AAAI Technical Report WS-11-18

Selecting Agents for Narrative Roles

Alexander Shoulson, Daniel Garcia, and Norman I. Badler
Department of Computer and Information Science
University of Pennsylvania
Philadelphia, PA 19104-6389, USA
{*shoulson, garciada, badler*}@*seas.upenn.edu*

Abstract

We present ongoing work on a system that accommodates player agency in a digital narrative with an external plot. We focus on key events that should occur in that storyline for dramatic effect, but do not explicitly specify the characters that should fill the roles needed for those events. Instead, we define them abstractly, with characteristics that the selected characters should have (including previous events they should have completed for eligibility), and rely on a Director construct to populate those roles from agents in the selection pool that fit those criteria. Agents begin as largely homogeneous, primordial entities that accumulate data and narrative value from the events in which they participate. This creates an environment that differentiates characters by the actions they perform, conferring worth onto characters that become important to the player based on their direct involvement in the plot. The focus, then, is on defining *a priori* the *what* of the narrative, while leaving it to the Director construct to decide at runtime exactly *who* among a distributed pool of agents carries it out.

Introduction

In real-time simulation or digital storytelling environments, interactive virtual agents play an important role in establishing player immersion and presence within the simulated setting. These agents can appear in small groups with high levels of visual and behavioral detail, or in much larger crowds with less individual dimensionality and focus. In the former scenario, high-fidelity agents can select actions based on complex motivations and planned narrative goals for complex behaviors, while in larger numbers agent behavior is usually based off of the emergent effects of path-finding and locomotion (Stocker et al. 2010). Designing and executing sophisticated behavior for large crowds poses an interesting challenge from both a computational and authorship standpoint, as having large numbers of autonomous agents represents a demanding and chaotic system that is difficult to control on a macroscopic level (Durupinar et al. 2011).

A strong base of work exists on animated large crowds of homogeneous pedestrians for simulating large-scale urban areas (Farenc et al. 1999; Haciomeroglu, Laycock, and Day 2008; Narain et al. 2009; Shao and Terzopoulos 2005), but this approach is only equipped to display agents walking from place to place according to expected population density distributions and human crowding behavior. While path-finding and locomotion are important for displaying a convincing crowd of people, we believe that more sophisticated behaviors are needed for truly immersive populated environments. A restaurant and a shopping mall, for instance, would not appear truly consistent without visitors eating and shopping, respectively. Displaying agents as simply walking through these purposeful environments is not sufficient. Furthermore, they do not develop the agents presented in their environments with the level of personification needed to drive a character narrative.

While agents could be made more sophisticated, driven by individual impulses, large-scale high fidelity agent simulation becomes demanding with respect to resources. Adaptive level of detail for agent simulation (Brom, Šerý, and Poch 2007) can help mitigate the computational cost, but does not significantly reduce the effort required to author the individualized behaviors of heterogeneous characters. Games like The Sims 3 (Aaby 2011) successfully simulate large numbers of agents with their own needs and motivations, but to tell a story with an externally-specified plot would require character goals to be more narrative-driven than human needs such as hunger, tiredness, or the drive for success or acquisition of resources. These systems lend themselves well to environments where a story emerges from the characters' interactions (i.e., the player can superimpose a narrative on the sequence of events as presented), but they do not as easily support digital narratives where a dramatic arc is at least partially specified *a priori* and then maintained by the program.

Instead, we are developing a system that differentiates agents by the actions they perform and narrative roles they fill. We centralize narrative development in a controlling entity called a Director, similar to those of (Bates 1992) and (Magerko et al. 2004). Our Director system, however, is responsible for populating an environment with "primordial", undifferentiated agents and then selecting some of them to perform tasks important to the storyline encoded and provided by the author. We codify agent actions in a manner similar to the Smart Event paradigm, each defined by the author so as to comprise a lexicon of transitory agent behaviors. As agents participate in events, they accumulate data and can satisfy the prerequisites for subsequent, dependent events. An agent in a park having bought a hot dog from a vendor cart as part of a ”purchase” Smart Event would now be eligible to be selected for a second Smart Event which is designed for an agent to sit on a bench and have them eat a meal. This causality can be applied for short-term enforcement of consistency, or to develop long-term narrative arcs over the course of a digital story. This accumulation of information is monotonic, so that as agents become important to the Director by virtue of their participation in events, we retain the information that differentiates them (aside from transitory qualities, such as being hungry or holding a suitcase, that expire when no longer valid). This follows traditional progression: the scene is first established on the whole, and then characters are introduced as they perform the actions that make them significant.

The question then becomes how we select agents to fill these narrative roles and acquire importance from the perspective of the story. We present our ongoing work on this topic in the rest of this paper. First, we discuss the event-driven architecture that enables the higher-level decisions upon which our system relies, and then motivate the factors involved in populating those events. We end with a discussion and road map for the future.

Event-Driven Behavior

Since we primarily differentiate agents by their actions (as perceived by the player), our system centralizes behavior to the events in which agents particpiate. Agents possess their own behavior logic, but on their own agents do not perform very sophisticated actions (restricted to, for example, wandering around, checking a wristwatch, reading a book, and so on). Complex behaviors (and multi-agent interaction) is managed by a lexicon of events, each with its own data structure representing the decision process and control logic dictating what an agent should do at each step while involved in that event. For instance, to display two agents haggling over the purchase of an item, one would involve those two agents in a bartering event, which would first instruct both agents to display greeting animations, and then in turn tell each agent to play the animations of gestures indicating negotation. The control of each agent would be passed to the event itself, which has full authority over each agent for the duration.

Centralizing behavior logic to the specific events that occur simplifies authoring in two ways. For one, the author can fully control each agent directly in the event, as opposed to being forced to design agents so that they respond to a chain of messages passed by one to the other. Secondly, once one or more agents are selected to perform the event and the event is instantiated over them, no further intervention is needed on the part of the Director construct to steer that event – the progression of the event is entirely managed by the event itself.

The Director construct, then, is responsible for selecting an appropriate event from the event lexicon, identifying which available agents (agents that are not currently involved with other, more important events) should perform it, and then dispatching the instantiated event to them. The agents perform the event in a location specified by the Director, and when the events end, the agents return to their baseline idle behaviors and return to the pool of available agents to be selected for another event. The Director does not need to be involved in the event once it has been instantiated, though events can be overridden by one another based on a priority system.

There are several ways for deciding *which* event to execute. For background characters maintaining the ambience of a scene, we can select events based on a probability distribution of desired events for a location. A park, for instance, should have a number of people sitting on the grass, others reading books on benches or riding bicycles – all in various malleable proportion. For an overarching story, a planner can determine events to dispatch to characters based on character goals and factors such as the dramatic intensity of a sequence of events. What is important is that the lexicon of events is the same between these two decision systems. Background characters driven by probability distributions receive the same data structures to execute their events as the pivotal plot characters for which events are being planned. As a result, agents can move between these two groups: an agent that at one point was reading a newspaper in the background of a scene for ambiance can suddenly be “promoted” to interacting with the player in a story-driving fashion as a pivotal plot character without any interchange of the agent’s control structure. The only fundamental difference between unsubstantiated background agents and principal characters is that the latter group has performed actions that personify them in the eyes of the player, and so we retain data from those actions, when performed, for causality.

Towards Adaptive Narrative

We primarily encounter two types of interactive narrative, each with different allowances for player agency. Our definitions are guided by, but differ slightly from those of (El-Nasr 2007) in order to underscore a specific distinction. The first type is what we refer to as *emergent* narrative, in which the story progresses in a manner similar to a soap opera – characters interact in small-scale interpersonal dramas, but no global story arc is maintained. Emergent narratives are realized in systems like UNIVERSE (Lebowitz 1984) and game series like The Sims (Aaby 2011), where character assets and drives perpetually feed a generative plot with no external objectives. In emergent narrative, player agency is intuitively accounted for because the characters operate reflexively on local motivations. If a player interferes with a character, then that character’s assets or drives will be accordingly modified and the character will react based on its current state and goals. There is no need to worry about a deviation from some story arc because the effects of the player interference become part of that story itself.

The second type of interactive narrative is what we term *driven* narrative, in which a designer specifies an overarching plot to which the narrative at least loosely adheres in order to tell an intended story. Here, player agency is usually considered to be directly at odds with the constraints of the story we wish to tell (or “make”, as in (Knoller 2010)). We must give a player some manner of control so that she feels invested in the progression and outcome of the narrative, but cannot allow the player to violate the preconditions necessary to advance the story. Fortunately, an event-centric authoring approach gives us the methodology to approach this problem. We can design narrative based on the types of incidents we want to occur, while allowing the system to specify at runtime the manner in which the event occurs and which characters are involved. This resource flexibility enables us to permit more player agency, since we can exploit the fluidity of assets available to the Director.

While selecting which event we instantiate at a given time is important, we instead focus here on on selecting the agents to fulfill each of the roles in the event. Picking agents selectively allows us to enforce microscopic agent-specific consistency, but gives the Director more difficulty in populating the events it wishes to invoke. Conversely, picking agents without discrimination allows the Director to populate more events at the cost of reasonable agent causality. We further define that scale as the "1-∃-∀" scale, as illustrated in Fig. 1. On the bottom axis, the "1" represents selecting a unique, individual agent in the environment, such as an agent with a particular name or identifier, to fill a role. If we focus intently on this agent, we can model all necessary state attributes and ensure that the agent acts only in a way consistent with its previous set of actions. This leans towards a full-simulation approach, but by making any event in the narrative explicitly dependent on this agent, we lose the ability to adapt to this agent's removal from the environment or occupation with another important event.

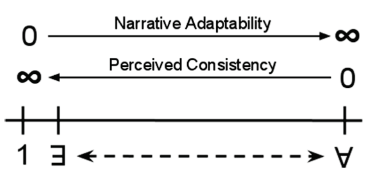

Figure 1: "1-∃-∀" scale of narrative adaptability by agent selection paradigm.

The $\exists - \forall$ range represents a sliding scale of selectivity. On the less adaptive side, the $\exists$ represents a degree of restrictiveness such that there may be only one or very few agents in the environment. More restrictiveness means that we can afford to only pick agents that could most plausibly be involved in the intended event (say, selecting an agent to display drinking coffee from among a pool of only those agents we've displayed purchasing coffee), and so the player is less likely to notice inconsistencies. However, this could cause situations where no agents fit the expected criteria and the intended event cannot be executed, thus impeding on the role of the Group Coordinator. On the surface this appears to only prevent the execution of a single event, but the consequences could be more severe if the event and its timing were crucial to plot progression. As we slide in event selectivity towards unrestricted admission of agents ("∀"), we experience the opposite dilemma. Unsurprisingly, this delineates a tradeoff. It is left to the scenario designer to choose a point along the axis depending on context but with the knowledge of this compromise.

Of course, event selection can become more specific over time. We begin with very few predesignated principal characters, as it is possible for agents to begin as undifferentiated background characters and later acquire importance by performing pivotal events. In this way, we attribute data to characters as they acquire value to the narrative. If an essential plot character is somehow removed from the environment and we still need to accomplish an event critical to the plot, it may be necessary to "promote" a background character to fulfill an equivalent narrative role. This is possible because events are interchangeably populated with types of characters, rather than specific agents in the environment. In this situation, we monotonically increase the amount of data stored about that agent. That is, as agents become more vital to the plot (we keep a history of the key narrative events in which the agent participated), we attribute more information to them. Once we do so, however, that information must persist so that we do not later contradict it. Because we are limited in the amount of data we can feasibly store about each agent, it is important to very carefully select the agents worth promoting with persistent data in this fashion.

As agents acquire importance, it may be necessary to retroactively develop them as well. There are methods for developing this information – (Lucie Kučerová and Kadlec 2010) discusses planning the long-term history (major life events) of an agent based on its current status in the simulation, while Alibi Generation (Sunshine-Hill and Badler 2010) focuses more on the short-term. In the latter case, however, we already maintain to some extent a short history of recent dispatched events for an agent, which can be used for justifying later actions performed.

The process for agent selection, then, proceeds as follows. For an agent to fill a narrative role in an event, we define the prerequisites for that role to be filled (these may include having completed other events). When we instantiate that event, we search for any agent that meets the prerequisites and weight the search by other factors, such as proximity and visibility. If a suitable match is found, that agent is selected for that role in the event, the event executes, and we keep a history of any residual data stored in the agent from that event (including its successful completion). If no match is found, we select the closest candiate to the profile. If that closest match needs to execute a chain of events before being eligible for the current role (these are broad-scale, in the sense that "getting hired" precedes "getting a promotion"), we begin that agent along that course of prerequisite events. This relies on the use of a planner, but is handled more implicitly in the event dispatch itself. If no agent suitably matches the profile to begin with, we cancel that event and see if we can find an event with similar dramatic effect. This searching process leads to more robust narratives, since we can use the pool of available agents to better accommodate unforseen situations while still telling a story. Nowhere do we specify a particular individual for these narrative arcs. The author specifies the story based on *what* occurs, and the Director construct selects *who* realizes it. We believe this enables narrative where control is balanced between the player and the storyline in a world where characters can be interrupted and distracted by the player at any time.

Conclusions

While this represents ongoing work, the key idea is to define narrative incidents and constrain the simulation so that they must occur, while leaving flexibility in the manner in which they are executed. This is primarily accomplished through dynamically populating the roles of those incidents (i.e., the firefighters, bystanders, and victims of a ”building fire” event) based on the player-malleable state of the situation. Player agency, by its nature, precludes an author anticipating every possible sequence of events the player might initiate. Therefore, we must develop adaptive narratives that build upon the player's participation while still telling, to some degree, an author-specified plot.

We do not predetermine the agents intended to participate in the events we author, but still need agents to perform each role. As such, we define selection criteria based on the context of the event. For memorable roles that the player will remember, these are more stringent restrictions (the participating agent must have completed some other events, and have certain characteristics such as the proper 'costume'). For other roles, such as bystanders, or a waiter in a restaurant, we can permissively select agents convenient to the context. This distinction is delineated by the "1-∃-∀" scale we define, which illustrates the balance between consistency and robustness.

The system begins with a specified series of events, and a distribution of available "primordial" agents, and selects and applies appropriate agents for the roles of the specified behaviors. In performing the specified events, the agents become differentiated, both in the eyes of the player and in the data they accumulate as part of a plot-relevant event history (which then renders them eligible for some subsequent events). This way the plot can progress around the player – the particular fashion in which it proceeds influenced by the actions of the player and the agents she involves in those actions.

Acknowledgements

The research reported in this document/presentation 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/presentation 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

Aaby, C. 2011. The Sims - Part 1: How Playing With Dolls can Create Amazing Stories. *Gamasutra*.

Bates, J. 1992. Virtual Reality, Art, and Entertainment. In *PRESENCE: Teleoperators and Virtual Environments*, 133–138. MIT Press.

Brom, C.; Šerý, O.; and Poch, T. 2007. Simulation Level of Detail for Virtual Humans. In *Proceedings of the 7th international conference on Intelligent Virtual Agents*, IVA '07, 1–14. Berlin, Heidelberg: Springer-Verlag.

Durupinar, F.; Pelechano, N.; Allbeck, J.; Güdükbay, U.; and Badler, N. 2011. How the Ocean Personality Model Affects the Perception of Crowds. *Computer Graphics and Applications, IEEE* 31(3):22 –31.

El-Nasr, M. S. 2007. Interaction, Narrative, and Drama Creating an Adaptive Interactive Narrative Using Performance Arts Theories. *Interaction Studies* 8(2).

Farenc, N.; Schweiss, E.; Kallmann, M.; Aune, O.; Boulic, R.; and Thalmann, D. 1999. A Paradigm for Controlling Virtual Humans in Urban Environment Simulations. *Applied Artificial Intelligence* 14:69–91.

Haciomeroglu, M.; Laycock, R. G.; and Day, A. M. 2008. Dynamically Populating Large Urban Environments with Ambient Virtual Humans. *Computer Animation and Virtual Worlds* 19(3-4):307–317.

Knoller, N. 2010. Agency and the Art of Interactive Digital Storytelling. In *Proceedings of the Third Joint Conference on Interactive Digital Storytelling*, ICIDS'10, 264–267. Berlin, Heidelberg: Springer-Verlag.

Lebowitz, M. 1984. Creating Characters in a Story-Telling Universe. *Poetics* 13(3):171–194.

Lucie Kučerová, L. C. B., and Kadlec, R. 2010. Towards Planning the History of a Virtual Agent. In *The ICAPS 2010 Workshop on Planning in Games*.

Magerko, B.; Laird, J. E.; Assanie, M.; Kerfoot, A.; and Stokes, D. 2004. AI Characters and Directors for Interactive Computer Games. In *Proceedings of the 2004 Innovative Applications of Artificial Intelligence Conference*, 877–883. AAAI Press.

Narain, R.; Golas, A.; Curtis, S.; and Lin, M. C. 2009. Aggregate Dynamics for Dense Crowd Simulation. In *ACM SIGGRAPH Asia 2009 papers*, SIGGRAPH Asia '09, 122:1–122:8. New York, NY, USA: ACM.

Shao, W., and Terzopoulos, D. 2005. Autonomous pedestrians. In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, SCA '05, 19–28. New York, NY, USA: ACM.

Stocker, C.; Sun, L.; Huang, P.; Qin, W.; Allbeck, J. M.; and Badler, N. I. 2010. Smart events and primed agents. In *Proceedings of the 10th international conference on Intelligent virtual agents*, IVA'10, 15–27. Berlin, Heidelberg: Springer-Verlag.

Sunshine-Hill, B., and Badler, N. 2010. Perceptually Realistic Behavior Through Alibi Generation. In *Proceedings of the 6th Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE)*. AAAI Press.