Review report of the PhD thesis presented by Martin Černý

Reducing Complexity of AI in Open-World Games by Combining Search-based and Reactive Techniques

As the title of the thesis already suggests Martin Černý has explored ways to make AI techniques usable in Open-World Games (OWG). The complexity in OWG comes from two dimensions: there can be many agents with relative simple behavior that have to be coordinated and there can be agents that have complex behavior that has to be coordinated. Although both combinations lead to high complexity in the OWG, these complexities are of different types and also have to be dealt with in different ways.
The big contribution of this research is that it carefully analyzes the types of complexities in chapter 2 and then sets out several possible solutions to handle the complexity in different cases in chapter 5, 6 and 7.
In chapter 5 behavioral objects are introduced. They are interesting as they build on the idea of smart objects, but are more general. They do not encompass just physical objects in the environment, but can also be areas and quests. The general idea is that the entities in the environment carry some affordances that the characters can subscribe to and that execute some local behavior fitting in that situation. By extending the idea of smart objects to smart areas and quests there is a very natural and efficient way to coordinate the behavior of characters in an area such as a pub, a city square, etc. rather than the characters having to perform a complex distributed coordination the area takes over and solicits characters to fill in parts of its interaction script.
The conceptual ideas have also been carefully implemented where the candidate has shown how elements of OOP can be used, but also where OWG require some different solution. The environment in which the behavior objects are implemented uses behavior trees to program the behavior of the characters. There is a nice way in which the behavior objects can take over limited control of the character behavior by infusing some subtree in the behavior tree of the character. This leads to the following question:

> *Could you describe a more general method of implementing this type of control shift between the character and the behavior object such that if the OWG does not make use of behavior trees it still will work?*

In chapter 6 the argument is made that although behavior entities in all their forms are very useful to program all kinds of local intelligent behavior and coordinated interactions there is also a need for a global connection between behavior and content. Rather than having a smart area that covers the whole environment (which would not be very efficient) there is a need to be able to specify some constraints that tie the behavior of several characters in a specific situation. It seems that constraint satisfaction techniques are very suitable for this situation. Due to the specific context in which they are used some efficiency can be gained and the techniques are shown to work as well in OWG.

> *Although it is argued that the constraint satisfaction techniques are good to provide a global perspective and also why they are different from the smart areas, it seems the difference between the two are not so black and white as they are sketched. Is it possible to give a precise definition or heuristic when CSP should be used and when smart areas should be used?*

Finally in chapter 7 some techniques are discussed to handle cases in OWG with large rule complexity. The main problem in this case is that decisions are based on many parameters and have to be taken in real-time. Thus if a goal-based approach would be taken it is too slow although it solves some of the problems of the large parameter space. Thus in this chapter the idea of adversarial search is used to predict behavior of the environment given some own behavior. This can often be done quicker than actual planning from first principles. In the chapter a careful mechanism is developed to implement such an adversarial search in a fighting scenario. The limitation of the fighting scenario gives all kinds of heuristics that make the search very efficient. Some more simplifications are made in order to make the search efficient enough for real-time use. My question here becomes:

> *The presented technique is very specific for a certain type of fighting situation within a very simple environment. It is not clear which steps can and should be taken to generalize this technique for other situations. Thus how general is this applicable?*

In order to conclude my review I first want to indicate that I really appreciate this type of applied research. It tries to create a bridge between the theoretical work done in AI and its application in actual games. It means that many practical considerations are playing a role besides the theoretical ones. Some people see this as less “scientific” work than the purely theoretical work. However, I think it is of prime importance to test theory in practice and feed back the problems and successes in order to drive new theoretical research as well. In my opinion this thesis is a good example of applied research that is highlighting both the success and areas where more research is needed.

The presented thesis is well written and provides a state-of-the-art review as well as a rigorous presentation of the designed algorithms. The presented work has been in parts presented to the artificial intelligence and gaming scientific community, has been positively accepted and cited. Martin Černý demonstrated his capability to perform independent research work resulting in original results.

I recommend

accepting his thesis and awarding him a title of a doctor of philosophy (Ph.D).

Dr. Frank Dignum
Department of Information and Computing Sciences
Faculty of Science
Utrecht University
The Netherlands