What's Next? The New Era of Autonomous Virtual Humans

Mubbasir Kapadia, Alexander Shoulson, Cory D. Boatright, Pengfei Huang,
Funda Durupinar, and Norman I. Badler

University of Pennsylvania

Abstract. This paper identifies several key limitations in the representation, control, locomotion, and authoring of autonomous virtual humans that must be addressed to enter the new age of interactive virtual world applications. These limitations include simplified particle representations of agents which decouples control and locomotion, the lack of multimodal perception in virtual environments, the need for multiple levels of control granularity, homogeneity in character animation, and monolithic agent architectures which cannot scale to complex multi-agent interactions and global narrative constraints. We present this broad perspective with the objective of providing the stimulus for an exciting new era of virtual human research.

1 Introduction

The simulation of autonomous agents in large crowds has been the subject of widespread attention in many areas ranging from graphics and animation, robotics, urban planning, disaster and security simulation, to the visual effects and gaming industry. It is no surprise that this field has matured at an exponential rate with a plethora of new approaches pushing the frontiers of virtual human simulation.

One of the fundamental challenges which initially needed to be addressed was to efficiently simulate a very large number of agents that can sufficiently capture the macroscopic phenomena of crowds for interactive applications. Dictated by limitations in computational resources, seminal contributions [28] made certain simplifying assumptions in agent representation and control, many of which are prevalent even today.

The new era of visual computing provides at our disposal cutting edge hardware and dedicated graphical processing units making these assumptions and their resulting limitations largely outmoded. The challenge is no longer simulating large crowds. Instead, the next generation of interactive virtual world (IVW) applications demand functional, purposeful, heterogeneous autonomous virtual humans, that exhibit rich, believable interactions with their environment and other agents, with the far-reaching goal of complete immersion for the end users.

In this paper, we identify several underlying assumptions that are holding us back from entering into the new age of interactive virtual world applications.

M. Kallmann and K. Bekris (Eds.): MIG 2012, LNCS 7660, pp. 170–181, 2012.

Our aim is *not* to provide all the answers, but to provoke a strong discussion among researchers and practitioners alike, and provide the stimulus for an exciting era of virtual human research.

2 Traditional Autonomous Agent Model

There is a large breadth of research in simulating autonomous virtual humans which spans the spectrum of modeling the macroscopic phenomena of crowds, to developing monolithic agent architectures that model agent attention, memory, and inter-agent communication. We refer the readers to comprehensive reviews [26,36] of the field, and focus on medium to large scale agent interactions, while stressing on modeling agents as autonomous entities. These works address challenges in many areas, as illustrated in Figure 1, and described below.

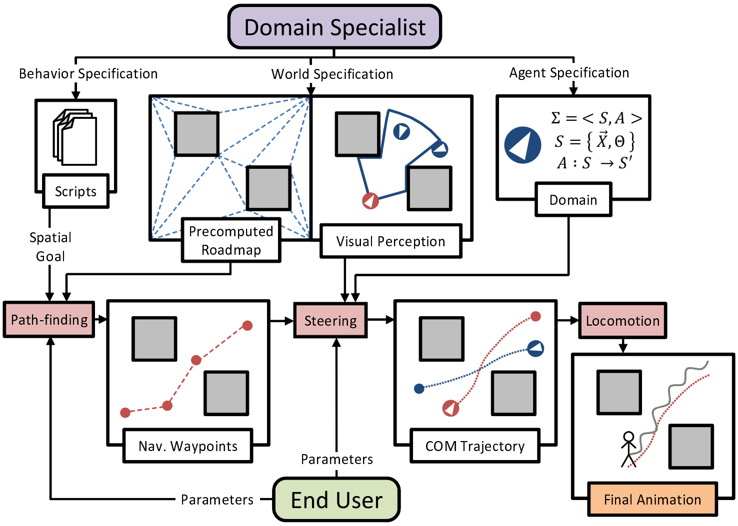

Fig. 1. The traditional approach to modeling agents as autonomous entities

Agent Representation. Domain experts are tasked with defining the configuration of the environment and the agents that populate it. The state and action space, or problem domain, of agents is greatly simplified to ensure that agent simulation is not a computational bottleneck. The state of an individual agent is represented as a point mass with a collision radius. The orientation of the disc may be modeled to indicate the facing direction of the agent. The action space of the agent has 2 or 3 degrees of freedom, with displacement along the forward vector and turning about the center. Lateral displacement (e.g., side-stepping) is rarely modeled. The possible action choices are pre-conditioned to satisfy constraints on maximum walking speed, linear and angular acceleration to conform to real humans. Agent objectives are desired spatial locations.

Pathfinding. Pathfinding is the process of finding a collision-free path from the current position of an agent to its target location, usually taking into account the static aspects of the environment. Global navigation approaches [9, 34, 35] precompute a roadmap of the global environment which is used for making efficient navigation queries, but generally regard the environment to be static.

Steering. Steering is the act of navigating an agent along the planned path while avoiding static as well as dynamic threats, relying on computationally efficient solutions for collision avoidance, by making locally optimal control decisions. Reactive approaches [16, 28] use a one-step lookahead to avoid collisions with most imminent threats. Predictive approaches [2, 11, 24, 32] approximate the trajectories of neighboring agents in choosing collision-free velocities.

Locomotion. Locomotion [13] animates a fully articulated virtual human to follow the trajectory output by steering, taking into account the abilities (affordances) of the human model.

Behavior Authoring. A behavior authoring framework provides the tools for users to specify and automatically generate complicated interactions between multiple agents in dynamic environments. Scripted approaches [19] are used to describe carefully orchestrated action sequences. Rule-based systems [27] describe behaviors as condition-action mappings while cognitive approaches [40] use decision networks to model knowledge and action selection in virtual agents. Authoring behaviors is usually the task of domain specialists, and end users are usually provided with parameters to tweak the simulation.

The virtual populace and the environment is visualized to immerse end users into these interactive virtual worlds. The pathfinding, steering, and locomotion layers produce homogenous agents, while heterogeneity amongst agents is usually injected at the authoring and visualization layer by generating variety in the objectives and appearance of agents.

3 Agent Representation

> *A particle-based agent representation cannot capture nuanced interactions that humans exhibit in confined and crowded situations.*

A particle-based agent representation enables efficient state-based queries, such as collision checks, and greatly reduces the branching factor of the action space to enable efficient control. However, a particle cannot capture all the capabilities of human locomotion such as side-stepping, careful foot placement, and even nuanced upper body motions. This reduced problem domain prevents the solution of many challenging scenarios such as oncoming agents in a narrow passage-way, where an agent may have to carefully step out of the way and re-orient its shoulders, allowing the other agent collision-free passage. The inverse of this is also true, where constraints imposed on a particle cannot capture bipedal locomotion constraints. This is evident in the artifacts observed when a superimposed virtual human tries to follow a particle trajectory.

Furthermore, a circular collision radius is a conservative estimate of the boundary of a real human and prohibits optimal packing density in crowded situations. The bounding volume of real humans constantly changes and deforms as they maneuver themselves around different obstacle and agent configurations. The representation of agent groups as deformable meshes [8] simulates passive dynamics between the crowd and environment, but does not capture the physical interactions between agents. The work in [33] represents the torso, shoulders, and feet as separate circular colliders, allowing tighter packing density for egress simulations, while footstep-based control facilitates richer steering capabilities like side-stepping.

> *Multi-modal perception in autonomous virtual humans is needed for more realistic behavior.*

Traditional agent models are equipped with visual perception queries such as line-of-sight ray casts and foveal cone intersections to *see* the world around it in order to make an informed control decision. Auditory perception is sorely lacking in current agent models.

Sound propagates differently from light, providing a richer set of perceptual options for an agent, including localization of an unseen event and the recognition or possible mis-identification of a sound signal. For example, a person may not be seen because of visual occlusion, but the persons footsteps or voice may still be heard. In a cocktail party, someone might not be able to see everyone else, but she can still tune to various conversations based on the topics that resonate with her interests. For crowd simulation in games, an acoustic perception sense can provide additional useful behavior constraints including possible goals (sound sources), avoidance regions (noisy areas), animated reactions (to a loud noise), or even navigation cues (such as hearing someone approaching around a blind corner).

Geometric sound propagation methods [3] model sound as rays and are unable to capture acoustic properties like diffraction. Cellular automata based approaches [14] can demonstrate diffraction, reflection, and refraction of sound, but don't scale well to large environments. A sound propagation model is needed that can sufficiently capture acoustic sound properties, while meeting real-time efficiency constraints.

Additionally, agents must be able to perceive and recognize (correctly or incorrectly depending on sound degradation) sound signals to trigger appropriate behavioral responses. Humans can efficiently distinguish between multiple convolved sound signals [1], which is difficult to compute. However, we can leverage work in human auditory perception [7] which concludes that amplitude, duration, and pitch are the principal components of environmental sound classification. These findings can help us identify simplified representations of sound signals which can be efficiently propagated and perceived in complex, dynamic virtual environments.

4 Navigation and Steering

> *Complex, dynamic environments require a tighter coupling between steering and navigation for more robust control of autonomous agents.*

Current pathfinding approaches [9, 34, 35] precompute roadmaps to efficiently query paths for static environments. Steering is responsible for following the path while factoring in dynamic world events. The next generation of interactive applications requires high-fidelity navigation in non-deterministic dynamic virtual worlds. The environments and agents may be constantly changing due to unpredictable events (e.g., other agents, dynamic obstacles, and human input), potentially invalidating computed paths. For example, a crowd of agents blocking an exit would require choosing an alternate route, requiring a re-plan.

There is also a need for a two-way coupling between steering and navigation where global paths can be efficiently repaired by accounting for the current dynamic world state, including other agents. Truly dynamic environments require the use of planning approaches [18] which can quickly return solutions with bounds on sub-optimality, and iteratively refine the solution while accommodating dynamic changes in the environment.

> *Different levels of control granularity are required depending on the configuration of obstacles and agents in the environment.*

Different situations require different granularity of control. An open environment with no agents and static obstacles requires only coarse-grained control while cluttered dynamic environments require fine-grained character control with carefully planned decisions that have spatial and temporal precision. Furthermore, some situations (e.g., potential deadlocks) may require explicit coordination between multiple agents.

The problem domain of interacting autonomous agents in dynamic environments is extremely high-dimensional and continuous, with infinite ways to interact with objects and other agents. Having a rich action set with a system that makes intelligent action choices facilitates robust, intelligent virtual characters, but at the expense of interactivity and scalability. Greatly simplifying the problem domain yields interactive virtual worlds with hundreds and thousands of agents that exhibit simple behavior.

One possible solution is to define a set of problem domains which spans the spectrum in terms of complexity of agent representation and fidelity of agent control. Multiple domains provide different trade-offs in performance and control granularity, requiring a control scheme that can efficiently work in multiple domains by choosing the appropriate domain depending on the current situation, and transitioning across multiple domains while leveraging solutions across domains.

The use of synthetic training data to learn steering policies for different subsets of scenarios can generalize well across the space of all possible agent interactions, while providing a solution that is scalable, yet efficient.

As virtual worlds grow both in complexity and interactivity, it is quickly becoming intractable to predict *a priori* the possible scenarios an agent will encounter during simulation. The next generation of IVW's need algorithms that are scalable not only in agent count, but experience as well.

Data-driven solutions are a natural fit to expanding an algorithm to handle new situations. There are two main types of data-driven crowd simulation, trained models [22] and database queries [17,37]. Both of these approaches suffer from two main limitations. First, trained models fit a single, monolithic model which must generalize over all the training data, requiring accuracy in some scenarios to be sacrificed for others. The lossless alternative of database querying leads to unwieldy amounts of data which become impractical to store and search. Second, these approaches use trajectories of real humans as training data, which poses logistical challenges for data acquisition and control over what will be observed.

These limitations can be overcome through the use of synthetic data and by clustering the scenarios into groups based on similarity. We need to identify subspaces in the set of all possible scenarios [12] for which a single trained model can fit well and learn to classify these subspaces at runtime to choose the appropriate model. These more focused models mitigate the underfitting problem as the virtual world's diversity increases. Once identified, these subspaces need training data. A synthetic source in the form of a steering oracle algorithm can be sampled as needed with full control and without the logistical problems of collecting real-world data.

5 Locomotion

There is need for a tighter, two-way coupling between simulation and locomotion.

There exists a uni-directional communication interface between steering and locomotion where steering outputs only a force or velocity vector to an animation system, without necessarily conforming to the constraints and capabilities of human-like movement. Steering trajectories may have discontinuous velocities, oscillations, awkward orientations, or may try to move a character unnaturally. A vector-based interface does not have enough information to indicate certain maneuvers such as side-stepping versus reorienting the torso, stepping backwards versus turning around, planting a foot to change momentum quickly, or carefully placing steps in exact locations.

This simplistic approach lacks control over how a character should truly animate and often results in visual artifacts when an animated virtual human follows a trajectory output by steering. The locomotion system is not always able to accurately follow the steering trajectory, causing deviations and possibly even collisions, and the steering layer may not always be aware of the abilities of the locomotion system.

To offset these limitations, we need to identify a more appropriate interface between locomotion and steering – such that there is a tighter, two-way coupling between these two layers. An animated virtual human should be able to accurately follow decisions output by the control layer, which should be constantly aware of the affordances and constraints imposed on the locomotion system.

> *Crowd heterogeneity can be greatly enhanced by adding expressivity and variety to the full-body motion of virtual characters.*

Heterogeneity in agents and agent interactions is a key challenge in the realization of immersive virtual environments. In addition to focusing on variations in the physical appearance of characters [21], existing approaches [5,6] characterize individual differences through psychological states by integrating emotion and personality models at the steering or behavior layers. The human body has been shown to be as expressive as the face, conveying information about a persons personality and social role. Body cues are the first to be perceived, especially at a distance when individuals are approaching to initiate social interaction [38]. Hence, expressive body motion manifests itself effectively in crowd simulation applications, where multiple agents are usually observed from a distance.

The animation of virtual characters that express their personalities through their body motions requires a careful understanding of the contribution of different personality factors to human motion. First, we need to characterize human personality as a set of mutually orthogonal factors to observe their effect on human motion in isolation, and in unison. For this purpose, the Five Factor Model [39] can be exploited as a comprehensive personality model which captures the different facets of personality. Second, it is necessary to recognize the key features of human motion to capture motion intention because understanding how human motion alters due to personality differences in the space of joint angles is most likely intractable. Third, we need to derive a mapping between the personality factors and this feature space using analytical or computational methods, which can be used for motion synthesis. The work in [4] procedurally generates synthetic gestures by parameterizing Laban Movement Analysis components [15] and provides a strong foundation for animating heterogeneous crowds.

6 Behavior Authoring

> *A monolithic autonomous agent model cannot sufficiently capture the space of all possible agent behaviors, and interactions between agents and objects.*

As we diversify the actions that virtual characters perform, we encounter a combinatorial explosion of those characters' potential interactions. If we model characters as monolithic autonomous agents, coordinating activities becomes a challenge of passing messages and sharing state between thousands of virtual peers, especially in prolonged interactions where actors take on roles. In such a system, introducing a new factor to the environment requires updating each character type to account for the newly introduced mechanic. This amplifies the design effort required to introduce new content or author interesting behaviors involving multiple agents.

In order to produce meaningful simulations, we aim to make our system accessible to content authors of various degrees of training. In an ideal IVW system, we want an untrained content creator to be able to design rich behaviors involving multiple characters interacting with one another and the environment. To facilitate this, we advocate an event-centric approach [30] that can temporarily operate multiple characters as if they were limbs of the same entity. When an actor is not involved in an interaction, it is *autonomous* with individual goals and objectives. Actors participating in an *event* suspend autonomy and are exclusively controlled by the event to ensure coordination between actors.

This event-centric model allows an author to design interactions in terms of what he or she wants the characters to do in a given situation. In contrast, a monolithic agent model would require the author to coordinate events by passing stimuli between two or more agents and determining the appropriate response emergently.

> *Next generation IVWs need to satisfy global narrative constraints while still providing an "open-world" to facilitate authoring story-driven simulations without severely limiting user agency.*

Simulations exist for an authored purpose, whether designed for training, teaching, or entertainment. In most cases, the simulation is driven by a sequence of events connected by a narrative structure. Training simulations operate on specific scenarios that the human participants must navigate, while video games evoke the escapist pleasure of being placed in a world driven by some conflict that must be resolved. What brings these environments to life in the user's mind is the underlying narrative thread where actions have consequences and both the user and the virtual characters are involved in a causal relationship with the rules of the system. Here, story is the crucial context that makes the characters compelling, the virtual world immersive, and the user's actions meaningful.

Most narrative systems tell a story with a centralized director that sends commands to the agents involved. The work in [10] uses an automated planner to satisfy global narrative constraints but lacks interactivity. Façade [20] features two conversational agents in a social drama, while systems like MIST [25] focus on numerous homogeneous agents. Since the user's actions cannot always be predicted, allowing the user unmitigated control within the story requires a system capable of adapting to unexpected states.

Both the system managing the story and the agents enacting it must be robust enough to recover if derailed by the user. One representation for this type of behavior is a story specified as a "soft plot", where key events selected by an author must occur in some sequence over the course of the simulation, but the system itself has multiple paths through. This does not fully prevent the user from blocking the plot, but allows more flexibility in adhering to the story and enables the system to afford the user more agency.

7 Conclusion

In this paper, we identify several key simplifications in the representation, control, animation, and behavior authoring of autonomous virtual humans that are holding us back from entering the new era of interactive virtual world applications. These are certain important research questions that must be addressed if we are to break away from these limiting barriers, and provide functional, purposeful, heterogeneous autonomous virtual humans (Figure 2).

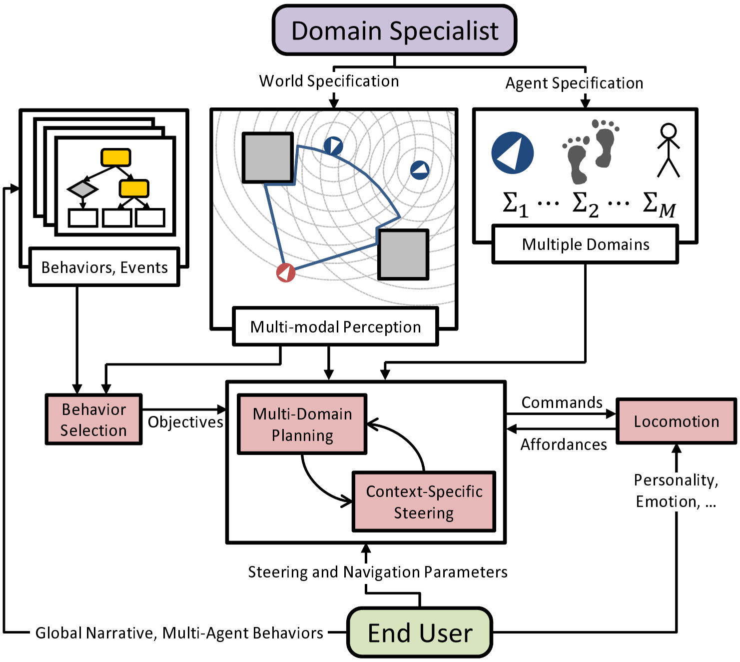

Fig. 2. A forward looking model for simulating autonomous virtual humans

There exists a large gap between a simple particle-based representation and a fully articulated virtual human, which limits agent steering, and provides insufficient information to allow an animated virtual human to follow the steering trajectory. Particle-based control does not conform to a bipedal agent model and this gross simplification results in controllers with conservative bounding volumes, insufficient degrees of freedom, as well as relaxed constraints on movement. We pose the need for a more complex agent representation that provides a high degree of control fidelity, while still meeting efficiency requirements. One possible solution is to define multiple domains of control and heuristically choose the appropriate domain, depending upon the current situation.

The one-way communication between navigation, steering, and locomotion severely limits agent control. For example, there may be locomotion constraints imposed by the scenario author which may prevent agents to steer with a high degree of dexterity. Dynamic changes in the environment constantly invalidate global paths, requiring constant plan refinement.

Visual perception alone cannot sufficiently capture the rich sensory modalities that humans are equipped with, thereby greatly limiting agent perception, and adversely effecting agent behavior. The ability to *hear* is an important sensory modality that needs to be integrated into existing agent architectures. Challenges include developing an efficient sound propagation framework that sufficiently captures the properties of audio and facilitates auditory perception.

The use of data-driven solutions to tackle difficult control problems is receiving a lot of traction due to its ability to provide generalized, scalable control in highly-dimensional, non-linear spaces where heuristic strategies cannot sufficiently capture the space, and optimization is generally intractable. Nevertheless, the application of data-driven methods [17, 22, 37] to crowds has been limited so far, mainly due to the difficulty in data acquisition and the extreme non-linearity of the space of all possible agent and obstacle interactions. To offset these limitations, we pose the use of synthetic training data and learning policies in sub-spaces with contextually similar scenarios for robust control.

Heterogeneity in multi-agent simulations can be greatly enhanced by adding expressivity and variety to the full-body motion of virtual characters. The challenge remains to provide an expressive parametrization of character locomotion that allows users to specify the personality, mood, and emotion of a populace. As we improve at controlling a character's body, the scope of behaviors these actors can perform greatly expands. With the right systems, we can create characters that can involve the user in meaningful scenarios capable of adapting to unanticipated input. Effectively designing for this space, however, requires the tools to enable untrained authors to create content from a narrative and personality standpoint, without confronting the complexity introduced by message passing, shared memory state, and emergent behavior.

> *Open source solutions for standard practices in steering, navigation, and character locomotion are essential to address difficult problems in simulating functional, purposeful, heterogeneous autonomous virtual humans.*

Perhaps, the most important ingredient to answering many of these challenging research questions is the provision of easily available solutions for many of the standard practices in virtual human simulation. Recast [23] and Tripath [9] provide libraries for generating navigation meshes of arbitrarily complex environments, facilitating efficient pathfinding. SteerSuite [31] is an open-source platform for developing and evaluating steering behaviors. SmartBody [29] is a modular, controller-based full-body character animation system.

We hope that the wider perspective we have taken with these suggested research directions will open the way to new topics and accelerated progress in simulating meaningful groups of virtual human characters.

Acknowledgements. The research reported in this document was performed in connection with Contract Number W911NF-10-2-0016 with the U.S. Army Research Laboratory. The views and conclusions contained in this document are those of the authors and should not be interpreted as presenting the official policies or position, either expressed or implied, of the U.S. Army Research Laboratory, or the U.S. Government unless so designated by other authorized documents. Citation of manufacturers or trade names does not constitute an official endorsement or approval of the use thereof. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation heron.

References

1. Bee, M., Micheyl, C.: The cocktail party problem: What is it? how can it be solved? and why should animal behaviorists study it? Journal of Comparative Psychology 122(3), 235 (2008)
2. van den Berg, J., Lin, M.C., Manocha, D.: Reciprocal velocity obstacles for real-time multi-agent navigation. In: ICRA, pp. 1928–1935. IEEE (2008)
3. Chandak, A., Lauterbach, C., Taylor, M., Ren, Z., Manocha, D.: Ad-frustum: Adaptive frustum tracing for interactive sound propagation. IEEE TVCG 14(6), 1707–1722 (2008)
4. Chi, D., Costa, M., Zhao, L., Badler, N.: The emote model for effort and shape. In: SIGGRAPH 2000, pp. 173–182. ACM (2000)
5. Durupinar, F., Pelechano, N., Allbeck, J., Gudukbay, U., Badler, N.I.: How the ocean personality model affects the perception of crowds. IEEE CGA 31(3), 22–31 (2011)
6. Guy, S.J., Kim, S., Lin, M.C., Manocha, D.: Simulating heterogeneous crowd behaviors using personality trait theory. In: SCA 2011, pp. 43–52. ACM (2011)
7. Gygi, B., Kidd, G., Watson, C.: Similarity and categorization of environmental sounds. Attention, Perception, & Psychophysics 69(6), 839–855 (2007)
8. Henry, J., Shum, H.P.H., Komura, T.: Environment-aware real-time crowd control. In: SCA 2012. ACM, New York (2012)
9. Kallmann, M.: Shortest paths with arbitrary clearance from navigation meshes. In: SCA 2010, pp. 159–168. Eurographics (2010)
10. Kapadia, M., Singh, S., Reinman, G., Faloutsos, P.: A behavior-authoring framework for multiactor simulations. IEEE CGA 31(6), 45–55 (2011)
11. Kapadia, M., Singh, S., Hewlett, W., Faloutsos, P.: Egocentric affordance fields in pedestrian steering. In: I3D, pp. 215–223. ACM (2009)
12. Kapadia, M., Wang, M., Singh, S., Reinman, G., Faloutsos, P.: Scenario space: characterizing coverage, quality, and failure of steering algorithms. In: SCA 2011, pp. 53–62 (2011)
13. Kovar, L., Gleicher, M., Pighin, F.: Motion graphs. In: SIGGRAPH 2002, pp. 473–482. ACM, New York (2002)
14. Kristiansen, U.: Viggen: Computational methods in acoustics (2010)
15. Laban, R.: The Mastery of Movement. Northcote House Publishers Ltd. (1971)
16. Lamarche, F., Donikian, S.: Crowd of virtual humans: a new approach for real time navigation in complex and structured environments. In: CGF 23 (2004)
17. Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by Example. Computer Graphics Forum 26(3), 655–664 (2007)

18. Likhachev, M., Ferguson, D.I., Gordon, G.J., Stentz, A., Thrun, S.: Anytime dynamic a*: An anytime, replanning algorithm. In: ICAPS, pp. 262–271 (2005)
19. Loyall, A.B., Bates, J., Mitchell, T.: Believable agents: Building interactive personalities. Tech. rep. (1997)
20. Mateas, M., Stern, A.: Facade: An experiment in building a fully-realized interactive drama. In: Game Developers Conference, GDC 2003 (2003)
21. McDonnell, R., Larkin, M., Dobbyn, S., Collins, S., O'Sullivan, C.: Clone attack! perception of crowd variety. In: SIGGRAPH 2008, pp. 26:1–26:8. ACM (2008)
22. Metoyer, R.A., Hodgins, J.K.: Reactive pedestrian path following from examples. In: CASA, vol. 20, pp. 149–156 (November 2003)
23. Mononen, M.: Recast: Navigation-mesh construction toolset for games (2009), `http://code.google.com/p/recastnavigation/`
24. Paris, S., Pettré, J., Donikian, S.: Pedestrian reactive navigation for crowd simulation: a predictive approach. In: Eurographics, vol. 26, pp. 665–674 (2007)
25. Paul, R., Charles, D., McNeill, M., McSherry, D.: MIST: An Interactive Storytelling System with Variable Character Behavior. In: Aylett, R., Lim, M.Y., Louchart, S., Petta, P., Riedl, M. (eds.) ICIDS 2010. LNCS, vol. 6432, pp. 4–15. Springer, Heidelberg (2010)
26. Pelechano, N., Allbeck, J.M., Badler, N.I.: Virtual Crowds: Methods, Simulation, and Control. Synthesis Lectures on Computer Graphics and Animation (2008)
27. Perlin, K., Goldberg, A.: Improv: a system for scripting interactive actors in virtual worlds. In: SIGGRAPH 1996, pp. 205–216. ACM, New York (1996)
28. Reynolds, C.W.: Flocks, herds and schools: A distributed behavioral model. In: SIGGRAPH 1987, pp. 25–34. ACM (1987)
29. Shapiro, A.: Building a Character Animation System. In: Allbeck, J.M., Faloutsos, P. (eds.) MIG 2011. LNCS, vol. 7060, pp. 98–109. Springer, Heidelberg (2011)
30. Shoulson, A., Badler, N.I.: Event-Centric Control for Background Agents. In: Si, M., Thue, D., André, E., Lester, J.C., Tanenbaum, J., Zammitto, V. (eds.) ICIDS 2011. LNCS, vol. 7069, pp. 193–198. Springer, Heidelberg (2011)
31. Singh, S., Kapadia, M., Faloutsos, P., Reinman, G.: An Open Framework for Developing, Evaluating, and Sharing Steering Algorithms. In: Egges, A., Geraerts, R., Overmars, M. (eds.) MIG 2009. LNCS, vol. 5884, pp. 158–169. Springer, Heidelberg (2009)
32. Singh, S., Kapadia, M., Hewlett, B., Reinman, G., Faloutsos, P.: A modular framework for adaptive agent-based steering. In: I3D, pp. 141–150. ACM (2011)
33. Singh, S., Kapadia, M., Reinman, G., Faloutsos, P.: Footstep navigation for dynamic crowds. CAVW 22(2-3), 151–158 (2011)
34. Sud, A., Gayle, R., Andersen, E., Guy, S., Lin, M., Manocha, D.: Real-time navigation of independent agents using adaptive roadmaps. In: VRST 2007, pp. 99–106 (2007)
35. Sung, M., Kovar, L., Gleicher, M.: Fast and accurate goal-directed motion synthesis for crowds. In: SCA 2005, pp. 291–300. ACM, New York (2005)
36. Thalmann, D.: Crowd simulation. In: Wiley Encyclopedia of Computer Science and Engineering (2008)
37. Torrens, P., Li, X., Griffin, W.A.: Building Agent-Based Walking Models by Machine-Learning on Diverse Databases of Space-Time Trajectory Samples. Transactions in GIS 15, 67–94 (2011)
38. Vinayagamoorthy, V., Gillies, M., Steed, A., Tanguy, E., Pan, X., Loscos, C., Slater, M.: Building Expression into Virtual Characters. In: Eurographics STAR (2006)
39. Wiggins, J.S.: The Five-Factor Model of Personality: Theoretical Perspectives. The Guilford Press (1996)
40. Yu, Q., Terzopoulos, D.: A decision network framework for the behavioral animation of virtual humans. In: SCA 2007, pp. 119–128. Eurographics (2007)