
Ambiguity aversion and the role of information in decision making and planning

Naftali Tishby

School of Computer Science and Engineering,
The Hebrew University of Jerusalem,
Ross Building,
Givat Ram Campus,
91904 Jerusalem, Israel
tishby@cs.huji.ac.il

Abstract

One of the most striking characterizations of life is the ability to efficiently extract information - through sensory perception, and exploit it - through behavior. There is a growing empirical evidence that information seeking is as important for optimal behavior as reward seeking. Yet our basic algorithms for describing planning and behavior, in particular reinforcement learning (RL), so far ignored this component.

In this talk I will describe new extensions of reinforcement learning that combine information seeking and reward seeking behaviors in a principle optimal way. I will argue that Shannon's information measures provide the only consistent way for trading information with expected future reward and show how the two can be naturally combined in the frameworks of Markov-Decision-Processes (MDP) and Dynamic Programming (DP). This new framework unifies techniques from information theory (like the Huffman source coding algorithm) with methods of optimal control (like the Bellman equation). We show that the resulting optimization problem has a unique global minimum and convergence (even that it lacks convexity). Moreover, the tradeoff between information and value is shown to be robust to fluctuations in the reward values by using the PAC-Bayes generalization bound, providing another interesting justification to its biological relevance.