Reusing Risk-Aware Stochastic Abstract Policies in Robotic Navigation Learning*

Valdinei Freire da Silva[1], Marcelo Li Koga[2],
Fábio Gagliardi Cozman[2], and Anna Helena Reali Costa[2]

[1] Universidade de São Paulo – Escola de Artes, Ciências e Humanidades – São Paulo, Brazil.
[2] Universidade de São Paulo – Escola Politécnica – São Paulo, Brazil.
{valdinei.freire,mlk,fcozman,anna.reali}@usp.br

Abstract. In this paper we improve learning performance of a risk-aware robot facing navigation tasks by employing transfer learning; that is, we use information from a previously solved task to accelerate learning in a new task. To do so, we transfer risk-aware memoryless stochastic abstract policies into a new task. We show how to incorporate risk-awareness into robotic navigation tasks, in particular when tasks are modeled as stochastic shortest path problems. We then show how to use a modified policy iteration algorithm, called AbsProb-PI, to obtain risk-neutral and risk-prone memoryless stochastic abstract policies. Finally, we propose a method that combines abstract policies, and show how to use the combined policy in a new navigation task. Experiments validate our proposals and show that one can find effective abstract policies that can improve robot behavior in navigation problems.

Keywords: Risk-Awareness, Memoryless Stochastic Abstract Policies, Transfer Learning

1 Introduction

Consider an agent that displays risk-awareness in the sense that she has preferences amongst risk-prone, risk-neutral and risk-averse attitudes. This agent may face a series of tasks; for instance a robot may face navigation problems in a variety of rooms, each one in a different day. It is reasonable to suppose that insights obtained in the solution of the first task will be useful in the solution of the second task, and so on. Hence the agent may be interested in learning not only the solution of the first task, but also an *abstract* description of the solution that may be used when solving the second task, and so on. How can this risk-aware agent find a suitable abstract policy? How can the agent reuse and transfer into new problems this abstract policy, and how useful can it be?

* This research was partly sponsored by FAPESP – Fundação de Amparo à Pesquisa do Estado de São Paulo (Procs. 11/19280-8, 12/02190-9, and 12/19627-0) and CNPq – Conselho Nacional de Desenvolvimento Científico e Tecnológico (Procs. 311058/2011-6 and 305395/2010-6).