Statistical Foundations of Virtual Democracy

Anson Kahng, Min Kyung Lee, Ritesh Noothigattu, Ariel Procaccia, and Alex Psomas

ICML 2019

Automating Ethical Decisions

Donors

GIANT EAGLE

WHOLE FOODS MARKET

EAST END

EINSTEIN BROS BAGELS

TRADER JOE'S

412 • FOOD RESCUE

Recipients

GREATER PITTSBURGH community food bank
A member of Feeding America™

THE SALVATION ARMY

MEALS ON WHEELS
deliver a difference

the YMCA

ALLEGHENY COUNTY HOUSING AUTHORITY

Automating Ethical Decisions

Donors

GIANT EAGLE

WHOLE FOODS MARKET

EAST END

EINSTEIN BROS BAGELS

TRADER JOE'S

412 • FOOD RESCUE

Recipients

GREATER PITTSBURGH community food bank
A member of Feeding America™

THE SALVATION ARMY

MEALS ON WHEELS
deliver a difference

the YMCA

ALLEGHENY COUNTY HOUSING AUTHORITY

How do you make this decision?

Which recipient **deserves** the food?

A Modest Proposal

Ask participants to cast a vote every time a decision needs to be made

Donor: GIANT EAGLE **Type of donation:**

WHOLE FOODS MARKET

THE SALVATION ARMY

412 FOOD RESCUE

MEALS ON WHEELS deliver a difference

GREATER PITTSBURGH community food bank *A member of Feeding America™* > MEALS ON WHEELS deliver a difference > ALLEGHENY COUNTY HOUSING AUTHORITY > THE SALVATION ARMY > the YMCA

Issue: we must consult participants <u>every time</u> a donation occurs!

Idea: what if we could **predict** how people would vote?

Virtual Democracy

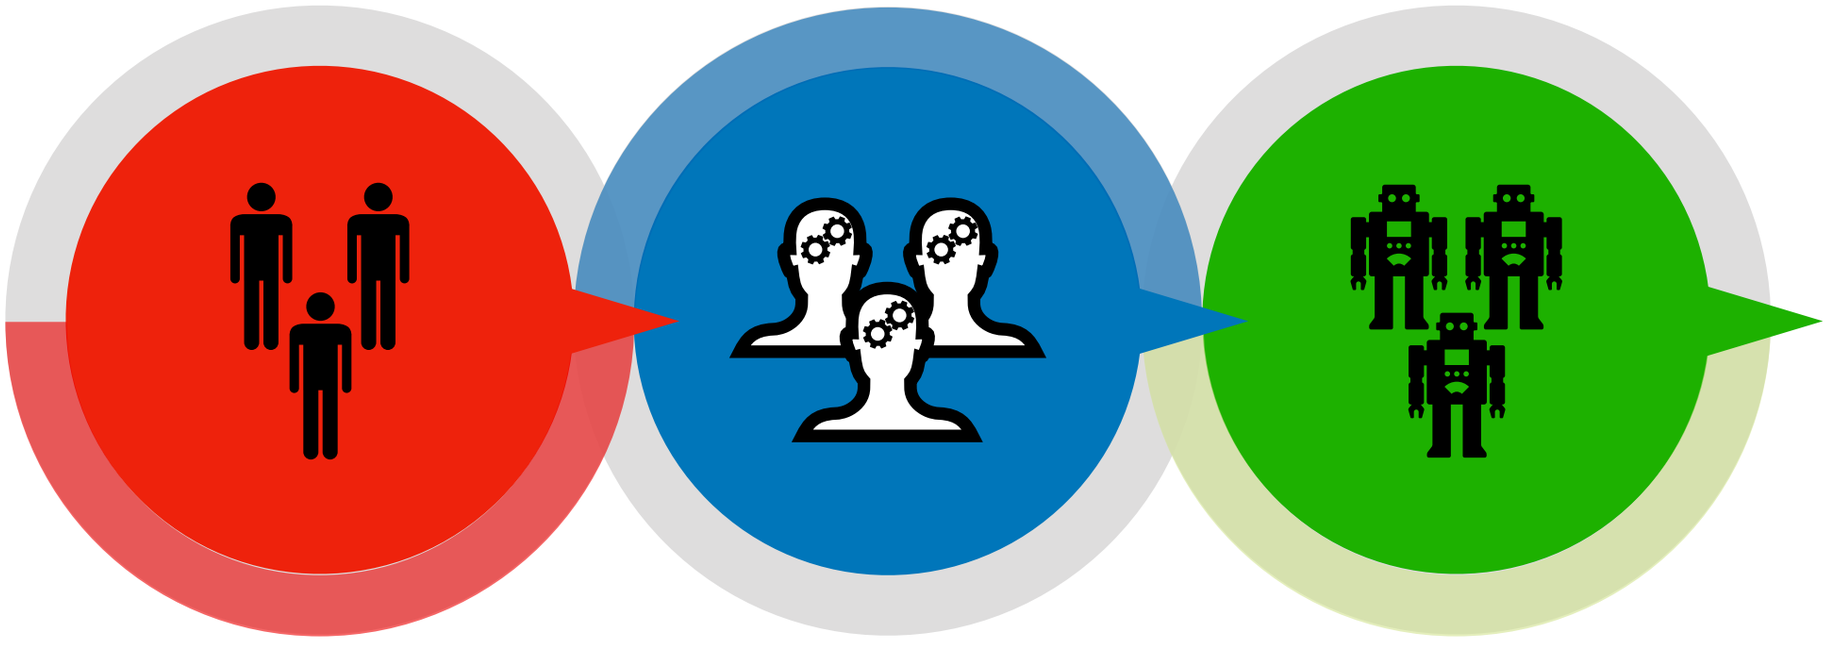

Data Collection Learning Aggregation

“Learn models of people, and let the models vote”

Data Collection

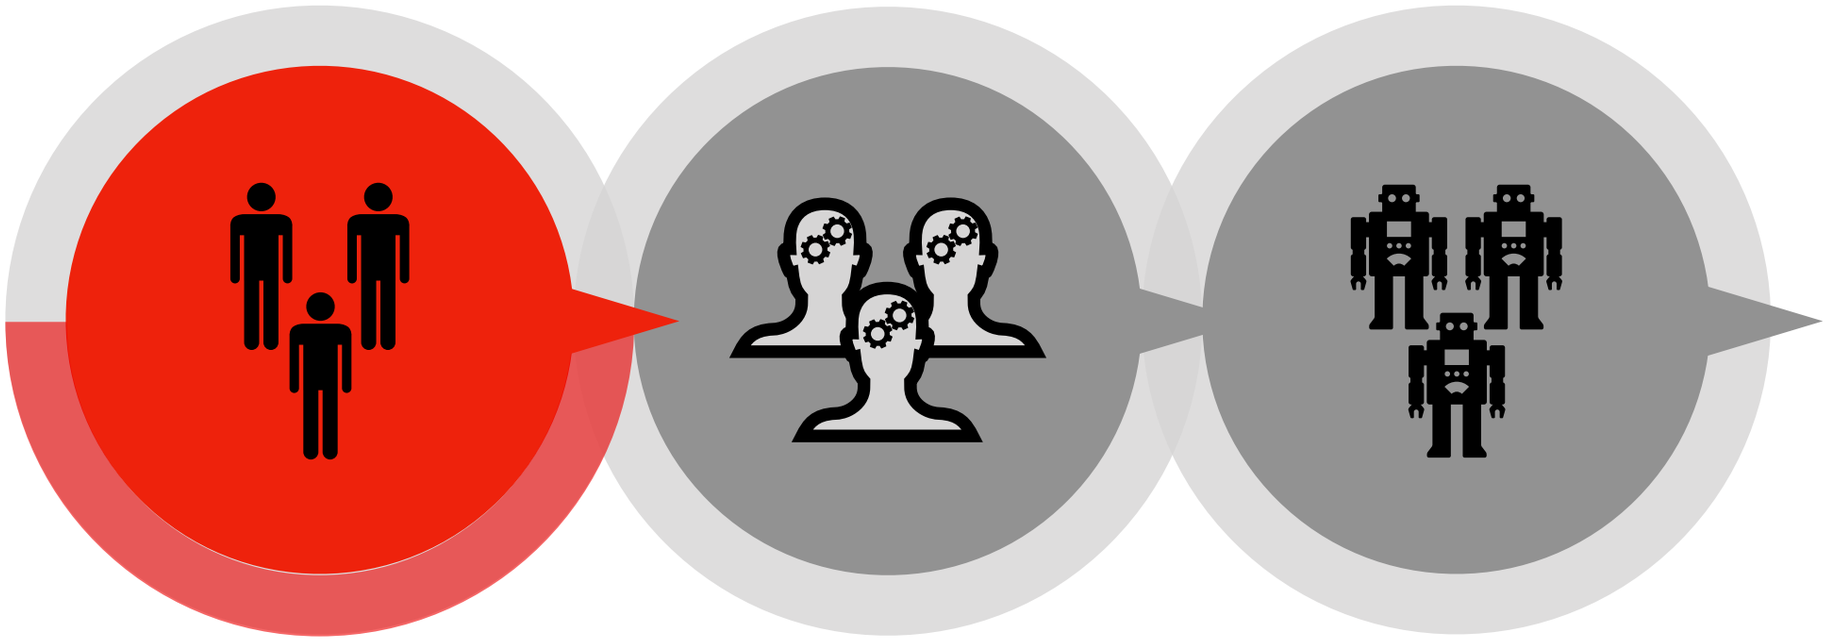

Data Collection Learning Aggregation

Use features identified by Lee et al. (2017) to collect pairwise comparisons of potential recipients

LEARNING

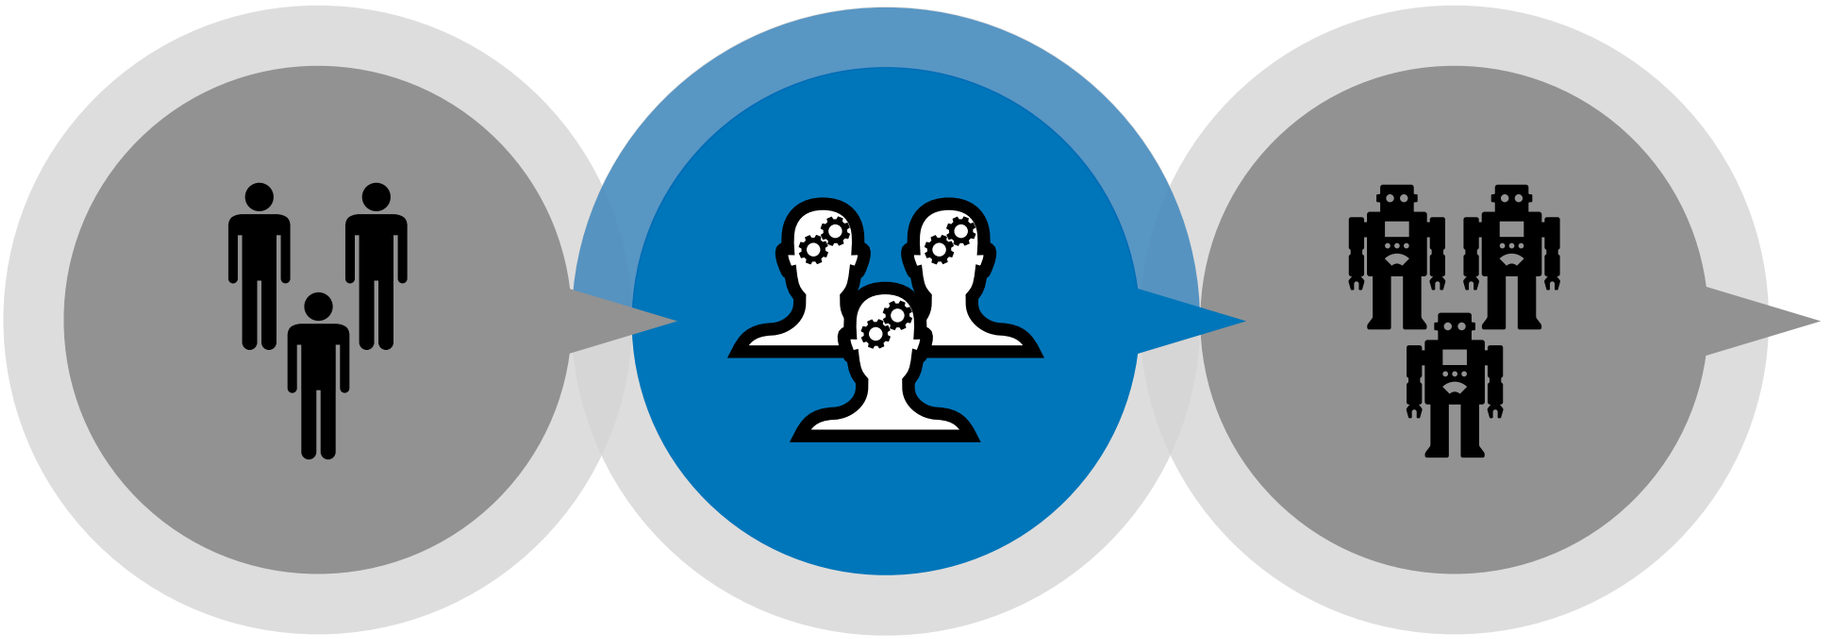

Data Collection **Learning** Aggregation

Learn models of participants that capture their reported preferences on pairwise comparisons; let models vote

Aggregation

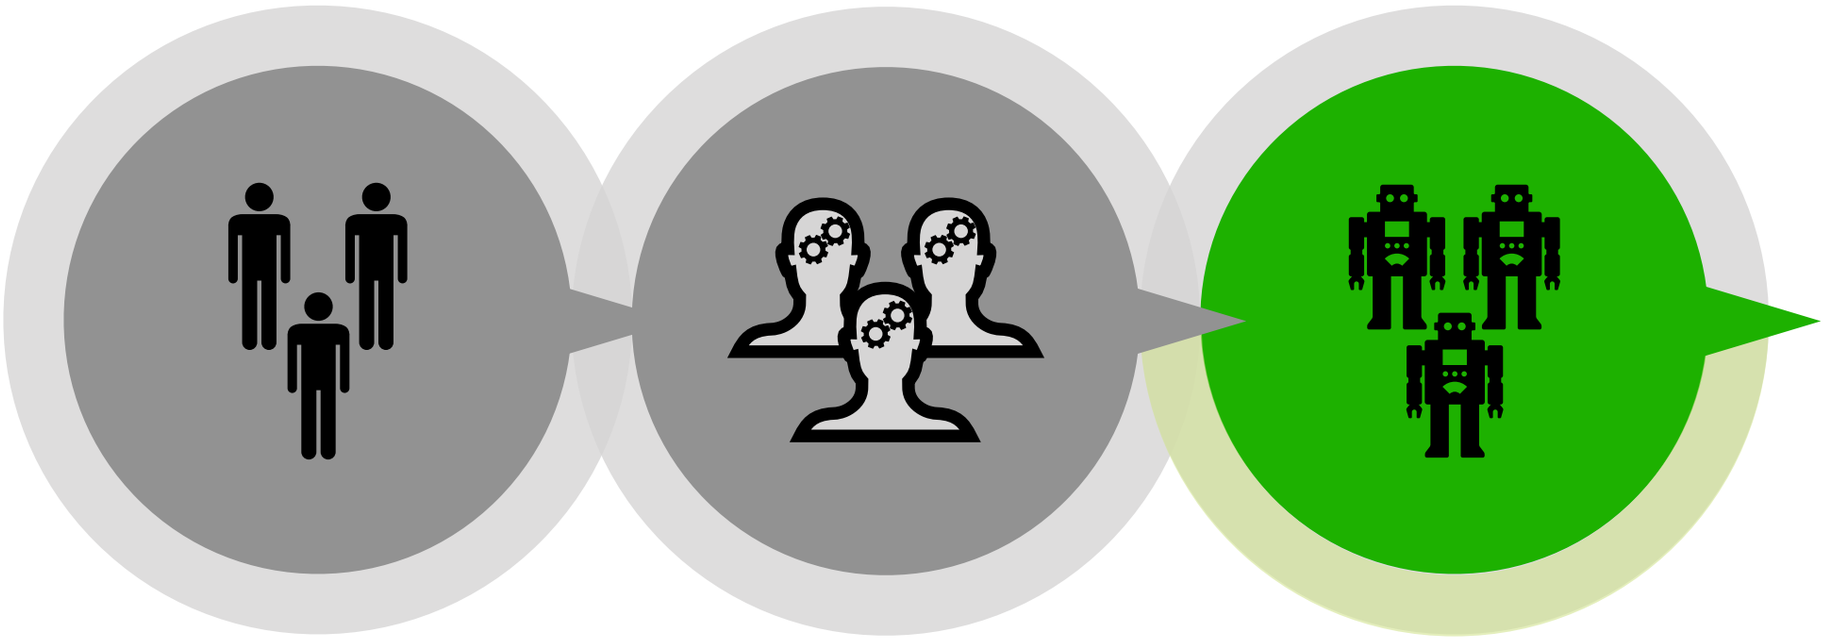

Data Collection Learning **Aggregation**

How do we aggregate these votes?

Aggregation

Fundamental question in virtual democracy:

Which voting rule should we use to aggregate votes?

Desideratum: robustness to machine learning errors

We want voting rules that are likely to output the **same** result on both true underlying preferences and noisy votes

Theoretical Results

Theorem: Borda Count is robust under Mallows noise

> If the difference between the true Borda scores of two alternatives is small, then the probability that Borda swaps them in the noisy ranking is exponentially small

Theorem: PMC rules are not robust under Mallows noise

> There always exists a profile with an acyclic pairwise majority graph, but whose noisy profile has an acyclic pairwise majority graph with a different topological ordering

Theoretical Results

Theorem: Borda Count is robust under Mallows noise

"Use Borda Count for virtual democracy"

Theorem: PMC rules are not robust under Mallows noise

"Don't use PMC rules for virtual democracy"