
AWS
re:Invent

AIM401

Replicate AWS DeepRacer architecture to conquer the track with SageMaker

Sahika Genc, Ph.D.

Principal Applied Scientist

AWS

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Recap

- ✓ Taken the 200- and 300-level courses on AWS DeepRacer
- ✓ Configured AWS DeepRacer vehicle in the Garage
- ✓ Developed AWS DeepRacer reward function
- ✓ Trained and evaluated AWS DeepRacer models

Agenda

Using AWS DeepRacer with Amazon SageMaker reinforcement learning (RL) notebooks

Modifying neural network architecture

Real-world examples of RL

New for 2021 on AWS DeepRacer



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

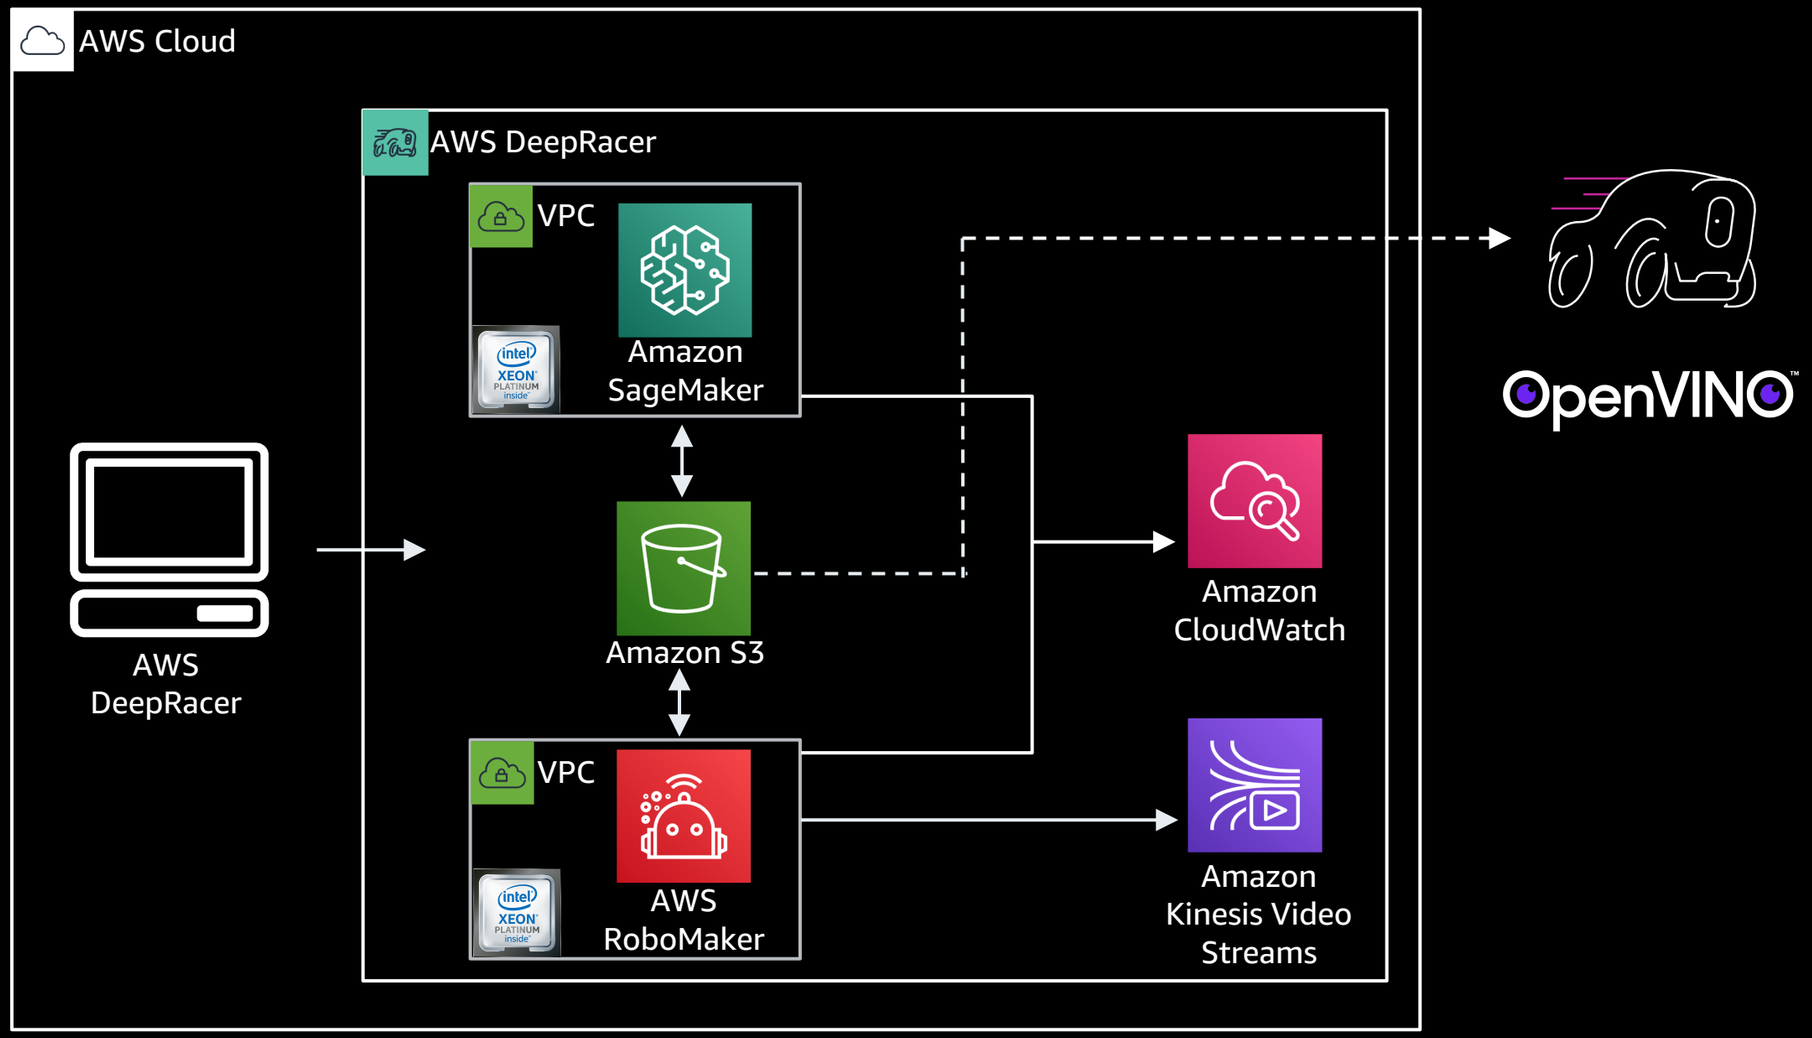
AWS DeepRacer with Amazon SageMaker RL notebooks



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS DeepRacer architecture diagram



Today's focus is the Amazon SageMaker notebook

What does Amazon SageMaker enable beyond the AWS DeepRacer console?

- Design and edit new worlds for navigation beyond the league
- Add new sensors beyond image-based cameras and LIDAR
- Multiple simulation applications to collect more experience in shorter time
- Evaluate models in previously unseen race tracks
- Modify neural network architectures beyond shallow neural networks

Learn pros / cons of using a deep neural network agent

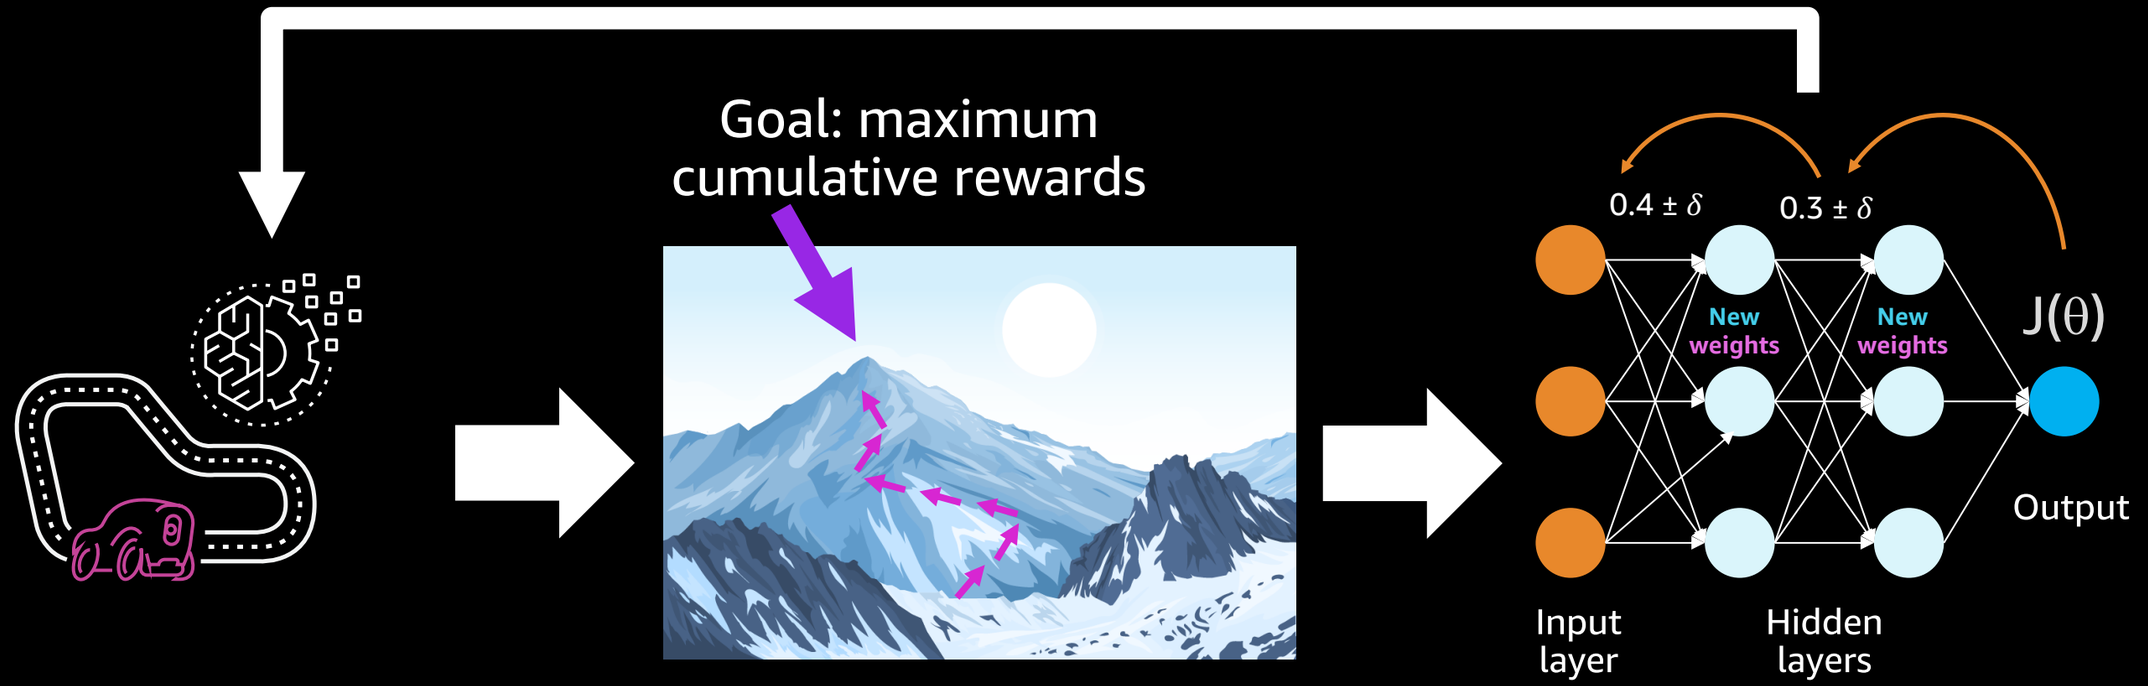
Modifying neural network architecture



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Rewards are maximized by a neural network



Gather episode
using current policy

Calculate the gradient of
estimated cumulative reward

Update the policy for
gradient ascent

**Neural network approximates agent behavior
by maximizing discounted cumulative rewards**

Deeper networks approximate discontinuous cumulative reward

Pros

Better approximation of the reward function leads to higher rewards and faster convergence in *number of episodes*

Approximates reward function with discontinuities with lower error (on-track / off-track) vs. gradually decaying penalties

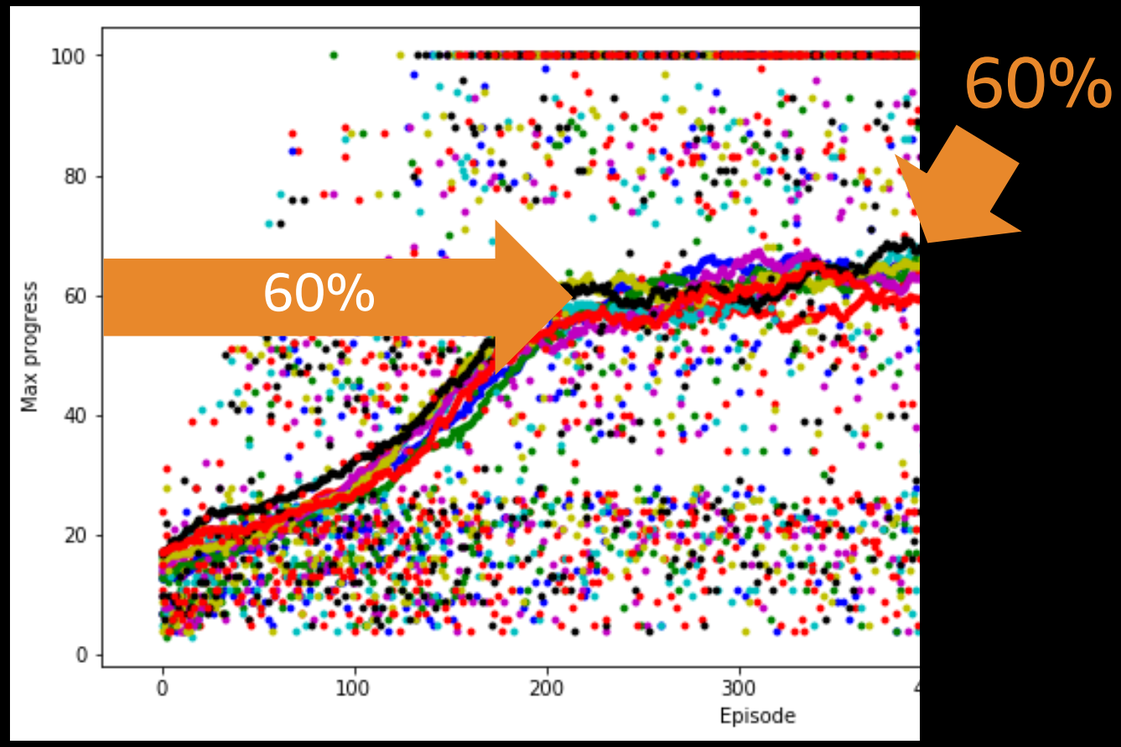
Cons

Potentially slower with CPU instances, hence, slower convergence *over time* → Select a GPU instance in Amazon SageMaker notebook

May not generalize to real world as well as the shallower neural networks

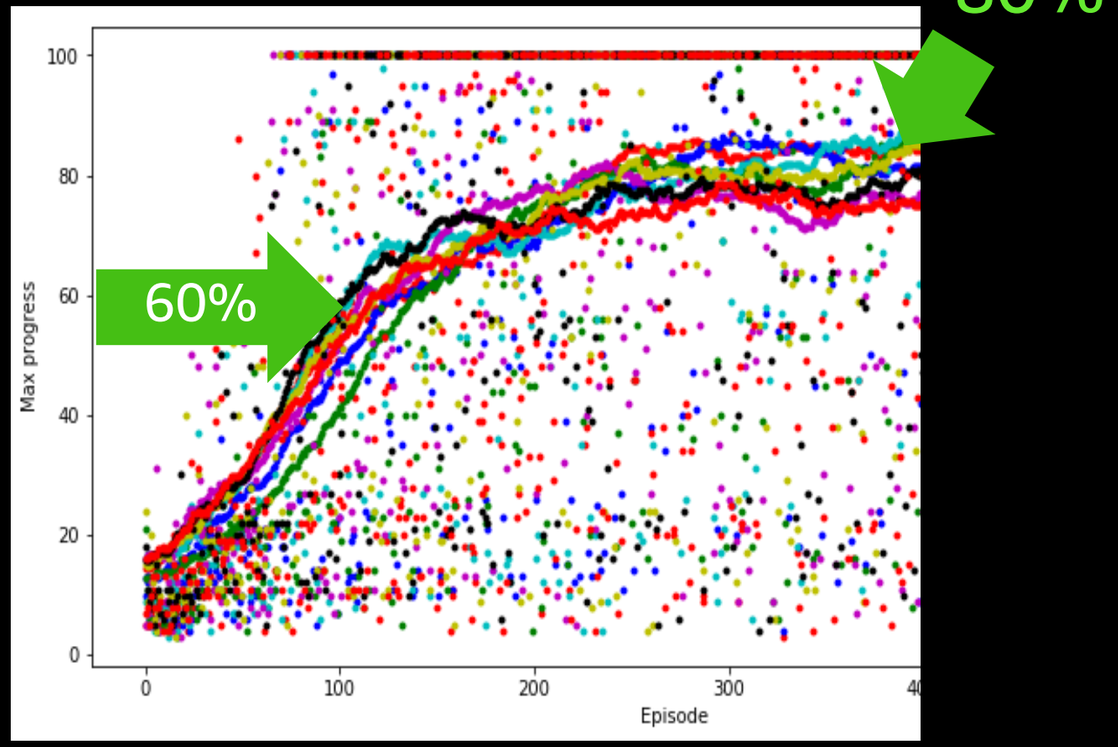
Neural network approximates discounted cumulative reward

Shallow



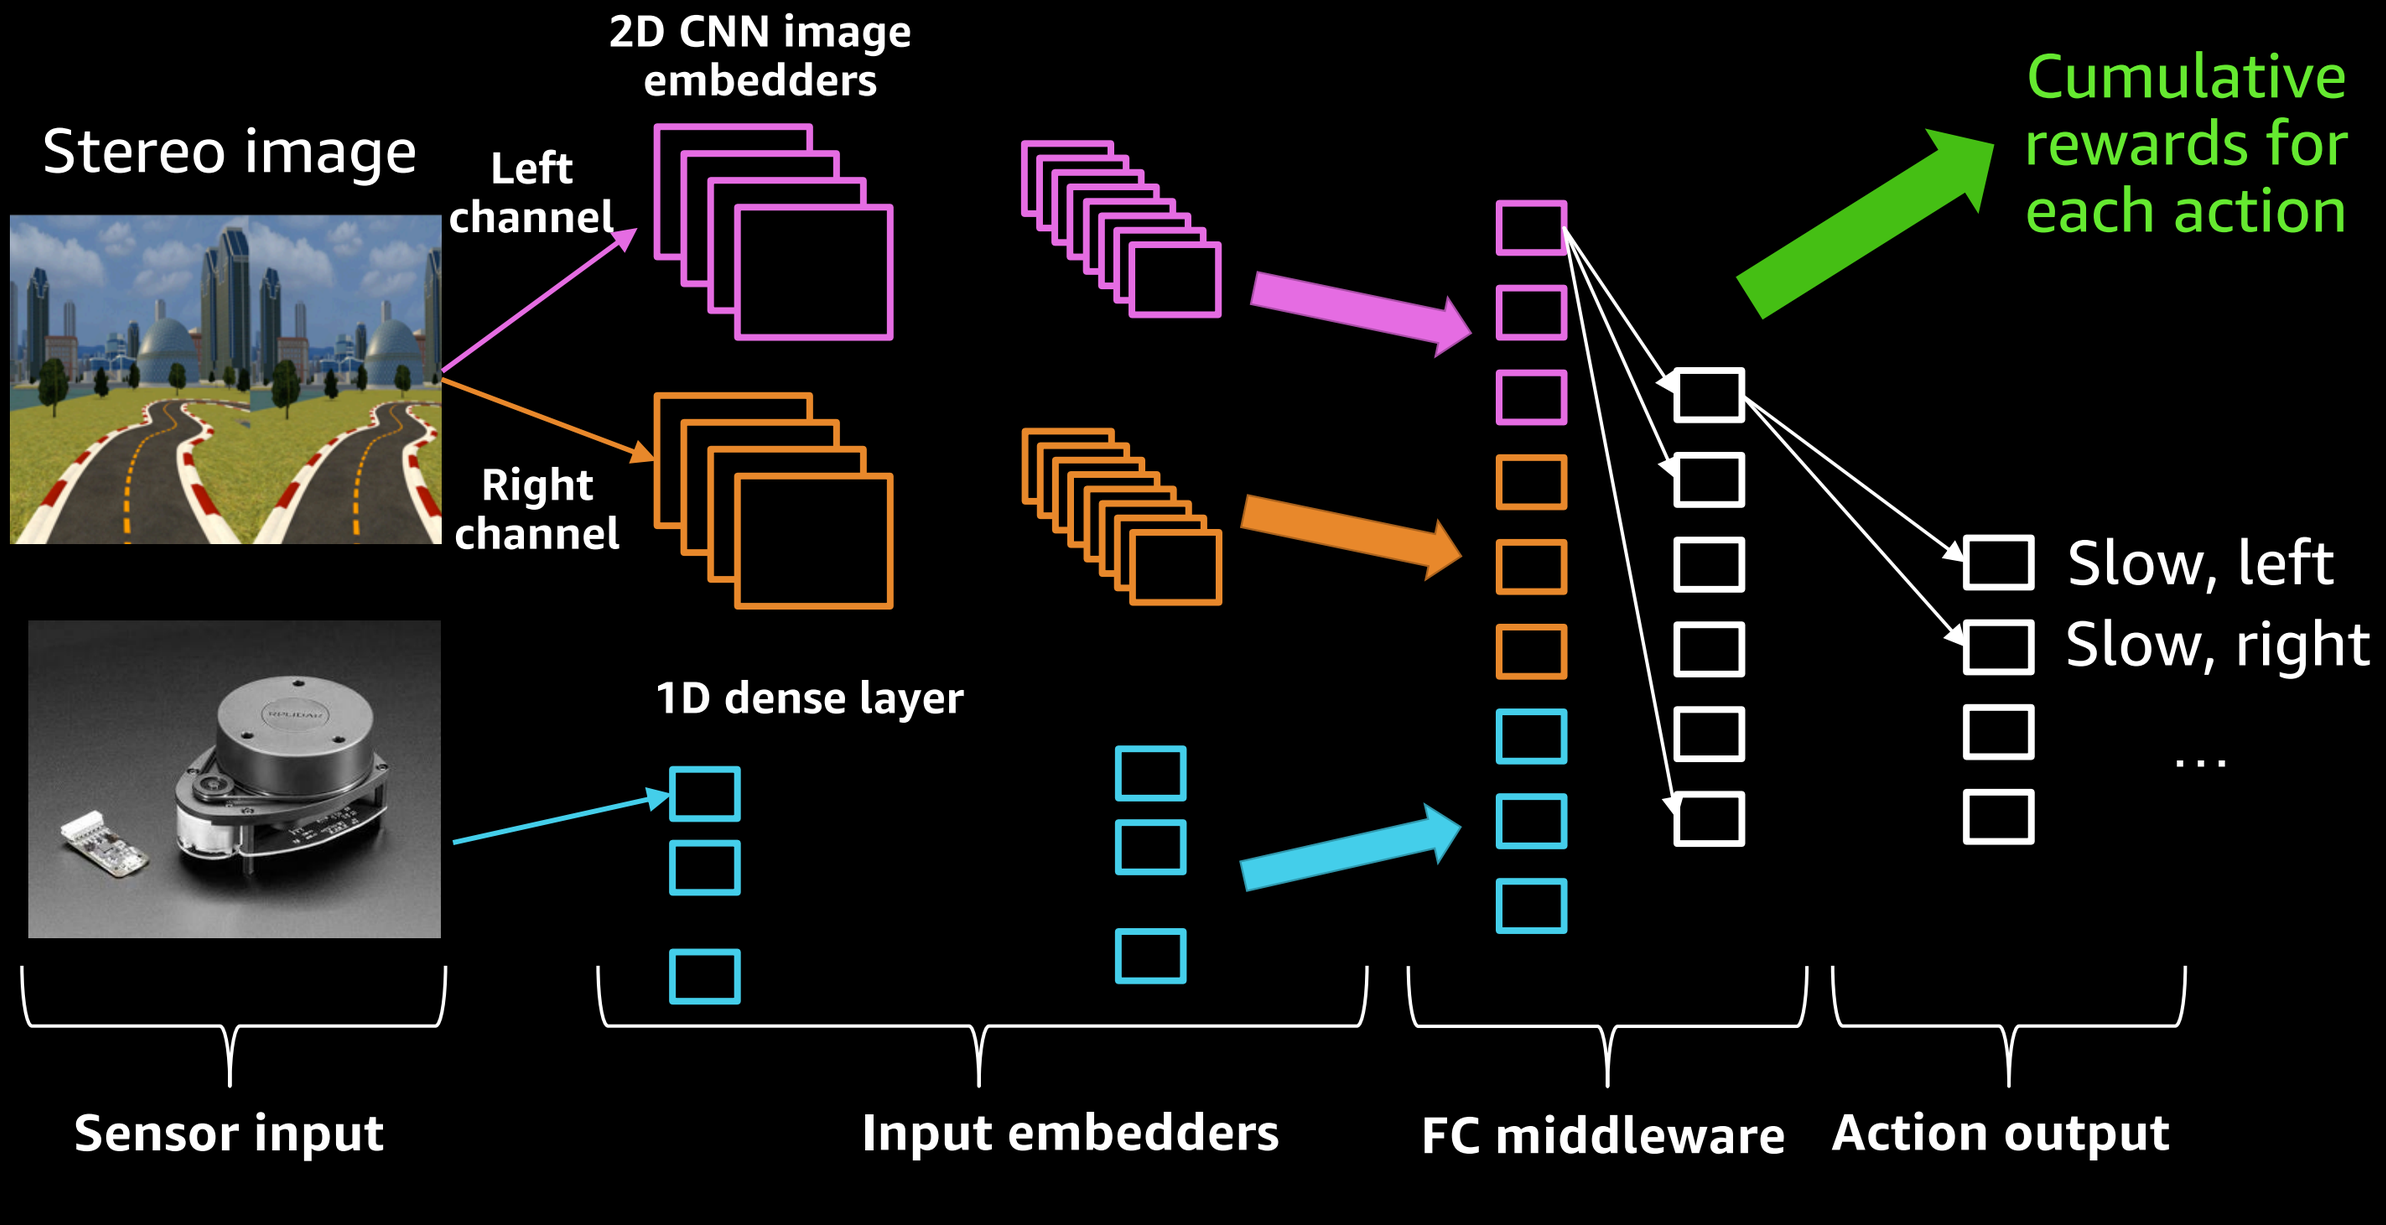
Still 60% progress after 400 episodes

Deep



80% progress after 400 episodes

How does the neural network architecture look?



How to update the neural network with Amazon SageMaker

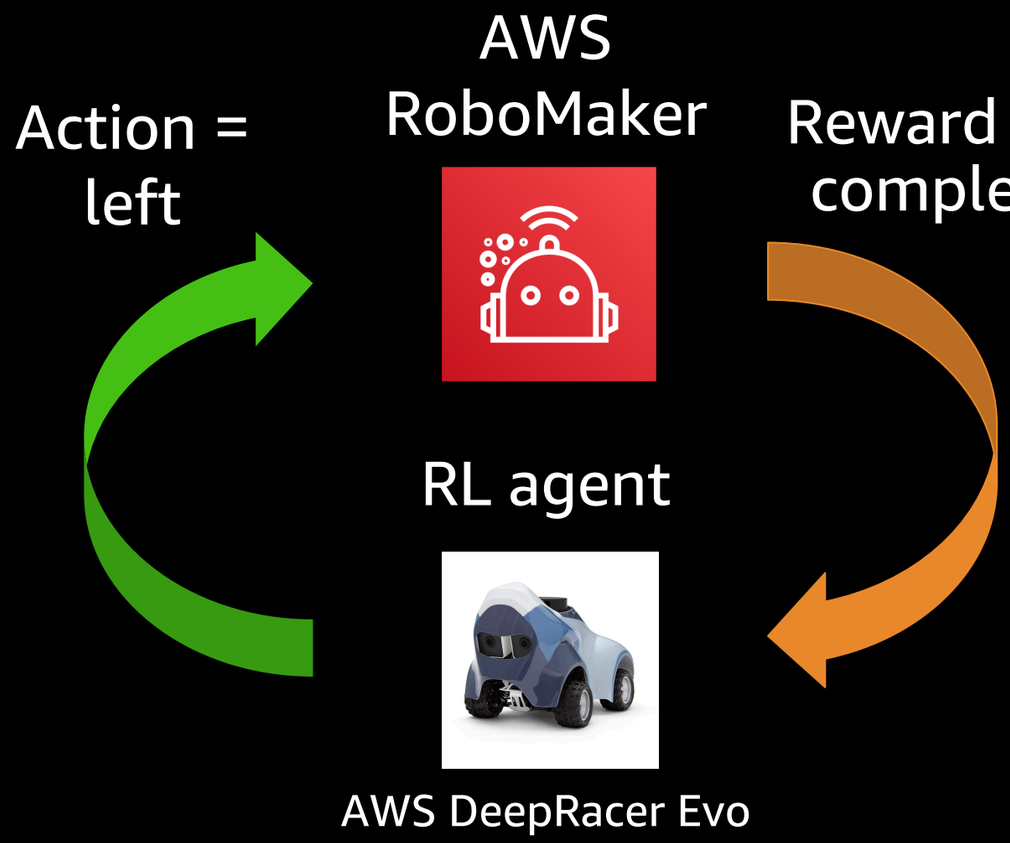
- To get started, go to http://www.deepracerworkshops.com/400-level-workshop/modifyneuralnetwork/
- Follow step-by-step instructions on the AWS DeepRacer 400-level workshop materials to quickly change your neural network with existing templates for shallow and deep neural networks

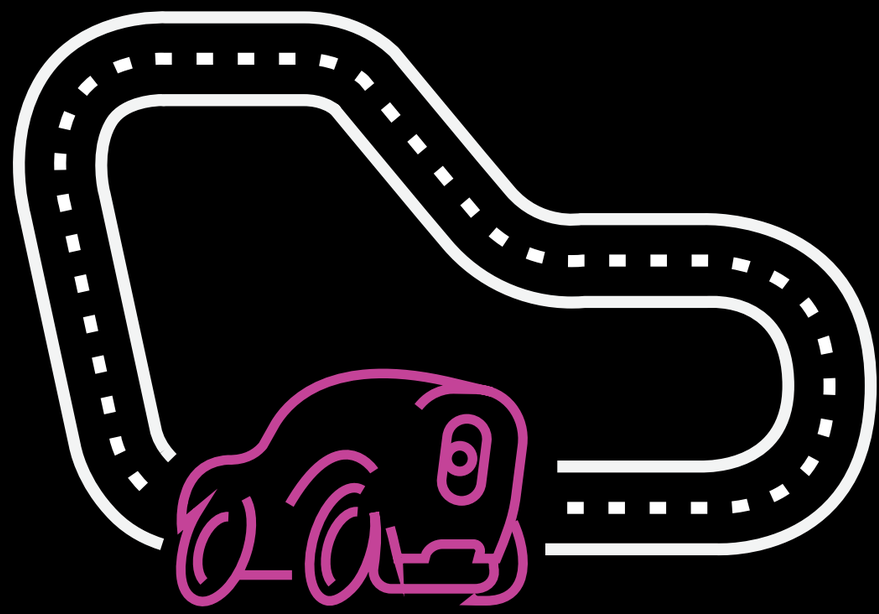
Real-world use cases for applying RL



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

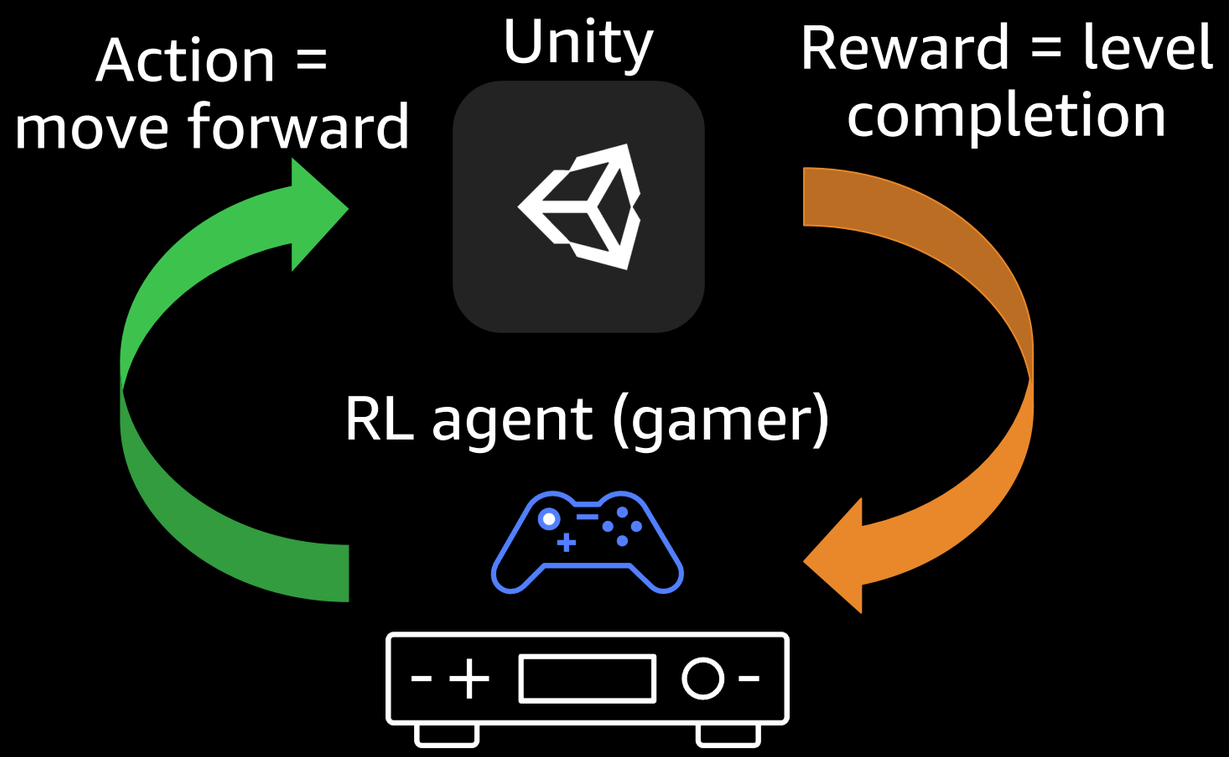
Recap: AWS DeepRacer learns to drive fast

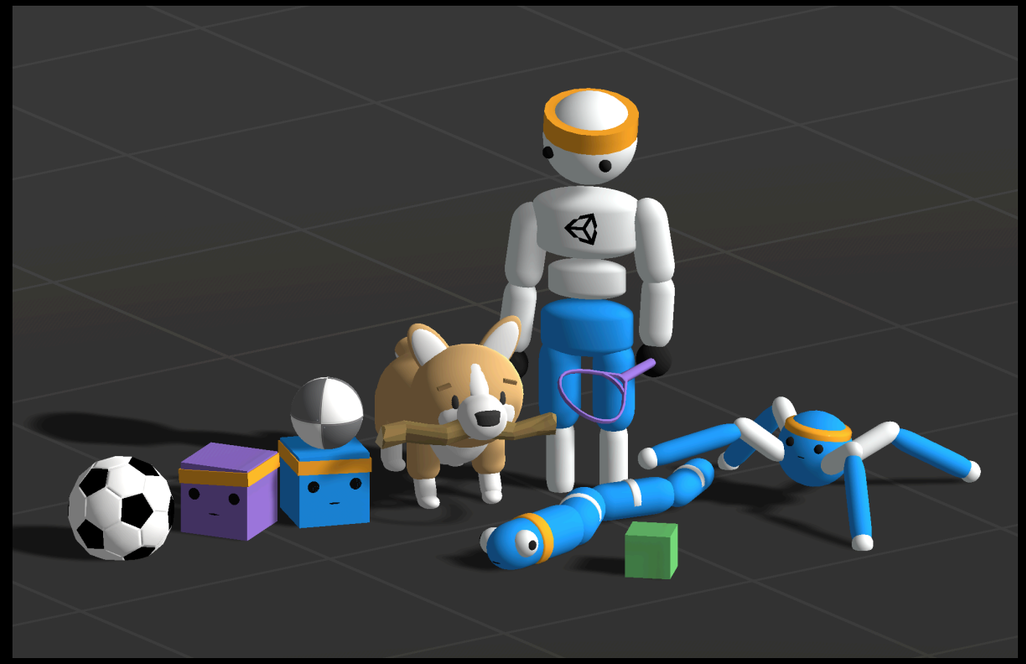




Agent takes actions and gets rewards for completing a lap without crashing

CAPCOM builds fun games faster with managed services on AWS

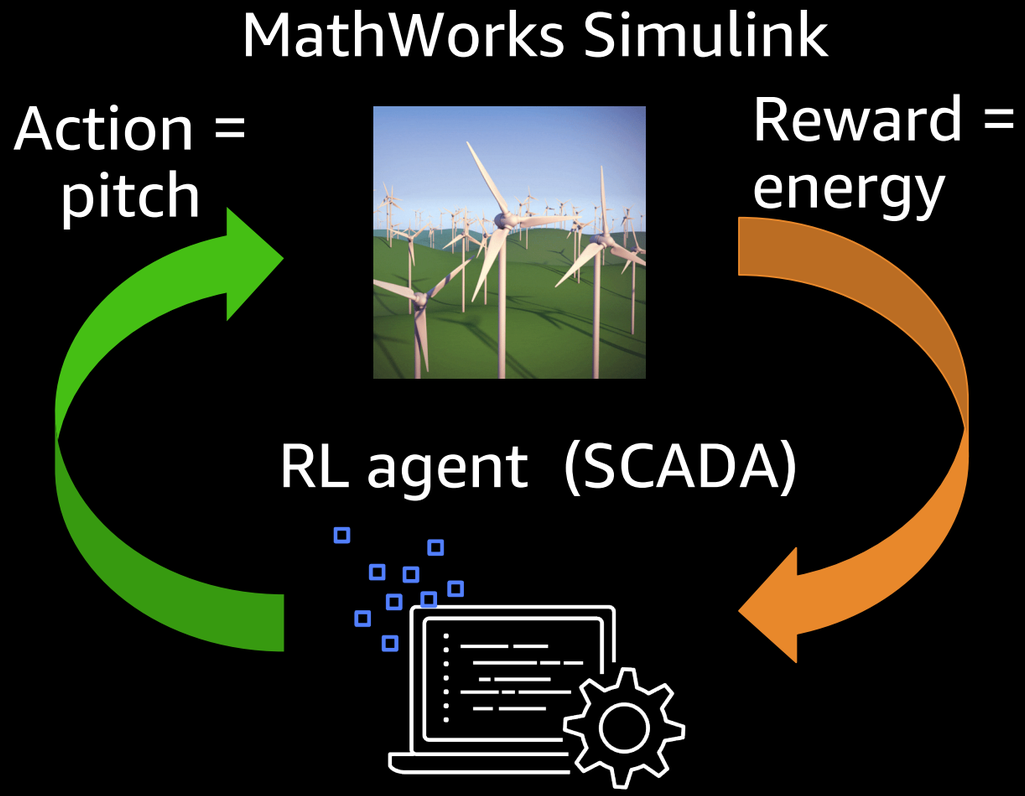


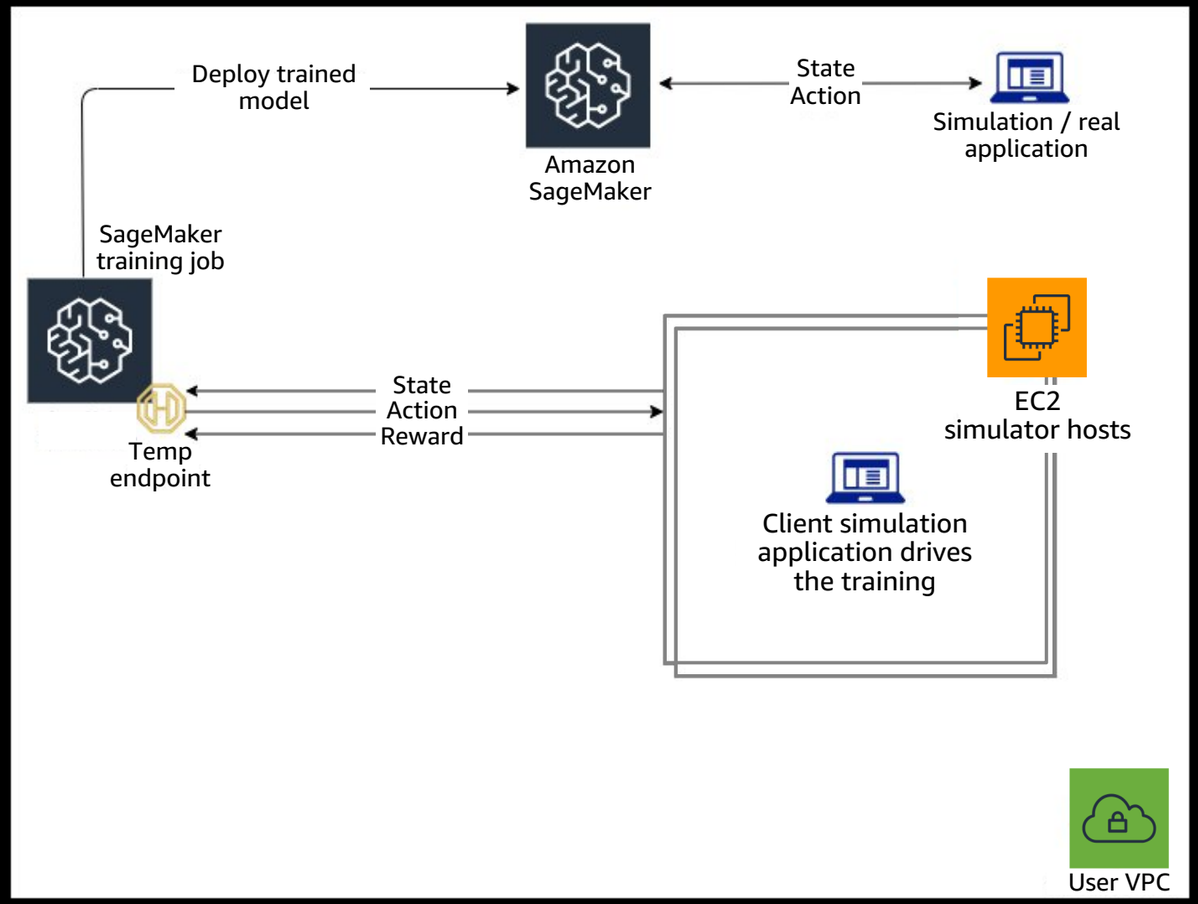


https://github.com/aws/amazon-sagemaker-examples/tree/master/reinforcement_learning/rl_unity_ray

Agent plays various levels of computer games and evaluates new games prior to production

MathWorks simulation environments enable industrial workflows





Agent learns to control the wind turbines to maximize the energy generation

Announcing new 2021 League features



Open and pro divisions

Digital rewards in console

Special sponsored events

Thank you!



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.