Title: "Categorical, Narrative, and Hybrid Behavior Generation in the GENIE Environment for Interactive Narrative Virtual Worlds"

Authors:

Craig A. Lindley
Starlab NV
Boulevard St.-Michel 47
B-1040 Brussels
Belgium
craig@starlab.net

Frank Nack
Centrum voor Wisekunde en Informatica (CWI)
Amsterdam, Netherlands
Frank.Nack@cwi.nl

Contact: Craig A. Lindley, Starlab NV, Boulevard St.-Michel 47, B-1040 Brussels, Belgium
craig@starlab.net

Topic: Content production: multimedia authoring tools, VR

Synopsis/Extended Abstract

The GEneric Narrative Interaction Environment (GENIE) is a conceptual architecture for the exploration of interactive narrative generation in virtual worlds. Specific research demonstrators implement different aspects of the GENIE architecture, with the longer term goal of implementing the full system. A particular feature of GENIE is the adoption of narrative, categorical, and hybrid sequencing strategies for model-based presentation generation in virtual worlds. Narrative generation is concerned with the preservation of spatio-temporal continuity of action within the agent behaviors and events in the virtual world. Categorical sequencing is concerned with the generation of associative sequences of events based upon patterns of semantic similarity and dissimilarity. Lindley and Nack (2000) have described nine additional hybrid strategies that combine narrative and categorical sequence generation.

Within the GENIE system, action or behavior generation in virtual worlds is regarded in terms of four levels of control: the world level, the object level, the agent level, and the discursive agent level. These levels and their components are identified at a level of authorship that lies between the represented world that a user of the media system experiences and the implementation architecture; this is the level at which *a priori* classes of represented entity can manifest different behaviors within a range of possible behaviors that is constrained by the behavior generation representations provided for a particular production. Behaviors within the architecture may be inherited from higher levels to lower levels of the model.

Each level is distinguished by a scale of encapsulation, and by a set of primitive components to which it may refer. For instance, world level scripts are not encapsulated within any kind of representational structure at a level smaller than the world level, while agent-level scripts are encapsulated within a representation of an agent class or instance. Each level may be realized by the interpretation of a more or less detailed script or system of scripts, by deliberative planning to achieve high level goals for that level of control, or by the execution of reaction rules in response to user actions or external state changes.

The distinction between planning approaches and behavioral approaches is based upon the extent to which action generation relies upon the creation of a model of (some part of) the world and its contents, as opposed to the creation of behaviors. A purely behavioral control system can only refer to behavioral goals, internal state representations, simulated sensory inputs, and behavioral outputs at a given level of the hierarchy. Alternatively, a more traditional planning-based approach, referred to by Brooks (1999) as the sense-model-plan-act, or SMPA, approach, relies upon the construction of a model of the external (virtual) environment as a basis for action determination, and may maintain a description of the current state of progress with respect to a complex hierarchical or extended linear action plan. Behaviors may also be scripted at higher levels, where the scripting language is closer to a natural language (eg. Goldberg, 1997). Hence there are different levels of authorship for behavioral control; a set of agents may be authored at a planning and/or behavioral level, and then the behaviors created for those agents provide the semantics for higher level scripts. These techniques can be combined, for instance, by using a script interpreter to read a

high level or natural language expression that is converted into one or more high level goals that are then passed by the interpreter to a deliberative planner, or to instantiate the parameters of, or activate, a set of reaction rules.

The world level of the hierarchy represents the definition of high level, global control of the virtual environment, including message transmission between and among objects and agents.

Beneath the world level, more variability can be achieved by scripting object and agent behaviors. Objects are generally passive, and most object behavior can be implemented by a generic physics or mechanics engine that does not require detailed scripting. Object scripting in this case may be limited to basic existential statements (that an object of a particular kind exists at a particular time and place under specific conditions), and parameterisation of object characteristics, such as size, mass, temperature and hardness.

Agents can be regarded as objects having autonomous behaviors. Autonomous behaviors are those that are not simply the result of generic physical laws within the virtual world, but originate within the encapsulated control architecture of the agent itself. Autonomous behaviors may range from simple reactive stimulus-response rules characteristic of behavioral systems, through complex sequences of behavior requiring higher level, deliberative AI-planning techniques, to behaviors expressed in a more natural language that must then be interpreted to form sequences of deliberative rules and/or behavioral directives.

Discursive agents are those having some behaviors that serve purely communicative or expressive functions. The distinction between discursive and non-discursive agents separates explicitly modeled and controlled communicative behavior from behavior represented instrumentally that has only implicit communicative functions. Discursive behaviors also potentially range from reactive responses (eg. using ELIZA-style language generation rules, or Expressivator-style gesture rules, Sengers, 1998) to the autonomous generation of stories (Bickmore and Cassell, 1999) and the ability to use a conversational input as a goal for complex planning (Cavazza et al, 1999). More advanced discursive capabilities might include the ability of a virtual agent to interact with another agent or the user through a more complex discourse structure, such as an extended conversation or explanations for problem solving, commentary, post-mortems or tutoring (eg. André and Rist, 2000, Lester et al, 2000, Cassell et al, 2000).

In effect, the distinctions between 11 different sequencing techniques, the four levels of encapsulation at which those techniques are applied (world, object, agent, and discursive agent), and the three levels of abstraction of behavior generation instructions (natural language like descriptions, declarative behavioral goals, and reaction rules) form a three dimensional matrix of behavior generation techniques. This matrix contains a total of one hundred and thirty two different behavior generation strategies that could in principle be used independently or in any set of combinations. Since any particular strategy can be either present or absent, the number of possible combinations of strategies is $n = 2^{132} - 1$. Each combination of strategies defines a different potential authoring environment for interactive cinema, and hence a different interactive cinematic form from the author's perspective (although the resulting presentations may not show a corresponding variability).

The analytical framework embodied in the GENIE architecture provides a basis for the classification and comparison of virtual world authoring systems. Ongoing research is addressing the incremental development of the GENIE authoring environment to provide this full set of interactive cinematic forms at a very generic level. Creating interesting interactive productions, however, will be knowledge intensive; in addition to the basic shell functions provided by the GENIE software, an increasing library of particular agent descriptions and behaviors will provide a foundation for generating productions of increasing interest and complexity. A key component of any such library will be world and/or agent classes that incorporate generic models of phenomena such as personality and subtext, allowing production authors to model the personalities of virtual characters directly and explicitly, together with specifying the evolving subtext of interactively generated action.