

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Deep Reinforcement Learning

Deep Learning — Unit 10

Dr. Jon Krohn

jon@untapt.com

Slides available at jonkrohn.com/talks

December 1st, 2018

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
Jeanne Calment
Culture Dish
Enlightenment
The View From Here

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Categories of A.I.

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 ANI: *Narrow*
- 2 AGI: *General*
- 3 ASI: *Super*

Categories of A.I.

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 ANI: *Narrow*
- 2 AGI: *General*
- 3 ASI: *Super*

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Categories of A.I.

- 1 ANI: *Narrow*
- 2 AGI: *General*
- 3 ASI: *Super*

Deep Learning Categories

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

Deep Learning Categories

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

Deep Learning Categories

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

Supervised Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- **have x and y**
- goal is to approximate function that predicts y with x
- e.g.: classification (MNIST digits, IMDB sentiment), regression (sales of a product, future value of an asset)

Supervised Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- have x and y
- goal is to approximate function that predicts y with x
- e.g.: classification (MNIST digits, IMDB sentiment), regression (sales of a product, future value of an asset)

Supervised Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- have x and y
- goal is to approximate function that predicts y with x
- e.g.: classification (MNIST digits, IMDB sentiment), regression (sales of a product, future value of an asset)

Unsupervised Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

Unsupervised Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

Unsupervised Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

Reinforcement Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



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

Reinforcement Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

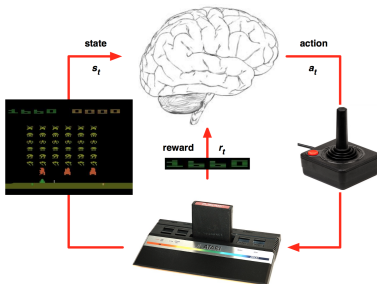
AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



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

Reinforcement Learning

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



- 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. (2013)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- “Playing Atari with Deep RL”
 - NIPS and arXiv paper
 - played seven Atari 2600 games
 - beat previous ML approaches on six
 - beat human expert on three

Deep Q-Learning

Mnih et al. (2013)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- “Playing Atari with Deep RL”
- NIPS and arXiv paper
- played seven Atari 2600 games
- beat previous ML approaches on six
- beat human expert on three

Deep Q-Learning

Mnih et al. (2013)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- “Playing Atari with Deep RL”
- NIPS and arXiv paper
- played seven Atari 2600 games
 - beat previous ML approaches on six
 - beat human expert on three

Deep Q-Learning

Mnih et al. (2013)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- “Playing Atari with Deep RL”
- NIPS and arXiv paper
- played seven Atari 2600 games
- beat previous ML approaches on six
- beat human expert on three

Deep Q-Learning

Mnih et al. (2013)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- “Playing Atari with Deep RL”
- NIPS and arXiv paper
- played seven Atari 2600 games
- beat previous ML approaches on six
- beat human expert on three

Deep Q-Learning

Mnih et al. (2015)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



[Atari Games]

Deep Q-Learning

Mnih et al. (2015)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



[Atari Games]

Deep Q-Learning

Mnih et al. (2015)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



[Atari Games]

AlphaGo

Silver et al. (2016)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



AlphaGo Zero

Silver et al. (2017a)

Intro



DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment
Culture Dish
Enlightenment
The View From Here



AlphaZero

Silver et al. (2017b)

Intro


DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

- Jeanne Calment
- Culture Dish
- Enlightenment
- The View From Here



AlphaZero

Silver et al. (2017b)

Intro


DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment
Culture Dish
Enlightenment
The View From Here



Real-World End-to-End

Levine et al. (2016)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment


The View From Here




(a) hanger



(b) cube



(c) hammer



(d) bottle

Environments

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 OpenAI Gym
- 2 DeepMind Lab
- 3 Unity

...or create your own with, e.g., the PyGame Learning Environment

Environments

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 OpenAI Gym
- 2 DeepMind Lab
- 3 Unity

...or create your own with, e.g., the PyGame Learning Environment

Environments

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 OpenAI Gym
- 2 DeepMind Lab
- 3 Unity

...or create your own with, e.g., the PyGame Learning Environment

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Essential Deep Q-Learning Theory

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

[whiteboarding]

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning**
- 4 Beyond DQN
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

The Cartpole Game

OpenAI

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



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

The Cartpole Game

OpenAI

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



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

The Cartpole Game

OpenAI

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



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

The Cartpole Game

OpenAI

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



- 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

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

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

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

[*cartpole* notebook]

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN**
- 5 The AI Revolution
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Genetic Algorithms

[Mar/I/O]

Ready-Made Agents

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



[Coach]

Automated Experimentation

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

[OpenAI Lab]

Human Imitation Learning

Warneken & Tomasello (2006)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

[videos]

Machine Imitation Learning

Finn, Levine & Abbeel

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

[videos]

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Software 2.0 Pros

Karpathy (2017)

- 1 Computational Homogeneity
- 2 Constant Running Time
- 3 Constant Memory Use
- 4 Portable Across Devices
- 5 Easy
- 6 Superior

[Medium]

Software 2.0 Pros

Karpathy (2017)

- 1 Computational Homogeneity
- 2 Constant Running Time
- 3 Constant Memory Use
- 4 Portable Across Devices
- 5 Easy
- 6 Superior

[Medium]

Software 2.0 Pros

Karpathy (2017)

- 1 Computational Homogeneity
- 2 Constant Running Time
- 3 Constant Memory Use
- 4 Portable Across Devices
- 5 Easy
- 6 Superior

[Medium]

Software 2.0 Pros

Karpathy (2017)

- 1 Computational Homogeneity
- 2 Constant Running Time
- 3 Constant Memory Use
- 4 Portable Across Devices
- 5 Easy
- 6 Superior

[Medium]

Software 2.0 Pros

Karpathy (2017)

- 1 Computational Homogeneity
- 2 Constant Running Time
- 3 Constant Memory Use
- 4 Portable Across Devices
- 5 Easy
- 6 Superior

[Medium]

Software 2.0 Pros

Karpathy (2017)

- 1 Computational Homogeneity
- 2 Constant Running Time
- 3 Constant Memory Use
- 4 Portable Across Devices
- 5 Easy
- 6 Superior

[Medium]

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 The AI Revolution**
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 **The AI Revolution**
Jeanne Calment
Culture Dish
Enlightenment
The View From Here

Jeanne Calment

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Jeanne Calment

(1875-1997 — i.e., 122 years)

21



1896

121



1996

Life in the Year 2140

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish


Enlightenment

The View From Here

Everybody's choice

Mobile-phone penetration

Per 100 people, worldwide



Source: ITU

Economist.com

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

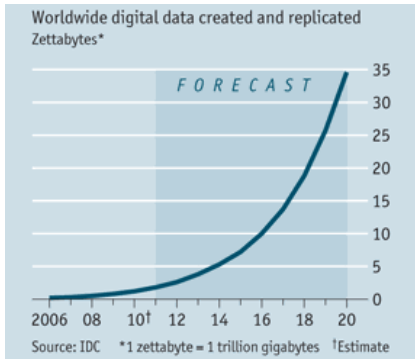
AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

The AI Revolution

Hasn't Even Begun

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Factors Driving A.I.

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Data, incl. Sensors Everywhere
- 2 Computing Power
- 3 Algorithms
- 4 Infrastructure

[Wait But Why]

Factors Driving A.I.

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Data, incl. Sensors Everywhere
- 2 Computing Power
- 3 Algorithms
- 4 Infrastructure

[Wait But Why]

Factors Driving A.I.

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Data, incl. Sensors Everywhere
- 2 Computing Power
- 3 Algorithms
- 4 Infrastructure

[Wait But Why]

Factors Driving A.I.

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Data, incl. Sensors Everywhere
- 2 Computing Power
- 3 Algorithms
- 4 Infrastructure

[Wait But Why]

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Categories of A.I.

- 1 ANI: *Narrow*
- 2 AGI: *General*
- 3 ASI: *Super*

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Categories of A.I.

- 1 ANI: *Narrow*
- 2 AGI: *General*
- 3 ASI: *Super*

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

Categories of A.I.

- 1 ANI: *Narrow*
- 2 AGI: *General*
- 3 ASI: *Super*

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Requires Many, Many Samples
- 2 Black Box
- 3 Doesn't Leverage Knowledge of the World
- 4 Correlation = Causation
- 5 Unintuitive / Embarrassing / Adversarial Failures

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Requires Many, Many Samples
- 2 Black Box
- 3 Doesn't Leverage Knowledge of the World
- 4 Correlation = Causation
- 5 Unintuitive / Embarrassing / Adversarial Failures

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Requires Many, Many Samples
- 2 Black Box
- 3 Doesn't Leverage Knowledge of the World
- 4 Correlation = Causation
- 5 Unintuitive / Embarrassing / Adversarial Failures

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Requires Many, Many Samples
- 2 Black Box
- 3 Doesn't Leverage Knowledge of the World
- 4 Correlation = Causation
- 5 Unintuitive / Embarrassing / Adversarial Failures

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Requires Many, Many Samples
- 2 Black Box
- 3 Doesn't Leverage Knowledge of the World
- 4 Correlation = Causation
- 5 Unintuitive / Embarrassing / Adversarial Failures

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 **The AI Revolution**
Jeanne Calment
Culture Dish
Enlightenment
The View From Here

Culture Dish

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment


Culture Dish

Enlightenment

The View From Here

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

Thiel & Masters (2014)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution


Jeanne Calment

Culture Dish


Enlightenment

The View From Here


RECURRENT COLLAPSE



EXTINCTION



PLATEAU



TAKEOFF



Thiel & Masters (2014)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution


Jeanne Calment

Culture Dish


Enlightenment

The View From Here


RECURRENT COLLAPSE



EXTINCTION



PLATEAU



TAKEOFF



Thiel & Masters (2014)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution


Jeanne Calment

Culture Dish


Enlightenment

The View From Here


RECURRENT COLLAPSE



EXTINCTION



PLATEAU



TAKEOFF



Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 **The AI Revolution**
Jeanne Calment
Culture Dish
Enlightenment
The View From Here

Pinker & Mack (2014)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution


Jeanne Calment

Culture Dish

Enlightenment


The View From Here

**PREVALENCE OF MASS KILLINGS
1945-2013**



Source: Early Warning Project, <https://ppparfwarwarning.worldpress.com/>; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2006). Available at SSRN: <https://ssrn.com/abstract=1034626> or <http://dx.doi.org/10.2339/arm-034626>

**HOMICIDE RATES IN THE US AND ENGLAND
1967-2013, AND THE WORLD, 2003-2012**




Sources: U.S. FBI Uniform Crime Reports, England (Including Wales) U.K. Office for National Statistics, World U.K. Office on Drugs and Crime, reported by U.N. Economic and Social Council's World Crime Trends and Emerging Issues and responses in 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.

**RATE OF BATTLE DEATHS IN ARMED CONFLICTS
1946-2013**



Source: Utopia Conflict Data Program Battle-Related Deaths Dataset 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, "Unsettled Status: A Statistical Analysis of Documentation of Killings in the Syria Arab Republic," Human Rights Data Analysis Group, http://www.hrdag.org/Documents/Countries/SY/HRDAG_SyriaReportAug2014.pdf, World population figures from U.S. Census Bureau, http://www.census.gov/population/international/data/worldpop/stable_population.php

**VICTIMIZATION OF CHILDREN IN THE US
1990-2012**



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 31, 2014. Victimization at school, Bureau of Economic Analysis (top.gov), using the National Crime Victimization Survey, Victimization Analysis Tool.



Pinker & Mack (2014)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment


The View From Here

**PREVALENCE OF MASS KILLINGS
1945-2013**




Source: Early Warning Project, <https://earlywarning.worpress.com/>; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Oct. 1, 2006). Available at SSRN: <https://ssrn.com/abstract=1033426> or <http://dx.doi.org/10.2339/arm-031426>

**HOMICIDE RATES IN THE US AND ENGLAND
1967-2013, AND THE WORLD, 2003-2012**




Sources: U.S. FBI Uniform Crime Reports, England (including Wales) U.K. Office for National Statistics, World U.K. Office on Drugs and Crime, reported by U.N. Economic and Social Council's World Crime Trends and Emerging Issues and responses in 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.

**RATE OF BATTLE DEATHS IN ARMED CONFLICTS
1946-2013**



Source: Utopia Conflict Data Program Battle-Related Deaths Dataset 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 Meisander 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. East, "Unsettled Status: Local Analysis of Documentation of Killings in the Syrian Arab Republic," Human Rights Data Analysis Group, http://www.hrdag.org/Documents/Countries/SY/HRDAG_LocalReportAug2014.pdf. World population figures from U.S. Census Bureau, http://www.census.gov/population/international/data/worldpop/stable_population.php

**VICTIMIZATION OF CHILDREN IN THE US
1990-2012**



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 12-17.

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)

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution


Jeanne Calment

Culture Dish

Enlightenment


The View From Here

**PREVALENCE OF MASS KILLINGS
1945-2013**




Source: Early Warning Project, <https://ppparfwarmping.worldpress.com/>; Ulfelder, Jay and Valentino, Benjamin, "Assessing Risks of State-Sponsored Mass Killing" (Feb. 1, 2006). Available at SSRN: <https://ssrn.com/abstract=1033426> or <http://dx.doi.org/10.2339/arm-033426>

**HOMICIDE RATES IN THE US AND ENGLAND
1967-2013, AND THE WORLD, 2003-2012**




Sources: U.S. FBI Uniform Crime Reports, England (including Wales) U.K. Office for National Statistics, World U.K. 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.

**RATE OF BATTLE DEATHS IN ARMED CONFLICTS
1946-2013**



Source: Utopia Conflict Data Program Battle-Related Deaths Dataset 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. East, "Underestimated: Statistical Analysis of Documentation of Killings in the Syria Arab Republic," Human Rights Data Analysis Group, http://www.hrdag.org/Documents/Countries/SY/HRDAG_SyriaReportAug2014.pdf. World population figures from U.S. Census Bureau, http://www.census.gov/population/international/data/worldpop/stable_population.php

**VICTIMIZATION OF CHILDREN IN THE US
1990-2012**




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 12-17.

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

Outline

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN

AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here

- 1 Introduction
- 2 Deep Q-Learning Theory
- 3 Cart Pole Deep Q-Learning
- 4 Beyond DQN
- 5 **The AI Revolution**
 - Jeanne Calment
 - Culture Dish
 - Enlightenment
 - The View From Here**

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



waitbutwhy.com

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



waitbutwhy.com

Intro

DQN In
Theory

DQN In
Practice

Beyond DQN


AI Revolution

Jeanne Calment

Culture Dish

Enlightenment

The View From Here



waitbutwhy.com