

Massachusetts Institute of Technology
16.412J/6.834J Cognitive Robotics

Problem Set 1

Cognitive Robot Types, Topics and Reasoning Methods

Lawrence Bush

BushL2@mit.edu, larrybush@ll.mit.edu

February 15, 2005

Part A: Topics of Fascination

In this section I will discuss the following three topics of interest in the area of reasoning applied to robotics:

Collaborative Task Execution
Resource Allocation
Learning

Collaborative Task Execution and Resource Allocation

Collaborative Task Execution and Resource Allocation are closely related. Therefore, I will address them together. The objective of collaboration is to produce a higher quality result than could be achieved alone. The objective of resource allocation is to increase efficiency or lower costs. These two concepts embody the quality – cost tradeoff of many operations research optimization problems.

The central idea is that teams or colonies of robots would need to be able to self organize in a way that allocates them in an efficient manner. The tasks that need to be done would change as time goes on.

For a given task, the robots need to decide who does what. For example, in a search and rescue operation, the robots need to spread out over an area to search for the missing person.

The way in which the robots are deployed will significantly affect the success of the mission. Suppose that prior knowledge indicates that the missing person is more likely to be in one location over another. Also, certain areas are easier (lower cost) to search than others. Consequently, if each robot operates independently, they may all choose the area that is easiest to search and has the highest probability of success while neglecting other areas. A more optimal solution can be found. Therefore, another way to collaborate is desirable for this application.

The robot allocation could be centrally planned and modeled as a Markov Decision Process (MDP) or constraint satisfaction problem or it could be constructed using a market-based approach. These are very interesting approaches and will be discussed more in Part D.

Learning

Learning is a very broad topic in artificial intelligence. The index of [Artificial Intelligence, A Modern Approach 2nd Edition](#), by Stuart Russell and Peter Norvig (AIMA) lists over 100 pages in the book that address the topic of learning. The broad reaching nature of the topic is captured in titles like “universal

reinforcement learner.” While that goal seems intractable, robots still need to be able to perceive reason and act in new and changing environments.

A variety of learning methods have been studied for many applications. Traditional learning algorithms have a narrow application scope. These approaches are very interesting, important and challenging. However, robots need to learn from experience in a more general sense. They need to learn to identify what is important, internalize this into their knowledge base and use it to make decisions in the future.

Two approaches to this problem are recurrent neural networks (<http://www.idsia.ch/~juergen/rnn.html>) and optimal ordered problem solver (<http://www.idsia.ch/~juergen/oops.html>). The referenced web site is flat out fascinating. With that said, I find this reasoning area to be overly ambitious. My approach to a problem this big is to nibble away at the edges. In short, I intend to choose a problem that is more specific and narrower in scope.

A Summary of Robotics

This section provides an overall view of robotics. I wrote this section in order to become familiar with the different types of robots, sensors, reasoning techniques and applications in the field. This section essentially summarizes Chapter 25 of AIMA.

This section is organized in the following order:

- Sensors
- Reasoning
- Types of Robots
- Applications

A robot is a feedback and control system that involves sensing and reasoning about its environments, which leads to a decision on how to act. The first three parts of this section are organized in the order in which they are used in a feedback and control system. The final part covers how robots can be applied.

Sensors

There are many kinds of sensors that a robot can use [1] to sense its environment:

Range Finders

- Laser Range Finder
- Sonar
- Global Positioning System
- Tactile Sensors (whiskers)

Imaging Sensors

- Camera
- Infrared

Other

- Temperature
- Odor
- Acoustic
- Velocity
- Light Intensity

The sensors that are chosen for the particular application become part of the feedback loop described above. This feedback loop is sometimes called perception. Perception is essentially a filtering task, which

involves sensing, updating the internal belief state and acting. For example, the important and well-studied localization problem involves sensing the robot's location (i.e. using a range finder), updating the robot's estimated location (i.e. Kalman Filter) and moving in the direction that leads to the goal.

Proprioceptive Sensors are another class of sensors that measure the internal state of the robot. Some examples of these are:

- Inertial (gyroscope)
- Force / torque sensors
- Odometer

While these measure the internal state of the robot, they can also be used to improve the robot's interpretation of its environment. For example, an inertial sensor will reduce the position uncertainty in a localization task.

Reasoning

After sensing the environment, a robot must update its internal state and make a decision about how to act. This reasoning step can have a variety of objectives, for example:

- Localization
- Mapping
- Task Planning
- Resource Allocation
- Generic Search
- Learning (in general)

Acting (Types of Robots)

The manifestation of the acting step depends on the type of robot and how it can affect its environment. Robots can be grouped into 3 types: manipulators, mobile and hybrid robots. Manipulators are able to change their environment through grasping, pushing, lifting etcetera. An assembly line robot is a popular example of this type of robot.

A mobile robot is able to move about via wheels or legs. Some examples of mobile robots are unmanned land vehicle, unmanned air vehicle, autonomous underwater vehicles, and planetary rovers.

A hybrid robot is a mobile robot equipped with manipulators, for example a humanoid robot.

Applications

A robot's ability to sense, reason and interact with its environment is application dependent. Robots are used in industry, exploration, health care and personal services. The following are some possible applications:

- Industry
 - Mining
 - Agriculture
 - Material Handling
 - Welding
 - Painting
 - Harvest

- Excavate Earth
- Transportation
 - Of goods between warehouses
 - Retrieving items from a warehouse
- Exploration
 - 3-D map of a mine
 - Mars
 - Robotic arms assist astronaut in deploying and retrieving satellites
 - Military “drones”
- Health care
 - Assist surgeons with instrument placement (heart, eyes, brain)
 - Aides for the elderly (robotic walkers, pill reminders)
- Personal Services (perform daily tasks)
 - Vacuum cleaners
 - Lawn mowers
 - Golf caddies
 - Snow removal
- Public Safety
 - Search and Rescue
 - Fire fighting
- Entertainment
 - Robot dog
 - Robotic soccer
- Human Augmentation
 - Walking machines
 - Robotic teleoperation (hepatic interface)

Current Research Areas

The above section provides a survey of the relevant areas of Robotics. This section serves to identify where the current research areas are. In order to do this, I surveyed the research accepted at the Proceedings of the Fourth International Cognitive Robotics Workshop, Valencia, Spain, August 23-24, 2004. I took the approach of summarizing every paper at the conference in order to find out what is important.

I review each paper in a necessarily short format, which includes a summary of the main components of the paper and a list of the application, type of reasoning and techniques used to solve the problem.

Paper: Specifying Failure and Progress Conditions in a Behavior-Based Robot Programming System,
Authors: F. Kabanza, K. B. Lamine

Uses Linear Temporal Logic to monitor the progress of a behavior-based robotic system toward a global goal.

Application:	Navigation
Form of Reasoning:	Progress Monitoring
Techniques:	Linear Temporal Logic

Paper: Robel: Synthesizing and Controlling Complex Robust Robot Behaviors,
Authors: B. Morisset, G. Infante, M. Ghallab, F. Ingrand

Uses a Markov Decision Process to model / learn the relationship between the supervision state and the proper action. The objective is a high level task, thus the actions are modeled using a Hierarchical Task Network.

Application: Navigation
Form of Reasoning: Progress Monitoring and Policy Learning
Techniques: MDP

Paper: Patterns in Reactive Programs,
Author: P. Haslum

Discusses analogous software design patterns and reactive task procedures.

Application: Unmanned Air Vehicle Navigation and target tracking
Form of Reasoning: Navigation and Tracking
Techniques: Object Oriented Programming and Design

Paper: Decision Theoretic Planning for Playing Table Soccer,
Authors: M. Tacke, T. Weigel, B. Nebel

Maximum expected utility action selection (decision theoretic planning) using forward simulation.

Application: Robotic Table Soccer (Foosball)
Form of Reasoning: Planning
Techniques: Forward Simulation

Paper: On-line Decision Theoretic Golog for Unpredictable Domains,
Authors: A. Ferrein, C. Fritz, G. Lakemeyer

This paper integrates a decision theoretic planner with real-time updatable location uncertainty.

Application: Robotic Soccer (RoboCup)
Form of Reasoning: Planning
Techniques: Decision Theoretic Planning

Paper: Learning Partially Observable Action Models,
Author: E. Amir

This paper presents principles and algorithms for learning and filtering a logical action model.

Application: Turning on a light.
Form of Reasoning: Learning
Techniques: MDP, POMDP, HMM, EM, STRIPS

Paper: Building Polygonal Maps from Laser Range Data,
Authors: J. Latecki, R. Lakaemper, X. Sun, D. Wolter

This paper presents a heavily feature based localization and mapping approach (to the exclusion of odometer data).

Application: Generic Navigation

Form of Reasoning: Localization and Mapping
Techniques: Shape Based Object Recognition

Paper: Hierarchical Voronoi-based Route Graph Representations for Planning, Spatial Reasoning and Communication,
Authors: J. O. Wallgrün

This paper presents a new graph based path planning approach.

Application: Office Robot
Form of Reasoning: Path Planning
Techniques: Spatial Representation, Graph Model

Paper: Schematized Maps for Robot Guidance,
Authors: D. Wolter, K-F Richter

This paper presents a multi-robot central localization, mapping and path planning system.

Application: Service Robots
Form of Reasoning: Navigation
Techniques: Shape Matching, Multi-robot learning

Paper: How can I, robot, pick up that object with my hand,
Authors: A. Morales, P.J. Sanz, A.P. del Pobil

This paper presents vision based feature selection learning for reliable grasp comparison and prediction.

Application: Service robots
Form of Reasoning: Grasping
Techniques: Nearest Neighbor Classification

Paper: On Ability to Automatically Execute Agent Programs with Sensing,
Authors: S. Sardina, G. DeGiacomo, Y. Lesperance, H. Levesque

This paper addresses the inconsistency of real time sensing uncertainty into an entailment and consistency-based paradigm.

Application: Lumberjacking
Form of Reasoning: Planning
Techniques: Situation Calculus (Golog)

Paper: Have Another Look: On Failures and Recovery in Perceptual Anchoring,
Authors: M. Broxvall, S. Coradeschi, L. Karlsson, A. Saotti

This paper addresses perceptual error recognition and recovery.

Application: Generic Navigation
Form of Reasoning: Knowledge Based Planning, Recovery from Perceptual Errors
Techniques: Anchoring

Paper: Flexible Interval Planning in Concurrent Temporal Golog,

Authors: A. Finzi, F. Pirri

This paper discusses resource allocation and the exploration / exploitation tradeoff.

Application: Search and Rescue team
Form of Reasoning: Multi-Agent Planning and Scheduling of Localization and Mapping
Techniques: Situation Calculus (Golog)

Paper: Imitation and Social Learning for Synthetic Characters,

Authors: D. Buchsbaum, B. Blumberg, C. Breazeal

This paper discusses a method to identify actions and goals through imitation.

Application: Watching Cartoons
Form of Reasoning: Action and Goal Identification, Imitation Learning
Techniques: Simulation theory

Paper: On Reasoning and Planning in Real-Time: An LDS-Based Approach,

Authors: M. Asker, J. Malec

This paper addresses the complexity of Labeled Deduction Systems for planning.

Application: Theoretical
Form of Reasoning: Planning
Techniques: Labeled Deductive Systems,

Paper: Exploiting Qualitative Spatial Neighborhoods in the Situation Calculus,

Authors: F. Dylla, R. Moratz

This paper integrates a spatial theory into situation calculus (Golog) for robot navigation.

Application: General robot navigation.
Form of Reasoning: Localization, Mapping and Navigation
Techniques: Dipole Calculus, Line Segment Calculus, and Situation Calculus

Part D: Researching a Critical Reasoning Method

In this section I discuss three papers in the area of multi-agent planning. The three papers are applicable to the resource allocation reasoning capability described in Part B. Specifically, the application is a robotic search and rescue team. This team would be capable of collaboratively searching a pre-defined area. The search plan would need to balance multiple priorities, robot abilities and initial conditions. The team would be connected via wireless communication. One of the reasoning capabilities that this would require is resource allocation.

The three papers each provide a different perspective on how to solve the multi-agent planning / resource allocation problem in a distributed system. I selected these three papers primarily because they represent the 3 main schools of thought in this area, which are:

- Distributed Markov Decision Process
- Distributed Constraint Satisfaction
- Market Based Framework

Title: Using Cooperative Mediation to Solve Distributed Constraint Satisfaction Problems
Authors: Mailler,Roger; and Lesser,Victor
Publication: Proceedings of Third International Joint Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2004)

Why I selected this paper:

I selected this paper because it is relevant to the problem (distributed resource allocation), it used one of the three primary ways of addressing the problem, it was in a respected publication (AAMAS 2004) and I was impressed by one of it's authors, Victor Lesser, who had a mind-boggling 7 papers accepted to the same conference.

Major contribution:

This paper presents a new algorithm (Asynchronous Partial Overlay) to solve the distributed constraint satisfaction problem. The major contribution involves a mediator agent, who takes over the process (effectively centralizing it) when the distributed constraint satisfaction (DCSP) fails to solve the problem. Depending on the situation, the mediator agent may centralize only a portion of the graph.

Computer simulations show that the new algorithm is faster than the current best algorithm (Asynchronous Weak-Commitment). A proof shows that the distributed portion knows if it is locally unsolvable, the local-centralized protocol is dead-lock free and the worst case fully centralized solution will always find a solution if one exists (or terminate if a solution does not exist).

Strengths:

The strengths of this presentation are the inclusion of algorithm pseudo code, the clearness of the example problem, and the combined use of simulation and mathematical proofs.

Weaknesses:

I found no notable weaknesses.

Title: Hoplites: A Market-Based Framework for Planned Tight Coordination in Multi-robot Teams
Authors: N. Kalra, D. Ferguson, and A. Stentz
Publication: Proceedings of the International Conference on Robotics and Automation, April, 2005.

Why I selected this paper:

I selected this paper because it is relevant to the problem (coordination in multi-robot teams), it used one of the three primary ways of addressing the problem (market based framework), it is a fascinating subject and has a cool name.

Major contribution:

The major contribution of this paper is a new market based coordination method that improved (over MVERT, P-MVERT and PC-MVERT) significantly the quality of the results.

With that said, the way in which the auction based coordination was orchestrated bears some explanation.

In most market-based resource allocation methods, each robot earns points, which measure how much the robot's action moved the team toward its goal. Each task also has a team resource cost. The robots can trade tasks through auctions and negotiations in order to get the most profitable tasks.

In the Hoplites method there are 2 modes, active and passive. In the passive mode, the agents formulate and trade **plans**, which inform each other of their intentions. This allows team members to make informed decisions. In the active mode, teammates publish plans that require participation to work. They then pay for that participation. These active plans do not always work, however, they are initiated when tight coordination is required to achieve for a more optimal solution.

Strengths:

The new algorithm improved significantly the quality of the coordination results. I was particularly impressed that the made an effort to optimize the other algorithms to the application before comparing them. The algorithm was also demonstrated in real robots, which is even more interesting.

Weaknesses:

I found no notable weaknesses.

Title: Multi-agent Planning with Factored MDPs
Authors: Carlos Guestrin, Daphne Koller, and Ronald Parr
Publication: Advances in Neural Information Processing Systems NIPS-14, 2001
Invited Talk: Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)

Why I selected this paper:

I selected this paper because it is relevant (multi-agent planning) and it was by an author (Daphne Koller) whose book on Bayesian Networks I have read and liked. The paper was also an invited talk in a respected conference (Proceedings of Third International Joint Conference on Autonomous Agents and Multi-Agent Systems), which gave it added credibility.

Major contribution:

This paper explains a new way to decompose the value function of an MDP in order to make the algorithm tractable in a distributed framework. This is an approximate method, which achieves close to the theoretically optimal results.

Essentially, you can use an MDP to find the optimal policy for a resource allocation problem. However, to do this in a distributed sense, each agent needs to know its portion of the overall system reward, in order for it to make a decision. This decomposition method provides new way to figure out each agent's portion.

Strengths:

The author explains some concepts well. The idea is very interesting.

Weaknesses:

This author explains things very well in her book. This algorithm is really interesting, however, it requires more than 7 pages to explain precisely.

Part E: A Simple Project For Your Cognitive Robot

Application: Search and rescue mission.
Reasoning: Collaborative Resource Allocation

Motivation

Resource allocation is an important form of reasoning because it is a step toward high-level reasoning. Essentially, resource allocation is deciding what to do next in the context of a society. The robots are deciding what they can do that will be of greatest benefit to the society as a whole. Regardless of whether this uses a cooperative or competitive motivation mechanism, the robots want to do the task that is most beneficial. For this project, I propose a decision domain that does not require a complicated knowledge base, which will allow us to do research in high-level decision-making.

Setup

My project involves the simulation of a simulated robotic search and rescue team. The team will consist of ground moving wheeled robots. The search area will be represented as a grid. The team will collaboratively search the pre-defined area. The search plan would need to balance differing robot abilities and initial conditions. The team would be connected via wireless communication. The reasoning capabilities that this would require are collaborative learning, resource allocation.

Algorithms

I have reviewed 3 different types of algorithms for solving this type of problem. Each of the algorithms is better than the ones that it was compared to. My proposal is to implement all 3 and compare them to each other, in our problem domain. We will then use the best of the 3 algorithms as our new starting point and improve upon it.

Note: The MDP algorithm cannot be implemented directly from the description in the paper. I would need to get clarification from the author on the details of the algorithm.

The teams will use the following collaboration schemes:

1. An auction based collaboration scheme similar to the one presented in “Hoplites: A Market-Based Framework for Planned Tight Coordination in Multi-robot Teams,” by N. Kalra, D. Ferguson, and A. Stentz.
2. The MDP based algorithm presented in “Multi-agent Planning with Factored MDPs,” by Carlos Guestrin, Daphne Koller, and Ronald Parr.
3. The constraint satisfaction based algorithm presented in “Using Cooperative Mediation to Solve Distributed Constraint Satisfaction Problems,” by Mailler, Roger; and Lesser, Victor.

Note: The probability that I will pursue this project is .75.

A summary of last year's projects:

This section provides a summary of the projects (on the web site) from last year's Cognitive Robotics class. I wrote this section to help formulate the scope and range of topics for an acceptable course project.

I review each paper in a necessarily short format, which includes a summary of the main components of the paper and a list of the application, type of reasoning and techniques used to solve the problem.

Paper: Model-based Programming for Cooperating Vehicles

Seung Chung
Robert Effinger
Thomas Léauté
Steven D. Lovell

This project studied cooperative path planning for autonomous UAV forest fire-fighting.

Application: Forest Fire-Fighting with Autonomous Cooperative UAVs
Form of Reasoning: Path Planning, Goal (Mission) Planning
Techniques: D*-Lite, Dynamic Backtracking, Strategy-driven Kino-dynamic Cooperative Path Planning, Receding Horizon Continuous Planner

Paper: Cooperative Q-Learning

Lars Blackmore
Steve Block

This project compared cooperative and independent Q-Learning for learning a navigation policy about a maze.

Application: Robot navigation
Form of Reasoning: Cooperative Navigation Policy Learning
Techniques: Cooperative Q-Learning, MDP

Paper: SLAM for Dummies

Morten Rufus Blas
Soren Riisgaard

This group implemented and wrote a tutorial for SLAM.

Application: Generic Localization and Mapping
Form of Reasoning: Localization and Mapping
Techniques: SLAM

Paper: Towards Visual SLAM in Dynamic Environments

Vikash Mansinghka

This project critiqued and developed some new concepts for Visual SLAM.

Application: Navigation (Localization and Mapping)
Form of Reasoning: Localization and Mapping
Techniques: Visual SLAM

Paper: Autonomous Visual Tracking Algorithms

Alexander Omelchenko

This project investigated the visual tracking of ground vehicles via the collaboration between an autonomous UAV and a rover (ground vehicle).

Application: Visual tracking of ground vehicles
Form of Reasoning: Tracking
Techniques: Image Processing: Contour Tracking and Center of Mass Tracking