

Deep Reinforcement Learning

Deep Learning — Unit 10

Dr. Jon Krohn

`jon@untapt.com`

December 16th, 2017

Outline

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Outline

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Outline

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Outline

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Outline

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Deep Learning Categories

Intro

Theory

In Practice

Applications

AI Revolution

- 1 supervised learning
- 2 unsupervised learning
- 3 reinforcement learning

Deep Learning Categories

Intro

Theory

In Practice

Applications

AI Revolution

- 1 supervised learning
- 2 unsupervised learning
- 3 reinforcement learning

Deep Learning Categories

Intro

Theory

In Practice

Applications

AI Revolution

- 1 supervised learning
- 2 unsupervised learning
- 3 reinforcement learning

Supervised Learning

Intro

Theory

In Practice

Applications

AI Revolution

- **have x and y**
- goal is to approximate function that predicts y with x
- e.g.: classification (MNIST digits, IMDB sentiment), regression

Supervised Learning

Intro

Theory

In Practice

Applications

AI Revolution

- have x and y
- goal is to approximate function that predicts y with x
- e.g.: classification (MNIST digits, IMDB sentiment), regression

Supervised Learning

Intro

Theory

In Practice

Applications

AI Revolution

- have x and y
- goal is to approximate function that predicts y with x
- e.g.: classification (MNIST digits, IMDB sentiment), regression

Unsupervised Learning

Intro

Theory

In Practice

Applications

AI Revolution

- have x without y labels
- goal is to learn data's underlying hidden structure
- e.g.: word vectors, GANs

Unsupervised Learning

Intro

Theory

In Practice

Applications

AI Revolution

- have x without y labels
- goal is to learn data's underlying hidden structure
- e.g.: word vectors, GANs

Unsupervised Learning

Intro

Theory

In Practice

Applications

AI Revolution

- have x without y labels
- goal is to learn data's underlying hidden structure
- e.g.: word vectors, GANs

Reinforcement Learning


Intro

Theory

In Practice

Applications

AI Revolution



- an *agent* takes action in an *environment*
- environment returns *state*, i.e., *reward* at t , state at $t + 1$
- repeat

Reinforcement Learning


Intro

Theory

In Practice

Applications

AI Revolution



- an *agent* takes action in an *environment*
- environment returns *state*, i.e., *reward* at t , state at $t + 1$
- repeat

Reinforcement Learning


Intro

Theory

In Practice

Applications

AI Revolution



- an *agent* takes action in an *environment*
- environment returns *state*, i.e., *reward* at t , *state* at $t + 1$
- repeat

Deep Q-Learning

Mnih et al. (2015)


Intro

Theory

In Practice

Applications

AI Revolution



[Atari Games]

Deep Q-Learning

Mnih et al. (2015)


Intro

Theory

In Practice

Applications


AI Revolution



[Atari Games]

Deep Q-Learning

Mnih et al. (2015)



[Atari Games]

Environments

Intro

Theory

In Practice

Applications

AI Revolution

[OpenAI Gym]

[Google DeepMind Lab]

Outline

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Introduction
- 2 Essential Theory**
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution

Essential Theory

Intro

Theory

In Practice

Applications

AI Revolution

- 1 **Markov Decision Processes**
- 2 Q-Learning
- 3 Policy Gradients
- 4 Actor-Critic Algorithm

[[GitHub Content](#)]

Essential Theory

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Markov Decision Processes
- 2 Q-Learning
- 3 Policy Gradients
- 4 Actor-Critic Algorithm

[GitHub Content]

Essential Theory

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Markov Decision Processes
- 2 Q-Learning
- 3 Policy Gradients
- 4 Actor-Critic Algorithm

[GitHub Content]

Essential Theory

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Markov Decision Processes
- 2 Q-Learning
- 3 Policy Gradients
- 4 Actor-Critic Algorithm

[GitHub Content]

Outline

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning**
- 4 Contemporary Applications
- 5 The AI Revolution

The Cartpole Game

OpenAI


Intro

Theory

In Practice

Applications

AI Revolution



- 1** cart position
- 2 cart velocity
- 3 pole angle
- 4 pole angular velocity

The Cartpole Game

OpenAI


Intro

Theory

In Practice

Applications

AI Revolution



- 1 cart position
- 2 cart velocity
- 3 pole angle
- 4 pole angular velocity

The Cartpole Game

OpenAI


Intro

Theory

In Practice

Applications

AI Revolution



- 1 cart position
- 2 cart velocity
- 3 pole angle
- 4 pole angular velocity

The Cartpole Game

OpenAI


Intro

Theory

In Practice

Applications

AI Revolution



- 1 cart position
- 2 cart velocity
- 3 pole angle
- 4 pole angular velocity

Installation

(outside my container)

- 1 <https://github.com/openai/gym>
- 2 anaconda
- 3 `(pip install gym[atari])`
- 4 `pip uninstall pygame`
- 5 `pip install pygame==1.2.4`
- 6 `jupyter notebook` in `nyc-ds-academy` directory

Installation

(outside my container)

- 1 <https://github.com/openai/gym>
- 2 `anaconda`
- 3 `(pip install gym[atari])`
- 4 `pip uninstall pygame`
- 5 `pip install pygame==1.2.4`
- 6 `jupyter notebook` in `nyc-ds-academy` directory

Installation

(outside my container)

- 1 `https://github.com/openai/gym`
- 2 `anaconda`
- 3 `(pip install gym[atari])`
- 4 `pip uninstall pygame`
- 5 `pip install pygame==1.2.4`
- 6 `jupyter notebook` in `nyc-ds-academy` directory

Installation

(outside my container)

- 1 `https://github.com/openai/gym`
- 2 `anaconda`
- 3 `(pip install gym[atari])`
- 4 `pip uninstall pygame`
- 5 `pip install pygame==1.2.4`
- 6 `jupyter notebook` in `nyc-ds-academy` directory

Installation

(outside my container)

- 1 `https://github.com/openai/gym`
- 2 `anaconda`
- 3 `(pip install gym[atari])`
- 4 `pip uninstall pygame`
- 5 `pip install pygame==1.2.4`
- 6 `jupyter notebook` in `nyc-ds-academy` directory

Installation

(outside my container)

- 1 `https://github.com/openai/gym`
- 2 `anaconda`
- 3 `(pip install gym[atari])`
- 4 `pip uninstall pygame`
- 5 `pip install pygame==1.2.4`
- 6 `jupyter notebook` in `nyc-ds-academy` directory

Intro

Theory

In Practice

Applications

AI Revolution

[*cartpole* notebook]

Outline

Intro

Theory

In Practice

Applications

AI Revolution

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications**
- 5 The AI Revolution

AlphaGO

Silver et al. (2016)



[AlphaGo Zero blog post]

Genetic Algorithms

[Mar/I/O]

Automated Experimentation

Intro

Theory

In Practice

Applications

AI Revolution

[OpenAI Lab]

Ready-Made Agents

[Coach]

Human Imitation Learning

Warneken & Tomasello (2006)

[videos]

Machine Imitation Learning

Finn, Levine & Abbeel

Intro

Theory

In Practice

Applications

AI Revolution

[videos]

Software 2.0

Karpathy (2017)

[Medium]

Intro

Theory

In Practice

Applications

AI Revolution

Outline

- 1 Introduction
- 2 Essential Theory
- 3 Cart Pole Deep Q-Learning
- 4 Contemporary Applications
- 5 The AI Revolution**

Jeanne Calment

Intro

Theory

In Practice

Applications

AI Revolution



Jeanne Calment

(1875-1997 — i.e., 122 years)

21



1896

121



1996

Life in the Year 2138

Intro

Theory

In Practice

Applications

AI Revolution



The AI Revolution

Hasn't Even Begun


Intro

Theory


In Practice

Applications

AI Revolution



Biodiversity during the Phanerozoic




Intro

Theory

In Practice

Applications

AI Revolution



- 1 data doubling every 18 months
- 2 processing power cost halving every two years
- 3 cheap sensors appearing everywhere
- 4 Deep Learning techniques refined in academia and in industry

- 1 data doubling every 18 months
- 2 processing power cost halving every two years
- 3 cheap sensors appearing everywhere
- 4 Deep Learning techniques refined in academia and in industry

- 1 data doubling every 18 months
- 2 processing power cost halving every two years
- 3 cheap sensors appearing everywhere
- 4 Deep Learning techniques refined in academia and in industry

- 1 data doubling every 18 months
- 2 processing power cost halving every two years
- 3 cheap sensors appearing everywhere
- 4 Deep Learning techniques refined in academia and in industry


Intro

Theory

In Practice

Applications

AI Revolution



Intro

Theory

In Practice

Applications

AI Revolution



Jeanne Calment

Intro

Theory

In Practice

Applications

AI Revolution



Thiel & Masters (2014)

Intro


Theory

In Practice


Applications

AI Revolution


RECURRENT COLLAPSE



EXTINCTION



PLATEAU



TAKEOFF



Thiel & Masters (2014)

Intro


Theory

In Practice


Applications

AI Revolution


RECURRENT COLLAPSE



EXTINCTION



PLATEAU



TAKEOFF



Thiel & Masters (2014)

Intro


Theory

In Practice


Applications

AI Revolution


RECURRENT COLLAPSE



EXTINCTION



PLATEAU



TAKEOFF




Intro

Theory


In Practice

Applications

AI Revolution




Jan-Dec Global Mean Temperature over Land & Ocean



Still lots to go around

Nuclear-warhead stockpile, '000




*Strategic Arms Limitation Treaty [†]Strategic Arms Reduction Treaty

[‡]Strategic Offensive Reductions Treaty [§]Includes 5,500 retired warheads awaiting dismantling


**Excludes retired warheads

Source: Bulletin of the Atomic Scientists


Pinker & Mack (2014)




Source: Early Warning Project, <https://earlywarning.worldpress.com/>; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2006). Available at SSRN: <http://ssrn.com/abstract=1033426> or <http://dx.doi.org/10.2339/arm-03-0626>



Sources: U.S. FBI Uniform Crime Reports, England (including Wales) U.K. Office for National Statistics, World U.N. Office on Drugs and Crime, reported by U.N. Economic and Social Council's World Crime Trends and Emerging Issues and responses to the field of crime prevention and social justice." Feb. 12, 2014, Figure 1. The percentages were converted to homicide rates by setting the 2012 rate at 6.2, the figure reported in the UNODC Global Study on Homicide 2013, Page 12.



Source: Utopia Conflict Data Program Battle-Related Deaths v5-2014, 1999-2013, "Best Estimates" Missing estimates for 2013 for the Syria civil war calculated from a revised ICDIP "Low Estimate" and a "High Estimate" from the Peace Research Institute Oslo (both obtained in consultation with Erik Meisler of ICDIP, The "High Estimate" is for May 2013-April 2014, rather than calendar year 2013, and comes from Page 9 of H. Price, A. Gendau, & P. Ball, "Underestimated: Statistical Analysis of Documentation of Killings in the Syria Arab Republic," Human Rights Data Analysis Group, <http://www.hrdag.org/Documenta/Countries/SY/HRDAG/updates/ReportAug2014.pdf>. World population figures from U.S. Census Bureau, <http://www.census.gov/population/international/data/worldpop/>




Rates for physical and sexual abuse are per 100,000 children younger than 18. Rates for violent victimization at school are per 10,000 children age 5-12.


Sources: Physical and sexual abuse, National Child Abuse and Neglect Data System, analyzed by David Brothman, 2014, in "Trends in Child Welfare," Presentation at the Carney Institute Policy Series, March 30, 2014. Victimization at school, Bureau of Economic Analysis (top.gov), using the National Crime Victimization Survey, Victimization Analysis Tool.




Pinker & Mack (2014)




Source: Early Warning Project, <https://earlywarning.worldpress.com/>; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2006). Available at SSRN: <http://ssrn.com/abstract=1034426> or <http://dx.doi.org/10.2339/arm-03-0626>



Sources: U.S. FBI Uniform Crime Reports, England (Including Wales): U.K. Office for National Statistics, World: U.N. Office on Drugs and Crime, reported by U.N. Economic and Social Council's World Crime Trends and Emerging Issues and responses to the field of crime prevention and social justice, Feb. 12, 2014, Figure 1. The percentages were converted to homicide rates by setting the 2012 rate at 6.2, the figure reported in the UNODC Global Study on Homicide 2013, Page 12.



Source: Utopia Conflict Data Program Battle-Related Deaths v5-2014, 1999-2013, "Best Estimates" Missing estimate for 2013 for the Syria civil war calculated from a revised ICDIP "Low Estimate" and a "High Estimate" from the Peace Research Institute Oslo (both obtained in consultation with Erik Meisler of ICDIP, The "High Estimate" is for May 2013-April 2014, rather than calendar year 2013, and comes from Page 9 of H. Price, A. Gendau, & P. Ball, "Underestimated: Statistical Analysis of Documentation of Killings in the Syrian Arab Republic," Human Rights Data Analysis Group, <http://www.hrdag.org/Documenta/Countries/SY/HRDAG/updates/ReportAug2014.pdf>. World population figures from U.S. Census Bureau, <http://www.census.gov/population/international/data/worldpop/>




Rates for physical and sexual abuse are per 100,000 children younger than 18. Rates for violent victimization at school are per 10,000 children age 5-12.


Sources: Physical and sexual abuse, National Child Abuse and Neglect Data System, analyzed by David Brothman, 2014, in "Trends in Child Welfare," Presentation at the Carney Institute Policy Series, March 30, 2014. Victimization at school, Bureau of Economic Analysis (top.gov), using the National Crime Victimization Survey, Victimization Analysis Tool.




Pinker & Mack (2014)




Source: Early Warning Project, <https://earlywarning.worldpress.com/>; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2006). Available at SSRN: <http://ssrn.com/abstract=1034426> or <http://dx.doi.org/10.2339/arm-03-0626>



Sources: U.S. FBI Uniform Crime Reports, England (including Wales): U.K. Office for National Statistics, World: U.N. Office on Drugs and Crime, reported by U.N. Economic and Social Council's World Crime Trends and Emerging Issues and responses to the field of crime prevention and social justice, Feb. 12, 2004, Figure 1. The percentages were converted to homicide rates by setting the 2012 rate at 6.2, the figure reported in the UNODC Global Study on Homicide 2013, Page 12.



Source: Utopia Conflict Data Program Battle-Related Deaths v5-2014, 1999-2013, "Best Estimates" Missing estimate for 2013 for the Syria civil war calculated from a revised ICDIP "Low Estimate" and a "High Estimate" from the Peace Research Institute Oslo (both obtained in consultation with Erik Meisler of ICDIP, The "High Estimate" is for May 2013-April 2014, rather than calendar year 2013, and comes from Page 9 of H. Price, A. Gendau, & P. Ball, "Underestimated: Statistical Analysis of Documentation of Killings in the Syria Arab Republic," Human Rights Data Analysis Group, <http://www.hrdag.org/Documenta/Countries/SY/HRDAG/updates/ReportAug2014.pdf>. World population figures from U.S. Census Bureau, <http://www.census.gov/population/international/data/worldpop/>




Rates for physical and sexual abuse are per 100,000 children younger than 18. Rates for violent victimization at school are per 10,000 children age 5-12.

Sources: Physical and sexual abuse, National Child Abuse and Neglect Data System, analyzed by David Brothman, 2014, in "Trends in Child Welfare," Presentation at the Carney Institute Policy Series, March 30, 2014. Victimization at school, Bureau of Economic Analysis (top.gov), using the National Crime Victimization Survey, Victimization Analysis Tool.




World population living in extreme poverty, 1820 to 2015



Data source: World Poverty in absolute numbers (Max Roser based on World Bank and Bourguignon and Morrisson (2002))

Literate and illiterate world population, 1800 to 2014




Data source: Literate World Population (Our World In Data based on OECD and UNESCO)

Number of world citizens living under different political regimes




The Polity IV score captures the type of political regime for each country on a range from -10 (full autocracy) to +10 (full democracy). Regimes that fall into the middle of this spectrum are called anocracies.




Data source: World Population by Political Regime they live in (by Our World In Data)


Global child mortality, 1800 to 2015


Share of the world population dying and surviving the first 5 years of life.



Data source: Global child mortality (since 1800) based on Gapminder and World Bank








Intro


Theory

In Practice

Applications

AI Revolution






Intro


Theory

In Practice


Applications

AI Revolution





waitbutwhy.com



waitbutwhy.com


Intro

Theory

In Practice

Applications

AI Revolution



waitbutwhy.com