

# MATLAB EXPO

## Reinforcement Learning Workflows for AI



*Naga Pemmaraju*  
*Application Engineering*



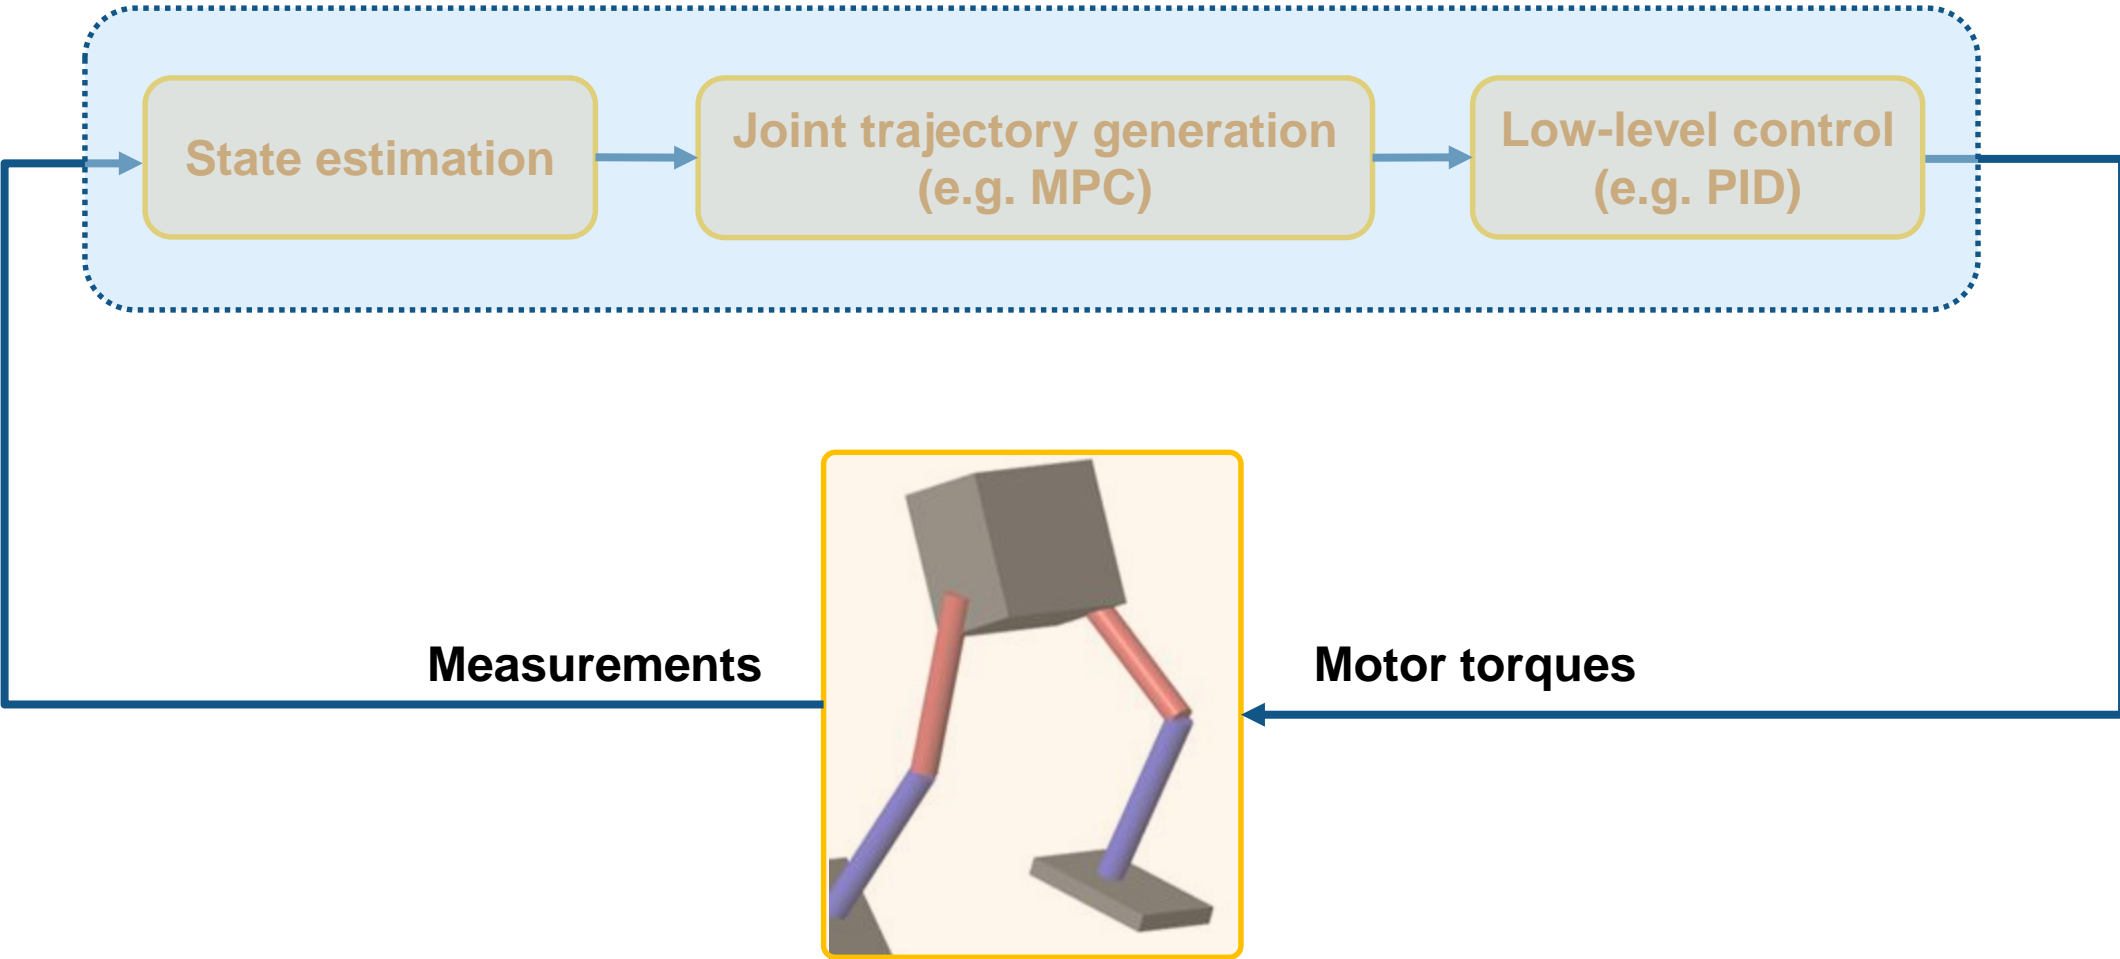
# Key Takeaways

- What is reinforcement learning and why should I care about it?
- How do I set up and solve a reinforcement learning problem?
- What are some common challenges?

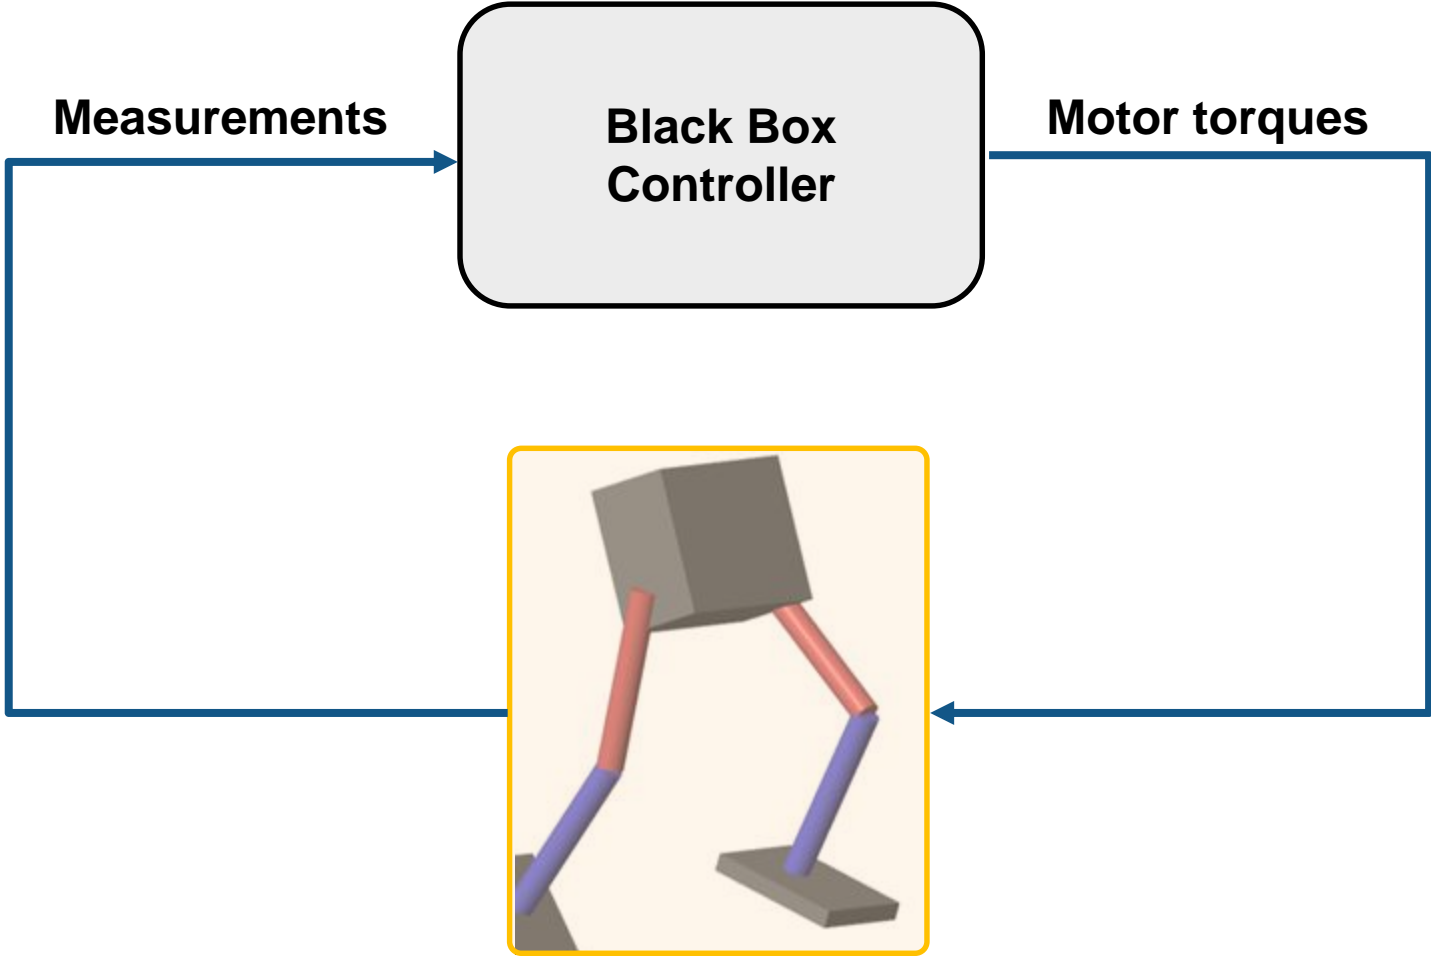
# Why Should You Care About Reinforcement Learning?






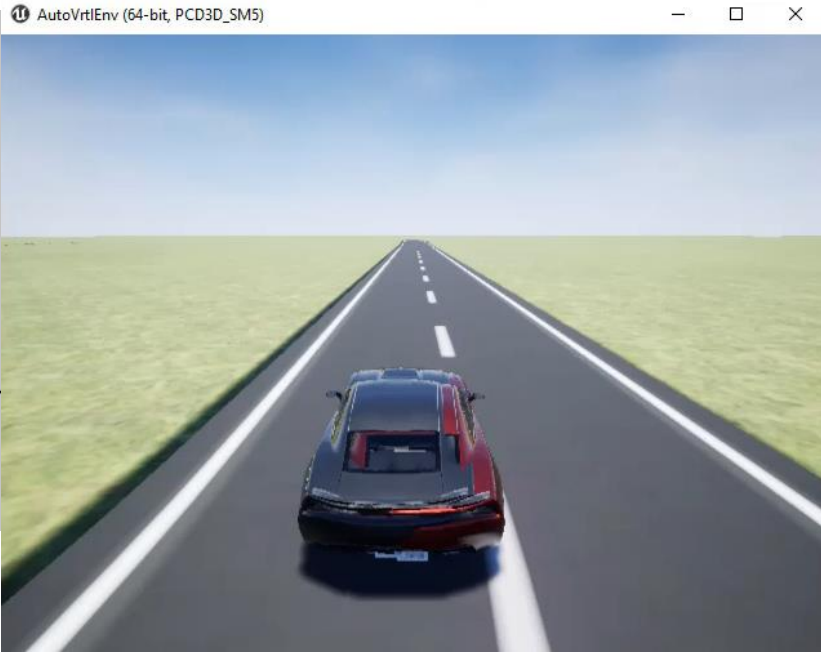

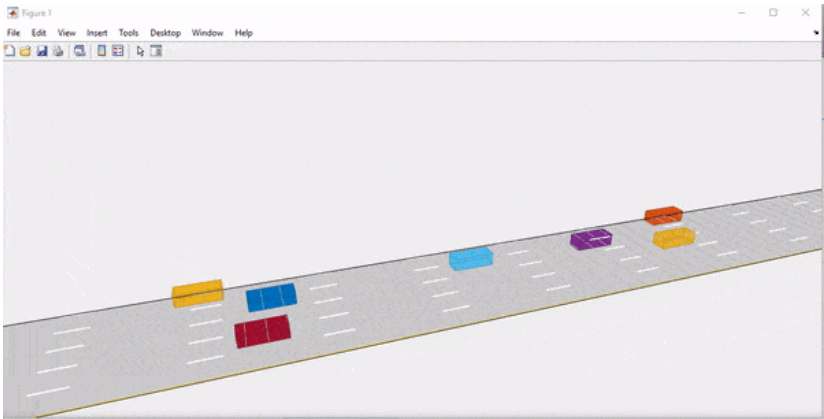
# One Approach Could Be...



# Any Alternatives?




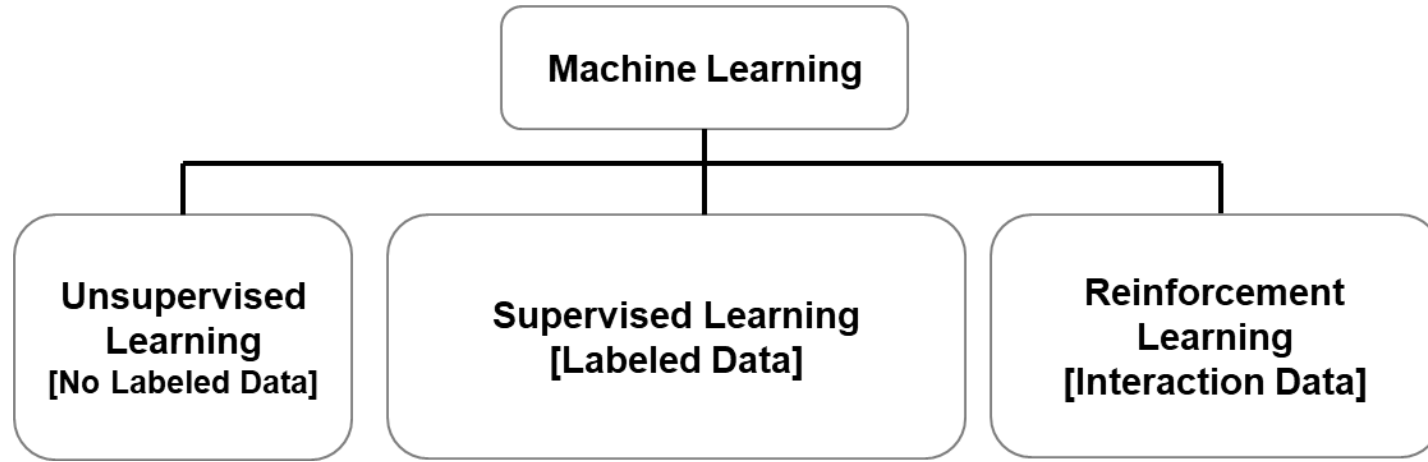
# Applications of Reinforcement Learning



# What is reinforcement learning?


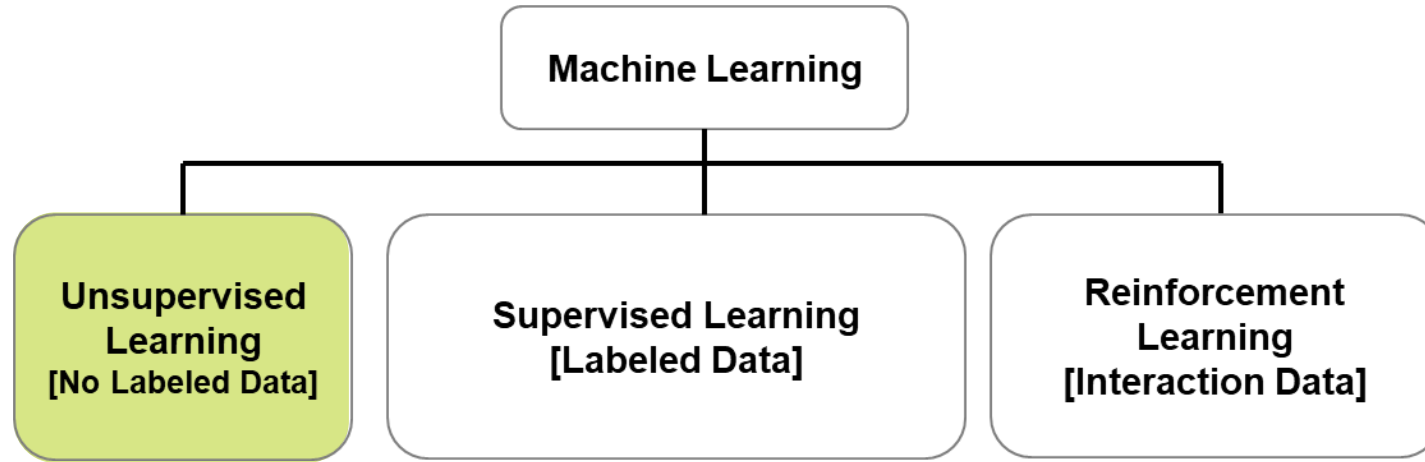
Type of machine learning that trains an **'agent'** through trial & error interactions with an **environment**

# Reinforcement Learning vs Machine Learning vs Deep Learning

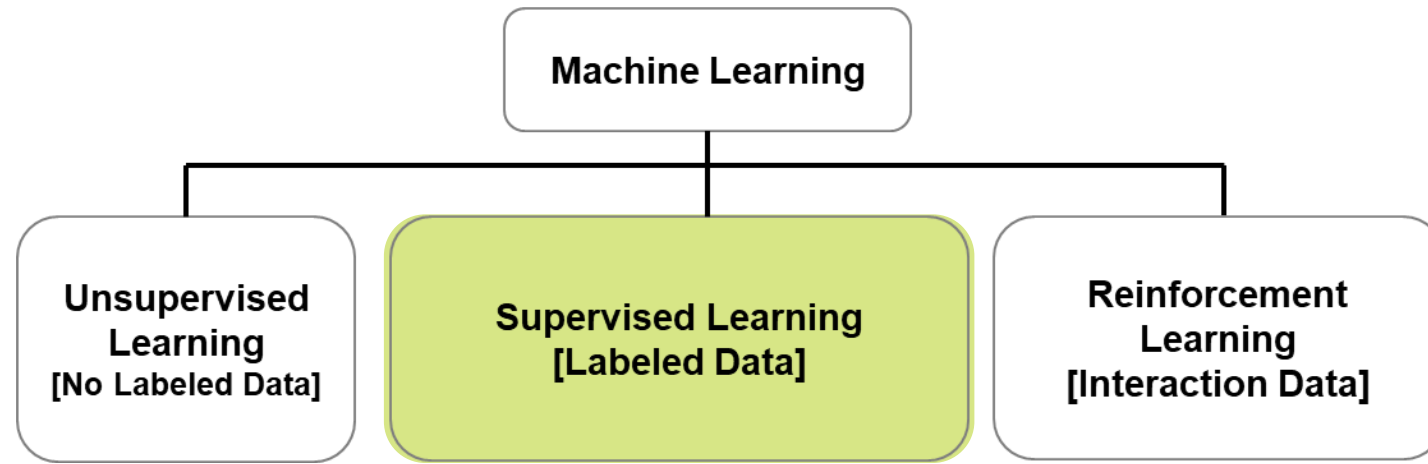




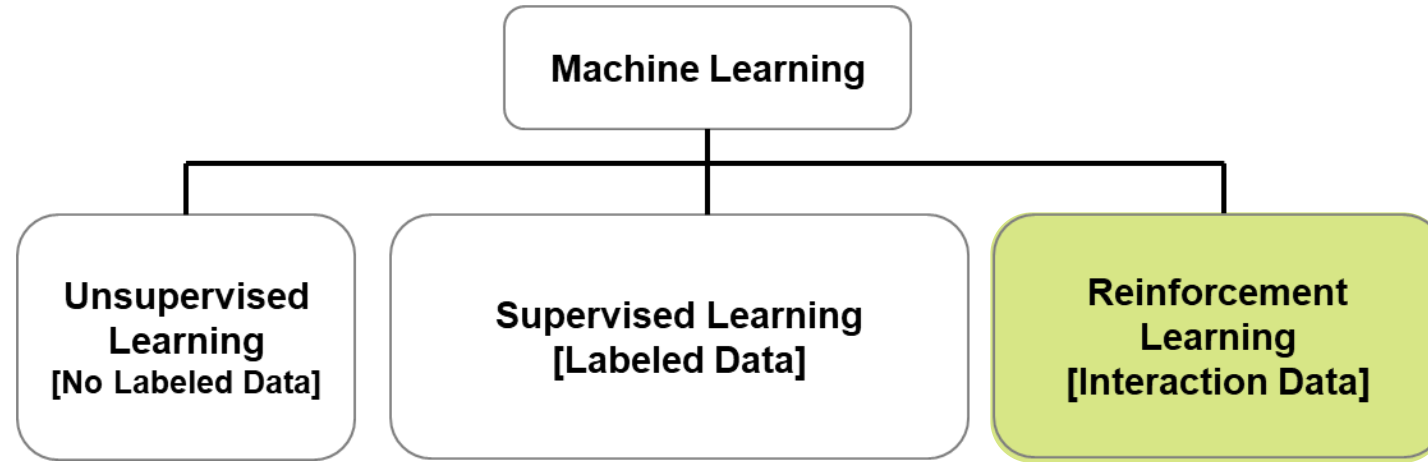
# Reinforcement Learning vs Machine Learning vs Deep Learning



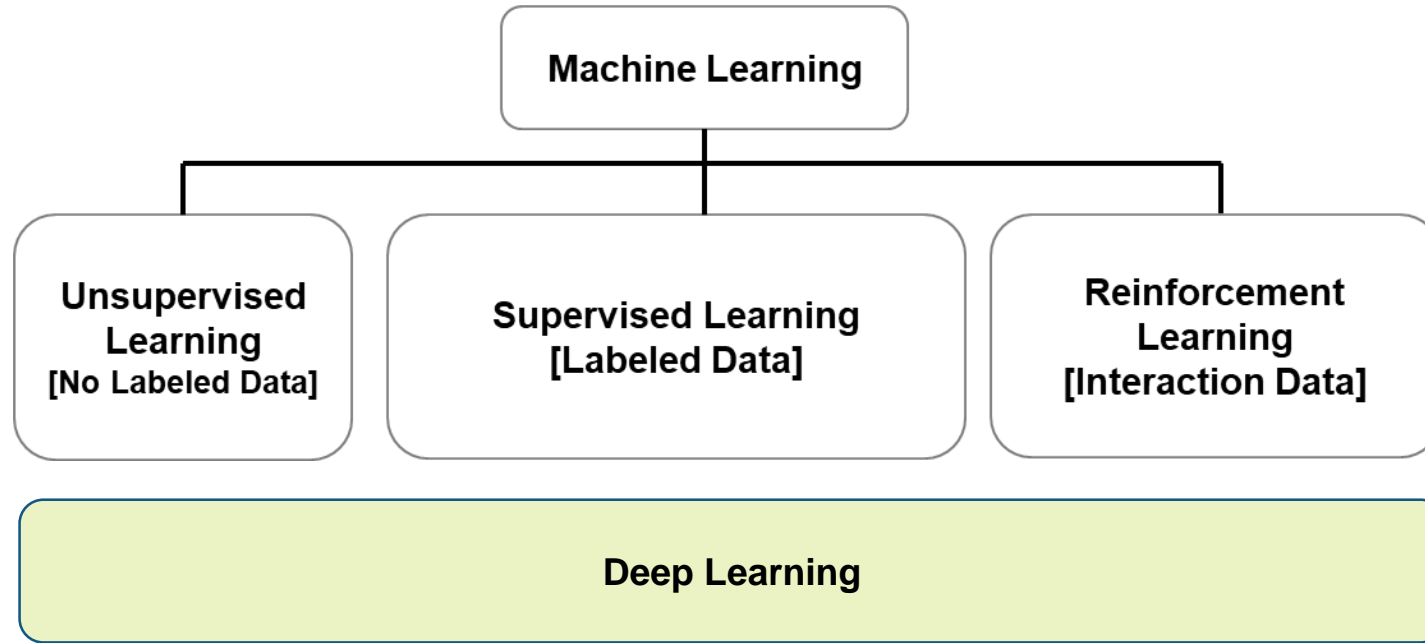
# Reinforcement Learning vs Machine Learning vs Deep Learning



# Reinforcement Learning vs Machine Learning vs Deep Learning



# Reinforcement Learning vs Machine Learning vs Deep Learning

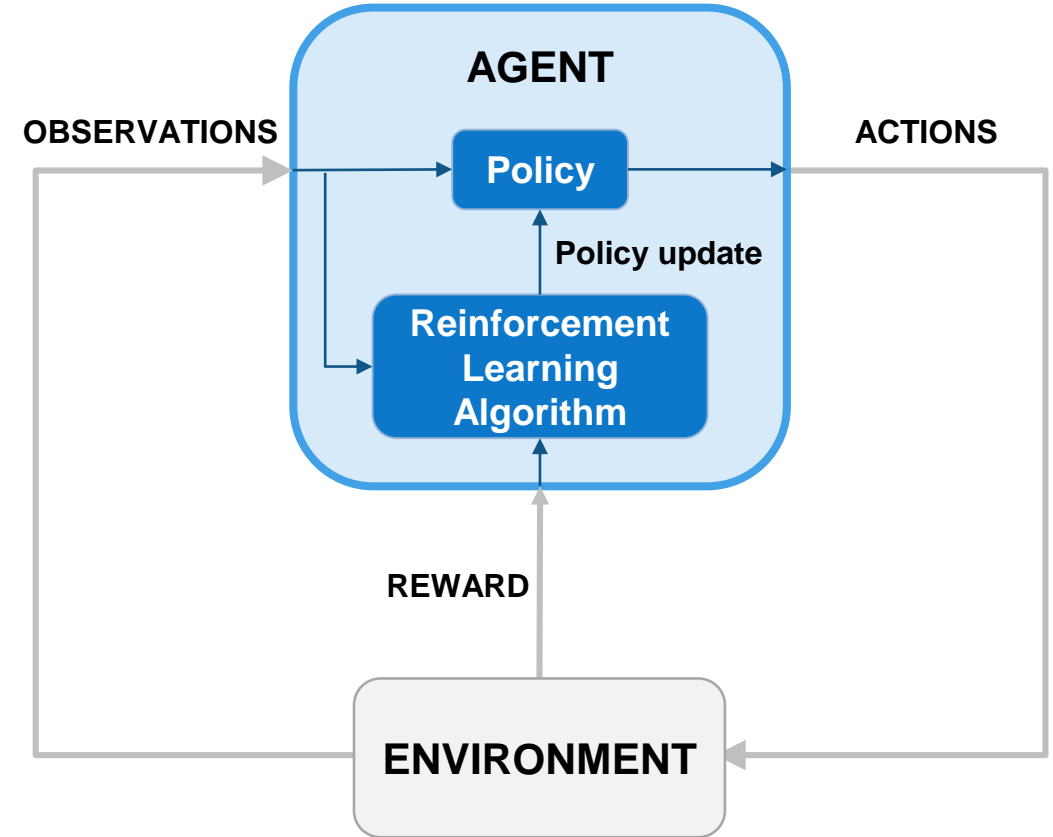
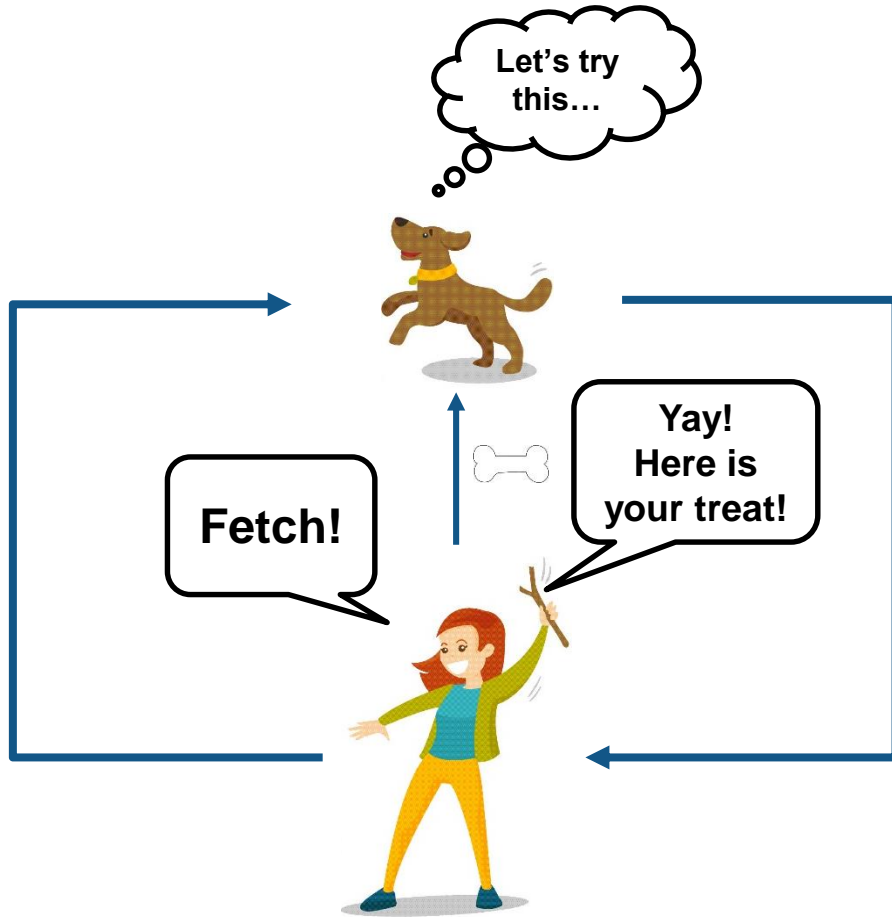


What about deep learning?

Complex reinforcement learning problems typically need deep neural networks  
*[Deep Reinforcement Learning]*

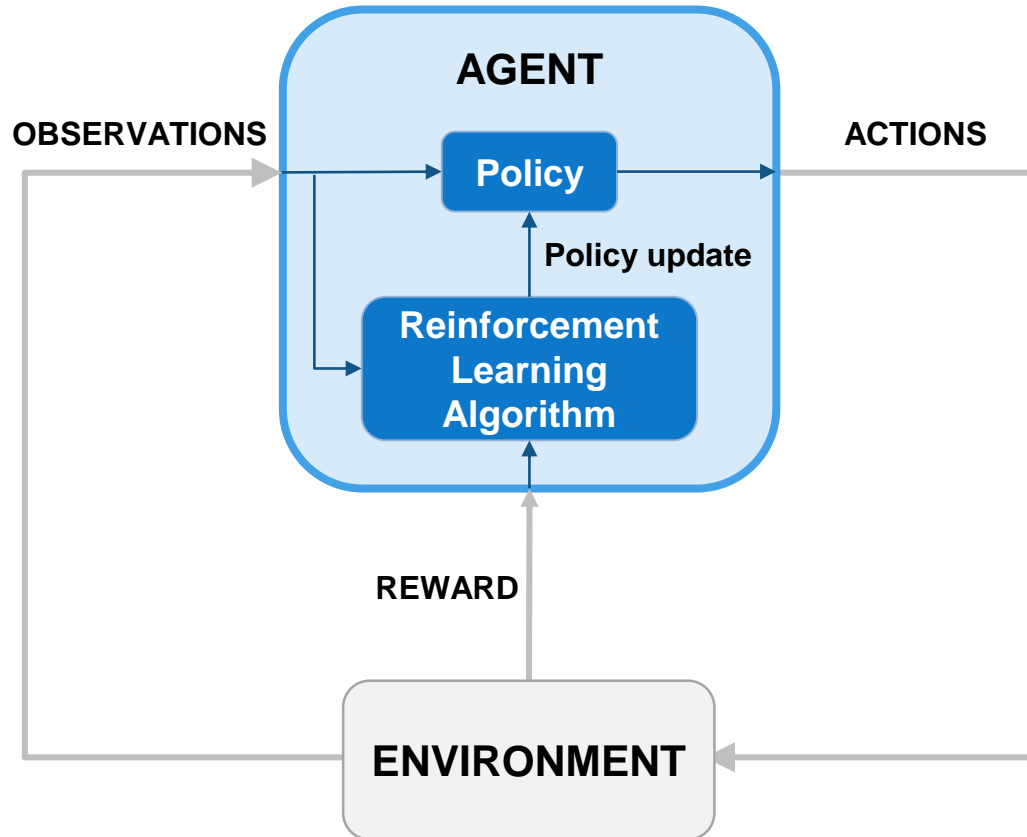
# How does reinforcement learning training work?

## Analogies with pet training



# Reinforcement Learning Concepts

## Training a self-driving car

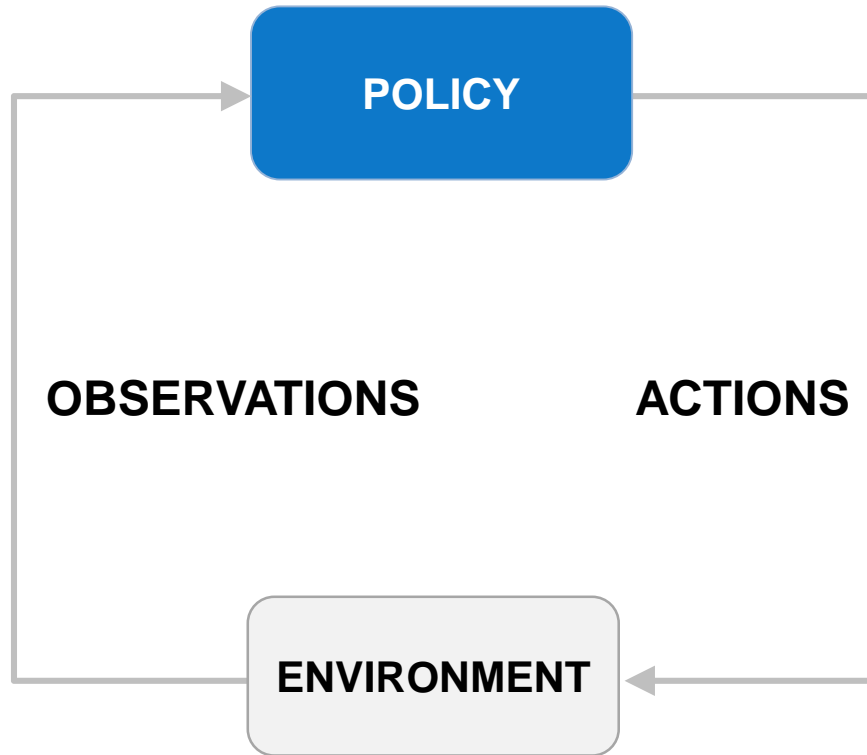


- Vehicle's computer...  
(**agent**)
- is reading sensor measurements from LIDAR, cameras,...  
(**observations**)
- that represent road conditions, vehicle position,...  
(**environment**)
- and generates steering, braking, throttle commands,...  
(**action**)
- based on an internal state-to-action mapping...  
(**policy**)
- that tries to optimize, e.g., lap time & fuel efficiency...  
(**reward**).
  
- The policy is updated through repeated trial-and-error by a **reinforcement learning algorithm**

# Reinforcement Learning Concepts

## Training a self-driving car

After training, only trained policy is needed

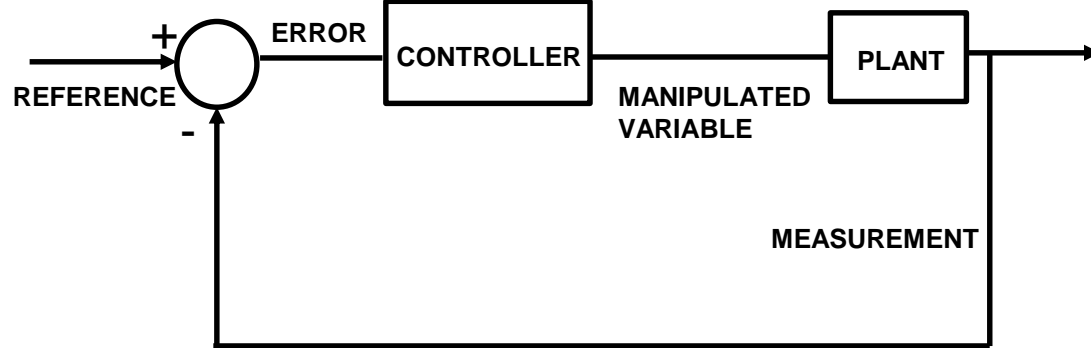


- Vehicle's computer uses the final state-to-action mapping... (**policy**)
- to generate steering, braking, throttle commands,... (**action**)
- based on sensor readings from LIDAR, cameras,... (**observations**)
- that represent road conditions, vehicle position,... (**environment**).

By definition, this trained policy is optimizing lap time & fuel efficiency

# Reinforcement Learning vs Controls

Control system



Adaptation mechanism

Error/Cost function

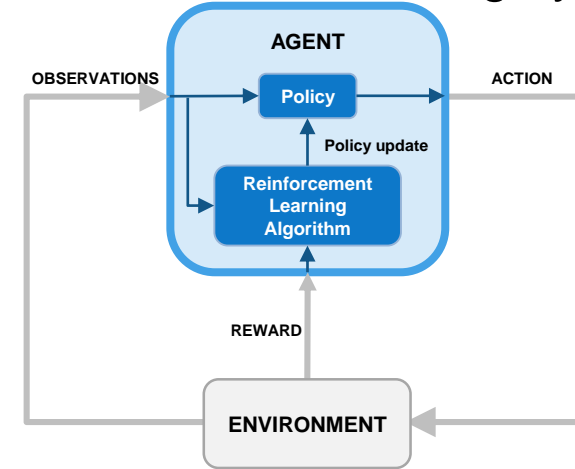
Manipulated variable

Measurement

Plant

Controller

Reinforcement learning system



RL Algorithm

Reward

Action

Observation

Environment

Policy

Reinforcement learning has parallels to **control system design**

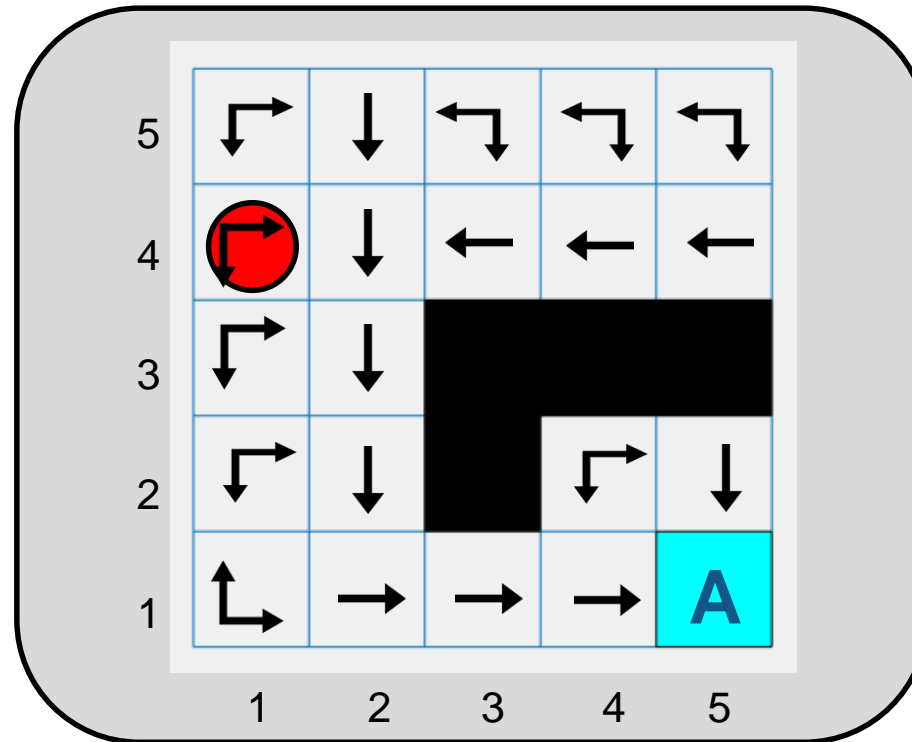



# Policy Representation and Deep Learning

## Representation options

- Look-up table
- Polynomials

Observations




Next action

Look-up tables **do not scale** well

# Policy Representation and Deep Learning

## Representation options

- Look-up table
- Polynomials
- **(Deep) neural networks**



Neural networks allow representation of **complex policies**

# How do I set up and solve a reinforcement learning problem?

# Reinforcement Learning Workflow


- Simulation models or real hardware
- Virtual models are safer and cheaper




Environment



Reward



Policy  
representation



Agent



Training



Deployment

- Deep network? Table? Polynomial?

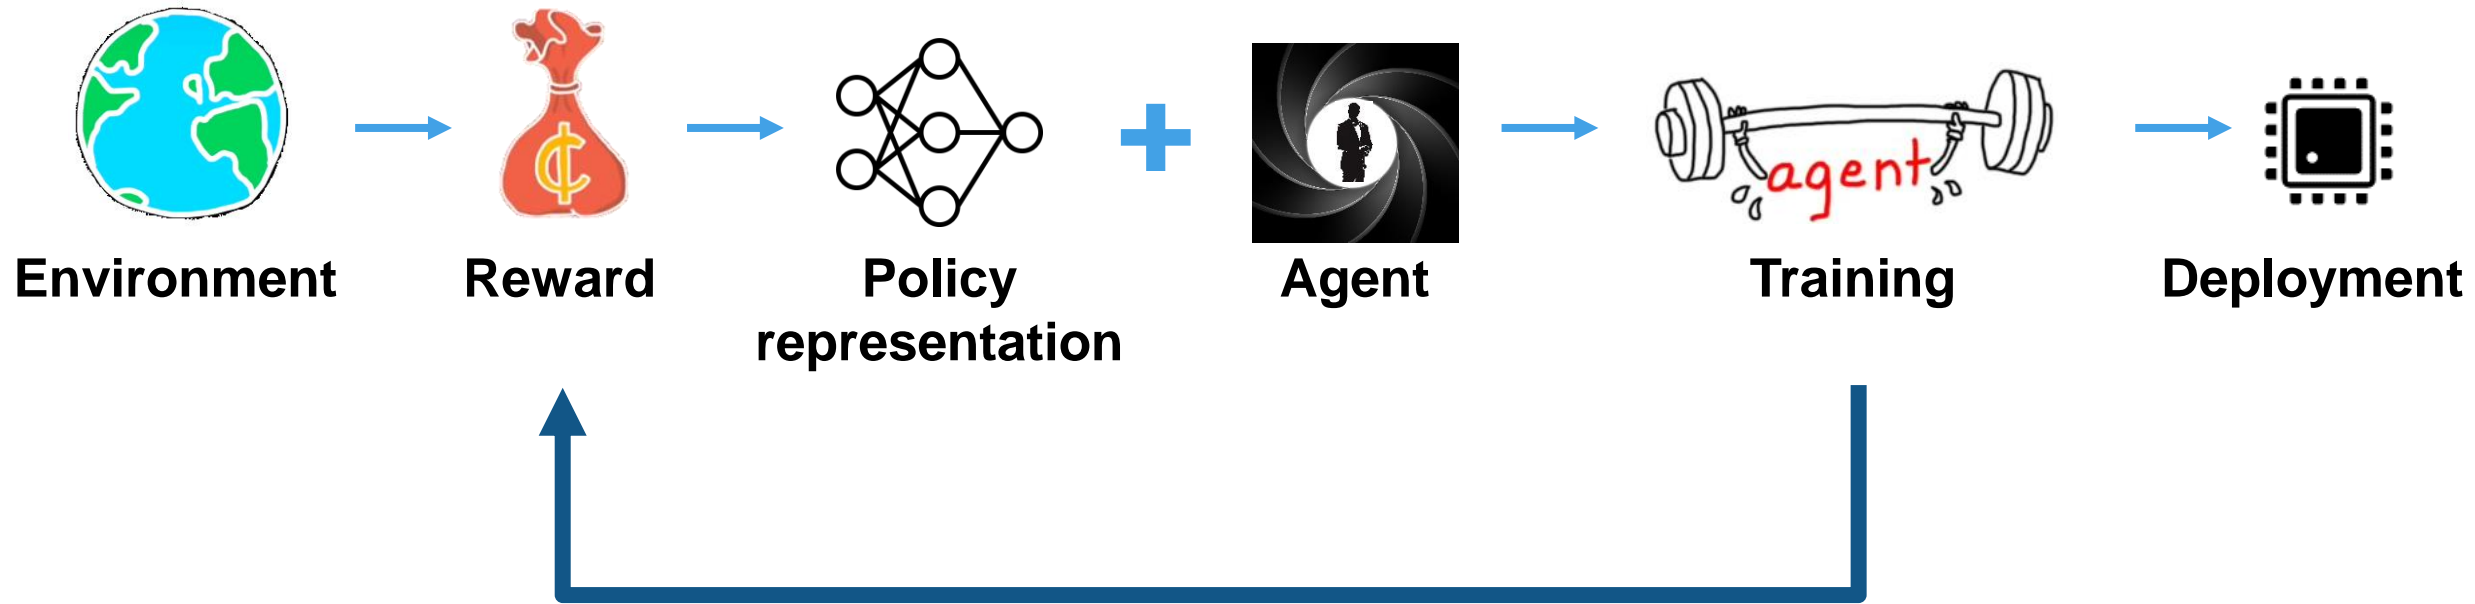
- Select training algorithm
- Tune hyperparameters

- Trained policy is a standalone function

- Numerical value that evaluates goodness of policy
- Reward shaping can be challenging

- Large number of simulations needed
- Parallel & GPU computing can speed up training
- Training could still take hours or days

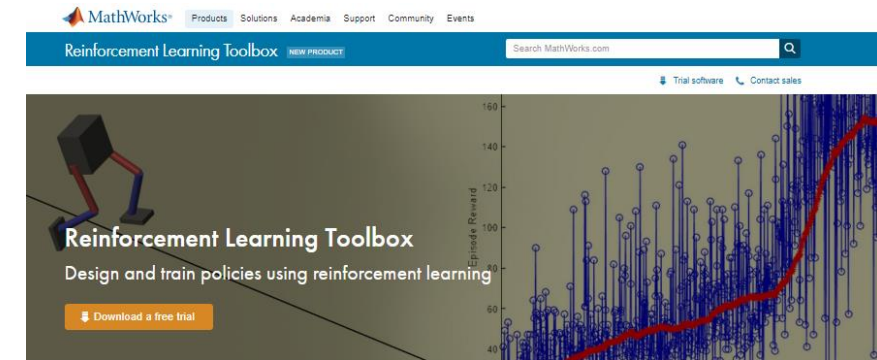
# Reinforcement Learning Workflow



# Reinforcement Learning Toolbox

## Introduced in R2019a

- Built-in and custom reinforcement learning algorithms
- Environment modeling in MATLAB and Simulink
  - Existing scripts and models can be reused
- Deep Learning Toolbox support for representing policies
- Training acceleration with Parallel Computing Toolbox and MATLAB Parallel Server
- Deployment of trained policies with GPU Coder and MATLAB Coder
- Reference examples for getting started

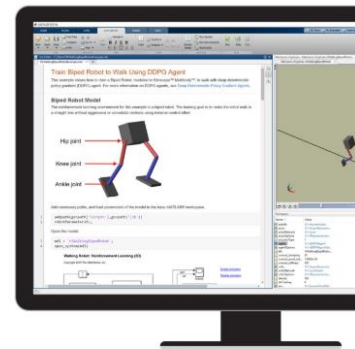


Reinforcement Learning Toolbox™ provides functions and blocks for training policies using reinforcement learning algorithms including DQN, A2C, and DDPG. You can use these policies to implement controllers and decision-making algorithms for complex systems such as robots and autonomous systems. You can implement the policies using deep neural networks, polynomials, or look-up tables.

The toolbox lets you train policies by enabling them to interact with environments represented by MATLAB® or Simulink® models. You can evaluate algorithms, experiment with hyperparameter settings, and monitor training progress. To improve training performance, you can run simulations in parallel on the cloud, computer clusters, and GPUs (with Parallel Computing Toolbox™ and MATLAB Parallel Server™).

Through the ONNX™ model format, existing policies can be imported from deep learning frameworks such as TensorFlow™, Keras and PyTorch (with Deep Learning Toolbox™). You can generate optimized C, C++, and CUDA code to deploy trained policies on microcontrollers and GPUs.

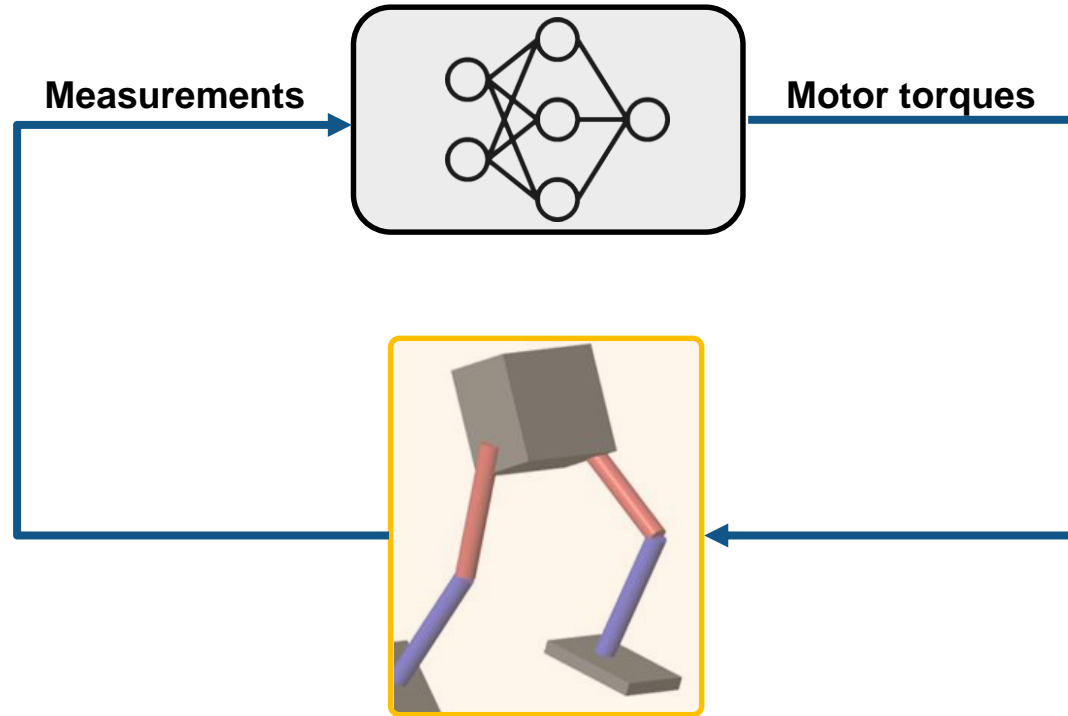
The toolbox includes reference examples for using reinforcement learning to design controllers for robotics and automated driving applications.



# Example: Walking Robot



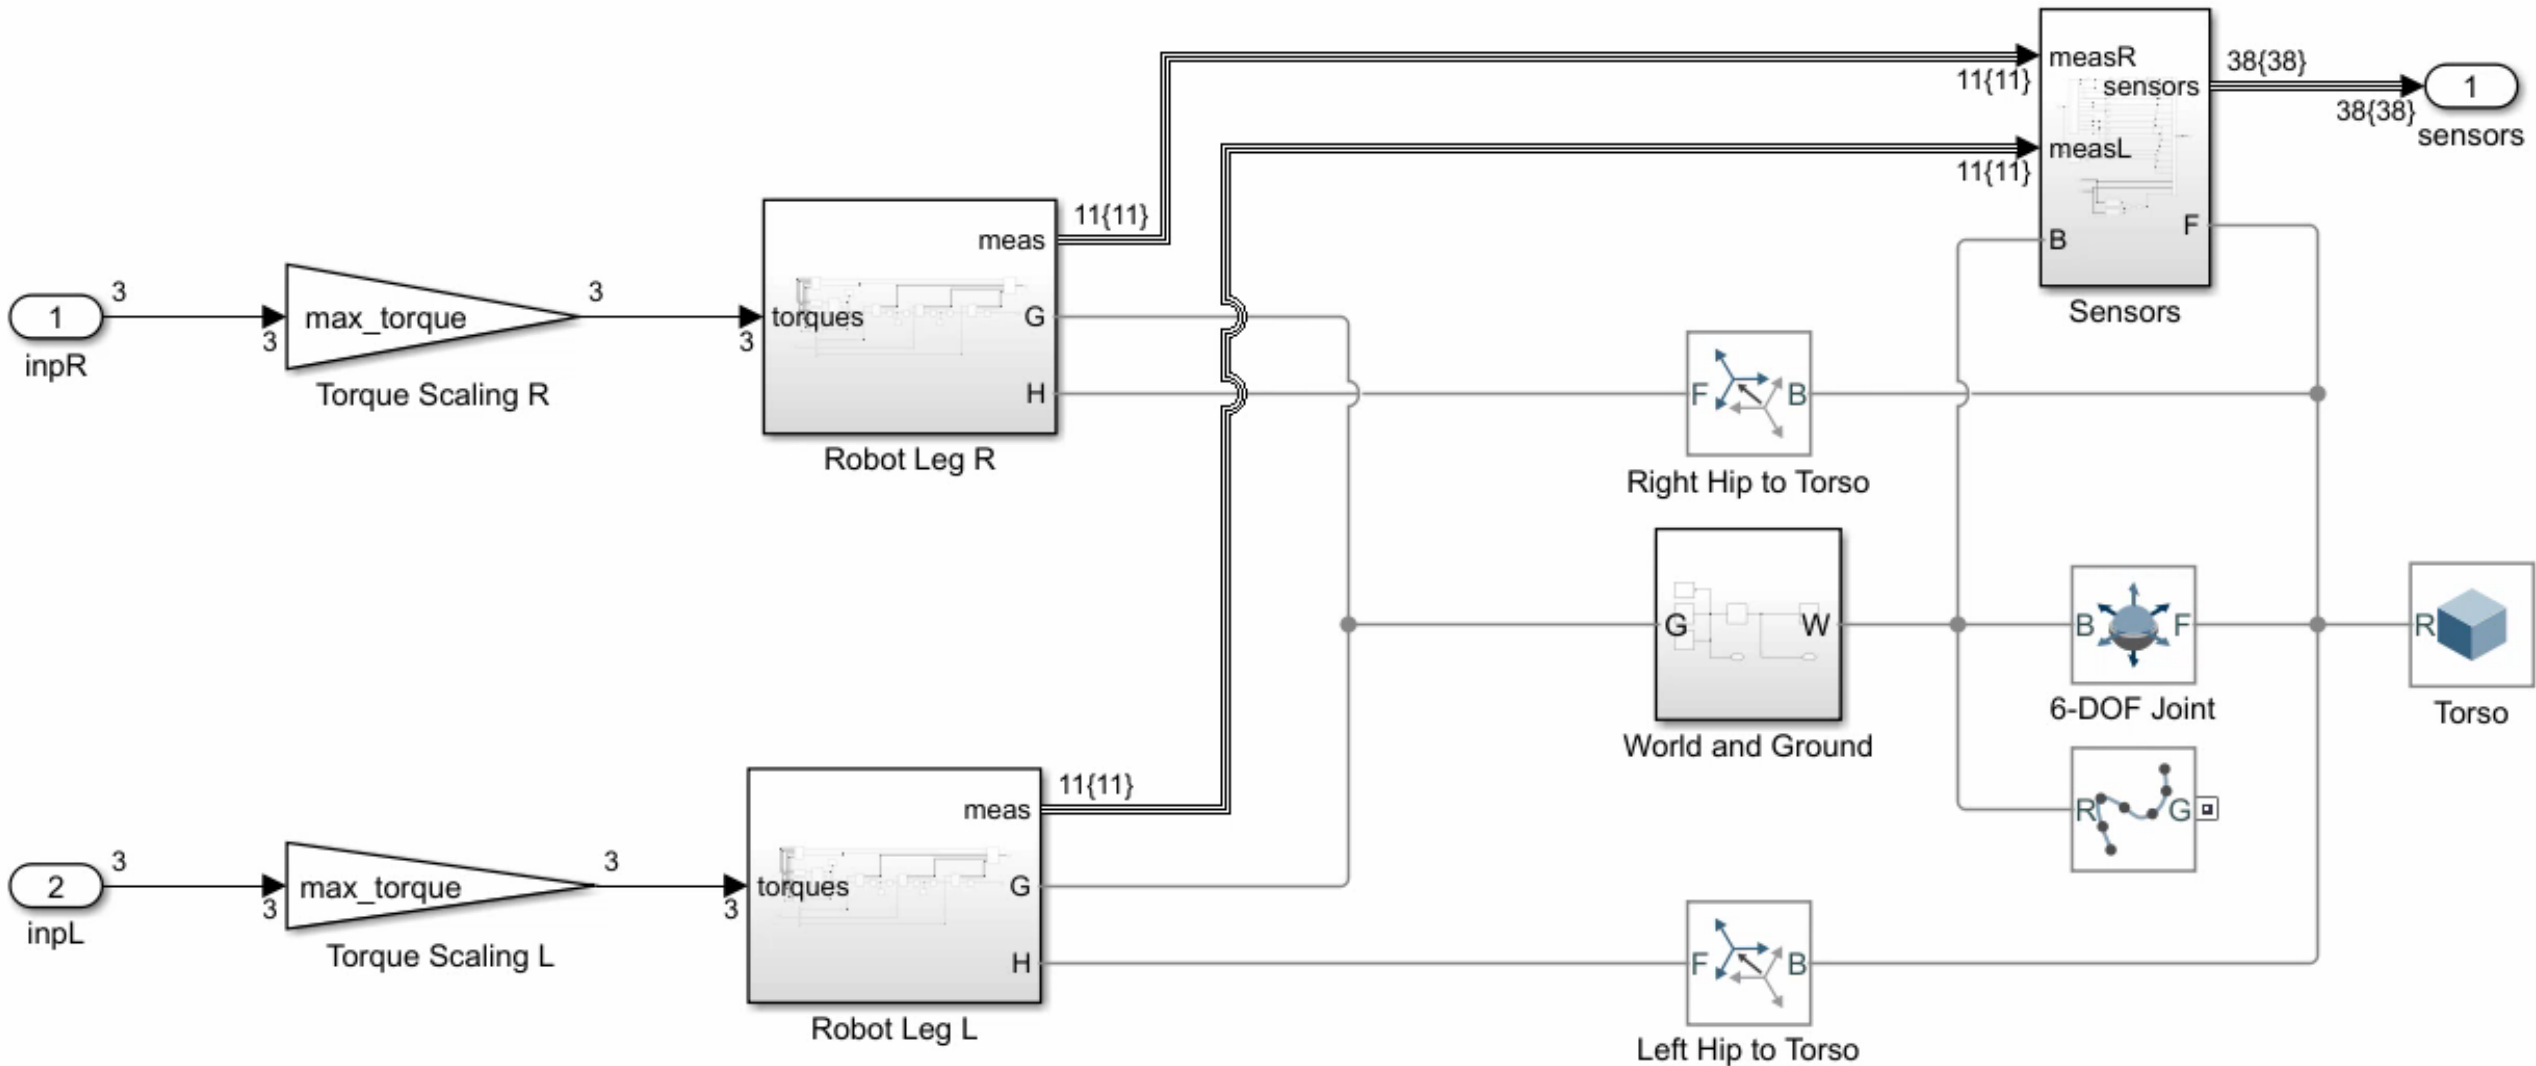
Control objective: Walk on a straight line



# Creating the Environment



rWalkingBipedRobot\_Template ▶ Walking Robot ▶



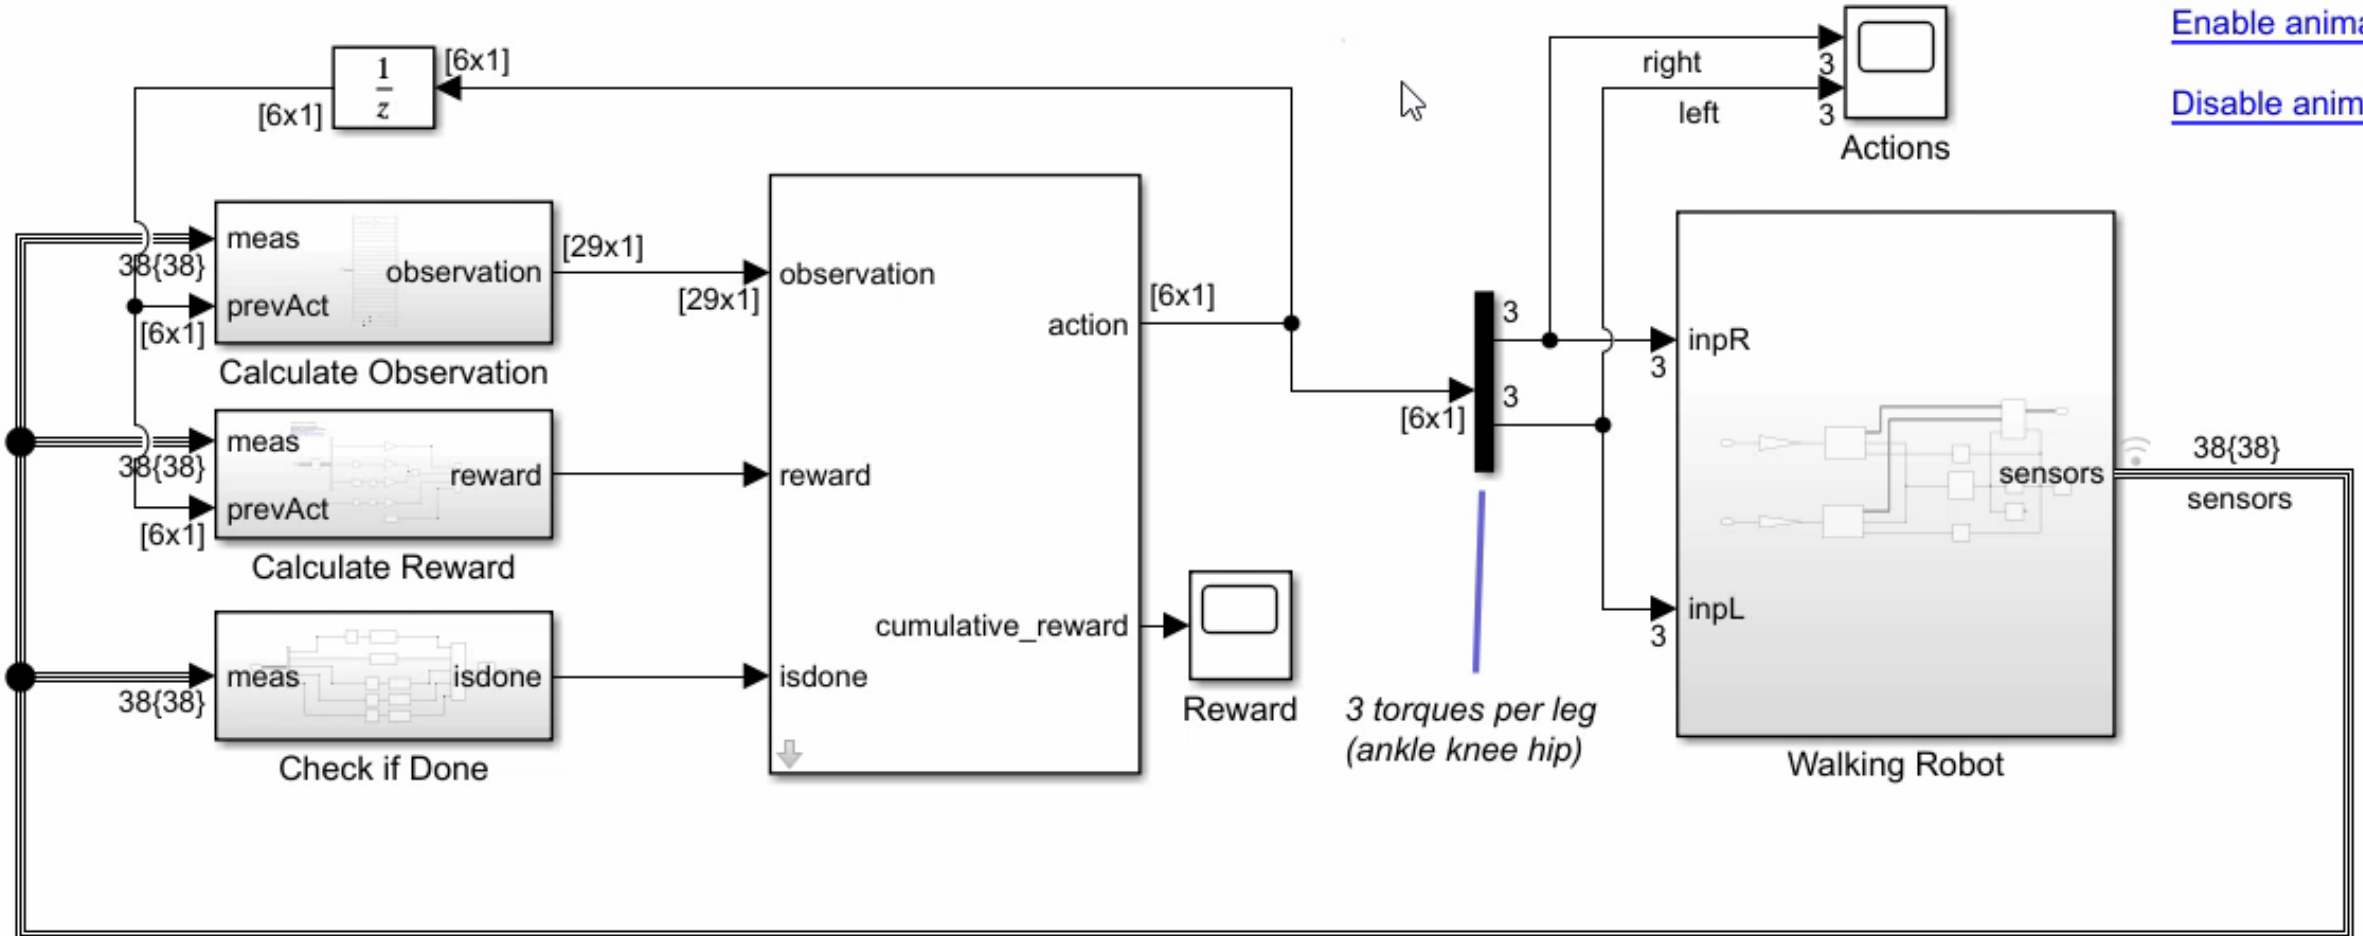


# Reward Shaping



## Walking Robot: Reinforcement Learning (2D)



Copyright 2019 The MathWorks, Inc.



[Enable animation](#)

[Disable animation](#)

## DESIGNER



## LAYER LIBRARY

Filter layers...

## INPUT





- imageInputLayer
- image3dInputLayer
- sequenceInputLayer
- roiInputLayer

## CONVOLUTION AND FULLY CONNECTED

- convolution2dLayer
- convolution3dLayer
- groupedConvolution2dLayer
- transposedConv2dLayer
- transposedConv3dLayer
- fullyConnectedLayer

## SEQUENCE

- lstmLayer
- hilstmLayer



## PROPERTIES

|                       |       |
|-----------------------|-------|
| Number of layers      | 7     |
| Number of connections | 6     |
| Input type            | Image |
| Output type           | None  |

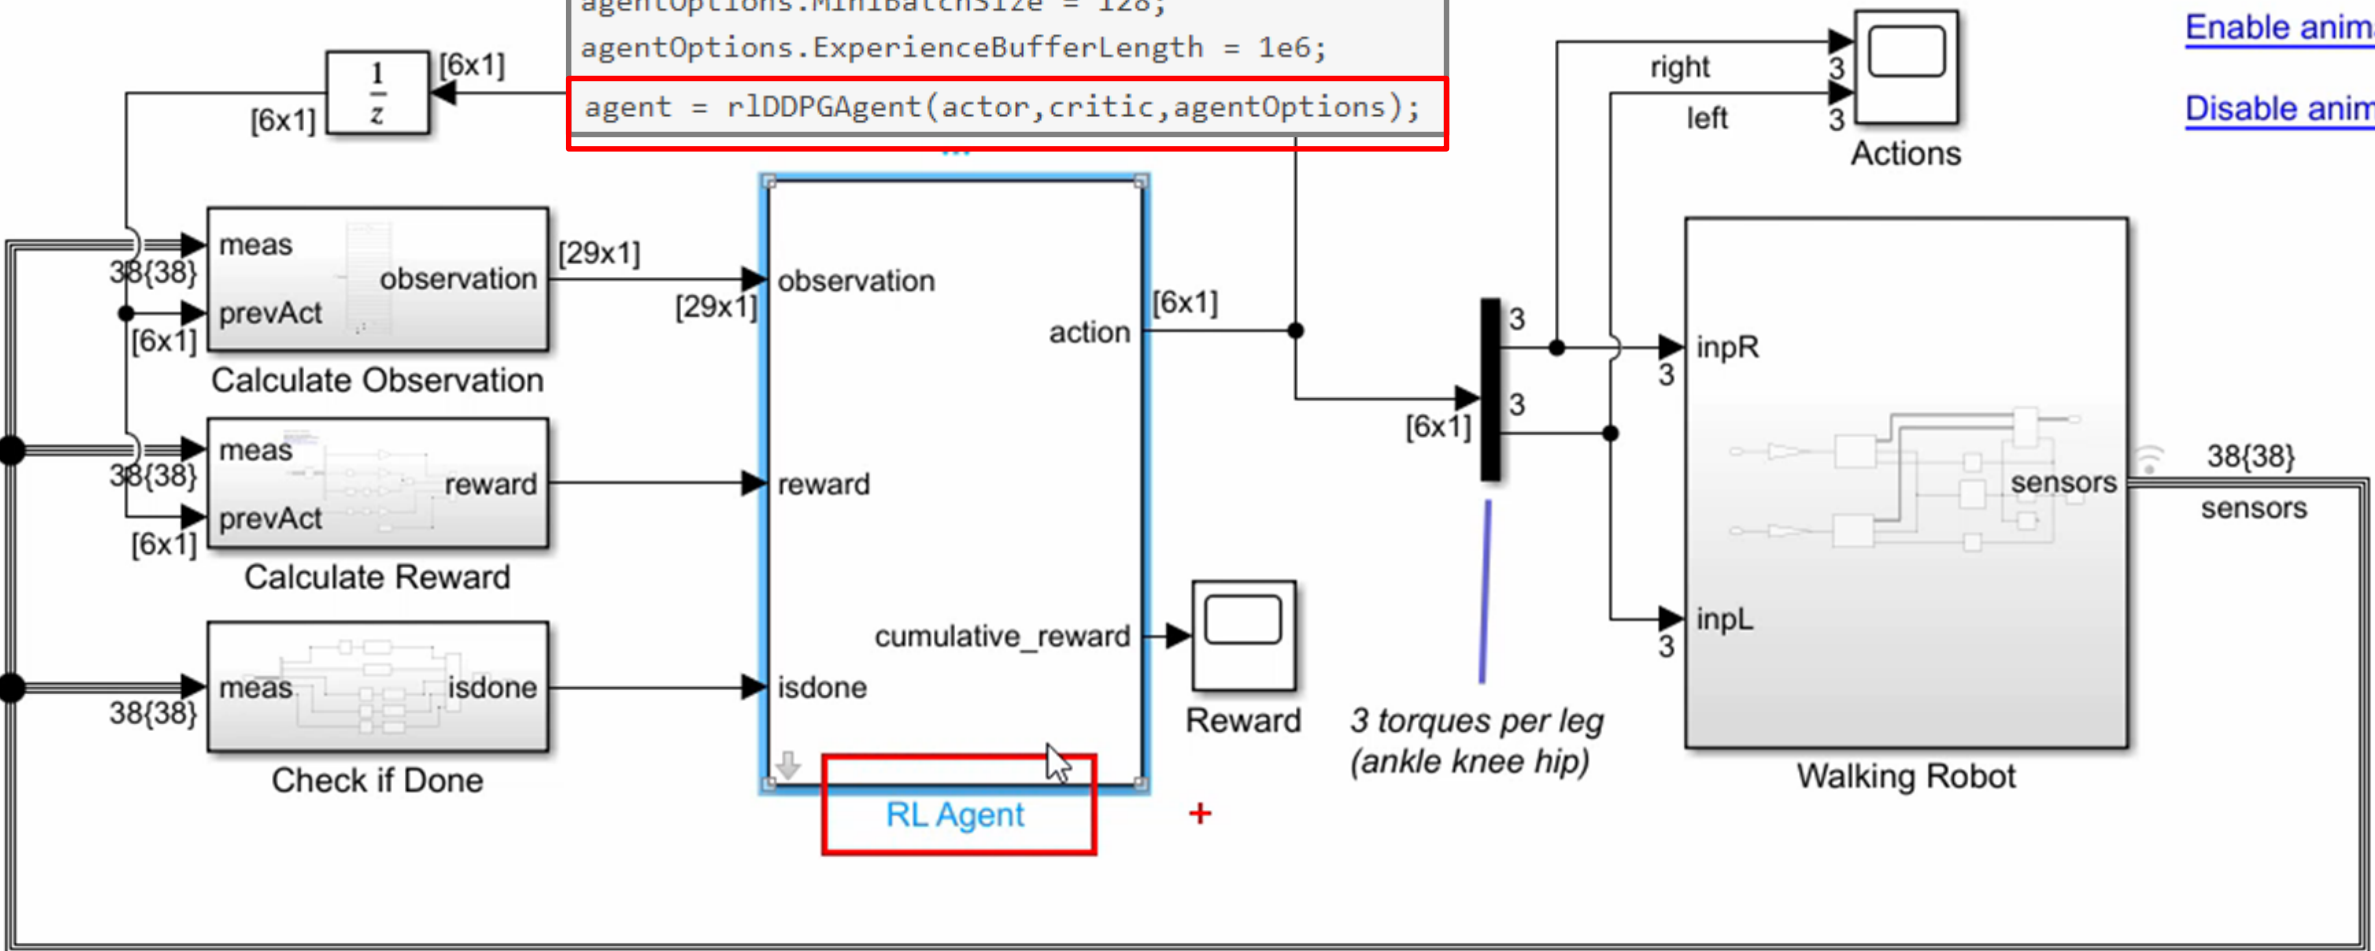
# Creating the Agent



## Walking Robot: Reinforcement Learning

Copyright 2019 The MathWorks

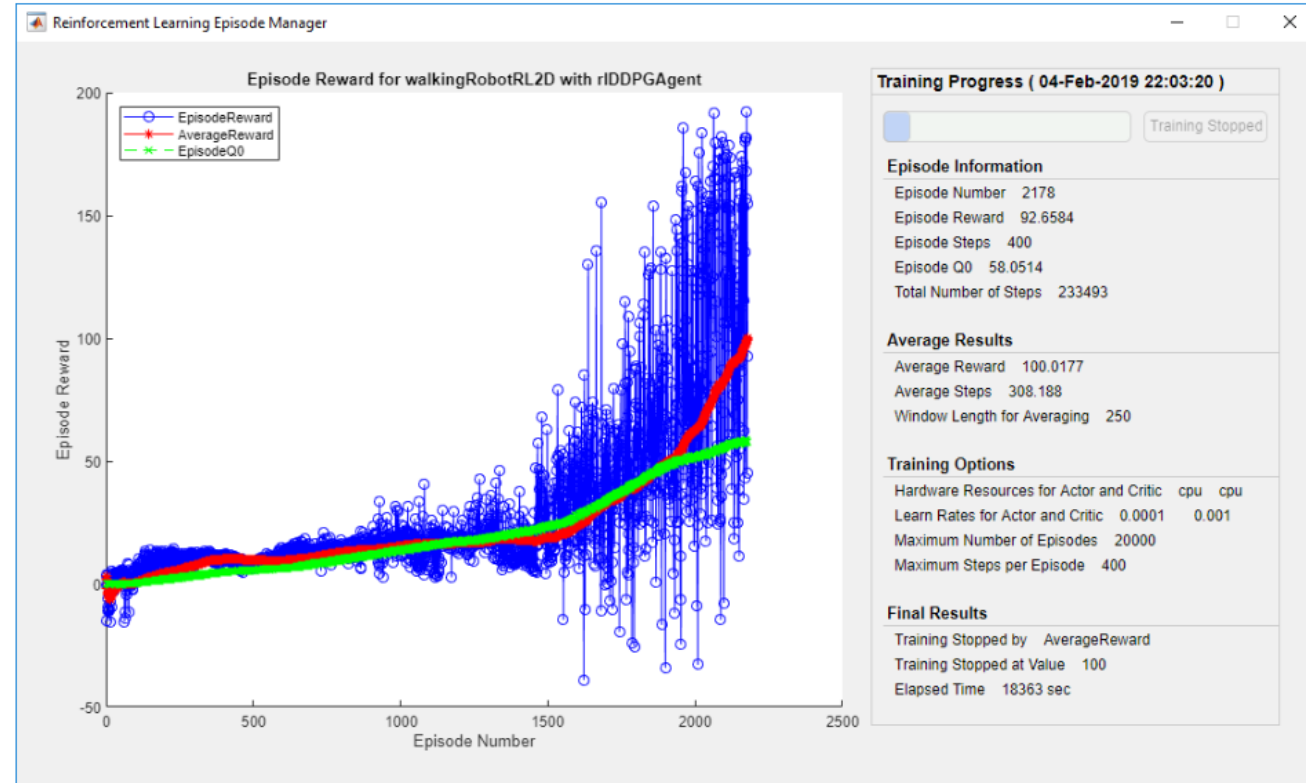
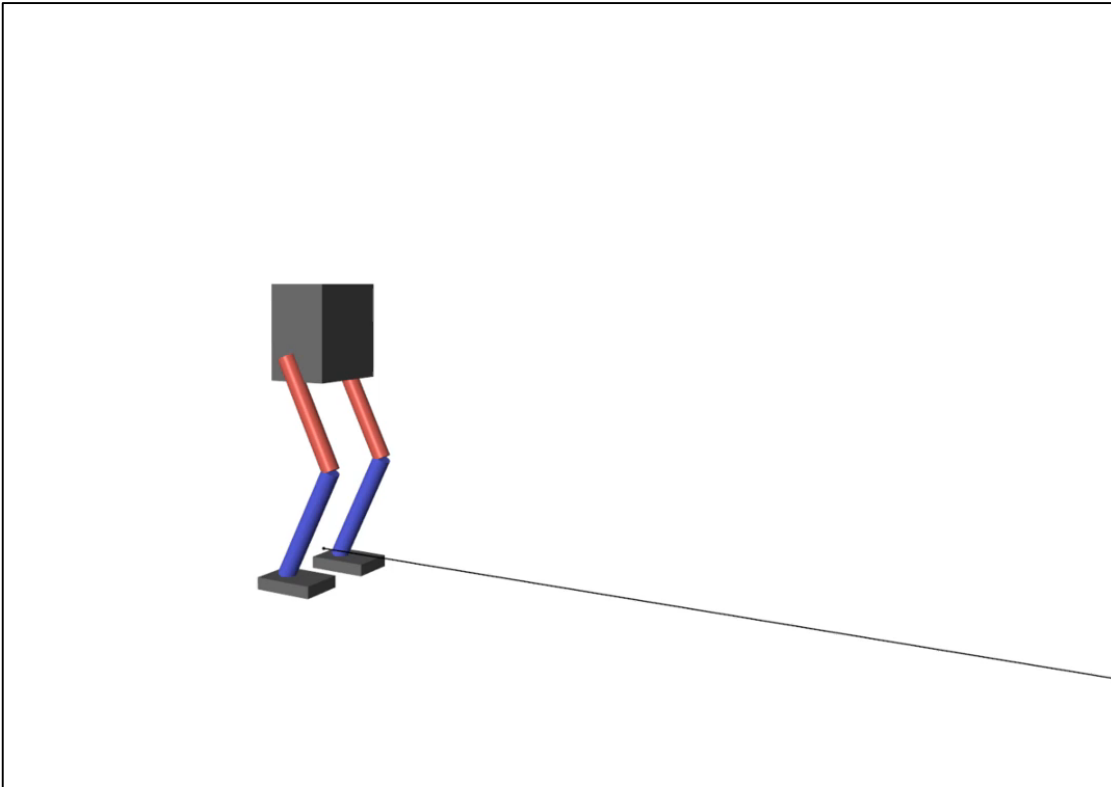
```
agentOptions = rlDDPGAgentOptions;  
agentOptions.SampleTime = Ts;  
agentOptions.DiscountFactor = 0.99;  
agentOptions.MinibatchSize = 128;  
agentOptions.ExperienceBufferLength = 1e6;  
agent = rlDDPGAgent(actor, critic, agentOptions);
```





[Enable animation](#)  
[Disable animation](#)



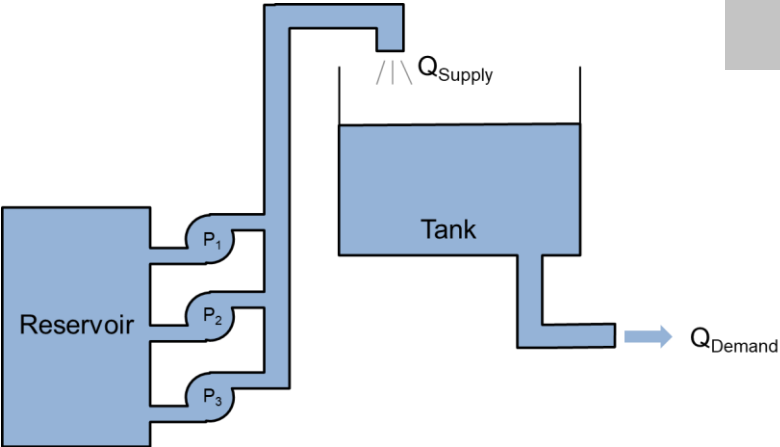


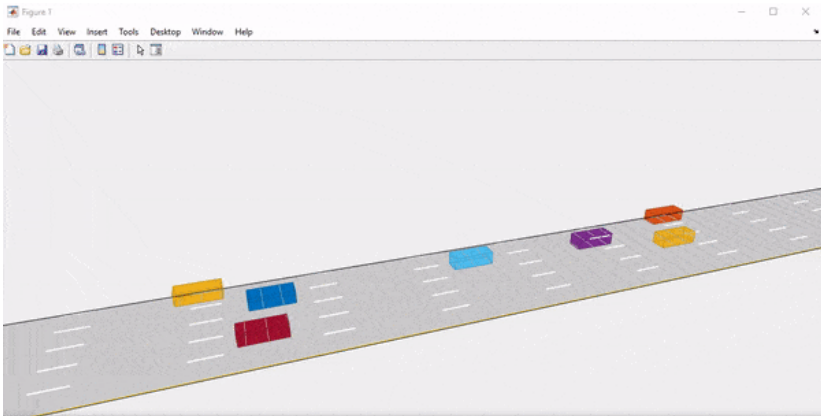
# Training the Agent

```
trainOpts.UseParallel = true;  
trainOpts.ParallelizationOptions.Mode = 'async';
```

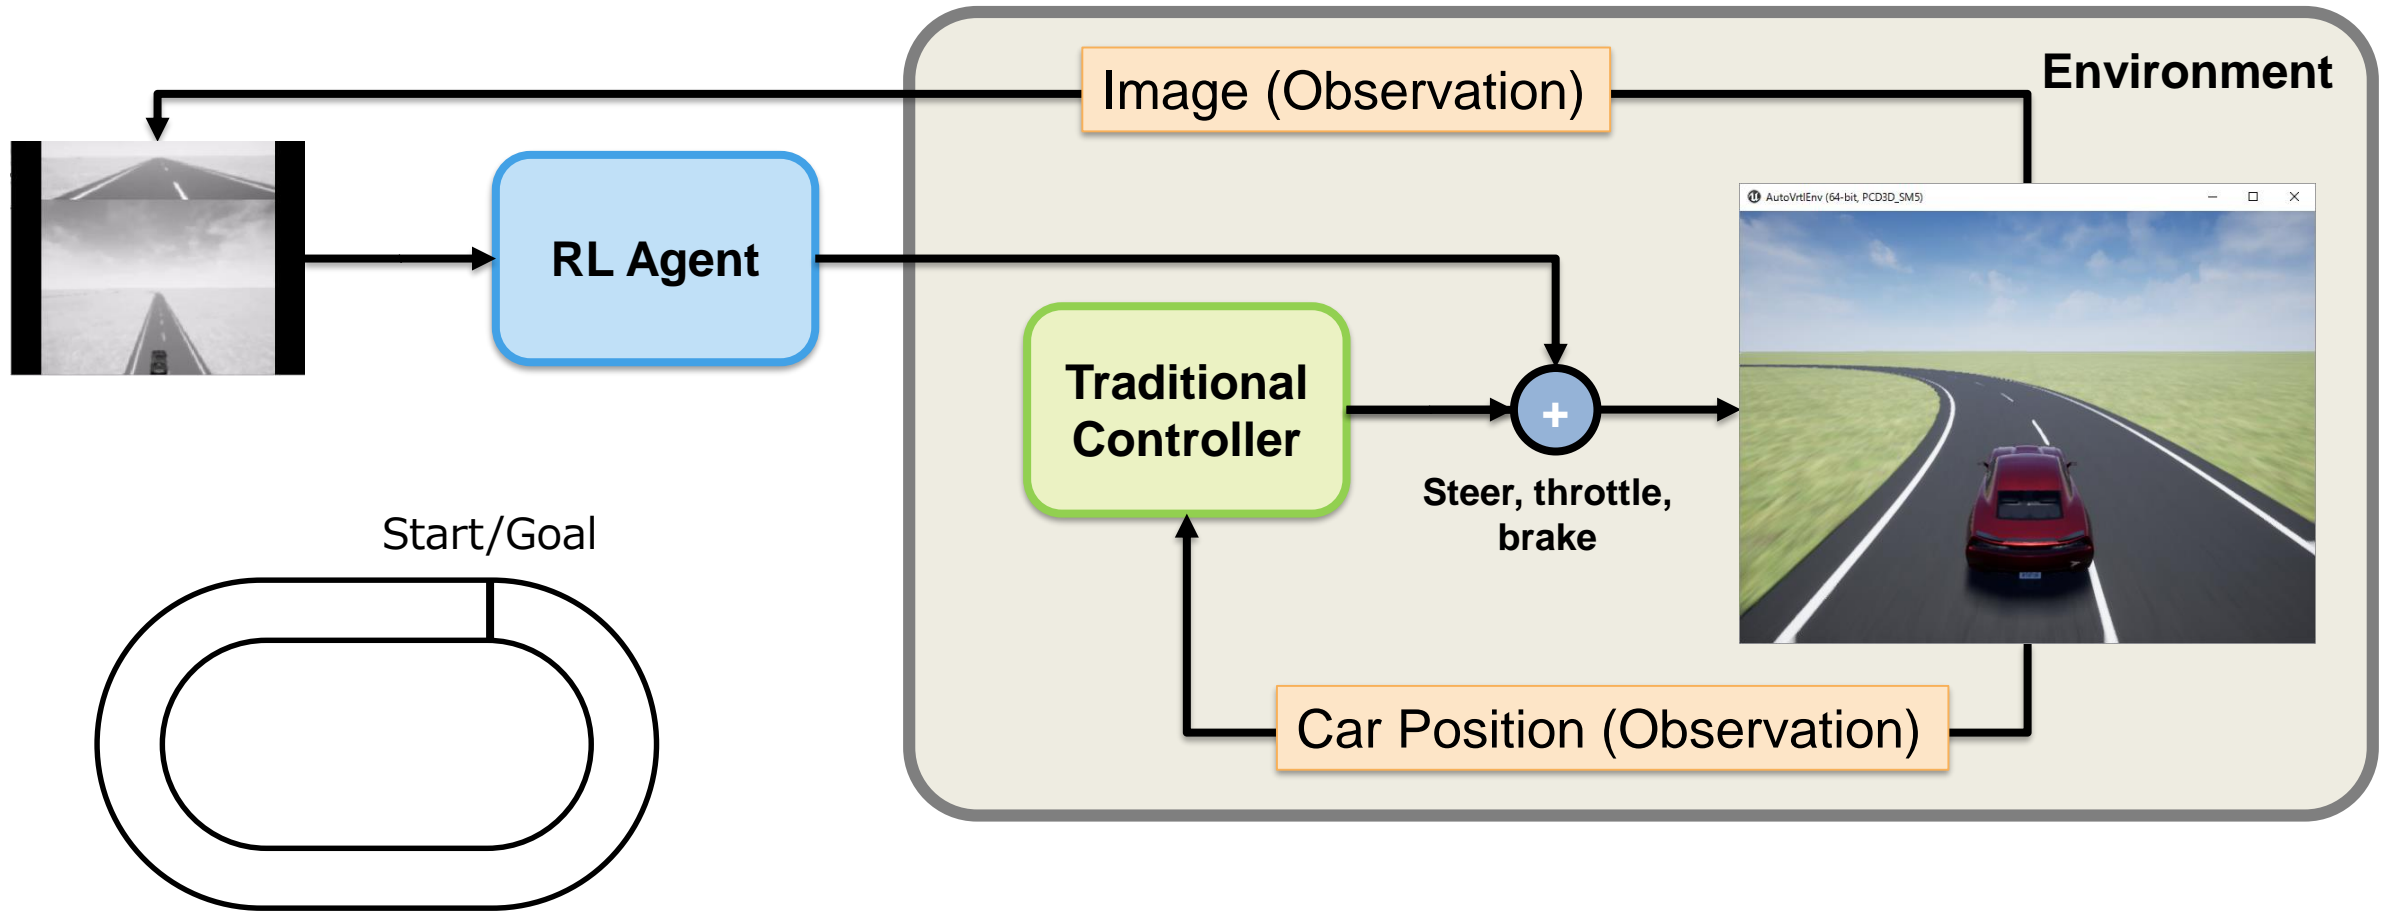




# Applications of Reinforcement Learning

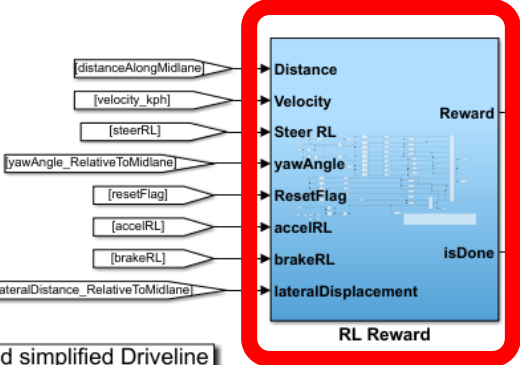
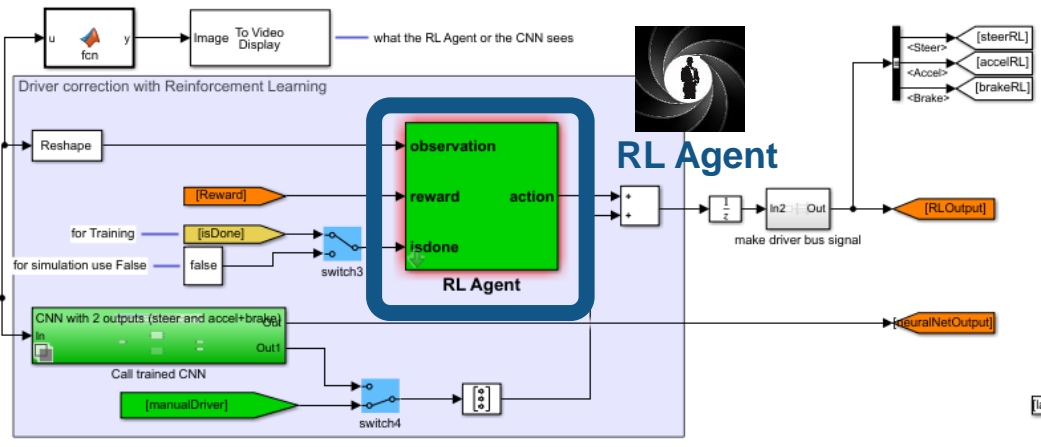





# Autonomous Driving Example



**Objective: Augment traditional controller with reinforcement learning to improve lap time**


# RL Training Environment




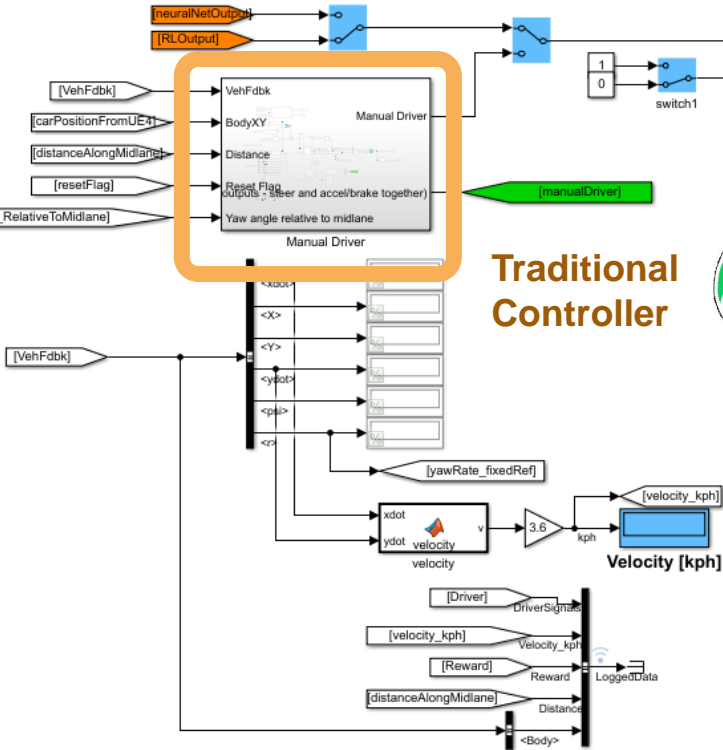
Reward



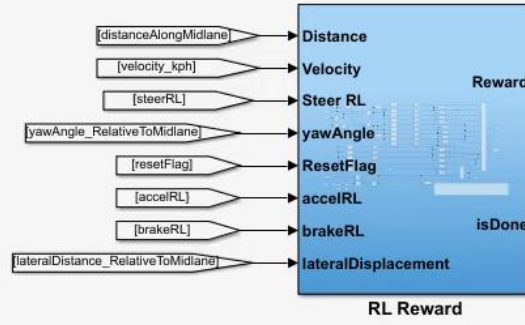



# VDBS Vehicle Model with 14 DOFs Body, Mapped Engine, and simplified Driveline



# Traditional Controller








Reinforcement Learning Episode Manager

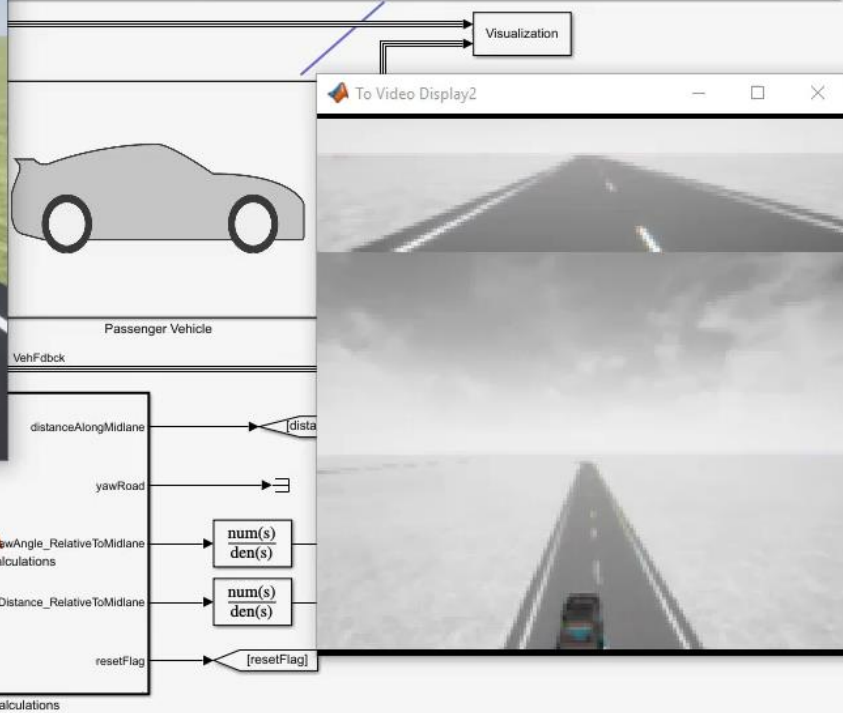
Episode Reward for OvalTrackVDBSR\_L\_9a with rDDPGAgent



**Training Options**  
 Hardware Resources for Actor and Critic cpu cpu  
 Learn Rates for Actor and Critic 0.001 0.0001  
 Maximum Number of Episodes 5000  
 Maximum Steps per Episode 167

**Final Results**  
 Training Stopped by ...  
 Training Stopped at Value ...  
 Elapsed Time ...

OBS Vehicle Model with 14 DOFs Body, Mapped Engine, and simplified Driveline

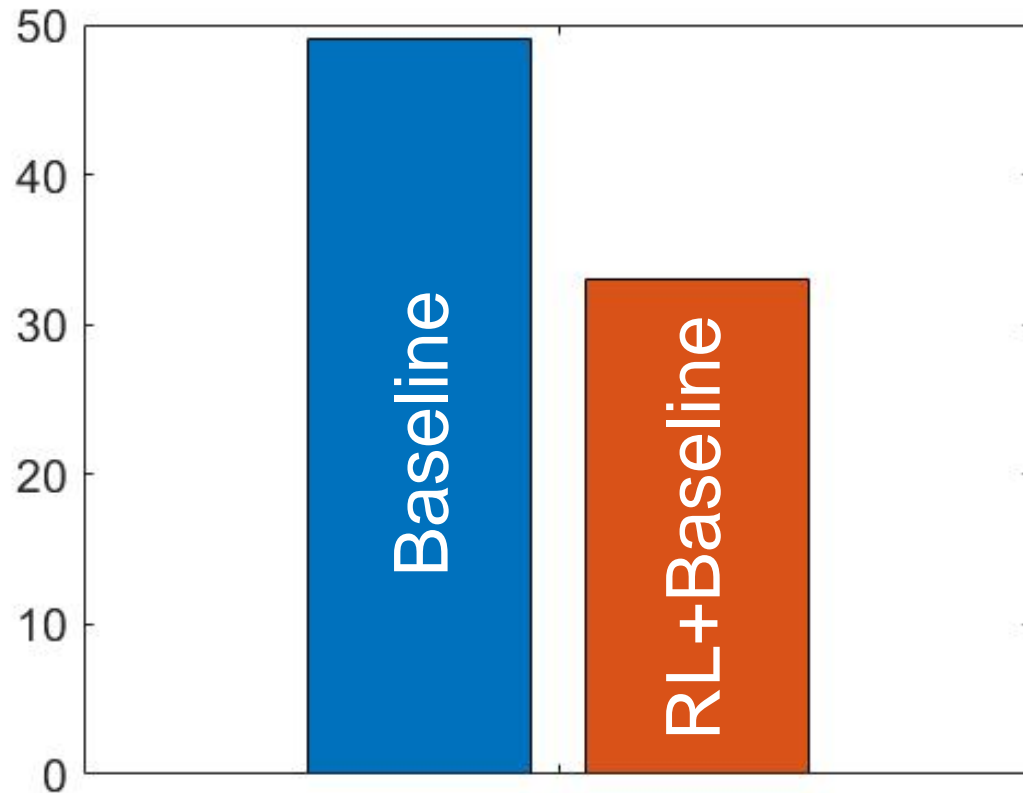


To Video Display2




# Results

Lap time (s)



30% performance improvement

Traditional controller + reinforcement learning



# Reference Applications in Documentation

- Controller Design
- Robotic Locomotion
- Lane Keep Assist
- Adaptive Cruise Control
- Imitation Learning

**Train DDPG Agent to Control Flying Robot**  
Train a reinforcement learning agent to control a flying robot.

**Train Biped Robot to Walk Using DDPG Agent**  
Hip joint  
Knee joint  
Ankle joint

**Train DDPG Agent for Adaptive Cruise Control**  
Ego Car with Adaptive Cruise Control (ACC)

**Train DQN Agent for Lane Keeping Assist**  
Train a reinforcement learning agent for a lane keeping assist application.

**Train DDPG Agent for Path Following Control**  
Train a reinforcement learning agent for a lane following application.

...learning agent control

# Pros & Cons of Reinforcement Learning

## Pros

- **No data** required before training
- **New possibilities** with AI for hard-to-solve problems
- Complex **end-to-end** solutions can be developed
- **Uncertain, nonlinear** environments can be used

## Cons

- Trained policies are **hard to verify** (no performance guarantees)
- Many trials/data points required (**sample inefficient**)
  - Training with real hardware can be expensive and dangerous
- Large number of **design parameters**
  - Reward signal
  - Network architectures
  - Training Hyperparameters

**Simulations** are key in Reinforcement Learning

# How Can MATLAB and Simulink Help?

## Challenges

- Trained policies are **hard to verify** (no performance guarantees)
- Many trials/data points required (**sample inefficient**)
  - Training with real hardware can be expensive and dangerous
- Large number of **design parameters**
  - Reward signal
  - Network architectures
  - Training Hyperparameters

## MATLAB® & SIMULINK®

- **Reuse** existing code and models for environments
- Use simulations for **policy verification**
  - Simulate extreme scenarios
- Run simulation trials **in parallel** to accelerate training
- Consult Reinforcement Learning Toolbox **examples**
  - Iterative tuning with simulations

# Key Takeaways

- What is reinforcement learning and why should I care about it?
- How do I set up and solve a reinforcement learning problem?
- What are some common challenges?

# Learn More

- Reference examples for controls, robotics, and autonomous system applications
- Documentation written for engineers and domain experts
- Tech Talk video series on Reinforcement Learning concepts
- Reinforcement Learning ebooks available at mathworks.com

**Train DDPG Agent to Control Flying Robot**  
Train a reinforcement learning agent to control a flying robot model.

**Train Biped Robot to Walk Using DDPG Agent**  
Train a reinforcement learning agent to control a biped walking robot model.

**Train DDPG Agent for Adaptive Cruise Control**  
Train a reinforcement learning agent to control a vehicle in an adaptive cruise control application.

**Train DDPG Agent for Lane Following Control**  
Train a reinforcement learning agent to control a vehicle in a lane following application.

**Reinforcement Learning Toolbox**  
Design and train policies using reinforcement learning.

Reinforcement Learning Toolbox™ provides functions and blocks for training policies using reinforcement learning algorithms including DQN, A2C, and DDPG. You can use these policies to implement controllers and decision-making algorithms for complex systems such as robots and autonomous systems. You can implement the policies using deep neural networks, polynomials, or look-up tables.

The toolbox lets you train policies by enabling them to interact with environments represented by MATLAB® or Simulink® models. You can evaluate algorithms, experiment with hyperparameter settings, and monitor training progress. To improve training performance, you can run simulations in parallel on the cloud, computer clusters, and GPUs (with Parallel Computing Toolbox™ and MATLAB Parallel Server™).

Through the ONNX™ model format, existing policies can be imported from deep learning frameworks such as TensorFlow™, Keras and PyTorch (with Deep Learning Toolbox™). You can generate optimized C, C++, and CUDA code to deploy trained policies on microcontrollers and GPUs.

The toolbox includes reference examples for using reinforcement learning to design controllers for robotics and automated driving applications.

**Getting Started**  
Learn the basics of Reinforcement Learning Toolbox

**MATLAB Environments**  
Model reinforcement learning environment dynamics using MATLAB

**Simulink Environments**  
Model reinforcement learning environment dynamics using Simulink models

**Policies and Value Functions**  
Define policy and value-function representations, such as deep neural networks and Q tables

**Agents**  
Create and configure reinforcement learning agents using common algorithms, such as SARSA

**Training and Validation**  
Train and simulate reinforcement learning agents

**Policy Deployment**  
Code generation and deployment of trained policies

**Agent-Environment Interaction Diagram:**  
The diagram shows an Agent (containing Policy and Reinforcement learning algorithm) interacting with an Environment (robot arm). The Agent sends actuator commands (Action) to the Environment. The Environment returns a Reward (+50) and Observation (state) to the Agent. The Observation (state) is used for Policy updates.