Michael Terry

16.412J/6.834J

2/16/05

Problem Set 1

A. Topics of Fascination

The first topic I would like to explore is probabilistic reasoning with Bayesian nets. I see that reasoning under situations of uncertainty is a very important area of AI, and it is apparent that these techniques are the standard approach for effectively modeling this uncertainty. I am particularly interested in learning how to model complex relationships between variables in spite of redundant links in Bayesian Nets. Also, I would like to study to the automatic generation of Bayesian Nets from large training set.

Secondly, I would like to investigate reasoning about opponents using game theory and Minimax search. Beyond the typical examples using a single opponent (ie chess), I would like to explore further the idea of reasoning in a world with potentially more than one cooperating and/or adversarial agent.

Finally, I would like to explore reasoning using first order logic. As described in Russell and Norvig, these concepts and methods seem to be very refined for basic objects, their properties, and the relationships between them. This is very useful for implementation of these properties in object-oriented data structures in any artificial intelligence application.

B. My Cognitive Robot

My cognitive robot is one that can sit down and play poker with expert-level opponents, and win the game consistently. To achieve that end, the robot must be endowed with a number of capabilities related to vision, mechanics and control, as well as a multitude of reasoning capabilities for strategy.

Simple object recognition would be required of the robot. In order to assess information about the game, the robot must be able to view partially obstructed chip stacks in front of its opponents, as well as "community" and "hole" cards that are dealt. A more sophisticated vision algorithm would incorporate information from opponents' mannerisms and body language. These critical pieces of information are known in the poker world as "tells".

In addition, the robot must be capable of stacking and counting chips in order to mange its own chip stack and place bets. To perform this function, pressure sensors and fine motor control of its "hands" would be required, along with the vision capabilities to determine chip denominations. These tasks are very similar to the classic AI situation of stacking blocks.

Finally, good strategy would require reasoning on many levels, with various levels of certainty regarding available information. I would like to focus on a three level model. At the highest level, the robot must be able to reason about its own hand given information about the community cards and opposition. Then, it must be able to extract information about its' opponents cards, given their tendencies and actions. Finally, it must make a decision about what to do regarding levels one and two. Against better opposition, the game could not be won by only thinking at two levels. For example, one must add another level of reasoning, which involves understanding what your opponent

thinks you have. However, I will argue that against average opposition, levels one and two are sufficient for a winning strategy.

C. Game Theory and Minimax Search in a Multi-Agent Environment

An ability that is critical to determining a good strategy is being able to look at future moves of opponents, to determine which move will allow the robot to maximize its Expected Value over the current round, hand, session, or lifetime. This is an extremely complex task that must be broken down into numerous subtasks. I will investigate the standard approach to performing these functions via Minimax Search.

D. Multi-Agent Game Theory and Minimax: Further Investigation

i.) Approximating Game-Theoretic Optimal Strategies in Full-scale Poker by Billings et. Al.

This paper's main contribution is a model for reducing the search space of Poker minimax through the use of "bucketing". This concept is very basic, and is used by most human players. Rather than describing a given situation using every detail, situation can be grouped into categories. You will often hear poker players describe these situations, such as "I had top pair top kicker", or "She turned a set on me". These situations frequently have their own notations (Top Pair Top Kicker = TPTK) in poker literature. I find that Billings et Al. have performed an appropriate grouping of situations based on a particular poker expert/writer, David Sklansky's rank of hands. However, they have oversimplified the problem to poker involving only one opponent. In real poker, you are very rarely up against one opponent, and so in some sense they've cheated by reducing

the minimax search space beyond what is useful. I intend to use their bucketing system for my search without the two-agent reduction.

ii.) *On Pruning Techniques for Multi-Player Games by* Sturtevant and Korf.

The major contribution of this paper is that it shows that minimax search against multiple opponents, known as MaxN, has very limited capacity to be pruned via alpha-beta and branch-and-bound pruning. Using examples from the games of Sergeant Major and Hearts, they have shown a way to turn multiplayer minimax search into a "paranoid" two-player situation, in which every opponent has formed a coalition against you. This is an oversimplification, probably induced by their bias toward the game of hearts. In most games, each opponent is usually out for each other as much as they are out for you. In hearts, however, when you are attempting to shoot the moon, at some point your opponents may form an explicit coalition against you. This situation is exclusive to hearts, and I cannot think of another game where this is relevant.

I think their point regarding the limited utility of pruning in Multi-Player minimax is valid. However, I do not necessarily agree with the usefulness of transforming a MaxN tree into a simple two-player model. Also, in their conclusion, they allude to the fact that the incorporation of domain knowledge is key to reducing the search space. I will work to reduce this multi-player search space by incorporating deterministic and probabilistic knowledge about the game in my robot's strategies.

iii.) *Deep Blue by Campbell et. Al.*

This paper describes and evaluates some of the design decisions regarding the design of Deep Blue. I believe the major contribution of the design of this system was

not necessarily the intelligence (or lack thereof) behind its reasoning, but with the positive press and attention this brought to the AI community.

I have mixed feelings regarding Deep Blue as an accomplishment in AI. In some sense, the fact that this machine beat a world-class player is an accomplishment in computing, but not necessarily intelligence. I liken this to a situation where our calculators reduce complex mathematical formulas in microseconds. Although this is impressive, most people do not consider this a marvel of intelligence.

The common understanding is that Deep Blue was a brute force attack on the search space. As we begin to address more complex information spaces with search, it becomes increasingly important to incorporate knowledge from the particular domain to reduce that space. This is essentially the theme that I introduced in the evaluation of the previous paper, and I hope to incorporate this notion into my implementation of strategy for games.

E. Simple Project: Win Small Stakes Holdem

Although the challenges of playing poker in a physical world are great, the most difficult task for my robot is actually determining a winning strategy. I would like to focus on this particular task, and the multiple levels of reasoning. In particular, my goal is to be able to win Small-Stakes/Low-Limit Texas Holdem. Although this game is characterized by poor opposition making numerous mistakes, it is by no means an easy game. However, I am optimistic given the body of literature and availability of large sets of training data that this can be done in the last third of the course. The likelihood of me choosing this project is very high.

References

[1] Billings, D., Burch, N., Davidson, A. Holte R., Schaeffer, J., Schauenberg, T., & Szafron, D. (2003). . Approximating Game-Theoretic Optimal Strategies in Full-scale Poker. *In Proceedings of the 2003 International Joint Conference on Artificial Intelligence.*

[2] Campbell, M., Hoane, A. J., & Hsu, F. (2002**).** Deep Blue. *Artificial Intelligence (134)*, pp. 57 – 83.

[3] Russell, S. & Norvig, P. (2003). Artificial Intelligence A Modern Approach. Pearson Education International.

[4] Sturtevant, N. R., & Korf, R. E. (2004). *On Pruning Techniques for Multi-Player Games.* Computer Science Department, University of California, Los Angeles.