NATIONAL ACADEMIES Sciences Engineering Medicine

TRB TRANSPORTATION RESEARCH BOARD

TRB Webinar: Equity in Artificial Intelligence Applications

May 22, 2023

11:00 AM – 12:30 PM

Steering Towards Equity: Harnessing Responsible AI for Inclusive and Sustainable Transportation Solutions

Dr. Kofi Nyarko

Professor, Electrical and Computer Engineering

Director, Center for Equitable AI/ML Systems

Morgan State University

CEAMLS MORGAN

CENTER FOR EQUITABLE AI & MACHINE LEARNING SYSTEMS

Overview of the Presentation

- The Current State of Transportation
- Understanding the Equity Issue in Transportation
- Role of AI in Addressing Transportation Equity
- The Concept of Responsible AI
- Responsible AI for Inclusive and Sustainable Transportation: Opportunities and Challenges
- Strategies and Best Practices for Harnessing Responsible AI
- Conclusion and Future Outlook

Current State of Transportation

- Increase in electric vehicles
- Increase in high-speed trains
- Potential for self driving cars

Background:
AI's Role in Transportation Solutions

- What Exactly is AI?
- Role of AI in transportation:
- Examples of AI in transportation:

Background: Sustainability Trend in Transportation Solutions

- What is sustainable transportation
- Role of AI in sustainable transportation
- Examples of AI in sustainable transportation

Understanding the Equity Issue in Transportation

- Accessibility
- Affordability
- Safety

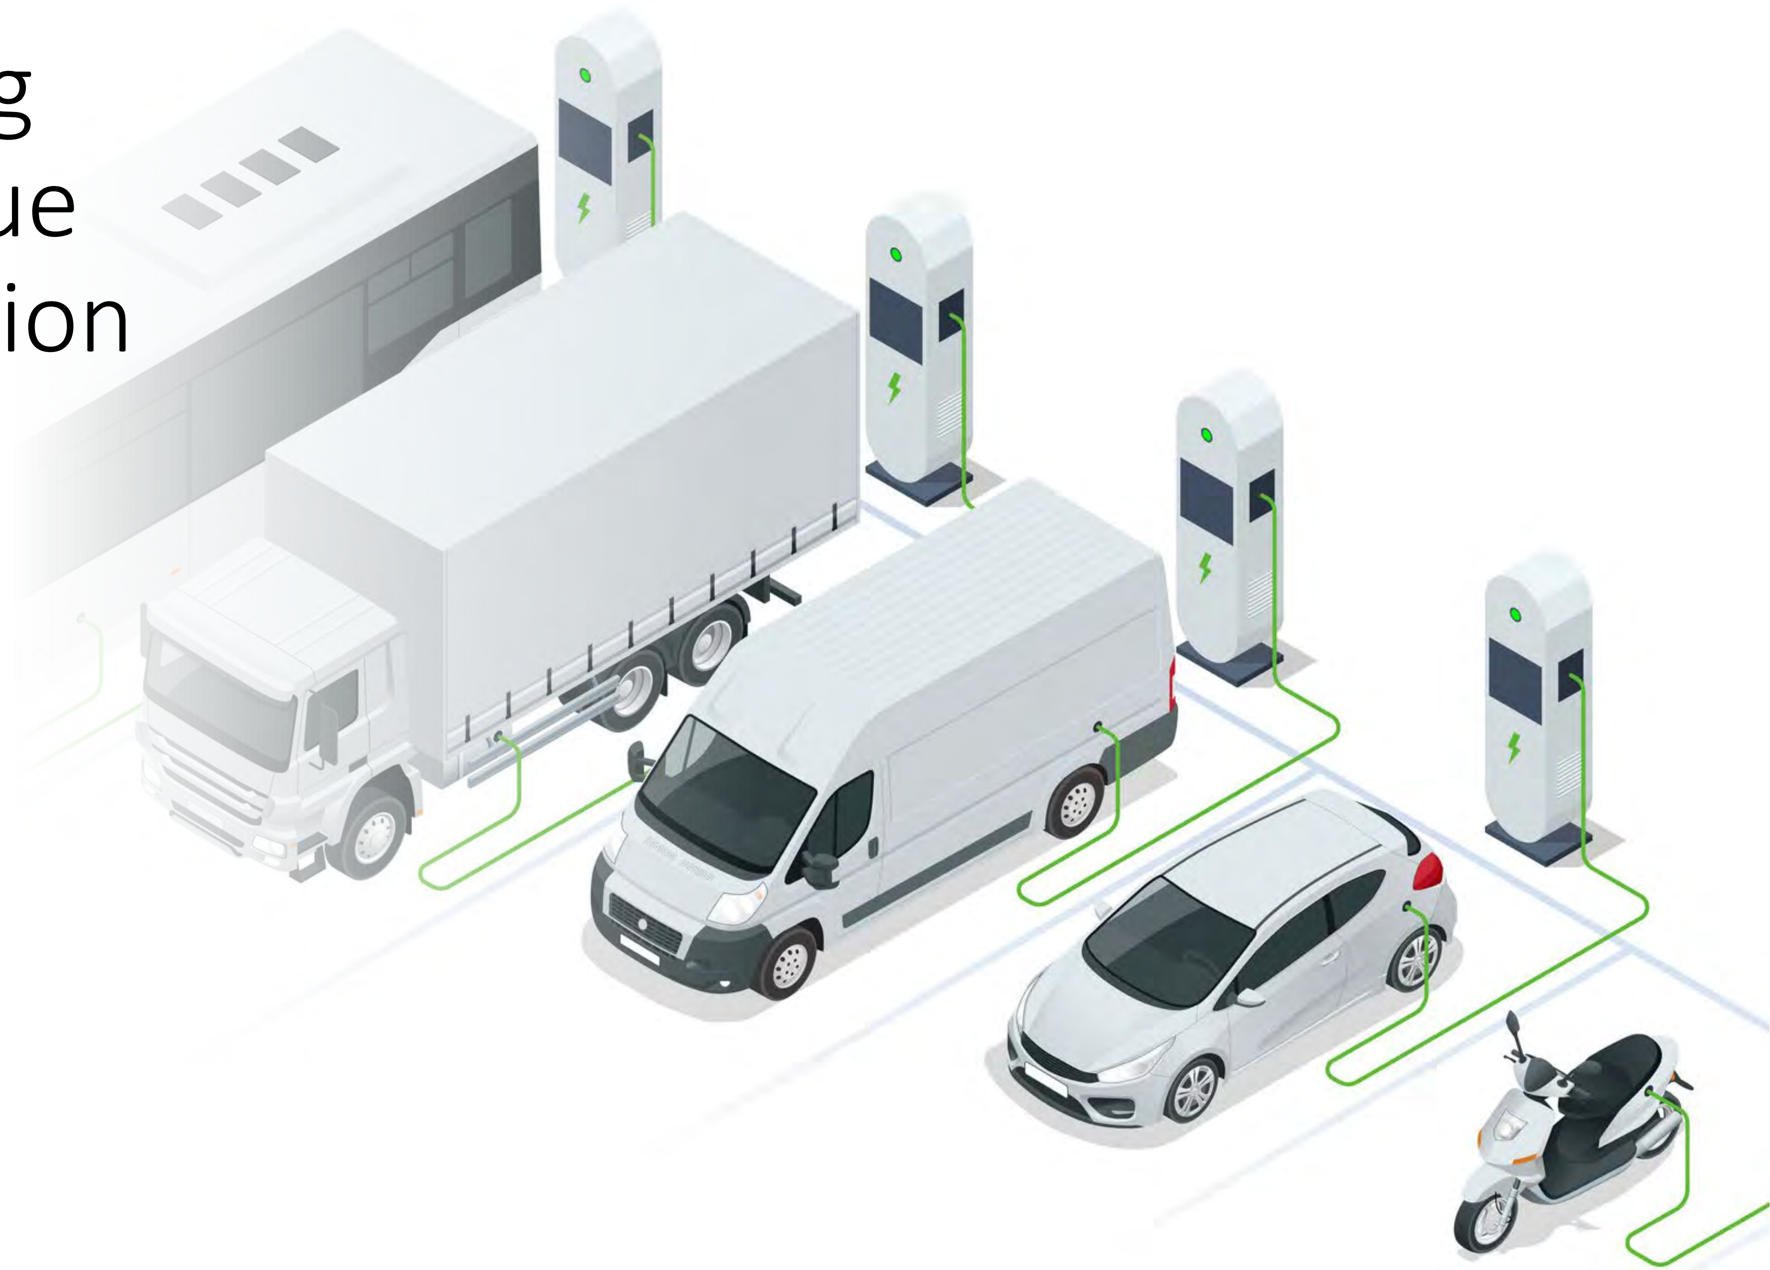

Understanding the Equity Issue in Transportation

- Certain communities are disproportionally impacted
- Rural areas
- Low-income urban areas
- Vulnerable groups

Understanding the Equity Issue in Transportation

- Case Study 1: Public transportation in rural areas
- Case Study 2: Inefficient public transportation in low-income urban areas
- Case Study 3: Accessibility and safety issues for vulnerable groups

Role of AI in Addressing Transportation Equity

- Optimized routing
- Demand-responsive transit
- Enhanced accessibility

Examples of Successful AI Applications in Enhancing Transportation Equity

- Example 1: Google's AI-powered Maps
- Example 2: Via Transportation
- Example 3: AI for Accessibility in Autonomous Vehicles

The Concept of Responsible AI

- What Does Responsible AI Mean?
- Key elements of Responsible AI
  - Transparency
  - Accountability
  - Fairness
  - Human oversight:

Importance of Ethics, Fairness, Transparency, and Accountability in AI

- Ethics in AI
- Fairness in AI
- Transparency and accountability in AI

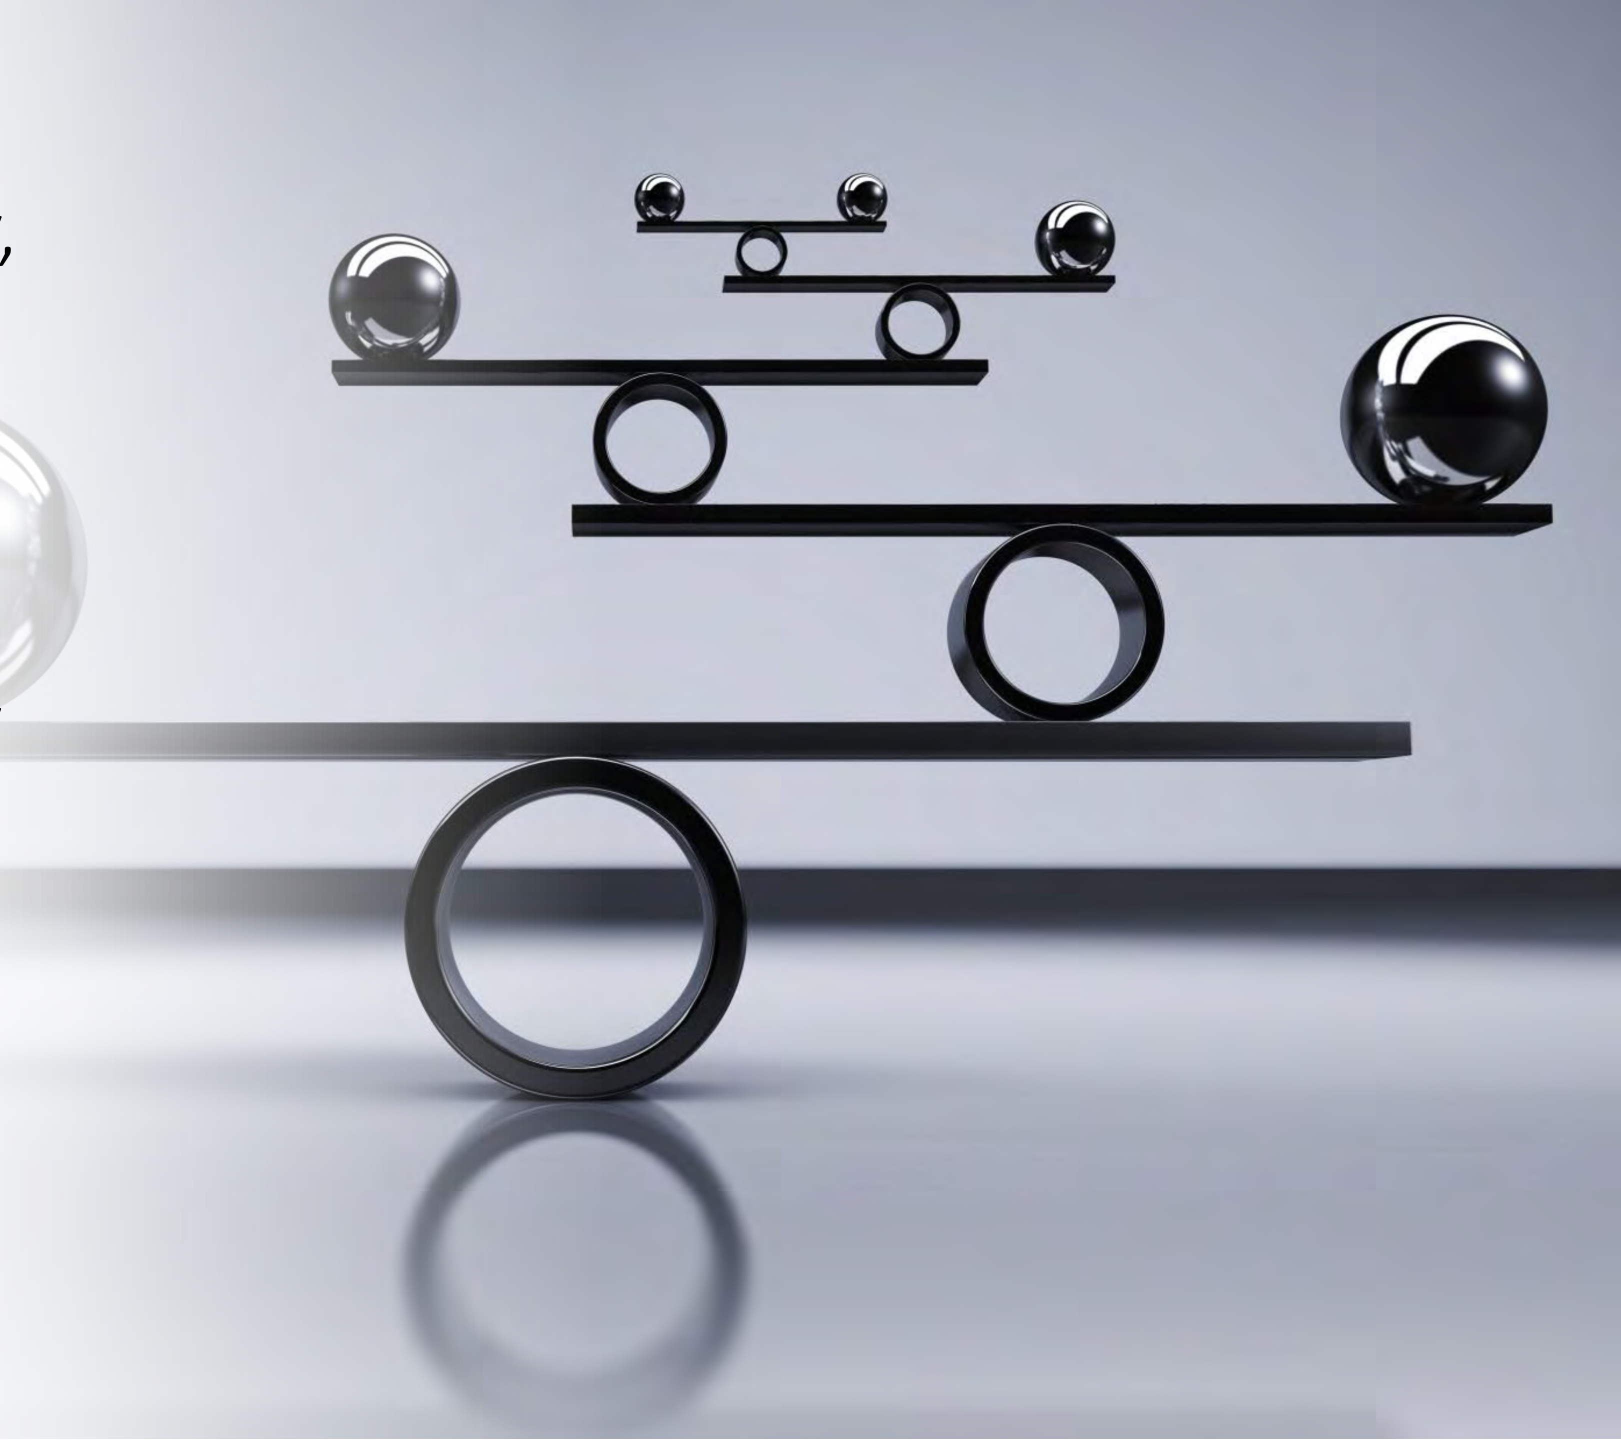

The Need for Responsible AI in Transportation Solutions

- Addressing bias and discrimination in transportation
- Ensuring transparency and accountability in AI-driven transportation
- Enhancing human oversight and control:

Responsible AI for Inclusive and Sustainable Transportation: Opportunities and Challenges

- Addressing biases in data
- Enhancing accessibility and affordability
- Promoting sustainability

Case Studies of Responsible AI's Success in Enhancing Transportation Equity

- Case Study 1: Citymapper
- Case Study 2: Moovit
- Case Study 3: Optibus

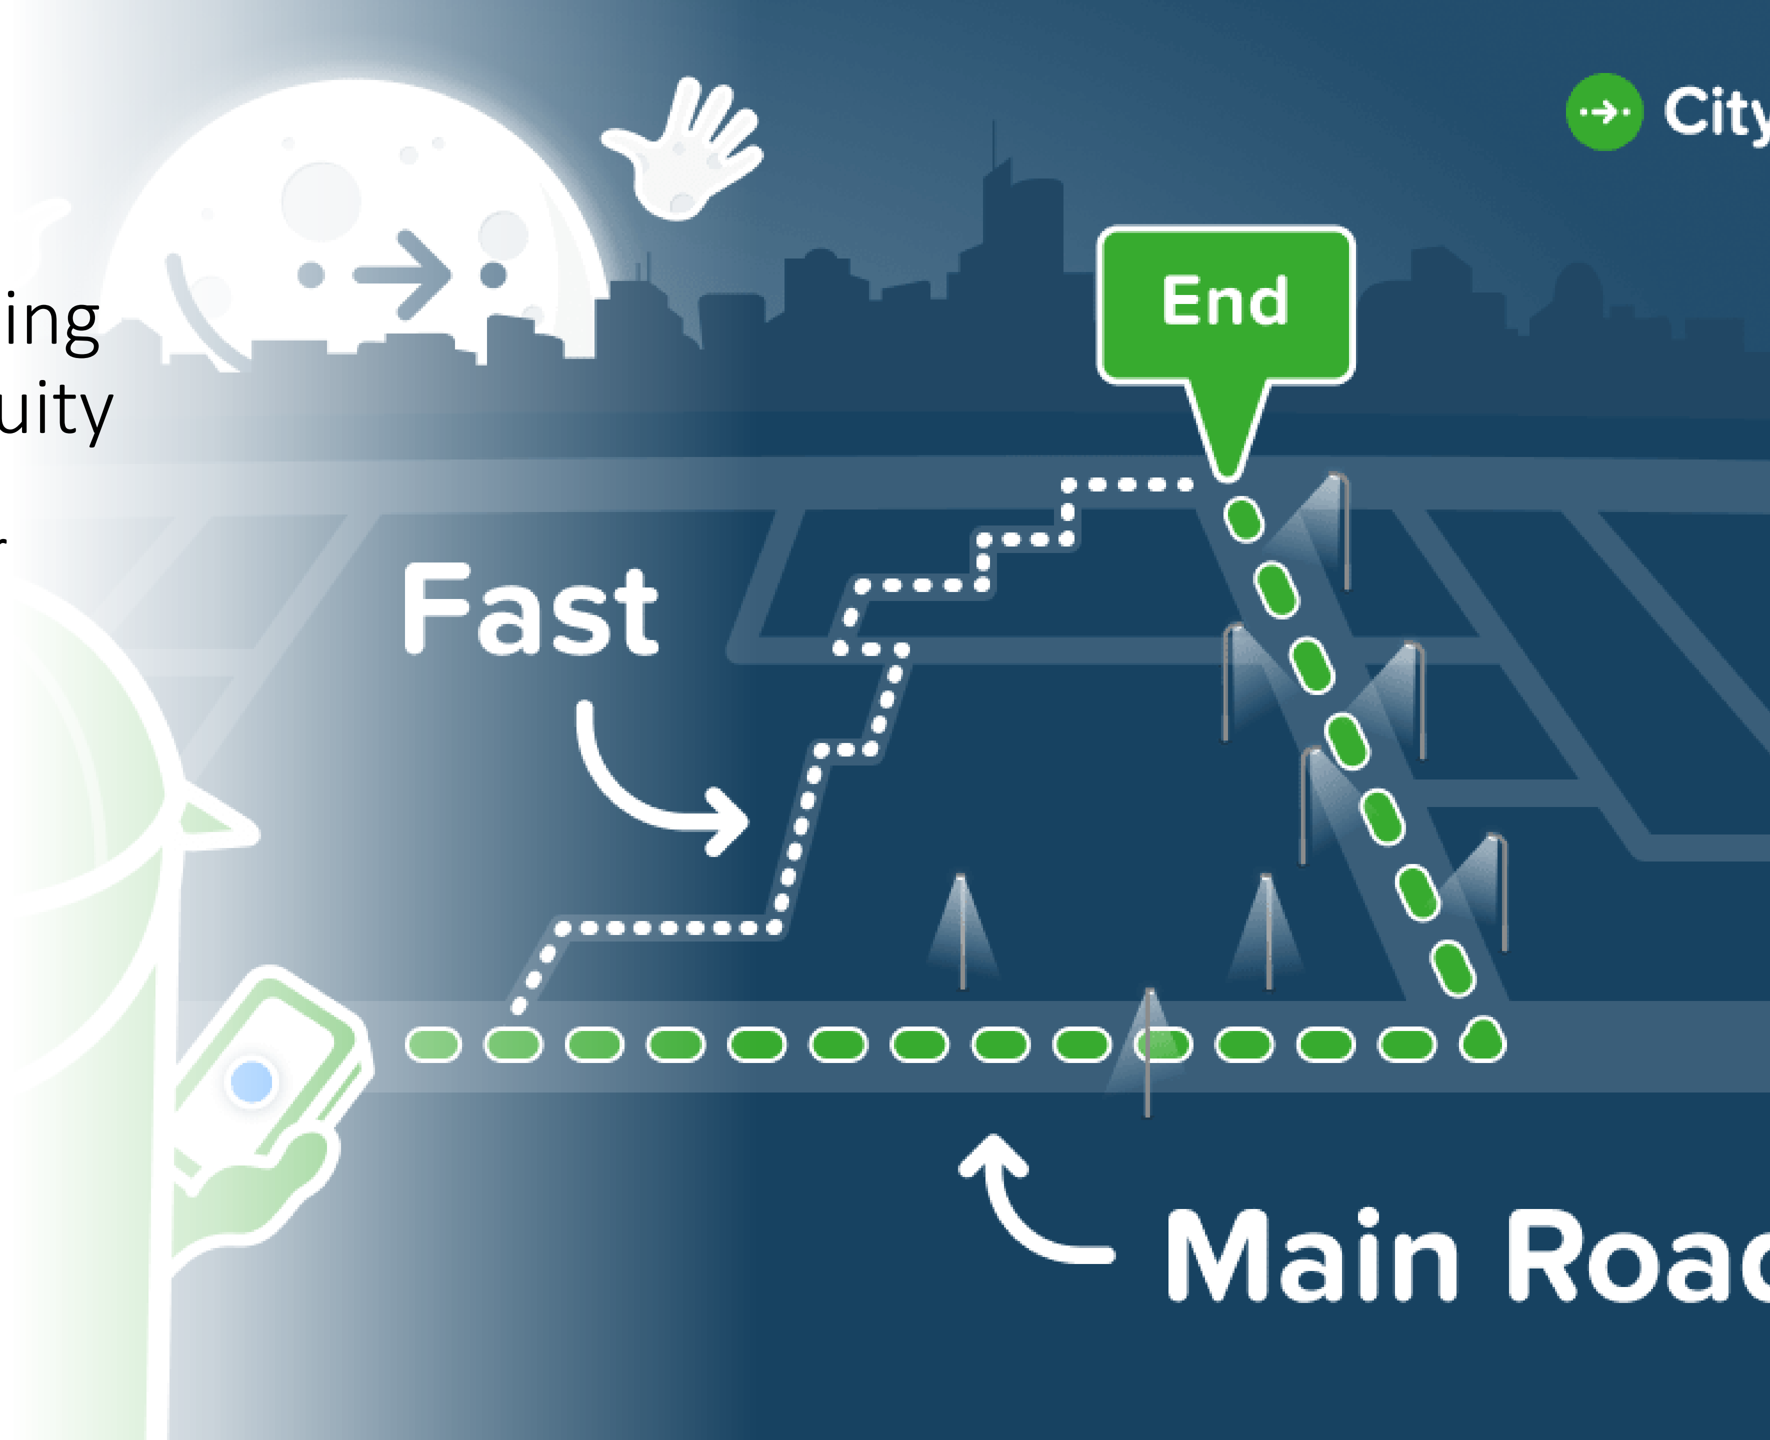

Potential Challenges in Implementing Responsible AI

- Biases in AI training data
- Lack of diverse representation
- Technical complexity and lack of transparency

Strategies and Best Practices for Harnessing Responsible AI

- Community involvement
- Diverse data sets
- Continual monitoring and adjustment

Best Practices in the Industry to Ensure Responsible AI Deployment

- Transparency in AI operations
- Accountability mechanisms
- Ethical guidelines and standards

Conclusion and Future Outlook

- Recap of the role of AI in transportation
- Recap of the importance of Responsible AI
- Recap of strategies for harnessing Responsible AI

The Potential Future of Transportation with Responsible AI

- Greater efficiency and accessibility
- Enhanced sustainability

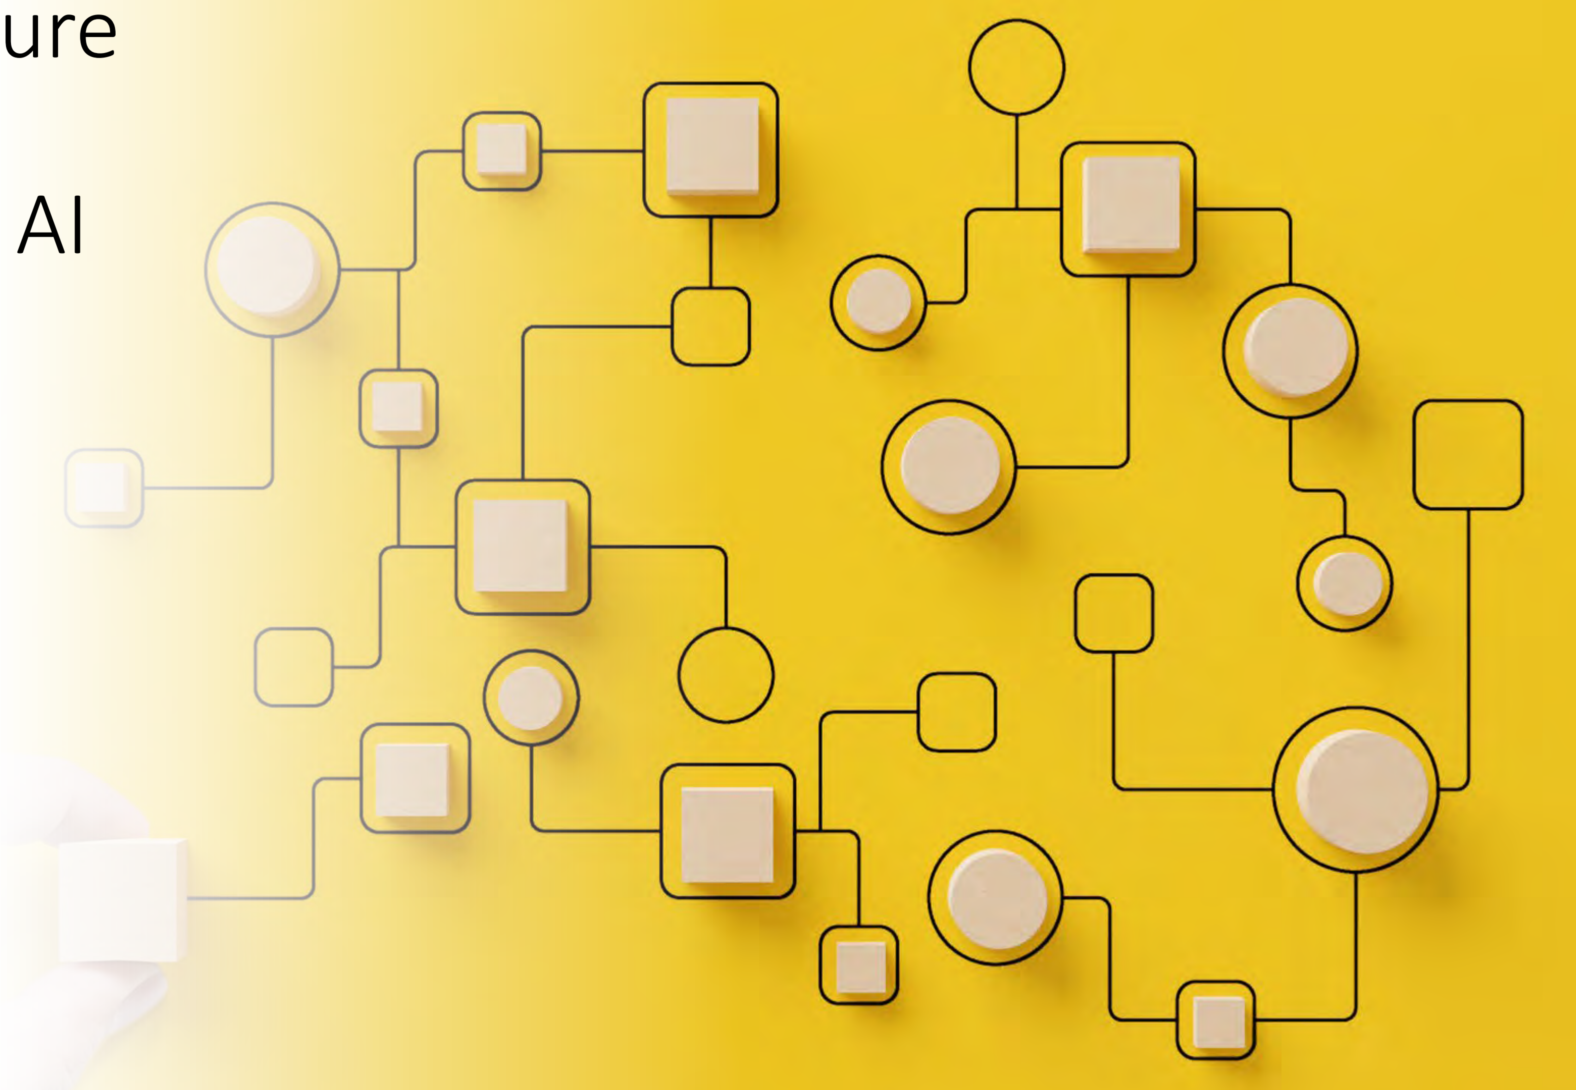

Final Thoughts on the Importance of Equity and Sustainability in Transportation Solutions

- The importance of equity:
- The role of AI in promoting equity and sustainability
- The importance of continued vigilance and commitment

Q&A Session

Building an Equitable and Intelligent Logistics Systems with Drones in Rural Areas

Dr. Ziping Wang

Associate professor of Information Science & Systems
Morgan State University

TRB Webinar: Equity in Artificial Intelligence Applications
May 22, 2023

Overview of the Topics

- Solving problems with equity
- Equity in problem-solving techniques

Delivery Experience in Rural Areas

- A real story during the pandemic
  - Elderly neighbors
  - On Daily Rx
- Zero accessibility to
  - prescription delivery (e.g., Walgreens)
  - grocery delivery (e.g., Amazon, InstaCart)
  - Food delivery (e.g. DoorDash)
  - ...

Same Day Delivery Accessibility

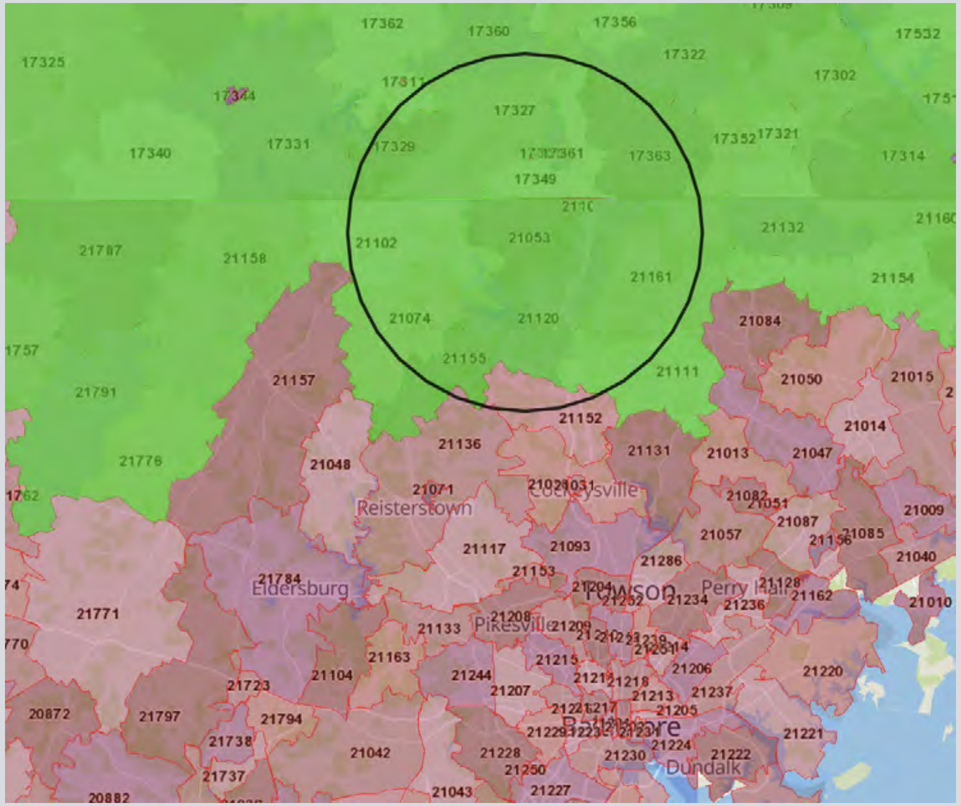

Amazon same-day delivery <u>unavailable</u> in Maryland's northern border areas and Pennsylvania's southern border areas (marked in green)

Population Density in Rural Areas

Zip codes are selected from the previous figure

	Rural areas						Suburban areas			Urban areas		
Zipcode	17327	17355	17361	21053	21074	21102	21093	21286	21204	21218	21213	21224
Population Density persons/sqml)	176	408	*1719*	147	380	277	1795	2290	3306	12259	8911	5043

- Rural areas usually have low population densities but large lands.
- Unique geographic features make truck delivery pricey and time inefficient.

Data source: www.mapszipcode.com

Voice from Residents in Rural Areas

A3. How would you rate the overall availability of Same Day delivery service (or Rapid delivery service in general) in your area?

- (•) Poor
- ( ) Fair
- ( ) Average
- ( ) Good
- ( ) Excellent
- ( ) Not sure

A4. Do you think Same Day delivery service (or Rapid delivery service in general) should be available to people wherever they reside?

- ( ) Not at all
- ( ) A little
- ( ) Rather
- ( ) Much
- (•) Very much

Data source: Ongoing survey conducted by the project team

Solving Problems

- Building equal delivery service for rural residents
  - How to measure equity?

Equity Measure

One possible formulation of the *system-adjusted equity index*

$$I_e = f(I_t, I_a, I_d)$$

where

$I_e \in [0,1]$: System equity index;

$I_t \in [0,1]$: Sub-equity index of travel time-saving *opportunity*;

$I_a \in [0,1]$: Sub-equity index of environmental amenity;

$I_d \in [0,1]$: Sub-equity index of delivery accessibility;

Solving Problems

- Building equal delivery service for rural residents
  - How to measure equity?
- Drone as an option?
  - Foreign success case such as Ziplines in Africa
  - Last mile drone/robot delivery in urban areas
  - Truck/Drone delivery cost structure
  - Equity-based incentives from government

$$S(I_e) = \begin{cases} \xi > 0, & if\ I_e \geq E \\ 0, & otherwise \end{cases}$$

E is a threshold of the equity index, predetermined by the government

  - AI algorithms to such complex problems
    - Equity-based Vehicle Routing Problem with Drones (VRP-D)

Representative Rural Area

Hereford zone, Maryland

	Zipcode	Population	Population density (per sqml)	Housing units
Officially Hereford zone	21120	6967	165	2540
	21111	4903	140	1883
	21053	3305	147	1223
	21161	5409	113	2074
Part of those two zipcode area are in Hereford zone	21102	parital		
	21074	parital		
Total		**20,584**		**7,720**

The Hereford Zone is an area in Northern Baltimore County, Maryland, constituting 20% of all of the land in Baltimore County.

Data source: https://www.unitedstateszipcodes.org/; USDA 2017 Census

Customer Locations & Road Network

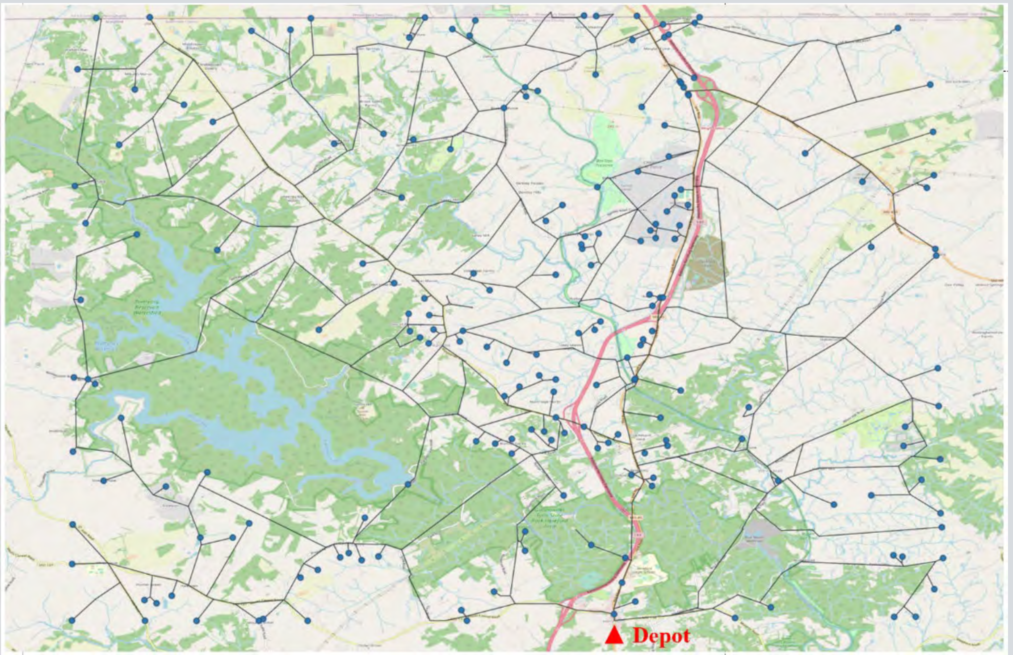

- Depot is in the southernmost of the network, near a UPS distribution center.
- The maximum load capacity of the trucks and drones are set to be 100 and 20 weight units, respectively.
- The threshold of the equity index *E* is set to 90%.

Truck/Drone Trajectory & Served Customer Locations (preliminary results)

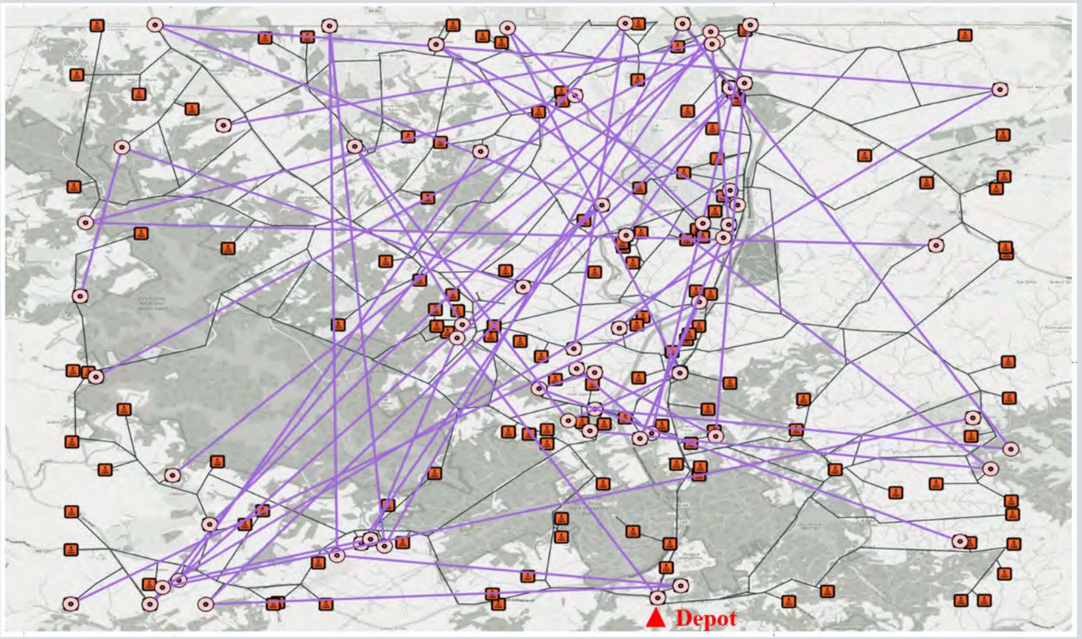

- Problem-solving techniques: Genetic algorithm
- Orange links are the truck delivery trajectory
- Blue links are the drone delivery trajectory
- 21 trucks were utilized; average load of each truck is 93 weight units

Truck/Drone Trajectory in Sample Case 1

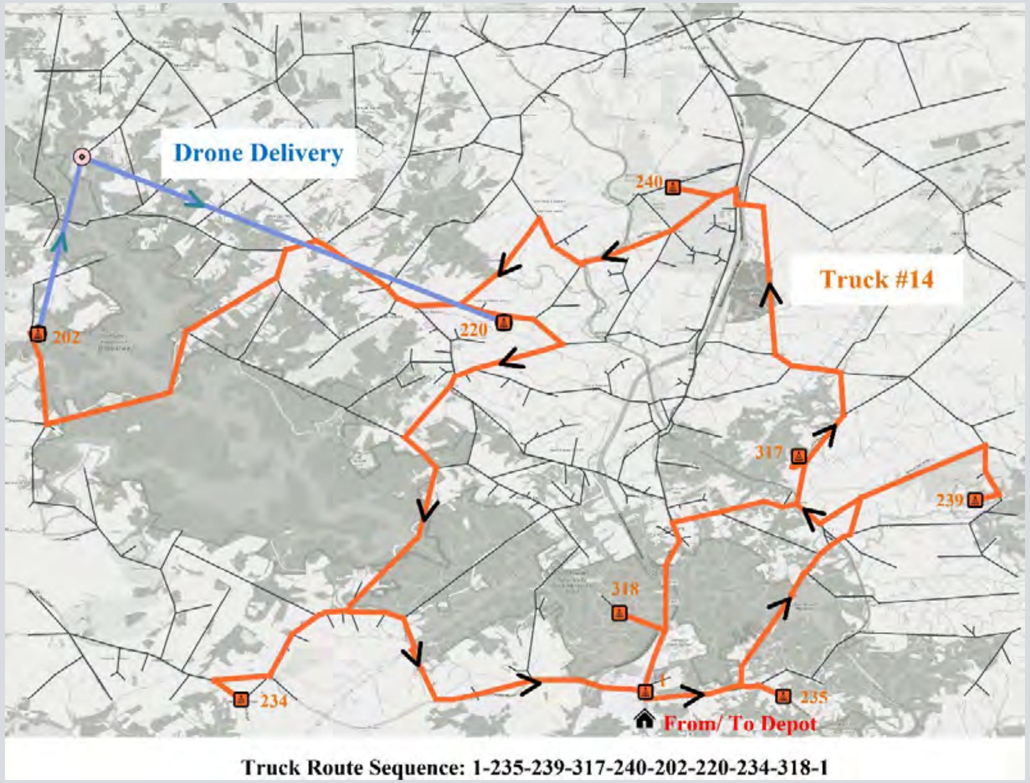

The trajectory of **Truck #14** and assigned drone deliveries in the road network

Truck/Drone Trajectory in Sample Case 2

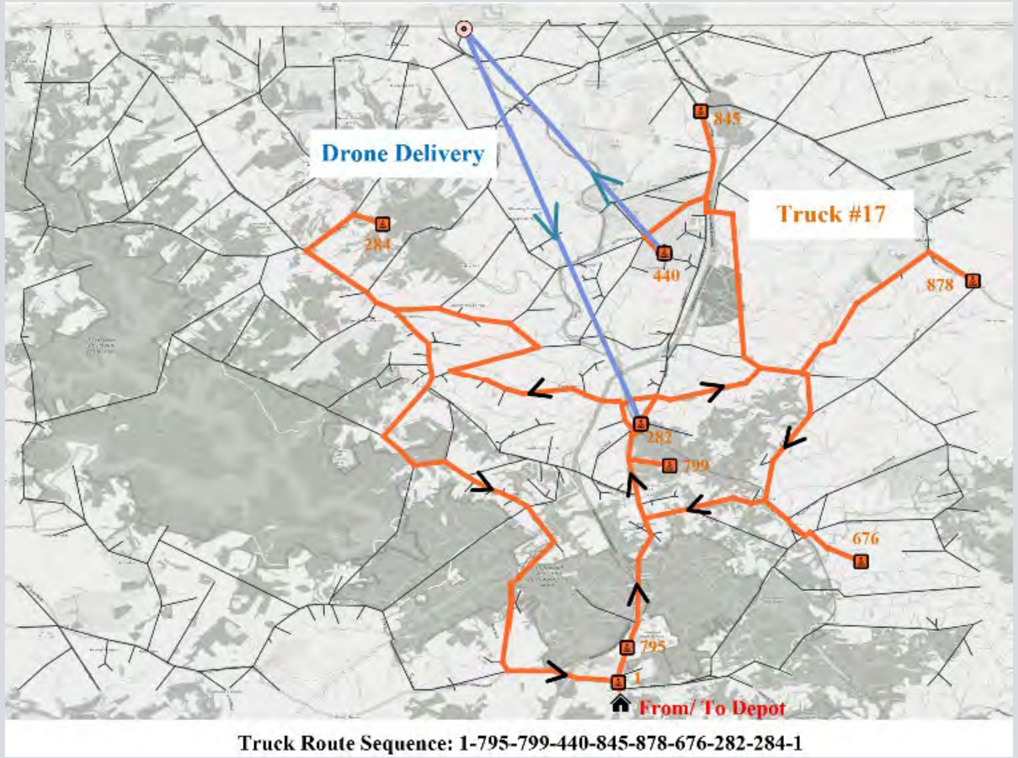

The trajectory of **Truck #17** and assigned drone delivery in the road network

Overview of the Topics

- Solving problems with equity
- Equity in problem-solving techniques

Revisiting Representative Rural Area

- The Hereford Zone is an area in Northern Baltimore County, Maryland, constituting 20% of all of the land in Baltimore County.
- 708 farms in Baltimore county.
- Farmers make up less than 5% of the total household in Hereford zone.

Data source: https://www.unitedstateszipcodes.org/; USDA 2017 Census

Two Selected Survey Responses from Farmers

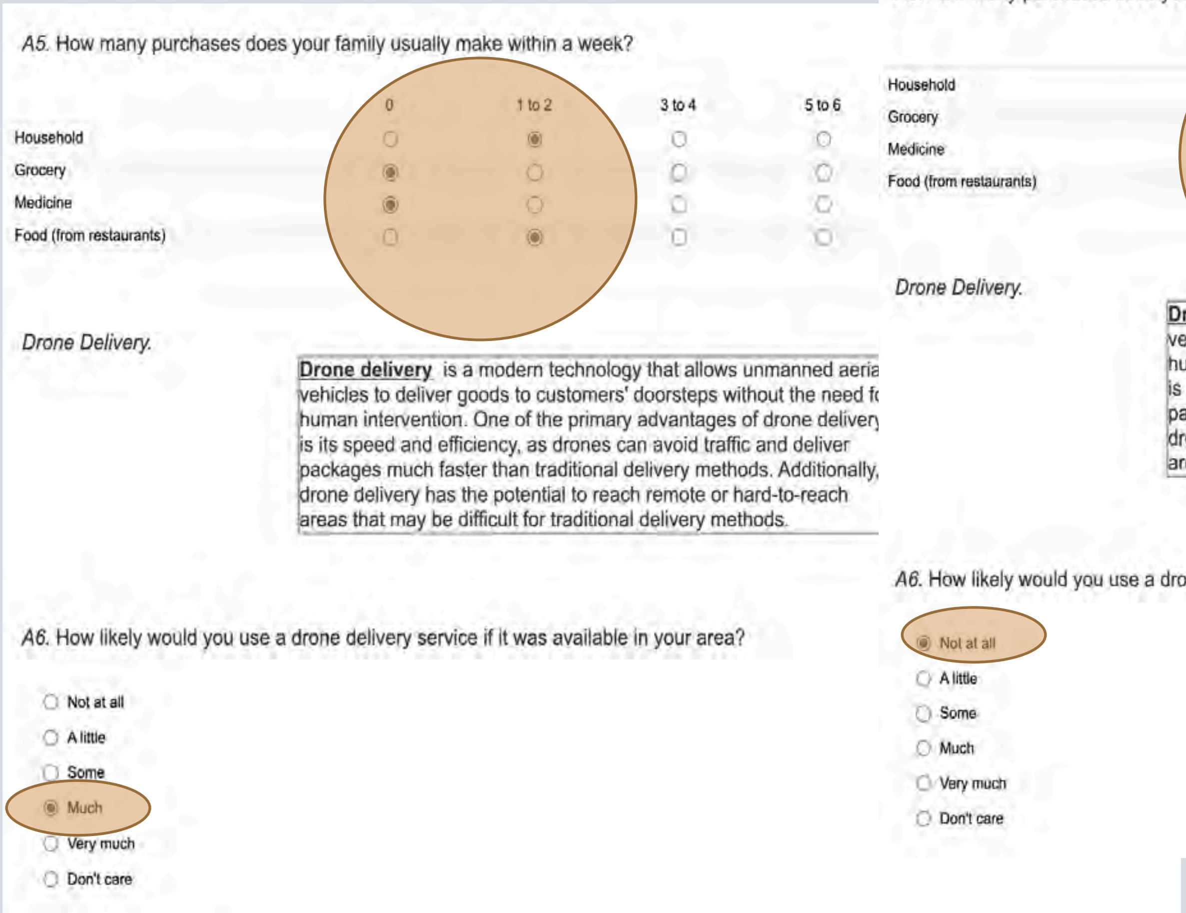

A5. How many purchases does your family usually make within a week?

	0	1 to 2	3 to 4	5 to 6
Household				
Grocery				
Medicine				
Food (from restaurants)				

Drone Delivery.

Drone delivery is a modern technology that allows unmanned aeria vehicles to deliver goods to customers' doorsteps without the need fo human intervention. One of the primary advantages of drone delivery is its speed and efficiency, as drones can avoid traffic and deliver packages much faster than traditional delivery methods. Additionally, drone delivery has the potential to reach remote or hard-to-reach areas that may be difficult for traditional delivery methods.

A6. How likely would you use a drone delivery service if it was available in your area?

- Not at all
- A little
- Some
- Much
- Very much
- Don't care

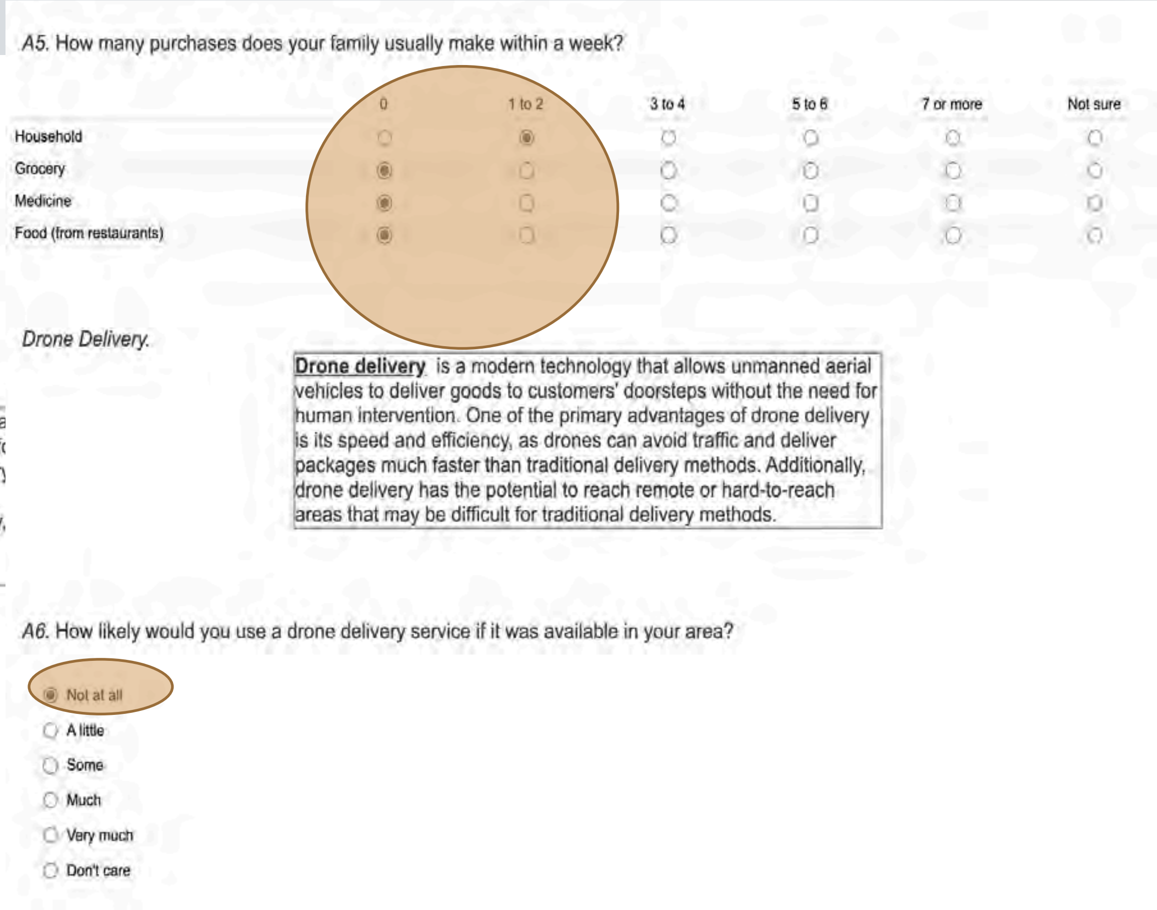

A5. How many purchases does your family usually make within a week?

	0	1 to 2	3 to 4	5 to 6	7 or more	Not sure
Household						
Grocery						
Medicine						
Food (from restaurants)						

Drone Delivery.

Drone delivery is a modern technology that allows unmanned aerial vehicles to deliver goods to customers' doorsteps without the need for human intervention. One of the primary advantages of drone delivery is its speed and efficiency, as drones can avoid traffic and deliver packages much faster than traditional delivery methods. Additionally, drone delivery has the potential to reach remote or hard-to-reach areas that may be difficult for traditional delivery methods.

A6. How likely would you use a drone delivery service if it was available in your area?

- Not at all
- A little
- Some
- Much
- Very much
- Don't care

Data source: Ongoing survey conducted by the project team

Examples of Bias When Using GA to Solve VRP-D

- **Chromosome bias**: In a genetic algorithm, the initial population of chromosomes can impact the final solution. Biased initial population towards certain routes or customers may lead to an unfair or suboptimal solution.
- **Fitness function bias**: The fitness function determines the selection of the fittest individuals for the next generation. This bias can be reinforced and lead to an unfair or suboptimal solution.
  - <u>Location bias</u>: neglecting customer stops or unwelcome intrusion into a farmland (no bias is bias)
  - <u>Objective bias</u>: cost only? Or with others such as equity, safety, sustainability, ect.
- **Crossover and mutation bias**: The crossover and mutation operators determine how new solutions are generated.
- **Resource bias**: multiple constraints are included.

Next Steps

- Developing a more comprehensive understanding of the delivery demand in Rural areas.
- Delving into the complexities of delivery equity.
- Revisiting the GA to solve VRP-D in a resounding way.
- Continuing efforts to promote equity in AI applications.

Thank you

Changing distributions and preferences in learning systems

Main takeaway

Many human-centric data sources are neither iid (in train + test) nor adversarial

modeling assumptions are critical for predicting behavior!

$$\mathscr{D} = \alpha \mathscr{D}_1 + (1 - \alpha) \mathscr{D}_2$$

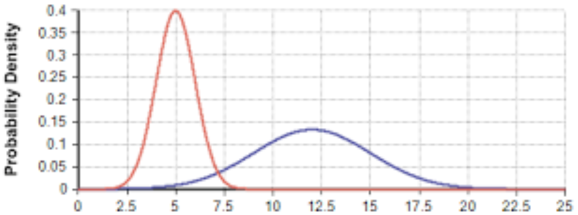

$\mathscr{D}_1$ $\mathscr{D}_2$

Training data and test data are (virtually never) distributed equally.

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

- On features, on labels…

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

- On features, on labels…

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

- On features, on labels…

Often have multiple data sources, where some

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

- On features, on labels…

Often have multiple data sources, where some

- may be unlabelled

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

- On features, on labels…

Often have multiple data sources, where some

- may be unlabelled
- may have auxilliary features

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

- On features, on labels…

Often have multiple data sources, where some

- may be unlabelled
- may have auxilliary features
- likely follow different distributions

Training data and test data are (virtually never) distributed equally.

Merely the passage of time leads to drift

  On features, on labels…

Often have multiple data sources, where some

  may be unlabelled

  may have auxilliary features

  likely follow different distributions

Hallmarks of non-iid data generating processes

Does the data have preferences over our predictions?

And do they have choices about which systems to provide which data?

Did our ML system change the way it samples as it operates?

Was the data collected at different times/locations/under different conditions?

Awesome collaborators

Awesome collaborators

Sarah Dean

Berkeley → UW

Cornell

Awesome collaborators

Sarah Dean
Berkeley →UW
Cornell

Mihaela Curmei
Berkeley

Awesome collaborators

Sarah Dean
Berkeley →UW
Cornell

Mihaela Curmei
Berkeley

Lilly Ratliff
UW ECE

Maryam Fazel
UW ECE

Setting: choosing between services.

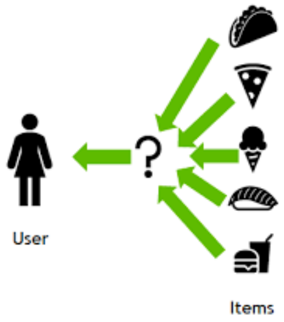

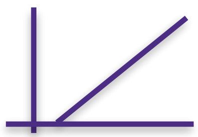

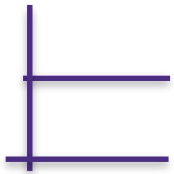

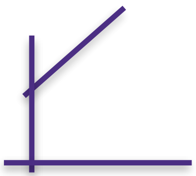

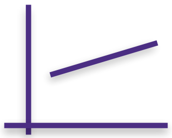

Setting: choosing between services.

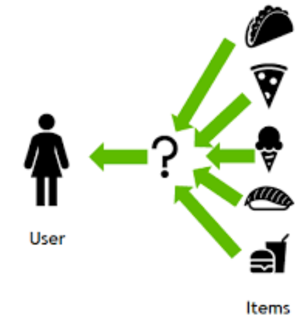

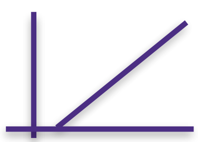

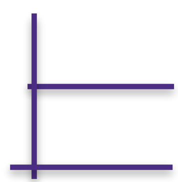

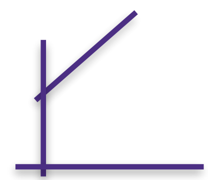

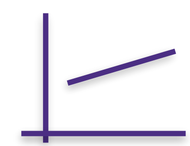

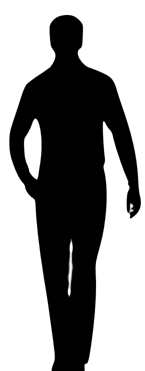

Setting: choosing between services.

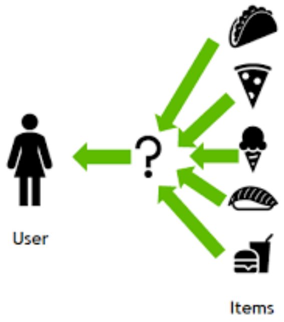

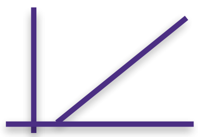

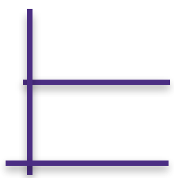

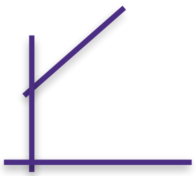

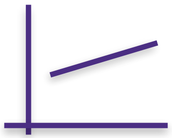

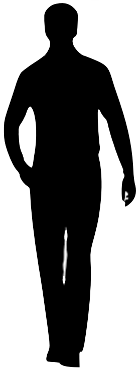

Setting: choosing between services.

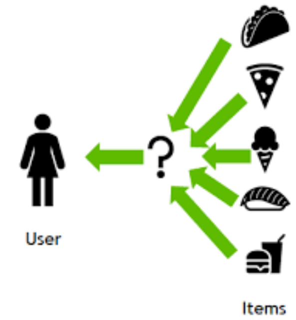

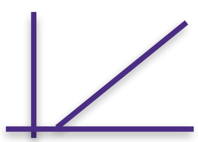

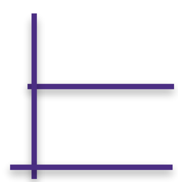

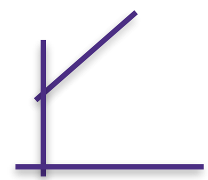

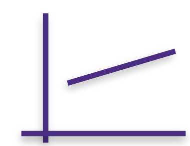

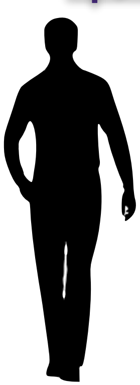

Setting: choosing between services.

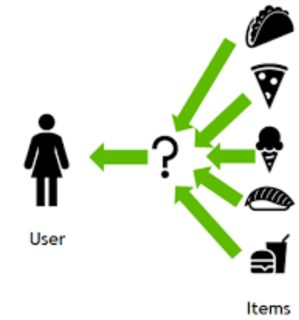

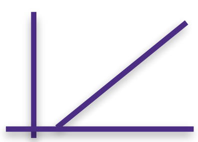

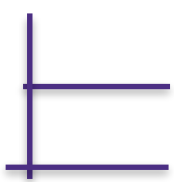

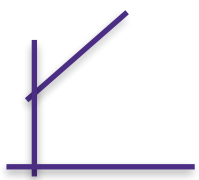

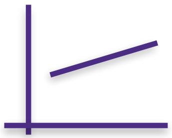

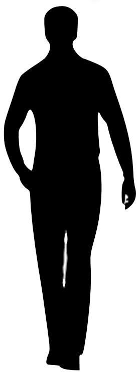

Setting: choosing between services.

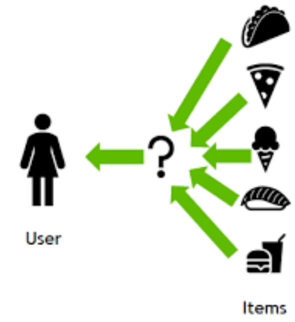

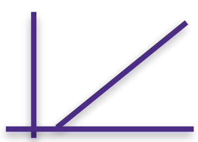

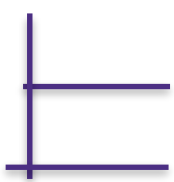

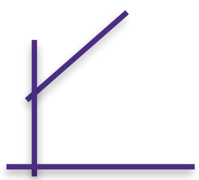

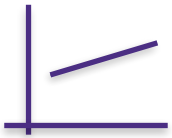

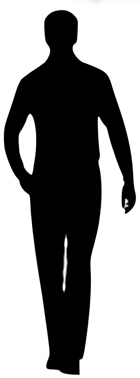

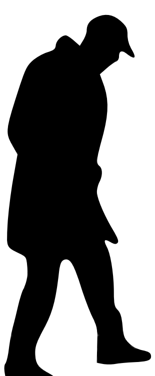

Setting: choosing between services.

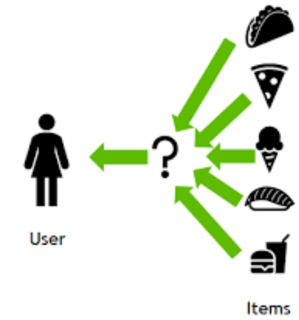

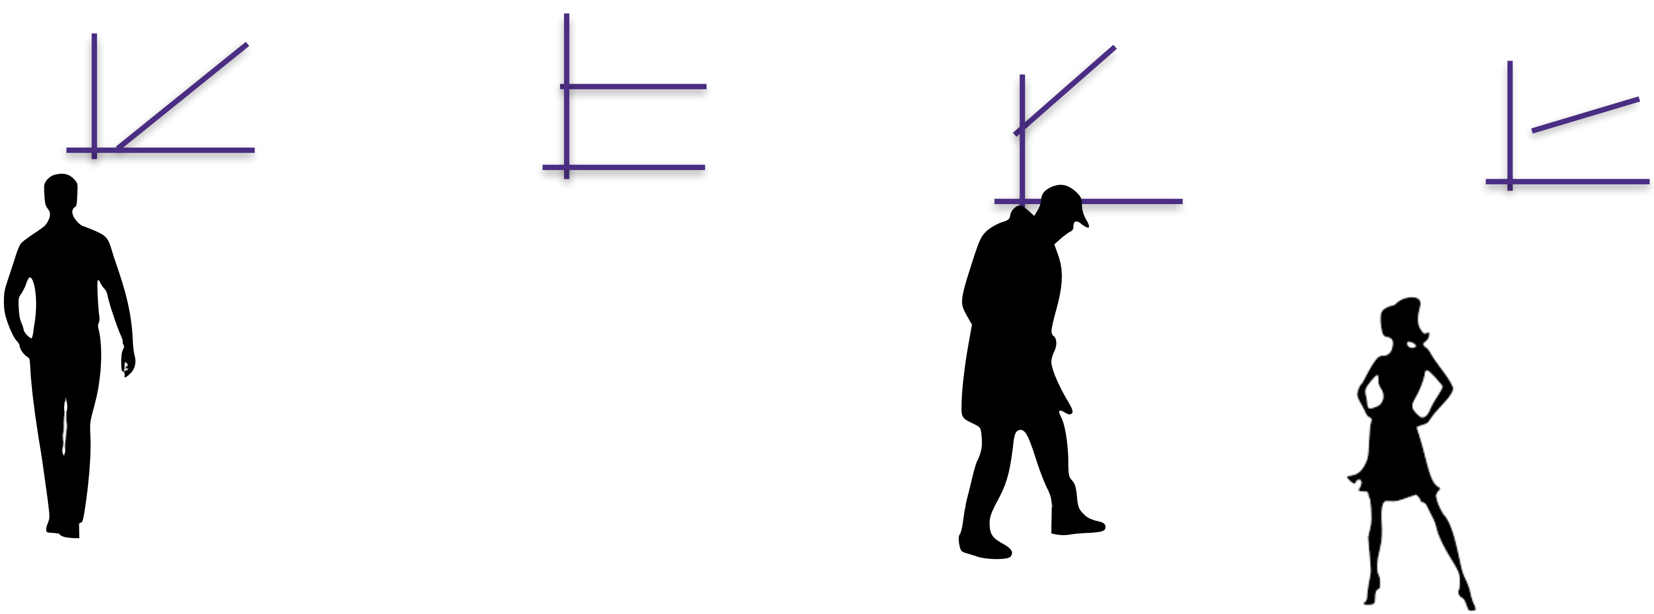

Setting: choosing between services.

Setting: choosing between services.

If populations choose between models based on their risk

Setting: choosing between services.

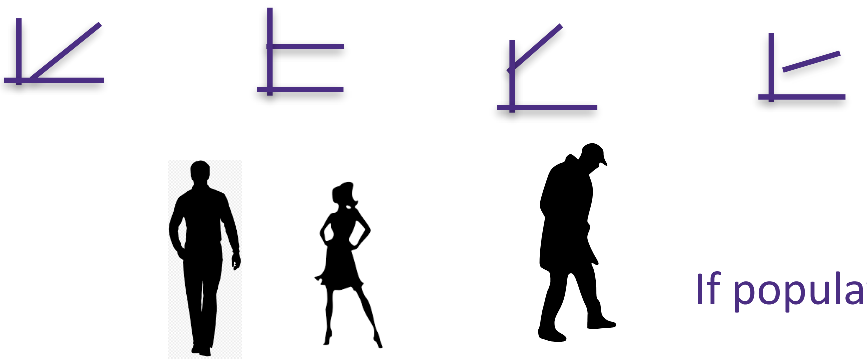

If populations choose between models based on their risk

And models are updated based on their current clientele

Setting: choosing between services.

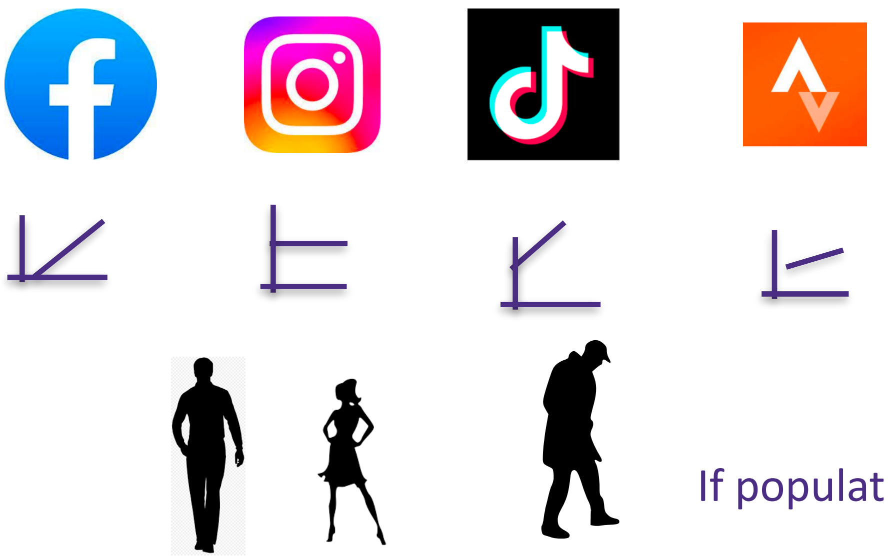

If populations choose between models based on their risk

And models are updated based on their current clientele

What can we say about the stable points of such systems?

Setting: choosing between services.

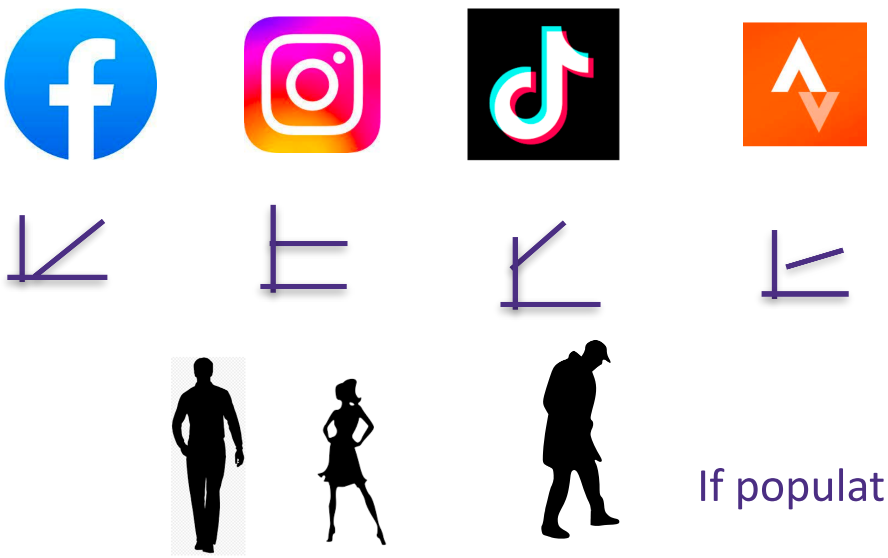

If populations choose between models based on their risk

And models are updated based on their current clientele

What can we say about the stable points of such systems?

Similar dynamics with a single model leads to representation disparity (Hashimoto et al., 2018; Zhang et al., 2019)

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha \mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

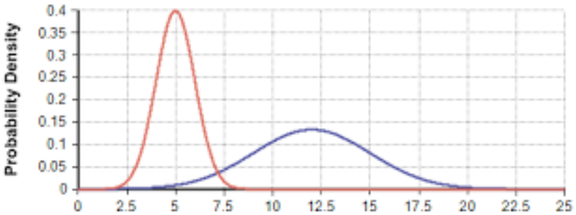

$\mathscr{D}_1$ $\mathscr{D}_2$

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff ‘22

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

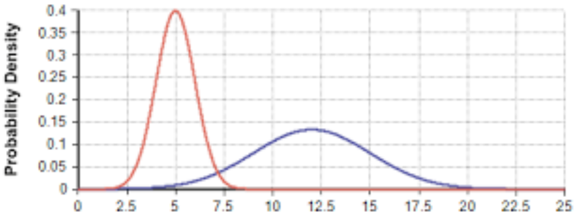

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff ‘22

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

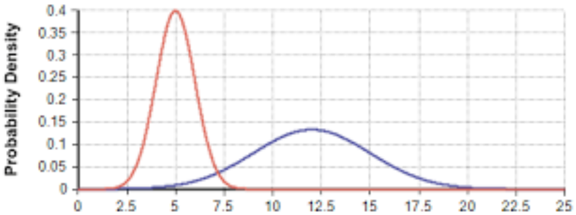

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Population $D_i \in [D_1, \ldots, D_n]$ look at $f_1, \ldots, f_m$'s performance on D_i

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff '22

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

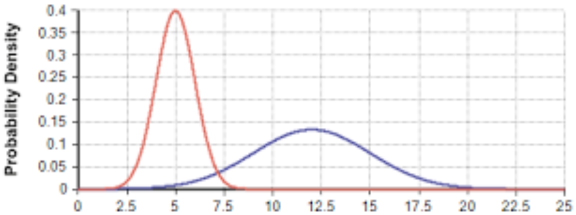

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Population $D_i \in [D_1, \ldots, D_n]$ look at $f_1, \ldots, f_m$'s performance on D_i

D_i shifts towards whichever model has smaller risk

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff '22

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

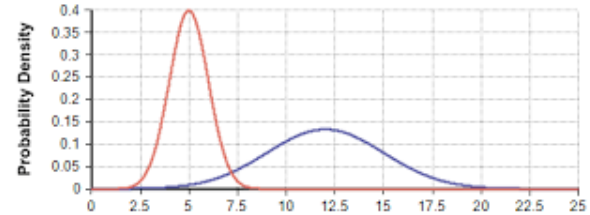

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Population $D_i \in [D_1, \ldots, D_n]$ look at $f_1, \ldots, f_m$'s performance on D_i

D_i shifts towards whichever model has smaller risk

These sorts of dynamics often lead to learner i having 100% of population i

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff ‘22

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

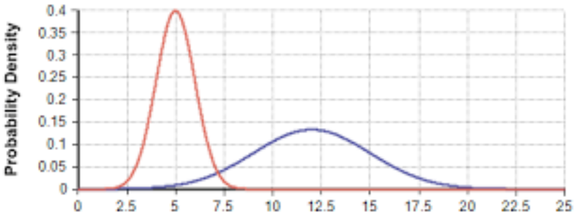

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Population $D_i \in [D_1, \ldots, D_n]$ look at $f_1, \ldots, f_m$'s performance on D_i

D_i shifts towards whichever model has smaller risk

These sorts of dynamics often lead to learner i having 100% of population i

"Segregated stable points."

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff '22

Non-iid DGP ⇒ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

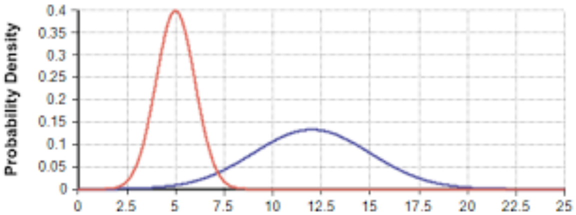

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Population $D_i \in [D_1, \ldots, D_n]$ look at $f_1, \ldots, f_m$'s performance on D_i

D_i shifts towards whichever model has smaller risk

These sorts of dynamics often lead to learner i having 100% of population i

"Segregated stable points."

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff '22

Non-iid DGP $\Rightarrow$ segregated populations

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

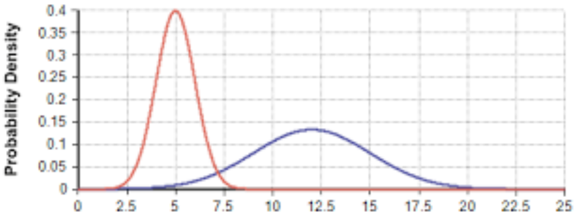

$\mathscr{D}_1$ $\mathscr{D}_2$

M learners pick regression hypotheses $f_1, \ldots, f_m$ to reduce risk on their current distribution

Population $D_i \in [D_1, \ldots, D_n]$ look at $f_1, \ldots, f_m$'s performance on D_i

D_i shifts towards whichever model has smaller risk

These sorts of dynamics often lead to learner i having 100% of population i

"Segregated stable points."

Curmei, Dean, Fazhel, **Morgenstern**, Ratliff '22

Participation dynamics

Example: linear regression with

- sub-populations 1 and 2
- learners 1 and 2

Participation dynamics

Example: linear regression with

- sub-populations ❶ and 2
- learners 1 and 2

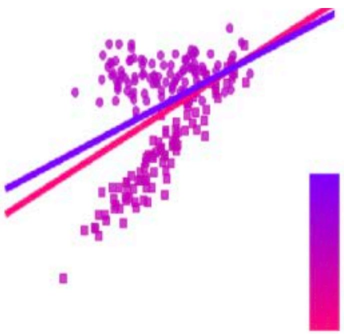

Participation dynamics

Example: linear regression with

- sub-populations 1 and 2
- learners 1 and 2

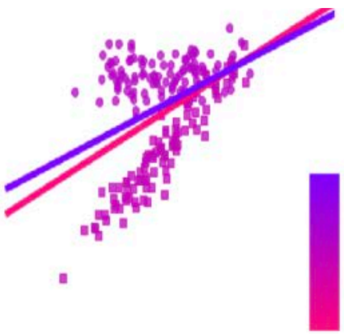

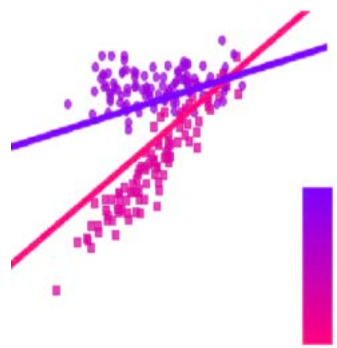

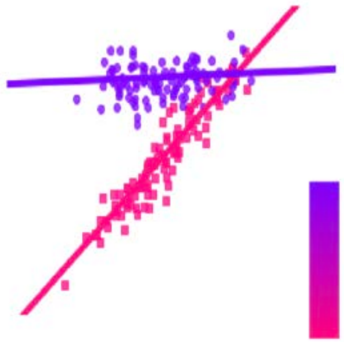

Participation dynamics

Participation dynamics

The *total risk* of a system is

$$R(\alpha, \Theta) = \sum_{i,j} \alpha_{i,j} R_i(\theta_j)$$

Participation dynamics

The *total risk* of a system is

$$R(\alpha, \Theta) = \sum_{i,j} \alpha_{i,j} R_i(\theta_j)$$

Theorem

When learners are risk reducing and
populations are risk reducing,
equilibria are local minimizers of total risk.

Participation dynamics

Participation dynamics

The *total risk* of a system is

$$R(\alpha, \Theta) = \sum_{i,j} \alpha_{i,j} R_i(\theta_j)$$

Participation dynamics

The *total risk* of a system is

$$R(\alpha, \Theta) = \sum_{i,j} \alpha_{i,j} R_i(\theta_j)$$

A *split market* has each population *i* participating at a single learner.

Participation dynamics

The *total risk* of a system is

$$R(\alpha, \Theta) = \sum_{i,j} \alpha_{i,j} R_i(\theta_j)$$

A *split market* has each population *i* participating at a single learner.

Theorem

When learners are risk minimizing in the limit,
populations are risk minimizing in the limit, and
R is strongly convex,
the asymptotically stable equilibria are split markets.

Hallmarks of non-iid data generating processes

Does the data have preferences over our predictions?

And did our predictions change their preferences?

Did our ML system change the way it samples as it operates?

Was the data collected at different times/locations/under different conditions?

Content recommendation systems

Informal goal: match users to content they like

Frameworks often used to model: matrix completion, MABs, variants on RL systems…

Content recommendation systems

Sarah Dean

Informal goal: match users to content they like

Frameworks often used to model: matrix completion, MABs, variants on RL systems…

Changing preferences in content recommendation

Changing preferences in content recommendation

Do users' preferences change as they interact with content? If so, how?

Changing preferences in content recommendation

Do users' preferences change as they interact with content? If so, how?

c1

Changing preferences in content recommendation

Do users' preferences change as they interact with content? If so, how?

c1

c2

Changing preferences in content recommendation

Do users' preferences change as they interact with content? If so, how?

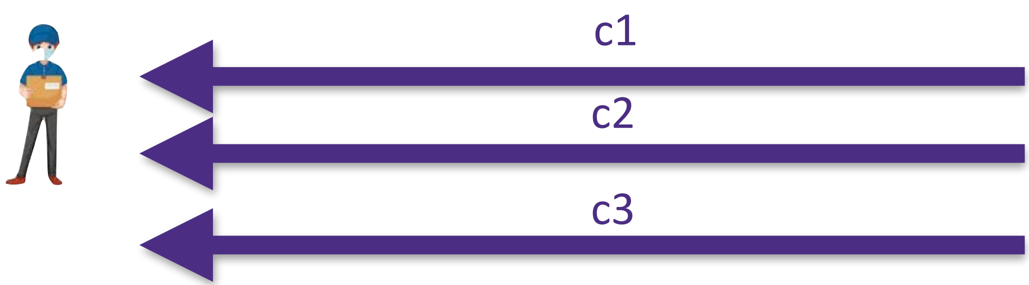

Changing preferences in content recommendation

Do users' preferences change as they interact with content? If so, how?

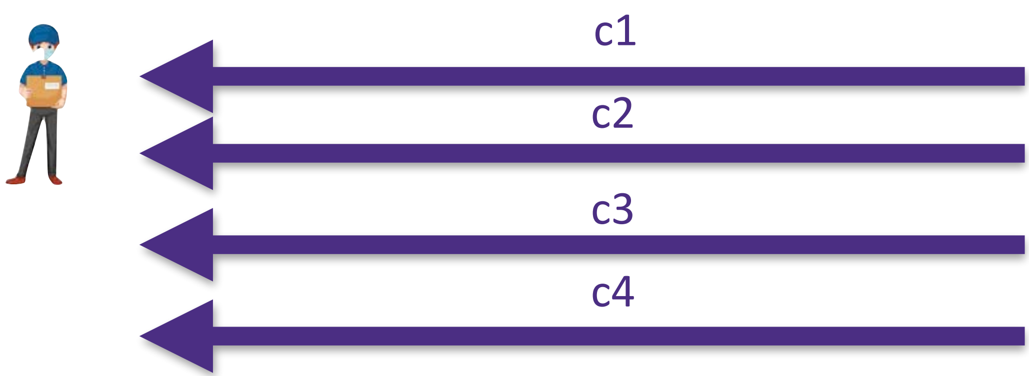

Changing preferences in content recommendation

Do users' preferences change as they interact with content? If so, how?

c1

c2

c3

c4

Evidence of change:

- Boredom
- Rabbit holes
- Polarization
- …

How do preferences change with consumption?

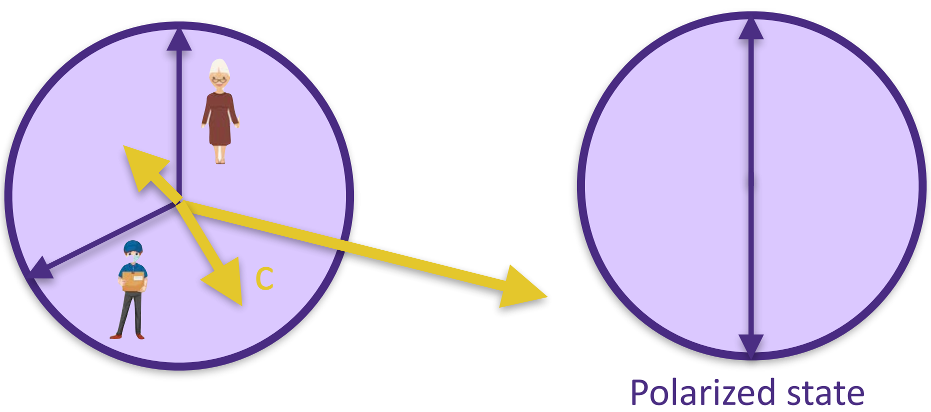

[1] Hązła, Jin, Mossel, Ramnarayan ’19
[2] Gaitonde, Kleinberg, Tardos ’21

How do preferences change with consumption?

Our preferences get drawn towards things we already like once we see it [1,2]

… and pushed away from content we already dislike when we see it.

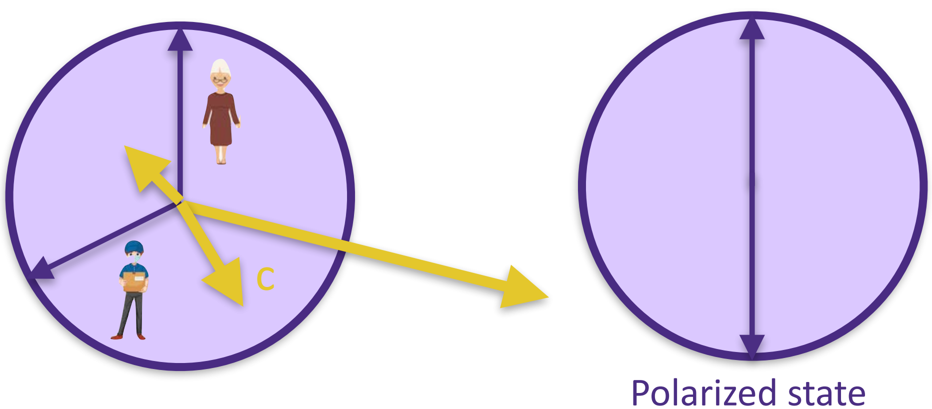

[1] Hązła, Jin, Mossel, Ramnarayan ’19

[2] Gaitonde, Kleinberg, Tardos ’21

How do preferences change with consumption?

Our preferences get drawn towards things we already like once we see it [1,2]

… and pushed away from content we already dislike when we see it.

Fix any sequence of content.

For many sets of users and content,
showing all users this same sequence leads to polarization.

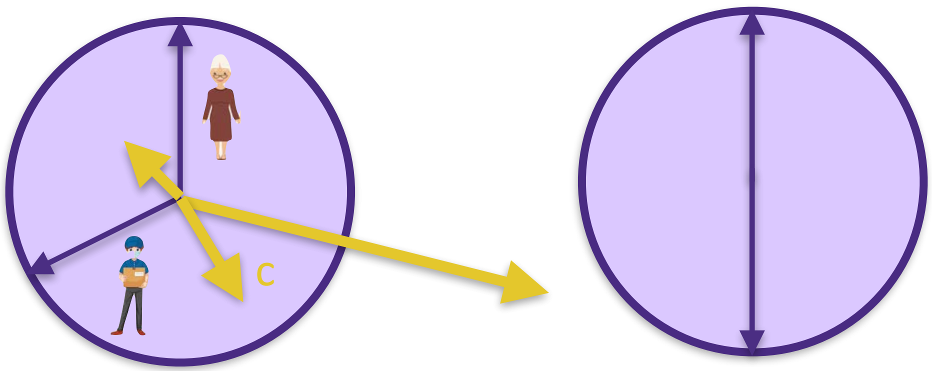

Polarized state

[1] Hązła, Jin, Mossel, Ramnarayan ’19
[2] Gaitonde, Kleinberg, Tardos ’21

How do preferences change with consumption?

Our preferences get drawn towards things we already like once we see it [1,2]

… and pushed away from content we already dislike when we see it.

Fix any sequence of content.

For many sets of users and content,
showing all users this same sequence leads to polarization.

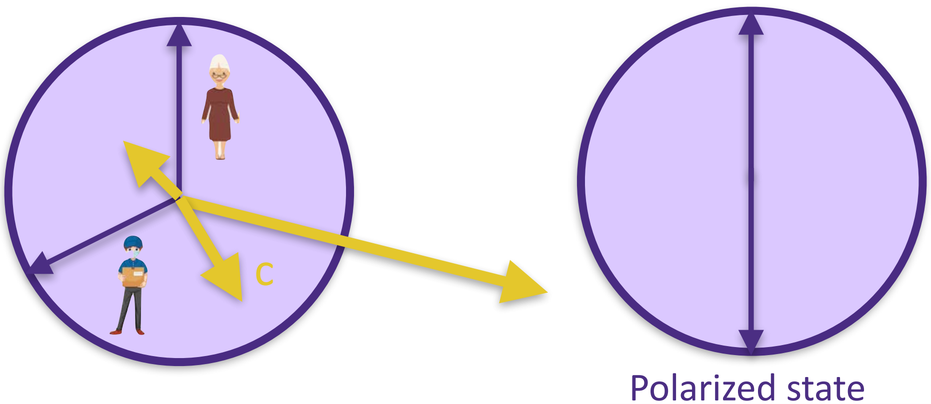

What about if we personalize content shown to each user?

[1] Hązła, Jin, Mossel, Ramnarayan '19
[2] Gaitonde, Kleinberg, Tardos '21

Does personalization preclude polarization?

Personalization is precisely a system where not all users consume the same content.

Does that mean that polarization may not occur in personalized systems?

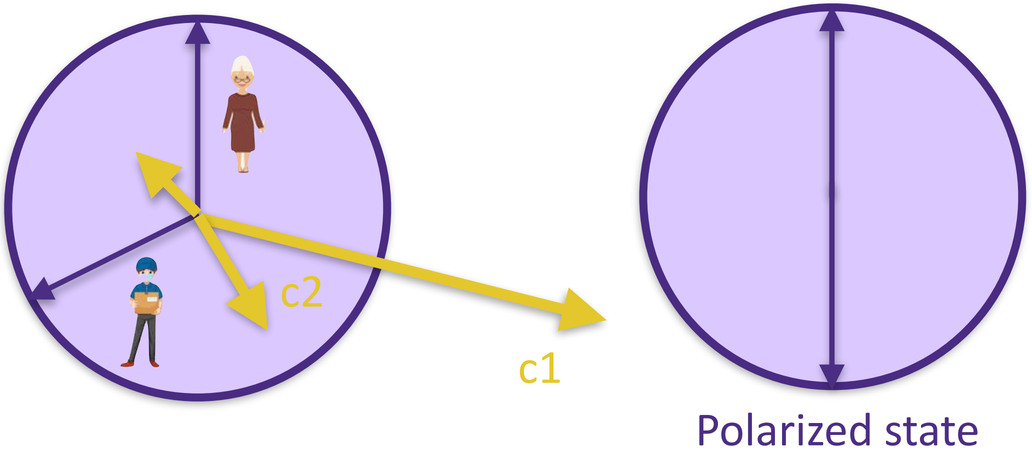

Does personalization preclude polarization?

Personalization is precisely a system where not all users consume the same content.

Does that mean that polarization may not occur in personalized systems?

Theorem:

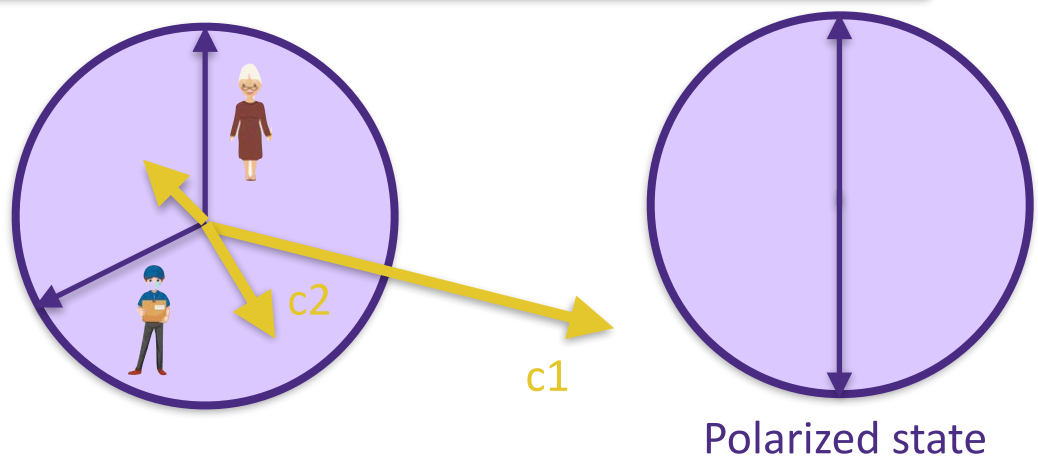

Does personalization preclude polarization?

Personalization is precisely a system where not all users consume the same content.

Does that mean that polarization may not occur in personalized systems?

Theorem:

Preferences need not be polarized from personalized content consumption.

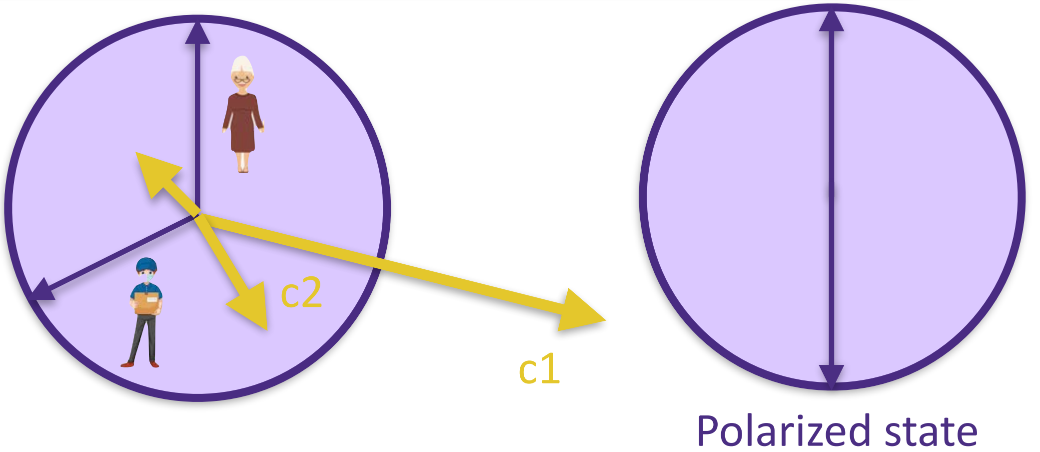

Does personalization preclude polarization?

Personalization is precisely a system where not all users consume the same content.

Does that mean that polarization may not occur in personalized systems?

Theorem:

Preferences need not be polarized from personalized content consumption.

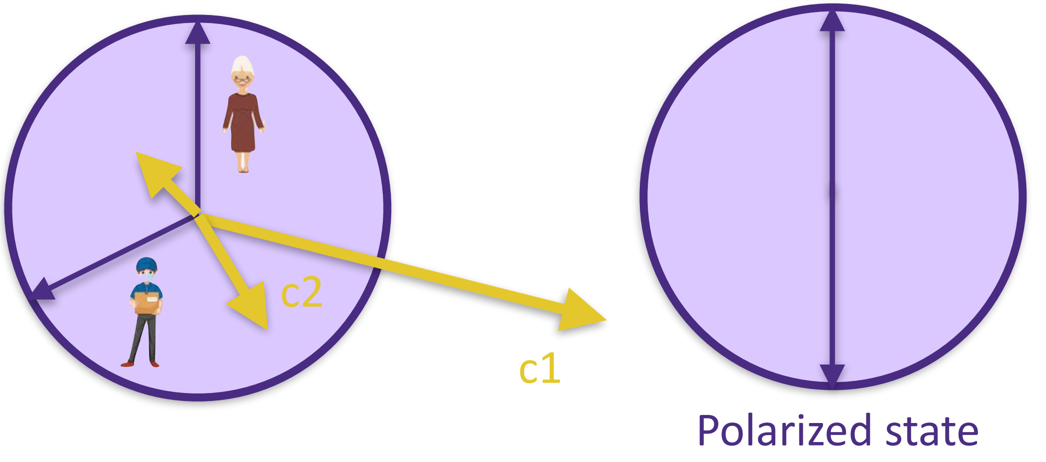

Does personalization preclude polarization?

Personalization is precisely a system where not all users consume the same content.

Does that mean that polarization may not occur in personalized systems?

Theorem:

Preferences need not be polarized from personalized content consumption.

Every preference will approach the *span* of the content set *V.*

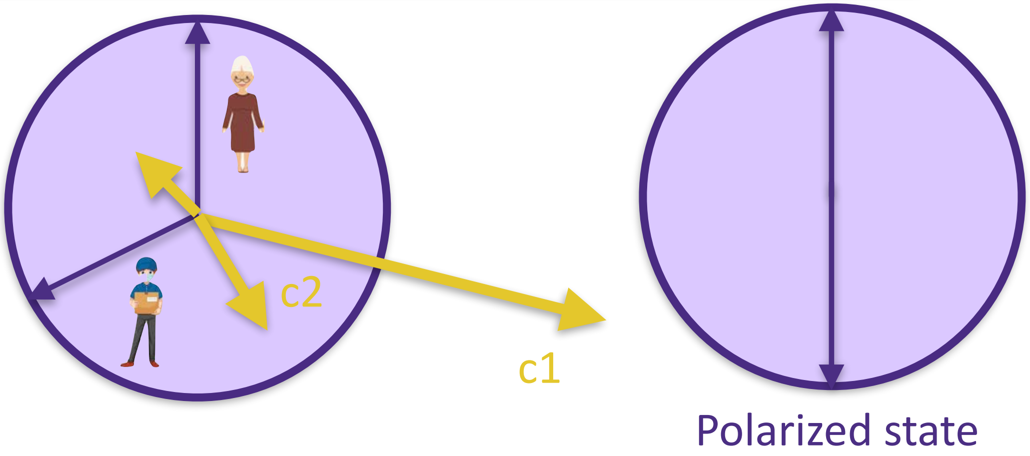

The right objective function?

One "standard" objective for a single user, user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

The right objective function?

One "standard" objective for a single user, user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

The right objective function?

One “standard” objective for a single user, user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

$$y_t = \langle u^t, v_t \rangle + w_t$$

The right objective function?

One “standard” objective for a single user,
user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

$$y_t = \langle u^t, v_t \rangle + w_t$$

Theorem:

The right objective function?

One "standard" objective for a single user,
user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

$$y_t = \langle u^t, v_t \rangle + w_t$$

Theorem:
Minimizing regret for user welfare is easy* if $\exists c : c, -c \in V$.

The right objective function?

One "standard" objective for a single user,
user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

$$y_t = \langle u^t, v_t \rangle + w_t$$

Theorem:
Minimizing regret for user welfare is easy* if $\exists c : c, -c \in V$.

The right objective function?

One “standard” objective for a single user,
user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

$$y_t = \langle u^t, v_t \rangle + w_t$$

Theorem:

Minimizing regret for user welfare is easy* if $\exists c : c, -c \in V$.

Many policies will have vanishing regret

The right objective function?

One “standard” objective for a single user,
user welfare

$$\max_{v_1,\ldots,v_T} \sum_t \langle u^t, v_t \rangle$$

Observe

$$y_t = \langle u^t, v_t \rangle + w_t$$

Theorem:
Minimizing regret for user welfare is easy* if $\exists c : c, -c \in V$.

Many policies will have vanishing regret
(… and they’re all pretty dumb. Show the same + content, over and over again, until your preferences perfectly align with that content.)

A more technically meaningful objective

For a single user, minimize their change in preferences.

Stat-Pref:

$$\min_{v_1,\ldots,v_T} \frac{\sum_t |u^t - u|}{T}$$

A more technically meaningful objective

For a single user, minimize their change in preferences.

Stat-Pref:

$$\min_{v_1,\ldots,v_T} \frac{\sum_t |u^t - u|}{T}$$

Theorem:

A more technically meaningful objective

For a single user, minimize their change in preferences.

Stat-Pref:

$$\min_{v_1,\ldots,v_T} \frac{\sum_t |u^t - u|}{T}$$

Theorem:

Exists a policy where E[Stat-Pref] → 0 with knowledge of *u* if $u \in \text{span}(\{\text{sgn}(u, v) \cdot v \,|\, v \in V\})$.

A more technically meaningful objective

For a single user, minimize their change in preferences.

Stat-Pref:

$$\min_{v_1,\ldots,v_T} \frac{\sum_t |u^t - u|}{T}$$

Theorem:

Exists a policy where E[Stat-Pref] → 0 with knowledge of *u* if
$u \in \text{span}(\{\text{sgn}(u, v) \cdot v \,|\, v \in V\})$.

A more technically meaningful objective

For a single user, minimize their change in preferences.

Stat-Pref:

$$\min_{v_1,\ldots,v_T} \frac{\sum_t |u^t - u|}{T}$$

Theorem:

Exists a policy where E[Stat-Pref] → 0 with knowledge of *u* if
$u \in \text{span}(\{\text{sgn}(u, v) \cdot v \mid v \in V\})$.

Without knowledge of *u*, sometimes possible to learn *u* from interactions first.

Main takeaway

Many human-centric data sources are neither
iid (in train + test) nor adversarial

modeling assumptions are critical for predicting behavior!

$$\mathscr{D} = \alpha\mathscr{D}_1 + (1-\alpha)\mathscr{D}_2$$

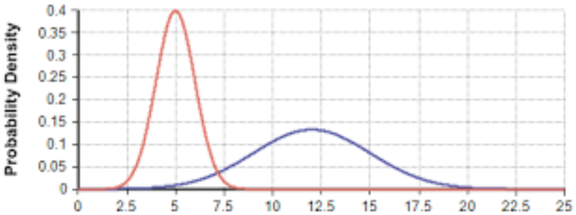

$\mathscr{D}_1$ $\mathscr{D}_2$