior, and often not enough about spatial reconstruction (D. Sanders, personal communication).



**Fig. 1.** Original Pennsylvania Station in New York City



**Fig. 2.** Large-scale simulation of human activity in the reconstructed Penn Station

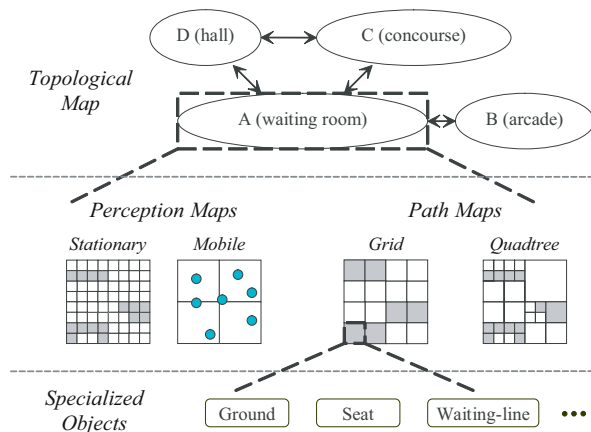
to visualize, test, and refine hypotheses about the usage of architectural structures in ancient times, such as the Colosseum in Rome [7].

Emulating human characters is a highly challenging problem in computer graphics [8]. Prompted by the seminal work of Reynolds [9], behavioral animation has been further developed to simulate artificial animals [10,11,12] and it has given impetus to an entire industry of applications for distributed (multiagent) behavioral systems that are capable of synthesizing flocking, schooling, herding, and other behaviors for lower animals, or in the case of human characters, crowd behavior. Numerous crowd interaction models have been developed [13,14,15,16] and work in this area continues.

Rather than “crowd simulation”, our work focuses on the modeling of intelligent virtual *individuals* that have non-trivial humanlike abilities suitable for urban environments. Our approach is inspired most heavily by the work of Tu and Terzopoulos [10] on artificial animals and by Funge *et al.* [17] on cognitive modeling for intelligent characters that can reason and plan their actions. Recently, we have further developed this comprehensive, artificial life approach and have adopted it for the first time to the case of virtual humans occupying large-scale urban environments [18].

### 3 Autonomous Pedestrian Simulation System

We have developed a realistic simulator of autonomous pedestrians in reconstructed environments. Our initial application of this model was to populate a “modern archaeological” reconstruction of the original Pennsylvania Train Station in New York City (Fig. 1) with hundreds of virtual pedestrians, demonstrating realistic human activity (Fig. 2). This section reviews our simulation system, whose details are presented in [18].



**Fig. 3.** Hierarchical environment model

### 3.1 Virtual Environment Model

As detailed in [19], we represent the virtual environment by a hierarchical collection of maps. Fig. 3 illustrates the environment model, which comprises (i) a topological map that represents the topological structure between different parts of the virtual world. Linked within this map are (ii) perception maps, which provide relevant information to perceptual queries, and (iii) path maps, which enable online path-planning for navigation. Finally, on the lowest level, are (iv) specialized objects that support quick and powerful perceptual queries.

In the *topological map*, nodes correspond to environmental regions and edges represent accessibility between regions. A walkable surface within a region is mapped onto a horizontal plane to enhance the simplicity and efficiency of environmental queries.

The *perception maps* are two types of grid maps that represent stationary environmental objects and mobile objects (usually virtual pedestrians), respectively. On the first type of perception map, a pedestrian emits rasterized “eye rays” to detect nearby static obstacles. On the second type, a sensing fan containing a constant number of neighbor cells are examined by the character to identify nearby pedestrians. Both perceptual algorithms are localized and are independent of the size of the world or the number of contained objects or pedestrians.

The *path maps* of each region consist of a quadtree map and a grid map. Quadtree maps support the execution of several variants of the  $A^*$  graph search algorithm, which are employed to compute quasi-optimal paths to desired goals. The efficient quadtree search algorithm enables pedestrians to do global, long-range path planning in an online fashion. On the other hand, grid maps are employed when detailed, short-range paths are needed. A typical example of their use is when a pedestrian is behind a chair or bench and must navigate around it in order to sit down.

The *specialized objects* at the lowest level of the environment hierarchy are able to provide answers to queries that cannot be handled directly by perception maps. For instance, a bench object keeps track of how many people are seated and where they sit. By querying nearby bench objects, weary virtual pedestrians are able to determine the available seat positions and decide where to sit without further reference to the perception maps. Other types of specialized objects include queues where pedestrians wait in line, purchase points where pedestrians can make a purchase, portals through which pedestrians can enter and exit, etc. These specialized objects make it easy for pedestrians to acquire higher level perceptual information from the virtual world.

### 3.2 Autonomous Pedestrian Model

Analogous to real humans, our synthetic pedestrians are fully autonomous. They perceive the virtual environment around them, analyze environmental situations, behave naturally, and plan their actions over global spatiotemporal scales. To this end, our autonomous human model incorporates appearance, motor, perception, behavior, and cognition sub-models, each of which we will review next.

**Appearance and Locomotion.** As an implementation of the appearance and motor levels of the character, we employ a human animation package called *DI-Guy*, which is commercially available from Boston Dynamics Inc. DI-Guy provides a variety of textured geometric human models together with a set of basic motor skills, such as strolling, walking, jogging, sitting, etc. Emulating the natural appearance and movement of human beings is a highly challenging problem and, not surprisingly, DI-Guy suffers from several limitations, mostly in their kinematic control of human motions. To ameliorate the visual defects, we have customized the motions of DI-Guy characters and have implemented a motor control interface to hide the details of the underlying kinematic layer of DI-Guy from our higher-level behavior modules, enabling the latter to be developed more or less independently.

**Perception.** In a highly dynamic virtual world, an autonomous intelligent character must have a keenly perceptive regard for the external world in order to interact with it effectively. The hierarchical world model is used extensively by pedestrians to perceive their environment, providing not only the raw sensed data (such as those obtained from perception maps), but also higher-level interpretations of perceived situations (such as those obtained from specialized objects) that are at least as important and useful to a pedestrian.

**Behavior.** The purpose of realistic behavioral modeling is to link perception to appropriate actions. We adopt a bottom-up strategy similar to [10], which uses primitive reactive behaviors as building blocks that in turn support more complex motivational behaviors, all controlled by an action selection mechanism.

At the lowest level, we developed six key reactive behavior routines that cover almost all of the obstacle avoidance situations that a pedestrian can encounter.

The first two are for static obstacle avoidance, the next three are for avoiding mobile objects (mostly other pedestrians), and the last one is for avoiding both. Given that a pedestrian possesses a set of motor skills, such as standing still, moving forward, turning in several directions, speeding up and slowing down, etc., these routines are responsible for initiating, terminating, and sequencing the motor skills on a short time scale guided by sensory stimuli and internal percepts. The routines are activated in an optimized sequential manner, giving each the opportunity to alter the currently active motor control command (speed, turning angle, etc.).

While the reactive behaviors enable pedestrians to move around freely, almost always avoiding collisions, navigational behaviors enable them to go where they desire. We developed several such routines—*passageway selection*, *passageway navigation*, *perception guided navigation*, *arrival-at-a-target navigation*, etc.—to address issues, such as the speed and scale of online path planning, the realism of actual paths taken, and pedestrian flow control through and around bottlenecks. Furthermore, to make our pedestrians more interesting, we have augmented their behavior repertoires with a set of non-navigational, motivational routines such as *select an unoccupied seat and sit*, *approach a performance and watch*, *queue at ticketing areas and purchase a ticket*, etc.

An action selection mechanism triggers appropriate behaviors in response to perceived combinations of external situations and internal affective needs represented by the mental state. For example, in a pedestrian whose thirst exceeds a predetermined threshold, behaviors will be triggered, usually through online planning, to locate a vending machine, approach it, queue if necessary, and finally purchase a drink. In case more than one need awaits fulfillment, the most important need ranked by the action selection mechanism receives the highest priority. Once a need is fulfilled, the value of the associated mental state variable decreases asymptotically back to its nominal value. We instantiate different classes of pedestrians suitable for a train station environment, each class having a specialized action selection mechanism, including commuters, tourists, law enforcement officers, and street performers.

**Cognition.** At the highest level of autonomous control, a cognitive model [17] yields a deliberative autonomous human agent capable of applying knowledge to conceive and execute intermediate and long-term plans. A stack memory model enables a pedestrian to “memorize”, “update”, and “forget” chains of goals. The stack couples the deliberative intelligence with the reactive behaviors, enabling a pedestrian to achieve its goals. For example, a commuter can enter the station, with the long-term goal of catching a particular train at a specific time. The cognitive model divides this complex goal into simpler intermediate goals, which may involve navigating to the ticket purchase areas to buy a ticket (which may involve waiting in line), navigating to the concourse area, possibly purchasing a drink if thirsty, sitting down to take a rest if tired, watching a street performance if interested, meeting a friend, and/or navigating to the correct stairs and descending to the proper train platform when the time comes.




Fig. 4. Original Petra site (a) and reconstructed 3D model of Petra Great Temple (b)

## 4 Application to Archaeology

As we have discussed, archaeology is a domain that can benefit from the application of autonomous virtual pedestrians, particularly for the purposes of visualizing human activity in computer reconstructions of archaeological sites and, potentially, for the testing and refinement of archaeological theories of site usage in ancient times.

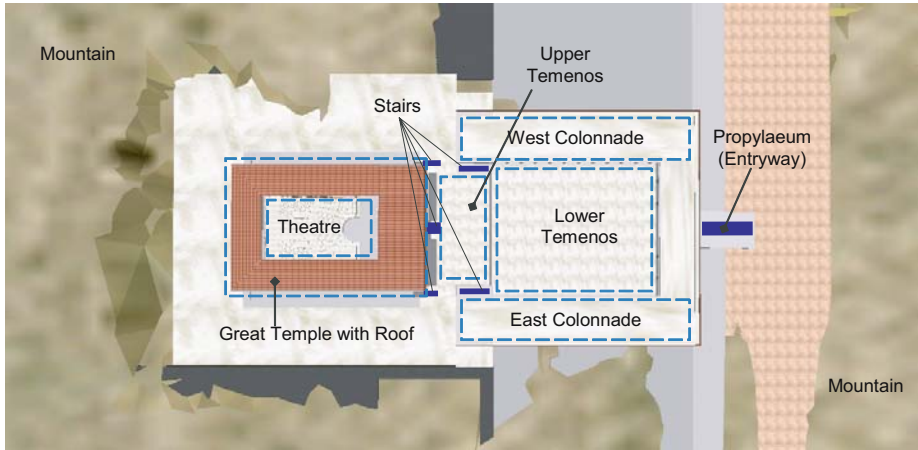
*Petra* (from the Greek word for “rock”) lies in a great rift valley east of Wadi Araba in Jordan about 80 kilometers south of the Dead Sea. As the principal city of ancient Nabataea in its heyday, Petra has a history that can be traced back over 3,000 years. The Great Temple of Petra is a major archaeological and architectural component of central Petra. It is now recognized by archaeologists to have been one of the main buildings of ancient Petra and it is thought to have been dedicated to the most important cult deity of the city. Unfortunately, earthquakes in ancient times demolished the temple and buried it under debris.

### 4.1 Petra Great Temple

The temple site has been explored since the 1890s. A major excavation recently carried out by an archaeological team from Brown University has unearthed many amazing structures and sculptures [20]. With the new findings, archaeologists, computer scientists, and artists have been collaborating to reconstruct the original appearance of the temple in a virtual 3D model [21]. Fig. 4(a) shows a photo of the temple site under excavation and Fig. 4(b) shows a rendered image of the reconstructed 3D model.

As the Temple has sustained severe damage and most of its parts have not yet been found. Many aspects of the Temple, from the layout of the missing structures to details of the sculptures to the purpose and use of the Temple and the theater inside, remain unclear. Archaeologists engage in scholarly speculation while studying the site.

Assimilating the reconstructed temple model within our autonomous pedestrian simulation system, we are able to visualize possible interactions between the



**Fig. 5.** Layout of the Petra Great Temple precinct

temple and its human occupants, thus potentially providing a simulation/visualization tool to archaeologists for use in their speculations about various functional aspects of the Temple.


#### 4.2 Temple Precinct Environmental Model

Fig. 5 shows the configuration of the various parts of the Temple Precinct. Located on the southern citadel hill of Petra, the Temple Precinct covers about 7560 square meters. We divide the entire space into 20 regions, including the Propylaeum (monumental entryway), the Lower Temenos (lower square with two colonnades on the east and west sides), the Upper Temenos (upper square), the Great Temple itself (measuring  $28\text{m} \times 42\text{m}$ ), the theater inside the temple, and the stairs connecting different regions, including those beneath the theater that lead to its auditorium. Objects such as columns and walls are automatically loaded during map initialization. This large model consists of over 410,000 triangles and consumes about 22 MB of memory for its geometry and textures. To prepare the model for use by our autonomous pedestrians, our environmental data structures consume an additional 65 MB of memory.

An issue of interest from the archaeological perspective is to determine how many people can sit in the theater and how efficiently they can enter and exit it. To this end, we developed a new specialized environment object to model the theater inside the temple. According to several user-specified parameters loaded upon initialization, the theater object defines the locations of the stage and the auditorium, including the arrangement of the seats. During simulation, the environment object keeps track of the size of the audience and where each member of the audience is sitting.

In our simulations, we set the theater parameters such that isle space is reserved for people to reach their seats, even when the auditorium is almost full.






(a)



(b)

**Fig. 6.** (a) Layout of the theater. Seats are color-coded from red (best seats) to green (worst seats). The spacing between neighboring seats is defined as  $0.6m$  laterally and  $1.0m$  in the anterior-posterior direction. The setting shown includes 201 seats. (b) Theater filled with 201 seated spectators.



**Fig. 7.** Autonomous pedestrians fill the Petra Great Temple theater. (a) Hundreds of pedestrians enter the Temple Complex through the Propylaeum. (b) On the lower Temenos, the stream of pedestrians bifurcates towards the east and west stairways. (c) Through two arched gates, each pedestrian enters the auditorium, selects a seating area, makes their way to their selected seat and sits down. (d) Spectators seated and attending to a speaker on stage.

Regions from which the stage area is largely occluded are excluded as possible seating areas. Fig. 6(a) illustrates the theater layout in our simulation tests.

### 4.3 New Behavior Routines

To function properly within the theater environment, our autonomous pedestrians required several new navigational and motivational behavior routines designed particularly for an auditorium-like environment; in particular,

1. a seat selection behavior to pick a seat in the auditorium,
2. a navigation behavior to reach the selected seat to sit down, and
3. an exit behavior to leave the auditorium.

The first routine belongs to a category of selection behaviors suitable for pedestrians. These behaviors address the resolution of inter-personal conflicts among

pedestrians for available resources. A pedestrian employs these behaviors to make a selection in order to mitigate the conflict while maximizing its own benefit. More specifically, the pedestrian will autonomously pick a seat based on availability, proximity, and stage-view quality. Fig. 6(a) shows seats color-coded according to their stage-view qualities. In the second routine, a pedestrian will regard other seated pedestrians as static obstacles and will appropriately apply detailed path planning and static obstacle avoidance routines to reach the selected seat. Finally, in the third routine, a pedestrian will simply follow pedestrians ahead in order to exit the auditorium.

#### 4.4 Simulation Example

An animation example that is typical of our simulation tests within the Petra Great Temple environment unfolds as follows: In the beginning, the simulation instantiates hundreds of pedestrians that enter the Temple Complex through the Propylaeum—its grand entryway. On the Lower Temenos, the stream of pedestrians bifurcates to approach the two stairways at the east and west sides. Proceeding to the Upper Temenos, the pedestrians enter the Great Temple via the three small staircases. Once inside the Temple, pedestrians approach the theater entrance stairs located on the east and west sides beneath the theater. Through two arched gates, each pedestrian autonomously enters the auditorium, selects a seating area and navigates to it, selects an empty seat, makes their way to it, and sits down. After all the spectators are seated, a “tour guide” comes to the center stage to enact the delivery of a short lecture about the Great Temple. At its conclusion, the spectators rise and start to evacuate the theater through the two narrow arched gates, which are hardly wide enough to accommodate two pedestrians side by side. Exiting the same way they entered, the pedestrians leave the Temple through the Propylaeum. Sample frames from the animation are shown in Fig. 7.

#### 4.5 Experimental Results

Our simulation experiments reveal that, given the structure of the theater, the fact that each row seems too narrow for the feet of someone seated in a row immediately above to fit comfortably behind the derriere of a person seated in the row below, and with an inter-personal distance (i.e., the distance between the centers of two spectators seated side by side) set to 0.6m laterally and 1.0m in the anterior-posterior direction, the theater can comfortably accommodate 201 spectators in the alternating-row seating arrangement shown in Fig. 6(b).

Given the alternating-row seating arrangement, it requires approximately 7–8 minutes for the audience of 201 to enter and fill the theater and approximately 5 minutes to exit and evacuate the theater. Note that the two arched stairways leading from underneath the theater to its auditorium are the only portals for people to enter and exit; thus, as expected, they become bottlenecks to pedestrian traffic.

## 5 Conclusion

We described a simulation system that incorporates a comprehensive decentralized, autonomous pedestrian model and a hierarchical environmental model. This simulator enables us to deploy a multitude of self-animated virtual pedestrians within large urban environments. We reviewed a virtual reconstruction of the original Pennsylvania Train Station in New York City, and reported on a large ancient outdoor environment, the archaeologically-reconstructed Petra Great Temple in Jordan. In the latter application, which was the novel contribution of this paper, we augmented our autonomous pedestrian model with new behavioral routines specialized for the theater environment within the Temple.

Our simulations suggest that the theater capacity is approximately 200 spectators, who can enter and be seated in approximately 7–8 minutes and exit in approximately 5 minutes. Interestingly, the archaeologist authorities on Petra with whom we have consulted in this project had estimated the theater capacity to be up to three times as many spectators. Our simulations have compelled them to reconsider their estimates. We are also considering seating arrangements alternative to the one shown in Fig. 6, which may require additional behavioral routines to implement more intricate social protocols.

Our simulation system has the ability to produce prodigious quantities of elaborate animation of pedestrians carrying out various individual and group activities suitable to their environment. Applying clever rendering techniques, such as those used in [22], we should be able to improve rendering performance eventually to achieve real-time animation rates. The pedestrian model promises to be of value in the visualization of and speculation about urban social life in the Great Temple precinct of ancient Petra. With this in mind, it will be necessary for us to further augment the behavioral repertoires of the autonomous pedestrians. We also intend to develop a satisfactory set of subconscious, reactive and deliberative head and eye motion behaviors for our virtual pedestrian model in order to make their upper bodies look less rigid during navigation. We will also augment the ability of pedestrians to interact with one another in small groups. This will require better modeling of personal space, territoriality, and familial behaviors.

## Acknowledgements

The research reported herein was supported in part by grants from the Defense Advanced Research Projects Agency (DARPA) of the Department of Defense and from the National Science Foundation (NSF) under Grant No. 0205477. The research has benefitted from collaboration with Brown University, in particular with Dr. Eileen Vote who created the reconstructed model of the Petra Great Temple (Daniel Feliz provided a recent version of the model) and with Professor David Cooper. We appreciate the extensive discussions about the Petra Great Temple and the occupancy of its theater with Dr. Paul C. Zimmerman and Professor Martha Joukowsky. Our thanks also go to Dr. Donald Sanders of

the Institute for the Visualization of History for his advice and feedback on the application of our pedestrian simulator to archaeology and for his comments on a draft. The majority of the research reported herein was done while the authors were at the Media Research Lab of the Courant Institute of Mathematical Sciences at New York University.

## References

1. Terzopoulos, D.: Artificial life for computer graphics. *Communications of the ACM* **42**(8) (1999) 32–42
2. Learning Sites: (2006) [www.learningsites.com](http://www.learningsites.com).
3. Digital Archaeology: (2006) [www.digital-archaeology.com](http://www.digital-archaeology.com).
4. The Institute for the Visualization of History: (2006) [www.vizin.com](http://www.vizin.com).
5. Ciechomski, P., Ulicny, B., Cetre, R., Thalmann, D.: A case study of a virtual audience in a reconstruction of an ancient Roman odeon in Aphrodisias. In: *Proc. 5th International Symposium on Virtual Reality (VAST 2004), Archaeology and Intelligent Cultural Heritage*, Eurographics Association (2004) 9–17
6. Papagiannakis, G., Foni, A., Magnenat-Thalmann, N.: Real-time recreated ceremonies in VR restituted cultural heritage sites. In: *Proc. CIPA 2003 XIXth International Symposium*. (2003)
7. Gutierrez, D., Seron, F., Frischer, B., Abernathy, D.: The Flavian Amphitheater (Colosseum) in Rome: An excellent people-mover? In: *Proc. Computer Applications and Quantitative Methods in Archaeology Conference*. (2005)
8. Badler, N., Phillips, C., Webber, B.: *Simulating Humans: Computer Graphics, Animation, and Control*. Oxford University Press (1993)
9. Reynolds, C.W.: Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics* **21**(4) (1987) 25–34 (*Proc. ACM SIGGRAPH 87 Conf.*).
10. Terzopoulos, D., Tu, X., Grzeszczuk, R.: Artificial fishes: Autonomous locomotion, perception, behavior, and learning in a simulated physical world. *Artificial Life* **1**(4) (1994) 327–351
11. Tomlinson, B., Downie, M., Berlin, M., Gray, J., Lyons, D., Cochran, J., Blumberg, B.: Leashing the alphawolves: Mixing user direction with autonomous emotion in a pack of semi-autonomous virtual characters. In: *Proc. ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA'02)*, ACM Press (2002) 7–14
12. Loyall, A.B., Reilly, W.S.N., Bates, J., Weyhrauch, P.: System for authoring highly interactive, personality-rich interactive characters. In: *Proc. ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA'04)*, ACM Press (2004) 59–68
13. Loscos, C., Marchal, D., Meyer, A.: Intuitive crowd behaviour in dense urban environments using local laws. In: *Theory and Practice of Computer Graphics*, IEEE Computer Society (2003) 122–129
14. Sung, M., Gleicher, M., Cheney, S.: Scalable behaviors for crowd simulation. *Comput. Graph. Forum* **23**(3) (2004) 519–528
15. Ulicny, B., de Heras Ciechomski, P., Thalmann, D.: Crowdbush: Interactive authoring of real-time crowd scenes. In: *Proc. ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA'04)*, ACM Press (2004) 243–252
16. Lamarche, F., Donikian, S.: Crowd of virtual humans: A new approach for real time navigation in complex and structured environments. *Comput. Graph. Forum* **23**(3) (2004) 509–518

17. Funge, J., Tu, X., Terzopoulos, D.: Cognitive modeling: Knowledge, reasoning and planning for intelligent characters. In: Proc. ACM SIGGRAPH 99 Conf. *Computer Graphics* Proceedings, Annual Conference Series, Los Angeles, CA (1999) 29–38
18. Shao, W., Terzopoulos, D.: Autonomous pedestrians. In: Proc. ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA'05), New York, NY, ACM Press (2005) 19–28
19. Shao, W., Terzopoulos, D.: Environmental modeling for autonomous virtual pedestrians. SAE Symposium on Digital Human Modeling for Design and Engineering. (2005) 1–8 (SAE Technical Paper 2005-01-2699).
20. Joukowsky, M.S.: Petra Great Temple Volume I: Brown University Excavations 1993-1997. E.A. Johnson Company, Providence, RI (1998)
21. Vote, E., Acevedo, D., Laidlaw, D.H., Joukowsky, M.: Discovering Petra: Archaeological analysis in VR. *IEEE Computer Graphics and Applications* **22**(5) (2002) 38–50
22. Dobbyn, S., Hamill, J., O'Connor, K., O'Sullivan, C.: Geopostors: A real-time geometry/impostor crowd rendering system. In: Proc. Symposium on Interactive 3D Graphics and Games (SI3D'05), Washington, DC, ACM Press (2005) 95–102