



Risk-Sensitive Reinforcement Learning via Policy Gradient Search

By Prashanth L. A., Michael C. Fu

Reinforcement learning (RL) is one of the foundational pillars of artificial intelligence and machine learning. An important consideration in any optimization or control problem is the notion of risk, but its incorporation into RL has been a fairly recent development. This monograph surveys research on risk-sensitive RL that uses policy gradient search.

The authors survey some of the recent work in this area specifically where policy gradient search is the solution approach. In the first risk-sensitive RL setting, they cover popular risk measures based on variance, conditional value at-risk and chance constraints, and present a template for policy gradient-based risk-sensitive RL algorithms using a Lagrangian formulation. For the setting where risk is incorporated directly into the objective function, they consider an exponential utility formulation, cumulative prospect theory, and coherent risk measures.

Written for novices and experts alike the authors have made the text completely self-contained but also organized in a manner that allows expert readers to skip background chapters. This is a complete guide for students and researchers working on this aspect of machine learning.

Pages: 170

List Price: \$99.00

Publication Date: June 2022

ISBN: 978-1-63828-026-2

Table of Contents

1. Introduction
 2. Markov Decision Processes
 3. Risk Measures
 4. Background on Policy Evaluation and Gradient Estimation
 5. Policy Gradient Templates for Risk-sensitive RL
 6. MDPs with Risk as the Constraint
 7. MDPs with Risk as the Objective
 8. Conclusions and Future Challenges
- Acknowledgements
References

Order at <https://nowpublishers.com/Order/BuyBook?isbn=978-1-63828-026-2> using discount code: 226485