Prospective Storytelling Agents

Gonçalo Lopes and Luís Moniz Pereira

Centro de Inteligência Artificial - CENTRIA
Universidade Nova de Lisboa, 2829-516 Caparica, Portugal
goncaloclopes@gmail.com
lmp@di.fct.unl.pt

Abstract. Prospective Logic Programming is a declarative framework supporting the specification of autonomous agents capable of anticipating and reasoning about hypothetical future scenaria. This capability for prediction is essential for proactive agents working with partial information in dynamically changing environments. The present work explores the use of state-of-the-art declarative non-monotonic reasoning in the field of interactive storytelling and emergent narratives and how it is possible to build an integrated architecture for embedding these reasoning techniques in the simulation of embodied agents in virtual three-dimensional worlds. A concrete graphics supported application prototype was engineered, in order to enact the story of a princess saved by a robot imbued with moral reasoning.

1 Introduction

Prospective Logic Programming (PLP) builds upon grounded theories of abduction and non-monotonic reasoning, and laid the prior foundations for combined developments in the fields of logic programming, AI, and cognitive science, so as to support an agent's prospection of its future and attending computational morality reasoning [3, 4].

Interactive storytelling and emergent narratives [1, 2] focuses on the automatic generation of non-linear dramatic storylines involving characters embodied in rich, dynamic worlds. In the sequel we show how PLP handled the specific requirements and challenges of the application mentioned in the Abstract.

2 Application Description

In order to test the basic PLP framework and the integration of a virtual environment for interactive storytelling, a simplified scenario was developed. In this fantasy setting, an archetypal princess is held in a castle awaiting rescue. The unlikely hero is an advanced robot, imbued with a set of declarative rules for decision making and moral reasoning. As the robot is asked to save the princess in distress, he is confronted with an ordeal. The path to the castle is blocked by a river, crossed by two bridges. Standing guard at each of the bridges are minions of the wizard which originally imprisoned the princess. In order to rescue the princess, he will have to defeat one of the minions to proceed.

Prospective reasoning is the combination of pre-preference hypothetical scenario generation into the future plus post-preference choices taking into account the imagined consequences of each preferred scenario. By reasoning backwards from this goal, the agent generates three possible hypothetical scenaria for action. Either it crosses one of the bridges, or it does not cross the river at all, thus negating satisfaction of the rescue goal. In order to derive the consequences for each scenario, the agent has to reason forwards from each available hypothesis. As soon as these consequences are known, meta-reasoning techniques can be applied to prefer amongst the partial scenaria [3].

In this case, a specific utility value was associated with each available hypothesis, quantifying the likelihood that the robot survives. Other things being equal, the robot will prefer scenaria which both maximize the likelihood of survival and the satisfaction of a greater number of goals. Also, as long as the likelihood of survival does not fall below a specified threshold, the robot will prefer scenaria which satisfy a greater number of goals. By relying on previous knowledge about the probability of defeating different types of minion, the robot computes the utility of surviving the crossing of each of the bridges. This knowledge store could be built incrementally by using past experiences in a number of ways which fall outside the scope of the current work (e.g. learning or case-based reasoning).

If the likelihood of survival is very low, the robot will choose not to rescue the princess. At this point, another complication was introduced: the possibility of endowing the robot with moral rules. The encoding of moral reasoning using declarative PLP techniques has been previously addressed in [4]. In this way, it is possible to encode a moral constraint that all princesses in distress must be saved. From this point on, the moral constraint will defeat the scenario on which the goal to rescue the princess is negated, regardless of survival utility.

Other moral constraints have also been explored. For instance, when choosing between facing a giant spider or a regular guard, the robot will choose the regular guard, as it presents a scenario with a higher survival utility. But the princess can then become angry that the robot has killed a man, and enforce a moral constraint that no humans can be killed. If however, in the next reiteration of the setting, both minions are humans, the constraint of always saving princesses in distress will conflict with that of never killing humans. By default, the robot will reason to choose the scenario that minimizes the number of violated constraints while maximizing the number of satisfied goals, so the princess will still be saved. Other options available were exploited to solve these conflicts, including the use of preference rules or meta-reasoning techniques.

3 Architecture

The system[1] exhibits a blend of imperative and declarative techniques. For graphics rendering, the Ogre3D library was used. Reasoning was implemented in the ACORDA framework [3] on top of XSB Prolog and Smodels. Additional implementation details provided in [3] and [4]. Integration was performed using the C# programming language and the .NET framework, by means of a wrapper[2] around XSB's external API.

[1] `http://centria.di.fct.unl.pt/~lmp/software/MoralRobot.zip` demo:`http://centria.di.fct.unl.pt/~lmp/publications/slides/padl10/moral_robot.avi`

[2] `http://sourceforge.net/projects/xsbdotnet/`

Several procedural routines were implemented to handle basic agent locomotion and perception. When necessary, the reasoning system was queried for goal satisfaction. During the reasoning process itself, queries to the procedural perception modules can be performed in order to probe the current state of the world. What is more, the robot is able to query the user for moral advice when finding itself in a conundrum. Also, actions chosen by declarative reasoning were procedurally simulated in the virtual environment. In this way, a full bidirectional coupling was achieved between simulation and declarative reasoning.

4 Conclusions and Future Work

We believe the present work, even in its prototypical working state, is a significant step forward in the application of state-of-the-art declarative reasoning techniques to the automatic generation of dramatic narratives in dynamic virtual environments. The coupling of sensors and actuators to a declarative non-monotonic reasoning model can easily ensure that changes in the virtual environment perceived by the agent can be incorporated in its knowledge base. PLP has been developed precisely for allowing knowledge to be constantly revised and updated, so this will not present any impediment as often happens when contradictory observations are updated to monotonic reasoning systems.

This robustness to novelty makes the system particularly useful for interactive storytelling techniques. Non-linear stories can be expressed and generated easily by coupling the knowledge updates to changing conditions, such as user determined actions and parameters. By reasoning on such conditions, the agents will naturally generate distinct scenaria and their interplay can mature into branching novel storylines. Traditional techniques used in interactive storytelling such as the integration of drama managers to control dramatic tension and story consistency can also be incorporated by designing their function around PLP rules.

This simple scenario already illustrates the interplay between different logic programming techniques and demonstrates the advantages gained by combining their distinct strengths. Namely, the integration of top-down, bottom-up, hypothetical, moral and utility-based reasoning procedures results in a flexible framework for dynamic agent specification. The open nature of the framework embraces the possibility of expanding its use to yet other useful models of cognition such as counterfactual reasoning and theories of mind.

References

[1] M. Cavazza, F. Charles, and S. J. Mead. Interacting with virtual characters in interactive storytelling. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 318–325, New York, NY, USA, 2002. ACM.

[2] C. Crawford. *Chris Crawford on Interactive Storytelling*. New Riders Games, 2004.

[3] L.M. Pereira and G. Lopes. Prospective logic agents. *International Journal of Reasoning-based Intelligent Systems (IJRIS)*, 3/4(1):200–208, 2009.

[4] L.M. Pereira and A. Saptawijaya. Modelling morality with prospective logic. *International Journal of Reasoning-based Intelligent Systems (IJRIS)*, 3/4(1):209–221, 2009.