Papers worth reading

Learning to learn by gradient descent by gradient descent

https://arxiv.org/abs/1606.04474v2

Google DeepMind

“We show how the design of an optimization algorithm can be cast as a learning problem, allowing the algorithm to learn to exploit structure in the problems of interest in an automatic way. “

Unifying Count-Based Exploration and Intrinsic Motivation

https://arxiv.org/abs/1606.01868v2

Google DeepMind

“We transform these pseudo-counts into intrinsic rewards and obtain significantly improved exploration in a number of hard games, including the infamously difficult MONTEZUMA’S REVENGE.“

Poisoning Attacks against Support Vector Machines

https://arxiv.org/abs/1206.6389v3

Biggio, Nelson, Laskov

“As we demonstrate, an intelligent adversary can, to some extent, predict the change of the SVM’s decision function due to malicious input and use this ability to construct malicious data.”

Listen, Attend and Spell

https://arxiv.org/abs/1508.01211v2

W. Chan Carnegie Mellon et al. Google Brain

“We present Listen, Attend and Spell (LAS), a neural network that learns to transcribe speech utterances to characters…without a dictionary or a language model“

Attend, Infer, Repeat: Fast Scene Understanding with Generative Models

https://arxiv.org/abs/1603.08575v3

G.E. Hinton et al. Google DeepMind

“We show that such models learn to identify multiple objects – counting, locating and classifying the elements of a scene – without any supervision“

Gated Feedback Recurrent Neural Networks

https://arxiv.org/abs/1606.04474v2

Yoshua Bengio et al. Université de Montreal, CIFAR

“We show how the design of an optimization algorithm can be cast as a learning problem, allowing the algorithm to learn to exploit structure in the problems of interest in an automatic way. “

ADVERSARIAL FEATURE LEARNING

https://arxiv.org/abs/1605.09782v6

Jeff Donahue, Trevor Darrell UC Berkeley & Philipp Krähenbühl UT, Austin

“We propose Bidirectional Generative Adversarial Networks (BiGANs) as a means of learning this inverse mapping, and demonstrate that the resulting learned feature representation is useful for auxiliary supervised discrimination tasks“

Gradient-based Hyperparameter Optimization through Reversible Learning

https://arxiv.org/abs/1502.03492v2

Google DeepMind

“We compute hyperparameter gradients by exactly reversing the dynamics of stochastic gradient descent with momentum.“

Mask R-CNN

https://arxiv.org/abs/1703.06870v1

Google DeepMind

“Our approach efficiently detects objects in an image while simultaneously generating a high-quality segmentation mask for each instance.“

ImageNet Classification with Deep Convolutional Neural Networks (AlexNet)

https://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks.pdf

Alex Krizhevsky Ilya Sutskever Geoffrey E. Hinton, University of Toronto

“The neural network, which has 60 million parameters and 650,000 neurons, consists of five convolutional layers, some of which are followed by max-pooling layers, and three fully-connected layers with a final 1000-way softmax.“

Reducing the Dimensionality of Data with Neural Networks

Science 28 Jul 2006: Vol. 313, Issue 5786, pp. 504-507

G. E. Hinton and R. R. Salakhutdinov

“We describe an effective way of initializing the weights that allows deep autoencoder networks to learn low-dimensional codes that work much better than principal components analysis as a tool to reduce the dimensionality of data.“

REINFORCEMENT LEARNING WITH UNSUPERVISED AUXILIARY TASKS

https://arxiv.org/abs/1611.05397v1

Jaderberg, Mnih et al. DeepMind

“We introduce an agent that also maximises many other pseudo-reward functions simultaneously by reinforcement learning….Our agent significantly outperforms the previous state-of-the-art on ... Labyrinth“

Decoupled Neural Interfaces using Synthetic Gradients

https://arxiv.org/abs/1611.05397v1

Jaderberg et al. DeepMind

“We demonstrate that in addition to predicting gradients, the same framework can be used to predict inputs, resulting in models which are decoupled in both the forward and backwards pass – amounting to independent networks which co-learn such that they can be composed into a single functioning corporation.“

DelugeNets: Deep Networks with Massive and Flexible Cross-layer Information Inflows

https://arxiv.org/abs/1611.05552v4

Jason Kuen, Xiangfei Kong, and Gang Wang, Nanyang Technological University, Singapore.

“We propose Deluge Networks (DelugeNets), a novel class of neural networks facilitating massive cross-layer information inflows from preceding layers to succeeding layers...efficiently established through cross-layer depthwise convolutional layers with learnable filters.“

Hierarchical Deep Reinforcement Learning: Integrating Temporal Abstraction and Intrinsic Motivation

https://arxiv.org/abs/1604.06057v2

Tejas D. Kulkarni, Karthik R. Narasimhan, Joshua B. Tenenbaum, Ardavan Saeedi MIT

“We present hierarchical-DQN (h-DQN), a framework to integrate hierarchical value functions, operating at different temporal scales, with intrinsically motivated deep reinforcement learning. A top-level value function learns a policy over intrinsic goals, and a lower-level function learns a policy over atomic actions to satisfy the given goals.“

Deep Generative Image Models using a Laplacian Pyramid of Adversarial Networks

https://arxiv.org/abs/1506.05751v1

Emily Denton Courant Institute New York University, et al Facebook AI Research

“We introduce a generative parametric model capable of producing high quality samples of natural images. Our approach uses a cascade of convolutional networks within a Laplacian pyramid framework to generate images in a coarse-to-fine fashion“

Deep learning

doi:10.1038/nature14539

Yann LeCun, Yoshua Bengio & Geoffrey Hinton

“Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.“

COMBATING DEEP REINFORCEMENT LEARNING’S SISYPHEAN CURSE WITH INTRINSIC FEAR

https://arxiv.org/abs/1611.01211v4

Zachary C. Lipton UC San Diego, et al. Microsoft Research

“We demonstrate unacceptable performance of deep Q-networks on two toy problems. We then introduce intrinsic fear, a method that mitigates these problems by avoiding dangerous states“

Long Short-Term Memory-Networks for Machine Reading

https://arxiv.org/abs/1601.06733v7

Jianpeng Cheng, Li Dong and Mirella Lapata, University of Edinburgh

“We propose a machine reading simulator which processes text incrementally from left to right and performs shallow reasoning with memory and attention.“