# Designing Reward Functions

CS 1840 Introduction to Reinforcement Learning

# Before we start

- You don’t need your laptops for this class
- Sit within talking distance to someone!

# Let’s design a ride share algorithm with RL!

States?

Actions?

# Let’s design a ride share algorithm with RL!

## States?

- Current demand
- Driver location
- Traffic (travel time)
- Current price
- ……

## Actions?

# Let's design a ride share algorithm with RL!

## States?

- Current demand
- Driver location
- Traffic (travel time)
- Current price
- ……

## Actions?

- Driver-customer match
- New price
- ……

# The reward function

When we studied MDPs, the reward function is pre-defined.

When we studied multi-armed bandits, the reward function is unknown, but also clearly defined by arm means.

# The reward function

When we studied MDPs, the reward function is pre-defined.

When we studied multi-armed bandits, the reward function is unknown, but also clearly defined by arm means.

## But in the real world.....

# Option 1: Maximize $$ paid to the company

a) Whose interests are favored? Whose are ignored?
b) Is there any additional information you need to answer (a)?
c) What do you think some likely impact on the society might be?

# Option 2: Maximize $$ paid to the driver per hour

a) Whose interests are favored? Whose are ignored?
b) Is there any additional information you need to answer (a)?
c) What do you think some likely impact on the society might be?

# Option 3: Maximize total rides provided per hour

a) Whose interests are favored? Whose are ignored?
b) Is there any additional information you need to answer (a)?
c) What do you think some likely impact on the society might be?

# Option 4: Minimize total distance traveled per day

a) Whose interests are favored? Whose are ignored?
b) Is there any additional information you need to answer (a)?
c) What do you think some likely impact on the society might be?

Optimal policy is value-laden!

# Create your reward function and justify using it

# Some principles ethicists have proposed:

- We should improve the lives of the least well-off

# Some principles ethicists have proposed:

- We should improve the lives of the least well-off
- We should aim to maximize the *total* benefit to the society

# Some principles ethicists have proposed:

- We should improve the lives of the least well-off
- We should aim to maximize the *total* benefit to the society
- We should aim for equal *marginal* gain in benefits

## Some principles ethicists have proposed:

- We should improve the lives of the least well-off
- We should aim to maximize the *total* benefit to the society
- We should aim for equal *marginal* gain in benefits

## Something to think about:

- Are some of the reward designs we looked at incompatible, or can they be combined?

# An Example

Equally-weighted linear combination of (1) $$ earned for the company and (2) $$ earned for the drivers (3) number of rides provided per hour.

# How might this have happened?

## Uber and Lyft pricing algorithms charge more in non-white areas

**TECHNOLOGY** 18 June 2020 , updated 19 June 2020

By Donna Lu


**Uber and Lyft seem to charge more for trips to and from neighbourhoods with residents that are predominantly not white**
Gado Images / Alamy

The algorithms that ride-hailing companies, such as Uber and Lyft, use to determine fares appear to create a racial bias.

By analysing transport and census data in Chicago, Aylin Caliskan and Akshat Pandey at The George Washington University in Washington DC have found that ride-hailing companies

# How might this have happened?

Could this happen in your algorithm? If so, how can you prevent it from happening?

## Uber and Lyft pricing algorithms charge more in non-white areas

**TECHNOLOGY** 18 June 2020 , updated 19 June 2020

By Donna Lu


**Uber and Lyft seem to charge more for trips to and from neighbourhoods with residents that are predominantly not white**
Gado Images / Alamy

The algorithms that ride-hailing companies, such as Uber and Lyft, use to determine fares appear to create a racial bias.

By analysing transport and census data in Chicago, Aylin Caliskan and Akshat Pandey at The George Washington University in Washington DC have found that ride-hailing companies

# Thank you!

Feedback:
bit.ly/3RHtlxy

Attendance:
bit.ly/3RcTC9T