Fair and Efficient Social Decision-Making

CSCI 699

Evi Micha

Introduction

- People
  - Instructor: Evi Micha (pmicha@usc.edu)

- Info
  - Course Page: https://evi-micha.github.io/Teaching/699/index.html
  - Discussion Board: https://piazza.com/usc/fall2024/csci69930174
  - Assignments and Recordings: https://brightspace.usc.edu/d2l/home/154713

- Meeting
  - KAP 145, Fri 1-4.20 pm
  - Questions? Schedule 1-1 meeting by emailing me

What?

Fair and

Efficient

Social Decision-Making

What?

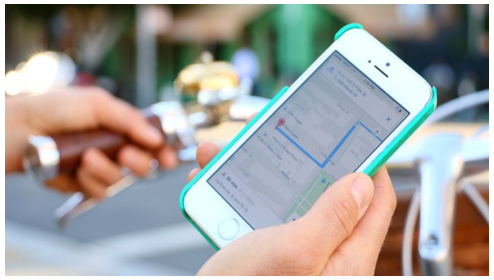

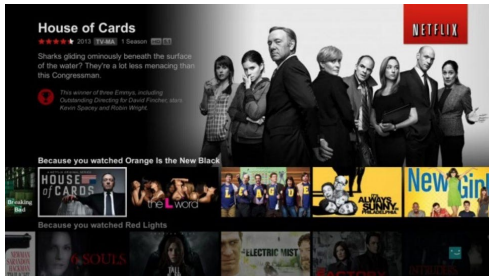

Design decision-making algorithms that treat people fairly, use limited resources efficiently, and foster social good

What?

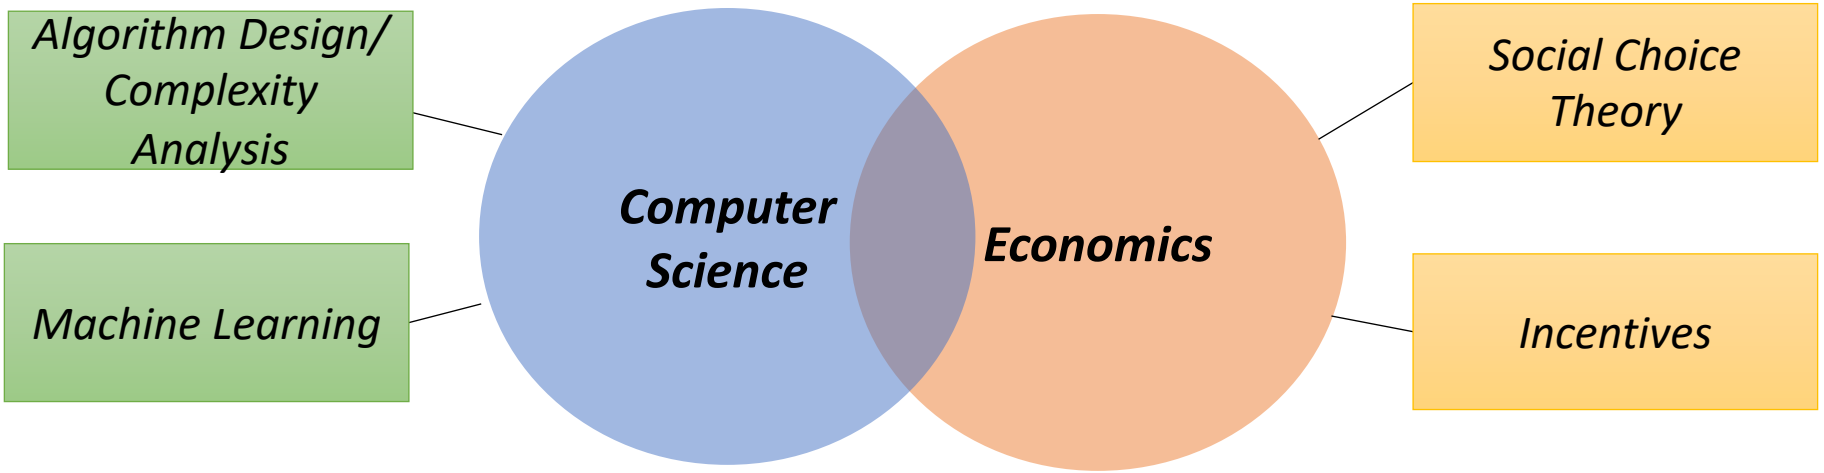
Algorithm Design/ Complexity Analysis
Machine Learning
Computer Science
Economics
Social Choice Theory
Incentives

Logistics

Optional Reference Textbooks

- Handbook of Computational Social Choice
  - By Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia
- Fairness and Machine Learning: Limitations and Opportunities
  - By Solon Barocas Moritz Hardt and Arvind Narayanan
- Online versions available on the course web page

Grading Policy

- 2 assignments: 40%
- Final project: 50%
- Class participation: 10%

Assignments

- Theoretical
  - They will require deriving intricate proofs

- We will assume…
  - Familiarity with abstract reasoning and proof techniques
  - Adequate familiarity of CS concepts (e.g., algorithm design, worst-case approximation, NP-hardness)
  - Adequate familiarity of math concepts (e.g., probability, statistics, linear algebra, calculus)
  - No prior background in economics

Assignments

- Individual assignments
  - Free to discuss with classmates or read online material
  - Must write solutions in your own words
    - Easier if you do not take any pictures/notes from the discussions

- Citation
  - For each question, you must cite the peer (write the name) or the online sources (provide links) referred, if any
  - Failing to do this is also plagiarism!

Other Policies

- "No Garbage" Policy
  - Borrowed from: Prof. Allan Borodin & Prof. Nisarg Shah (citation!)

  1. Partial marks for viable approaches
  2. Zero marks if the answer makes no sense
  3. 20% marks if you admit to not knowing how to solve

- 20% > 0% !!

Course Timeline

- (Approximate dates)
- ≈ Sep 27: HW1 posted
- ≈ Oct 18: HW1 due
- ≈ Oct 25: project proposal due (highly recommended to start earlier!)
- ≈ Nov 8: HW2 posted
- ≈ Nov 22: HW2 due
- ≈ Dec 6: Project presentations and project reports due

Course Project

- How?
  - Groups of 1-3
    - Larger groups are better
    - Find partners early, but maybe after the enrollment stabilizes
- What?
  - **Empirical:** Quantitative analysis of algorithms presented in class (or your own) using simulations or real data
  - **Theoretical:** Prove new observations about the algorithms or design new algorithms for a problem
  - **Ideal:** A bit of both

Project Topic

- From your own research area of interest
  - We'll introduce broad concepts that you may be able to apply to your own research area in order to find a project topic
  - E.g., fairness, allocation efficiency, preference elicitation, …
- From the course
  - I'll mention some open problems as we go along
  - You can also study realistic variants of problems that we see in class

Course Project: Timeline

- Find partners and think about a project idea
- Submission 1: Project proposal
  - Ideally 1 page but up to 2 pages excluding references
  - Outline of the idea, prior work, reasonable goals
- Mid-project meetings
  - Optional, 1-1 with me, 30-minute
- Class presentations
- Submission 2: Final project report
  - Up to 5 pages excluding references and appendix
  - Focus on quality academic writing

Introductions

Brief Introductions

- What to say?
  - Which program?
  - Which year?
  - Who are you working with (if any)?
  - What is your area of interest (if any)?
  - Anything else you'd like to share

Overview of the Course

Two Main Themes

Computational Social Choice

Algorithmic Fairness

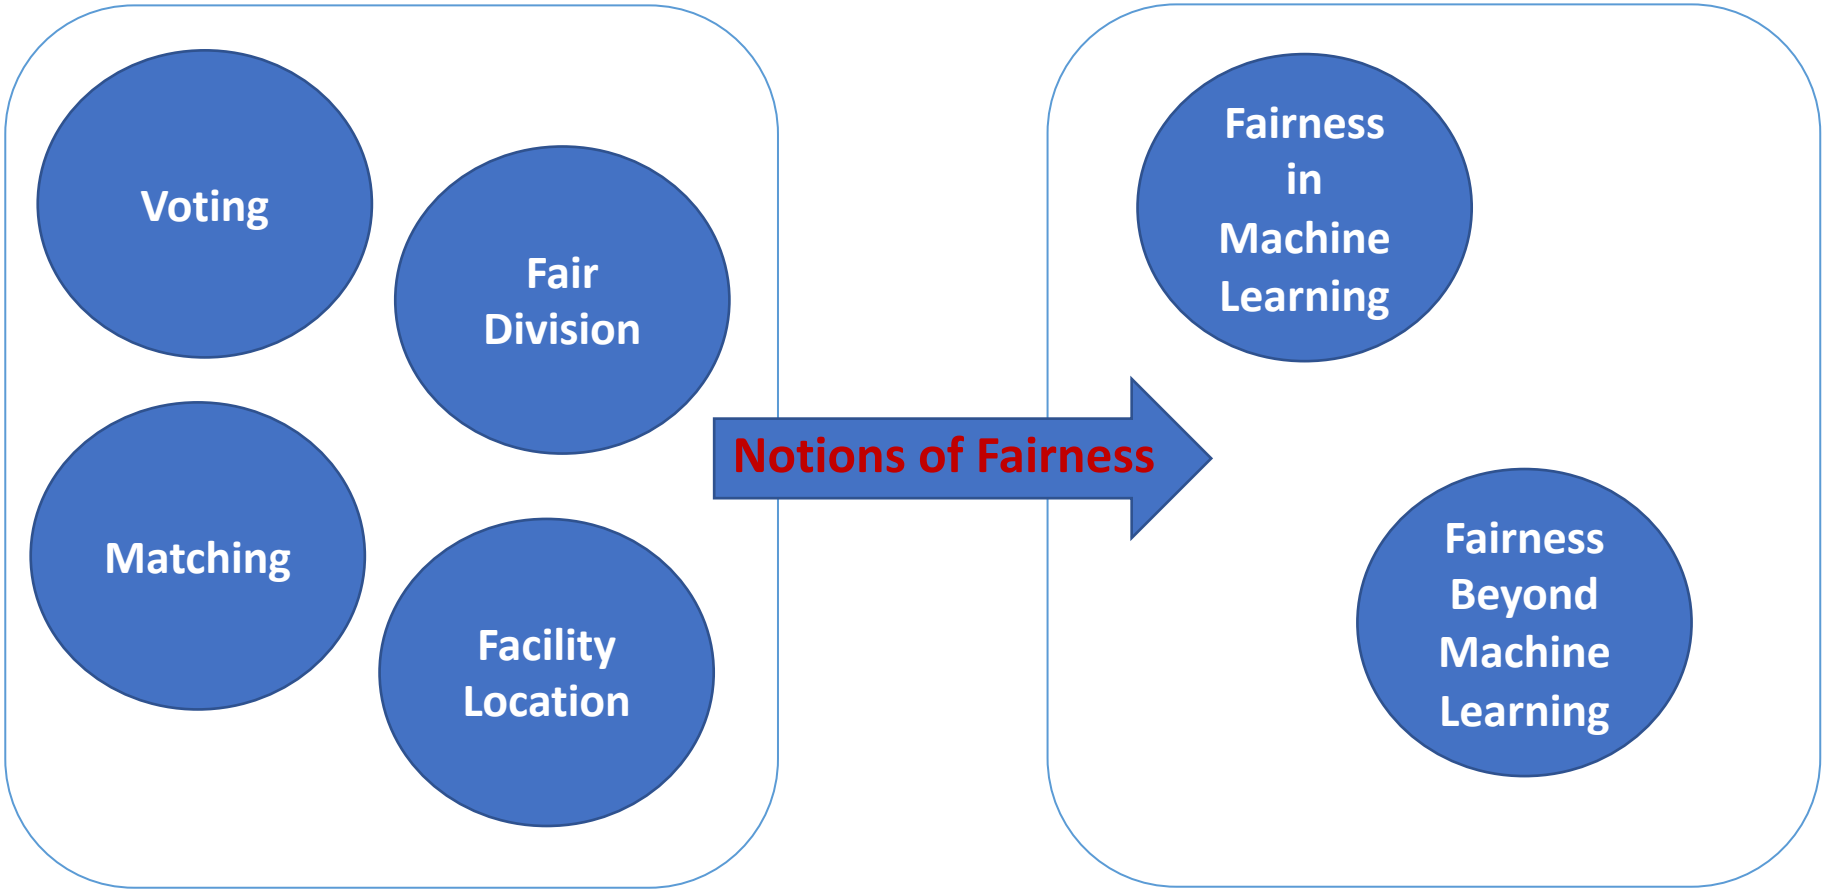

Voting

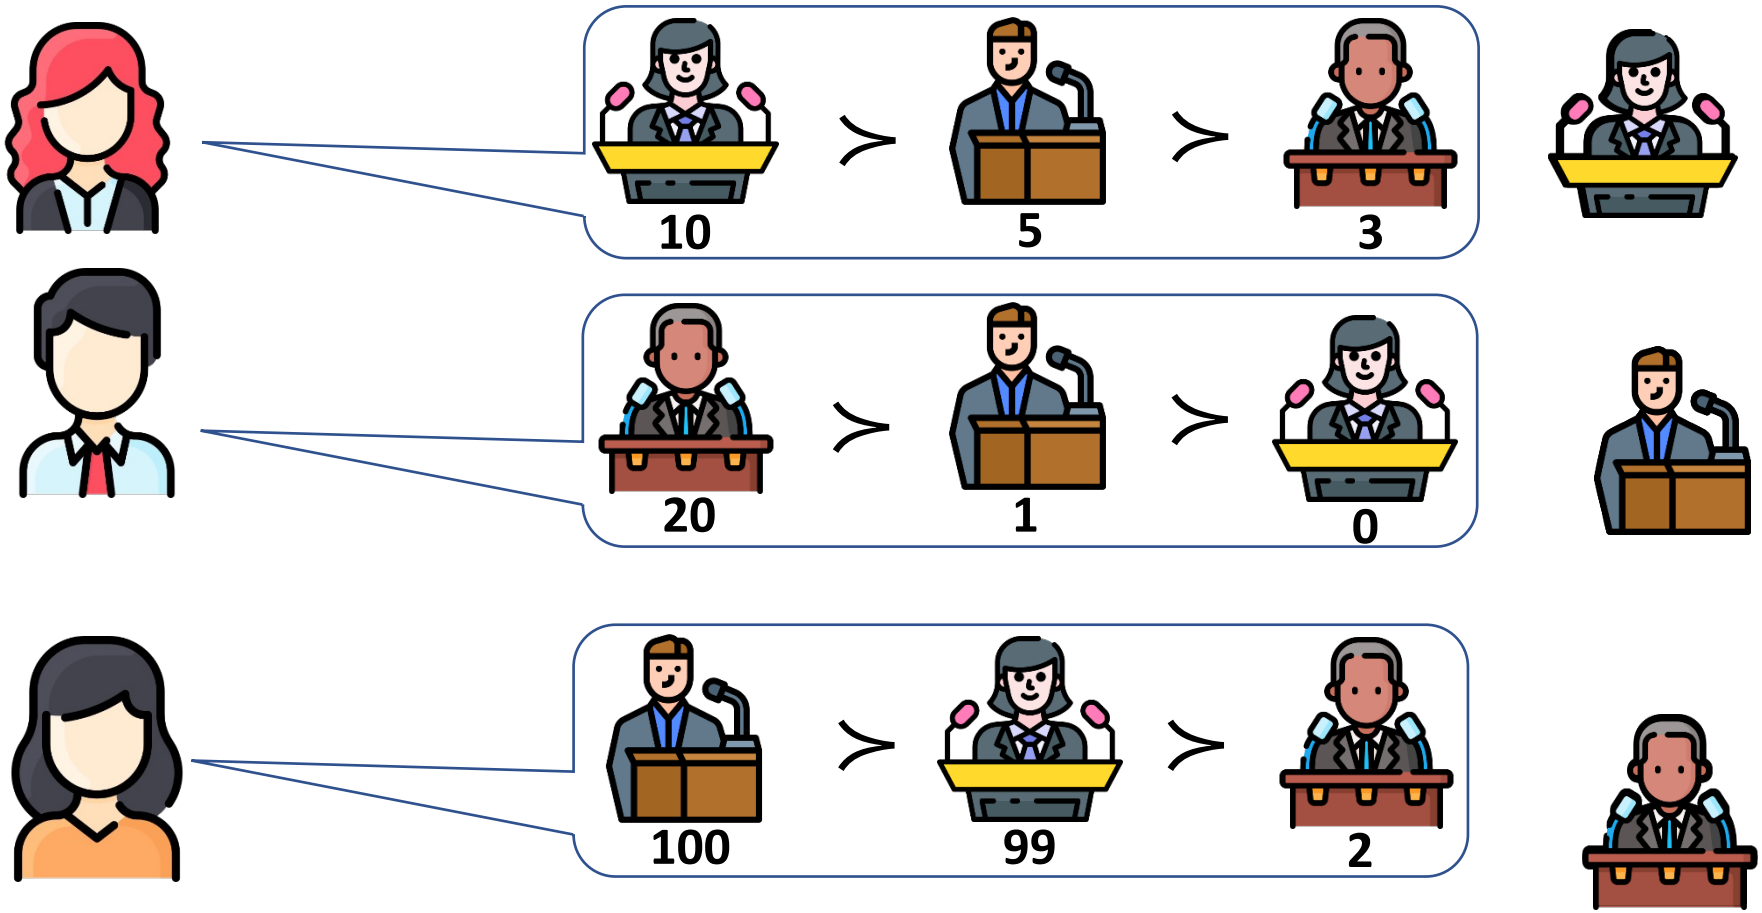

Q: How is the winner?

Example: Participatory Budgeting

Example: Human-AI Alignment

…AI alignment involves ensuring that an AI system's objectives match those of its designers…
(wikipedia)

Example: Human-AI Alignment

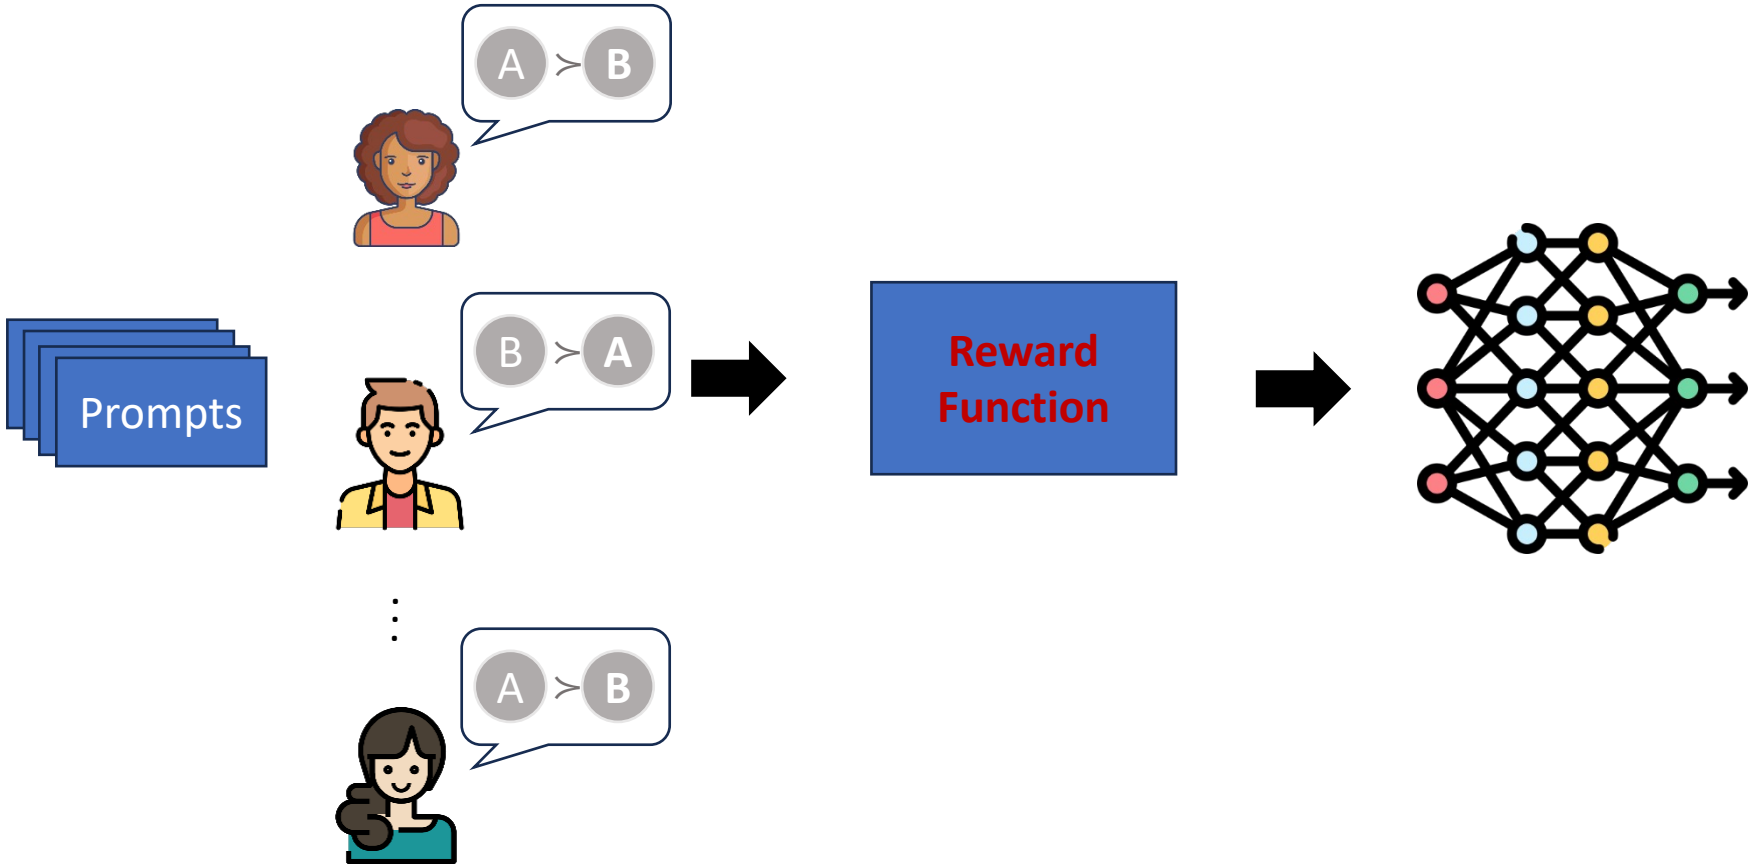

Fair Division

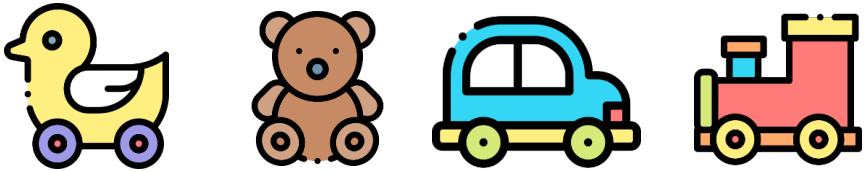

I prefer strawberry

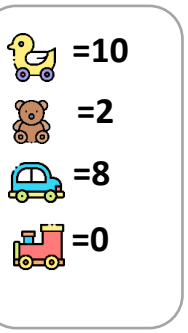

I love chocolate

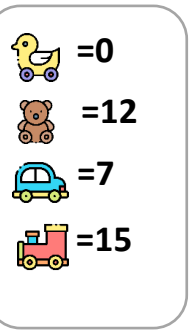

I don't like strawberry

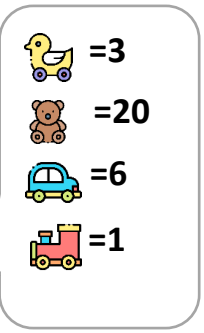

Q: What is a fair allocation?

Example: Citizens' Assemblies

Population

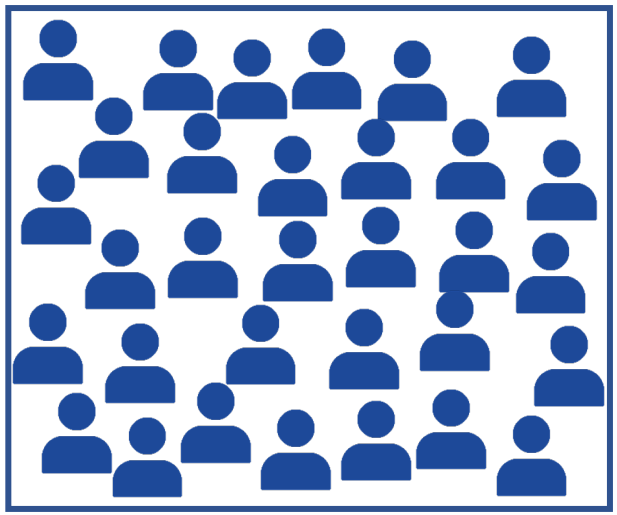

Panel

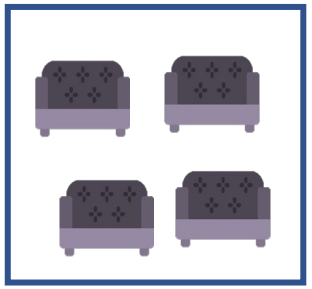

Facility Location

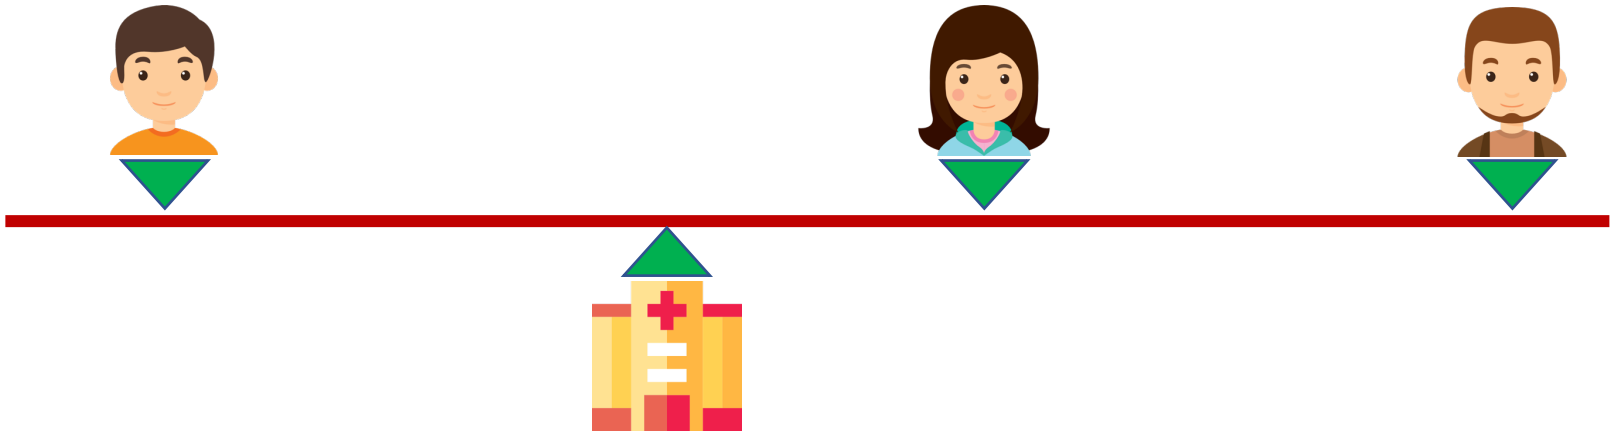

Q: What is the optimal hospital location?

Q: If we decide to choose the optimal location, will the agents really tell us where they live?

Image Courtesy: Freepik

Algorithmic Fairness

Fairness in Clustering

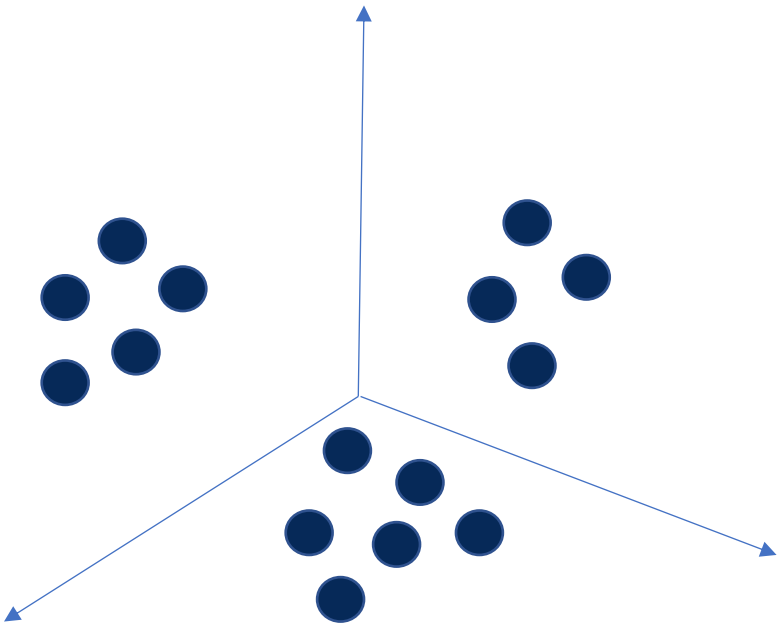

Q: What is a fair clustering solution for any subset of data points?

Fairness in RL

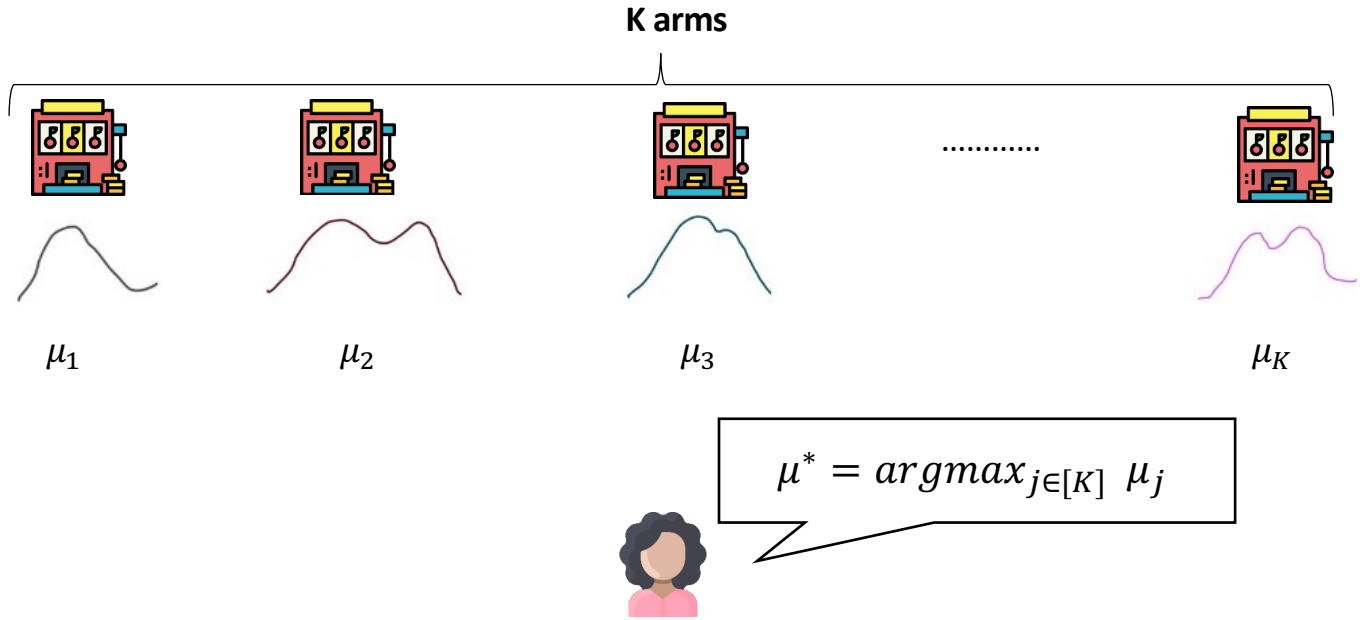

Exploration vs Exploitation

Fairness in RL

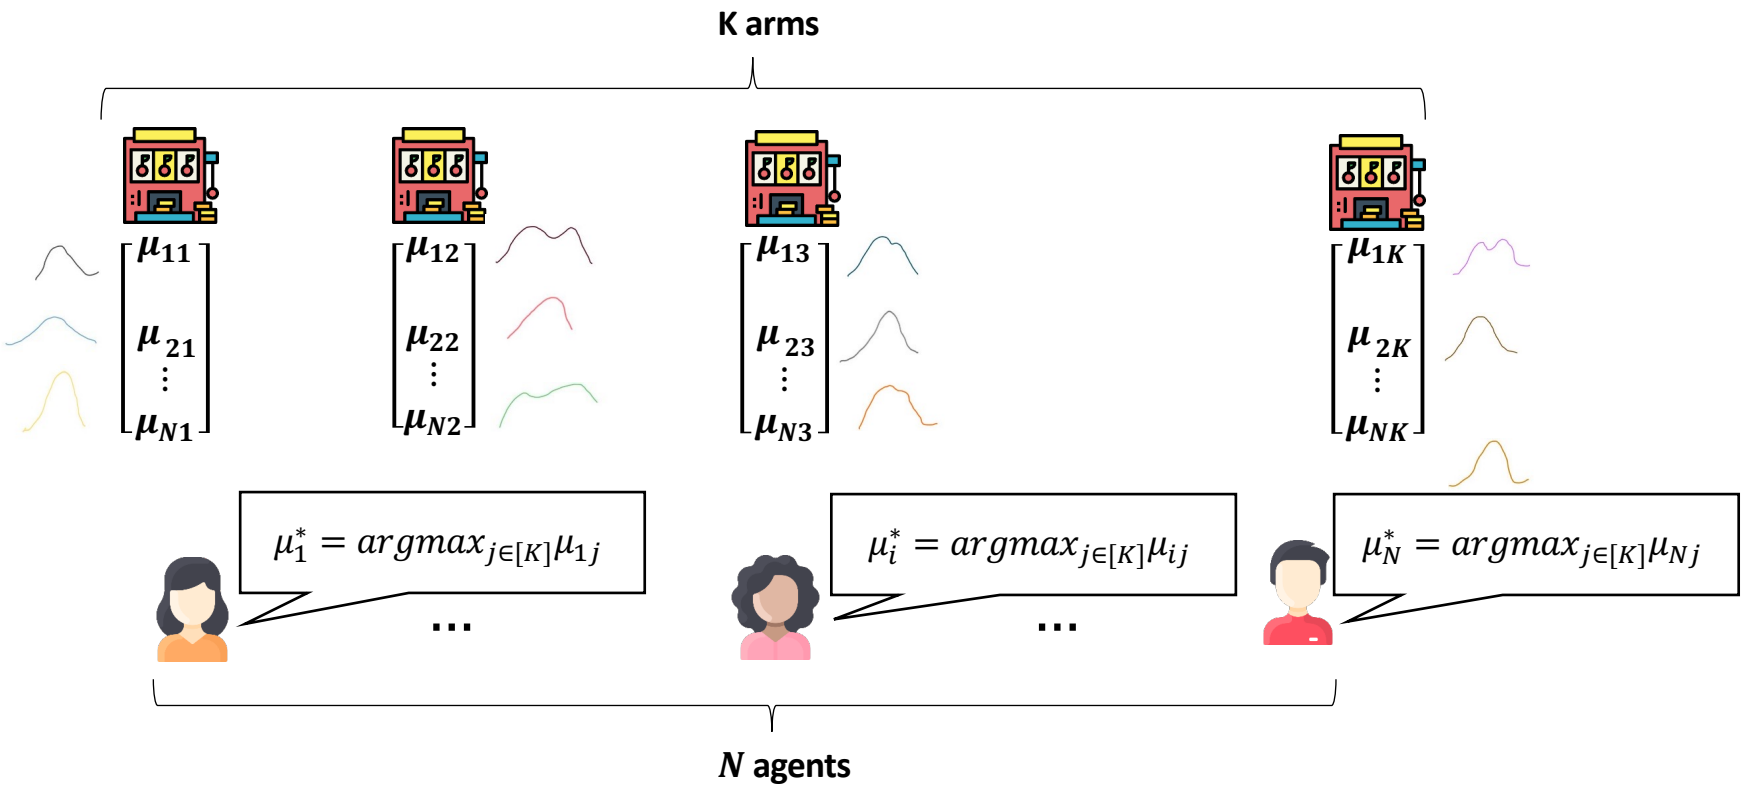

Q: What is a fair policy?

Peer Review

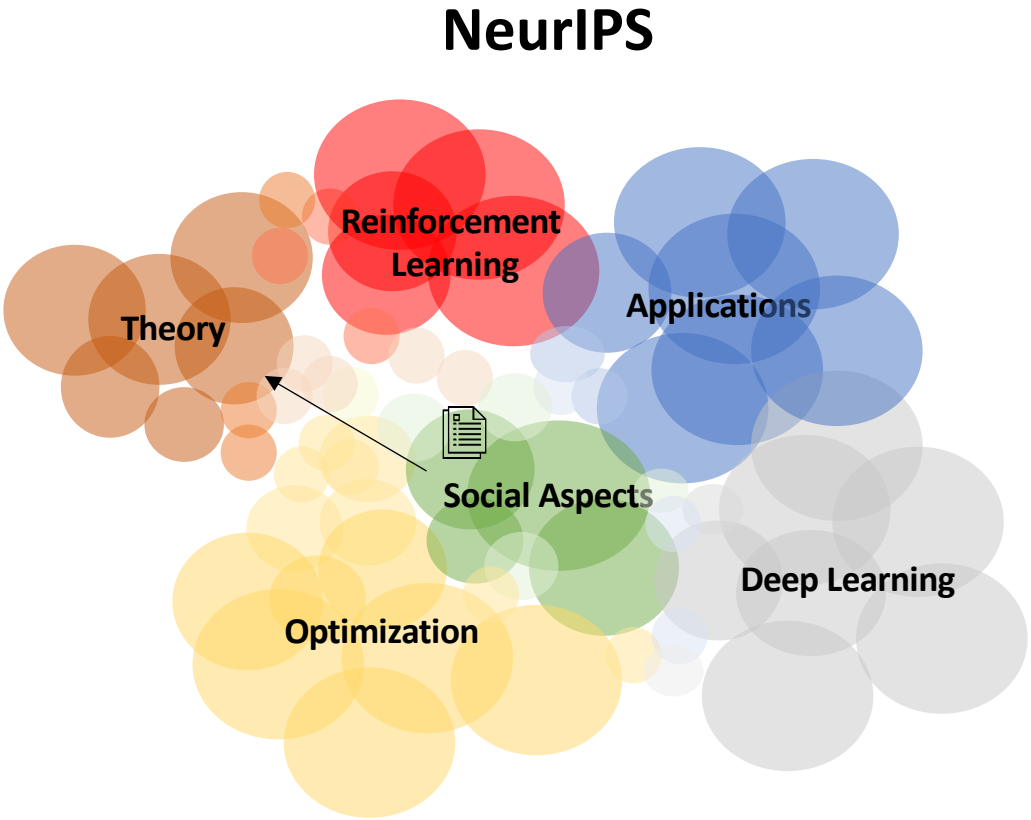

FAccT

COLT

ALT

Peer Review

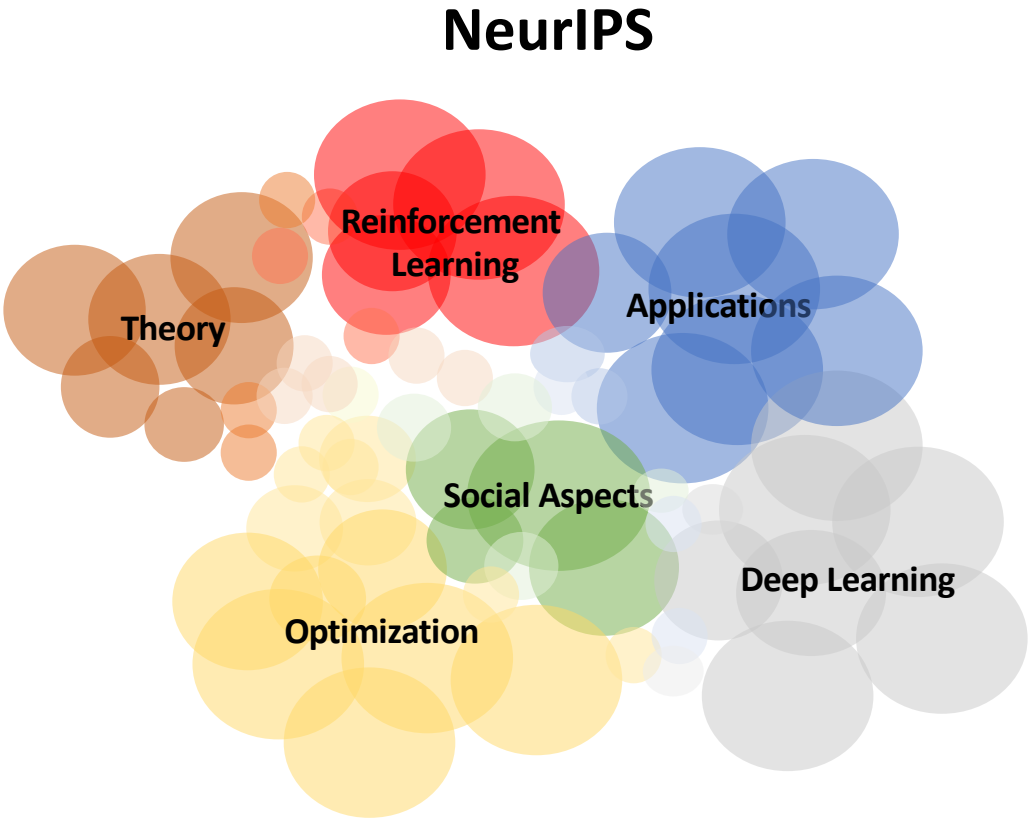
NeurIPS
Reinforcement Learning
Theory
Applications
Social Aspects
Deep Learning
Optimization

Peer Review

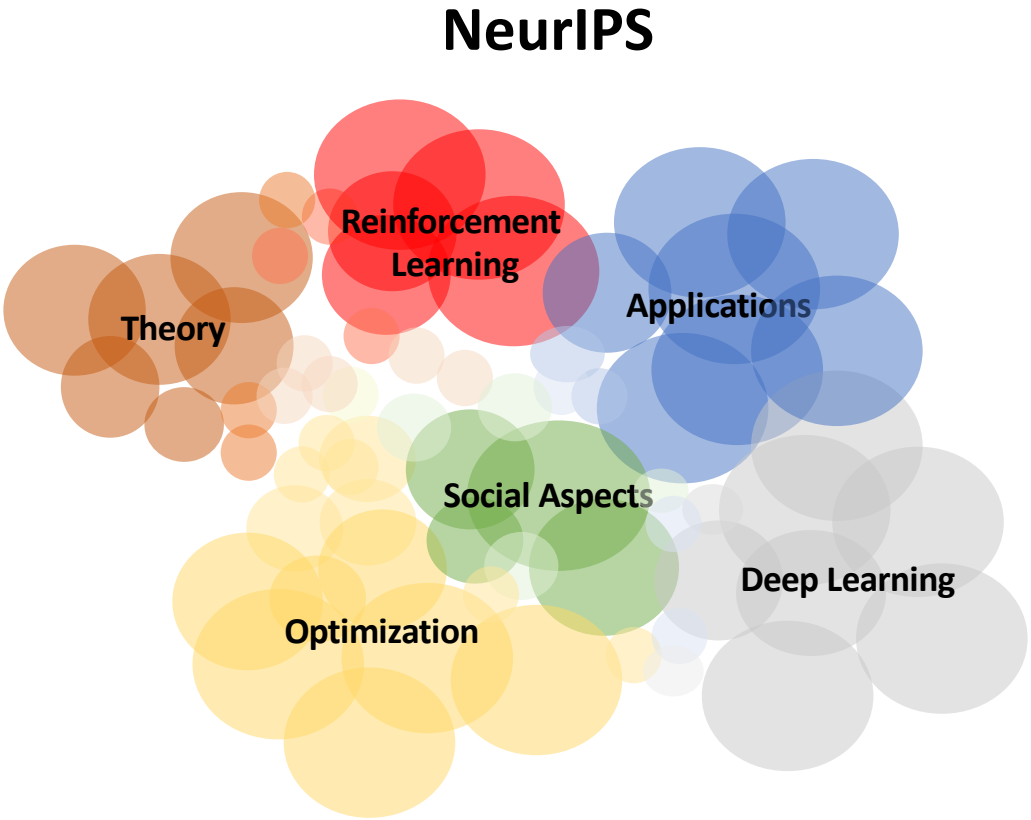

Q: Is it possible to create a reviewing procedure that prevents any subcommunity from benefiting by withdrawing from a large conference?