



REINFORCEMENT LEARNING COACH

Gal Novik

May 24th 2018

REINFORCEMENT LEARNING BACKGROUND

MACHINE LEARNING

SUPERVISED





Image → Class

UNSUPERVISED



Data Point → Cluster

REINFORCEMENT



State → Action

REINFORCEMENT LEARNING IN THE NEWS


2013

Playing Atari with Deep Reinforcement Learning

Volodymyr Mnih Koray Kavukcuoglu David Silver Alex Graves Ioannis Antonoglou
Daan Wierstra Martin Riedmiller
DeepMind Technologies



2016



How Google's AlphaGo Beat a Go World Champion

Inside a man-versus-machine showdown

CHRISTOPHER MOYER | MAR 28, 2016 | TECHNOLOGY

REINFORCEMENT LEARNING IN THE NEWS

2017

MIT
Technology
Review

Log in / Create an account Search Q

Past Lists+ Topics+ The Download Magazine Events More+

Subscribe


10 Breakthrough Technologies Years +

Reinforcement Learning

By experimenting, computers are figuring out how to do things that no programmer could teach them.

Availability: 1 to 2 years

by Will Knight



Now

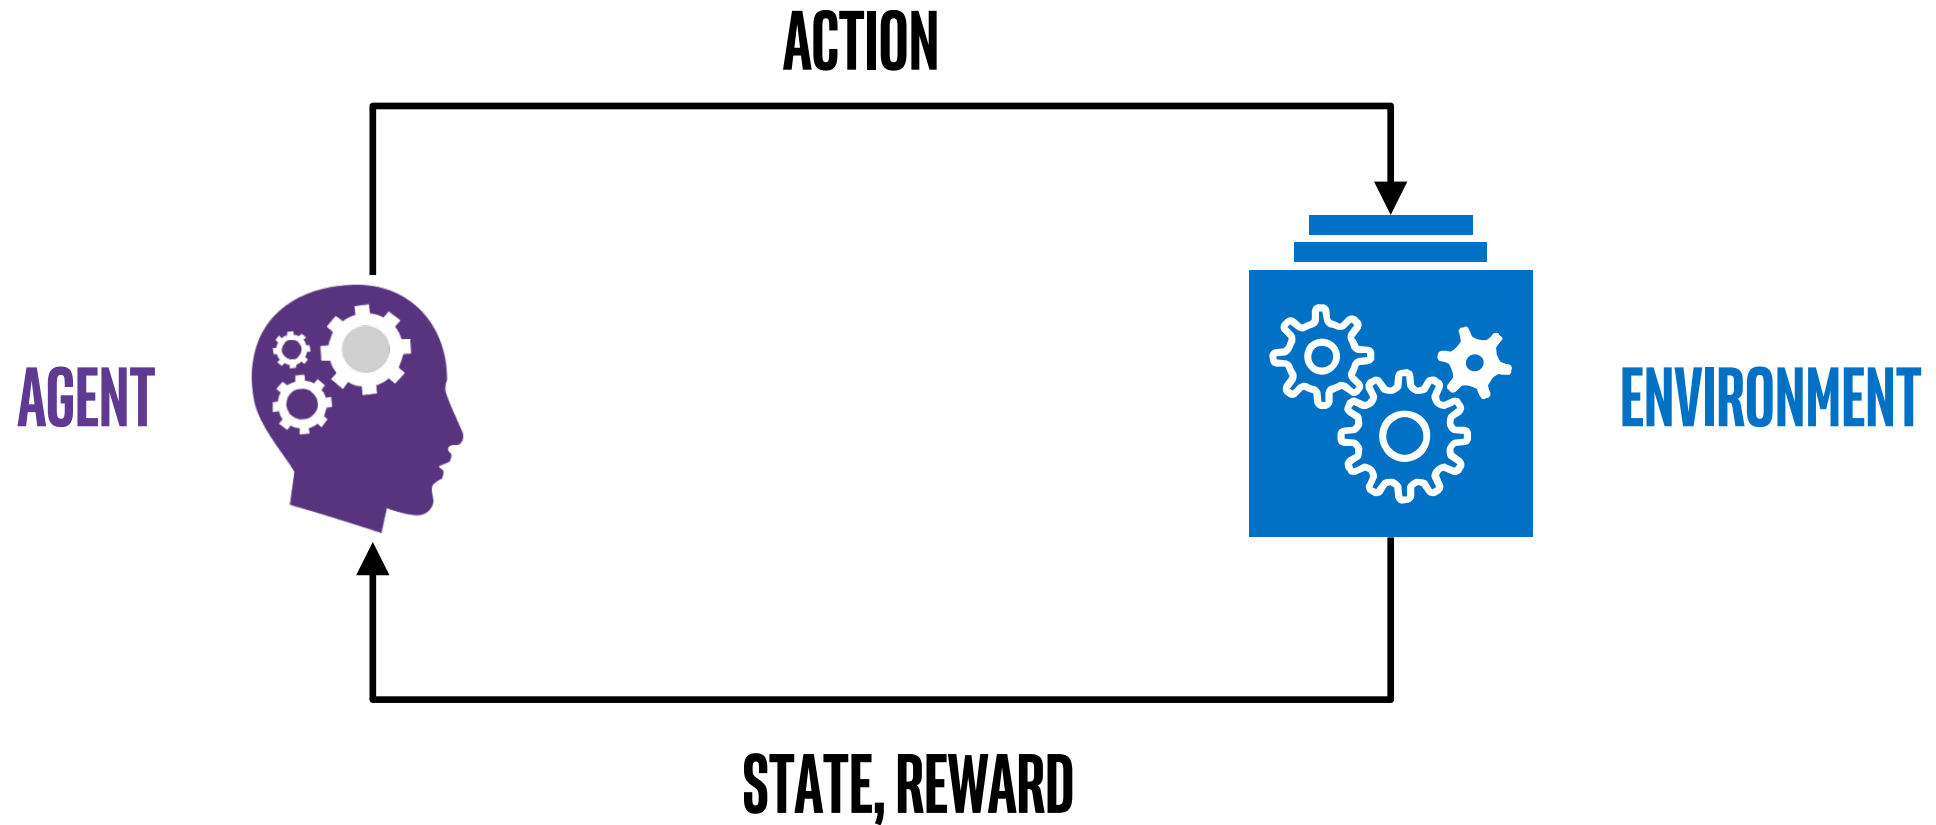
Deep Reinforcement Learning for Sepsis Treatment

A Deep Reinforcement Learning Chatbot


Composable Deep Reinforcement Learning for Robotic Manipulation

Tuomas Haarnoja¹ Vitchyr Pong¹ Aurick Zhou¹ Murtaza Dalal¹ Pieter Abbeel^{1,2} Sergey Levine¹ ¹Berkeley Artificial Intelligence Research, UC Berkeley, ²Open AI
{haarnoja, vitchyr, azhou42, mdalal, pabbeel, svlevine}@berkeley.edu

THE REINFORCEMENT LEARNING SETTING



THE REINFORCEMENT LEARNING SETTING



WHAT PROBLEMS IS IT GOOD FOR?


- + Very **generic**, can fit many problems
- + Clear **goal**, but optimal actions are unknown
- + Learning through **interaction** with a real/simulated environment
- Low sample efficiency
- Training instability

CHALLENGES

- ? Many RL algorithms. Which one should I use?
- ? How do I train an agent on my environment?
- ? How do I verify correctness? How do I debug?


COACH

Open source Python RL framework for AI agents development and training - <https://github.com/NervanaSystems/coach>



COACH

Open source Python RL framework for AI agents development and training - <https://github.com/NervanaSystems/coach>



COACH

Open source Python RL framework for AI agents development and training - <https://github.com/NervanaSystems/coach>



Potential users are students and researchers, or data scientists who would like to add a strong tool to their toolset

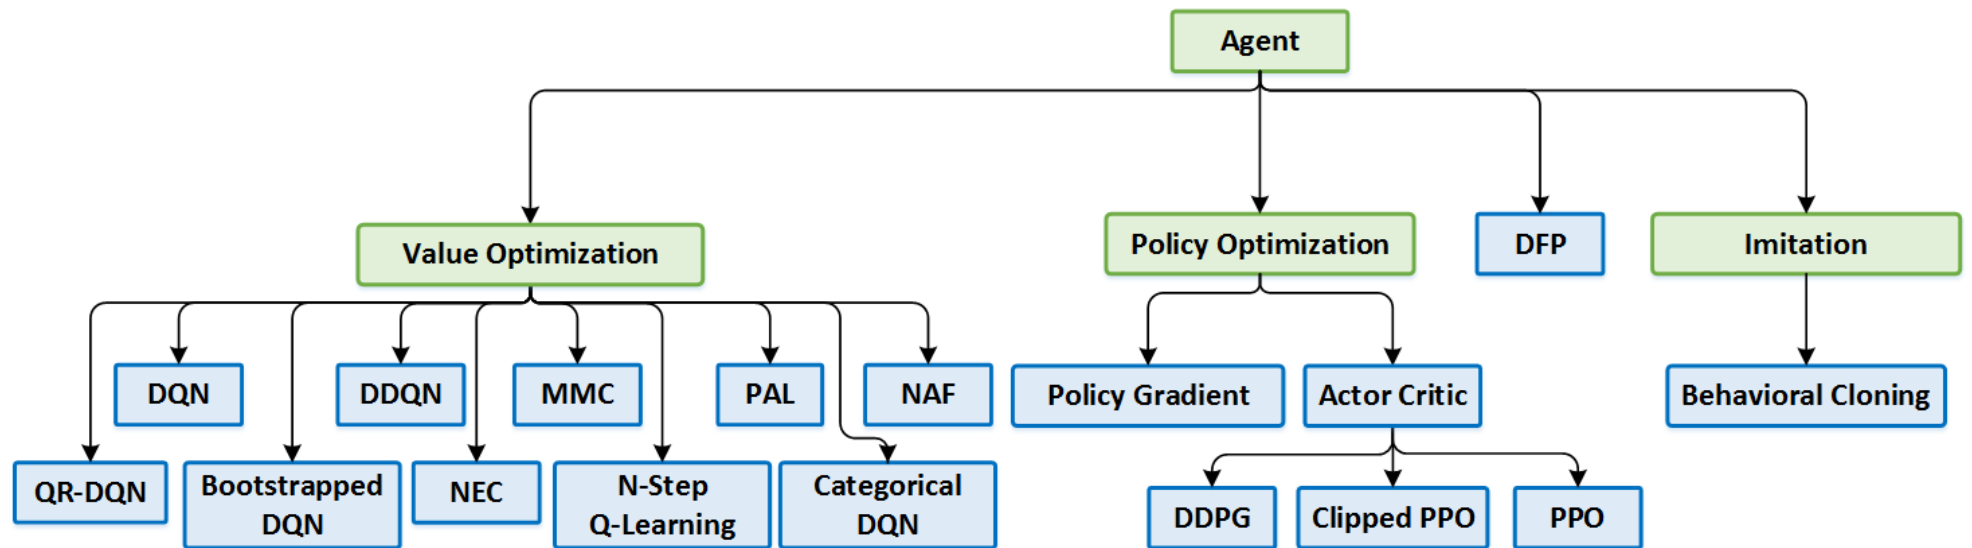


Uses the CPU optimized Tensorflow version for training agents efficiently on CPU




AGENTS

Implements many state-of-the-art algorithms

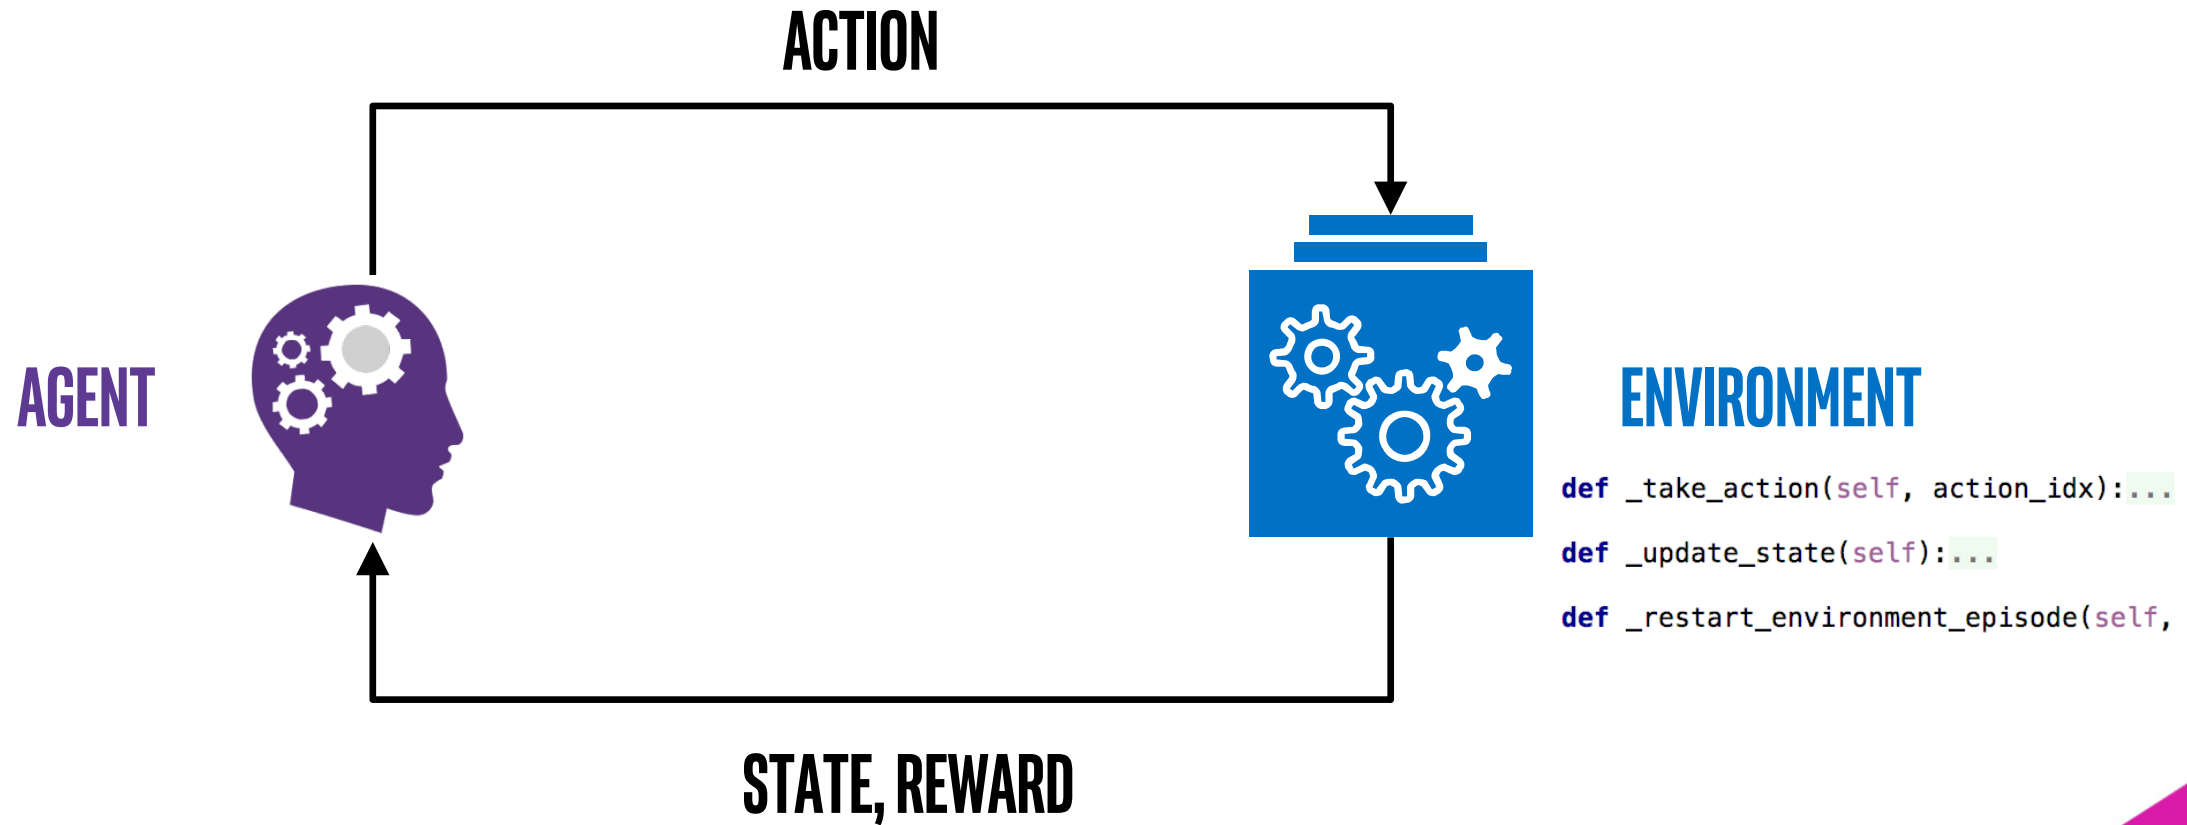


ENVIRONMENTS

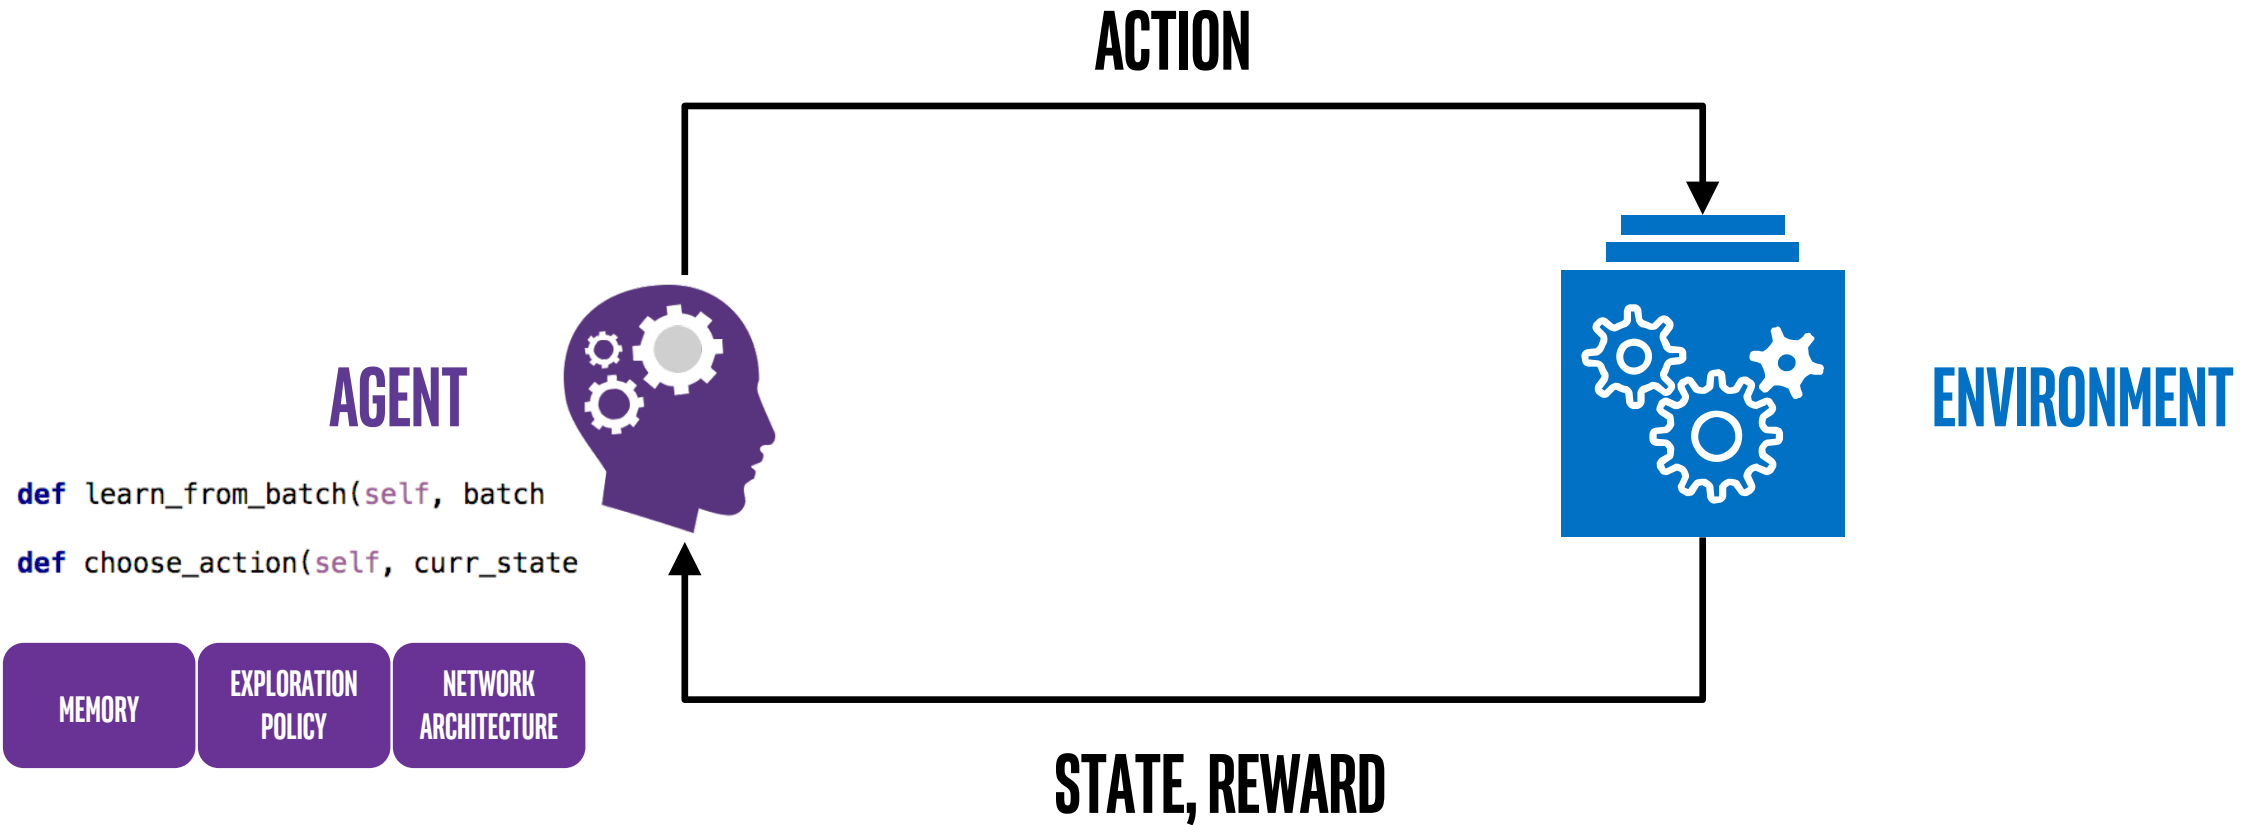
Integrated with multiple simulation environments



ADDING AN ENVIRONMENT




ADDING AN AGENT



DASHBOARD VISUALIZATION TOOL

Coach Dashboard



DEMO

FUTURE PLANS

- 🕒 Composite agents
- 🕒 Hierarchical RL
- 🕒 StarCraft II integration
- 🕒 New components and agents

WHAT NEXT?

- Ramp up
 - Book: Reinforcement Learning: An Introduction by Sutton & Barto
 - Videos: RL Course by David Silver on YouTube
 - Code: RL Coach on <https://github.com/NervanaSystems/coach>
- Build your own agent / simulator
- Contribute



WIN ►►►►
AWS DEEPLENS

AIDEVCONAPP.INTEL.COM



**Completing a session
evaluation in the
mobile app by
10:00 a.m. tomorrow
automatically enters
you in a drawing to win.**

**Copies of the complete
sweepstakes rules
are available at the
Concierge Desks.**

