Informatics Research Proposal

Arys E. Andreou 0459563
a.e.andreou@sms.ed.ac.uk

University of Edinburgh

Belief Revision in a Multi-Agent System

Problem Description

Most computer related systems have come to a point were they require a large amount of data in order to function as intended. This large amount of data is described as a knowledge base. It is evident that this knowledge base needs to be updated, irrespectively of the frequency in which this happens. In an environment where new things are being learned some times the new information or data that is being processed and later added to the existing knowledge base contradicts what is already known. In this case some actions needs to be taken in order to preserve consistency and also be efficient. This is the task of belief revision. The methods that can be used to solve this problem are numerous and their requirements vary. One approach is a system that will simply give precedence to the most recent information that is given. There are of course other more sophisticated approaches. Gärdenfors mentions some methodological problems that exist and proposes an approach to a belief revision system [1][2][3]. This approach is later tackled by Berendt and Smaill [1] using epistemic entrenchment.

Aim

This project will attempt to develop a multi-agent system that will be able to reason and revise its own knowledge-base in any given occasion, especially when a contradiction arises. The case in which there is a contradiction is what will be investigated the most and the real aim of the project. This means that the system should constantly have a concise knowledge-base with the minimum cost of information and potentially computational resources.

The task of this project will be to use and evaluate the various belief revision systems and then implement them in solving a problem of either navigation and planning or succeeding in a game scenario. A possible scenario, for navigation, is one in which a robot is in a labyrinth and without any map of the area must reach a target destination. All the robot has is its sensors and the knowledge of how to use them as well as a model of how the world works. I will assume in this case that there will be world model and not various internal models for various tasks.

Hypothesis

I hypothesize that the robot should be able to reason about what it receives from its sensors and assumes should be happening. In the case of a conflict of these two it should be able to reason and decide on what it should commit to its knowledge base. There should be no loss of information, or at least in the case where it is unavoidable it should be minimal. Thus the hypothesis of this project is that an agent using belief revision with epistemic entrenchment is better at collaborative problem solving than a system using

naive belief revision.

Objectives

Resulting from the project's aim there are some objectives that must be completed. These objectives are listed bellow.

1. Decide upon the best approach to solving this task. Which is the most efficient and which is the most rapid to develop. Weigh all the possible advantages and disadvantages of each one, later compare to the task at hand and conclude to which one is the best.
2. Develop the agents capable of interacting in the environment described in objective 1 and also incorporate the approach decided in objective 2
3. Implement the entire system to task at hand and tune to the best degree possible.
4. Design of multi-agent system incorporating belief revision.
5. Create an artificial and controllable problem configuration in which the agents will exist and interact. The reason for this sort of a problem configuration is that it will help test, troubleshoot and possibly compare with other similar systems.
6. Compare the existing system to others that already exist.

Timeliness and Feasibility

Considering that there might be a large demand for computational power, this is one reason why now with the advent of “cheap” supercomputers in the form of either clusters or grids such a task can be tackled within a sensible time frame. Another potential aspect for this projects timeliness is that given its highly theoretical nature its applications are numerous and uncertain at the same time. It can either be implemented in computer games or in future robotic systems that will require such an attribute.

Given that there is a large computational demand for a complex system that could potentially emerge from the project's description it will not be done in this specific case. The time constraint that is imposed is the most determining factor for this. There will be effort done though in the implementation of an as realistic and complicated system as needed for the problem description.

The degree of feasibility depends on how complicated the system will be. According to which cases are to be considered the degree of difficulty will be decided. There is research done on this topic so it is feasible, what remains to be seen is whether this particular implementation will succeed.

A reason for actually attempting such a task is its implementation in future robotic systems. The outcome of this project could not possibly affect this field in the short term but could potentially be a forerunner for future research that can spawn interesting applications. The optimal would be a case in which a robot could reason about its environment and gather knowledge autonomously. Based on the information extracted from the environment it would be able to reason over contradictions that might occur with its existing knowledge base and update it which would later result in potential activity that has not been predetermined programmatically.

Methodology

In order to implement and test the approach proposed a problem description must be formulated. This description will either be that of a game setting or a simulation of a robot trying to find its way around a maze. Given some predefined rules that are observable and controllable the proposed algorithm can be evaluated. The use of agents can either be done in a way that they will collaborate or that they will compete or a combination of both.

There is a plethora of platforms that can be used to develop the agents. Some of these are JADE and Mozart/Oz [4][5]. These are the two platforms that will be considered of which only one will be decided upon. There are several strengths and weaknesses to both that need to be evaluated prior to making a decision. Some of these characteristics are the large support for JADE which is highly attractive and on the other hand Mozart/Oz has the benefit of being faster in development [4][5]. Whether one would perform better than the other in this implementation is not known so they cannot be compared based on that.

Time Plan

Throughout the duration of the project's completion the various intermediate stages are the following. The estimated time required for their completion is shown in brackets.

1. Develop agents (1.5 months)
2. Setup of the environment rules and its implementation (2 weeks)
3. Experimenting on the possible approaches and concluding to one (3 weeks)
4. Implementation of the approach (3 weeks)
5. Fine tuning of the entire system (2 weeks)
6. Compare with alternatives (2 weeks)
7. Final report (3 weeks)

Most parts will be running concurrently with the first except for those that depend on its outcome. Namely these are 4,5 and 6.

Evaluation

In order to be able to measure or at least estimate the success of the project the outcome will be compare with other successful approaches. In the case where the robot in the maze is used a possible comparison method will be the estimation of this approaches' efficiency compared to others such as to naive belief revision strategies. Alternatively a similar approach will be used if the case will be that of a game setting.

References:

1. Berendt B., Smaill A. “Computing Belief Revision”. Research Paper 609, DAI, Edinburgh, 1992
2. Gärdenfors, P. “Knowledge in Flux”, MIT Press, 1988.
3. Gärdenfors, P. “Belief Revision”, Cambridge University Press, 1992.
4. Java Agent Development Framework, http://jade.tilab.com/
5. Mozart Programming System, http://www.mozart-oz.org/
6. Nute, D: “Defeasible Reasoning”, in Fetzer, J. (ed), Aspects of Artificial Intelligence, Kluwer, 1988.