Asynchronous Methods for Deep Reinforcement Learning

Type	Conference Paper
Author	Volodymyr Mnih
Author	Adria Puigdomenech Badia
Author	Mehdi Mirza
Author	Alex Graves
Author	Timothy Lillicrap
Author	Tim Harley
Author	David Silver
Author	Koray Kavukcuoglu
URL	http://proceedings.mlr.press/v48/mniha16.html
Pages	1928-1937
Date	2016/06/11
Accessed	26.3.2018, 15:04:32
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	We propose a conceptually simple and lightweight framework for deep reinforcement learning that uses asynchronous gradient descent for optimization of deep neural network controllers. We present as...
Proceedings Title	International Conference on Machine Learning
Date Added	26.3.2018, 15:04:32
Modified	26.3.2018, 15:04:32

Attachments

- Full Text PDF
- Snapshot

Averaged-DQN: Variance Reduction and Stabilization for Deep Reinforcement Learning

Type	Conference Paper
Author	Oron Anschel
Author	Nir Baram
Author	Nahum Shimkin
URL	http://proceedings.mlr.press/v70/anschel17a.html
Pages	176-185
Date	2017/07/17

Accessed	26.3.2018, 15:17:04
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	Instability and variability of Deep Reinforcement Learning (DRL) algorithms tend to adversely affect their performance. Averaged-DQN is a simple extension to the DQN algorithm, based on averaging p...
Proceedings Title	International Conference on Machine Learning
Short Title	Averaged-DQN
Date Added	26.3.2018, 15:17:04
Modified	26.3.2018, 15:17:04

Attachments

- Full Text PDF
- Snapshot

Deciding How to Decide: Dynamic Routing in Artificial Neural Networks

Type	Conference Paper
Author	Mason McGill
Author	Pietro Perona
URL	http://proceedings.mlr.press/v70/mcgill17a.html
Pages	2363-2372
Date	2017/07/17
Accessed	26.3.2018, 15:11:43
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	We propose and systematically evaluate three strategies for training dynamically-routed artificial neural networks: graphs of learned transformations through which different input signals may take ...
Proceedings Title	International Conference on Machine Learning
Short Title	Deciding How to Decide
Date Added	26.3.2018, 15:11:43
Modified	26.3.2018, 15:11:43

Attachments

- Full Text PDF
- Snapshot

Deep Reinforcement Learning from Human Preferences

Type	Book Section
Author	Paul F Christiano
Author	Jan Leike
Author	Tom Brown
Author	Miljan Martic
Author	Shane Legg
Author	Dario Amodei
Editor	I. Guyon
Editor	U. V. Luxburg
Editor	S. Bengio
Editor	H. Wallach
Editor	R. Fergus
Editor	S. Vishwanathan
Editor	R. Garnett
URL	http://papers.nips.cc/paper/7017-deep-reinforcement-learning-from-human-preferences.pdf
Publisher	Curran Associates, Inc.
Pages	4299–4307
Date	2017
Accessed	26.3.2018, 15:22:29
Library Catalog	Neural Information Processing Systems
Book Title	Advances in Neural Information Processing Systems 30
Date Added	26.3.2018, 15:22:29
Modified	26.3.2018, 15:22:29

Attachments

- NIPS Full Text PDF
- NIPS Snapshort

Deep Reinforcement Learning with Double Q-Learning

Type	Conference Paper
Author	Hado van Hasselt
Author	Arthur Guez
Author	David Silver
URL	http://dl.acm.org/citation.cfm?id=3016100.3016191
Series	AAAI'16

Place Phoenix, Arizona
Publisher AAAI Press
Pages 2094–2100
Date 2016
Accessed 26.3.2018, 15:06:01
Library Catalog ACM Digital Library
Abstract The popular Q-learning algorithm is known to overestimate action values under certain conditions. It was not previously known whether, in practice, such overestimations are common, whether they harm performance, and whether they can generally be prevented. In this paper, we answer all these questions affirmatively. In particular, we first show that the recent DQN algorithm, which combines Q-learning with a deep neural network, suffers from substantial overestimations in some games in the Atari 2600 domain. We then show that the idea behind the Double Q-learning algorithm, which was introduced in a tabular setting, can be generalized to work with large-scale function approximation. We propose a specific adaptation to the DQN algorithm and show that the resulting algorithm not only reduces the observed overestimations, as hypothesized, but that this also leads to much better performance on several games.
Proceedings Title Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence
Date Added 26.3.2018, 15:06:01
Modified 26.3.2018, 15:06:01

Dueling Network Architectures for Deep Reinforcement Learning

Type Conference Paper
Author Ziyu Wang
Author Tom Schaul
Author Matteo Hessel
Author Hado Hasselt
Author Marc Lanctot
Author Nando Freitas
URL http://proceedings.mlr.press/v48/wangf16.html
Pages 1995-2003
Date 2016/06/11
Accessed 26.3.2018, 15:08:07
Library Catalog proceedings.mlr.press
Conference Name International Conference on Machine Learning
Language en
Abstract In recent years there have been many successes of using deep representations in reinforcement learning. Still, many of these applications use conventional architectures, such as convolutional netwo...
Proceedings Title International Conference on Machine Learning

Date Added	26.3.2018, 15:08:07
Modified	26.3.2018, 15:08:07

Attachments

- Full Text PDF
- Snapshot

EX2: Exploration with Exemplar Models for Deep Reinforcement Learning

Type	Book Section
Author	Justin Fu
Author	John Co-Reyes
Author	Sergey Levine
Editor	I. Guyon
Editor	U. V. Luxburg
Editor	S. Bengio
Editor	H. Wallach
Editor	R. Fergus
Editor	S. Vishwanathan
Editor	R. Garnett
URL	http://papers.nips.cc/paper/6851-ex2-exploration-with-exemplar-models-for-deep-reinforcement-learning.pdf
Publisher	Curran Associates, Inc.
Pages	2577–2587
Date	2017
Accessed	26.3.2018, 15:21:22
Library Catalog	Neural Information Processing Systems
Book Title	Advances in Neural Information Processing Systems 30
Short Title	EX2
Date Added	26.3.2018, 15:21:22
Modified	26.3.2018, 15:21:22

Attachments

- NIPS Full Text PDF
- NIPS Snapshort

Human-level control through deep reinforcement learning

Type	Journal Article
Author	Volodymyr Mnih
Author	Koray Kavukcuoglu
Author	David Silver
Author	Andrei A. Rusu
Author	Joel Veness
Author	Marc G. Bellemare
Author	Alex Graves
Author	Martin Riedmiller
Author	Andreas K. Fidjeland
Author	Georg Ostrovski
Author	Stig Petersen
Author	Charles Beattie
Author	Amir Sadik
Author	Ioannis Antonoglou
Author	Helen King
Author	Dharshan Kumaran
Author	Daan Wierstra
Author	Shane Legg
Author	Demis Hassabis
URL	https://www.nature.com/articles/nature14236
Rights	2015 Nature Publishing Group
Volume	518
Issue	7540
Pages	529-533
Publication	Nature
ISSN	1476-4687
Date	2015-02
DOI	10.1038/nature14236
Accessed	26.3.2018, 15:04:56
Library Catalog	www.nature.com
Language	en
Abstract	The theory of reinforcement learning provides a normative account1, deeply rooted in psychological2 and neuroscientific3 perspectives on animal behaviour, of how agents may optimize their control of an environment. To use reinforcement learning successfully in situations approaching real-world complexity, however, agents are confronted with a difficult task: they must derive efficient representations of the environment from high-dimensional sensory inputs, and use these to generalize past experience to new situations. Remarkably, humans and other animals seem to solve this problem through a harmonious combination of reinforcement learning and hierarchical sensory processing systems4,5, the former evidenced by a wealth of neural data revealing notable parallels between the

phasic signals emitted by dopaminergic neurons and temporal difference reinforcement learning algorithms3. While reinforcement learning agents have achieved some successes in a variety of domains6,7,8, their applicability has previously been limited to domains in which useful features can be handcrafted, or to domains with fully observed, low-dimensional state spaces. Here we use recent advances in training deep neural networks9,10,11 to develop a novel artificial agent, termed a deep Q-network, that can learn successful policies directly from high-dimensional sensory inputs using end-to-end reinforcement learning. We tested this agent on the challenging domain of classic Atari 2600 games12. We demonstrate that the deep Q-network agent, receiving only the pixels and the game score as inputs, was able to surpass the performance of all previous algorithms and achieve a level comparable to that of a professional human games tester across a set of 49 games, using the same algorithm, network architecture and hyperparameters. This work bridges the divide between high-dimensional sensory inputs and actions, resulting in the first artificial agent that is capable of learning to excel at a diverse array of challenging tasks.

Date Added 26.3.2018, 15:04:56
Modified 26.3.2018, 15:04:56

Attachments

- Snapshot

Imagination-Augmented Agents for Deep Reinforcement Learning

Type Book Section
Author Sébastien Racanière
Author Theophane Weber
Author David Reichert
Author Lars Buesing
Author Arthur Guez
Author Danilo Jimenez Rezende
Author Adrià Puigdomènech Badia
Author Oriol Vinyals
Author Nicolas Heess
Author Yujia Li
Author Razvan Pascanu
Author Peter Battaglia
Author Demis Hassabis
Author David Silver
Author Daan Wierstra
Editor I. Guyon

Editor	U. V. Luxburg
Editor	S. Bengio
Editor	H. Wallach
Editor	R. Fergus
Editor	S. Vishwanathan
Editor	R. Garnett
URL	http://papers.nips.cc/paper/7152-imagination-augmented-agents-for-deep-reinforcement-learning.pdf
Publisher	Curran Associates, Inc.
Pages	5690–5701
Date	2017
Accessed	26.3.2018, 15:23:00
Library Catalog	Neural Information Processing Systems
Book Title	Advances in Neural Information Processing Systems 30
Date Added	26.3.2018, 15:23:00
Modified	26.3.2018, 15:23:00

Attachments

- NIPS Full Text PDF
- NIPS Snapshort

Improving Stochastic Policy Gradients in Continuous Control with Deep Reinforcement Learning using the Beta Distribution

Type	Conference Paper
Author	Po-Wei Chou
Author	Daniel Maturana
Author	Sebastian Scherer
URL	http://proceedings.mlr.press/v70/chou17a.html
Pages	834-843
Date	2017/07/17
Accessed	26.3.2018, 15:15:31
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	Recently, reinforcement learning with deep neural networks has achieved great success in challenging continuous control problems such as 3D locomotion and robotic manipulation. However, in real-wor...
Proceedings Title	International Conference on Machine Learning
Date Added	26.3.2018, 15:15:31

Modified 26.3.2018, 15:15:31

Attachments

- Full Text PDF
- Snapshot

Interpolated Policy Gradient: Merging On-Policy and Off-Policy Gradient Estimation for Deep Reinforcement Learning

Type	Book Section
Author	Shixiang Gu
Author	Tim Lillicrap
Author	Richard E Turner
Author	Zoubin Ghahramani
Author	Bernhard Schölkopf
Author	Sergey Levine
Editor	I. Guyon
Editor	U. V. Luxburg
Editor	S. Bengio
Editor	H. Wallach
Editor	R. Fergus
Editor	S. Vishwanathan
Editor	R. Garnett
URL	http://papers.nips.cc/paper/6974-interpolated-policy-gradient-merging-on-policy-and-off-policy-gradient-estimation-for-deep-reinforcement-learning.pdf
Publisher	Curran Associates, Inc.
Pages	3846–3855
Date	2017
Accessed	26.3.2018, 15:22:19
Library Catalog	Neural Information Processing Systems
Book Title	Advances in Neural Information Processing Systems 30
Short Title	Interpolated Policy Gradient
Date Added	26.3.2018, 15:22:19
Modified	26.3.2018, 15:22:19

Attachments

- NIPS Full Text PDF
- NIPS Snapshort

Neural Episodic Control

Type	Conference Paper
Author	Alexander Pritzel
Author	Benigno Uria
Author	Sriram Srinivasan
Author	Adrià Puigdomènech Badia
Author	Oriol Vinyals
Author	Demis Hassabis
Author	Daan Wierstra
Author	Charles Blundell
URL	http://proceedings.mlr.press/v70/pritzel17a.html
Pages	2827-2836
Date	2017/07/17
Accessed	26.3.2018, 15:13:17
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	Deep reinforcement learning methods attain super-human performance in a wide range of environments. Such methods are grossly inefficient, often taking orders of magnitudes more data than humans to ...
Proceedings Title	International Conference on Machine Learning
Date Added	26.3.2018, 15:13:17
Modified	26.3.2018, 15:13:17

Attachments

- Full Text PDF
- Snapshot

Neural Optimizer Search with Reinforcement Learning

Type	Conference Paper
Author	Irwan Bello
Author	Barret Zoph
Author	Vijay Vasudevan
Author	Quoc V. Le
URL	http://proceedings.mlr.press/v70/bello17a.html
Pages	459-468
Date	2017/07/17
Accessed	26.3.2018, 15:14:04

Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	We present an approach to automate the process of discovering optimization methods, with a focus on deep learning architectures. We train a Recurrent Neural Network controller to generate a string ...
Proceedings Title	International Conference on Machine Learning
Date Added	26.3.2018, 15:14:04
Modified	26.3.2018, 15:14:04

Attachments

- Full Text PDF
- Snapshot

Opponent Modeling in Deep Reinforcement Learning

Type	Conference Paper
Author	He He
Author	Jordan Boyd-Graber
Author	Kevin Kwok
Author	Hal Daumé Iii
URL	http://proceedings.mlr.press/v48/he16.html
Pages	1804-1813
Date	2016/06/11
Accessed	26.3.2018, 15:07:46
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	Opponent modeling is necessary in multi-agent settings where secondary agents with competing goals also adapt their strategies, yet it remains challenging because of strategies’ complex interaction...
Proceedings Title	International Conference on Machine Learning
Date Added	26.3.2018, 15:07:46
Modified	26.3.2018, 15:07:46

Attachments

- Full Text PDF
- Snapshot

Playing Atari with Deep Reinforcement Learning

Type	Journal Article
Author	Volodymyr Mnih
Author	Koray Kavukcuoglu
Author	David Silver
Author	Alex Graves
Author	Ioannis Antonoglou
Author	Daan Wierstra
Author	Martin Riedmiller
URL	http://arxiv.org/abs/1312.5602
Publication	arXiv:1312.5602 [cs]
Date	2013-12-19
Extra	arXiv: 1312.5602
Accessed	26.3.2018, 15:05:26
Library Catalog	arXiv.org
Abstract	We present the first deep learning model to successfully learn control policies directly from high-dimensional sensory input using reinforcement learning. The model is a convolutional neural network, trained with a variant of Q-learning, whose input is raw pixels and whose output is a value function estimating future rewards. We apply our method to seven Atari 2600 games from the Arcade Learning Environment, with no adjustment of the architecture or learning algorithm. We find that it outperforms all previous approaches on six of the games and surpasses a human expert on three of them.
Date Added	26.3.2018, 15:05:26
Modified	26.3.2018, 15:05:26

Tags:

Computer Science - Learning

Notes:

Comment: NIPS Deep Learning Workshop 2013

Attachments

- arXiv:1312.5602 PDF
- arXiv.org Snapshot

Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation

Type	Book Section
Author	Yuhuai Wu
Author	Elman Mansimov
Author	Roger B Grosse
Author	Shun Liao
Author	Jimmy Ba
Editor	I. Guyon
Editor	U. V. Luxburg
Editor	S. Bengio
Editor	H. Wallach
Editor	R. Fergus
Editor	S. Vishwanathan
Editor	R. Garnett
URL	http://papers.nips.cc/paper/7112-scalable-trust-region-method-for-deep-reinforcement-learning-using-kronecker-factored-approximation.pdf
Publisher	Curran Associates, Inc.
Pages	5279–5288
Date	2017
Accessed	26.3.2018, 15:22:44
Library Catalog	Neural Information Processing Systems
Book Title	Advances in Neural Information Processing Systems 30
Date Added	26.3.2018, 15:22:44
Modified	26.3.2018, 15:22:44

Attachments

- NIPS Full Text PDF
- NIPS Snapshort

Shallow Updates for Deep Reinforcement Learning

Type	Book Section
Author	Nir Levine
Author	Tom Zahavy
Author	Daniel J Mankowitz
Author	Aviv Tamar
Author	Shie Mannor
Editor	I. Guyon

Editor	U. V. Luxburg
Editor	S. Bengio
Editor	H. Wallach
Editor	R. Fergus
Editor	S. Vishwanathan
Editor	R. Garnett
URL	http://papers.nips.cc/paper/6906-shallow-updates-for-deep-reinforcement-learning.pdf
Publisher	Curran Associates, Inc.
Pages	3135–3145
Date	2017
Accessed	26.3.2018, 15:21:55
Library Catalog	Neural Information Processing Systems
Book Title	Advances in Neural Information Processing Systems 30
Date Added	26.3.2018, 15:21:55
Modified	26.3.2018, 15:21:55

Attachments

- NIPS Full Text PDF
- NIPS Snapshort

Zero-Shot Task Generalization with Multi-Task Deep Reinforcement Learning

Type	Conference Paper
Author	Junhyuk Oh
Author	Satinder Singh
Author	Honglak Lee
Author	Pushmeet Kohli
URL	http://proceedings.mlr.press/v70/oh17a.html
Pages	2661-2670
Date	2017/07/17
Accessed	26.3.2018, 15:16:12
Library Catalog	proceedings.mlr.press
Conference Name	International Conference on Machine Learning
Language	en
Abstract	As a step towards developing zero-shot task generalization capabilities in reinforcement learning (RL), we introduce a new RL problem where the agent should learn to execute sequences of instructio...
Proceedings Title	International Conference on Machine Learning

Date Added 26.3.2018, 15:16:12
Modified 26.3.2018, 15:16:12

Attachments

- Full Text PDF
- Snapshot