DECISION-THEORETIC PLANNING UNDER RISK-SENSITIVE PLANNING OBJECTIVES

A Dissertation
Presented to
The Academic Faculty

by

Yaxin Liu
刘亚新

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

College of Computing
Georgia Institute of Technology
April 2005

DECISION-THEORETIC PLANNING UNDER RISK-SENSITIVE PLANNING OBJECTIVES

Approved by:

Dr. Sven Koenig, Co-advisor
College of Computing
Georgia Institute of Technology

Dr. Craig Tovey, Co-advisor
College of Computing
Georgia Institute of Technology

Dr. Anton Kleywegt
School of Industrial and Systems Engineering
Georgia Institute of Technology

Dr. Frank Dellaert
College of Computing
Georgia Institute of Technology

Dr. Ashok Goel
College of Computing
Georgia Institute of Technology

Dr. Richard Goodwin
IBM T.J. Watson Research Center

Date Approved: April 11, 2005

To Zhijie and Emily.

Hofstadter's Law: It always takes longer than you think, even when you take into account Hofstadter's Law.

— Douglas R. Hofstadter

ACKNOWLEDGEMENTS

Many thanks to my advisor, Sven Koenig, for helping me through the process. Thanks to my co-advisor, Craig Tovey, for many encouraging conversations. Thanks to Richard Goodwin, for his guidance of my summer internships at IBM and for shaping the direction of this work. Thanks to Anton Kleywegt, for his encouragement and support from early on. Thanks to Frank Dellaert, for his advice on writing. And thanks to Ashok Goel, for helping in various ways.

Thanks to my fellow students, David Furcy, for long-term encouragement and collaboration; to Jimmy Davies, for moral support; to Chip Mappus, for much help with English; and to Patrick Yaner, for his bookshelf.

Thanks to Peter Stone, for allowing me to finish the thesis at UT Austin.

Special thanks to my friends Yuan Chen and Xiaoling Cao, for helping taking care of Emily.

I am forever in debt to Zhijie for the many years of enduring, and to Emily for not-so-many years. Lastly and not leastly, I appreciate the help and support from my parents and parents-in-law.

This work is partially supported by an IBM Fellowship.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ALGORITHMS

SUMMARY

Risk attitudes crucially affect human decision making preferences, especially in scenarios where huge wins or losses are possible, as exemplified by planetary rover navigation, oilspill response, and business applications. Decision-theoretic planners therefore need to take risk aspects into account to serve their users better. However, most existing decision-theoretic planners use simplistic planning objectives that are risk-neutral. The thesis research is the first comprehensive study of how to incorporate risk attitudes into decision-theoretic planners and solve large-scale planning problems represented as Markov decision process models. The thesis consists of three parts.

The first part of the thesis studies risk-sensitive planning in the case where exponential utility functions are used to model risk attitudes. In this case, there exists an optimal plan that maps states to actions. I show that existing decision-theoretic planners can be transformed to take risk attitudes into account. My approach is more general than previous approaches, which transform the planning tasks rather than the planning algorithms. The transformed algorithms bear visual resemblance to the original algorithms but special treatment may be needed to ensure their validity. Moreover, different transformations are needed if the transition probabilities are implicitly given, namely, temporally extended probabilities and probabilities given in a factored form. I show how the transformations and their variants can be applied to various decision-theoretic planners.

The second part of the thesis studies risk-sensitive planning in the case where general nonlinear utility functions are used to model risk attitudes. In this case, there does not in general exist an optimal plan that maps states to actions, and an optimal plan must take into consideration the accumulated rewards starting from the initial state. I show that a state-augmentation approach can be used to reduce a risk-sensitive planning problem to a risk-neutral planning problem with an augmented state space. I further use a functional

interpretation of value functions and approximation methods to solve the planning problems efficiently with value iteration. I also develop an exact method for solving risk-sensitive planning problems where one-switch utility functions are used to model risk attitudes.

The third part of the thesis studies risk sensitive planning in case where arbitrary rewards are used. In this case, many of the basic properties are unknown, including the existence and finiteness of optimal expected utilities. I propose a spectrum of conditions that can be used to constrain the utility function and the planning problem so that the optimal expected utilities exist and are finite. These conditions are the basis for the further development of computational procedures. I prove that the existence and finiteness properties hold for stationary plans, where the action to perform in each state does not change over time, under different sets of conditions.

I use two running examples to demonstrate that risk-sensitive planners can be easily created from their risk-neutral counterparts, and the resulting optimal plans are qualitatively different from optimal plans under a risk-neutral planning objective.