


# Reinforcement Learning: Leveraging Deep Learning for Controls

Christoph Stockhammer

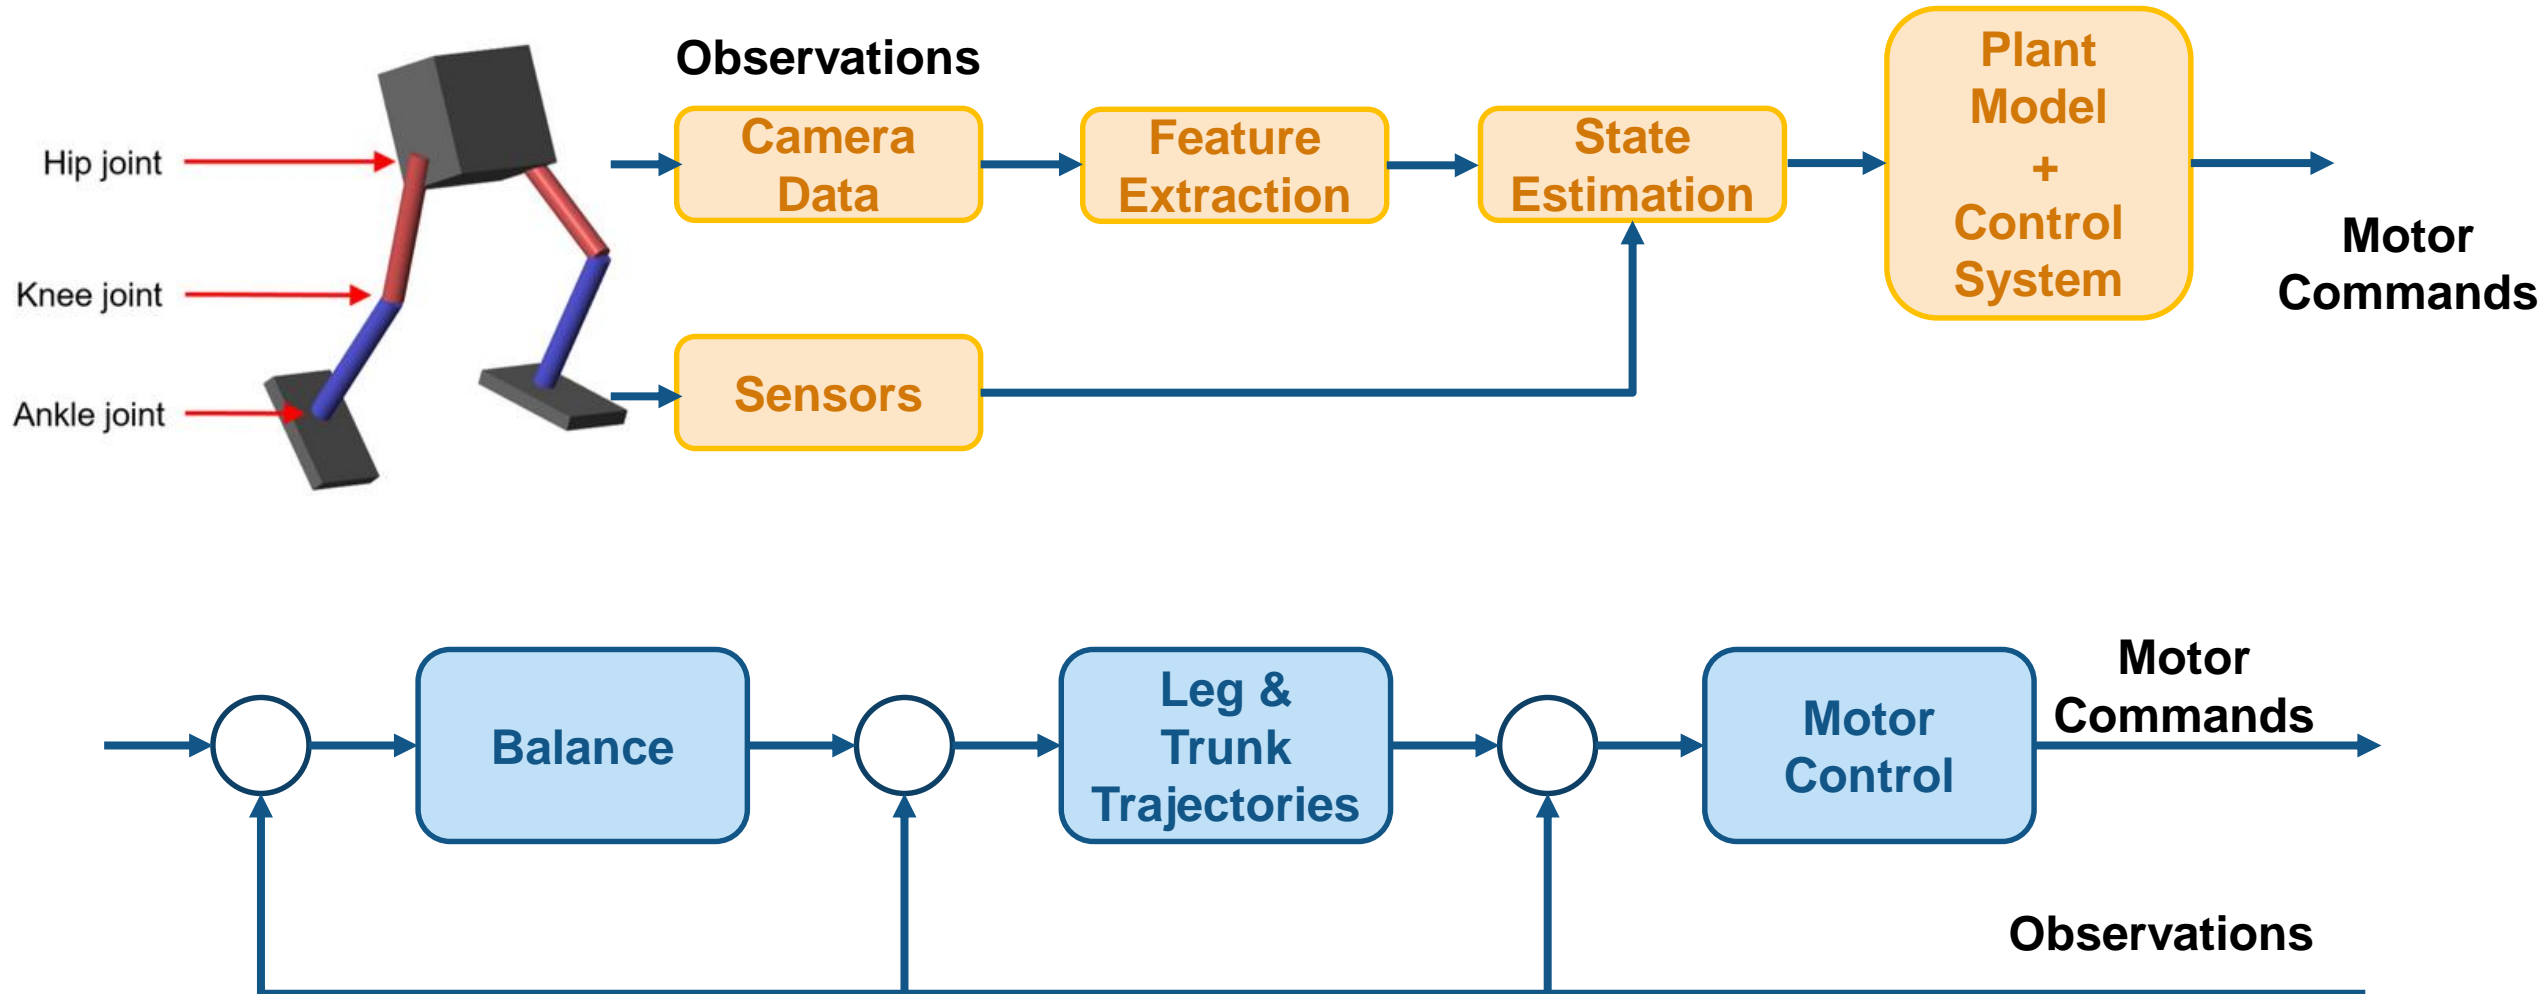
MathWorks Application Engineering



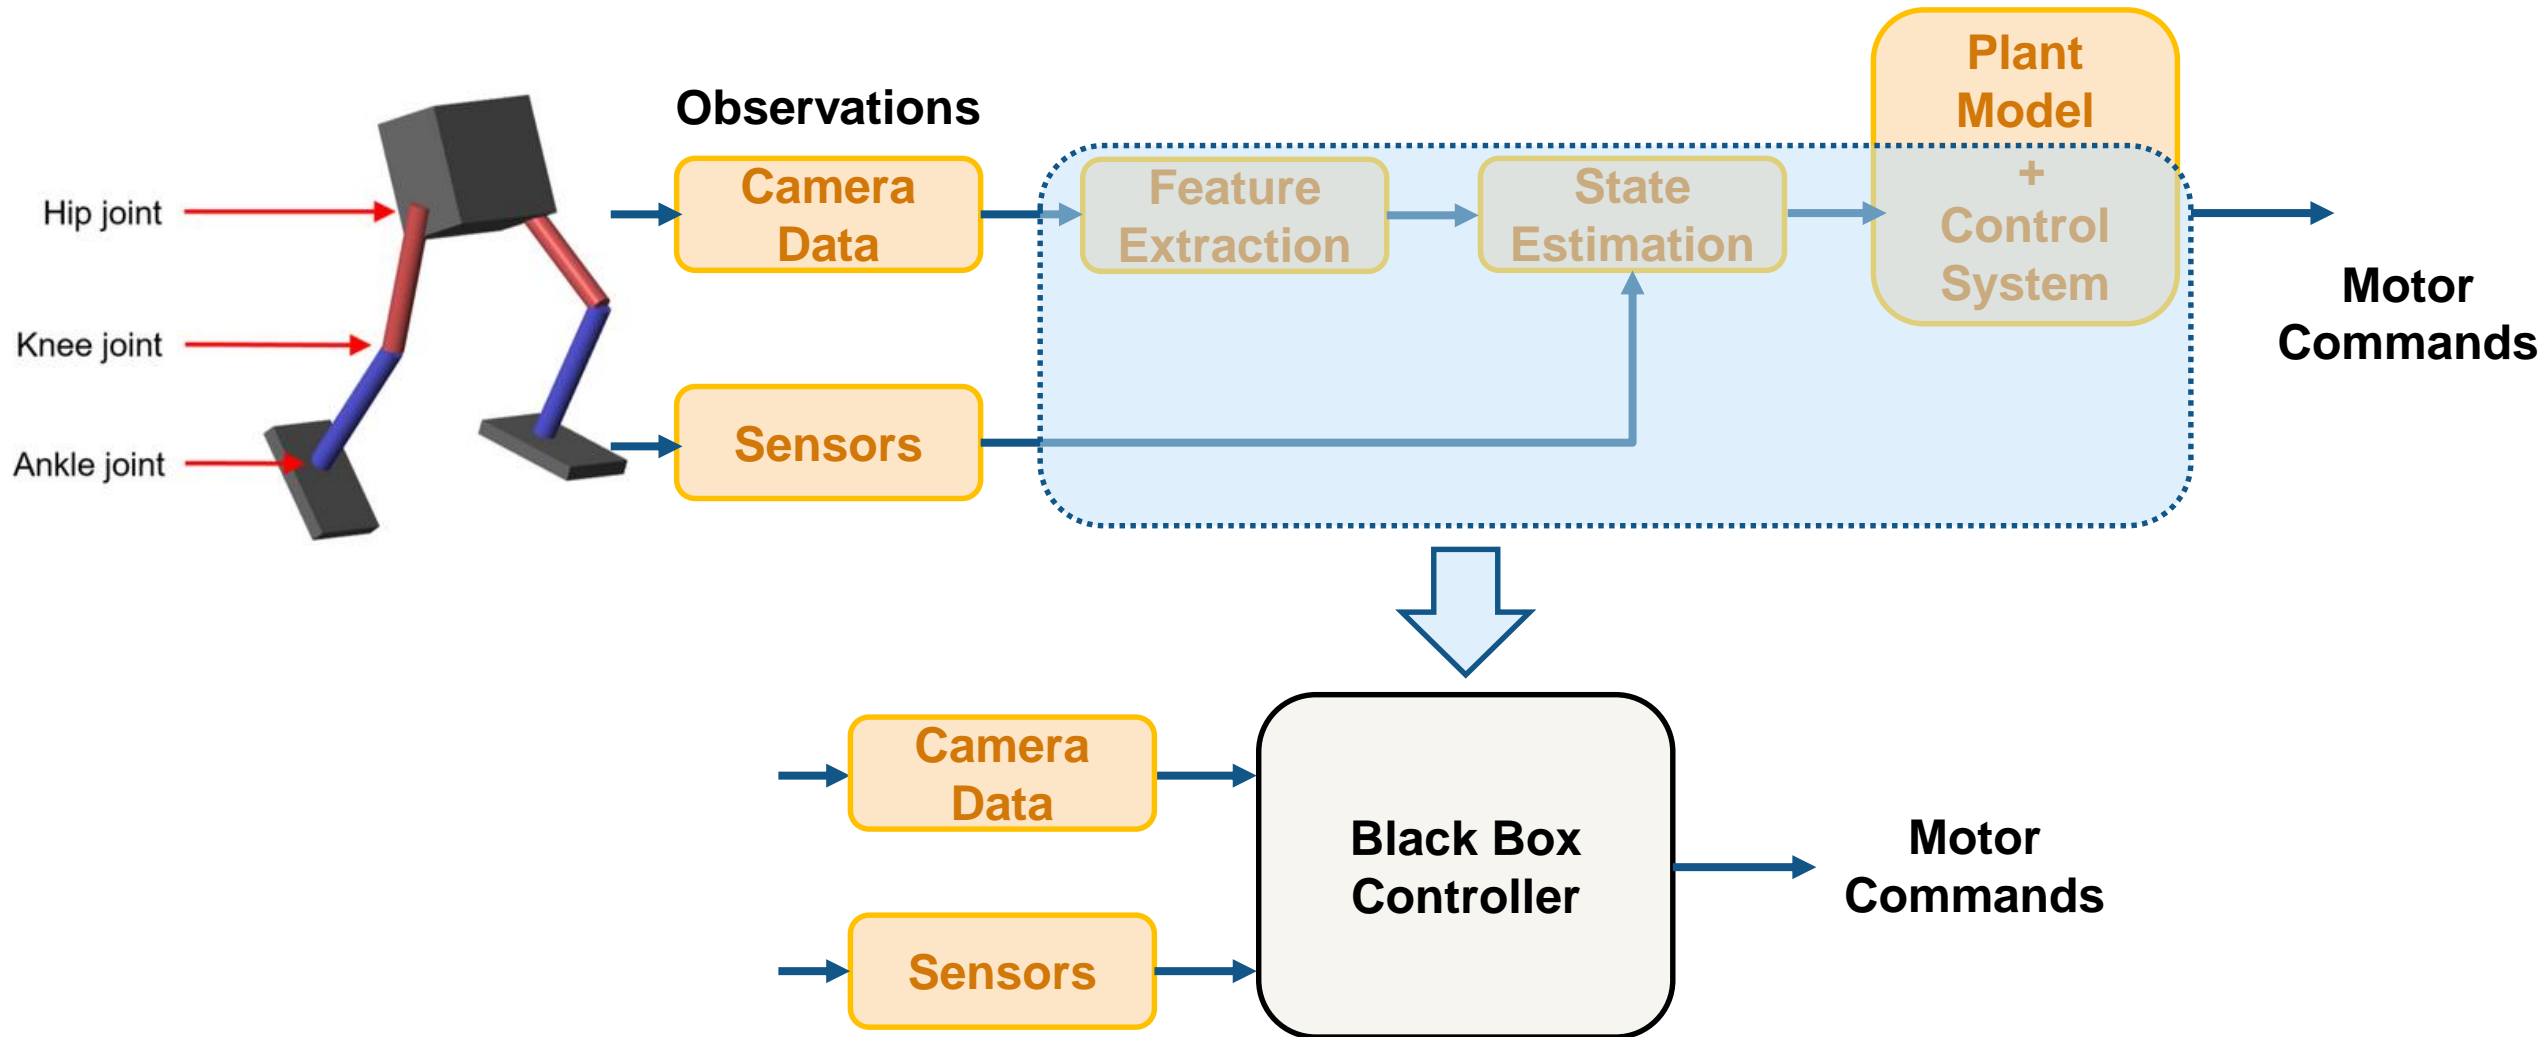
Goal: We hope you walk away knowing the answer to these questions

- What is reinforcement learning and why should I care about it?
- How do I set it up and solve it? [*from an engineer's perspective*]
- What are some benefits and drawbacks?


# Let's try to solve this problem the traditional way




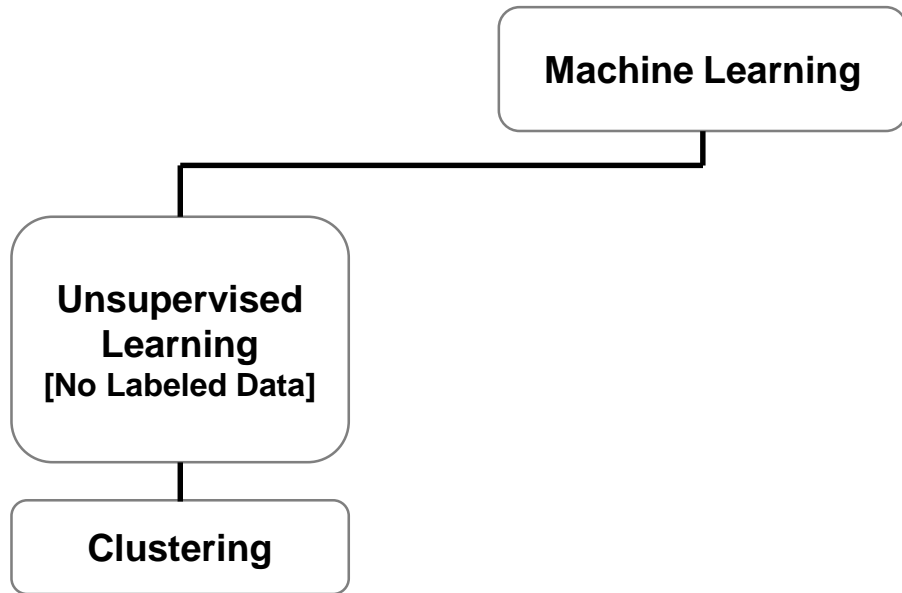
# What is the alternative approach?




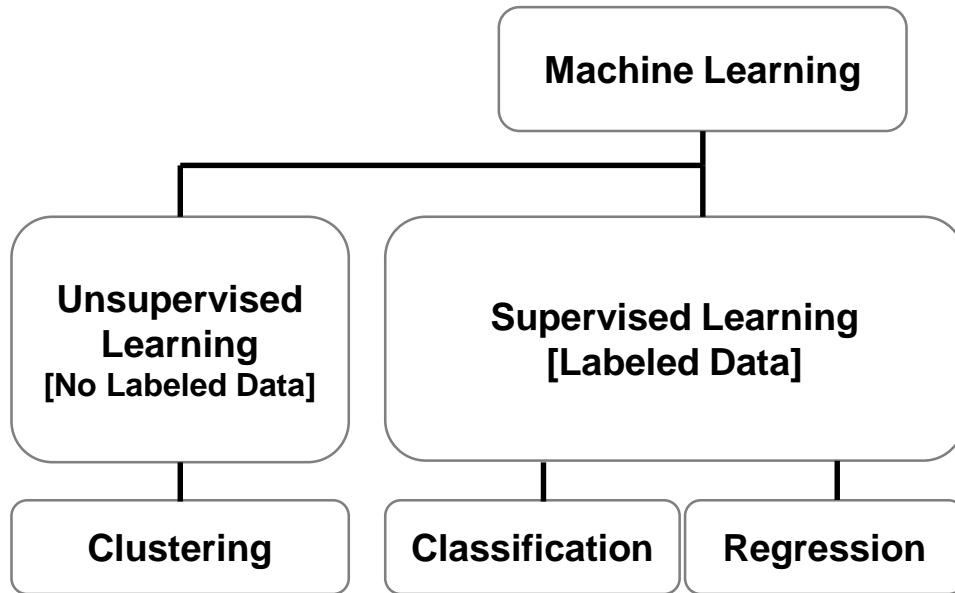
# What is reinforcement learning?



# Reinforcement Learning vs Machine Learning vs Deep Learning

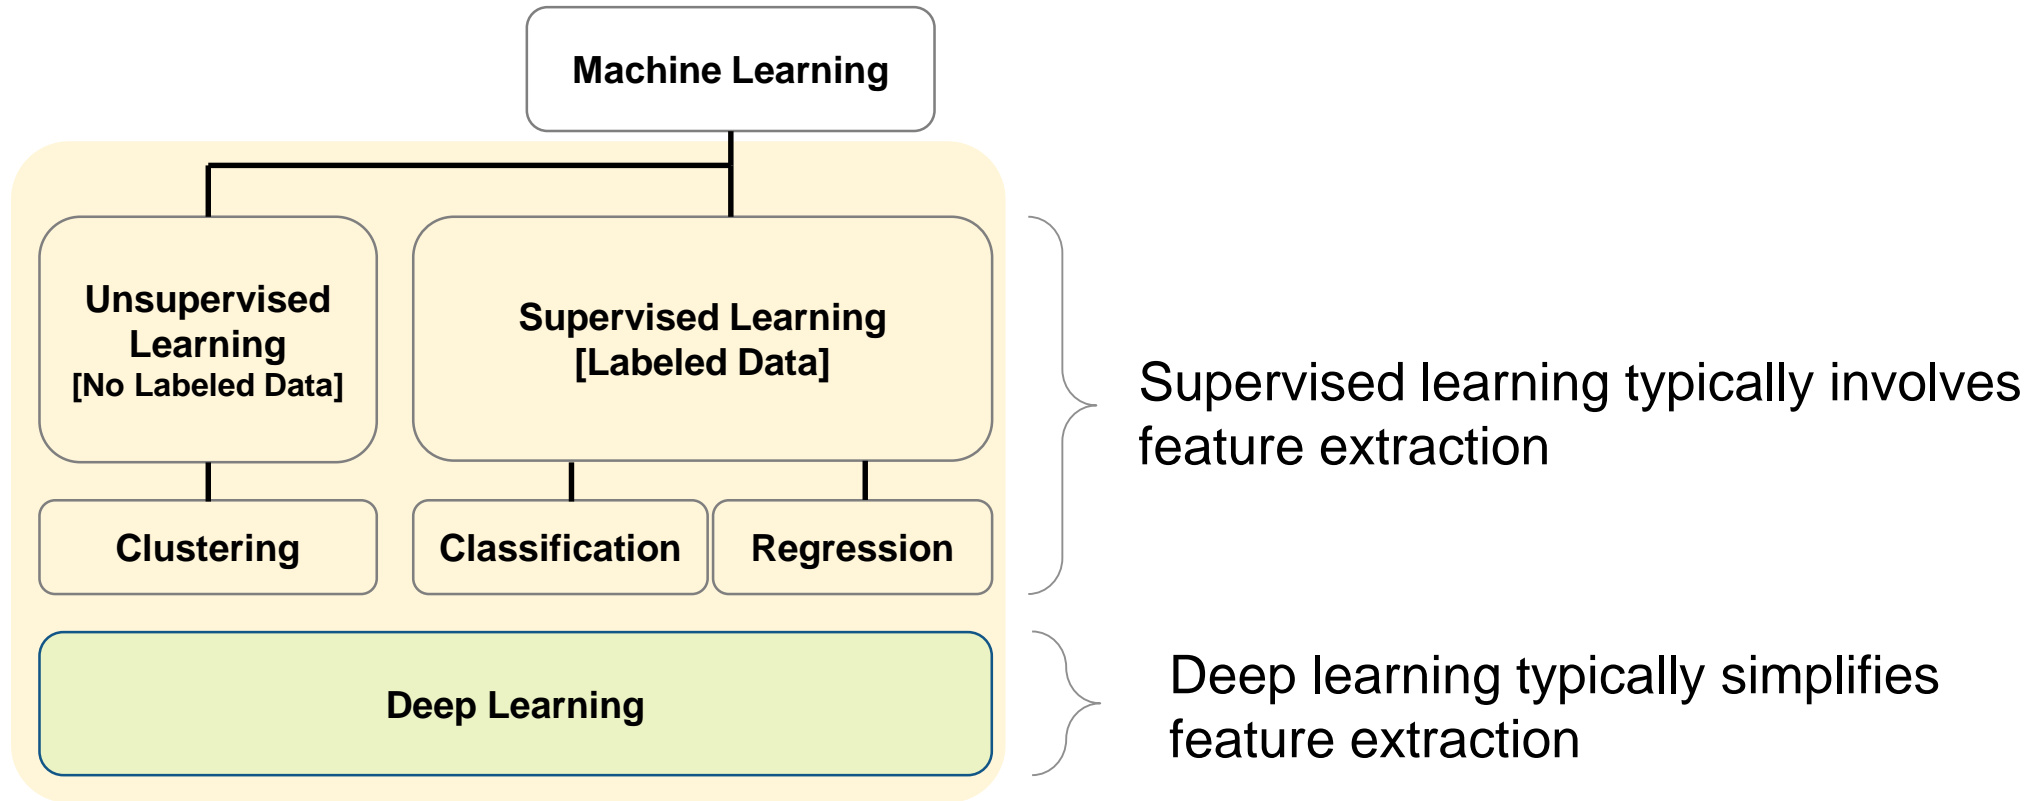


# Reinforcement Learning vs Machine Learning vs Deep Learning

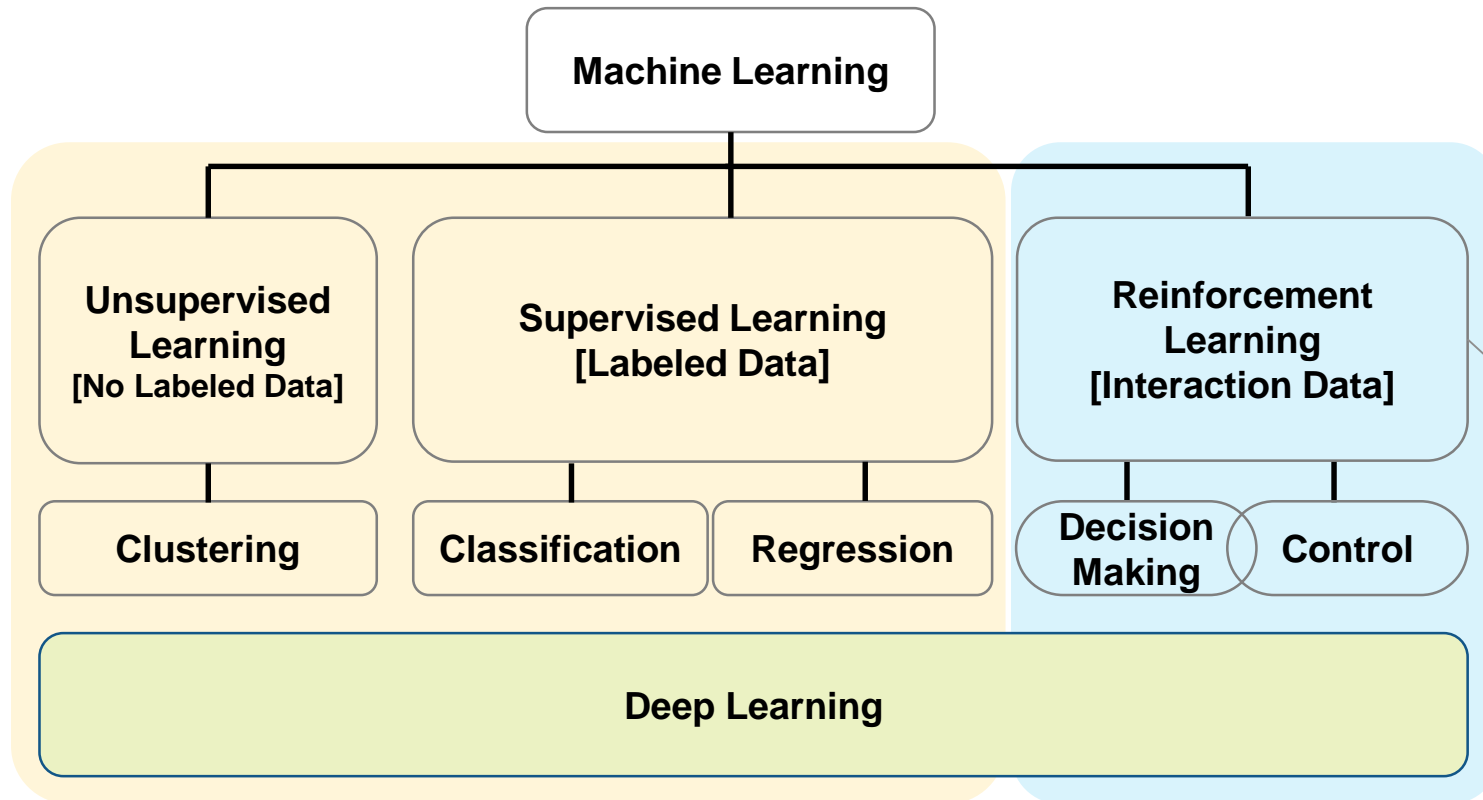




# Reinforcement Learning vs Machine Learning vs Deep Learning



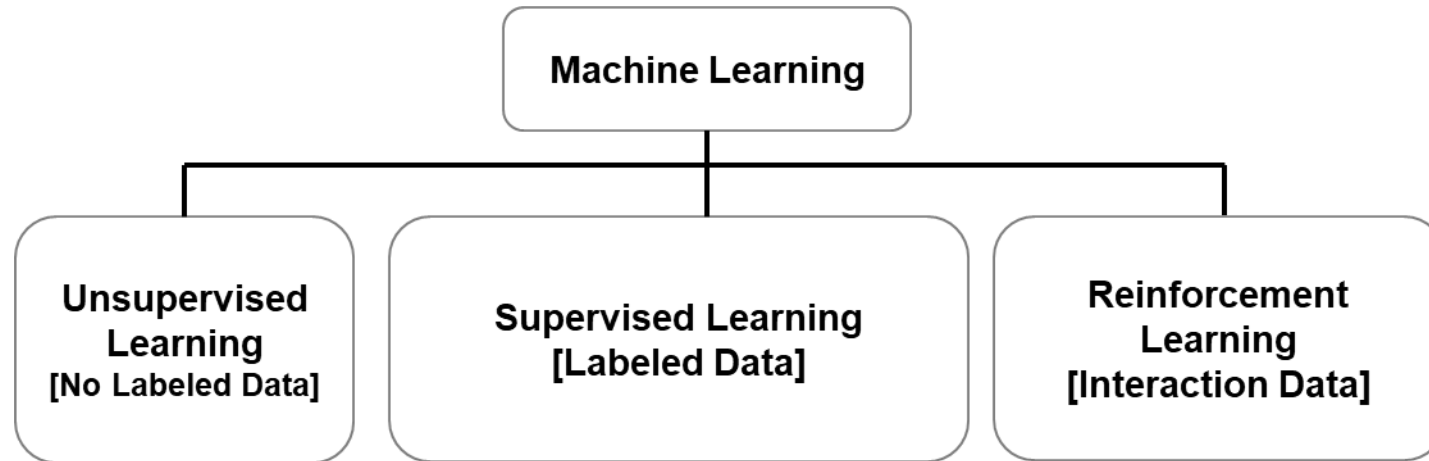
# Reinforcement Learning vs Machine Learning vs Deep Learning





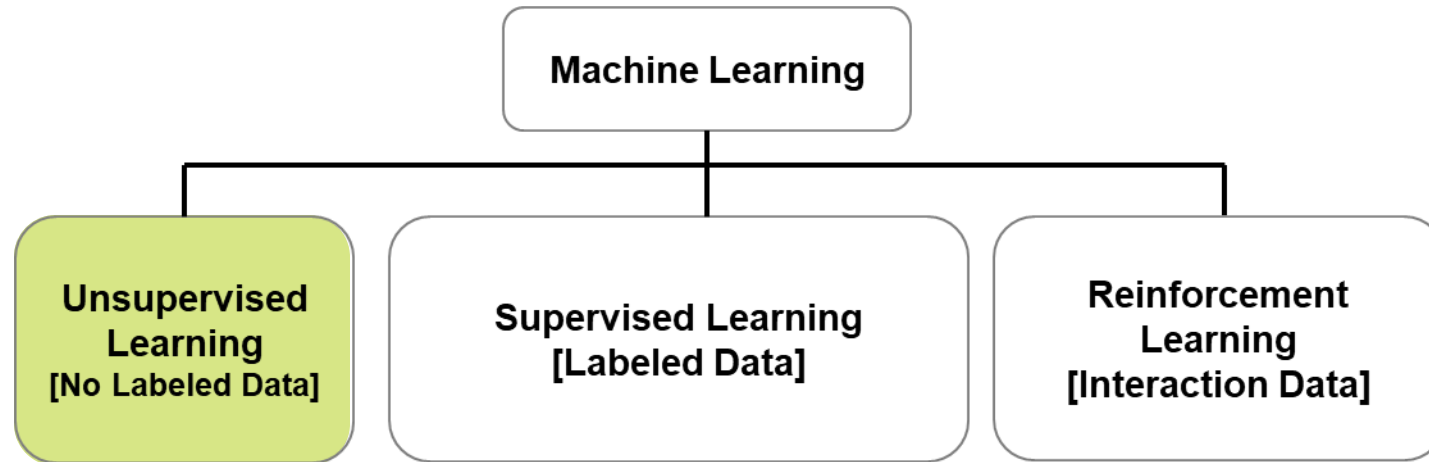
## Reinforcement learning:

- Learning through trial & error [*interaction*]
- Complex problems typically need deep learning [*Deep Reinforcement Learning*]
- It's about learning a **behavior** or accomplishing a **task**

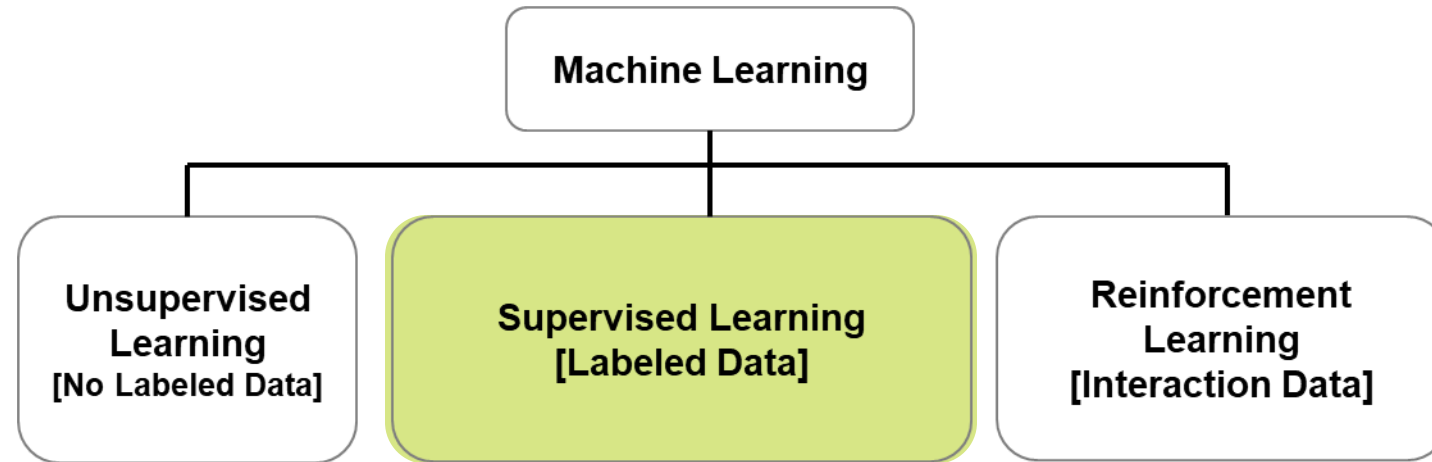
# Reinforcement Learning vs Machine Learning vs Deep Learning



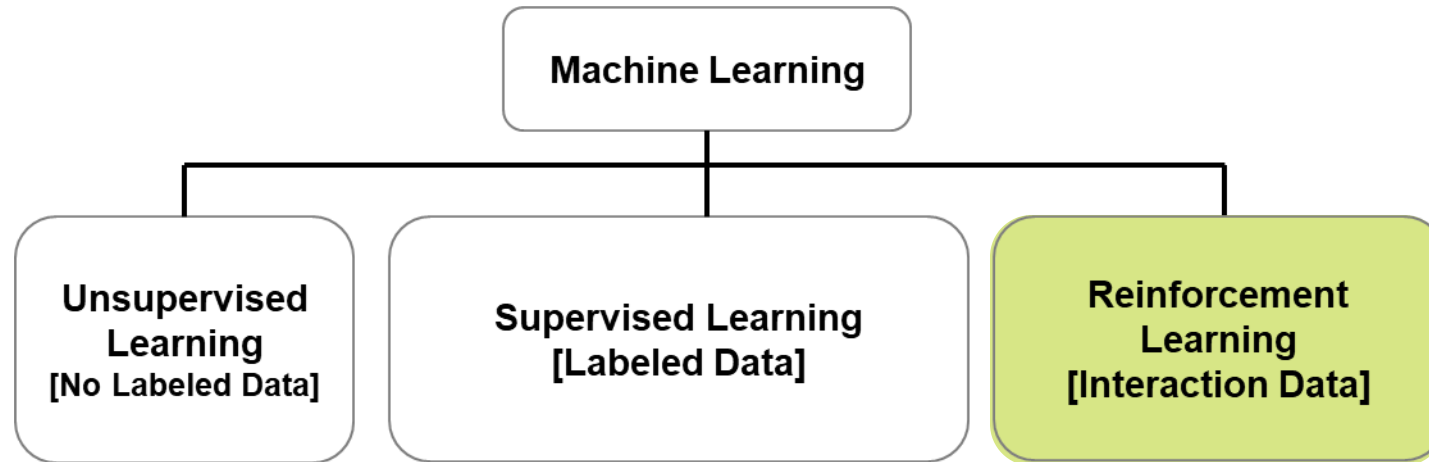
# Reinforcement Learning vs Machine Learning vs Deep Learning



# Reinforcement Learning vs Machine Learning vs Deep Learning




# Reinforcement Learning vs Machine Learning vs Deep Learning



# A Practical Example of Reinforcement Learning


## Training a Self-Driving Car



- Vehicle's computer learns how to drive...  
(**agent**)
- using sensor readings from LIDAR, cameras,...  
(**state**)
- that represent road conditions, vehicle position,...  
(**environment**)
- by generating steering, braking, throttle commands,...  
(**action**)
- based on an internal state-to-action mapping...  
(**policy**)
- that tries to get you from A to B without an accident while possibly optimizing driver comfort & fuel efficiency...  
(**reward**).
- The policy is updated through repeated trial-and-error by a **reinforcement learning algorithm**

# A Practical Example of Reinforcement Learning

## A Trained Self-Driving Car Only Needs A Policy To Operate

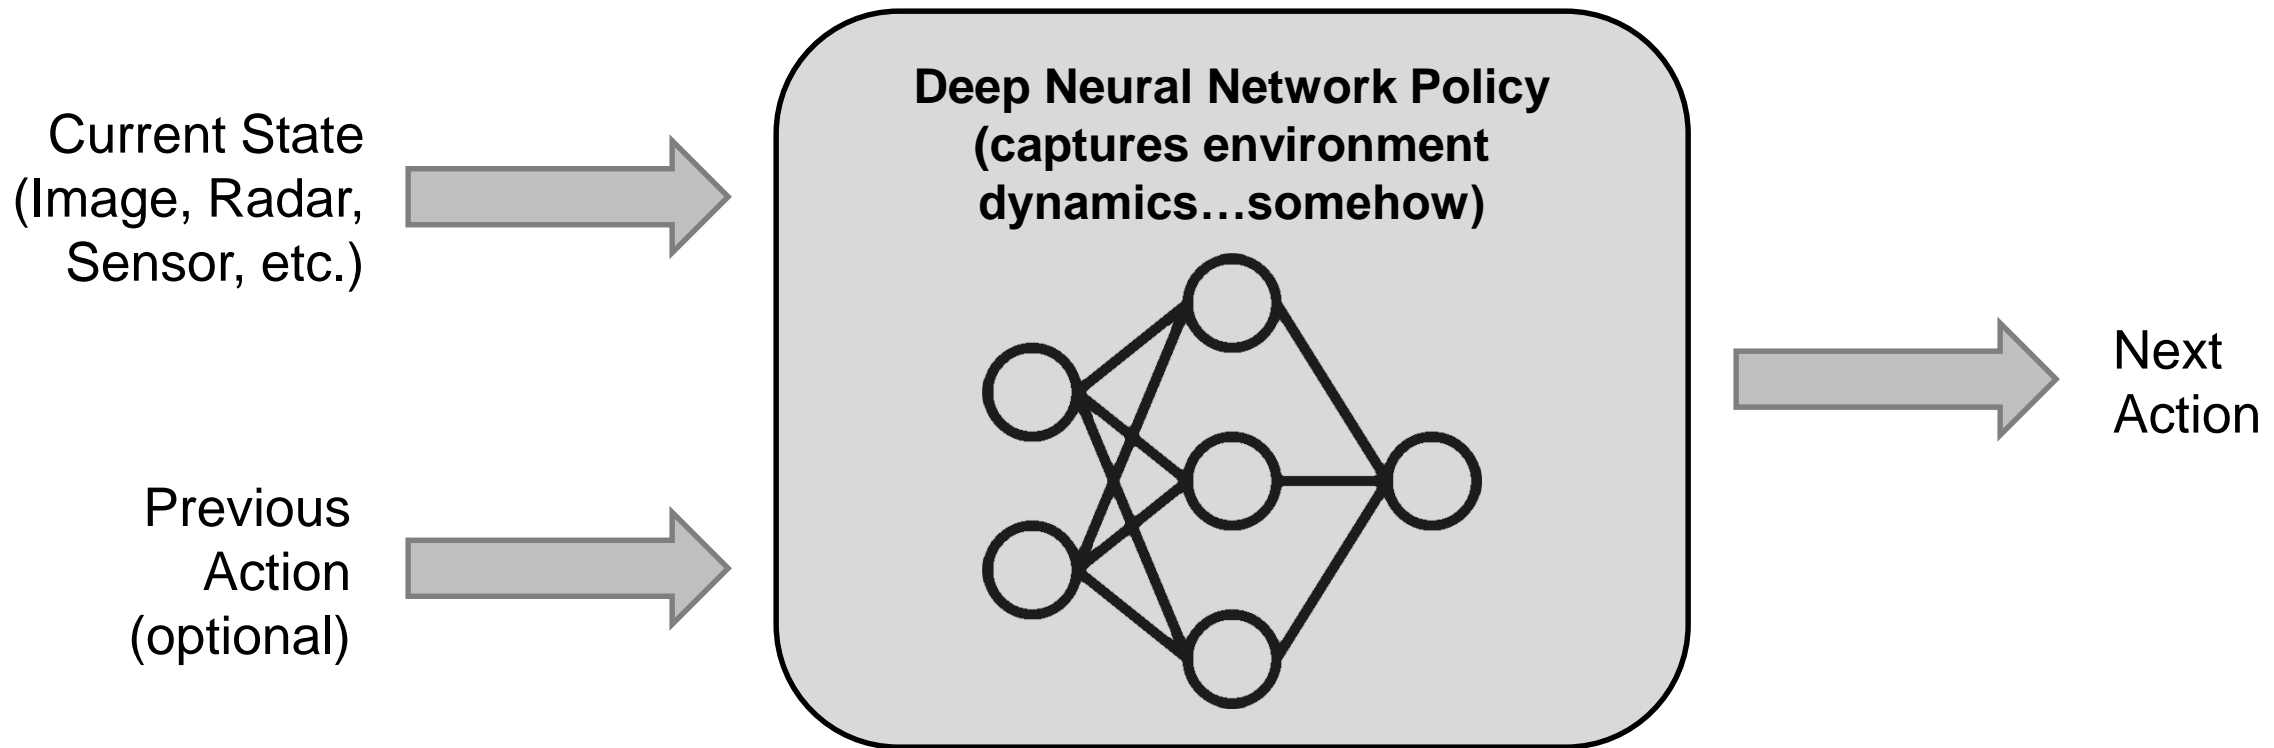


- Vehicle's computer uses the final state-to-action mapping... (**policy**)
- to generate steering, braking, throttle commands,... (**action**)
- based on sensor readings from LIDAR, cameras,... (**state**)
- that represent road conditions, vehicle position,... (**environment**)

**By definition, this trained policy puts into practice what is in the reward function**




A deep neural network trained using reinforcement learning is a black-box model that determines the best possible action




By representing policies using deep neural networks, we can solve problems for **complex, non-linear systems** (continuous or discrete) by directly using **data that traditional approaches cannot use easily**

# How do I set it up and solve it?

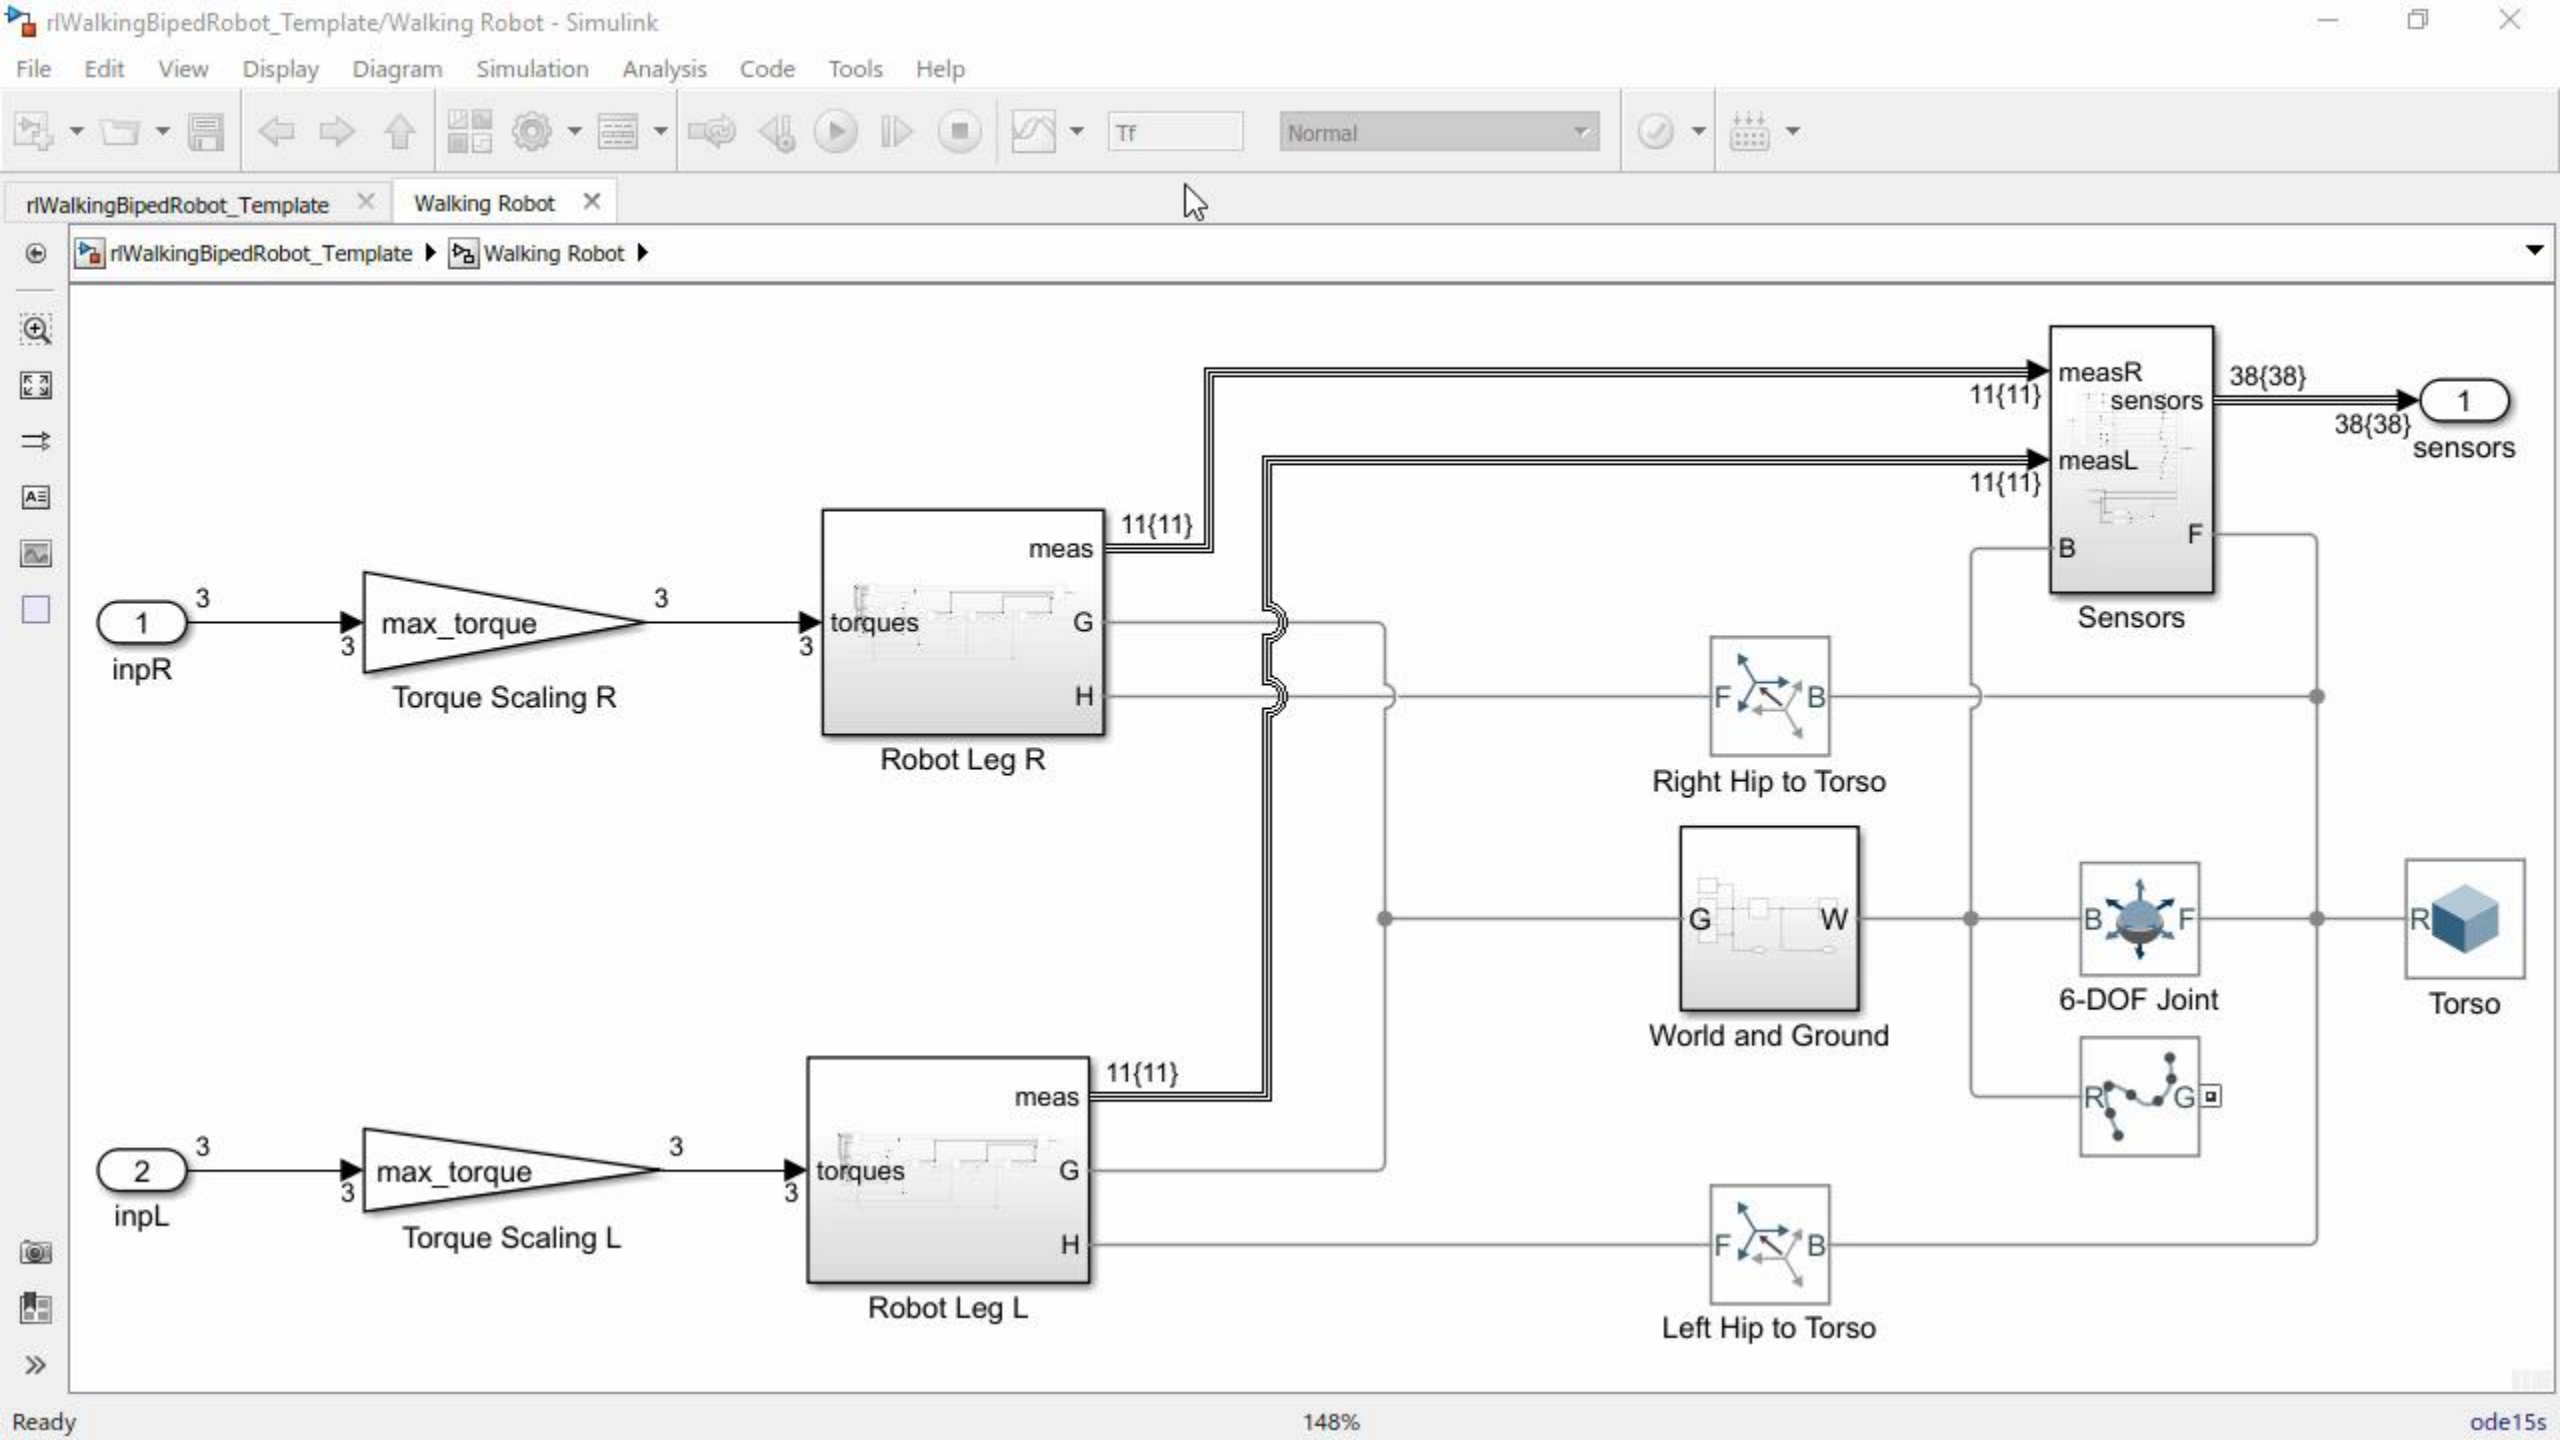


# Steps in the Reinforcement Learning Workflow





# Environment






# Steps in the Reinforcement Learning Workflow





# Reward Function




## Long-Term Reward (Q)

- $Q(S_0, a_1) = +105$  (optimal)
- $Q(S_0, a_2) = +60$
- $Q(S_0, a_3) = +30$

The logic you used right now to decide you should take action  $a_1$  is a policy

# Policy & Agent







- $T_s$  is the sample time of the environment.
- $T_f$  is the final simulation time of the environment.

## Create Environment Interface

Create the observation specification.

```
5 numObs = 29;  
6 obsInfo = rlNumericSpec([numObs 1]);  
7 obsInfo.Name = 'observations';
```

Create the action specification.

```
8 numAct = 6;  
9 actInfo = rlNumericSpec([numAct 1], 'LowerLimit', -1, 'UpperLimit', 1);  
10 actInfo.Name = 'foot_torque';
```

Create the environment interface for the walking robot model.

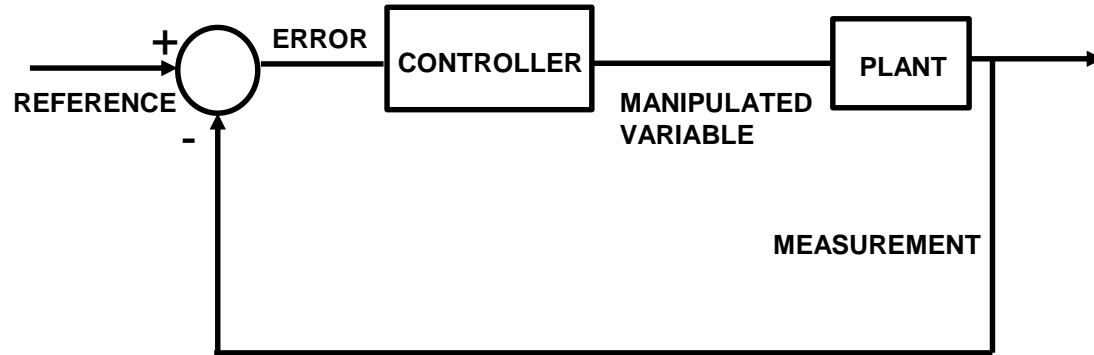
```
11 blk = [mdl, '/RL Agent'];  
12 env = rlSimulinkEnv(mdl, blk, obsInfo, actInfo);  
13 env.ResetFcn = @(in) walkerResetFcn(in, upper_leg_length/100, lower_leg_length/100, h/100);
```

## Create DDPG agent

A DDPG agent approximates the long-term reward given observations and actions using a critic value function representation. A DDPG agent decides which action to take given observations using an actor representation. The actor and critic networks for this example are inspired by [2].

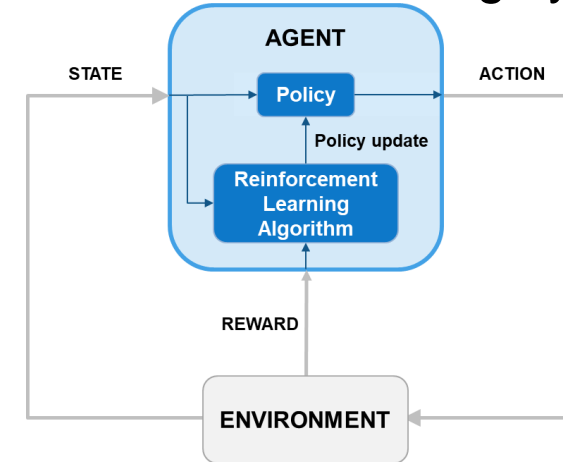
# Reinforcement Learning vs Controls

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**



# When would you use Reinforcement Learning?

	Controller Capability	Computational Cost in Training/Tuning	Computational Cost in Deployment
PID	Low	Low	Low
Model Pred Control	High	Low	High
Reinforcement Learning	Very High	High	Medium

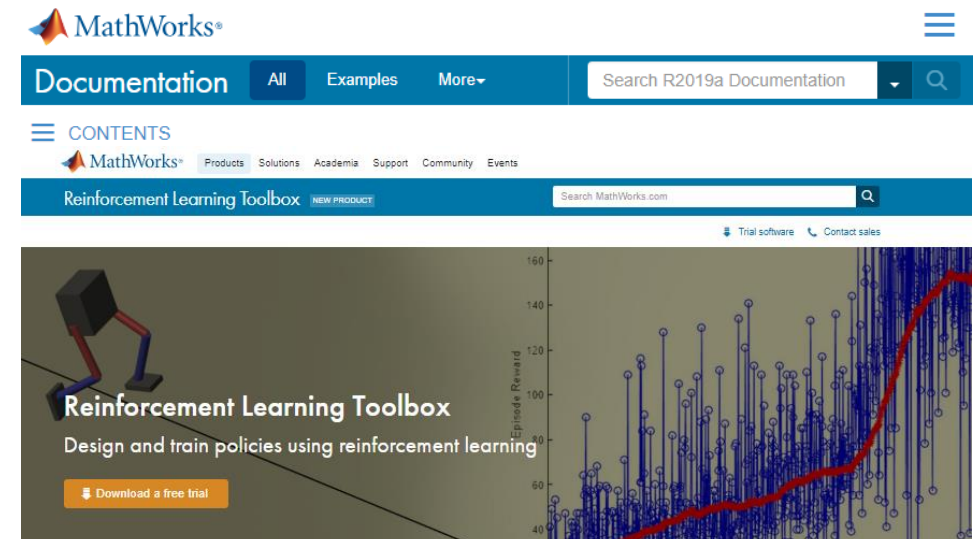
Reinforcement learning might be a good fit if

- An environment model is available (trial & error on hardware can be expensive), and
- Training/tuning time is **not** critical for the application, and
- Uncertain environments or nonlinear environments

# Reinforcement Learning Toolbox

## New in R2019a

- Built-in and custom algorithms for reinforcement learning
- Environment modeling in MATLAB and Simulink
- Deep Learning Toolbox support for designing policies
- Training acceleration through GPUs and cloud resources
- Deployment to embedded devices and production systems
- 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.

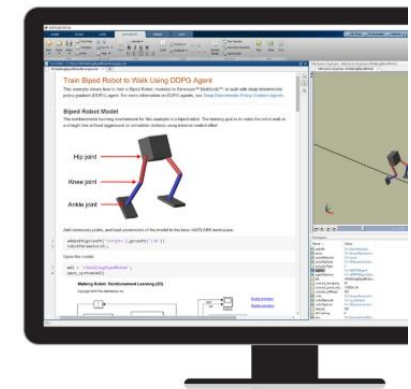
The toolbox includes reference examples for using reinforcement learning to design controllers for robotics and automated driving applications.

### Training and Validation

Train and simulate reinforcement learning agents

### Policy Deployment

Code generation and deployment of trained policies



# Automotive Applications

- Controller Design
- Lane Keep Assist
- Adaptive Cruise Control
- Path Following Control
- Trajectory Planning

**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 modeled in Simscape Multibody.

**Train DDPG Agent for Adaptive Cruise Control**

Train a reinforcement learning agent for an adaptive cruise control application.


**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**


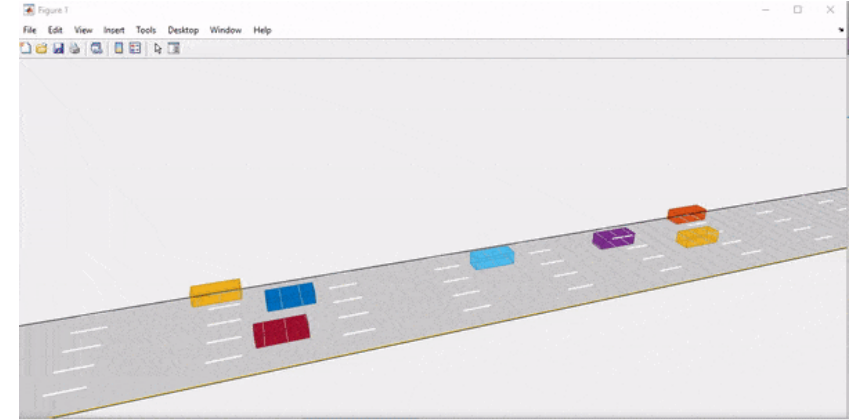
Train a reinforcement learning agent for a lane following application.

# Takeaways



# Simulation and virtual models are a key aspect of reinforcement learning

- Reinforcement learning needs **a lot** of data (*sample inefficient*)
  - Training on hardware can be prohibitively expensive and dangerous
- Virtual models allow you to simulate conditions hard to emulate in the real world
  - This can help develop a more robust solution
- Many of you have already developed MATLAB and Simulink models that can be reused



# Resources

- [Examples](#) for automotive and autonomous system applications
- [Documentation](#) written for engineers and domain experts
- [Tech Talk video series](#) on reinforcement learning concepts for engineers



**Thank You!**

