Lecture - How to Write a Literature Review

Guiliang Liu

The Chinese University of Hong Kong, Shenzhen

DDA4230: Reinforcement Learning
Course Page: [Click]

Steps to write a Literature Review

- Find a topic.
- Find the related scientific papers.
- Summarize these papers.
- Conclude the review with your thoughts.

Find a Topic

When you find the topic, there are several things you need to know. The topic must be related to course content (Reinforcement Learning).

Some basic candidate topics include:

- Multi-Arm Bandit Algorithms.
- Deep Q-Learning Algorithms.
- Experience Relay Algorithms.
- Policy Gradient and Optimization Algorithms.
- Actor-Critic Methods.

Find a Topic

When you find the topic, there are several things you need to know. The topic must be related to course content (Reinforcement Learning).

Some advanced candidate topics include:

- Monte-Carlo Tree Search
- Inverse Reinforcement Learning.
- Multi-Agent Reinforcement Learning.
- Decision Transformer.
- Reinforcement Learning from Human Feedback.

Find the related scientific papers

Use Google Scholar.

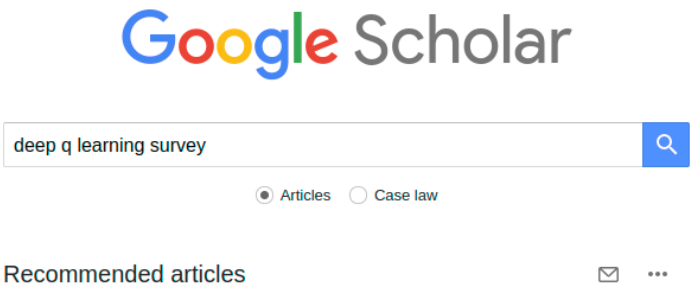

Find the related scientific papers

Use Google Scholar.

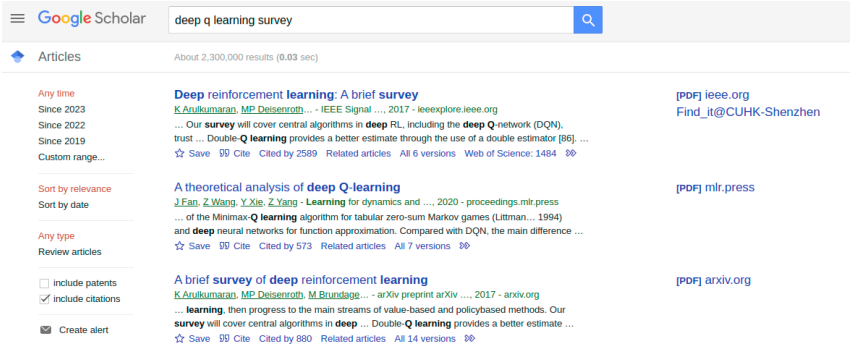

香港中文大學(深圳)
The Chinese University of Hong Kong, Shenzhen

Find the related scientific papers

Open the survey paper.

A Brief Survey of Deep Reinforcement Learning

Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, Anil Anthony Bharath

***Abstract*—Deep reinforcement learning is poised to revolutionise the field of AI and represents a step towards building autonomous systems with a higher level understanding of the visual world. Currently, deep learning is enabling reinforcement learning to scale to problems that were previously intractable, such as learning to play video games directly from pixels. Deep reinforcement learning algorithms are also applied to robotics, allowing control policies for robots to be learned directly from camera inputs in the real world. In this survey, we begin with an introduction to the general field of reinforcement learning, then progress to the main streams of value-based and policy-based methods. Our survey will cover central algorithms in deep reinforcement learning, including the deep Q-network, trust region policy optimisation, and asynchronous advantage actor-critic. In parallel, we highlight the unique advantages of deep neural networks, focusing on visual understanding via reinforcement learning. To conclude, we describe several current areas of research within the field.**

I. Introduction

is to cover both seminal and recent developments in DRL, conveying the innovative ways in which neural networks can be used to bring us closer towards developing autonomous agents. For a more comprehensive survey of recent efforts in DRL, including applications of DRL to areas such as natural language processing [106, 5], we refer readers to the overview by Li [78].

Deep learning enables RL to scale to decision-making problems that were previously intractable, i.e., settings with high-dimensional state and action spaces. Amongst recent work in the field of DRL, there have been two outstanding success stories. The first, kickstarting the revolution in DRL, was the development of an algorithm that could learn to play a range of Atari 2600 video games at a superhuman level, directly from image pixels [84]. Providing solutions for the instability of function approximation techniques in RL, this work was the first to convincingly demonstrate that RL agents

Find the related scientific papers

Find the relevant content and collect the papers (5-10 papers).

B. Q-Function Modifications

Considering that one of the key components of the DQN is a function approximator for the Q-function, it can benefit from fundamental advances in RL. van Hasselt [148] showed that the single estimator used in the Q-learning update rule overestimates the expected return due to the use of the maximum action value as an approximation of the maximum *expected* action value. Double-Q learning provides a better estimate through the use of a double estimator [148]. Whilst double-Q learning requires an additional function to be learned, later work proposed using the already available target network from the DQN algorithm, resulting in significantly better results with only a small change in the update step [149]. A more

香港中文大學(深圳)
The Chinese University of Hong Kong, Shenzhen

Summarize these papers

First of all,read the papers carefully!

A scientific literature review should have the following sections:

- Abstract.
- Introduction.
- Method 1, include motivation, algorithm, performance, shortages/limitations.
- Method 2, include motivation, algorithm, performance, shortages/limitations.
- ...
- Method N, include motivation, algorithm, performance, shortages/limitations.

香港中文大學(深圳)
The Chinese University of Hong Kong, Shenzhen

Conclude the Review with Your Thoughts

Write a conclusion section with your ideas and thoughts, please include the following things:

- The main contributions of these methods.
- The limitation and unsolved problems.
- Future works that can potentially solve these problems.
- Potentially real-world application of these methods.

Important Notes

- Write the review with Latex (the format file can be downloaded from the course page). The online editor Overleaf is recommended.
- If you don't know Latex, read the tutorial (Learn LaTeX in 30 minutes) or ask TA.
- The maximum page length is 10.
- No coding is required, and submit the PDF.

Question and Answering (Q&A)