INNOVATE

AI/ML EDITION

24 February 2022

Start your engines with AWS DeepRacer

Calvin Ngo

Developer Specialist Solutions Architect
AWS

Agenda

1. Machine Learning on AWS
2. Introducing AWS DeepRacer
3. Introduction to Reinforcement Learning
4. AWS DeepRacer Console
5. Demo
6. Additional resources

Machine Learning on AWS

The AWS ML Stack

BROADEST AND MOST COMPLETE SET OF MACHINE LEARNING CAPABILITIES

AI SERVICES

SPECIALIZED

Business processes
- Amazon Personalize
- Amazon Forecast
- Amazon Fraud Detector
- Amazon Lookout for Metrics

Search
- Amazon Kendra

Code + DevOps
- Amazon CodeGuru
- Amazon DevOps Guru

Industrial
- Amazon Monitron
- Amazon Lookout for Equipment
- Amazon Lookout for Vision

Healthcare
- Amazon HealthLake
- Amazon Comprehend Medical
- Amazon Transcribe Medical

CORE

Text & Documents
- Amazon Translate
- Amazon Comprehend
- Amazon Textract

Chatbots
- Amazon Lex

Speech
- Amazon Polly
- Amazon Transcribe
- Amazon Transcribe Call Analytics

Vision
- Amazon Rekognition
- AWS Panorama

ML SERVICES

Label data

SAGEMAKER CANVAS
No-code ML for business analysts

SAGEMAKER STUDIO LAB
Learn ML

AMAZON SAGEMAKER STUDIO IDE
Prepare data | Store features | Detect bias | Build with notebooks | Train models | Tune parameters | Deploy in production | Explain predictions | Manage & monitor

← **CI/CD** →

Manage edge devices

ML FRAMEWORKS & INFRASTRUCTURE

PyTorch, Apache MXNet, TensorFlow | Amazon Elastic Compute Cloud (Amazon EC2) | CPUs | GPUs | AWS Inferentia | AWS Trainium | Habana Gaudi | FPGA | Elastic inference

aws

Building your team's skills

HANDS-ON LEARNING

AWS DeepLens
Deep learning

AWS DeepRacer
Reinforcement learning

AWS DeepComposer
Generative AI

Introducing AWS DeepRacer

How can we put reinforcement learning in the hands of all developers?

Literally

aws

Under the hood

- 1:18 4WD scale car
- Intel Atom processor
- Intel distribution of OpenVINO toolkit
- Stereo Camera (4MP)
- 360-degree 12-meter scanning radius LIDAR sensor
- System memory: 4 GB RAM
- 802.11ac Wi-Fi
- Ubuntu 16.04.3 LTS
- ROS kinetic

OpenVINO™

Get hands-on experience with reinforcement learning

AWS DeepRacer
Evo

Get hands-on experience with reinforcement learning

AWS DeepRacer Evo

3D-racing simulator

Get hands-on experience with reinforcement learning

AWS DeepRacer Evo

3D-racing simulator

AWS DeepRacer League

Get hands-on experience with reinforcement learning

AWS DeepRacer Evo

3D-racing simulator

AWS DeepRacer League

Community races

Introduction to Reinforcement Learning

Reinforcement learning in context of AI

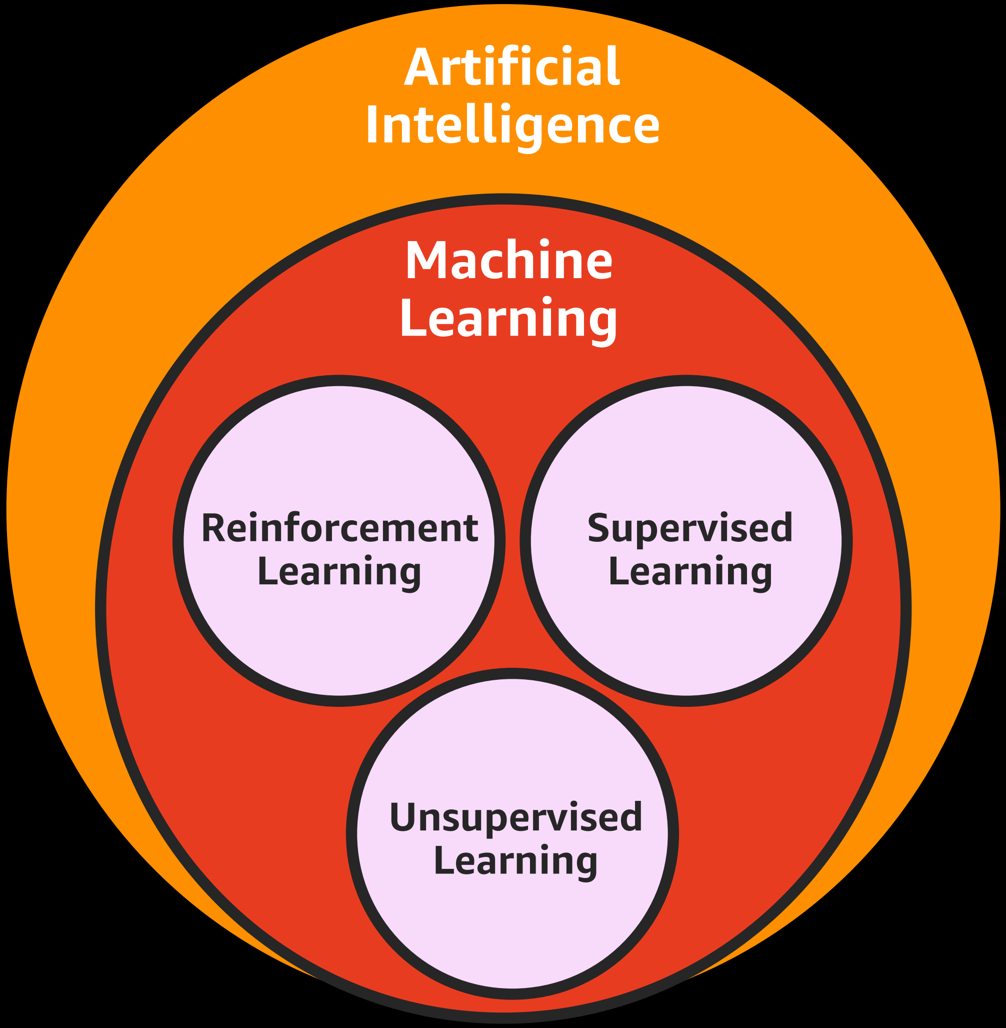

Real world reinforcement learning

Reward positive behavior

Don’t reward negative behavior

The result!

Reinforcement learning use cases

AUTONOMOUS CARS

FLEET LOGISTICS

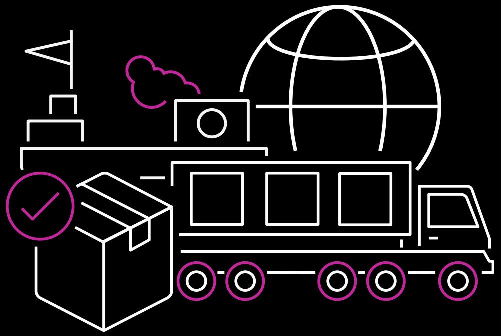

FINANCIAL TRADING

DATA CENTER COOLING

Reinforcement learning terminology

AGENT

ENVIRONMENT

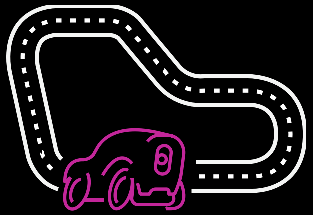

STATE

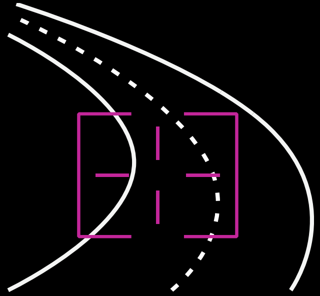

ACTION

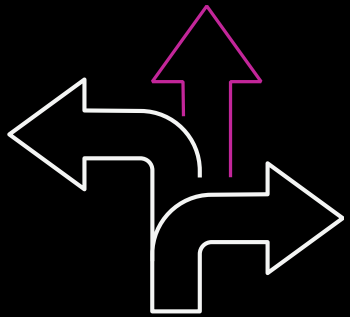

REWARD

EPISODE

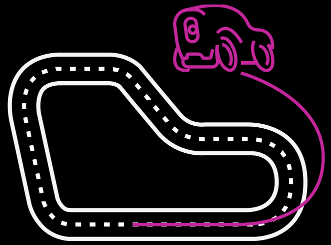

What is reinforcement learning?

MODEL

AGENT

ENVIRONMENT

ACTION

GOAL

Reward function in a grid race

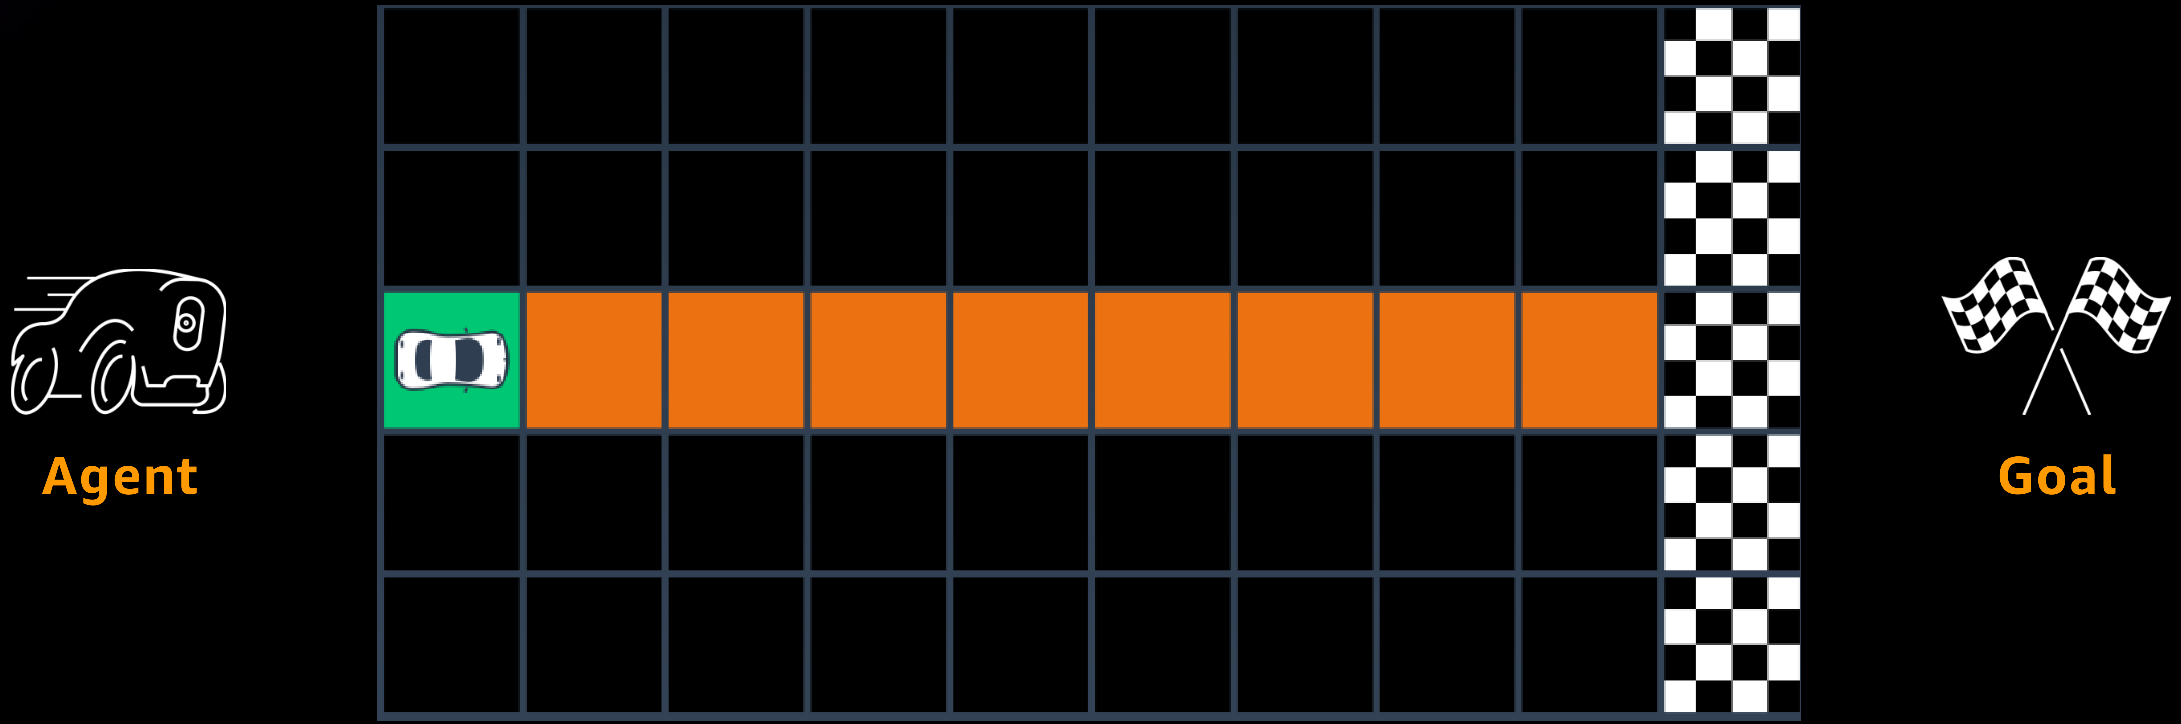

Incentivize center-line driving

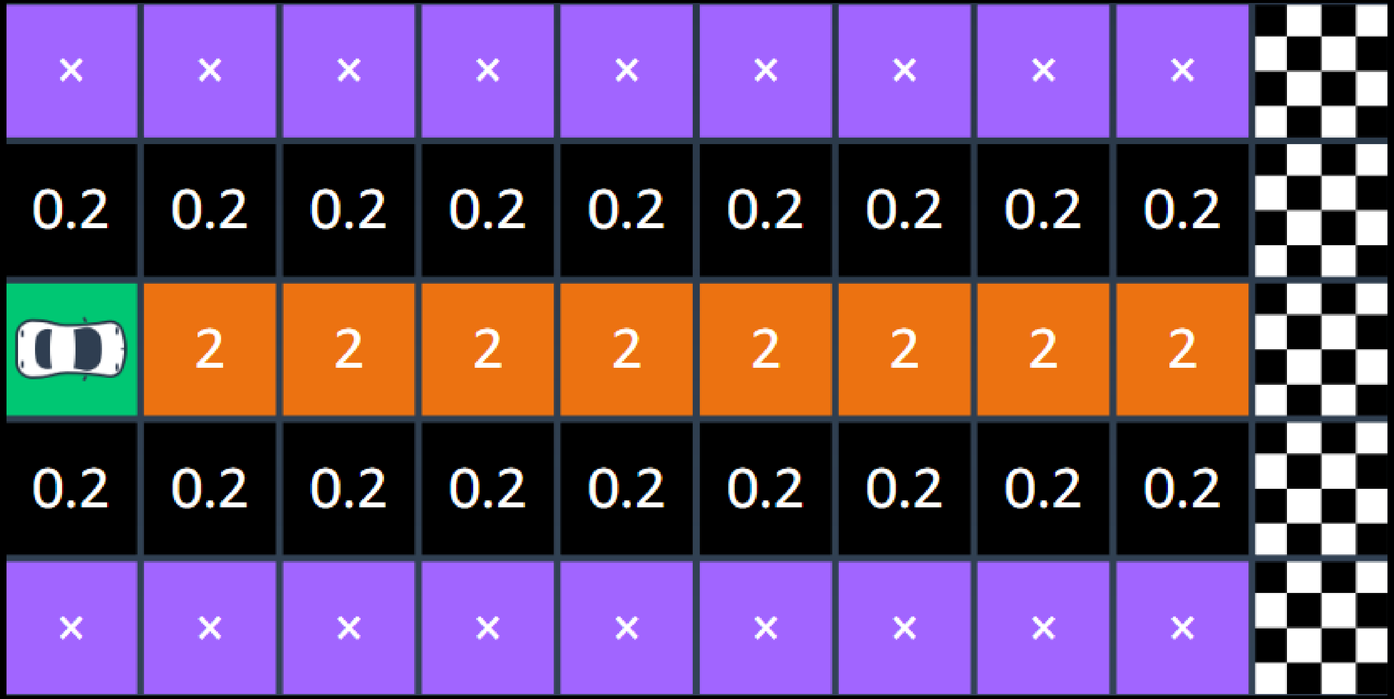

Iterate, iterate and converge

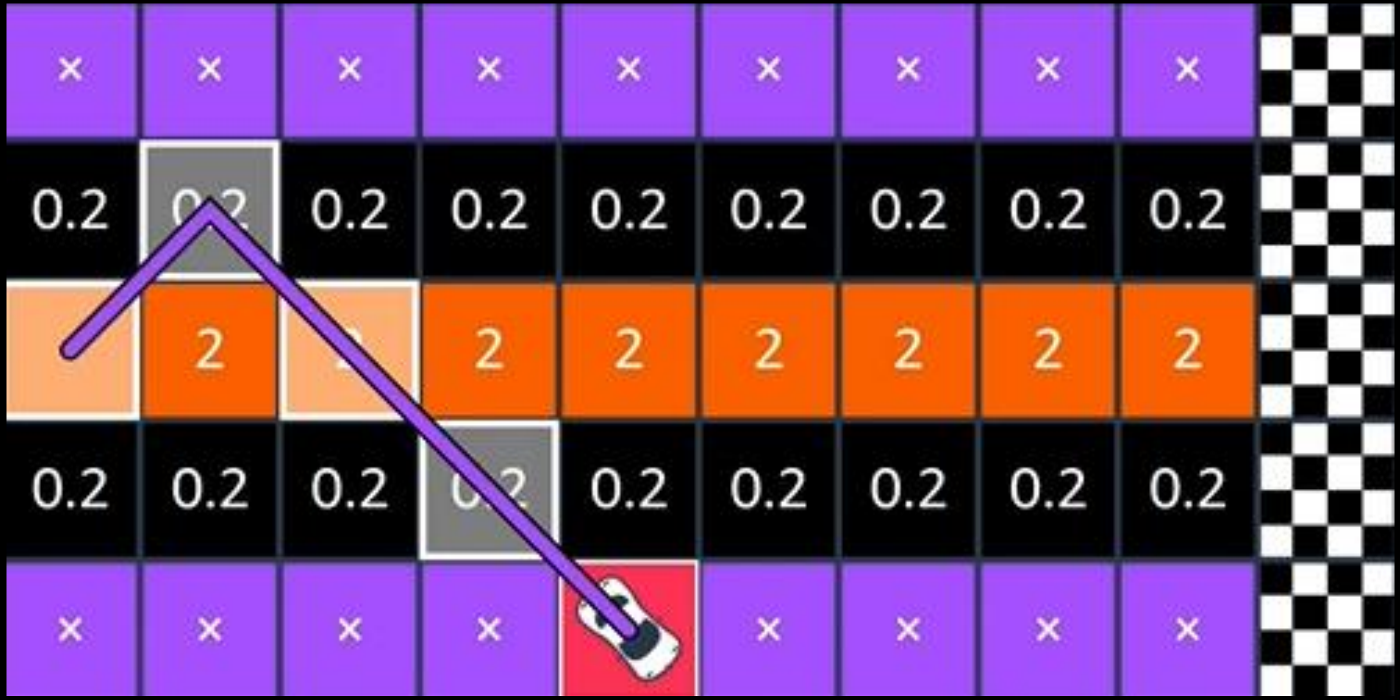

Exploration vs. Exploitation

EXPLORATION

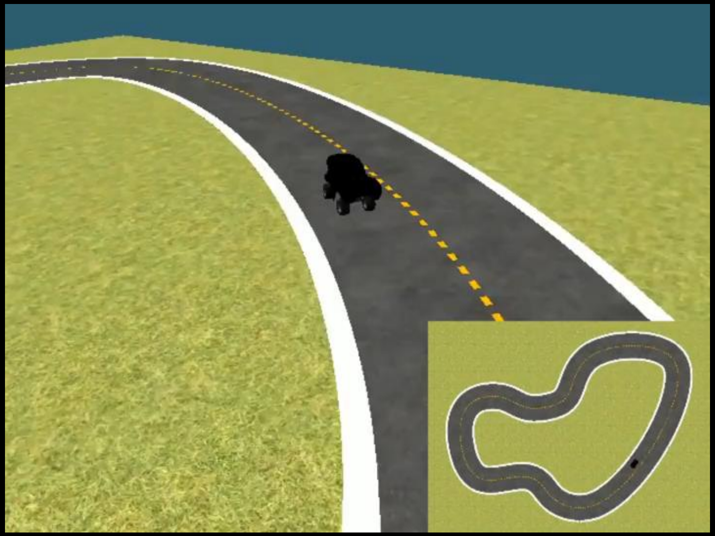

EXPLOITATION

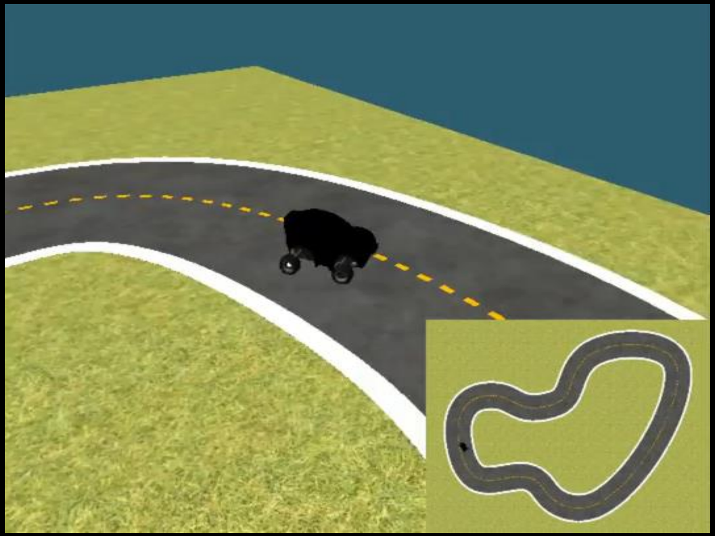

How does learning happen?

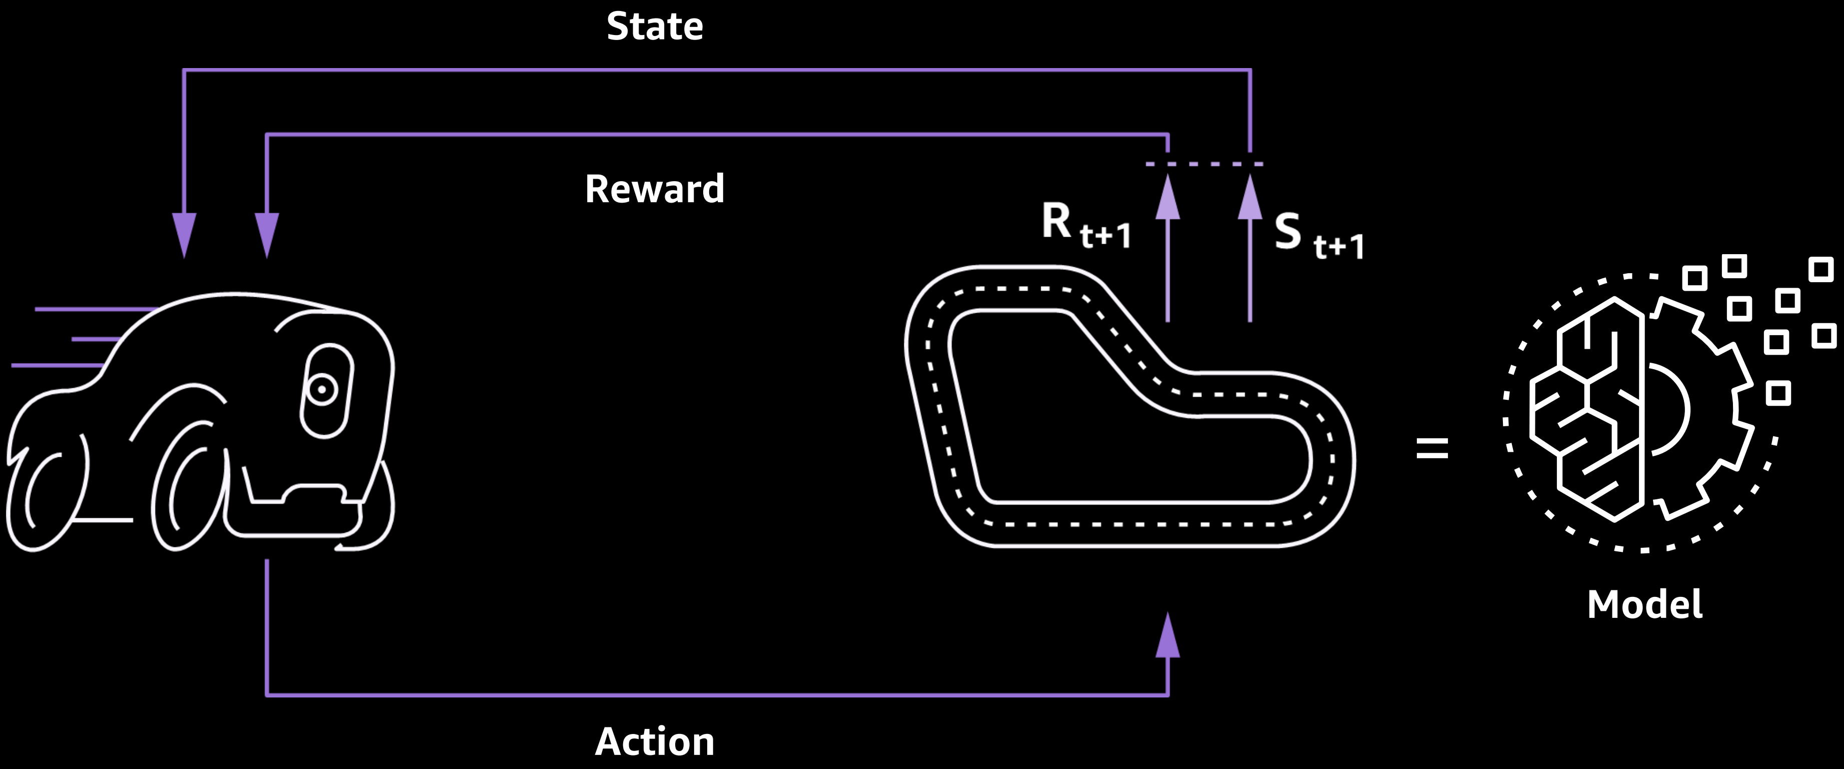

How does learning happen?

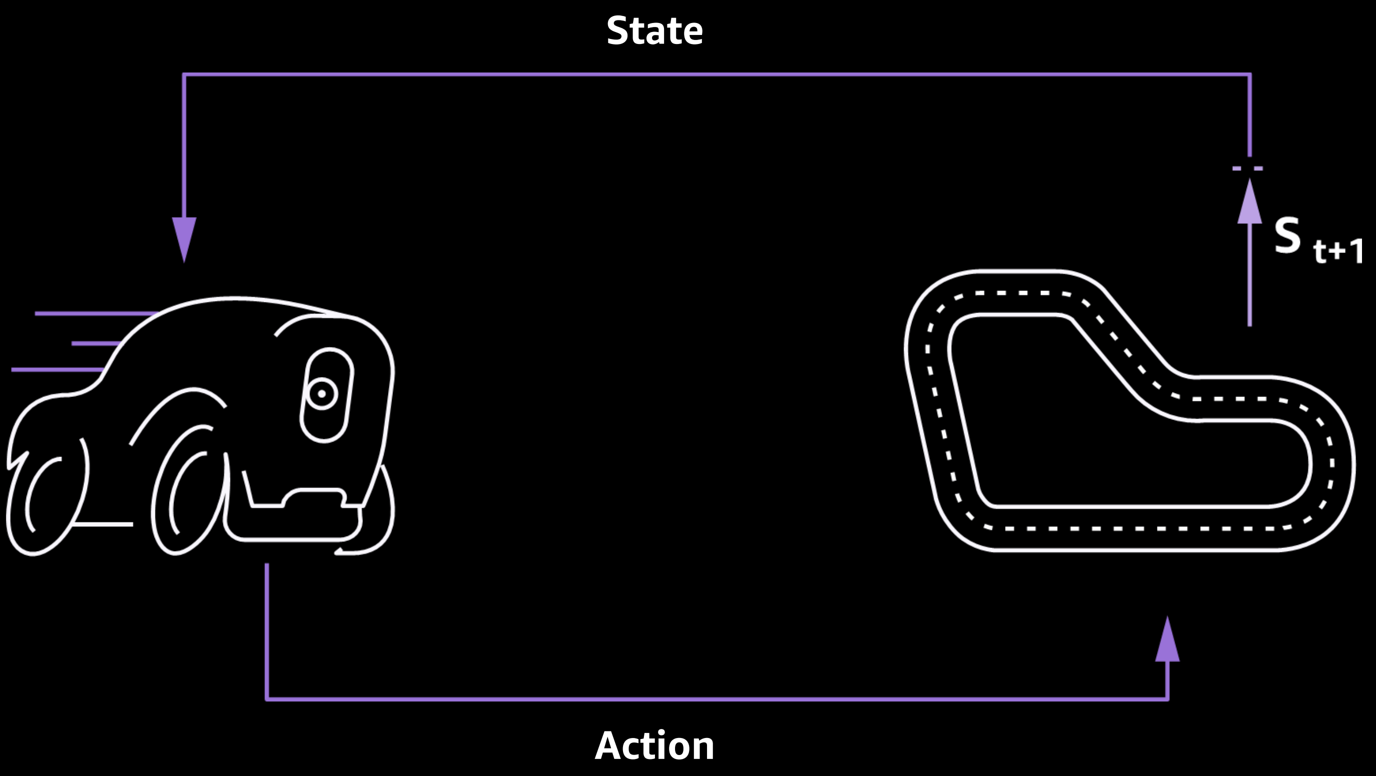

How does learning happen?

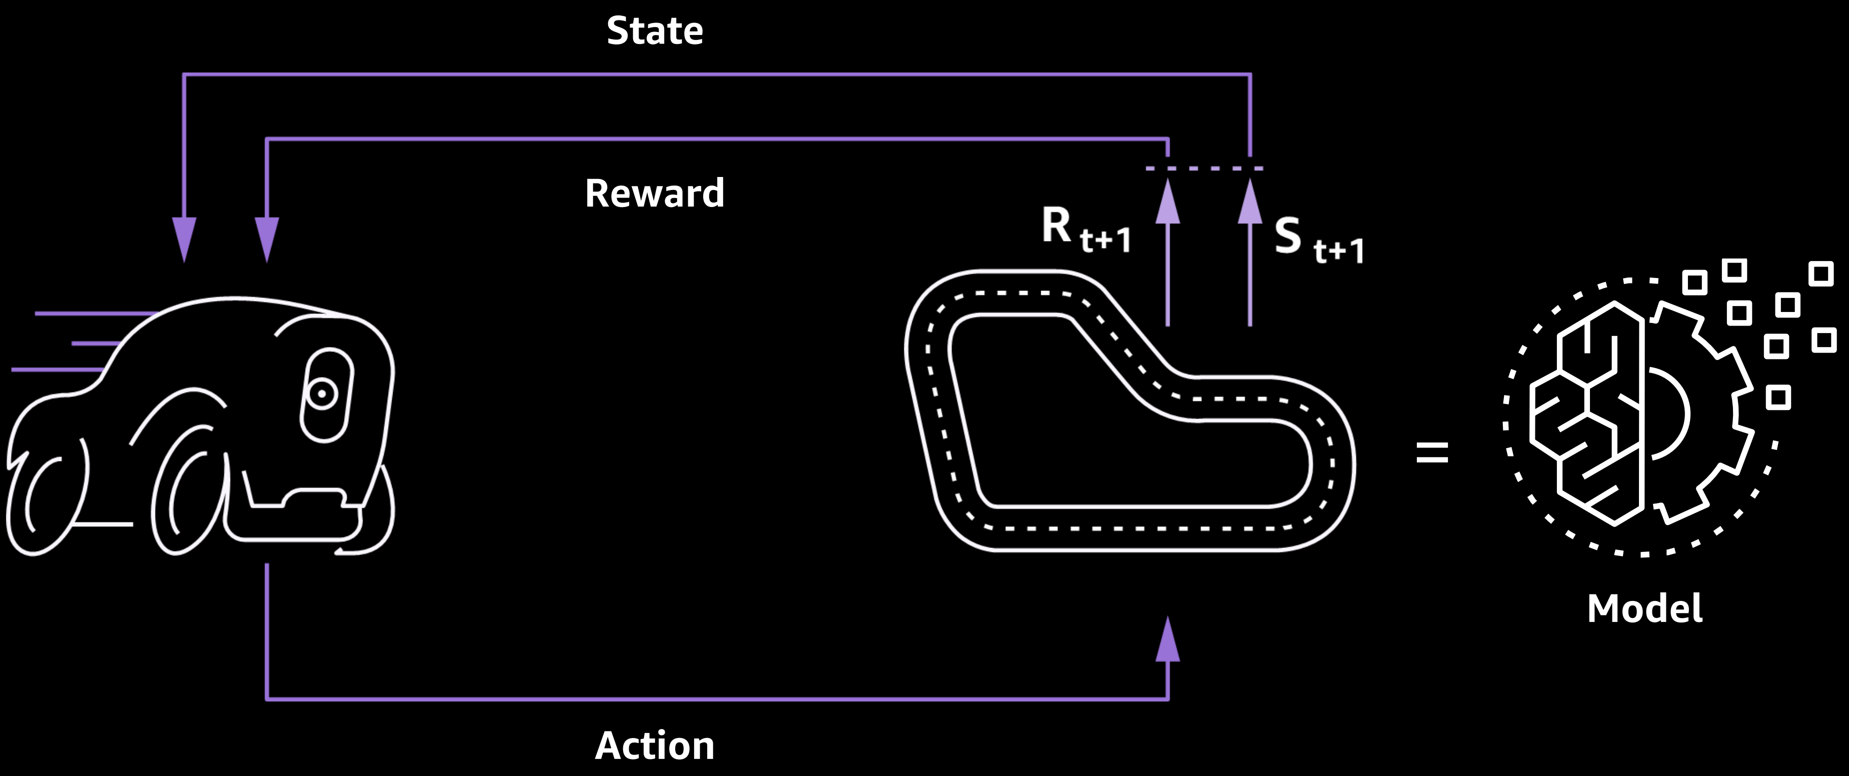

AWS DeepRacer neural network architecture

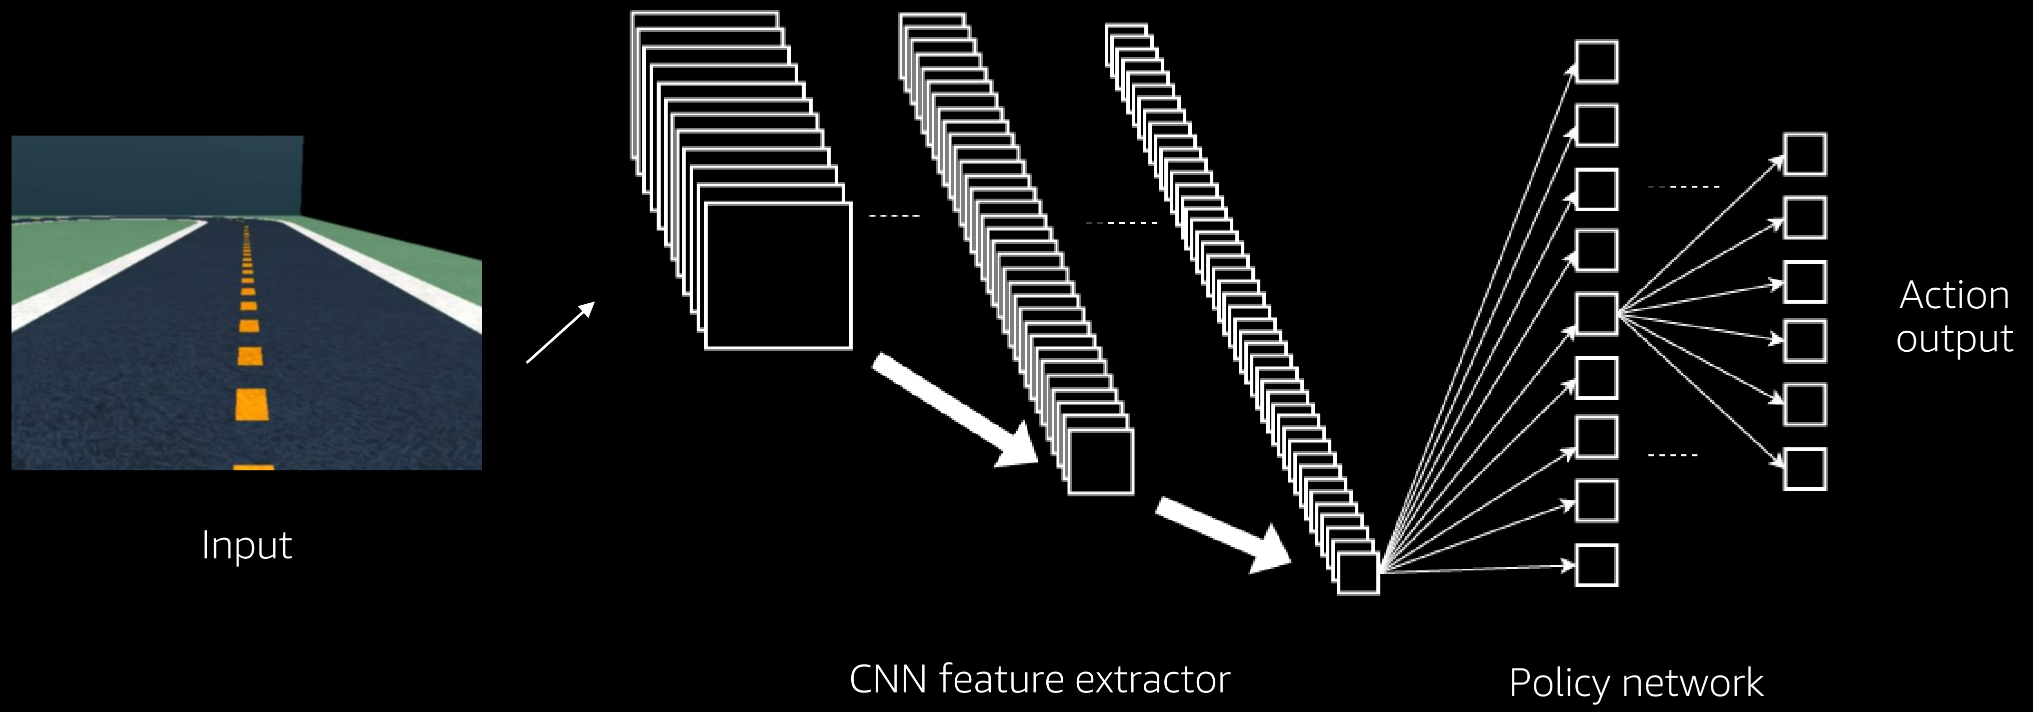

AWS DeepRacer neural network architecture

Input

AWS DeepRacer neural network architecture

Input

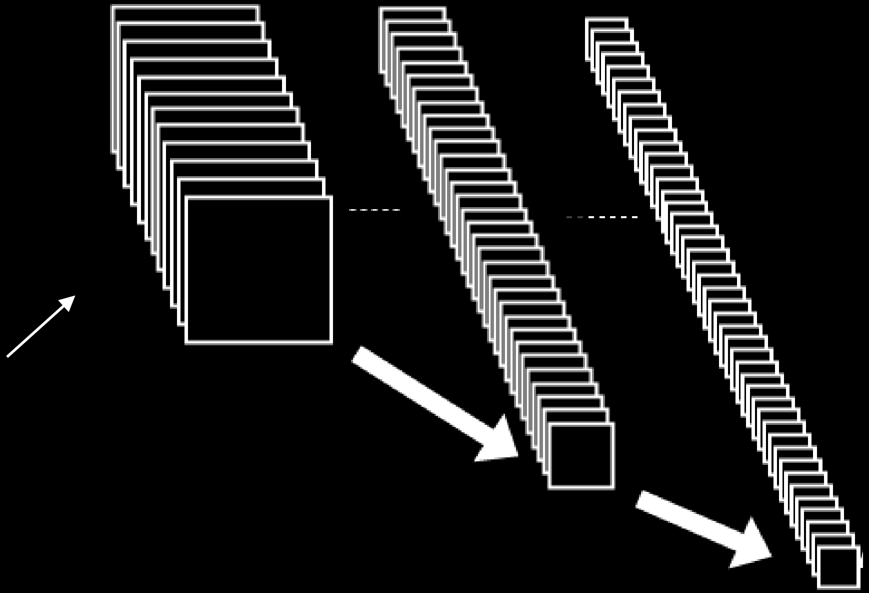

Convolutional neural network (CNN)
feature extractor

AWS DeepRacer neural network architecture

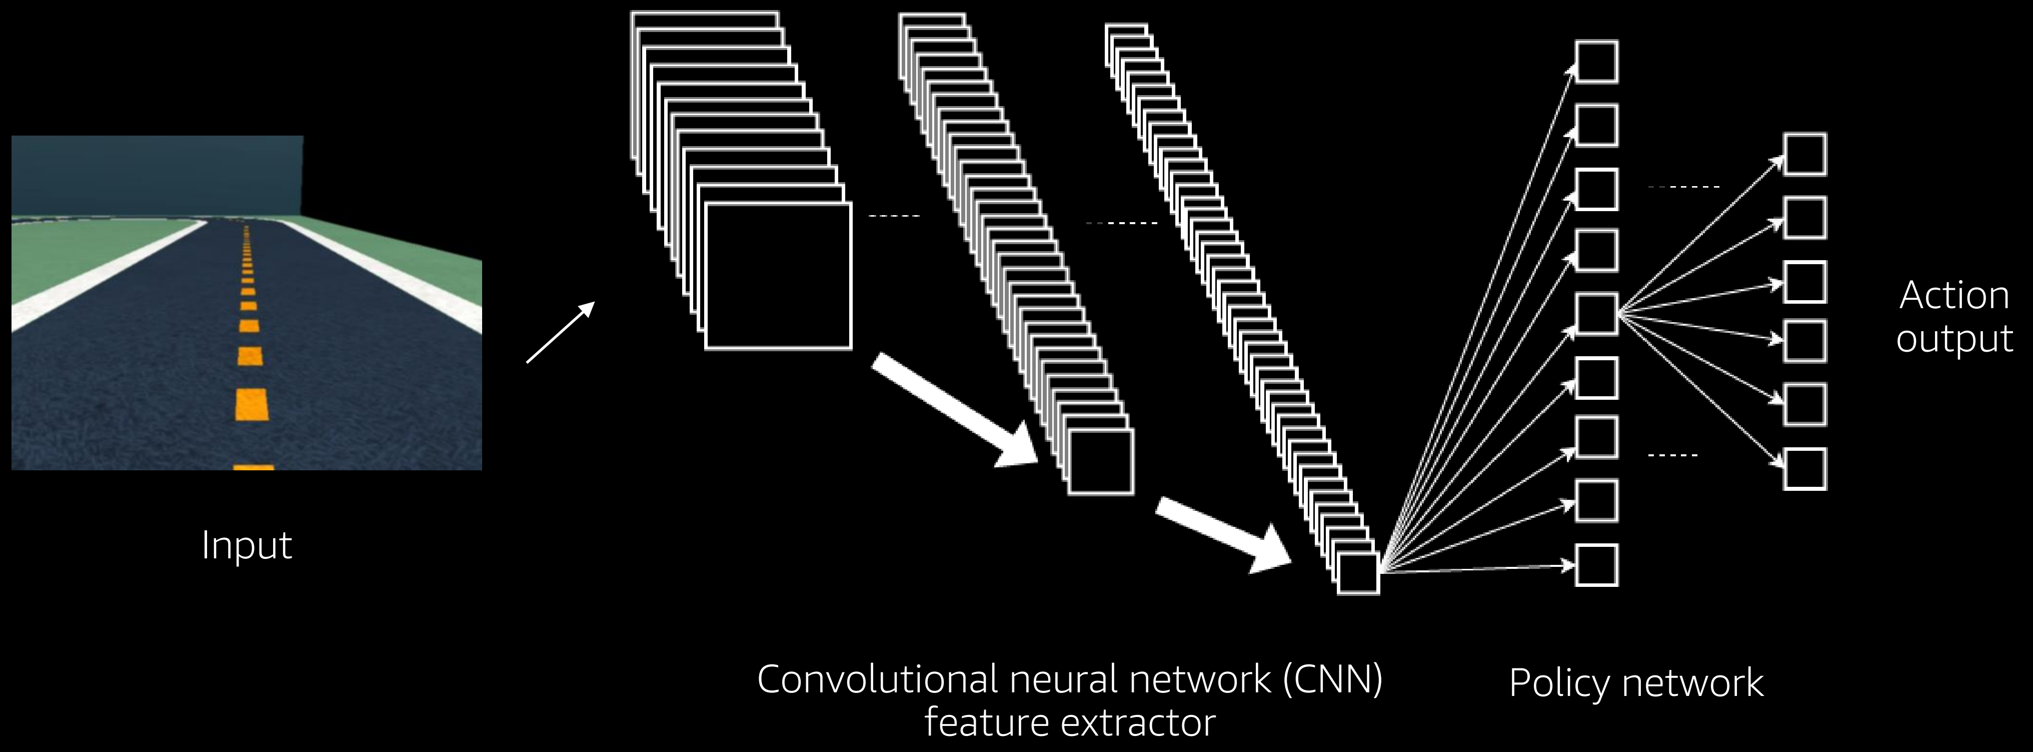

AWS DeepRacer console

AWS DeepRacer simulator architecture

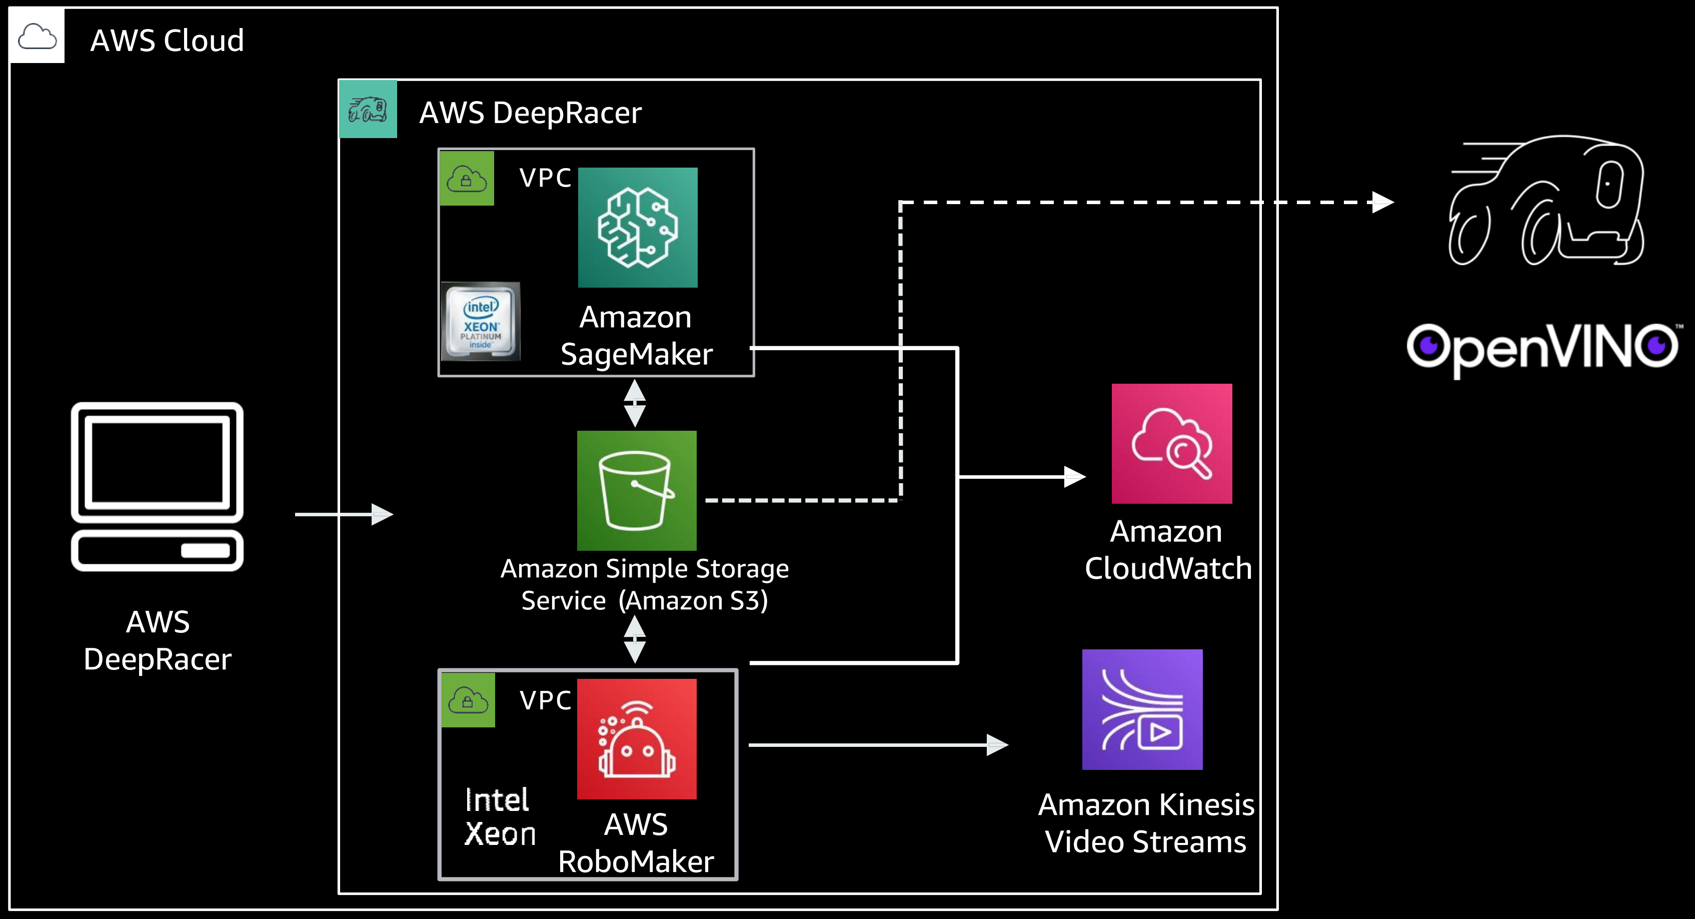

Programming your own reward function

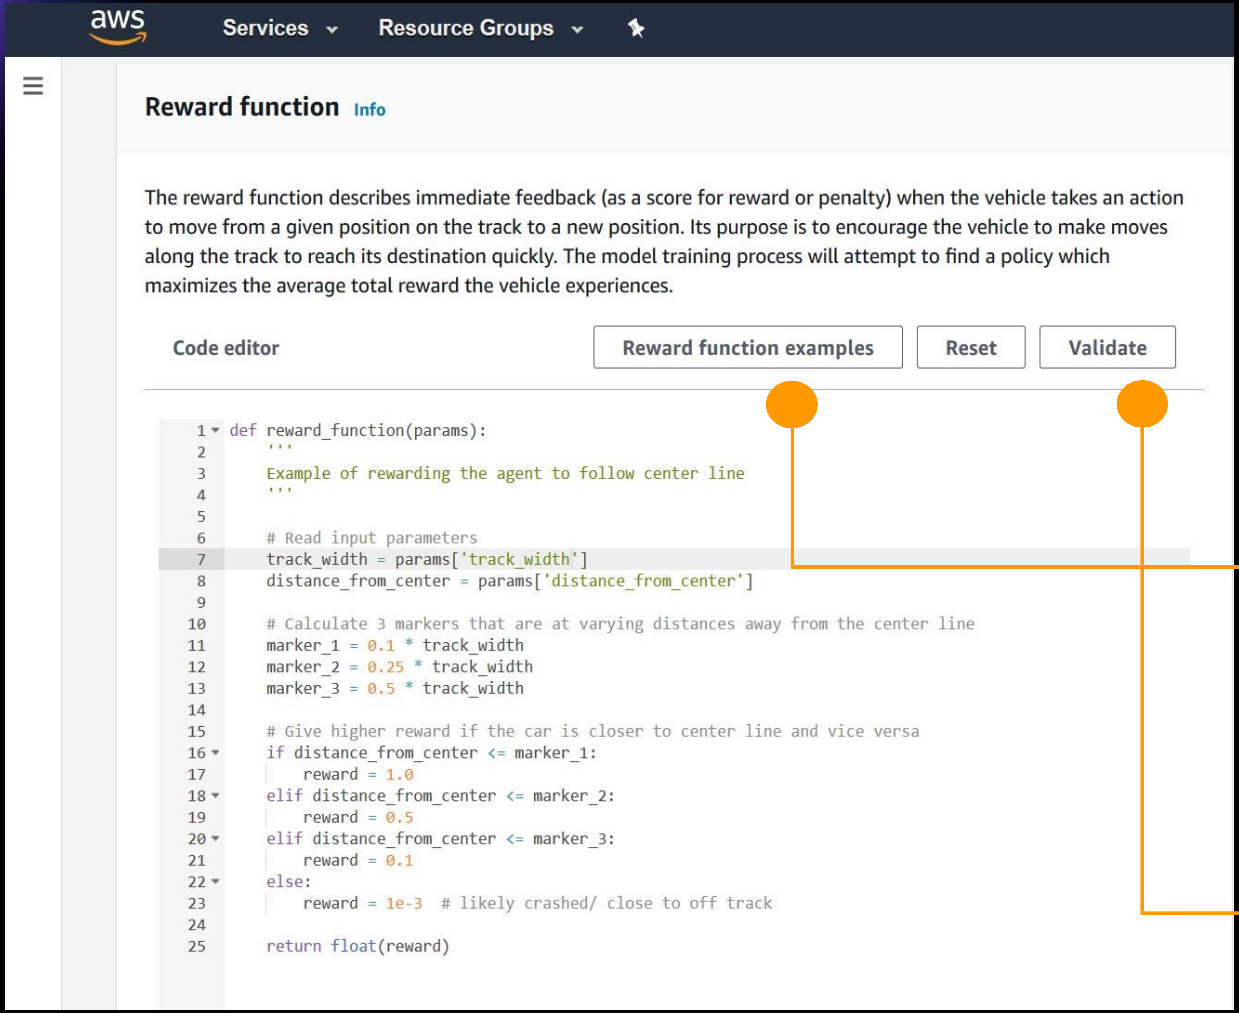

Reward function

The reward function describes immediate feedback (as a score for reward or penalty) when the vehicle takes an action to move from a given position on the track to a new position. Its purpose is to encourage the vehicle to make moves along the track to reach its destination quickly. The model training process will attempt to find a policy which maximizes the average total reward the vehicle experiences.

```python
def reward_function(params):
    '''
    Example of rewarding the agent to follow center line
    '''

    # Read input parameters
    track_width = params['track_width']
    distance_from_center = params['distance_from_center']

    # Calculate 3 markers that are at varying distances away from the center line
    marker_1 = 0.1 * track_width
    marker_2 = 0.25 * track_width
    marker_3 = 0.5 * track_width

    # Give higher reward if the car is closer to center line and vice versa
    if distance_from_center <= marker_1:
        reward = 1.0
    elif distance_from_center <= marker_2:
        reward = 0.5
    elif distance_from_center <= marker_3:
        reward = 0.1
    else:
        reward = 1e-3  # likely crashed/ close to off track

    return float(reward)
```

Code editor – Python 3 syntax

Three example reward functions

Code validation via AWS Lambda

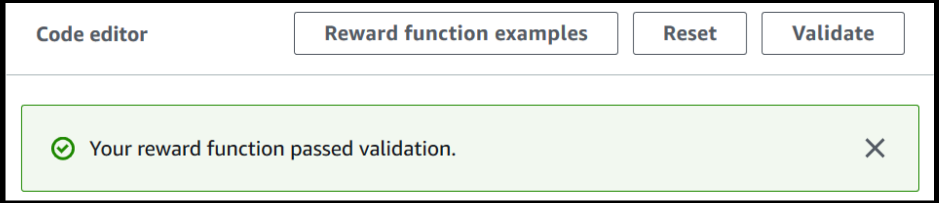

Your reward function passed validation.

Track components

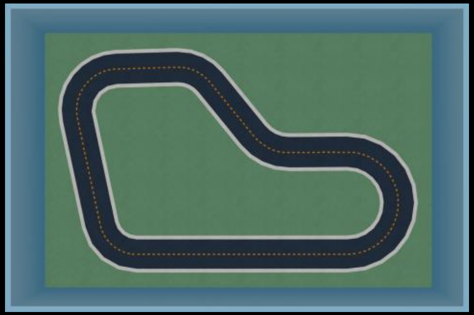

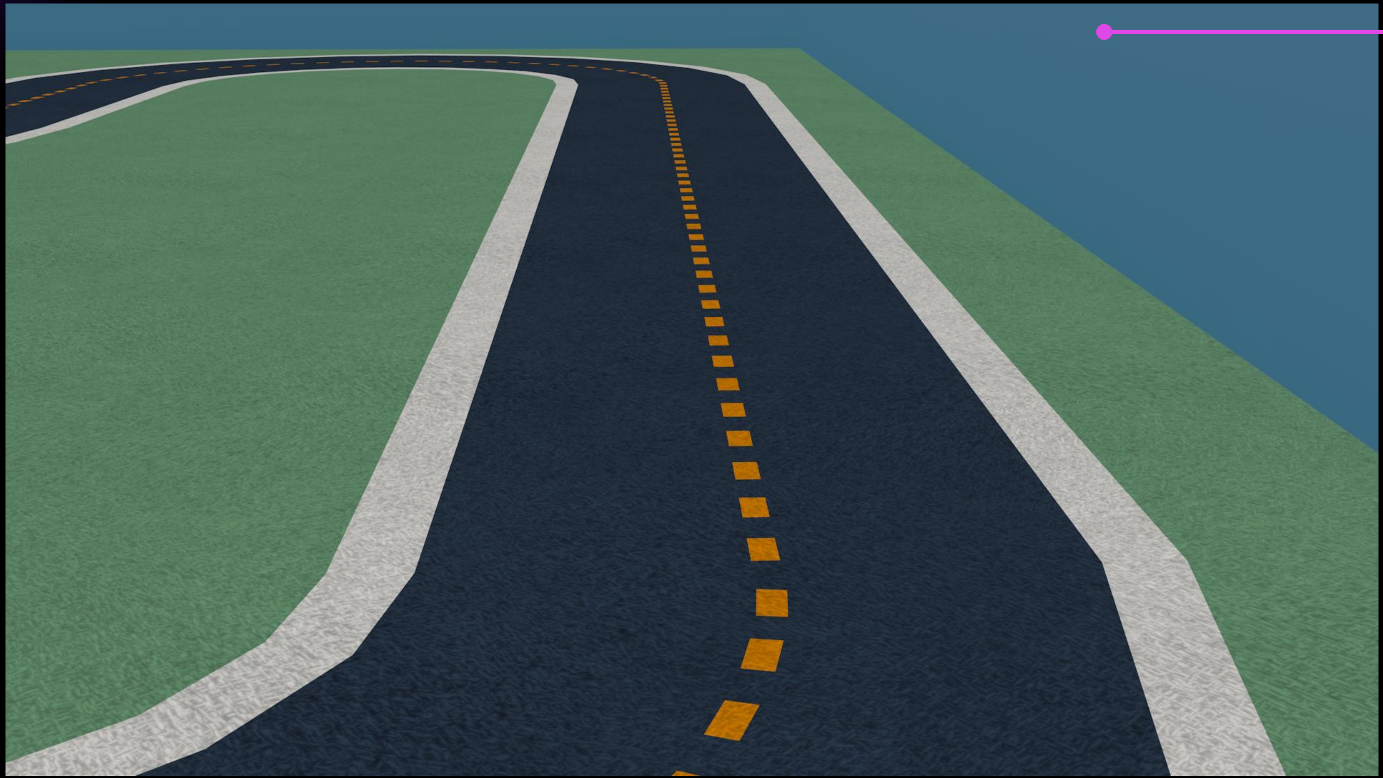

Track wall

Track components

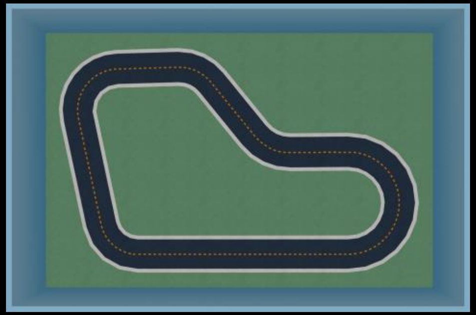

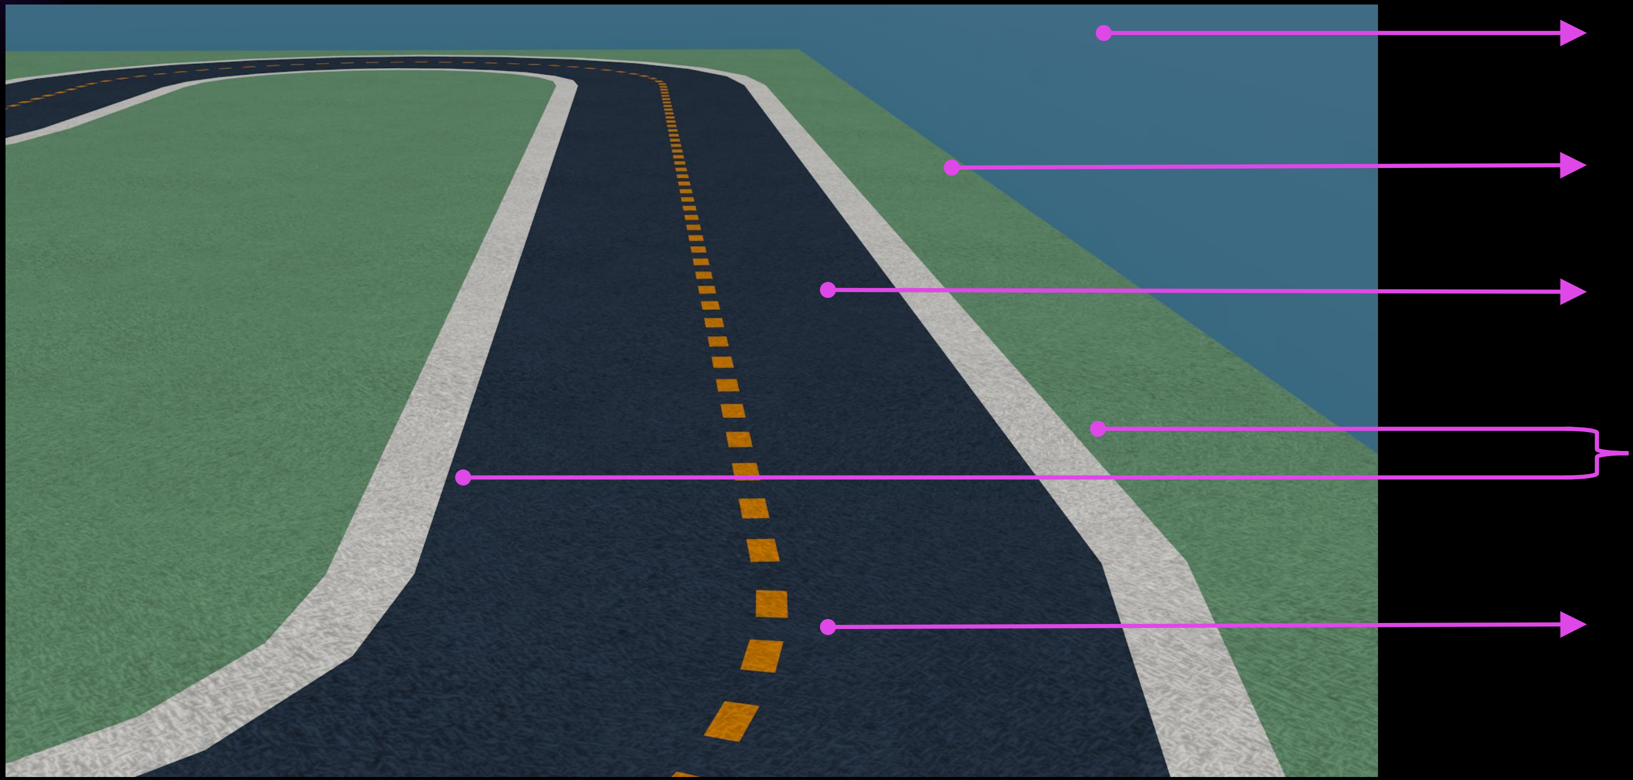

Track components

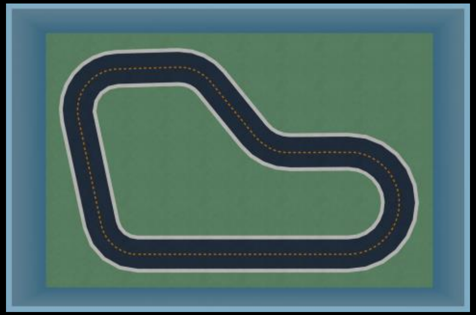

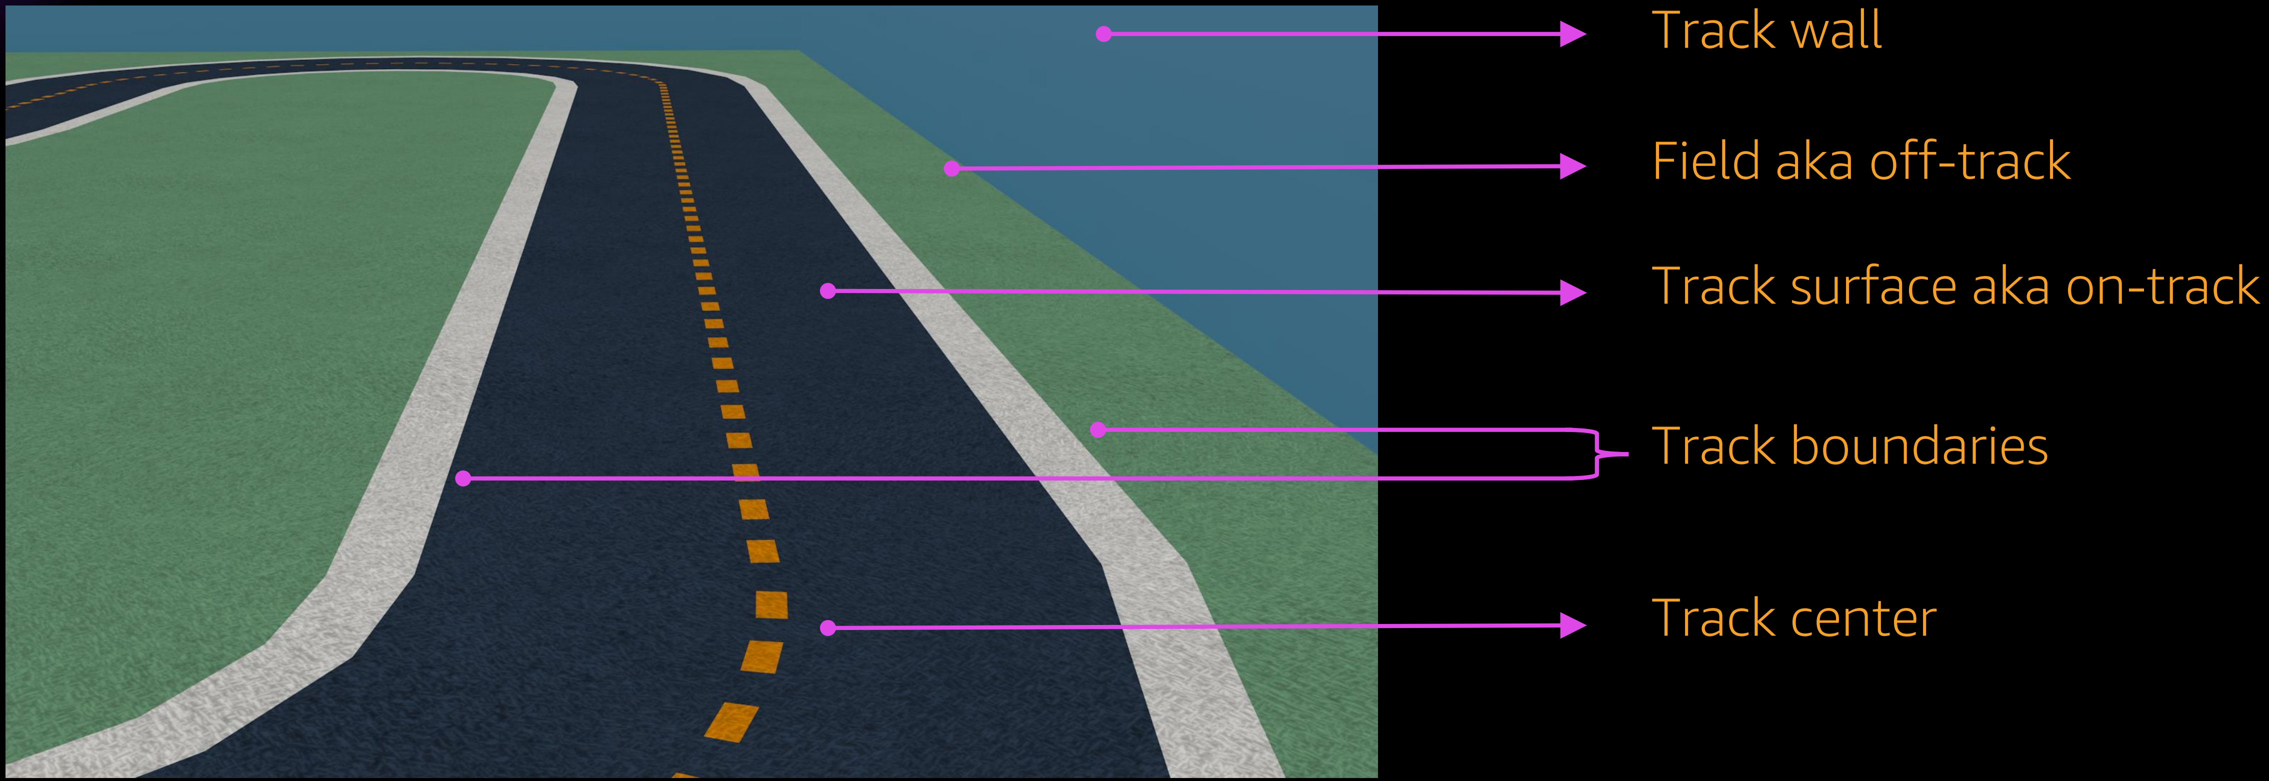

Coordinates system and track waypoints

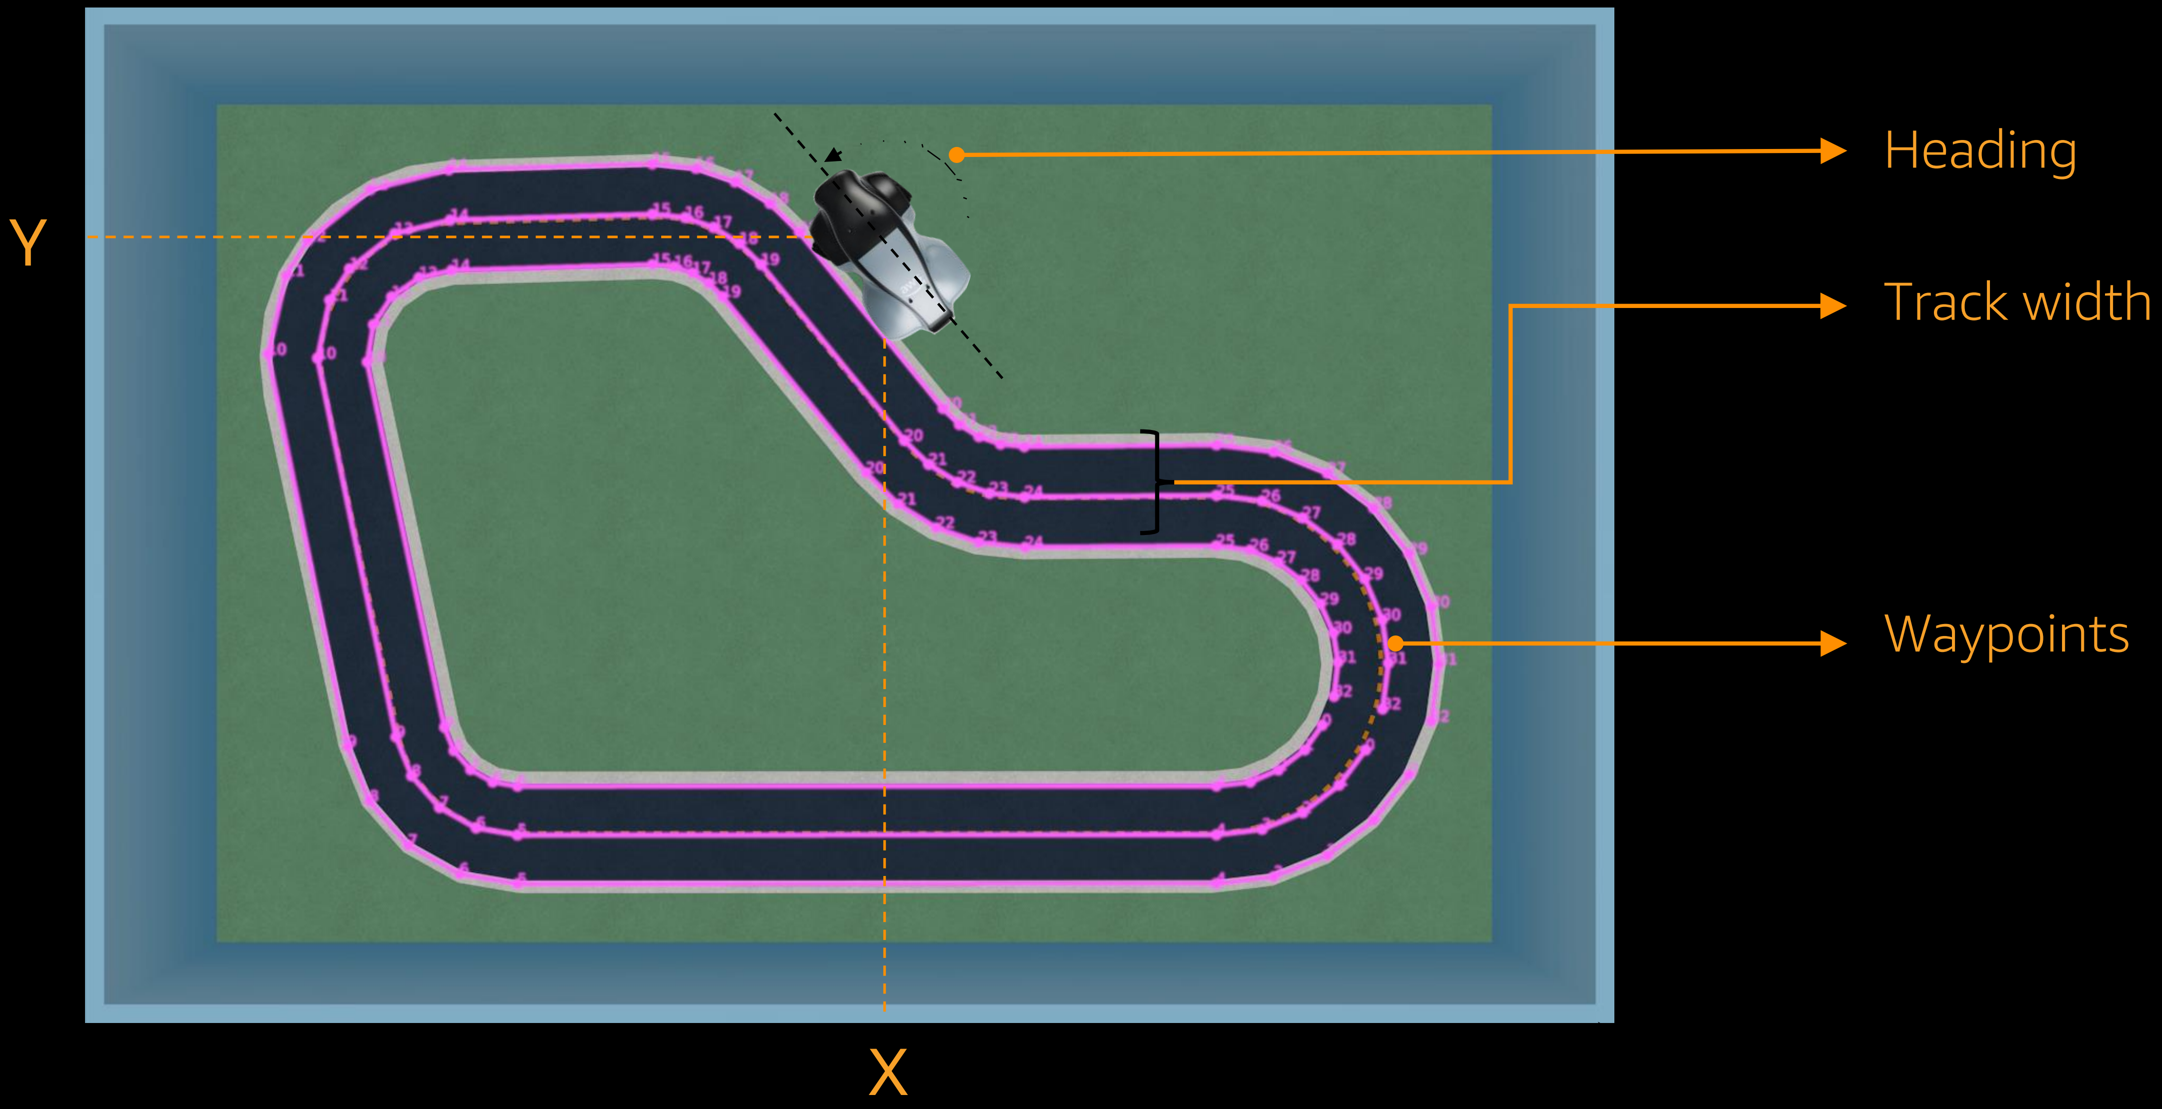

Example parameter – heading

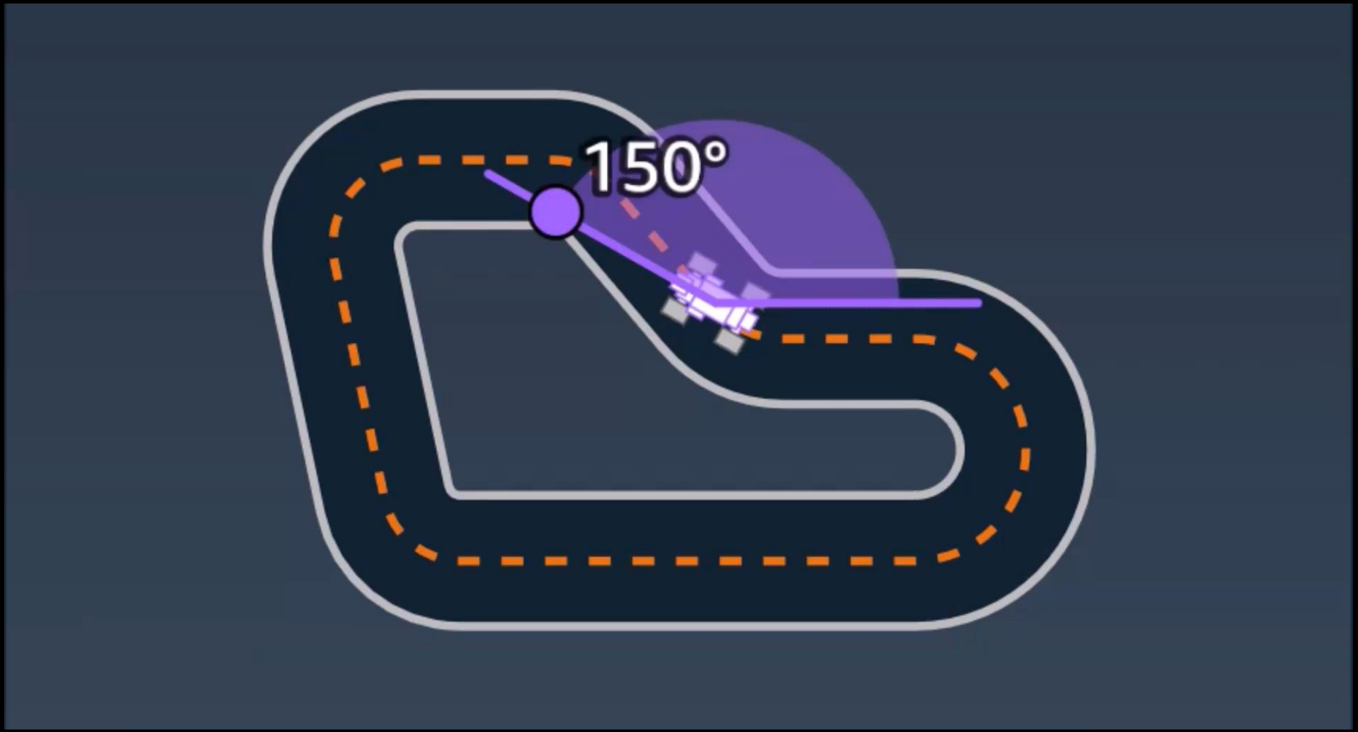

Example parameter – all_wheels_on_track

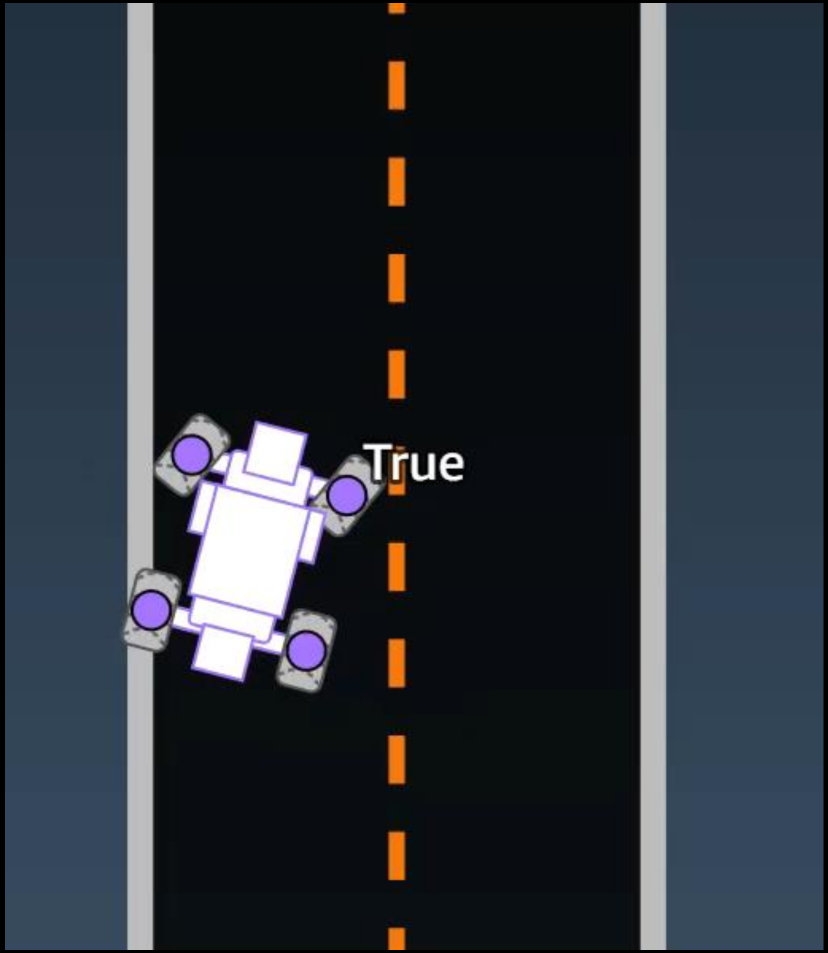

Example parameter – distance_from_center

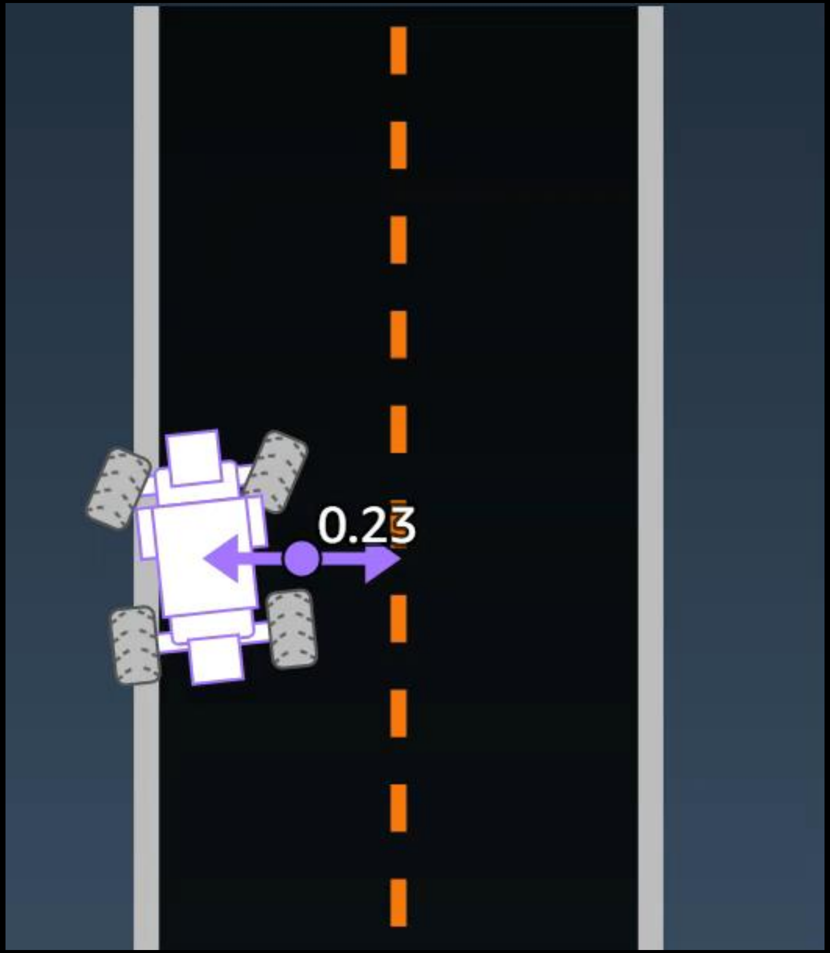

Customize your agent's sensor in the garage

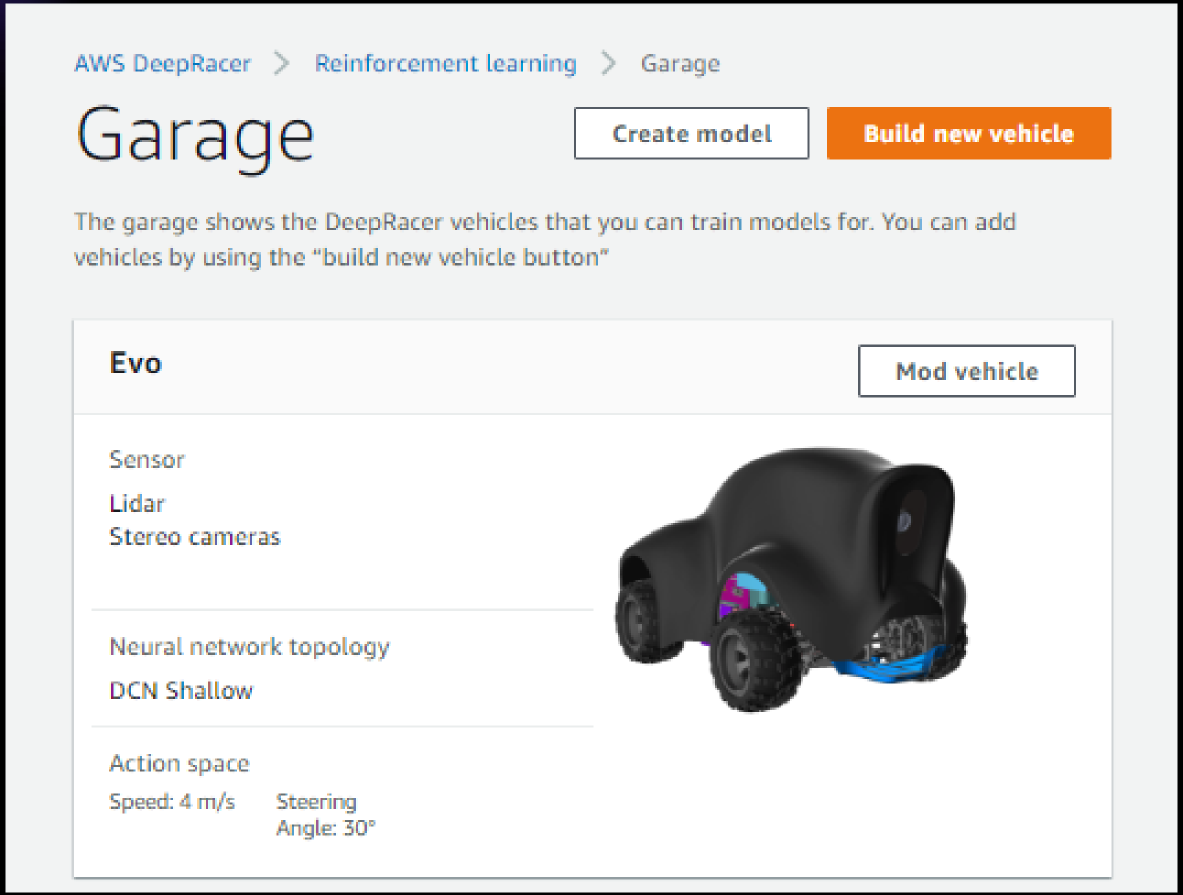

AWS DeepRacer > Reinforcement learning > Garage

Garage

The garage shows the DeepRacer vehicles that you can train models for. You can add vehicles by using the "build new vehicle button"

Evo

Sensor

Lidar
Stereo cameras

Neural network topology

DCN Shallow

Action space

Speed: 4 m/s

Steering Angle: 30°

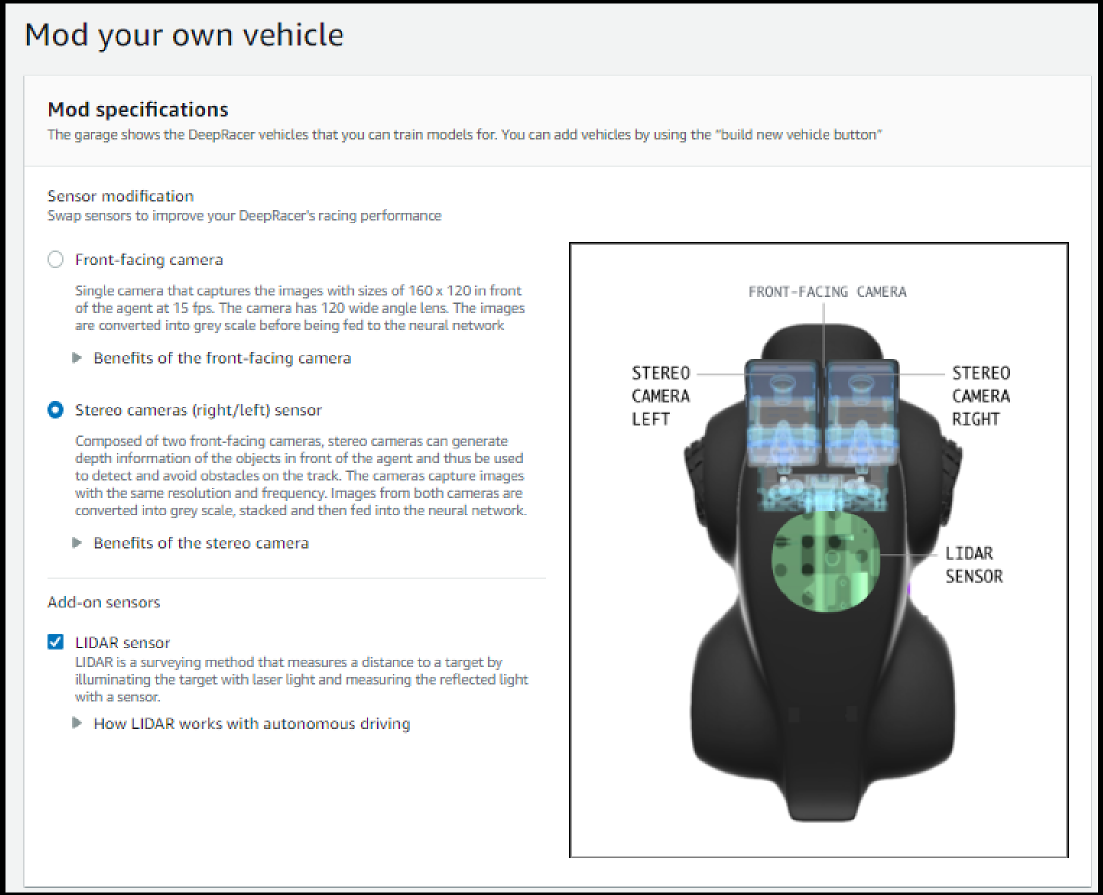

Mod your own vehicle

Mod specifications

The garage shows the DeepRacer vehicles that you can train models for. You can add vehicles by using the "build new vehicle button"

Sensor modification

Swap sensors to improve your DeepRacer's racing performance

○ Front-facing camera

Single camera that captures the images with sizes of 160 x 120 in front of the agent at 15 fps. The camera has 120 wide angle lens. The images are converted into grey scale before being fed to the neural network

▸ Benefits of the front-facing camera

[✓ Stereo cameras (right/left) sensor]

Composed of two front-facing cameras, stereo cameras can generate depth information of the objects in front of the agent and thus be used to detect and avoid obstacles on the track. The cameras capture images with the same resolution and frequency. Images from both cameras are converted into grey scale, stacked and then fed into the neural network.

▸ Benefits of the stereo camera

Add-on sensors

[✓ LIDAR sensor]

LIDAR is a surveying method that measures a distance to a target by illuminating the target with laser light and measuring the reflected light with a sensor.

▸ How LIDAR works with autonomous driving

Action space

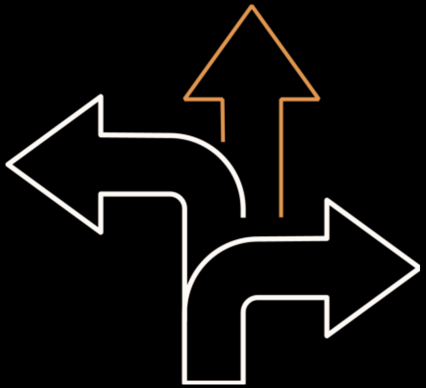

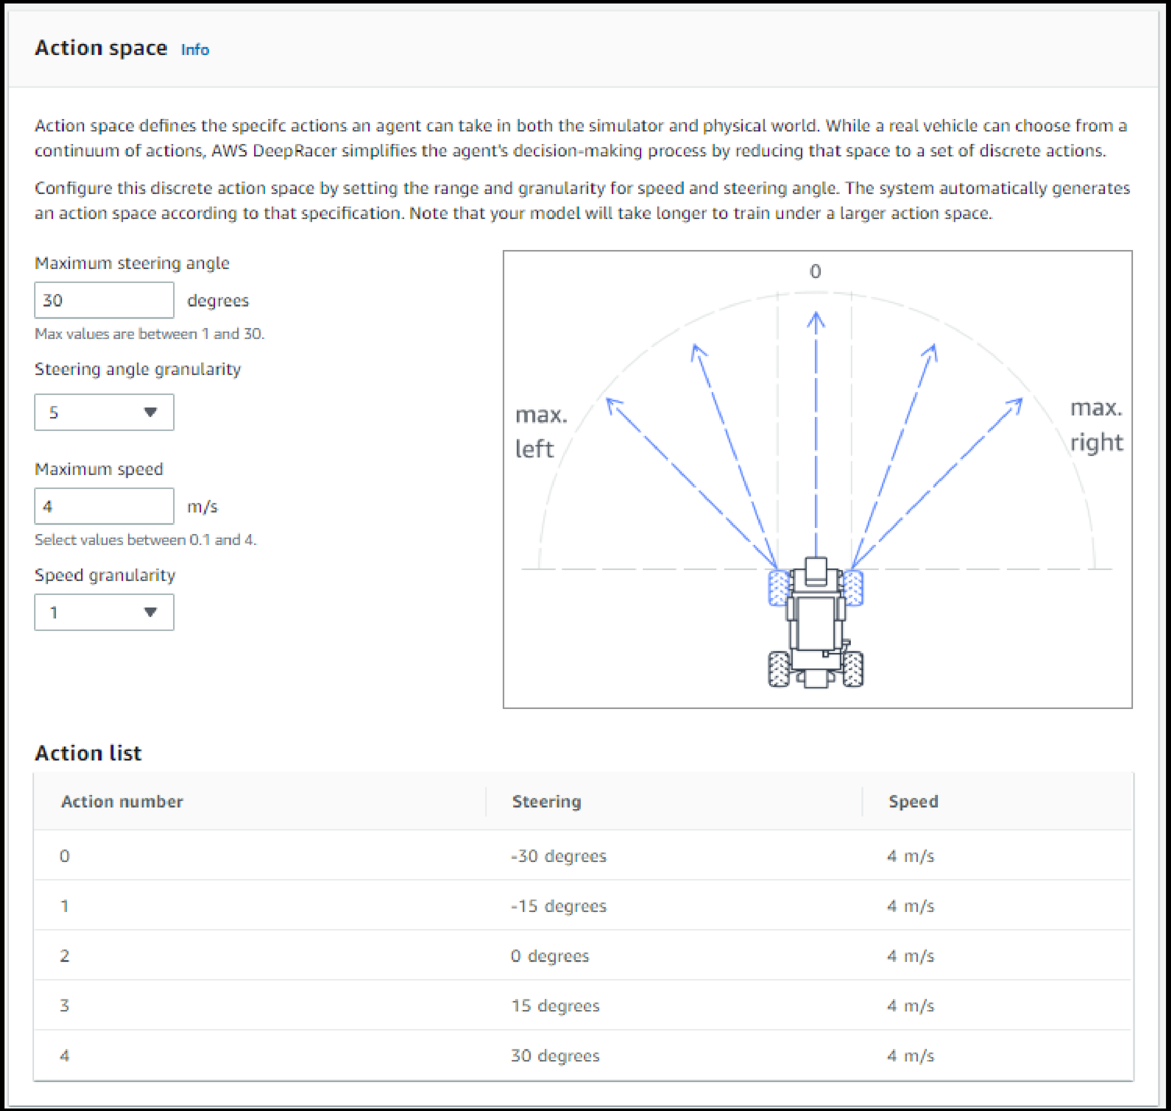

Action space Info

Action space defines the specifc actions an agent can take in both the simulator and physical world. While a real vehicle can choose from a continuum of actions, AWS DeepRacer simplifies the agent's decision-making process by reducing that space to a set of discrete actions.

Configure this discrete action space by setting the range and granularity for speed and steering angle. The system automatically generates an action space according to that specification. Note that your model will take longer to train under a larger action space.

Maximum steering angle: 30 degrees
Max values are between 1 and 30.

Steering angle granularity: 5

Maximum speed: 4 m/s
Select values between 0.1 and 4.

Speed granularity: 1

Action list

Action number	Steering	Speed
0	-30 degrees	4 m/s
1	-15 degrees	4 m/s
2	0 degrees	4 m/s
3	15 degrees	4 m/s
4	30 degrees	4 m/s

Use Python modules for your reward function

Consider using Python built-in modules such as numpy, scipy, and shapely to reduce the heavy lifting of waypoint math:

```
from shapely.geometry import Point, Polygon
from shapely.geometry.polygon import LinearRing, LineString

track = LinearRing(params[‘waypoints’])

first_object_location = Point(params[‘objects_location’][0])
object_distance_from_center_line = first_object_location.distance(track)
```

Demo

AWS DeepRacer League

https://console.aws.amazon.com/deepracer/home?region=us-east-1#league

Additional resources

- AWS DeepRacer Slack community: http://join.deepracing.io/
- GitHub: https://github.com/aws-samples/aws-deepracer-workshops/
- Free video course: https://www.aws.training/Details/eLearning?id=32143
- Tips: https://aws.amazon.com/deepracer/racing-tips/
- Intel distribution of OpenVINO toolkit: https://software.intel.com/en-us/openvino-toolkit
- AWS Developer Acceleration twitch channel - https://www.twitch.tv/devaxconnect

Visit the AI & Machine Learning resource hub for more resources

Dive deeper into these resources, get inspired and learn how you can use AI and machine learning to accelerate your business outcomes.

- The machine learning journey e-book
- 7 leading machine learning use cases e-book
- A strategic playbook for data, analytics, and machine learning e-book
  Accelerate machine learning innovation with the right cloud services & infrastructure e-book
- Choosing the right compute infrastructure for machine learning e-book
- Improving service and reducing costs in contact centers e-book
- Why ML is essential in your fight against online fraud e-book
- … and more!

Visit resource hub

https://bit.ly/3mwi59V

AWS Machine Learning (ML) Training and Certification

AWS is how you build machine learning skills

Courses built on the curriculum leveraged by Amazon's own teams. Learn from the experts at AWS.

aws.training/machinelearning

Flexibility to learn your way

Learn online with on-demand digital courses or live with virtual instructor-led training, plus hands-on labs and opportunities for practical application.

explore.skillbuilder.aws/learn

Validate your expertise

Demonstrate expertise in building, training, tuning, and deploying machine learning models with an industry-recognized credential.

aws.amazon.com/certification

Thank you for attending AWS Innovate – AI/ML Edition

We hope you found it interesting! A kind reminder to **complete the survey.**
Let us know what you thought of today's event and how we can improve the event experience for you in the future.

- aws-apj-marketing@amazon.com
- twitter.com/AWSCloud
- facebook.com/AmazonWebServices
- youtube.com/user/AmazonWebServices
- slideshare.net/AmazonWebServices
- twitch.tv/aws

Thank you!

Calvin Ngo