Driessens, Kurt; **Džeroski, Sašo**

Integrating guidance into relational reinforcement learning. (English) Zbl 1079.68084

Mach. Learn. 57, No. 3, 271-304 (2004).

Summary: Reinforcement learning, and Q-learning in particular, encounter two major problems when dealing with large state spaces. First, learning the Q-function in tabular form may be infeasible because of the excessive amount of memory needed to store the table, and because the Q-function only converges after each state has been visited multiple times. Second, rewards in the state space may be so sparse that with random exploration they will only be discovered extremely slowly. The first problem is often solved by learning a generalization of the encountered examples (e.g., using a neural net or decision tree). Relational Reinforcement Learning (RRL) is such an approach; it makes Q-learning feasible in structural domains by incorporating a relational learner into Q-learning. The problem of sparse rewards has not been addressed for RRL. This paper presents a solution based on the use of "reasonable policies" to provide guidance. Different types of policies and different strategies to supply guidance through these policies are discussed and evaluated experimentally in several relational domains to show the merits of the approach.

MSC:

68T05 Learning and adaptive systems in artificial intelligence

Keywords:

Reinforcement learning; relational learning; guided exploration

Full Text: DOI