Papers

Tom Rochette <tom.rochette@coreteks.org>

November 2, 2024 — 36c8eb68

The following lists papers I've read and reviewed or made notes for.

1 Artificial General Intelligence

1.1 Alex Graves

- Automated Curriculum Learning for Neural Networks (2017)
- Multi-Dimensional Recurrent Neural Networks (2007)
- Neural Turing Machines (2014)
- Offline Handwriting Recognition with Multidimensional Recurrent Neural Networks (2009)

1.2 Andrew Barto

- Intrinsically Motivated Learning of Hierarchical Collections of Skills (2004)

1.3 Andrew Carlson

- Toward an Architecture for Never-Ending Language Learning (2010)

1.4 Ashish Vaswani

- Attention Is All You Need (2017)

1.5 Barret Zoph

- Neural Architecture Search with Reinforcement Learning (2016)

1.6 Bradly Stadie

- Third-Person Imitation Learning (2017)

1.7 Burr Settles

- A Trainable Spaced Repetition Model for Language Learning (2016)

1.8 Carlos Florensa

- Reverse Curriculum Generation for Reinforcement Learning (2017)

1.9 Cristian Bucila

- Model Compression (2006)

1.10 David Silver

- Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm (2017)
- Mastering the Game of Go with Deep Neural Networks and Tree Search (2016)
- Mastering the game of Go without human knowledge (2017)
- The Predictron: End-To-End Learning and Planning (2017)

1.11 Dzmitry Bahdanau

- Neural Machine Translation by Jointly Learning to Align and Translate (2015)

1.12 Eric Laukien

- Feynman Machine: The Universal Dynamical Systems Computer (2016)

1.13 Geoffrey Hinton

- Distilling the Knowledge in a Neural Network (2015)

1.14 Greg Linden

- Amazon.com Recommendations - Item-to-Item Collaborative Filtering (2003)

1.15 Ian Goodfellow

- Generative Adversarial Nets (2014)

1.16 J. R. Quinlan

- Induction of Decision Trees (1986)

1.17 Jacob Devlin

- RobustFill: Neural Program Learning under Noisy I/O (2017)

1.18 Karl Friston

- The free-energy principle: a unified brain theory? (2010)

1.19 Kelvin Xu

- Show, Attend and Tell: Neural Image Caption Generation with Visual Attention (2015)

1.20 Ken Kansky

- Schema Networks: Zero-shot Transfer with a Generative Causal Model of Intuitive Physics (2017)

1.21 Leo Breiman

- Bagging Predictors (1996)

1.22 Levente Kocsis

- Bandit based Monte-Carlo Planning (2006)

1.23   Łukasz Kaiser

- One Model To Learn Them All (2017)

1.24   Manuel Lopes

- The Strategic Student Approach for Life-Long Exploration and Learning (2012)

1.25   Matej Balog

- DeepCoder: Learning to Write Programs (2016)

1.26   Matteo Hessel

- Rainbow: Combining Improvements in Deep Reinforcement Learning (2017)

1.27   Max Jaderberg

- Reading Text in the Wild with Convolutional Neural Networks (2014)

1.28   Miltiadis Allamanis

- SmartPaste: Learning to Adapt Source Code (2017)

1.29   Misha Denil

- Programmable Agents (2017)

1.30   Nal Kalchbrenner

- Grid Long Short-Term Memory (2015)

1.31   Neil Rabinowitz

- Machine Theory of Mind (2018)

1.32   Oriol Vinyals

- Starcraft II: A New Challenge for Reinforcement Learning (2017)

1.33   Paul Christiano

- Deep Reinforcement Learning from Human Preferences (2017)

1.34   Ralf Herbrich

- Learning and Generalization: Theoretical Bounds (2001)

1.35   Ronen Brafman

- R-max – A General Polynomial Time Algorithm for Near-Optimal Reinforcement Learning (2002)

1.36   Sercan Arik

- Deep Voice: Real-time Neural Text-to-Speech (2017)
- Deep Voice 2: Multi-Speaker Neural Text-to-Speech (2017)

1.37 Théodore Bluche

- Scan, Attend and Read: End-to-End Handwritten Paragraph Recognition with MDLSTM Attention (2016)

1.38 Thomas Anthony

- Thinking Fast and Slow with Deep Learning and Tree Search (2017)

1.39 Tom Mitchell

- Never-ending learning (2018)

1.40 Tomas Mikolov

- Efficient Estimation of Word Representations in Vector Space (2013)

1.41 Volodymyr Mnih

- Human-level control through deep reinforcement learning (2015)
- Playing Atari with Deep Reinforcement Learning (2013)

1.42 Wei Ping

- Deep Voice 3: 2000-Speaker Neural Text-to-Speech (2017)

1.43 Xuan-Bach Le

- History Driven Program Repair (2016)

1.44 Yingfei Xiong

- Precise Condition Synthesis for Program Repair (2017)

1.45 Yoav Freund

- A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting (1995)

1.46 Yoshua Bengio

- Curriculum Learning (2009)
- The Consciousness Prior (2017)

1.47 Yuxuan Wang

- Tacotron: Towards End-to-End Speech Synthesis (2017)