



INSTITUTE OF COMPUTER SCIENCE

Academy of Sciences of the Czech Republic

Martin Holeňa

Pod Vodárenskou věží 2, 182 07 Praha 8, phone: +420 266052921, fax: +420 286585789, e-mail: martin@cs.cas.cz
web: www.cs.cas.cz/~martin

Review of the PhD thesis

Controlling Virtual People

by Jakub Gemrot

The reviewed thesis investigates the specification of decision making as a part of the behavior of intelligent virtual agents (IVA), focusing on the productivity of behavior coding. Due to the crucial role that IVA aka non-player characters have in the application area of video games, this topic is very important and timely.

The thesis is, apart from the Introduction and Conclusion, structured into 11 chapters. Chapter 1 defines the concept of IVA and explains the difference between IVA in video games and non-game applications.

Chapter 2 formalizes IVA behavior modeling, dividing the model into reasoning and decision making. It introduces a concept very important for the remainder of the thesis – behavior plan development cost, which allows evaluating how economically, in terms of language primitives needed, a particular behavior plan implements a particular decision-making pattern.

The ideas and results of the thesis that I appreciate most are contained in Chapters 3 to 5. The first of them is the idea of a decision space and results concerning its relationships to a deterministic Turing machine, presented in Chapter 3. I like the parallel to state space search, as well as the inherently hierarchical structure of the decision space, which will later be used for code reusability.

Chapter 4 introduces another important notion – behavior path. With it available, it is already possible to separate reasoning from decision-making, but also to model transitions, switching and IVA sub-behavior.

Chapter 5 analyses the two main existing approaches to IVA behavior specification – scripting and hierarchical finite-state machines. What I appreciate most in this chapter is an attempt to formalize stack-based computation, leading to the concepts of a passive stack, typical for standard languages, and active stack, corresponding to hierarchical finite-state machines.

Chapter 6 is a bridge between the theory behind decision spaces, and the definition of a new language for AVI behavior specification. It introduces an extensible tree-like reactive engine,

which could serve as a basis for interpreting the proposed language. From the subsequent chapters, however, it is not clear whether it really plays this role.

Chapter 7 presents the main application result of the thesis – the definition of a new behavior design language (BDL), which generalizes both scripting and behavior specification by means of hierarchical finite-state machines. BDL combines the decision-making advantages of both mentioned approaches and allows mixing them in a behavior specification most economically.

Chapter 8 shows that BDL, in addition, generalizes also the action-selection mechanisms in three existing agent languages from academia – AgentSpeak(L), GOAL and SPOSH, as well as behavior trees, used in the game development industry. Due to the detailed treatment of those languages, this is the most extensive chapter of the thesis (94 pages; in fact, I have reviewed several theses that were shorter than this single chapter :)).

Chapter 9 recalls several decision-making patterns that are difficult to model in the existing languages, and shows how they can be modelled in BDL. In particular, delaying behavior switch-out and some transition behaviors are addressed.

In Chapter 10, a BDL example is elaborated.

Finally, Chapter 11 reports several controlled experiments comparing some of the languages usable for IVA behavior specification. They are compared from the point of view of productivity through economical development, i.e., the feature that has been the driving force of the research reported in the thesis.

In each of the chapters 3–9 and 11, there are some new contributions of the author. From the point of view of the whole thesis, I consider its main results to be the following:

1. The concepts of a decision space and of decision points, as a theoretical foundation for IVA decision making. Their distinctive features are that they do not enforce any decision making schema upon the development and structuring of behavioral plans and that they support the encapsulation of decision-making code.
2. The emphasis on productivity through economical development.
3. A general mechanism for performing switches between behaviors based on the propagation of signals through running behavior plans.
4. A new behavior design language, which is highly versatile, generalizing all commonly used action-selection mechanisms, and in this way differing from existing languages for IVA behavior specification.

From the point of view of novelty, I appreciate most the first result. All of them are, however, valuable contributions to the area of IVA development.

Due to the already mentioned role of virtual agents in video games, the results of the thesis, and most directly the BDL, have a great application potential in that area. Actually, it was this area that motivated the research reported in the thesis – Jakub Gemrot is one of the main authors of the Pogamut platform.

As for formal features of the thesis, I appreciate the care that the author devoted to making it easy to read and to navigate. Every chapter starts with a kind of a mini-abstract and ends with

a summary of its content. The text is accompanied with plenty of elucidating figures, and all concepts that could be somehow formalized, are explained in 72 specific definitions. Unfortunately, the author did not always succeed to keep the definitions consistent. However, the number of inconsistencies, as well as the number of confusing or grammatically incorrect formulations is very low in view of the extent of the thesis.

The comprehensiveness of the thesis enabled the author to present rich and manifold information. In this context, I miss only two marginal points mentioned below. It would be nice if the author could briefly comment on them during the thesis defence.

1. Video games are particular software systems, and the design of any software system usually starts with some kind of conceptual or at least logical modeling, which if relevant includes also behavior modeling. Therefore, I wonder how suitable or unsuitable (and why) for modeling IVA behavior are the traditional models for software systems modeling, in particular:
 - various variants of Petri nets, both deterministic and stochastic ones;
 - Markov finite state models, both with discrete time (Markov chains) and with continuous time;
 - behavior variants of the entity-relationship model.
2. When presenting the results of experiments in Chapter 11, an unnecessarily restricted information has been provided, namely only mean values and standard deviations. To assess the importance of the obtained results, also their statistical significance should be reported.

To sum up, the submitted thesis clearly shows that Jakub Gemrot is able to perform creative, independent and systematic research of a large extent, definitely fulfilling the high standards required for the PhD degree at the Charles University.

Prague, September 1, 2017

Martin Holeňa