


Reinforcement Learning



[Reinforcement Learning_ 下载链接1](#)

著者:Richard S. Sutton

出版者:A Bradford Book

出版时间:2018-11-13

装帧:精装

isbn:9780262039246

The significantly expanded and updated new edition of a widely used text on reinforcement learning, one of the most active research areas in artificial intelligence.

Reinforcement learning, one of the most active research areas in artificial intelligence, is a computational approach to learning whereby an agent tries to maximize the total amount of reward it receives while interacting with a complex, uncertain environment. In Reinforcement Learning, Richard Sutton and Andrew Barto provide a clear and simple account of the field's key ideas and algorithms. This second edition has been significantly expanded and updated, presenting new topics and updating coverage of other topics.

Like the first edition, this second edition focuses on core online learning algorithms, with the more mathematical material set off in shaded boxes. Part I covers as much of reinforcement learning as possible without going beyond the tabular case for which exact solutions can be found. Many algorithms presented in this part are new to the second edition, including UCB, Expected Sarsa, and Double Learning. Part II extends these ideas to function approximation, with new sections on such topics as artificial neural networks and the Fourier basis, and offers expanded treatment of off-policy learning and policy-gradient methods. Part III has new chapters on reinforcement learning's relationships to psychology and neuroscience, as well as an updated case-studies chapter including AlphaGo and AlphaGo Zero, Atari game playing, and IBM Watson's wagering strategy. The final chapter discusses the future societal impacts of reinforcement learning.

作者介绍:

Richard S. Sutton is Professor of Computing Science and AITF Chair in Reinforcement Learning and Artificial Intelligence at the University of Alberta, and also Distinguished Research Scientist at DeepMind.

目录:

[Reinforcement Learning 下载链接1](#)

标签

强化学习

机器学习

人工智能

RL

计算机科学

数学

MachineLearning

计算机

评论

只读了部分章节 写的很清楚

可以交流一下exercise的答案?

学习一个

看了前两部分

读了前几章。从动态规划讲到MC，TD，Q learning，但是还是不大好理解Q learning为啥work

强化学习必看书，还是草稿的时候从头到尾看了一遍，至少应该再看一遍。

看到最后当哲学书看了LOL...

RL的经典教科书。

不好读。

强化学习的书不多

主要是培养一些intuition 可以配合david silver的课程
<https://www.youtube.com/playlist?list=PLqYmG7hTraZDM-OYHWgPebj2MfCFzFObQ>

真的很好看 个人认为比花书好看，自洽 自成一体，公式简洁。当然 也有卖idea的现象

写的很系统，把之前散落在读论文里的各个知识点都串起来了，推荐，另外有随书的代码在github上

正确的元策略应该有一定的随机性

终于全部读完，非常好！堪称强化学习“圣经”。这周末出去，下周回来再做详细书评

配合cousera上reinforcement learning specialization食用更佳

随着上课看了下， 有时间再把后几张case看下

研究生靠啃这个毕业的

第二版是98年第一版基础上修订的，在原有基础上增加了一些神经网络相关算法。但原有知识架构还是比较老的，不过非常适合刚接触RL的新手。我只看了前三章基础，后面选择性看了几个重要的章节。前三章虽然基础，但内容非常引人入胜，语言逻辑清晰便于理解，从最基础的MDP讲起到Monte Carlo方法、到经典dynamic programming，后面章节再逐渐到用的最广的Q-learning和policy gradient等方法，一步一步加深完善，便于系统学习知识。Richard Sutton的RL经典，推荐！

[Reinforcement Learning_ 下载链接1](#)

书评

[<http://incompleteideas.net/book/the-book-2nd.html>] 有 [第二版的PDF](<http://incompleteideas.net/book/bookdraft2018jan1.pdf>)[]，还有 [Python实现](<https://github.com/ShangtongZhang/reinforcement-learning-an-introduction>)]。

可以在线阅读，还不错的 我还没仔细读，先把网址公布出来，大家一起学习
<http://webdocs.cs.ualberta.ca/~sutton/book/ebook/the-book.html>

这是一本极好的书，不仅能使你对强化学习有精确、透彻的理解，更能够提升你的思维层次。
接触人工智能领域6年多了，用过统计学习和深度学习做过一些项目。目前，David Silver的教学视频已经过完，这本书读到了第10章（第二版）。下面说一下个人浅陋的理解。 目前应用最广泛的监...

[Reinforcement Learning_ 下载链接1](#)