Machine Learning (ML) with SPIKE Prime

Teacher Workshop
May 10th, 2022

Tufts University | School of Engineering
Center for Engineering Education and Outreach

mta modern teaching aids

While you wait fo us to get started…

Share "***why are you here?***" or "***what do you hope to learn today?***" in the chat…

Visit this workshop's website:

bit.ly/MLwithSPIKE

Adjust your robot build to make sure all the cables are well secured!

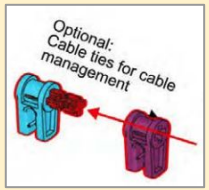

Agenda

- Introduction to Professor Ethan (*Dr. E*) Danahy from Tufts University
  - Examples of how Dr. E uses SPIKE Prime to teach Intro to Engineering at the university level
- Introduction to Machine Learning (ML)
- Supervised Learning Overview
  - **Hands-on with SPIKE Prime:** training Supervised Classification via Nearest Neighbor (NN)
- Unsupervised Learning Overview
- Reinforcement Learning Overview
  - **Hands-on with SPIKE Prime:** training Q-Learning for Finite Markov Decision Process (FMDP)
- Next Steps and Further Exploration

Today's Workshop Facilitators

Ethan (*Dr. E*) Danahy
Research Associate Prof.
Tufts University

Milan Dahal
Doctoral Student
Tufts University

Tufts UNIVERSITY | School of Engineering

Tufts UNIVERSITY | School of Engineering
Center for Engineering Education and Outreach

Tufts UNIVERSITY

School of Engineering

Center for Engineering Education and Outreach

Tufts CEEO is an interdisciplinary center dedicated to creating the next generation of problem solvers, kindergarten through college, through engineering education. With educational research, tool and technology development, and global outreach initiatives, we work to bridge the divide between research and practice.

TEACHER ENGINEERING EDUCATION PROGRAM
TUFTS UNIVERSITY

Tufts CEEO's online program to earn graduate credits in engineering education!

Special "*Secondary Educator*" track focusing on high-school topics and technologies…

Now enrolling for September 2022 courses! Visit http://teep.tufts.edu

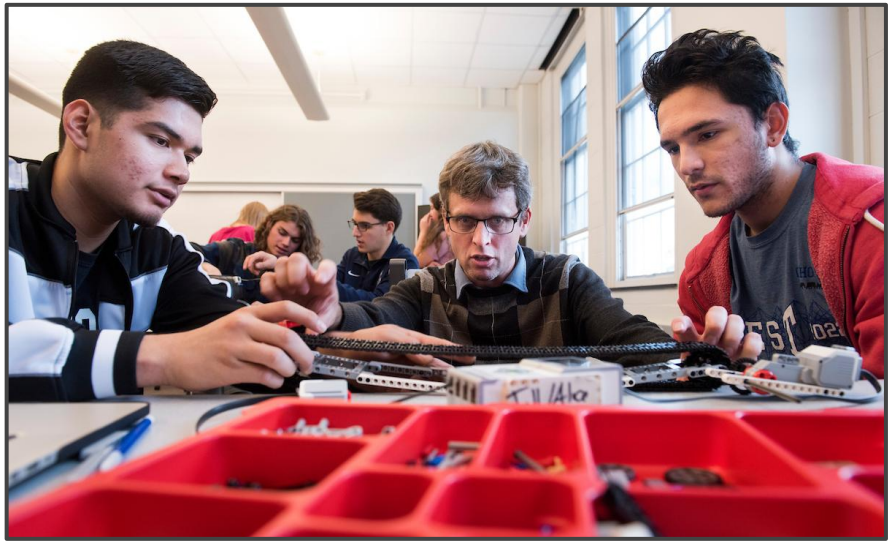

Project 01: Getting to Know You

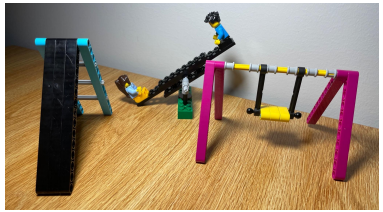



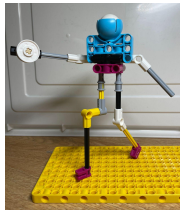

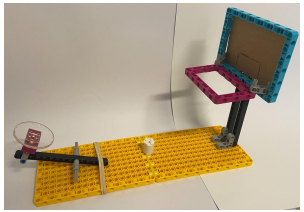

Project 03: Biomimicry

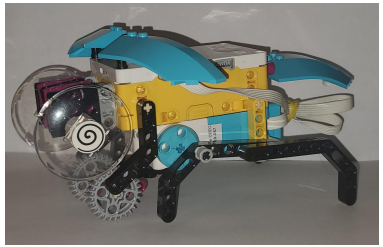
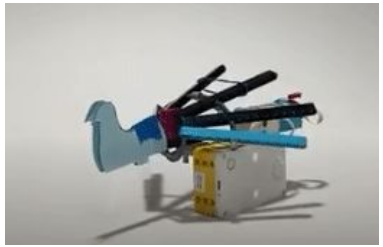
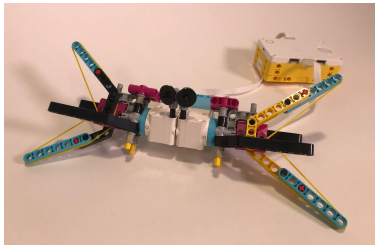
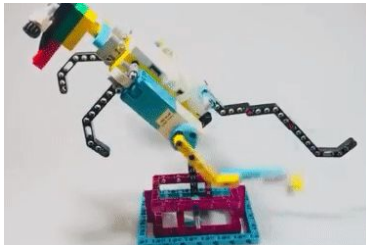
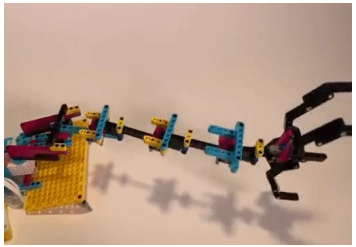

Project 08: Astronaut Tools

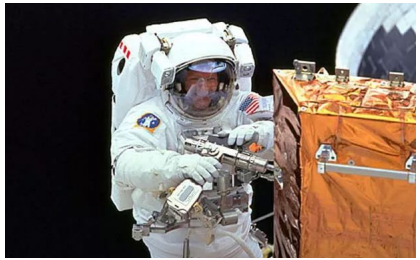
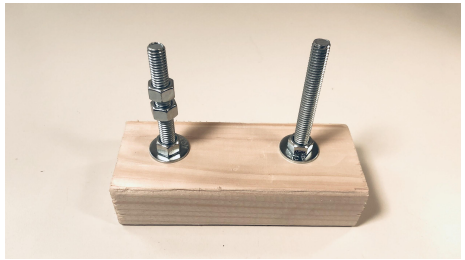
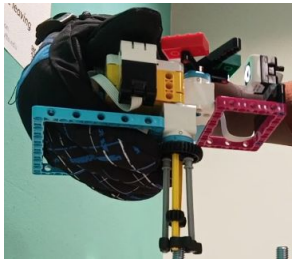
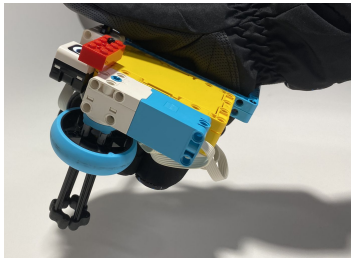
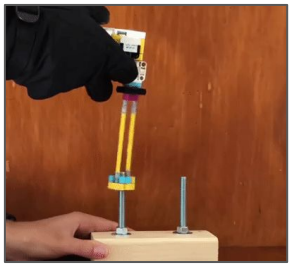

Project 10: Haunted House

Final Project: Playful Creations

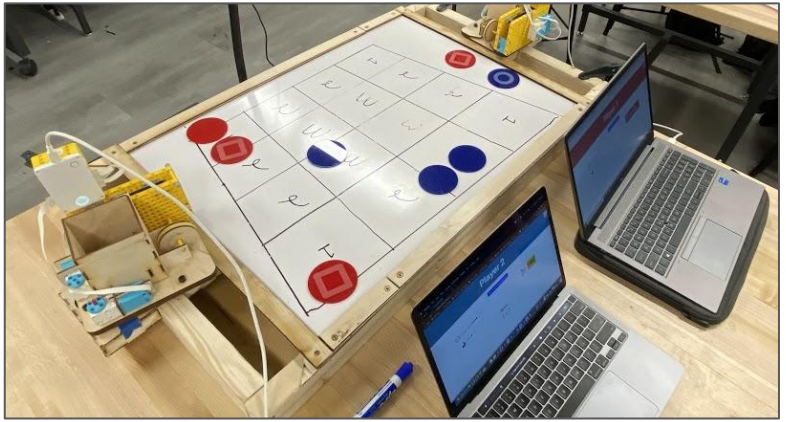

Introduction to Engineering: Semester Design (*1st half*)

Project Attributes	Project 01: Getting to Know You	Project 02: Silly Walks	Project 03: Biomimicry	Project 04: Top Spinner	Project 05: Spirograph Drawing Machine	Project 06: Simple Car	Project 07: Draw a Square	Project 08: Astronaut Tools	Project 09: Robotic Arm	Project 10: Haunted House
Context/Client										
Programming Structures (Loops, etc)							*optional*			
Sensors							*optional*			
Interaction with World										
Contest or Pressure Deliverable Deadline										
Simple Programming										
Gears							*optional*	*optional*	*optional*	*optional*
Motors										
Building										
Duration	week-long	in-class	week-long	in-class	week-long	in-class	week-long	week-long	week-long	two weeks

Tufts UNIVERSITY | School of Engineering

Artificial Intelligence (AI) and Machine Learning (ML)

Interact with Systems

Someone else builds the system
Someone else trains system
We interact with the system

Train Systems

Someone else builds the system
We train the system
We interact with the system

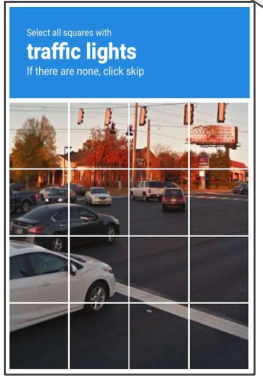

Create Systems

We create the system
We train the system
We interact with the system

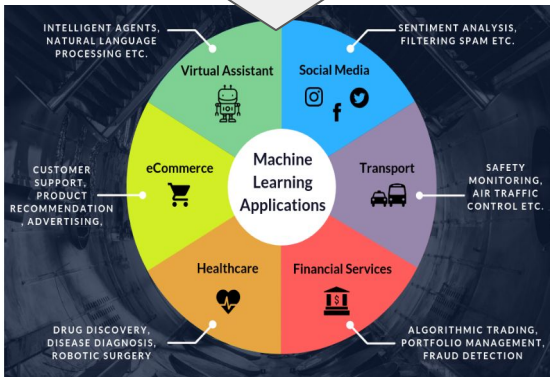

Machine Learning (ML)

Subset of Artificial Intelligence (AI)

Machine Learning (ML): allows a machine to automatically learn from past data without programming explicitly.

In ML, we teach machines with data to perform a particular task and give an accurate result. Machine Learning often has a limited scope, can perform only those specific tasks for which it is trained, and is mainly concerned about accuracy and patterns.

Three main types of Machine Learning:

- ***Supervised Learning***
- ***Unsupervised Learning***
- ***Reinforcement Learning***

Supervised Learning vs *Unsupervised Learning*

Supervised Learning

Uses ***labeled*** data sets

Goal: classifying data or predicting outcomes accurately

Classification: accurately assign test data into specific categories.

Algorithms we'll explore:

- Nearest Neighbor (NN Classification)
- K-Nearest Neighbor (KNN Classification)

Unsupervised Learning

Analyze and cluster ***unlabeled*** data sets

Goal: discover hidden patterns in data without the need for human intervention

Clustering: automatically grouping unlabeled data based on their similarities or differences.

Algorithm we'll explore:

- K-Means Clustering

Supervised Classification via Nearest Neighbor (NN)

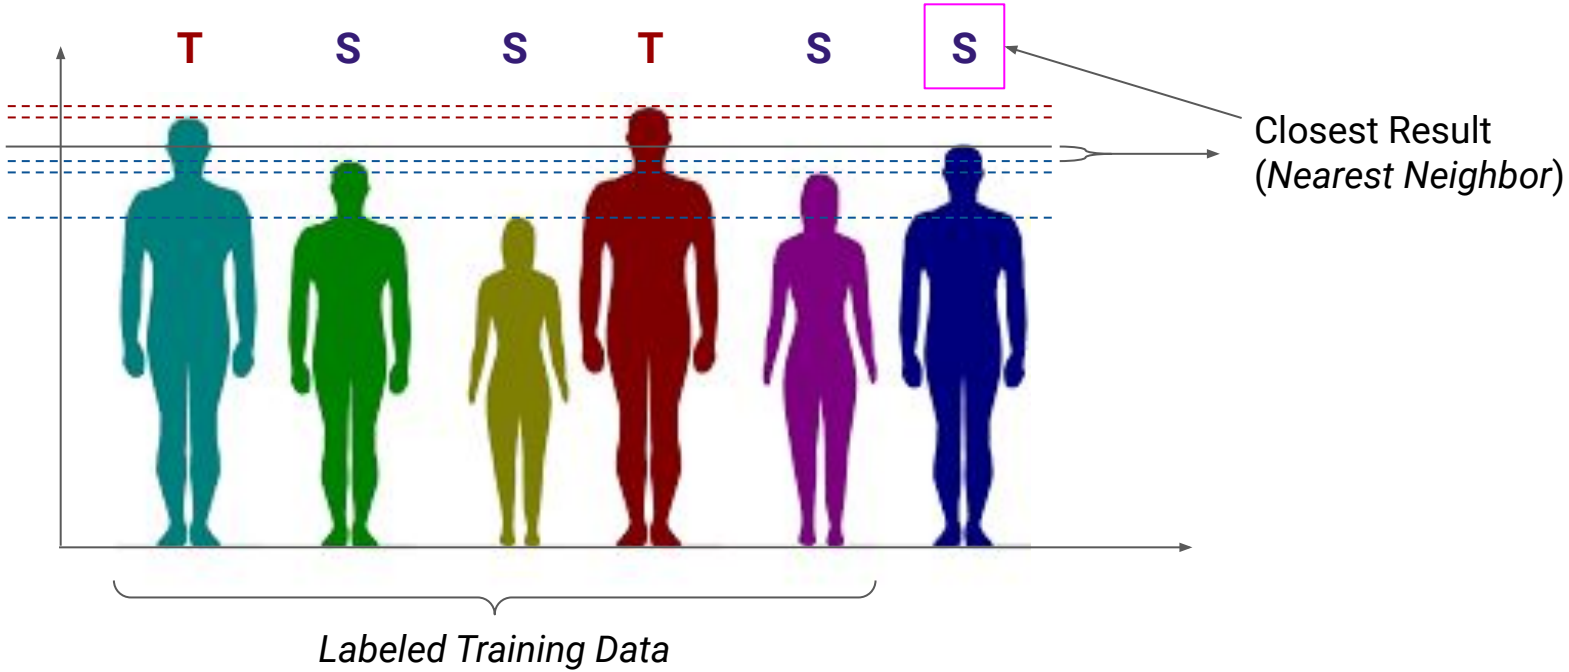

Hands-on Activity: ***Supervised Classification***

- Attach **DistanceSensor** to "Port E"
- Attach **ColorSensor** to "Port F"
- *Note 1:* can also use **ForceSensor**
- *Note 2:* can customize Port values within code

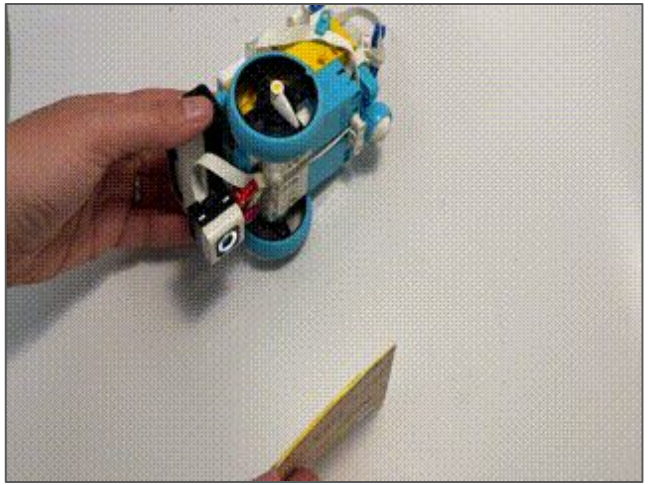

Hub's Right Button: add training data to supervised classification

Hub's Left Button: make predictions based on new data

Supervised Classification via Nearest Neighbor (NN)

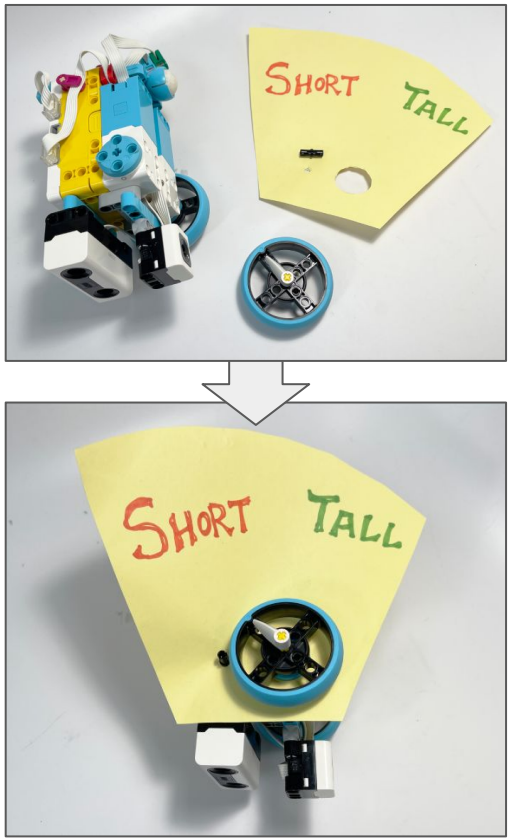

https://www.youtube.com/watch?v=e1MvAgCR2Cw

The Story of Punxsutawney Phil

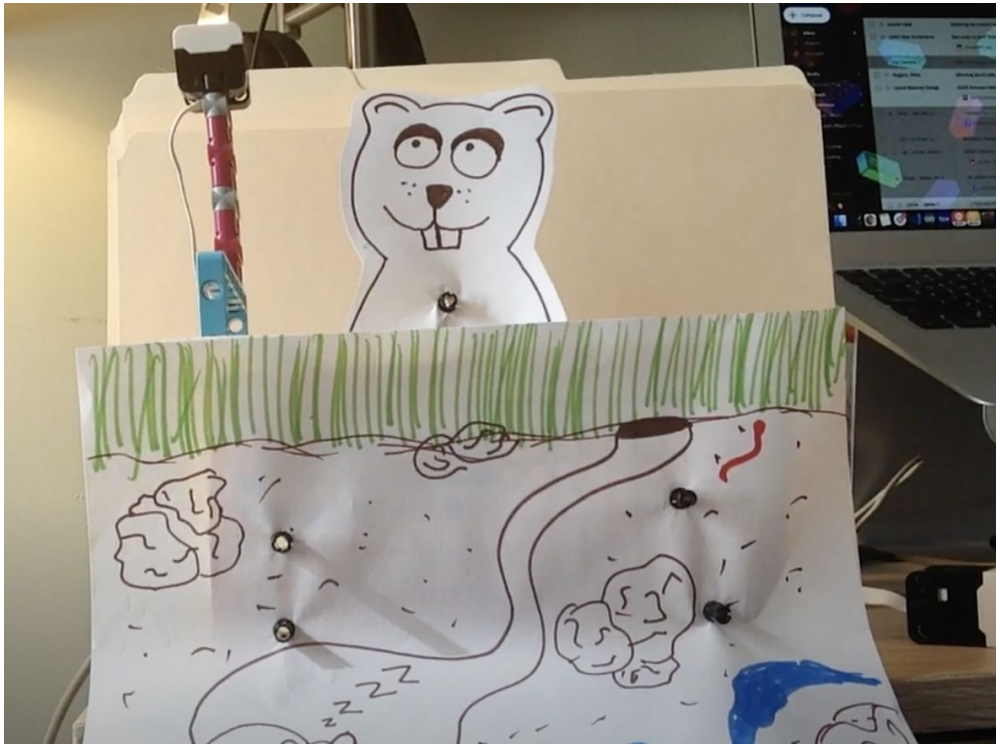

https://www.youtube.com/watch?v=0aMVeASDi70

Unsupervised Learning: **K-Means Clustering**

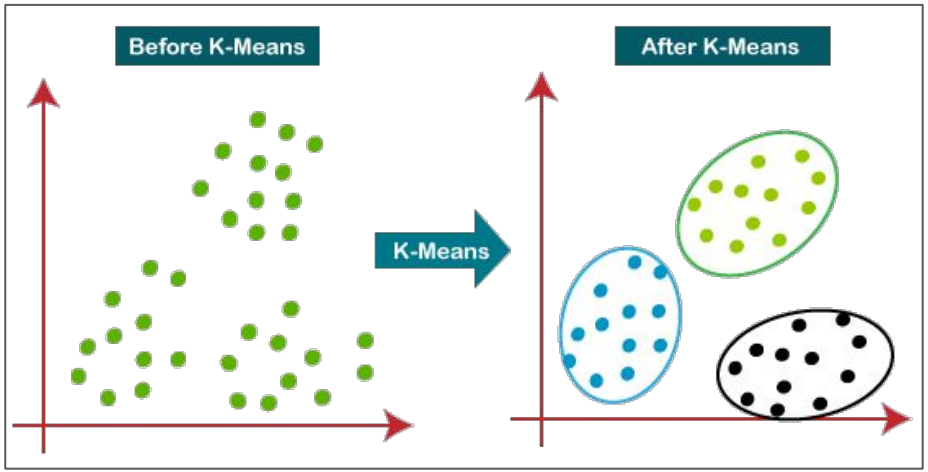

“K”: how many groupings you want

“Means”: the average value of all the data in the group

“Cluster”: data-points organized with the nearest mean

- “cluster center” or “cluster centroid”

K-Means Clustering Algorithm

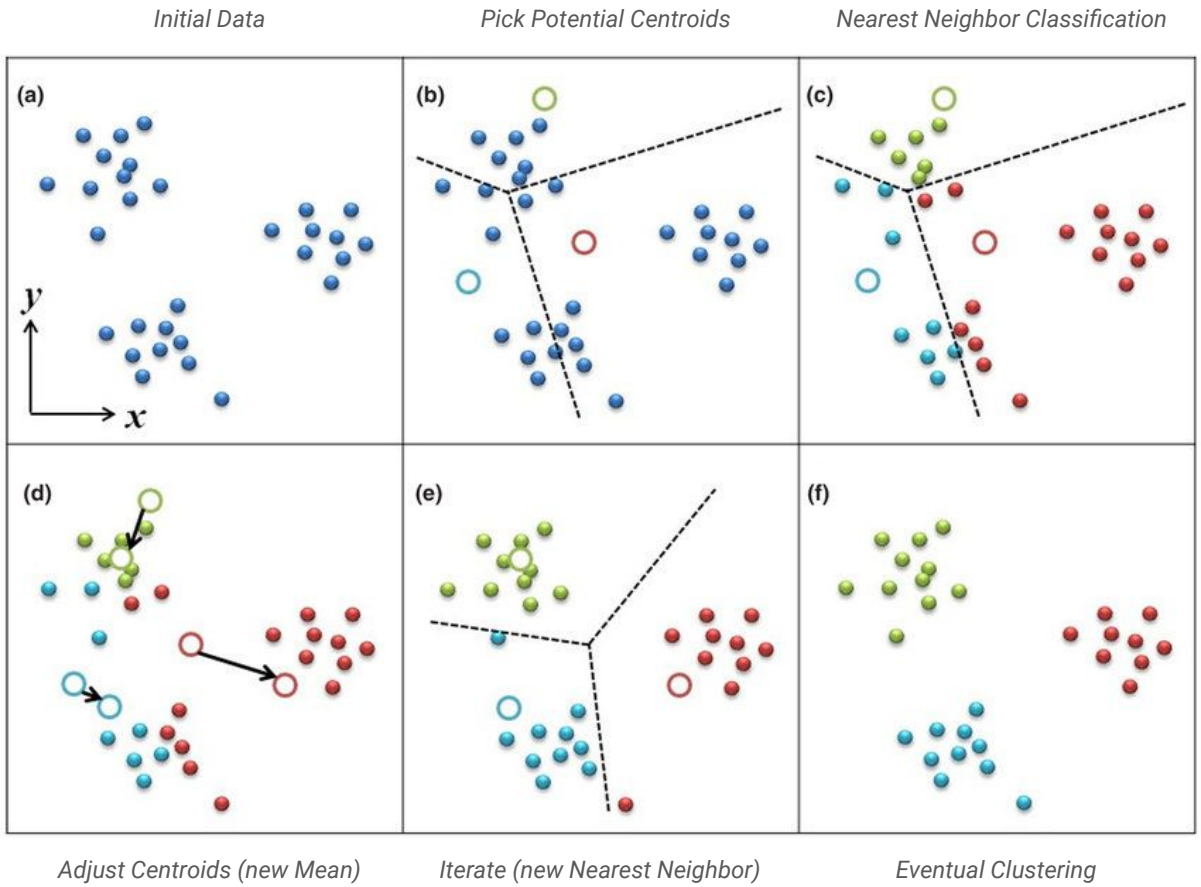

Reinforcement Learning

Reinforcement Learning (RL) is a machine learning technique that trains an algorithm (agent) by evaluating a current situation (state), taking an action, and receiving feedback (reward) from the environment after each act.

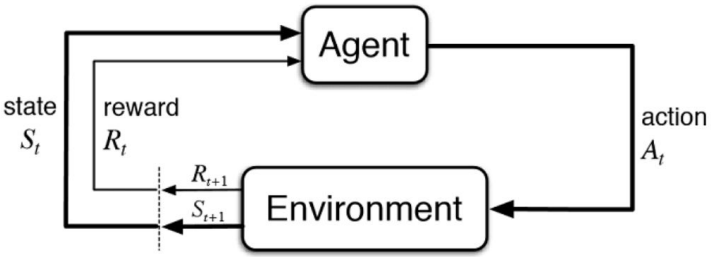

RL algorithms learn best through many attempts and failures. Short-term rewards lead to cumulative, long-term success.

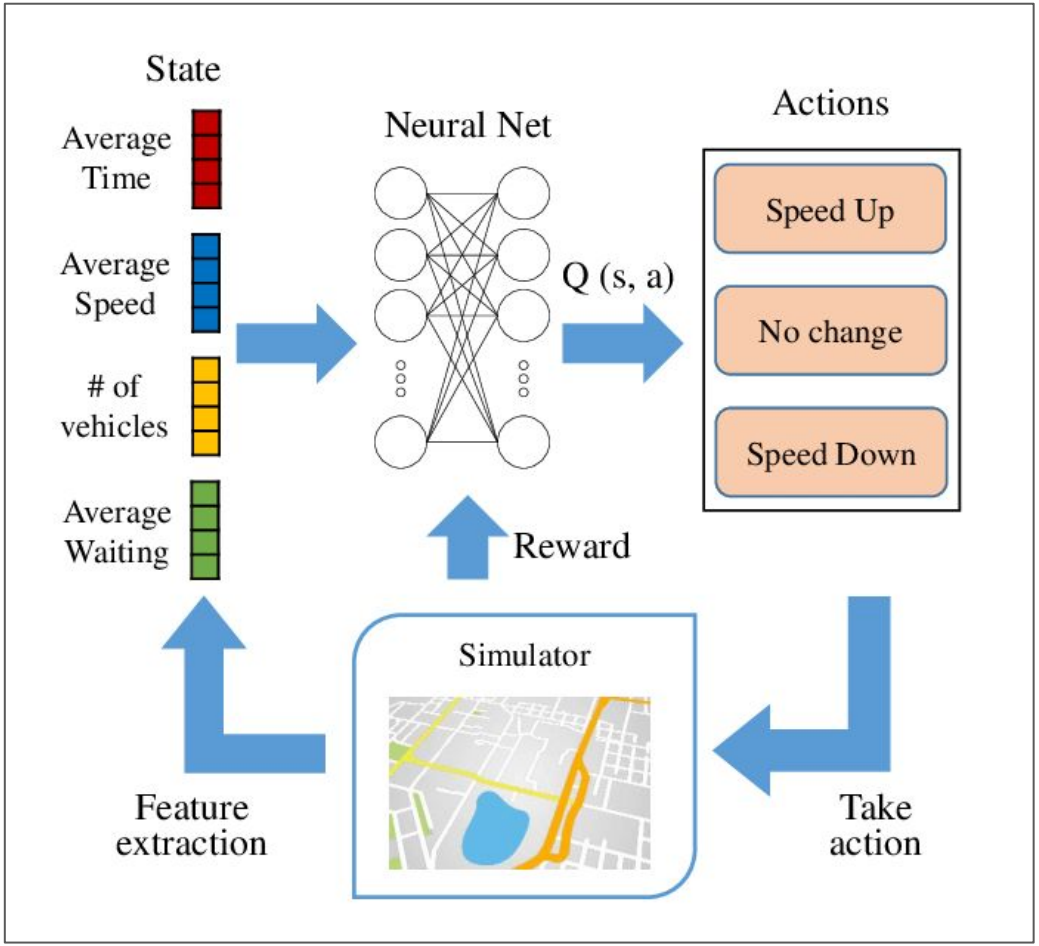

Deep Reinforcement Learning from "*Joint Modeling of Dense and Incomplete Trajectories for Citywide Traffic Volume Inference*" by Tang, et al (2019)

Reinforcement Learning Terminology

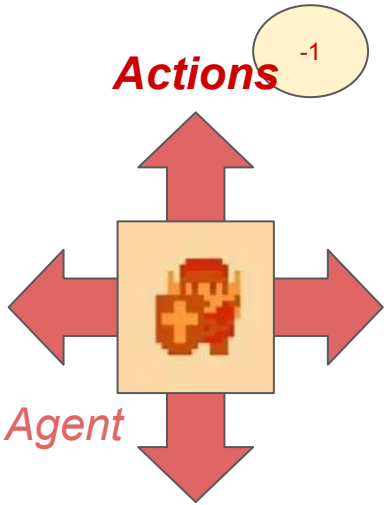

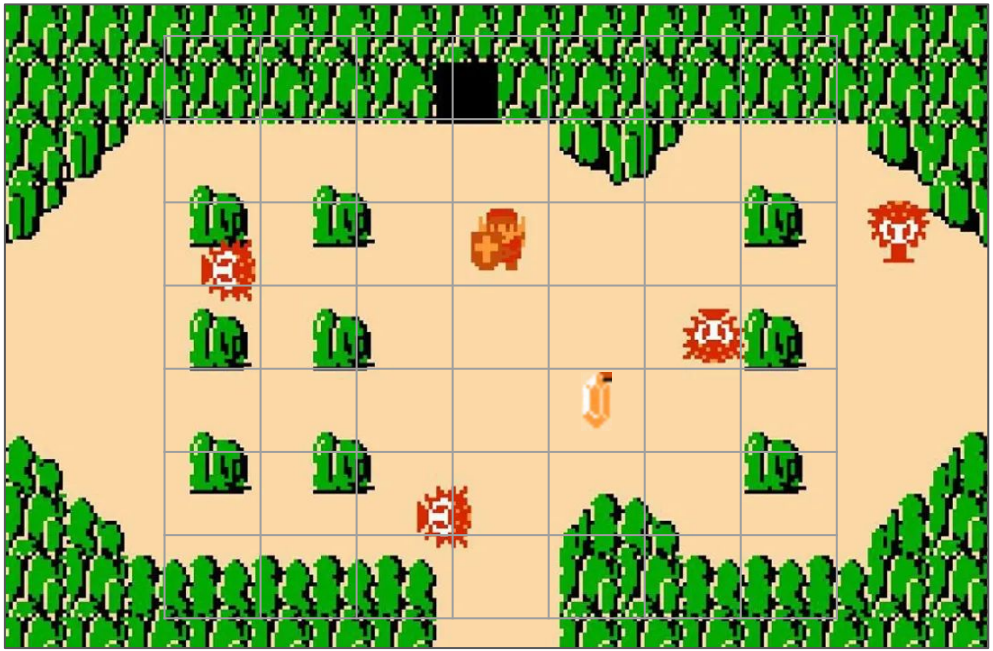

Environment (and States)

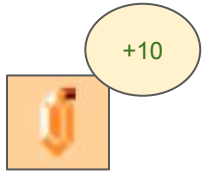

Rewards

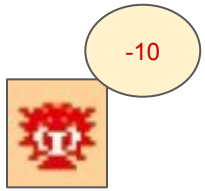

Penalties
(negative rewards)

Reinforcement Learning Terminology

Agent

Exploration

- Collection of ***episodes*** (aka trials)
- Each ***episode*** has series of ***steps***
- At each ***step***:
  - Pick and perform ***action,*** determine new ***state***
  - Assign ***reward*** (or penalties)
- ***Epsilon-Greedy***: don't always repeat the same actions; randomly explore new possibilities!

Training the Model

Exploitation

- Use the trained model to allow the ***agent*** to fully engage with the ***environment***.
- For each new ***state*** reached, look up optimal ***action*** previously determined by highest ***reward***, and execute it.

Testing the Model

Decisions to Make:

- Assign one reward at end of episode? Or at each step?
- How much training/exploration (num of episodes) is enough?

Our Exploration:
Silly Walks

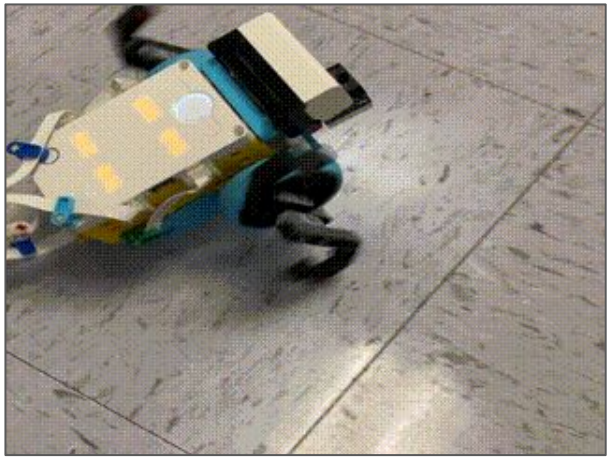

Challenge: build a robot that moves without wheels

These robots rarely go straight!

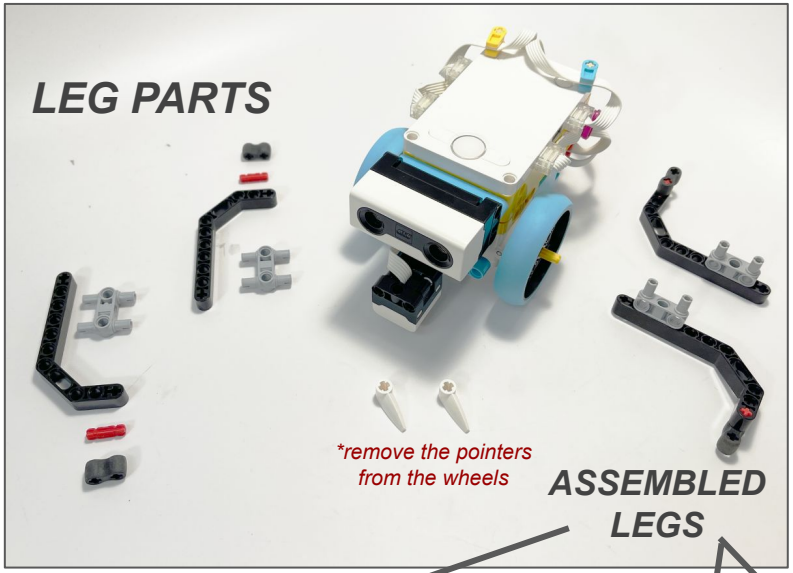

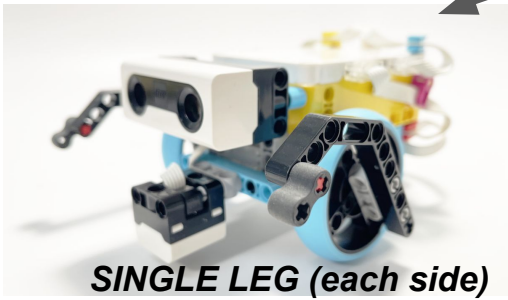

SINGLE LEG (each side)

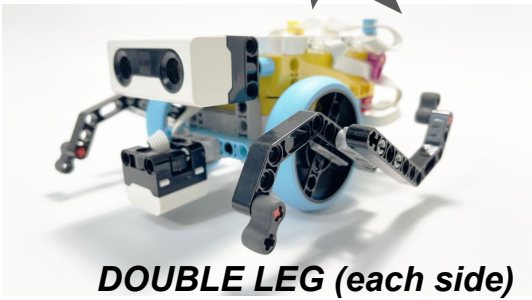

DOUBLE LEG (each side)

Our Exploration: ***Silly Walks***

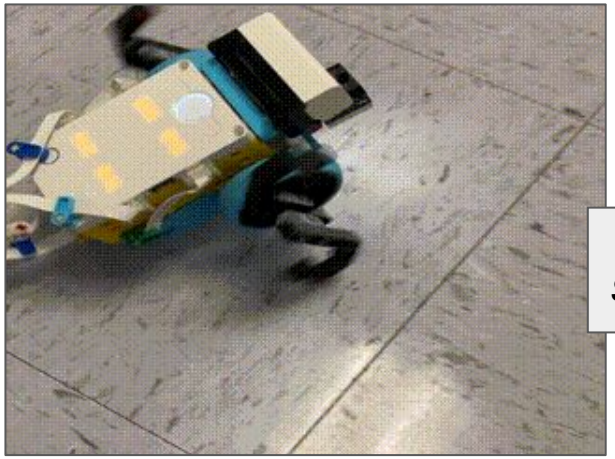

Challenge: build a robot that moves without wheels

Make a ***Smart Walker***

State: *Off-Center Angle*

Agent: *Robot*

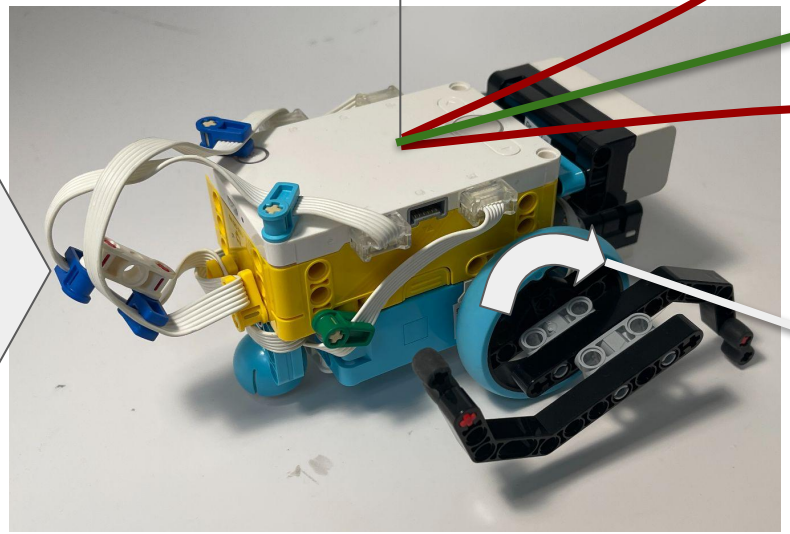

Action: *Motor Speeds*

Environment: Table/*Floor*

These robots rarely go straight!

Reinforcement Learning Algorithm: ***Q-Learning***

Actions: *Different Motor Speeds*

Motor 1	Motor 2	State -2	State -1	State 0	State 1	State 2
Slow	*Slow*	0	0	0	0	0
Slow	*Medium*	0	0	0	0	0
Slow	*Fast*	0	0	0	0	0
Medium	*Slow*	0	0	0	0	0
Medium	*Medium*	0	0	0	0	0
Medium	*Fast*	0	0	0	0	0
Fast	*Slow*	0	0	0	0	0
Fast	*Medium*	0	0	0	0	0
Fast	*Fast*	0	0	0	0	0

Q-Values
(aka ranking of the best actions for each state)

Reinforcement Learning Algorithm: ***Q-Learning***

		-10	-2	+10	-2	-10
Motor 1	**Motor 2**	***State -2***	***State -1***	***State 0***	***State 1***	***State 2***
Slow	*Slow*	0	0	0	0	0
Slow	*Medium*	0	0	0	0	0
Slow	*Fast*	0	0	0	0	0
Medium	*Slow*	0	0	0	0	0
Medium	*Medium*	0	0	0	0	0
Medium	*Fast*	0	0	0	0	0
Fast	*Slow*	0	0	0	0	0
Fast	*Medium*	0	0	0	0	0
Fast	*Fast*	0	0	0	0	0

Exploration (*Training*):

- Pick an **action** (randomly or based on previous data)
- Compare previous-**state** to new-**state**
- Identify **reward** based on new-**state**
- Update **Q-Values** in **Q-Table** for that **state**/**action** pair
  - Values updated based on *Bellman Equation*

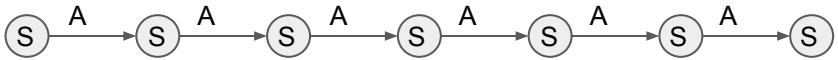

Exploitation (*Testing*):

- Determine current **state**
- Look up highest **Q-Value** and determine optimal **action**
- Execute new **action**
- Determine resulting new **state**

Hands-on Activity: **Reinforcement Learning**

“Smart Walkers”

Before you begin:

- Make sure your cables are organized
- Build new "legs" for your robot
- Have a big open space to train your robot

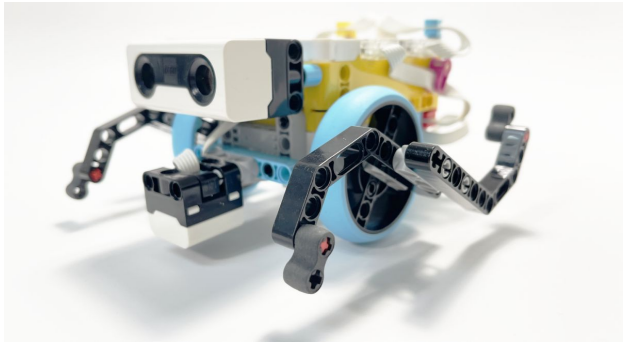

Note:

- Keep robot from hitting objects or flipping over
- Have faith! (It won’t seem like it’s working at first.)
- It might get worse before it gets better (as it explores environment and new actions)
- You might have to train it through 10 to 15 different episodes to get a good result!

Next Steps and Further Exploration

SPIKE Prime AI Puppy

Series of five activities to work through learning Artificial Intelligence and Machine Learning:

- Nearest Centroid Classification
- 1-Dimensional K-Nearest Neighbor (KNN)
- 3D KNN Algorithm
- Linear Regression and/or Reinforcement Learning
- Image Processing via Teachable Machines

Placemat Instructions:

https://www.ceeoinnovations.org/RoboticsPlayground/playlists/SpikeAIPuppy.html

CEEO Innovation's Robotics Playground

A huge collection of SPIKE Prime (and other platform) activities detailed in simple open-ended "Placemat" format.

Includes a number of Artificial Intelligence and Machine Learning explorations.

https://www.ceeoinnovations.org/RoboticsPlayground

Advanced SPIKE Prime workshops by MTA

Modern Teaching Aids (MTA) Australia is hosting a number of new workshops around implementing advanced topics with the SPIKE Prime platform, for secondary teachers to explore and bring into their classrooms.

https://www.teaching.com.au

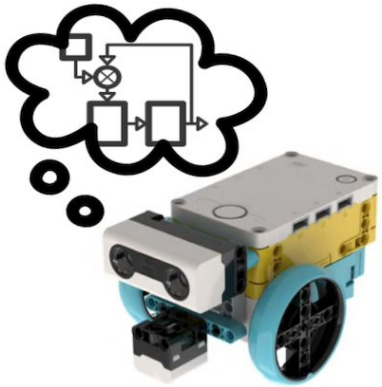

Thank You!

Tufts UNIVERSITY

School of Engineering
Center for Engineering Education and Outreach

mta modern teaching aids