Multiagent reasoning for social impact: Results from deployments for public health and conservation

CRCS Center for Research on Computation and Society at Harvard School of Engineering and Applied Sciences

Google Research

MILIND TAMBE

Director, Ctr for Research on Computation & Society

Harvard University

Director "AI for Social Good"

Google Research India

@MilindTambe_AI

AI & Multiagent Systems Research for Social Impact

Public Health

Conservation

Public Safety and Security

Optimize Our Limited Intervention Resources

Optimizing Limited Intervention Resources

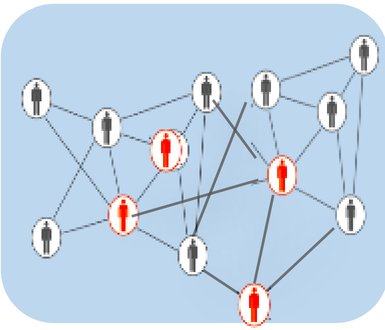

Social Networks & Bandits

Public Health

Green security games

Conservation

Stackelberg security games

Public Safety & Security

Three Common Themes

Non-profits, Data-to-deployment pipeline, Multiagent systems

Empower non-profits to use AI tools; avoid being gatekeepers to AI4SI technology

LOS ANGELES LGBT CENTER

S.P.Y safe place for youth

my friend's place

armman

ASHOKA TRUST FOR RESEARCH IN ECOLOGY & THE ENVIRONMENT

Wildlife Conservation Society

WWF

Air Shepherd The Lindbergh Foundation

Swasti THE HEALTH CATALYST

Wildlife Conservation Trust

SMART

everwell

KhushiBaby

Three Common Themes
Non-profits, Data-to-deployment pipeline, Multiagent systems

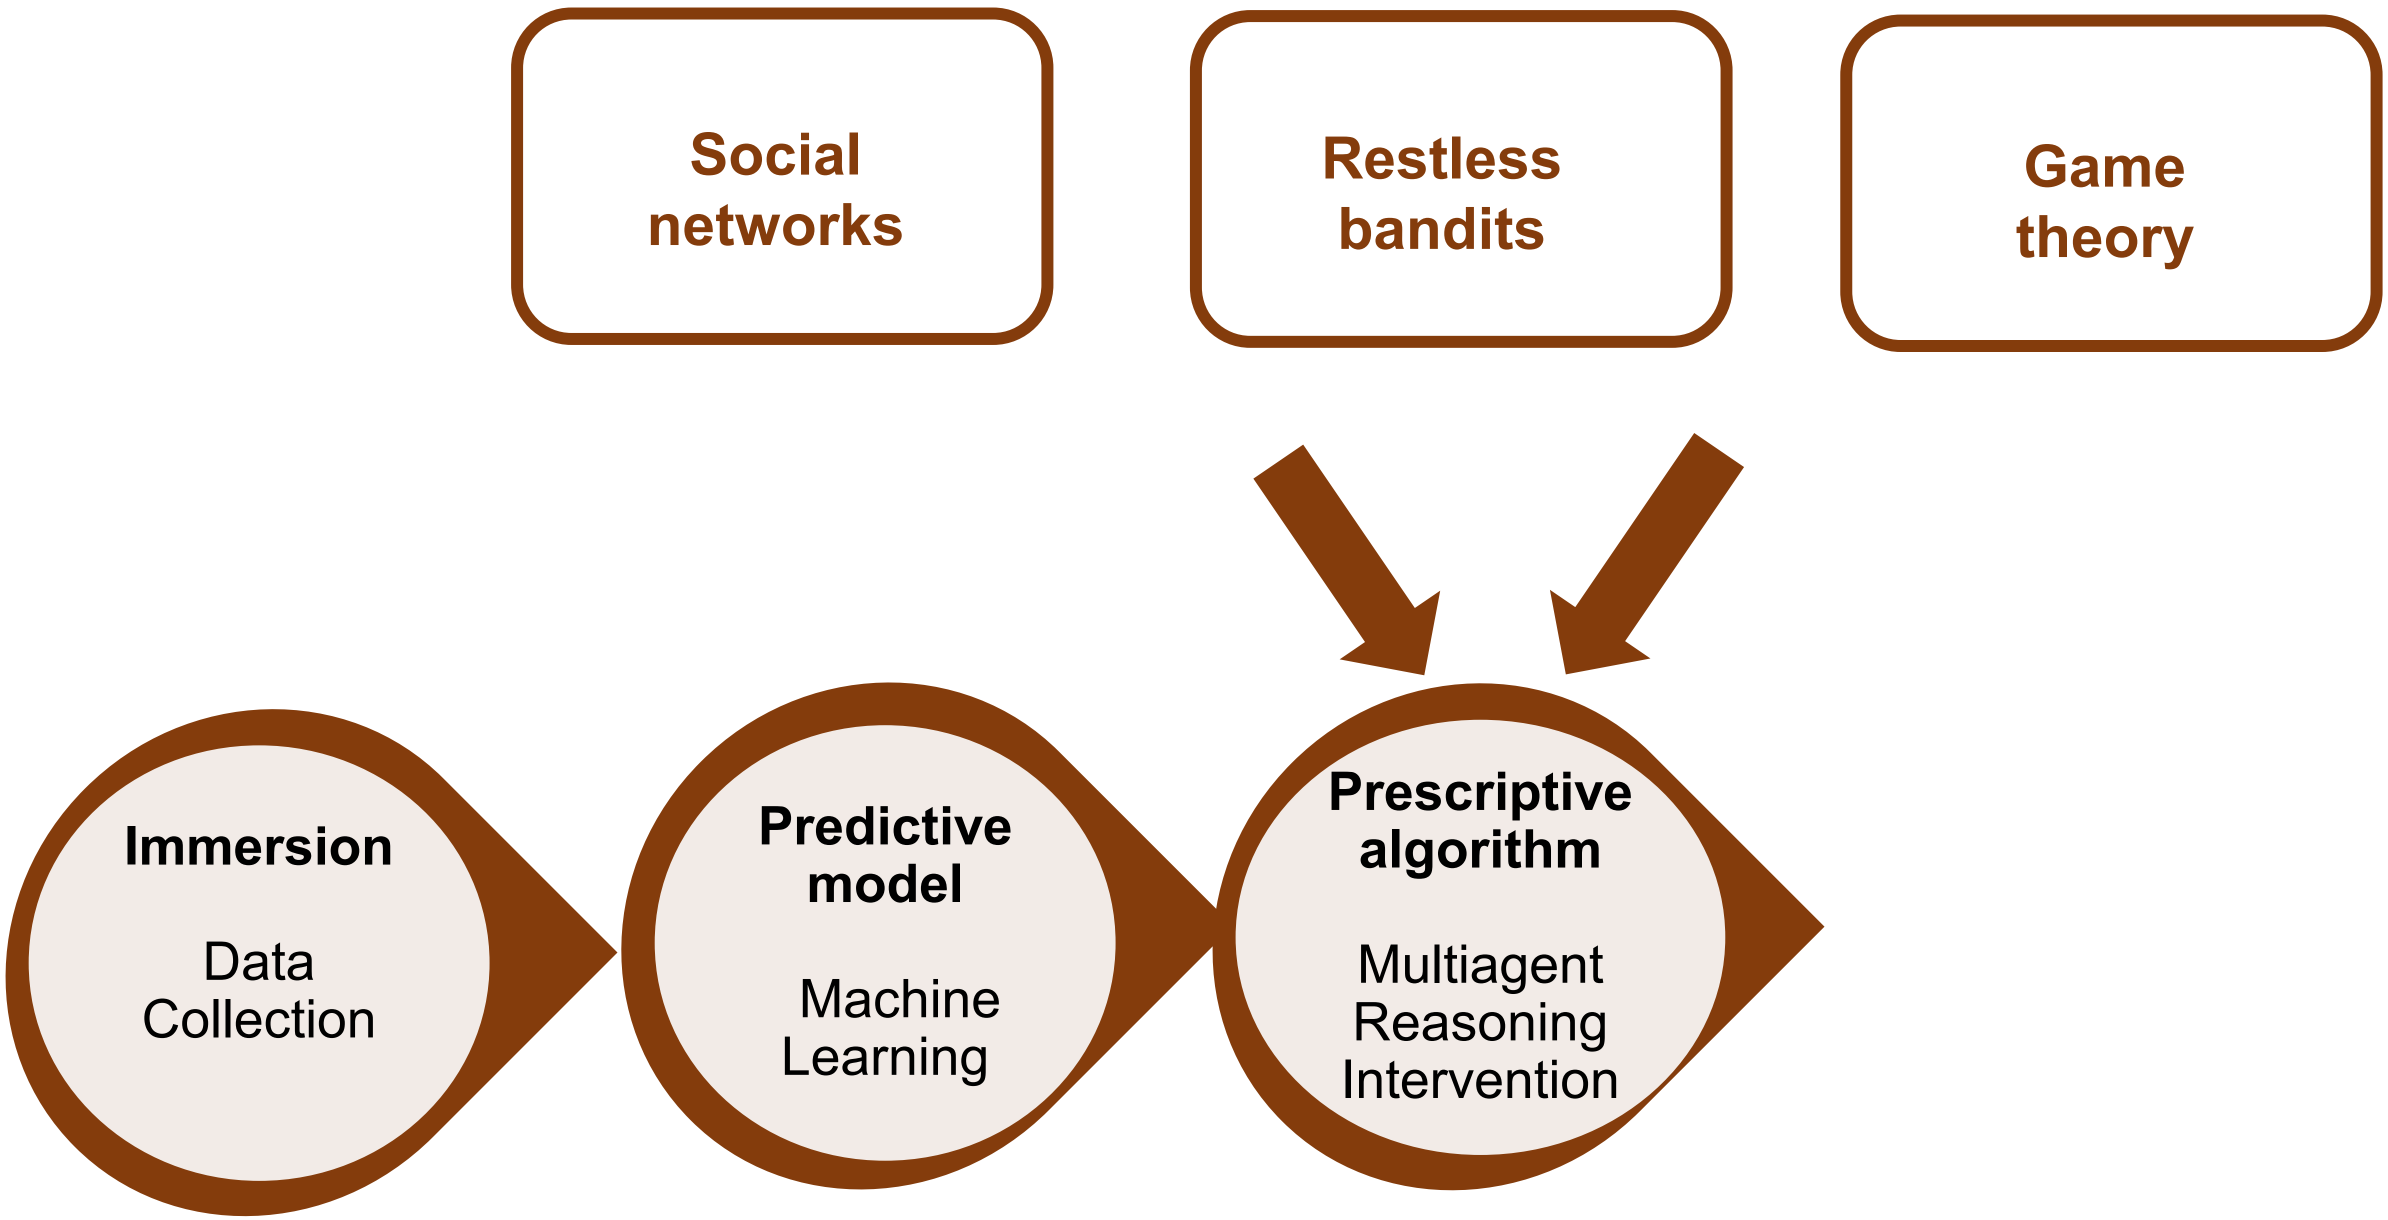

Three Common Themes

Multiagent systems, Data-to-deployment pipeline, Interdisciplinary partnerships

Field test & deployment: Social impact is a key objective

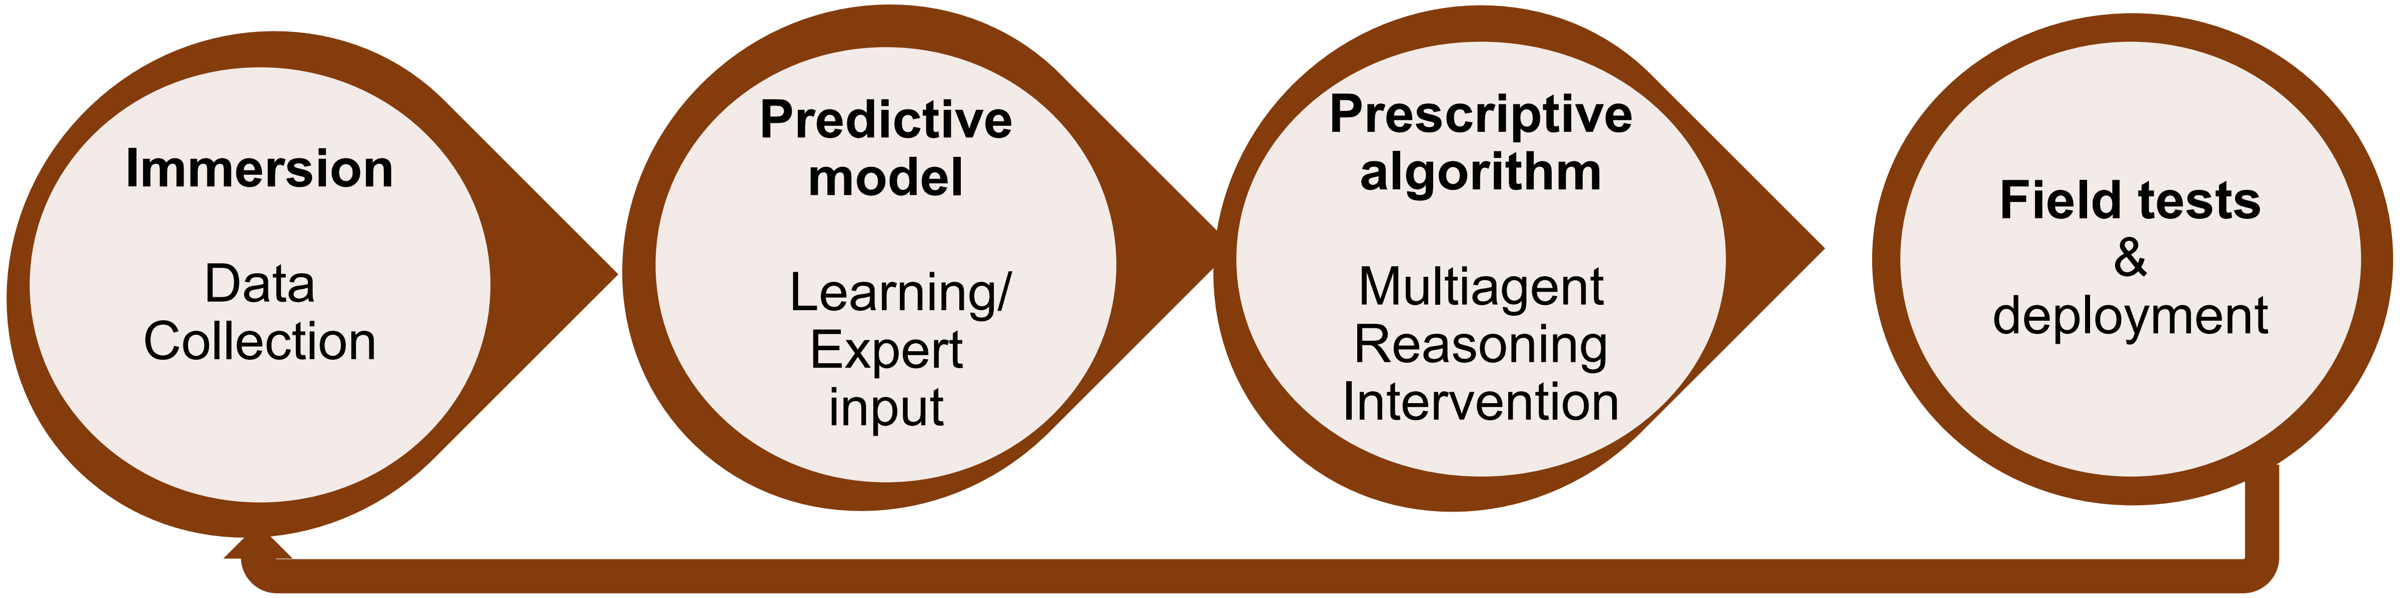

Outline

Public Health

- *Health information dissemination: Social networks*
- *Health program adherence: ML & Bandits*
- *COVID-19: Agent-based modeling*

Conservation

- Cover papers from 2017-now [AAMAS, AAAI, IJCAI, NeurIPS…]
- PhD students & postdocs highlighted

Information dissemination & behavior change
Optimizing Limited Intervention (Social Worker) Resources

Prevent HIV in youth experiencing homelessness: HIV 10x housed population

- ***Shelters**: Limited number of peer leaders to spread HIV information in social networks*
- *"Real" face-to-face interactions; not Facebook etc*

Influence Maximization in Social Networks

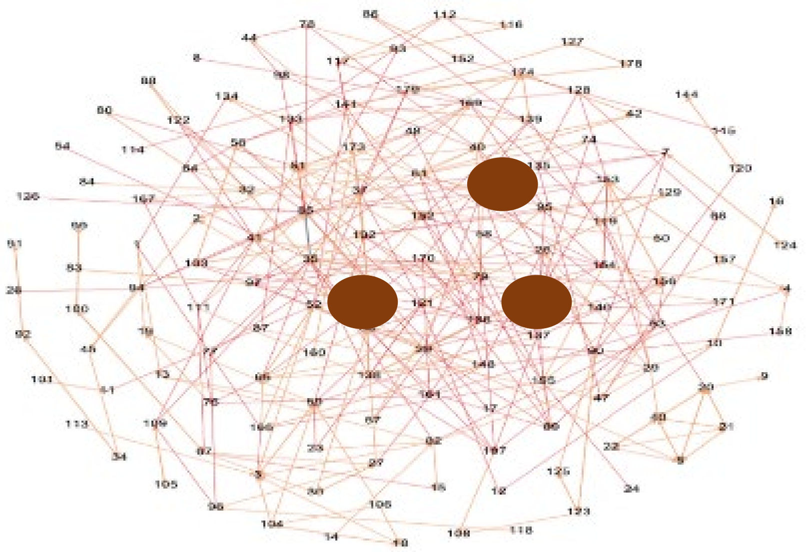

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
  - Assume: Independent cascade model of information spread
- Objective:
  - Maximize expected number of influenced nodes

Influence Maximization in Social Networks
Three Key Challenges Combined Together

- Uncertainty in propagation probability over edges
- Multi-step dynamic policies to handle peer leader “no shows”
- Unknown social network, limited query budget to SAMPLE network

Challenge 1: Uncertainty in Real-world Physical Social Networks

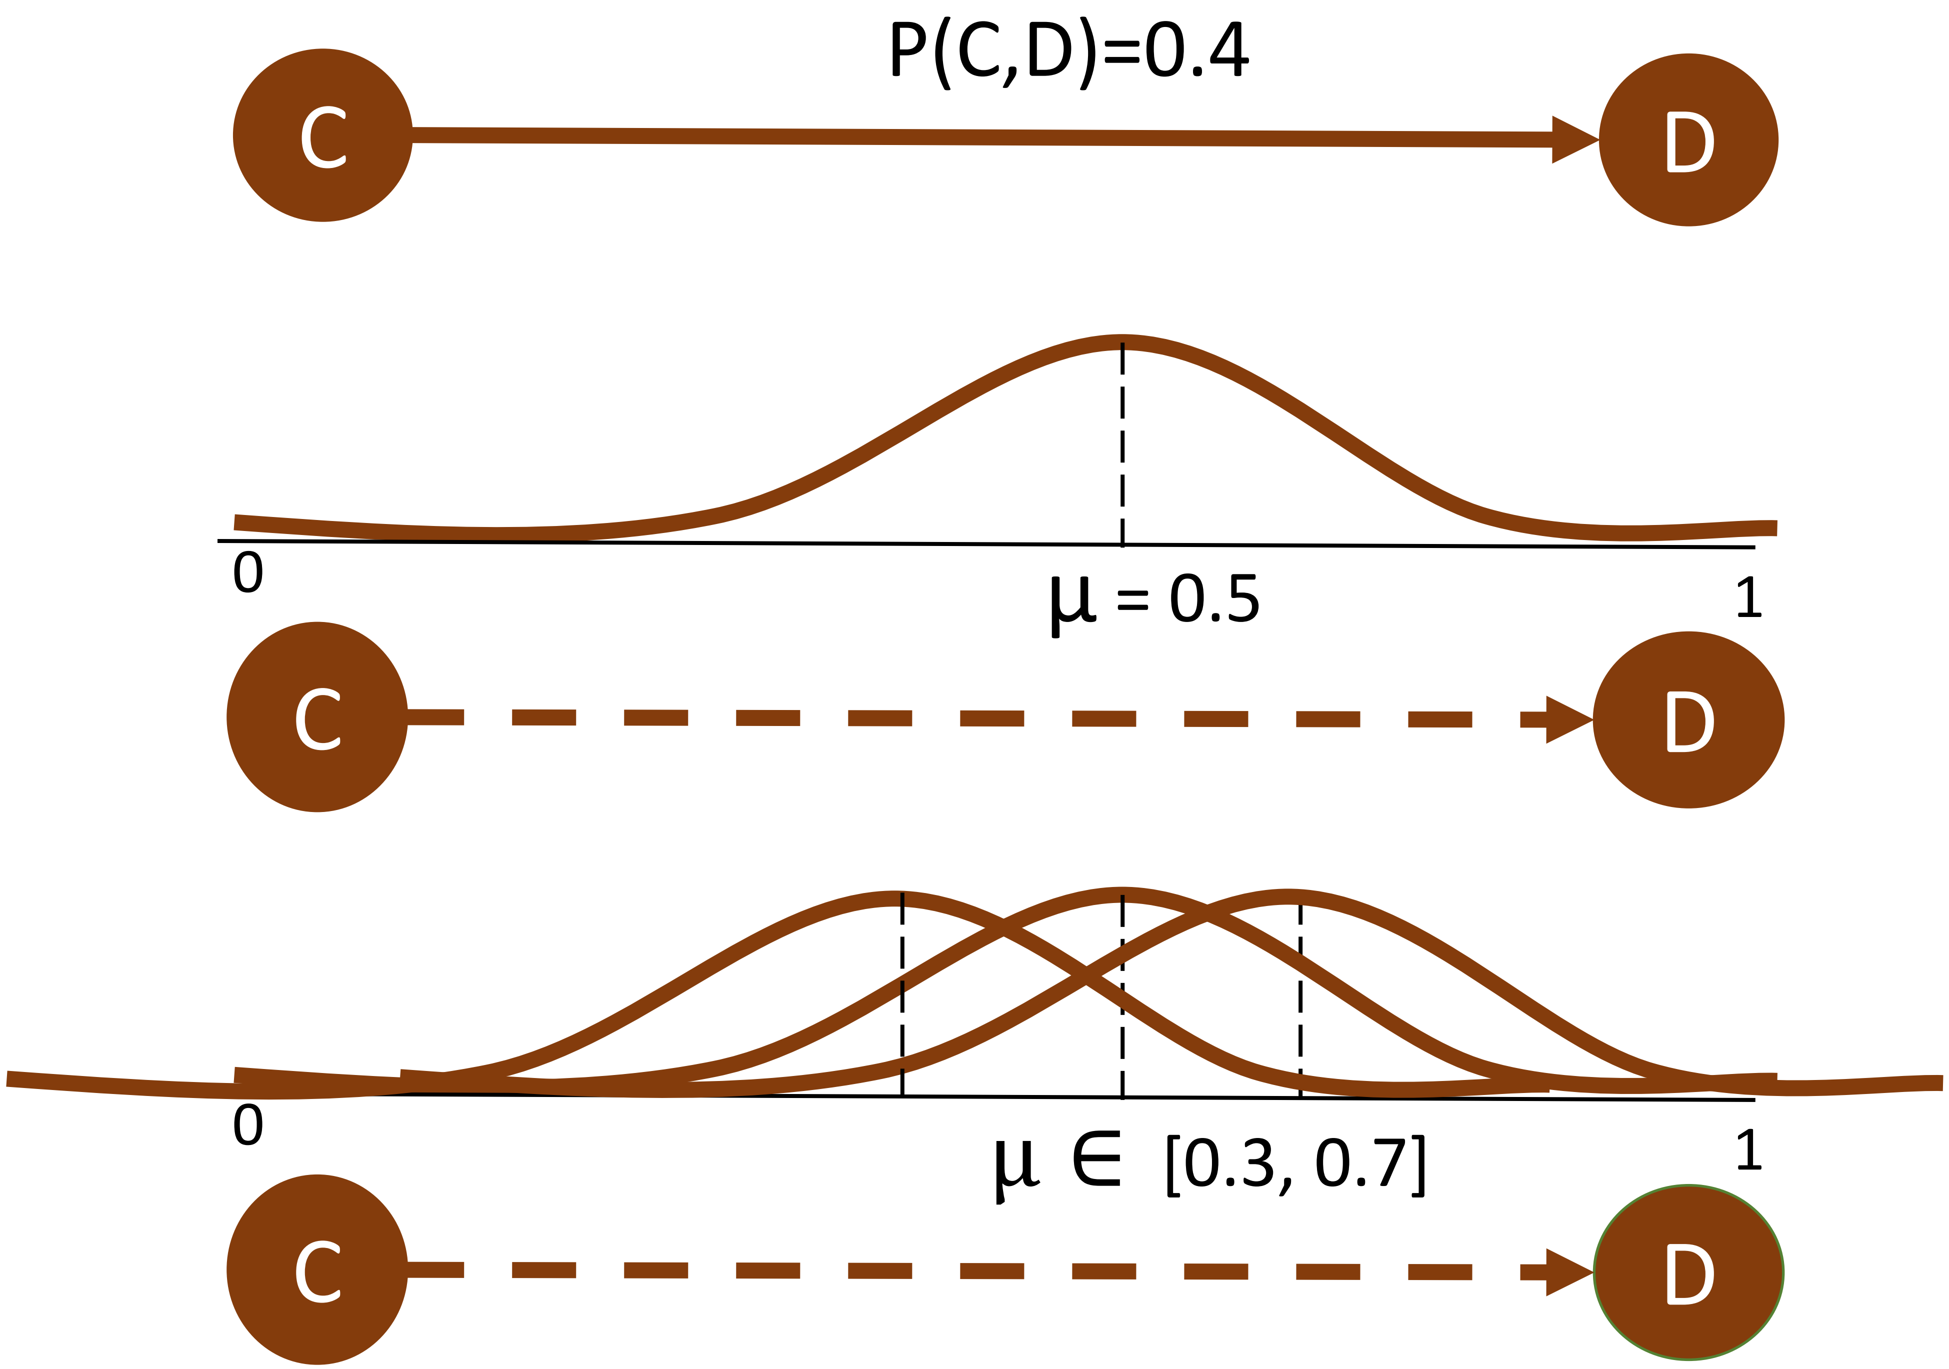

Wilder

Robust Influence Maximization
(AAMAS 2017)

- Worst case parameters: a zero-sum game against nature

$$\max_{x \in \Delta^{|P|}} \min_{\mu,\sigma} \sum x_p \frac{(Outcome(p))}{OPT(\mu,\sigma)}$$

Algorithm

Choose Peer Leaders p ∈ P generating mixed strategy "$x \in \Delta^{|P|}$"

vs

Nature

Chooses parameters μ,σ

HEALER Algorithm
Robust Influence Maximization
(AAMAS 2017)


Wilder

Theorem: Converge with approximation guarantees

- Equilibrium strategy despite exponential strategy spaces: Double oracle

Nature

Influencer

	Params #1	Params #2	Params #3
Policy #1	0.8, -0.8	0.3, -0.3	0.4, -0.4
Policy #2	0.7, -0.7	0.5, -0.5	0.6, -0.6
Policy #3	0.6, -0.6	0.4, -0.4	0.7, -0.7

Influencer's oracle

	Params #1	Params #2
Policy #1	0.8, -0.8	0.3, -0.3
Policy #2	0.7, -0.7	0.5, -0.5
Policy #3	0.6, -0.6	0.4, -0.4

Nature's oracle

	Params #1	Params #2	Params #3
Policy #1	0.8, -0.8	0.3, -0.3	0.4, -0.4
Policy #2	0.7, -0.7	0.5, -0.5	0.6, -0.6
Policy #3	0.6, -0.6	0.4, -0.4	0.7, -0.7

Challenge 2: Contingency-Aware Influence Maximization
(AAMAS 2018a)

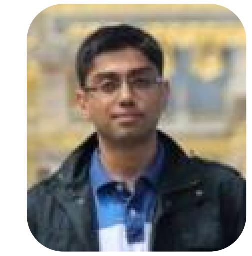
Yadav

K = 4
1^{st} time step

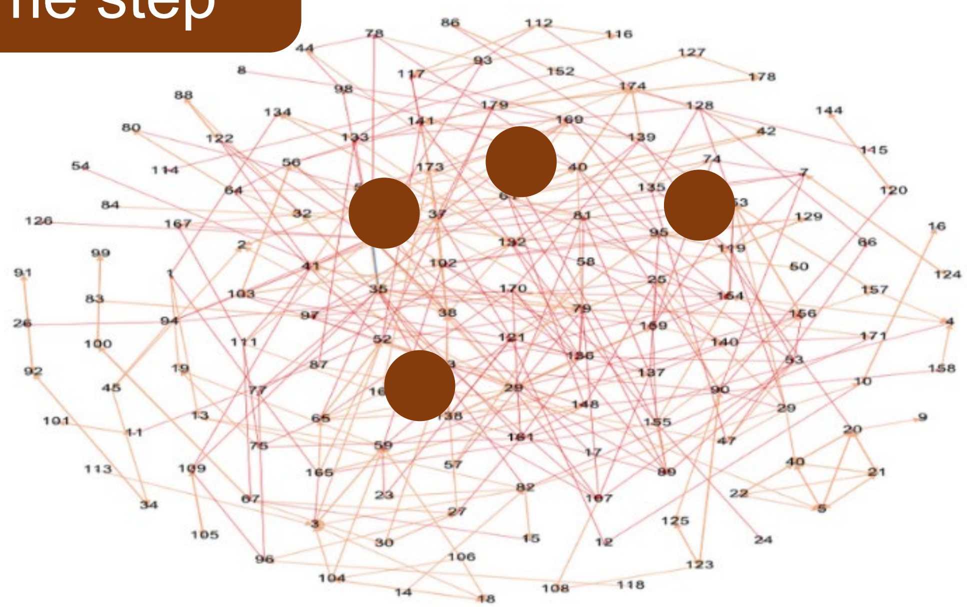

K = 4
2^{nd} time step

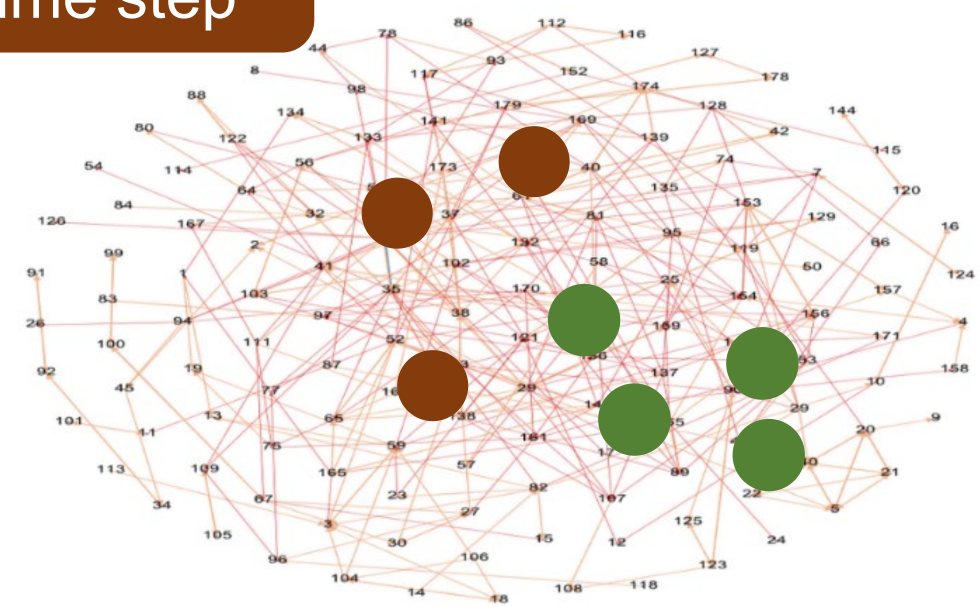

HIDDEN STATE

Action

Choose nodes

POMDP Policy

Observation: Node presence

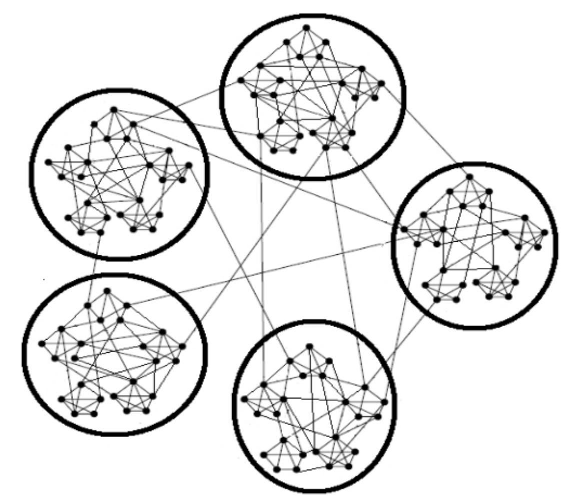

Partition POMDPs:
Exploit community structure

Challenge 3:
Sampling Networks: Exploratory Influence Maximization
(AAAI 2018)

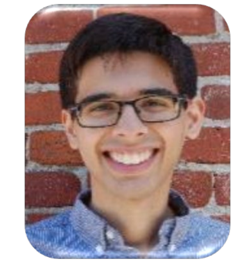

Wilder

Theorem: For community-structured graphs(*), sampling algorithm obtains a constant-factor approximation to the optimal influence spread using polylog(n) queries.

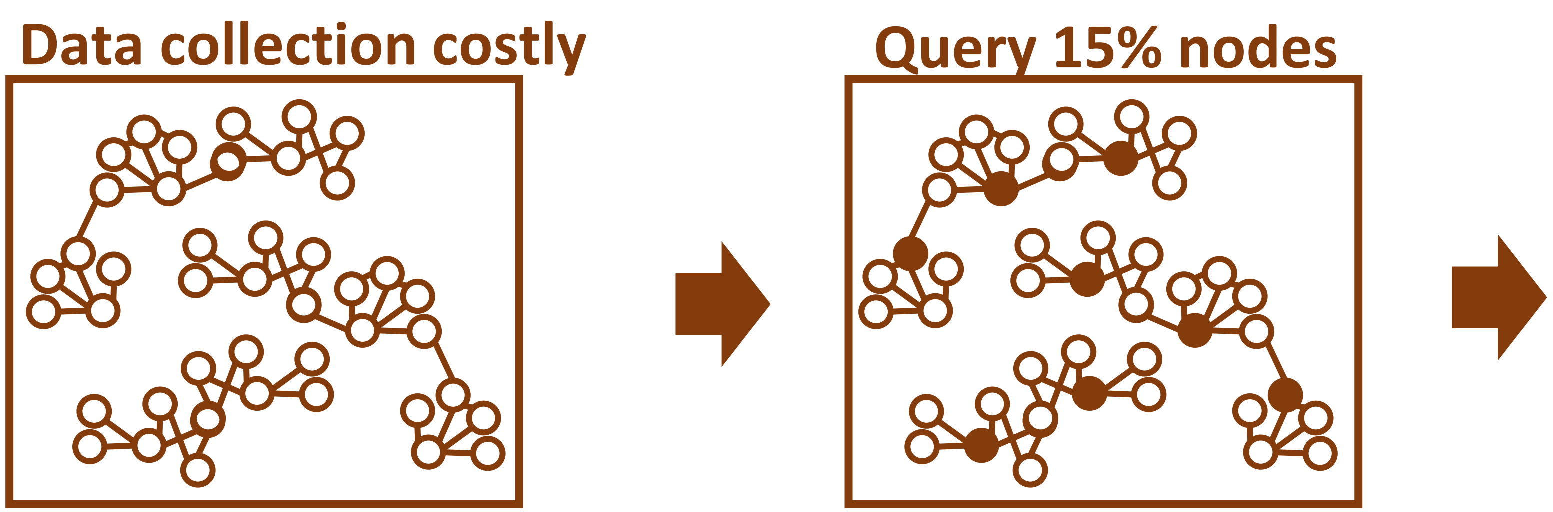

Sampling Algorithm

Sample node randomly & estimate size of its community;
Choose seeds from largest K communities

- Query 15% of nodes in the population
- Output K peer leader nodes to spread influence
- Perform similar to OPT, best influence spread with full network

(*)Community structured: drawn from a stochastic block model

Date: 5/30/2021

"Sampling-HEALER"
Pilot tests with Homeless Youth

(IJCAI 2018)

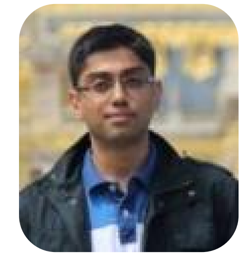

Yadav

Wilder

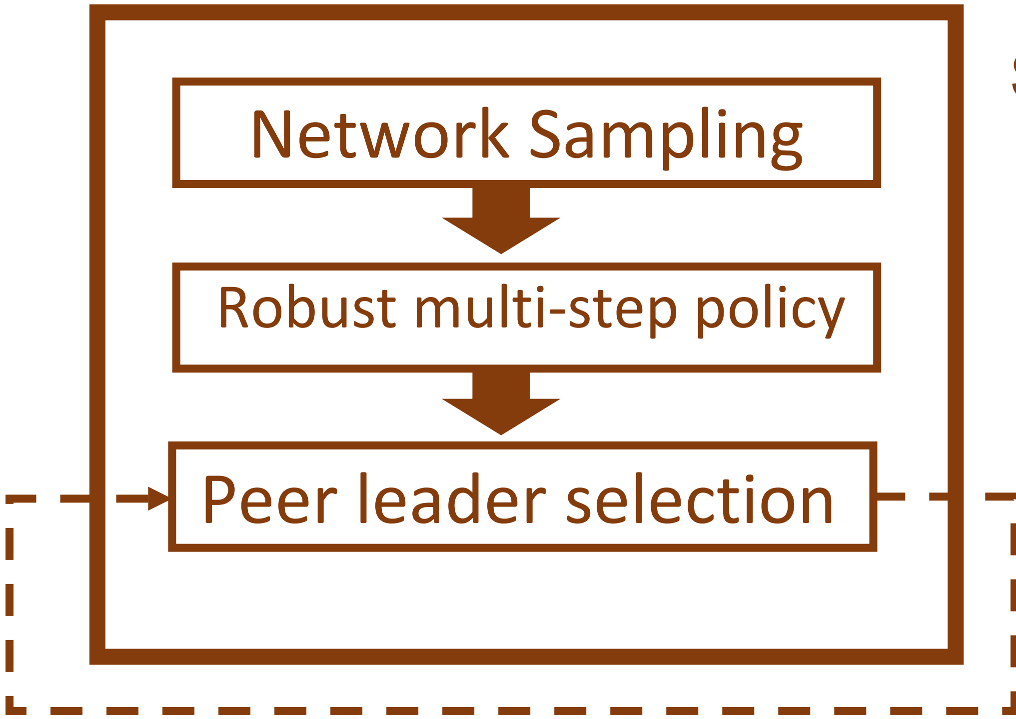

Sampling-HEALER

Observe peer leaders present/absent

12 peer leaders

Sampling HEALER (Sampled Network)	HEALER (Full Network)	DEGREE CENTRALITY (Full Network)
60 youth	62 youth	55 youth

Results: Pilot Studies

(Journal of Society of Social Work & Research 2018)


Yadav


Wilder

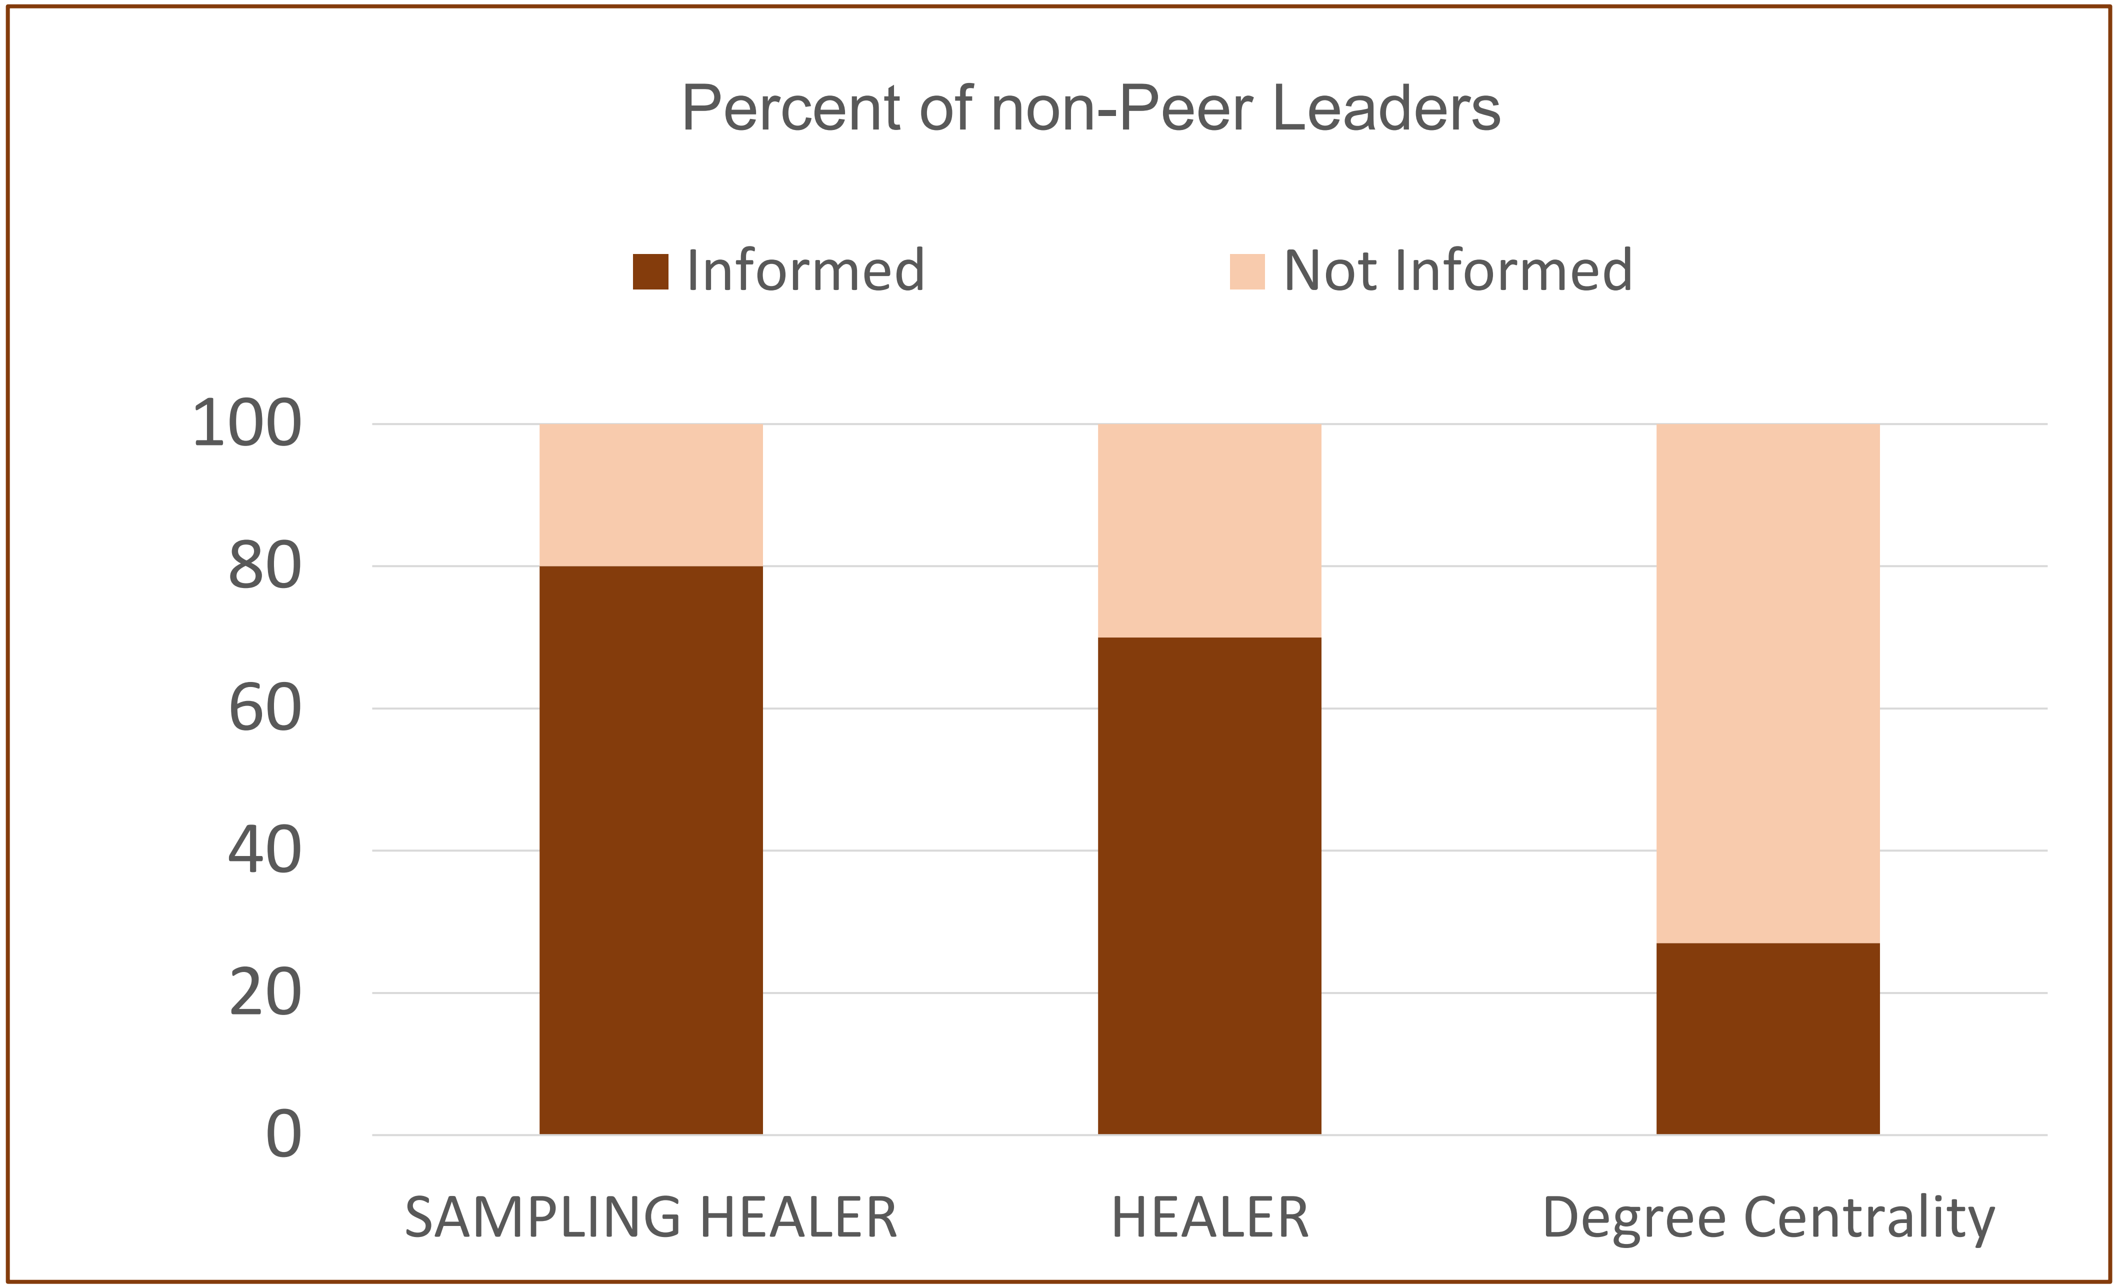

Results of 750 Youth Study [with Prof. Eric Rice]
Actual Change in Behavior?
(AAAI 2021)

First large-scale application of influence maximization for public health

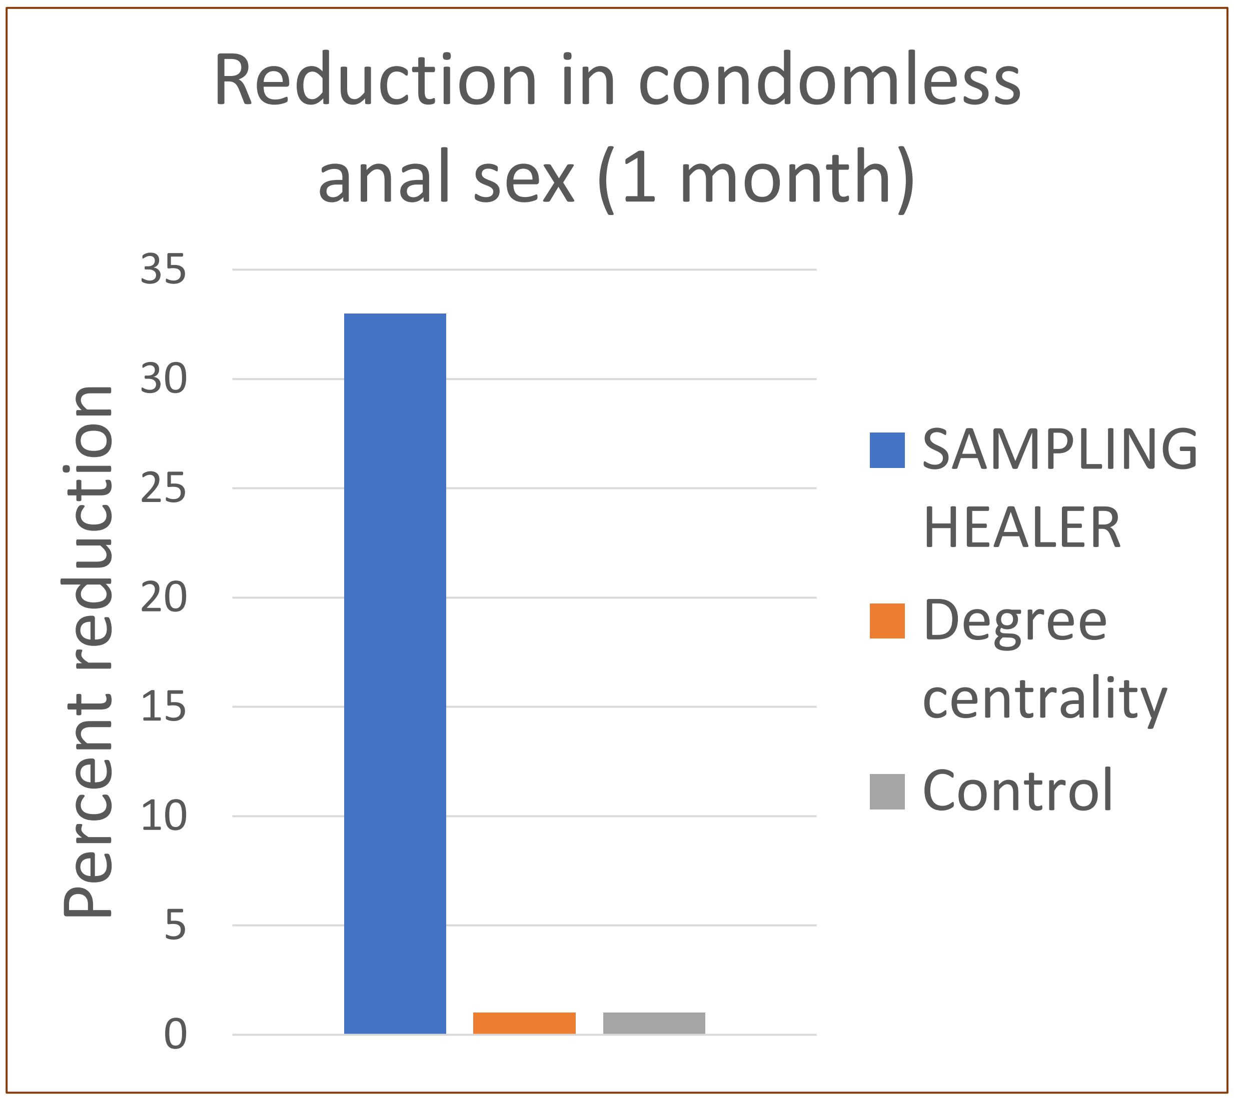

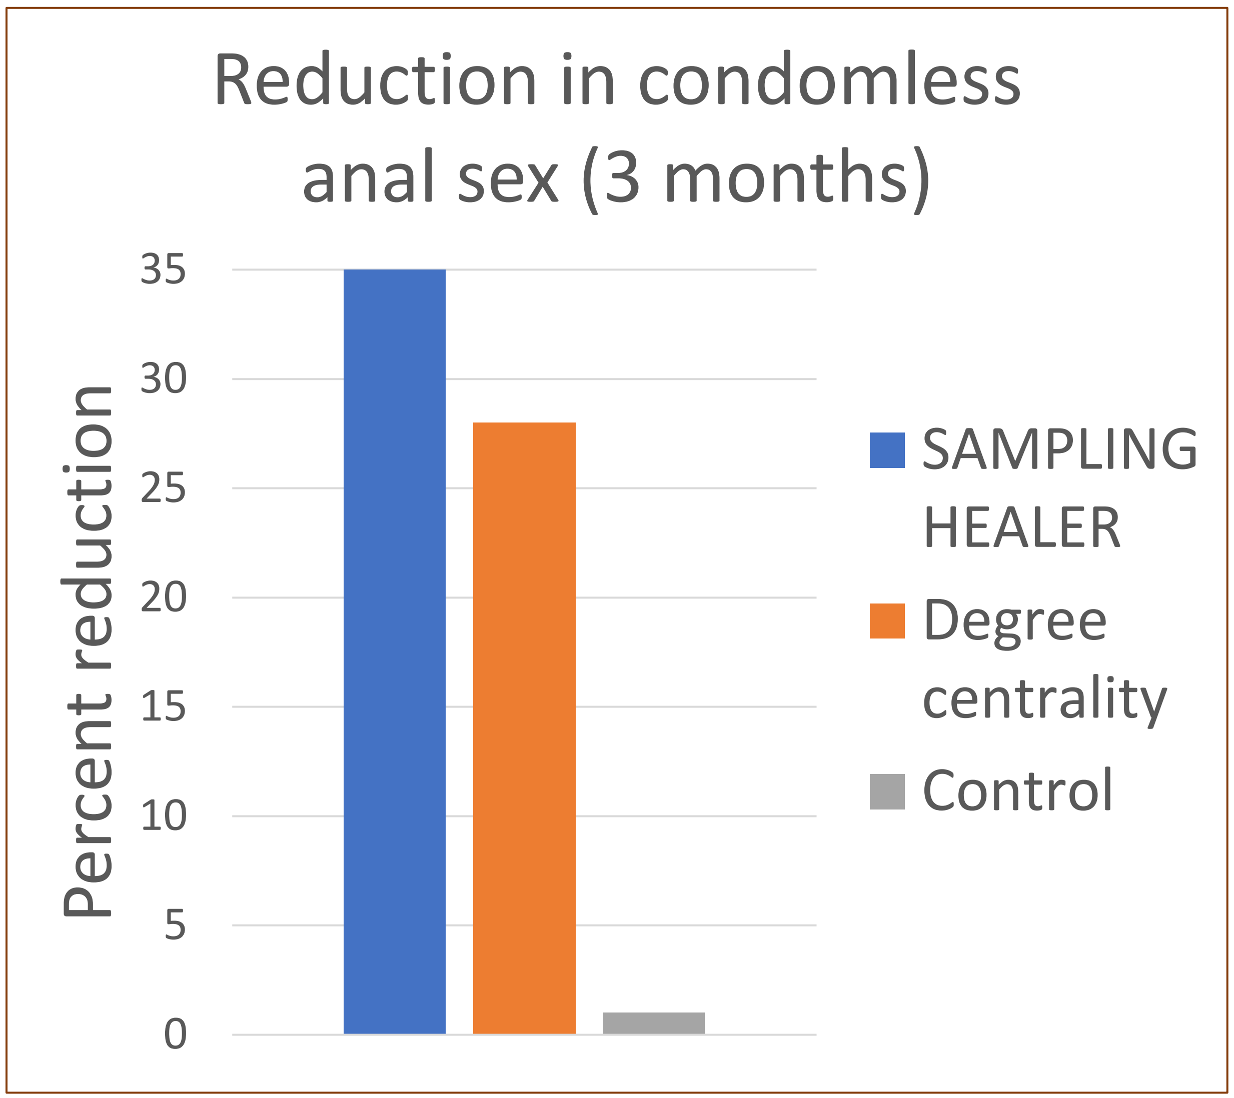

Results of 750 Youth Study [with Prof. Eric Rice]

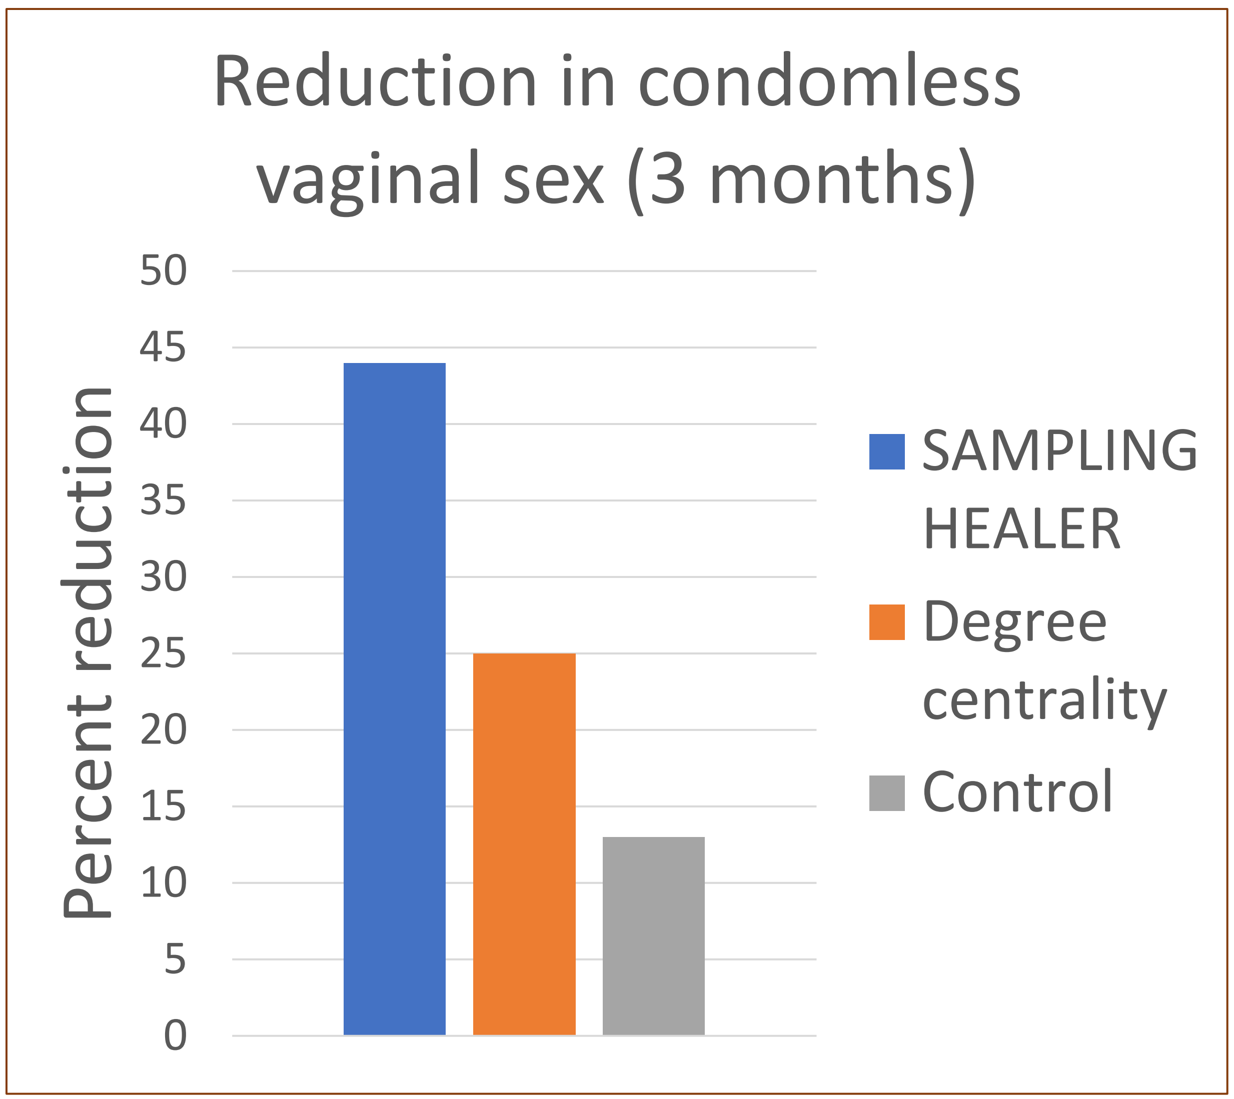

What our collaborators are saying:

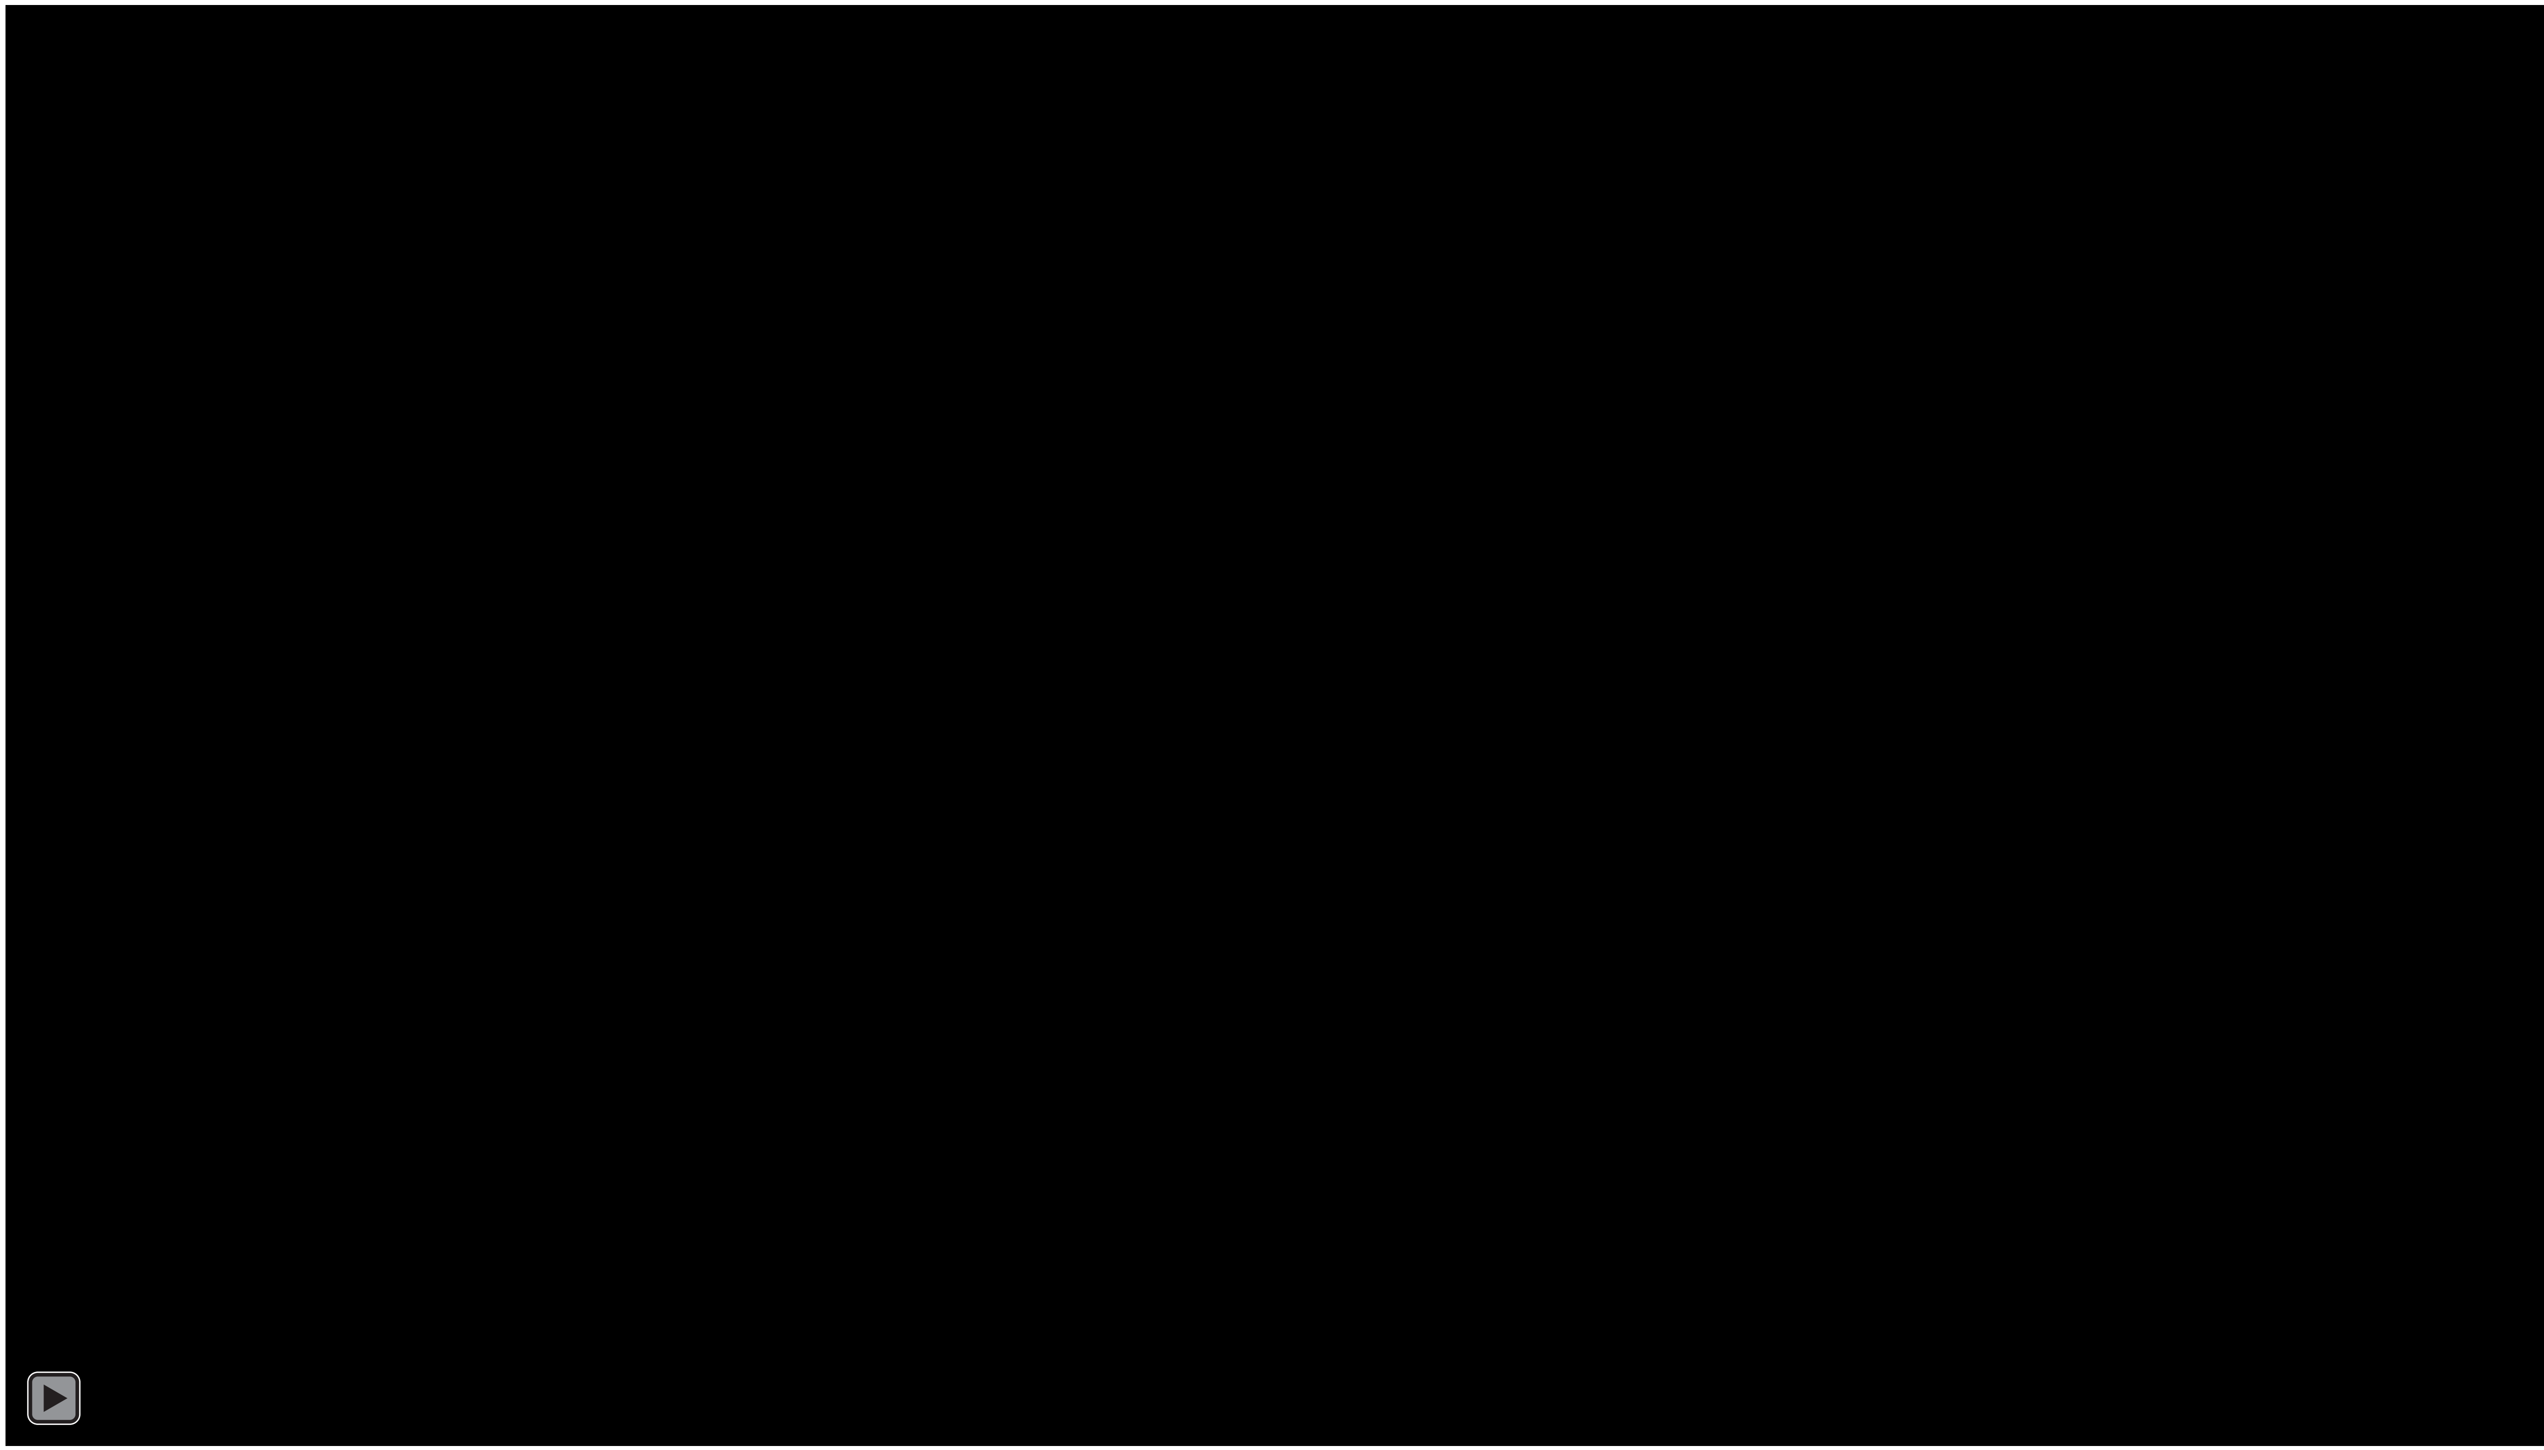

Next Steps: Fairness in Influence Maximization
(NeurIPS 2019, IJCAI 2019, AAAI 2021)

Rahmattalabi

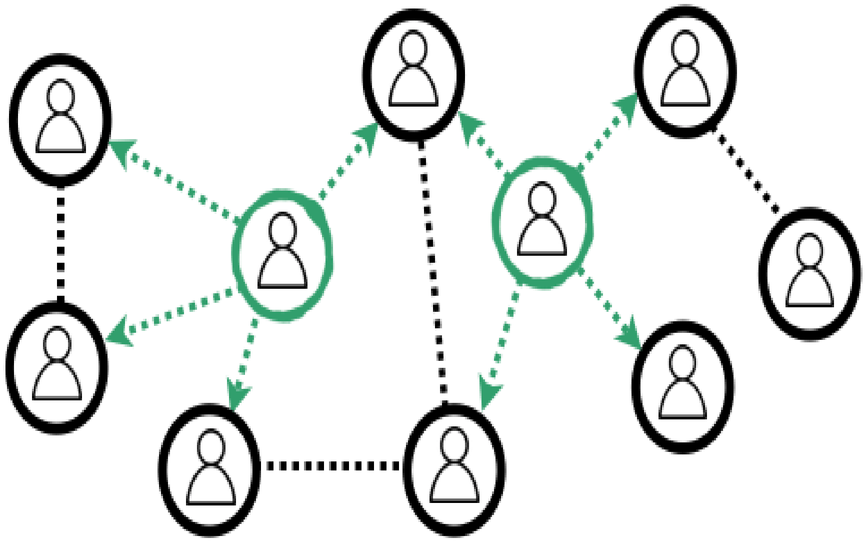

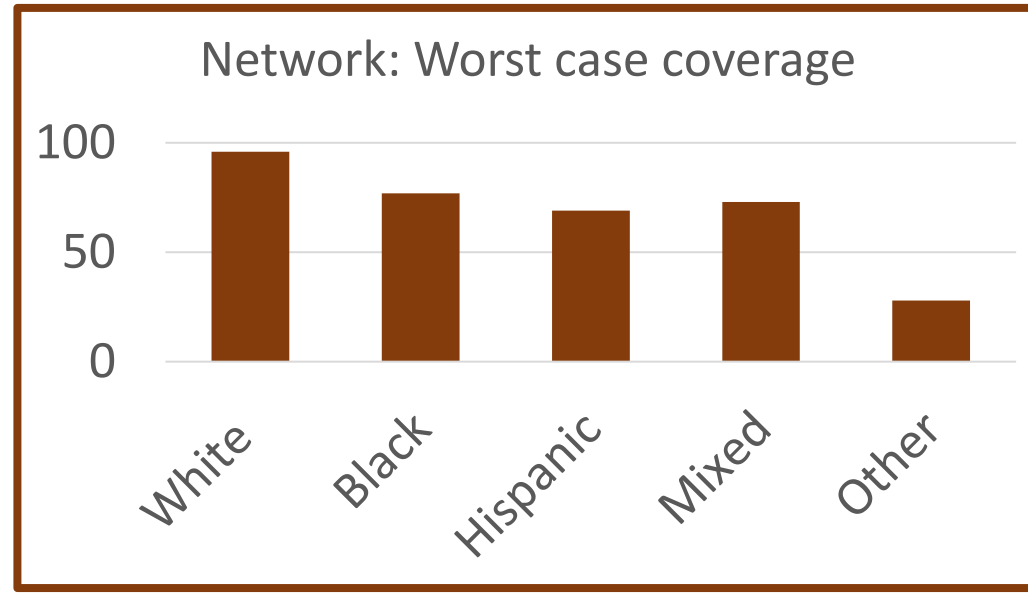

Robust graph covering, maximize worst case coverage → Disparity

Maxmin fairness:
NeurIPS2019

$$\min_{c \in C} u_c(A) \geq \gamma$$

Υ: Max of minimum utility for any community

Diversity constraints:
IJCAI2019

$$u_c(A) \geq U_c$$

U_c: Constraint from cooperative game theory

Inequity aversion:
AAAI 2021

$$W_\alpha\left(u(A)\right)$$

α controls fairness tradeoff; policymaker has choice

Next steps:RL for Influence Maximization in Social Networks

(with B. Ravindran & team, AAMAS 2020)

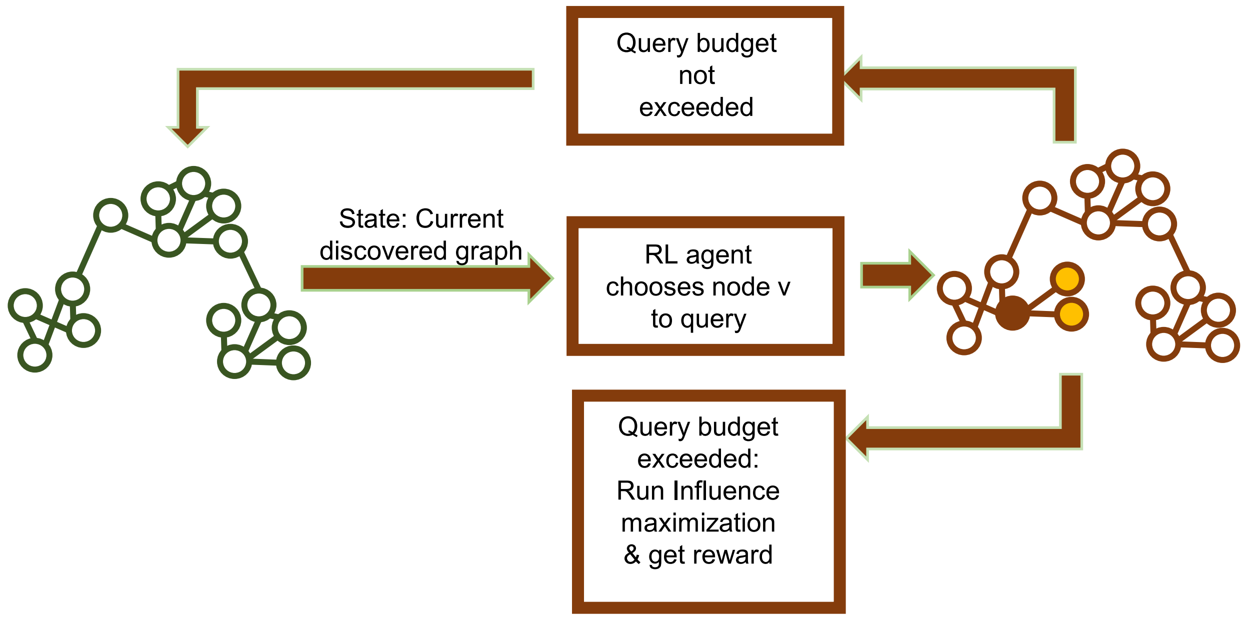

Network Family		Improve %
Rural		23.76
Animal		26.6
Retweet		19.7
Homeless		7.91

Outline

Public Health

- *Health information dissemination: Social networks*
- *Health program adherence: ML & Bandits*
- *COVID-19: Agent-based modeling*

Conservation

Intervention Reasoning: Active Adherence Monitoring

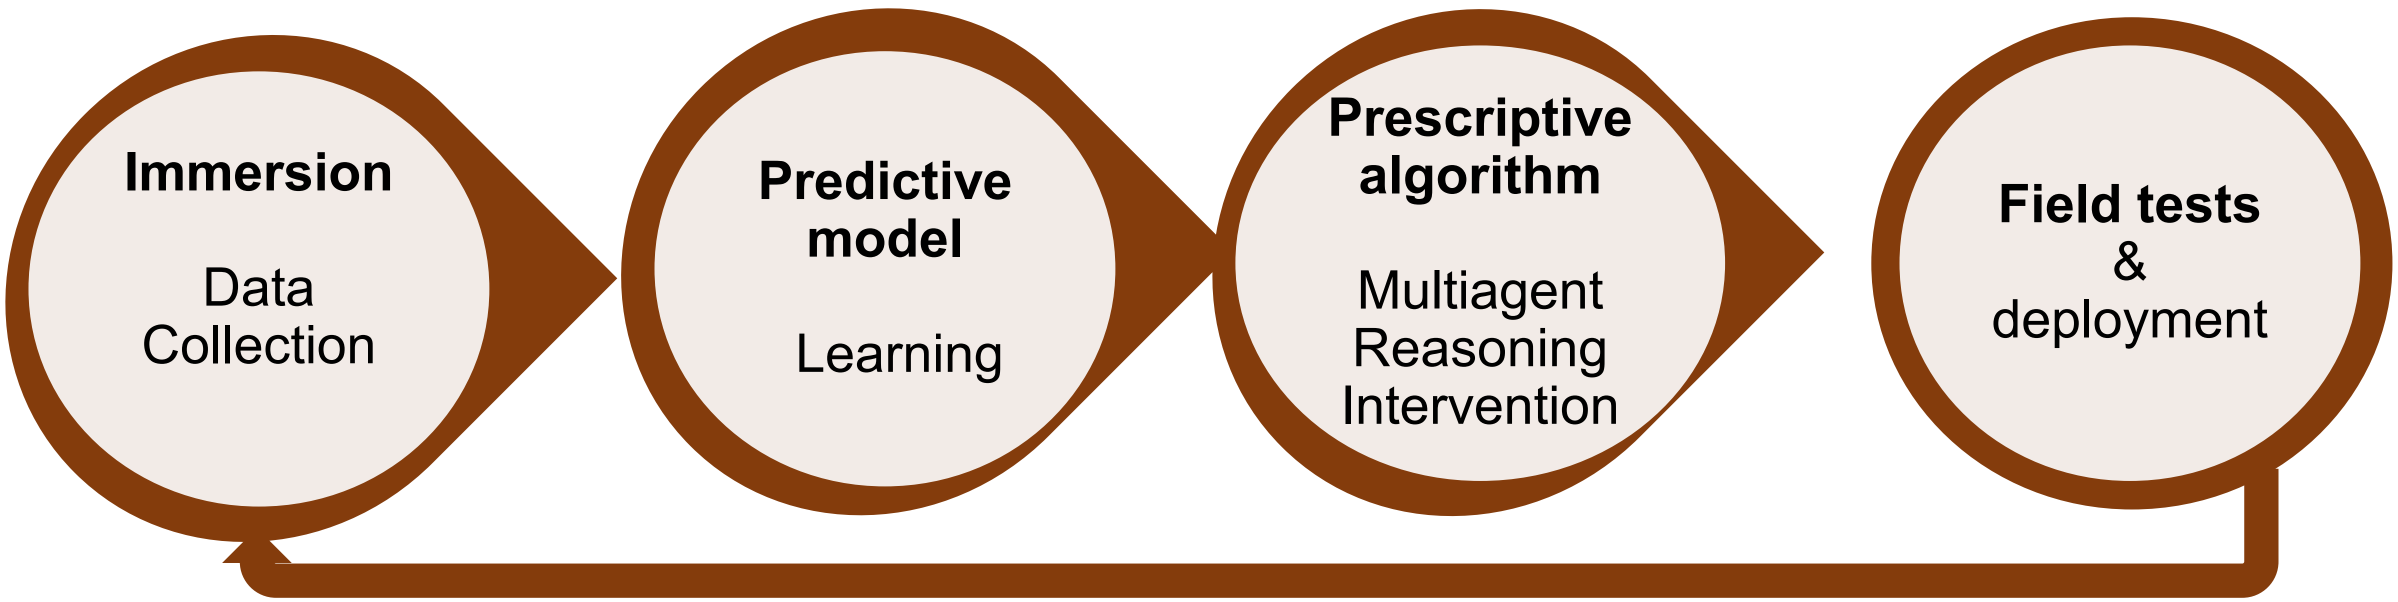

Health Program Adherence
Maternal & Child Care in India

(Under submission)

Woman dies in childbirth every 15 min; 4 of 10 children too thin/short

armman

(Advancing Reduction in Mortality and Morbidity of Mothers, Children and Neonates)

Delivering India's Future

18 Million women

Weekly 3 minute call
to new/expecting moms

mMitra: Significant benefits
2.2 million women enrolled

- ➢ *Unfortunately, significant fraction low-listeners or drop-outs*

Passive Adherence Monitoring
Maternal & Child Care in India

(with B Ravindran IIT Madras)

Classifier to predict beneficiaries drop out? So ARMMAN can focus interventions

- *Results of pilot with 18000 beneficiaries: High precision, recall, accuracy*
- *Field trial with 8000 beneficiaries: Call intervention helps*

- ***Prediction software in use to help 300,000 beneficiaries in mMitra***

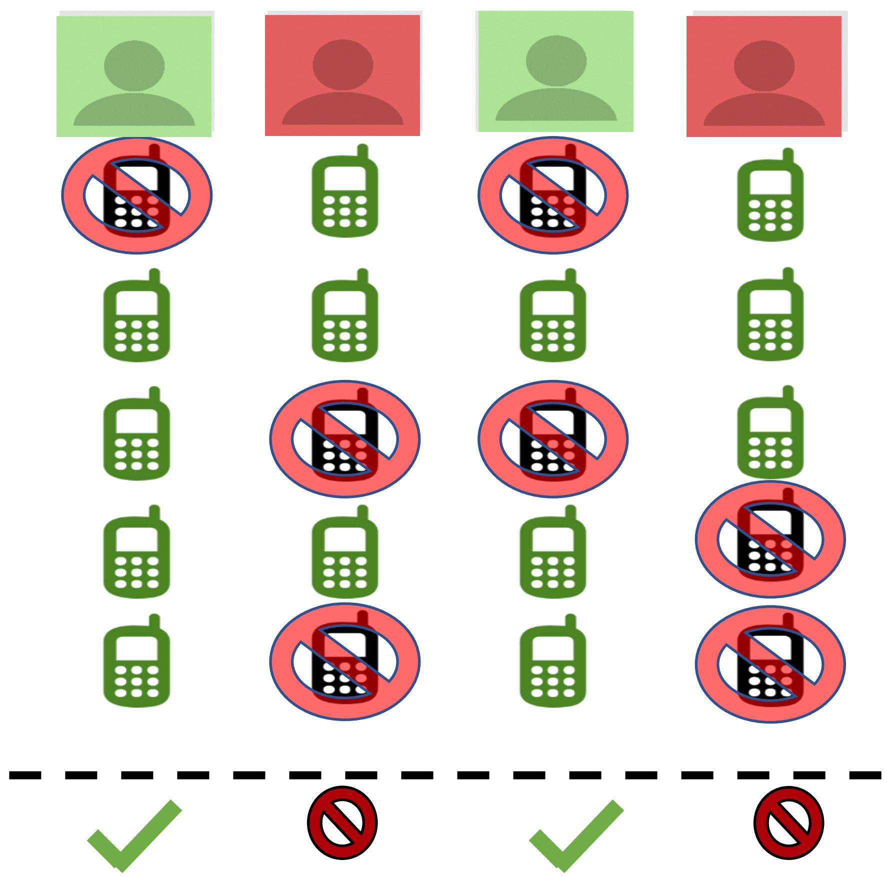

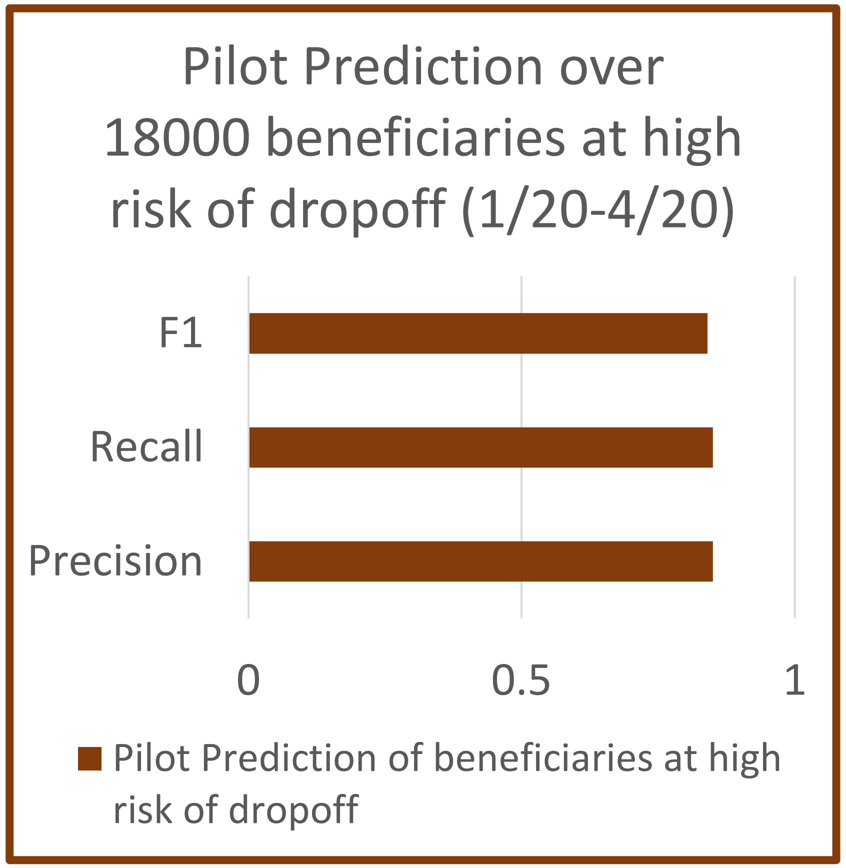

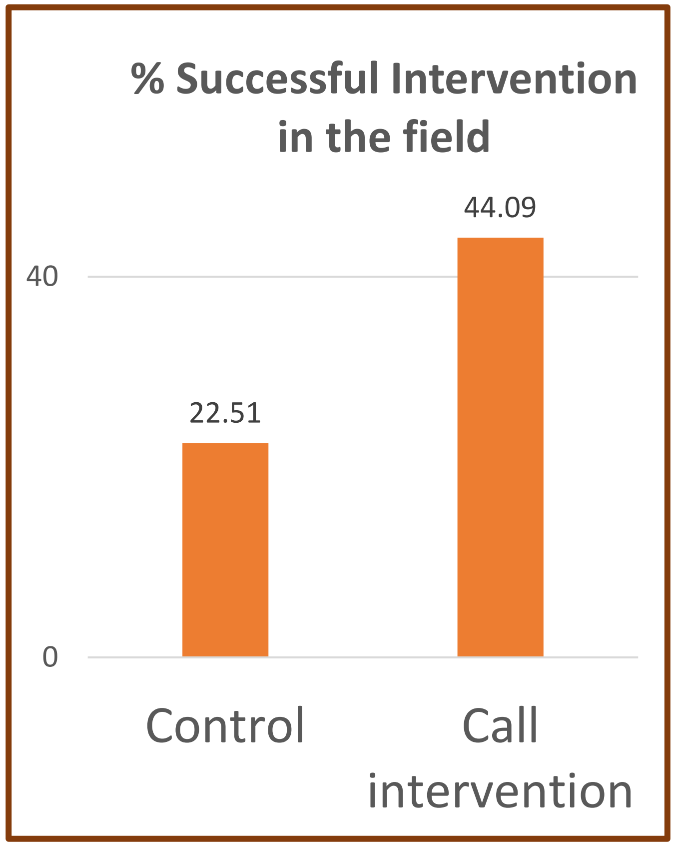

Passive Adherence Monitoring Preventing Tuberculosis in India

(KDD 2019)

Killian

Tuberculosis (TB): ~500,000 deaths/year, ~3M infected in India

- *Predict adherence risk from phone call patterns for early intervention?*

TB Treatment
6 months of pills

Track adherence
via daily phone calls

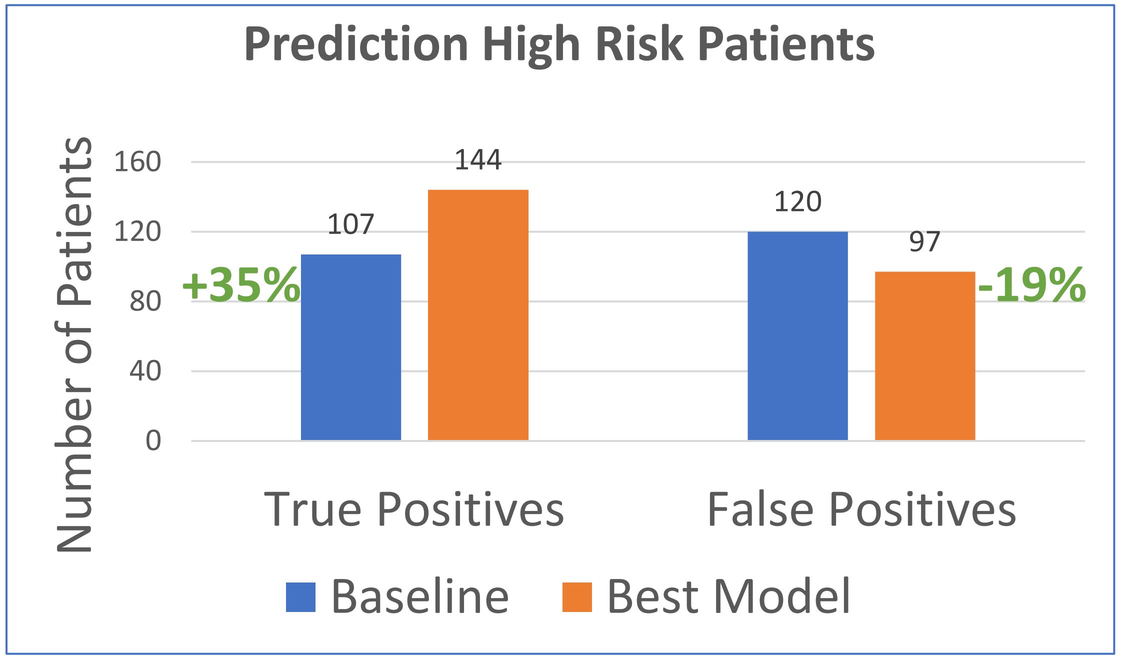

Results Mumbai, India:
15,000 patients, 1.5 Million calls

everwell

Intervention Reasoning: Active Adherence Monitoring

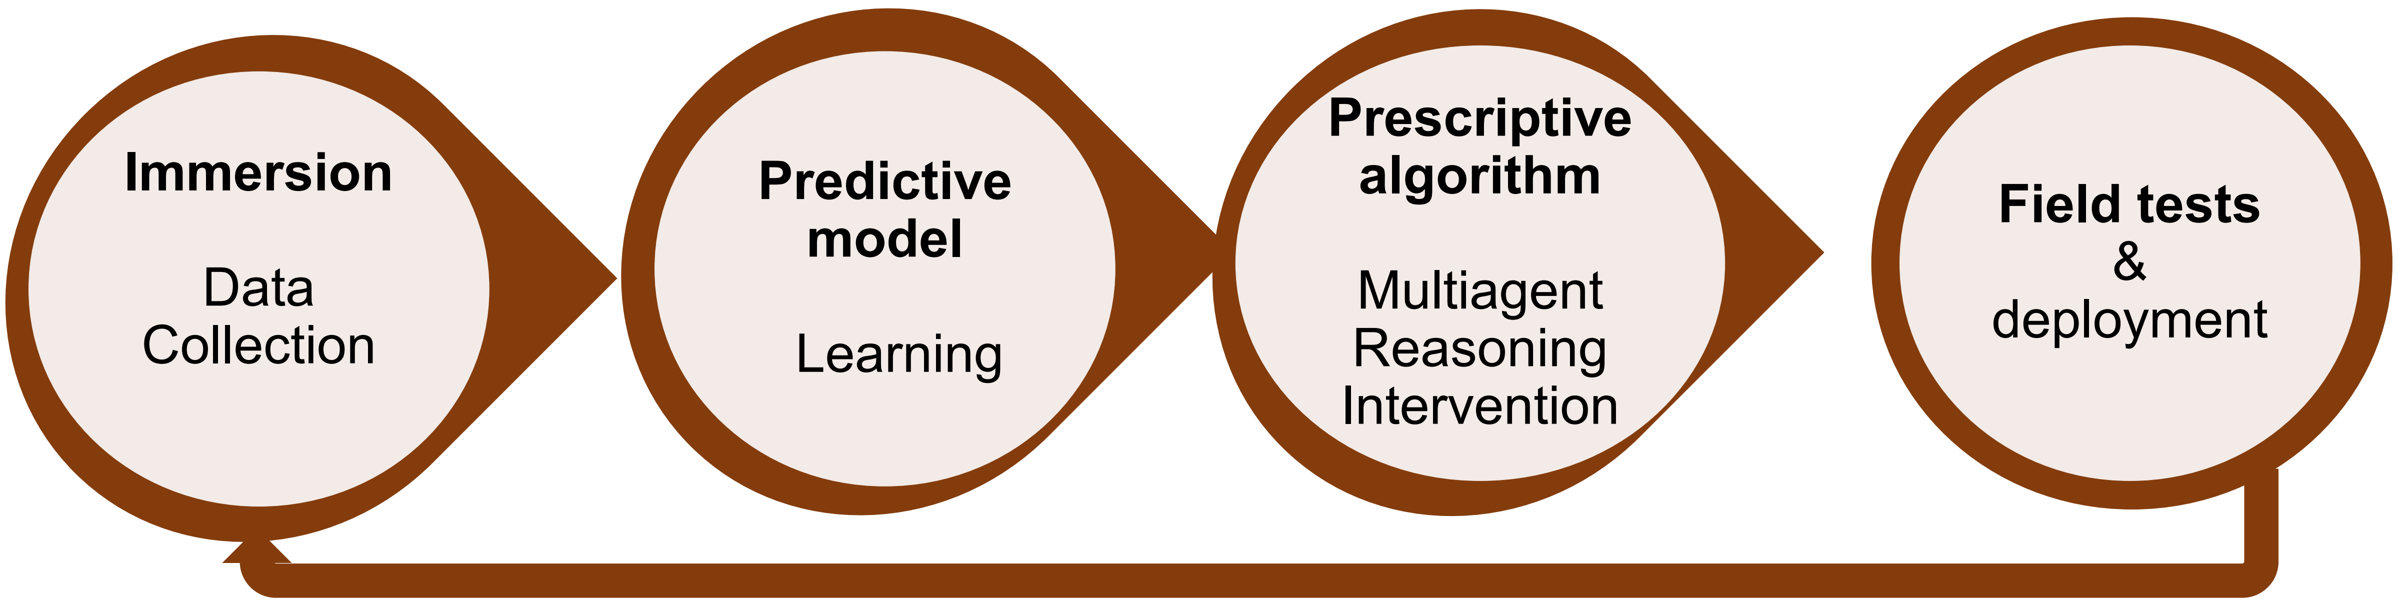

Intervention Scheduling with Scarce Data: Active Adherence Monitoring

(NeurIPS 2020)

Mate

Killian

Challenge:

- *Large number of patients (N)*
- *Can only call K patients per day.*
- *Which K?*

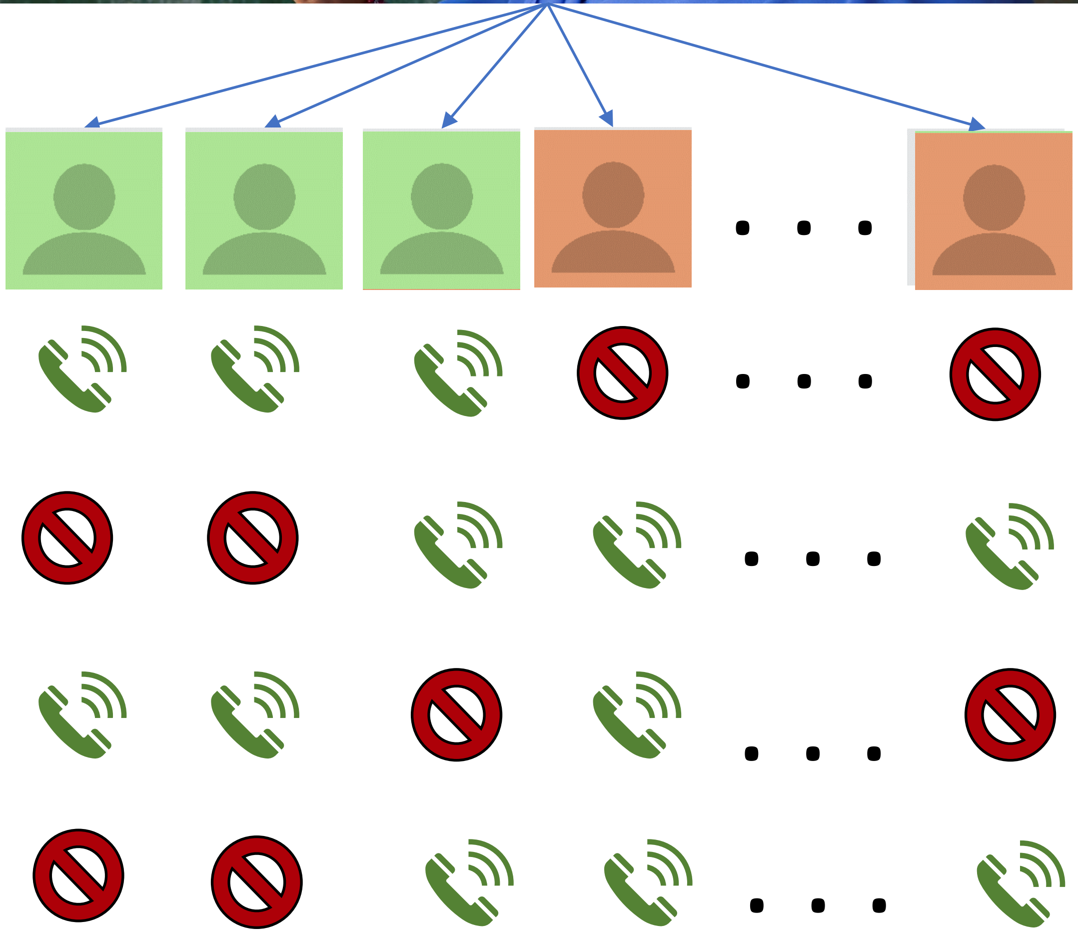

Approach: Restless bandit

- *Each arm (patient): POMDP*
- *Each POMDP: binary latent state {0, 1}*
- *0= not-adhering; 1= adhering*
- *Two actions*

Goal: Policy for K patients to call per day

Photo Credit: IntraHealth International (CC BY-NC-SA 3.0 via https://www.intrahealth.org/)

Intervention Scheduling with Scarce Data: Collapsing Bandits

(NeurIPS 2020)

Mate

Killian

Theorem (Whittle Index): Collapsing bandits are Indexable if threshold policies are optimal.

When arm not played (patient not called)

- *No observation*
- *Instead, compute belief of adherence*

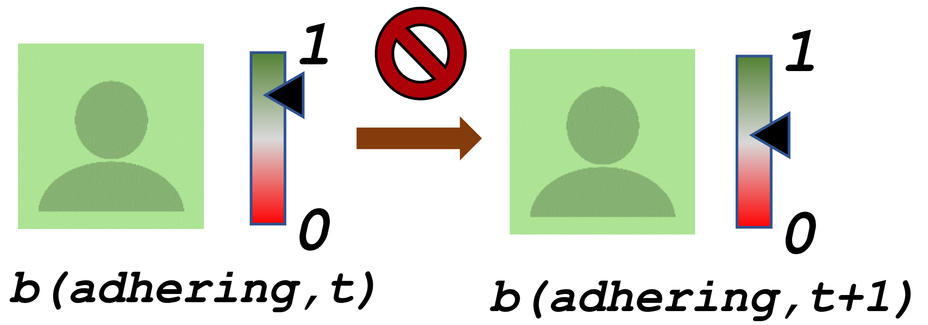

When arm is played: Uncertainty collapse

- *Observe current state*
- *Belief of adhering next round*

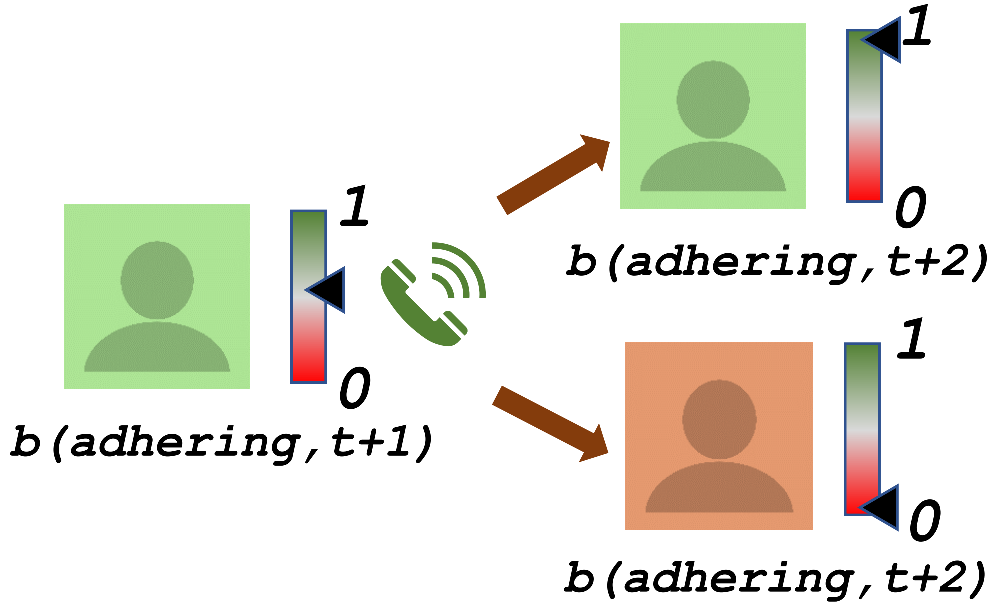

- ***Exploit "collapsing" for fast algorithm: Fixed number of belief states***

New Fast Algorithm: Collapsing Bandits

Mate

Killian

- ***Orders of magnitude speedup with no solution quality loss***
- ***ORANGE = Best baseline***
- ***Blue = Our model***

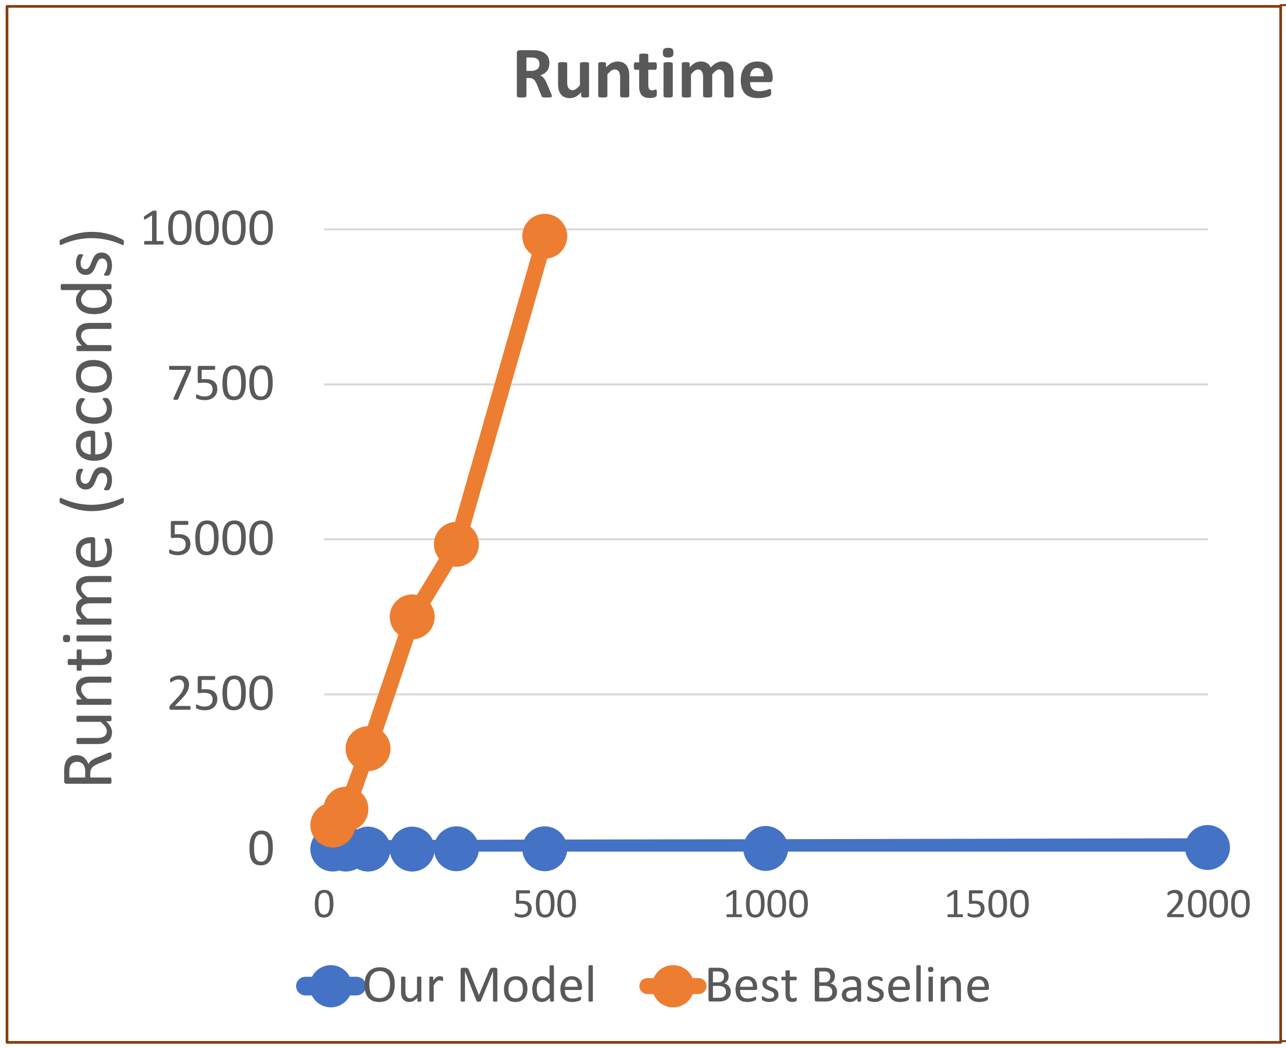

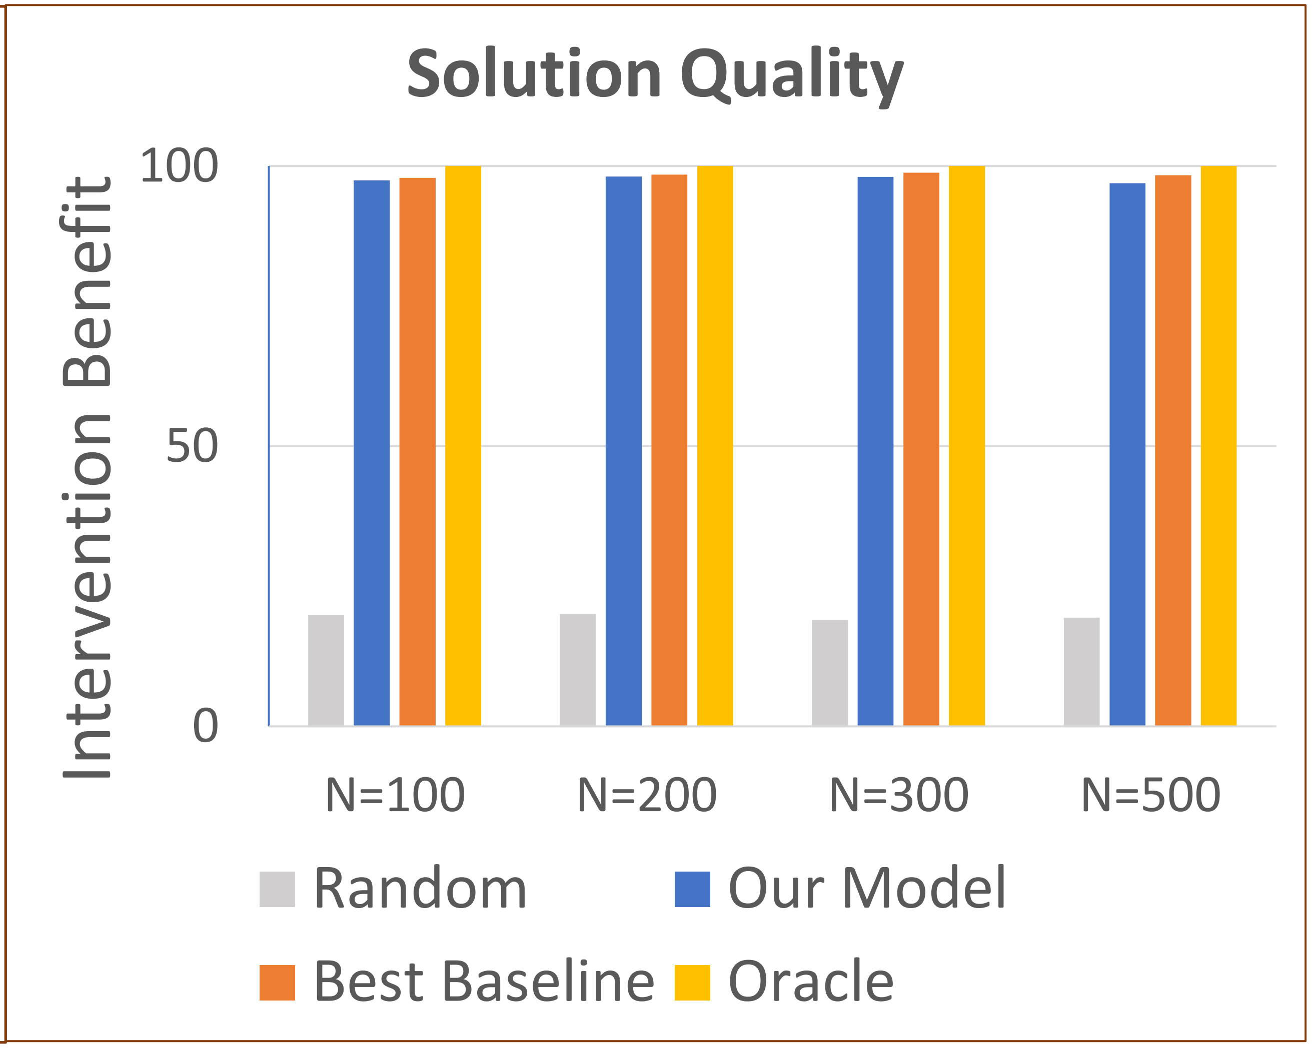

New Directions in Restless Bandits

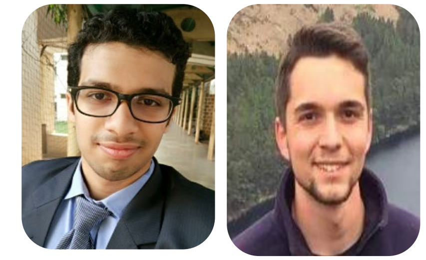

Mate Killian Biswas

Fast algorithms for extending to:

- *Multiple action types*

  (AAMAS 2021a)

- *Risk aware restless bandits*

  (AAMAS 2021b)

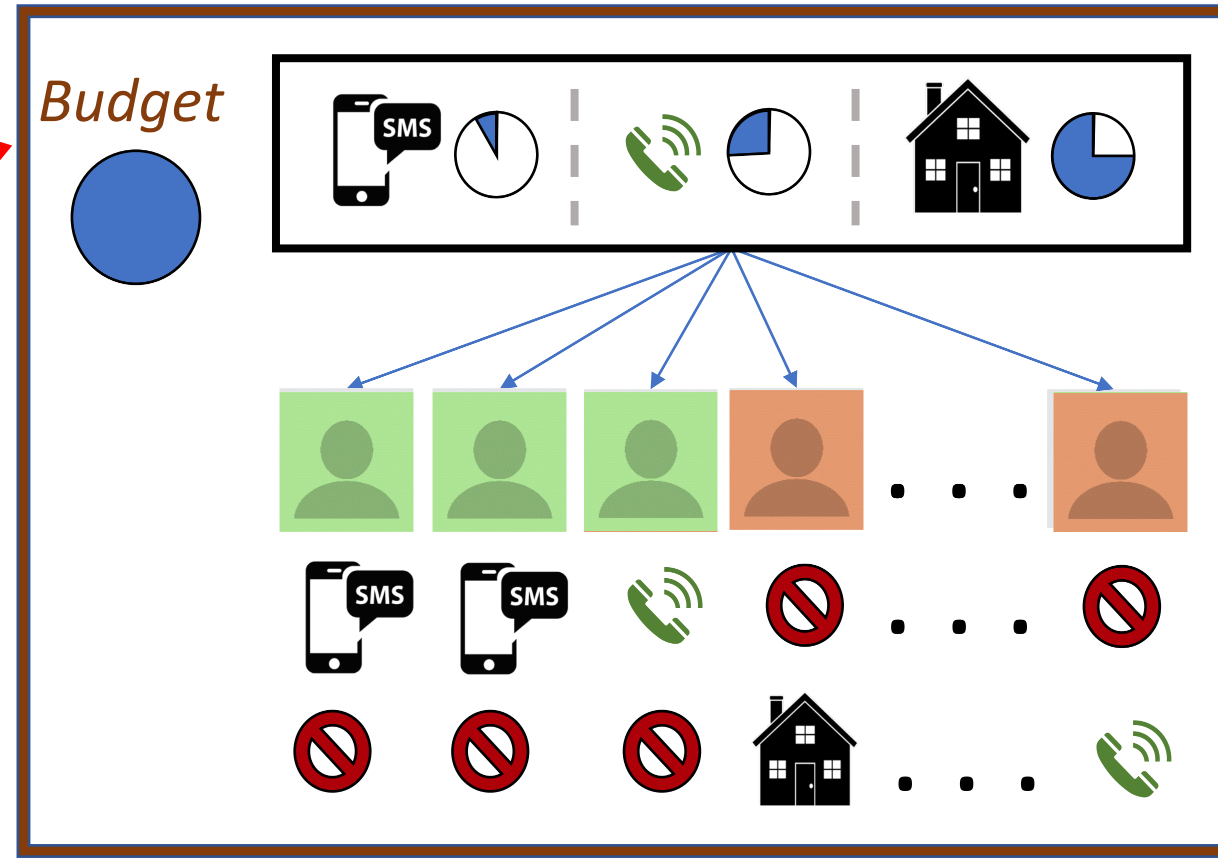

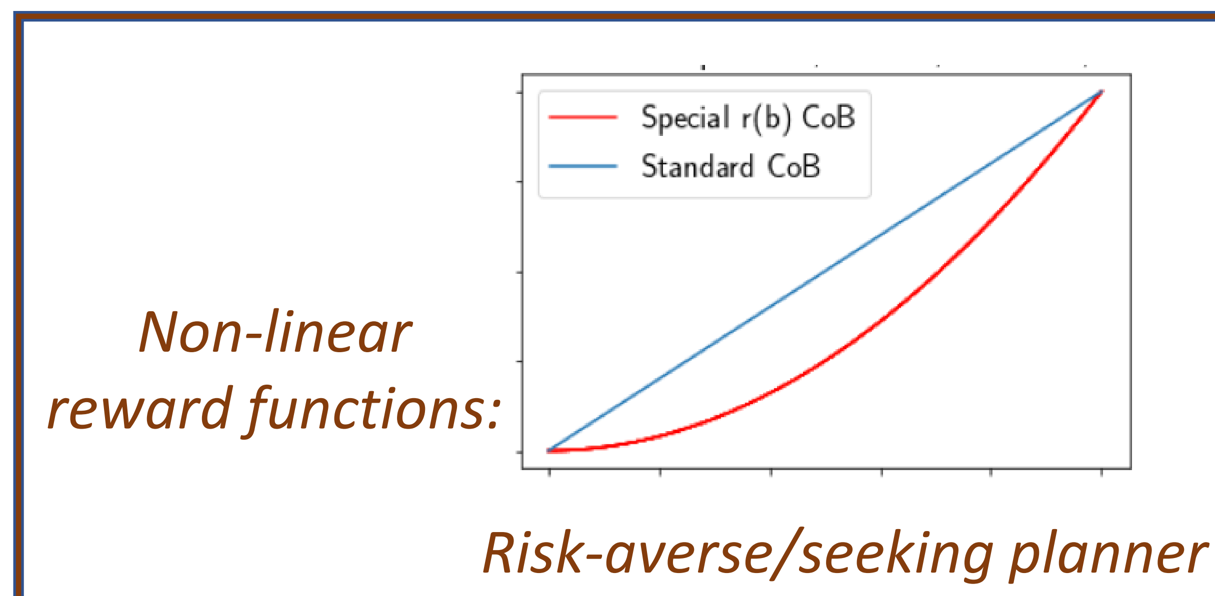

Online learning

- *Learning policies via Index Q-Learning*

  (AAMAS 2021c)

New Directions in Restless Bandits

Restless bandits for intervention:

- *10000 subject trial*

Outline

Public Health

- *Health information dissemination: Social networks*
- *Health program adherence: ML & Bandits*
- *COVID-19: Agent-based modeling*

Conservation

COVID-19: Agent-based Simulation Model

Wilder

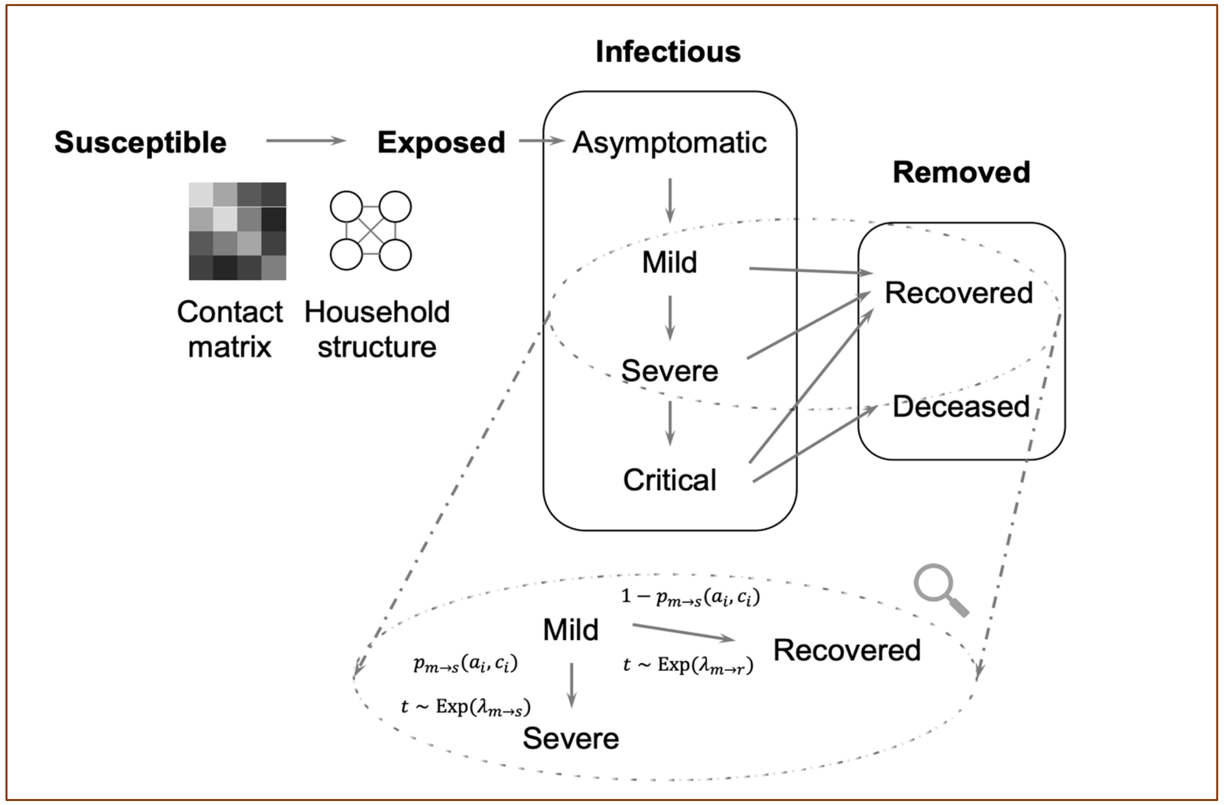

PNAS Proceedings of the National Academy of Sciences of the United States of America

RESEARCH ARTICLE

Modeling between-population variation in COVID-19 dynamics in Hubei, Lombardy, and New York City

Bryan Wilder, Marie Charpignon, Jackson A. Killian, Han-Ching Ou, Aditya Mate, Shahin Jabbari, Andrew Perrault, Angel N. Desai, Milind Tambe, and Maimuna S. Majumder

PNAS October 13, 2020 117 (41) 25904-25910; first published September 24, 2020;
https://doi.org/10.1073/pnas.2010651117

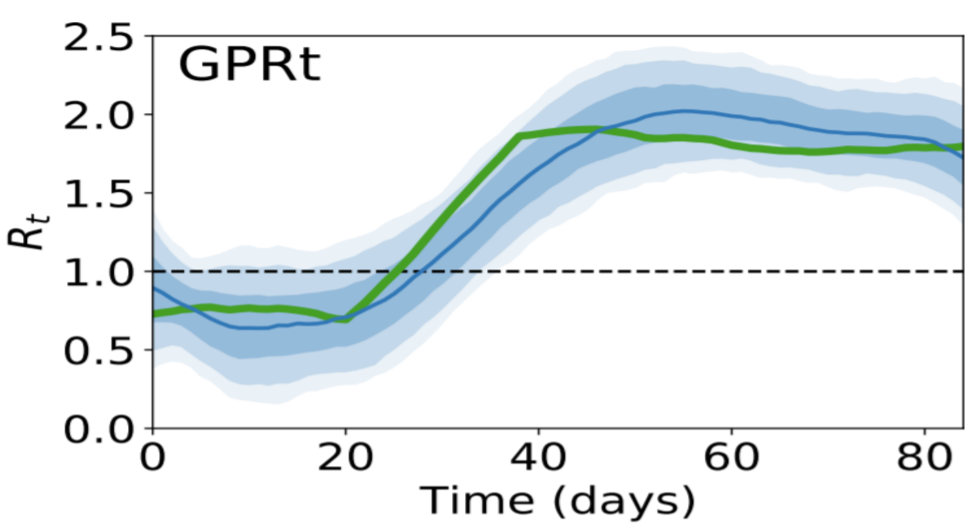

Tracking disease outbreaks from sparse data with Bayesian inference

Bryan Wilder,[1] Michael Mina[2], Milind Tambe[1]
[1] John A. Paulson School of Engineering and Applied Sciences, Harvard University
[2] T.H. Chan School of Public Health, Harvard University
bwilder@g.harvard.edu, mmina@hsph.harvard.edu, milind_tambe@harvard.edu

COVID Testing Policy: Accuracy vs Ease

(Science Advances, 2020) with Prof. Michael Mina


Wilder

- Range of tests entering market, varying sensitivity/cost: Quantity vs Quality?
  - qRT-PCR ("gold standard"): Detect viral concentration of **10^3**/mL, \$50-100
  - Antigen strip ("Less sensitive"): **10^6**/mL, \$3-5

Rapid turnaround time & frequency more critical than sensitivity for COVID-19 surveillance

Less sensitive; Cheap & fast turnaround **More sensitive; Costly & slow turnaround**

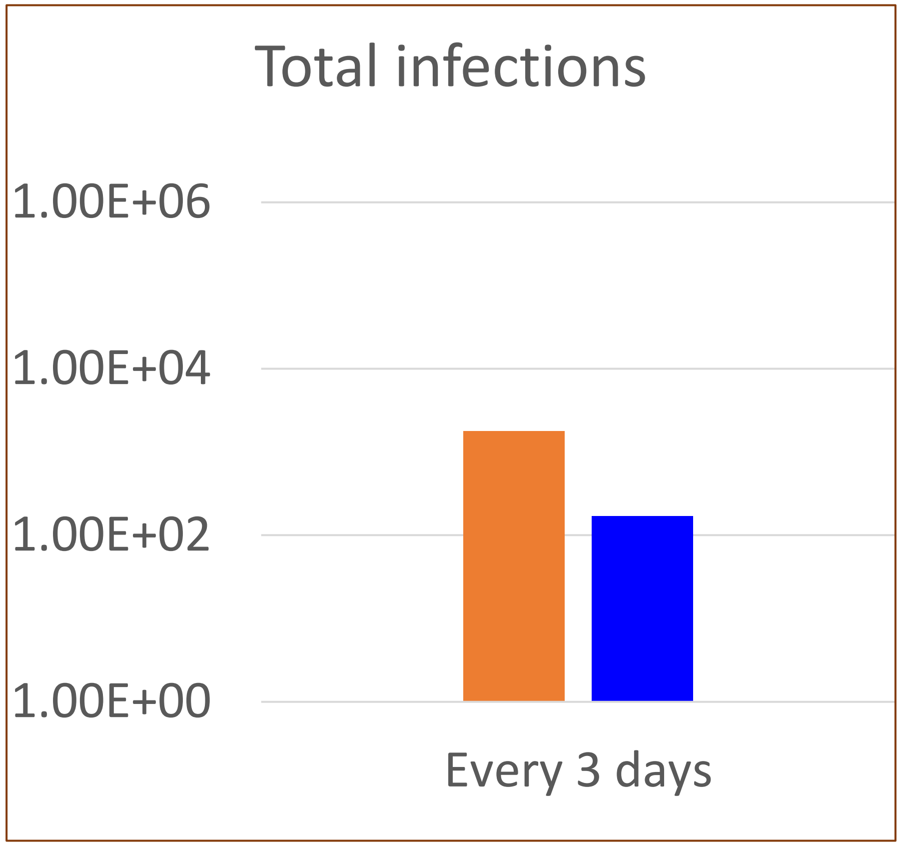

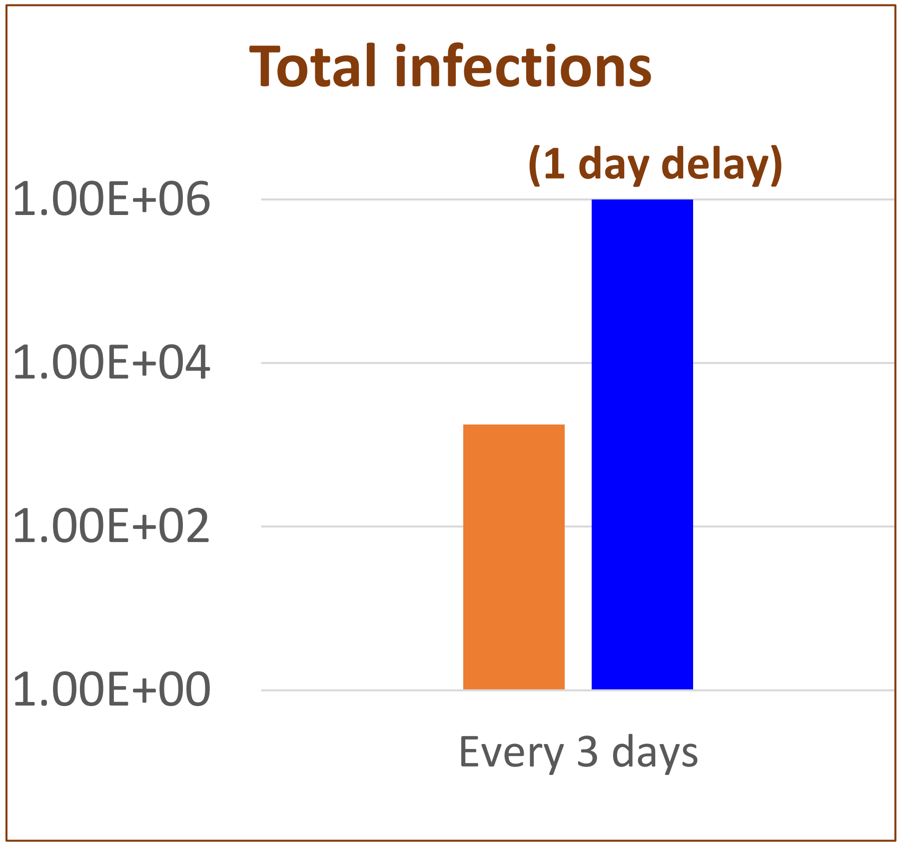

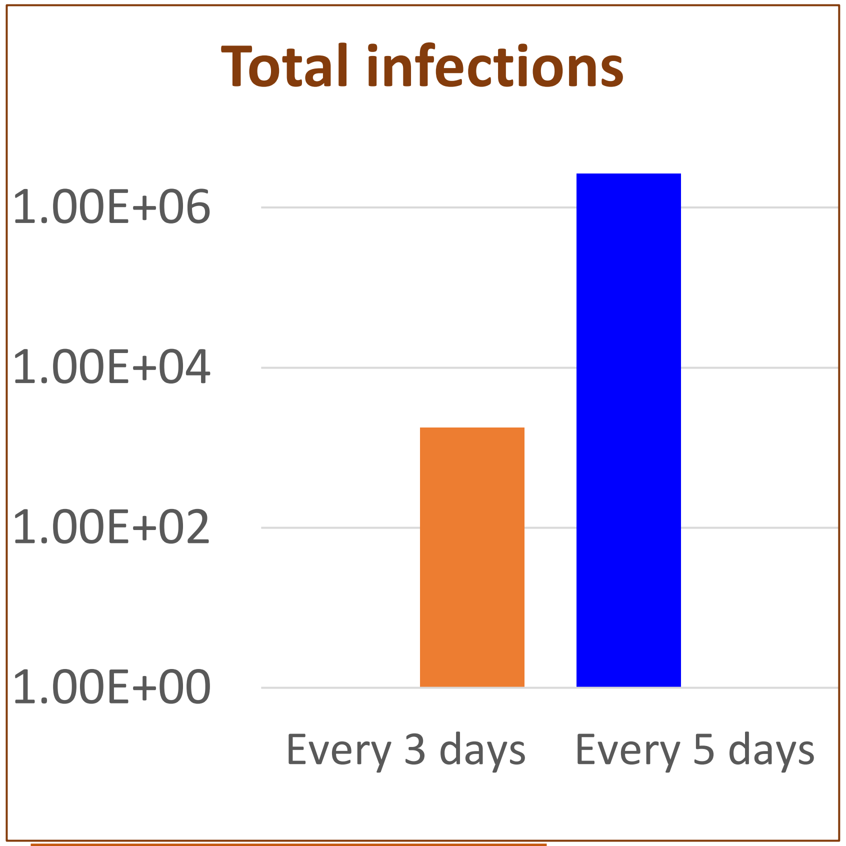

COVID Testing Policy: Impact

- Covered in NYT, WaPo, Time, The Atlantic, The Hill, etc
- Allowed epi collaborators to advocate to FDA/CDC

Outline

Public Health

Conservation

- *Protect wildlife, forests, fisheries: Game-focused learning*
- *Integrating real time data for protection: Signaling games*

Protecting Conservation Areas: Green Security Games
(IJCAI 2015)


Fang

Snare or Trap

Wire snares

From Stackelberg Security Games to Green Security Games
(IJCAI 2015)

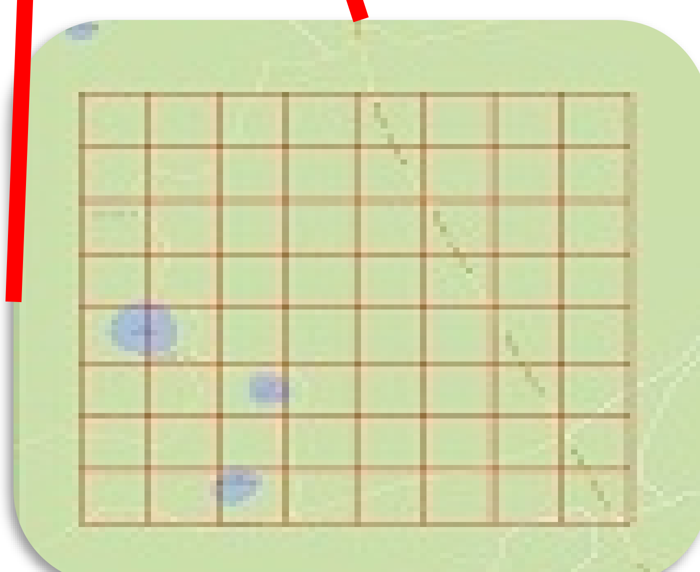

- ***Stackelberg security games (SSG)***
- ***With boundedly rational poachers***
- ***Learn adversary response model at targets "i"***

	Area1	Area2
Area1	4, -3	-1, 1
Area2	-5, 5	2, -1

Learning Adversary Response Model: Uncertainty in Observations

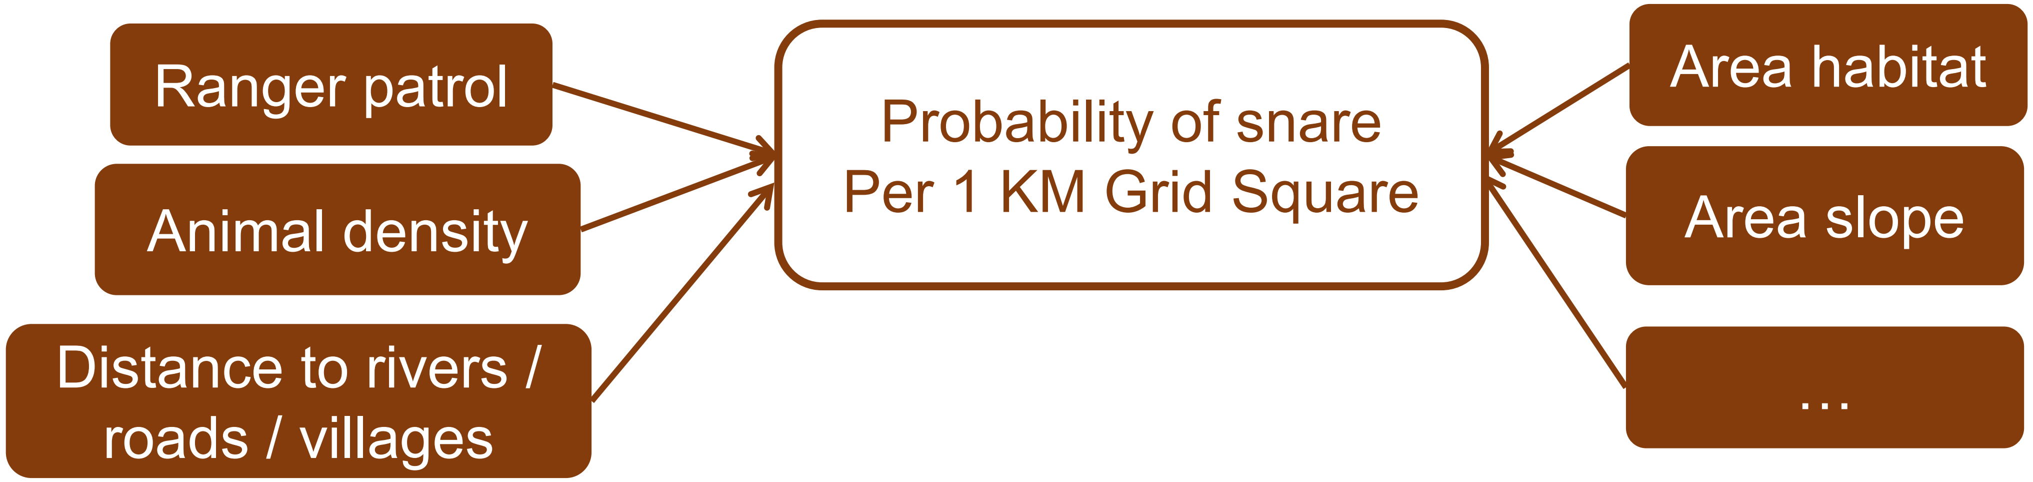

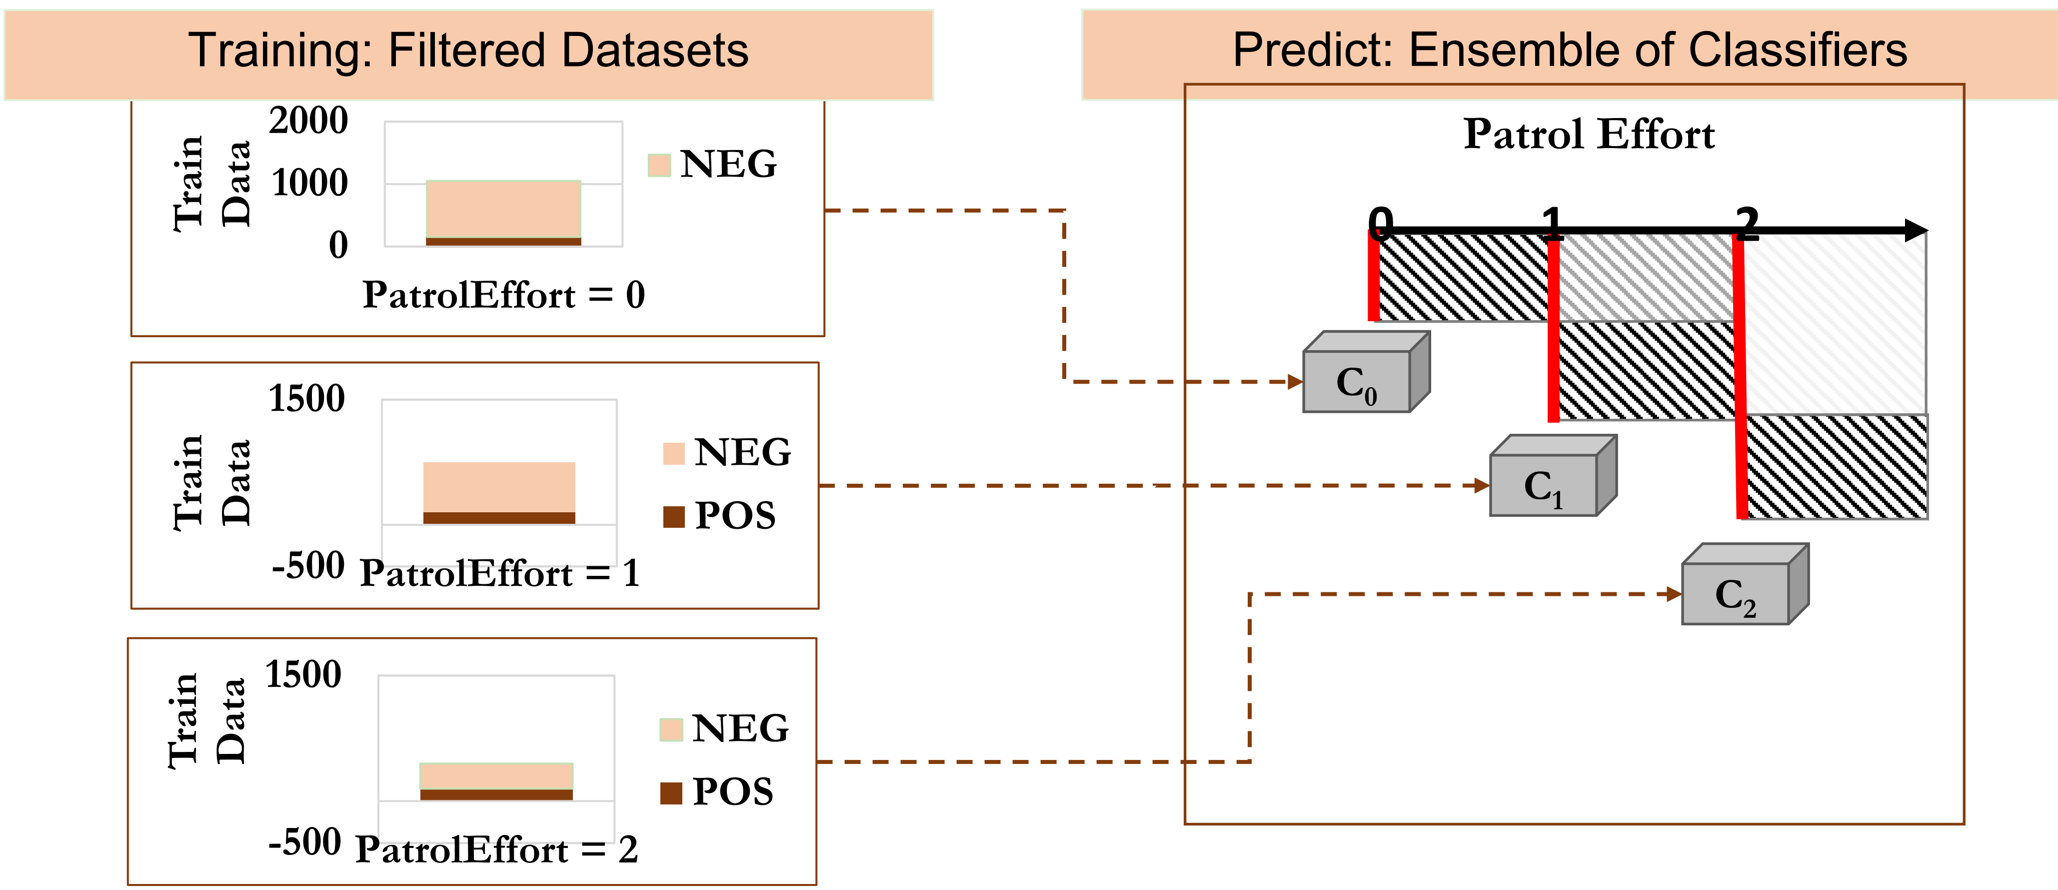

PAWS: First Pilot in the Field
(AAMAS 2017)


Ford


Gholami

- Two 9-sq.km areas, infrequent patrols

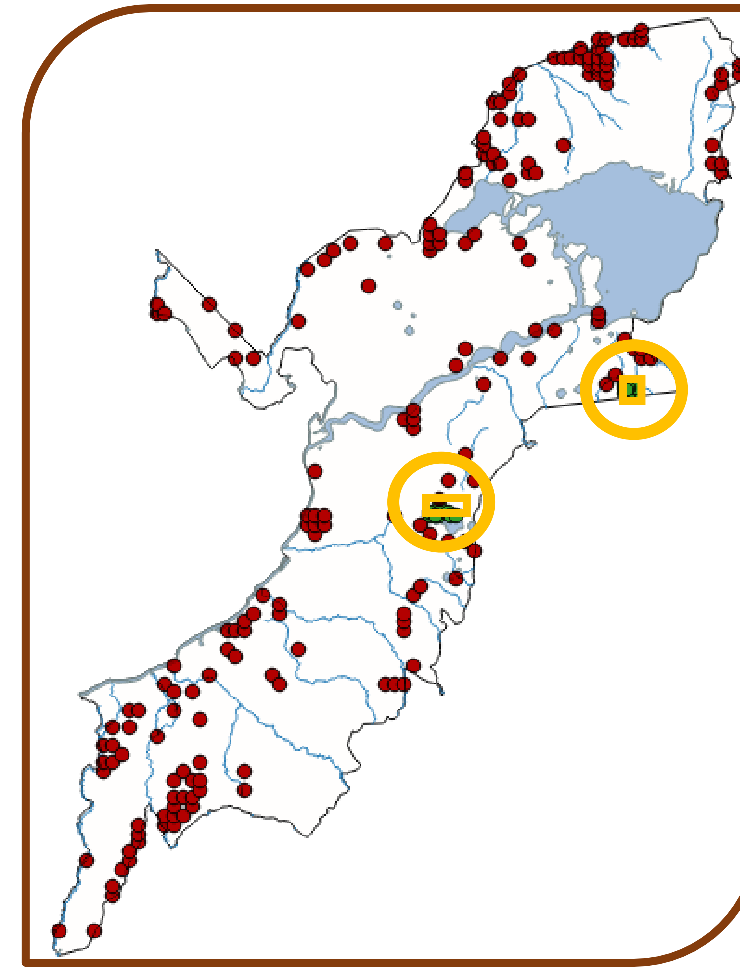

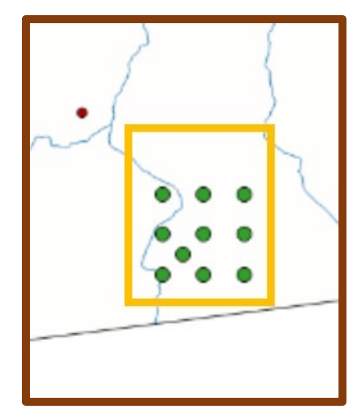

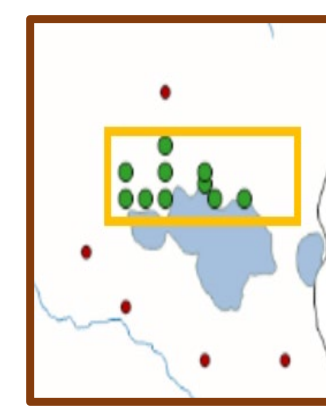

- Poached elephant
- 1 elephant snare roll
- 10 Antelope snares

PAWS Predicted High vs Low Risk Areas: 3 National Parks, 24 areas each, 6 months

(ECML PKDD 2017, ICDE 2020)

Xu

Gholami

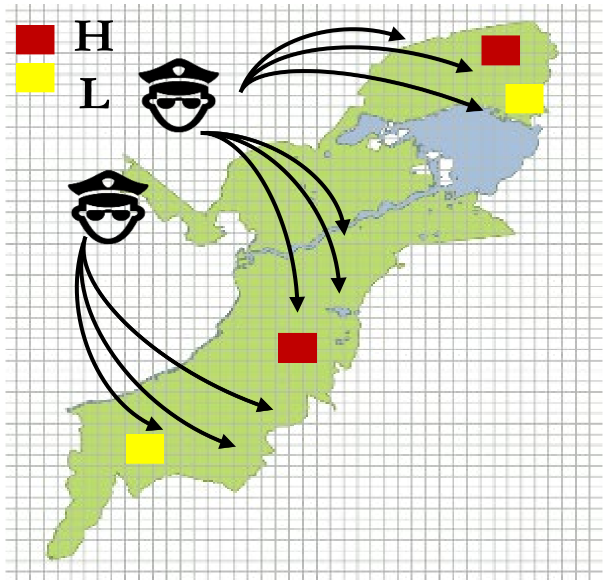

Queen Elizabeth National Park

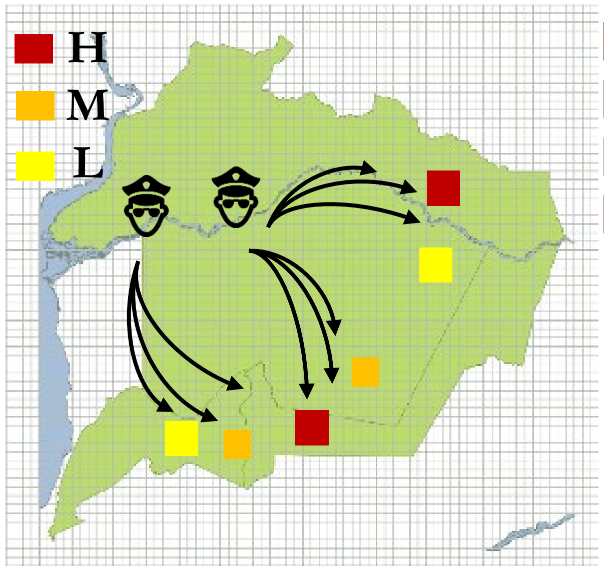

Murchison Falls National Park

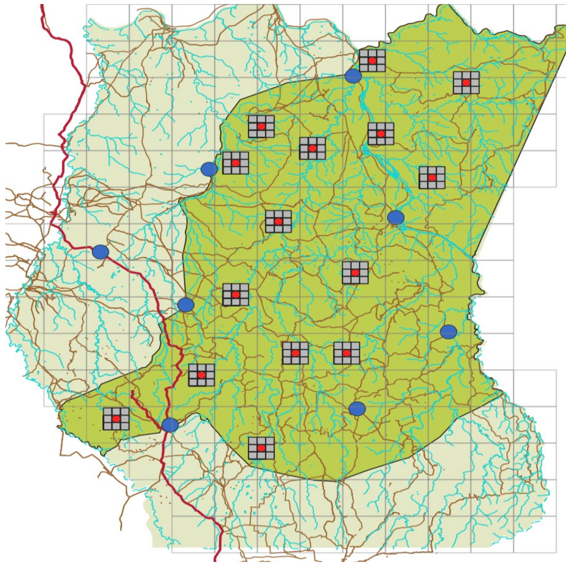

Srepok Wildlife Sanctuary

Snares per patrolled sq. KM

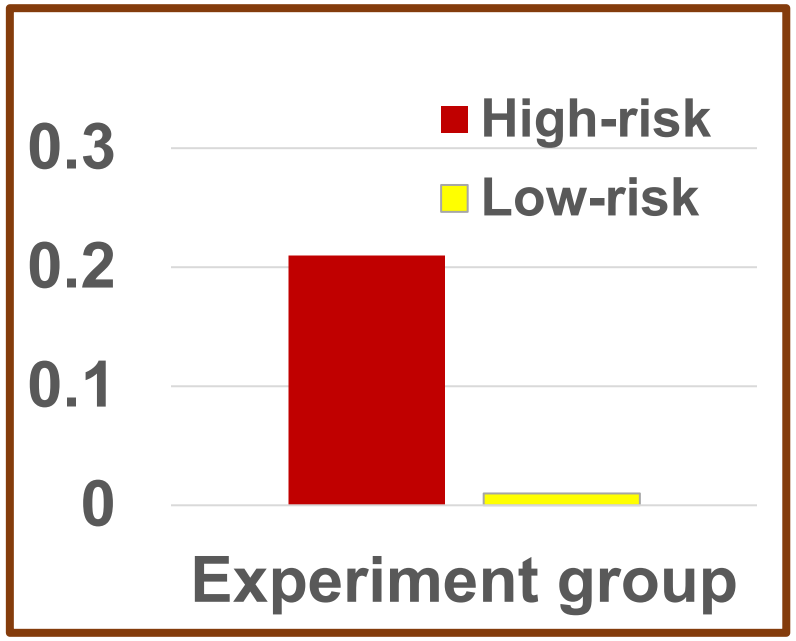

Snares per patrolled sq. KM

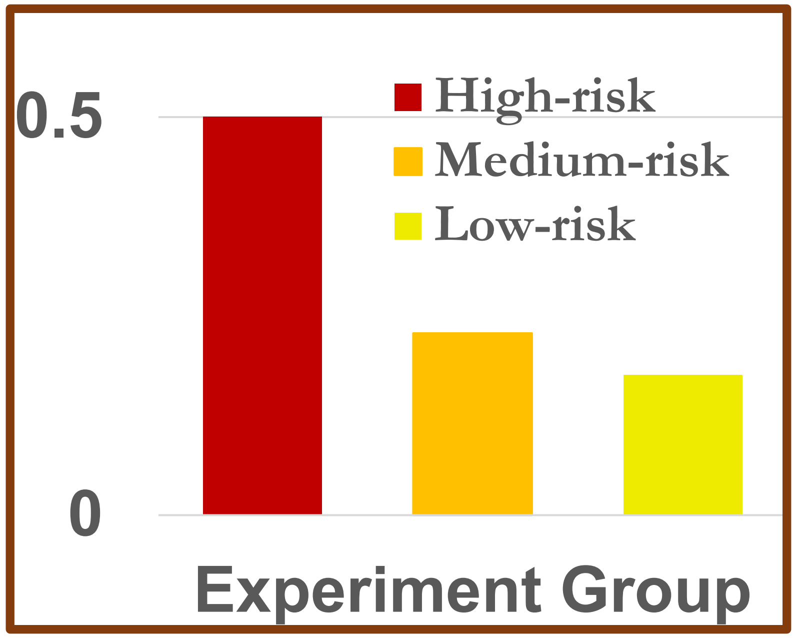

Snares per patrolled sq. KM

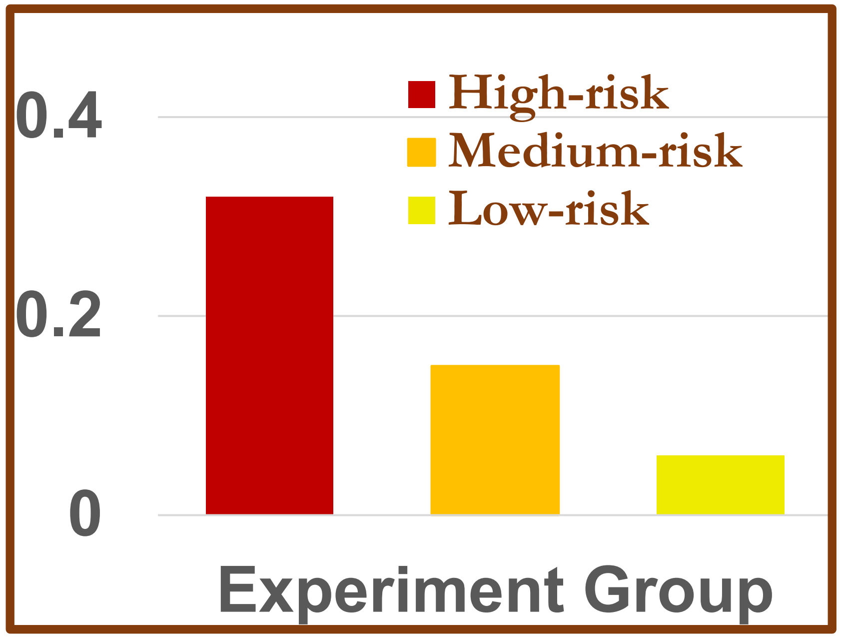

PAWS Real-world Deployment
Cambodia: Srepok Wildlife Sanctuary

(ICDE 2020)

Xu

2019 PAWS: *521 snares/month*

vs

2018: *101 snares/month*

2021 PAWS

1,000 snares found in March

PAWS GOES GLOBAL with SMART platform!!

SMART

WWF

PEACE PARKS FOUNDATION · THE GLOBAL SOLUTION

WCS

Protect Wildlife
800 National Parks
Around the Globe

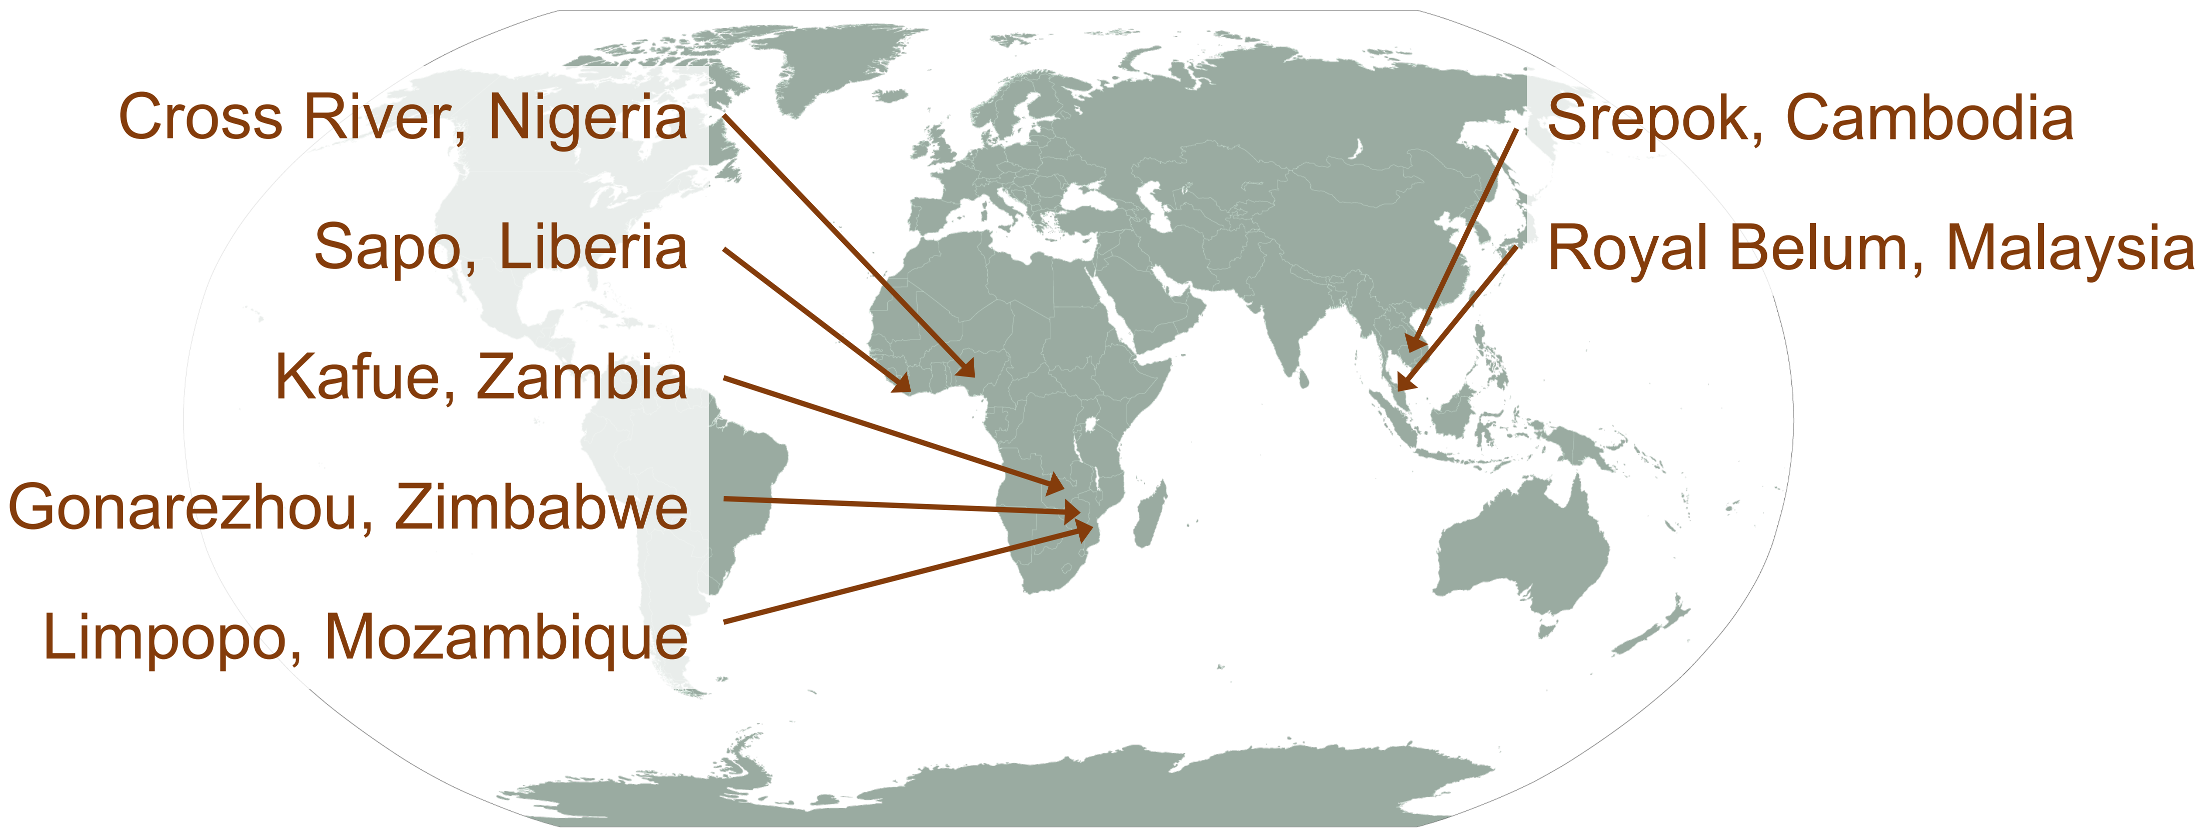

Direction #1: Integrating Real-Time "SPOT" Information

(IAAI 2018)

Bondi

Air Shepherd
The Lindbergh Foundation

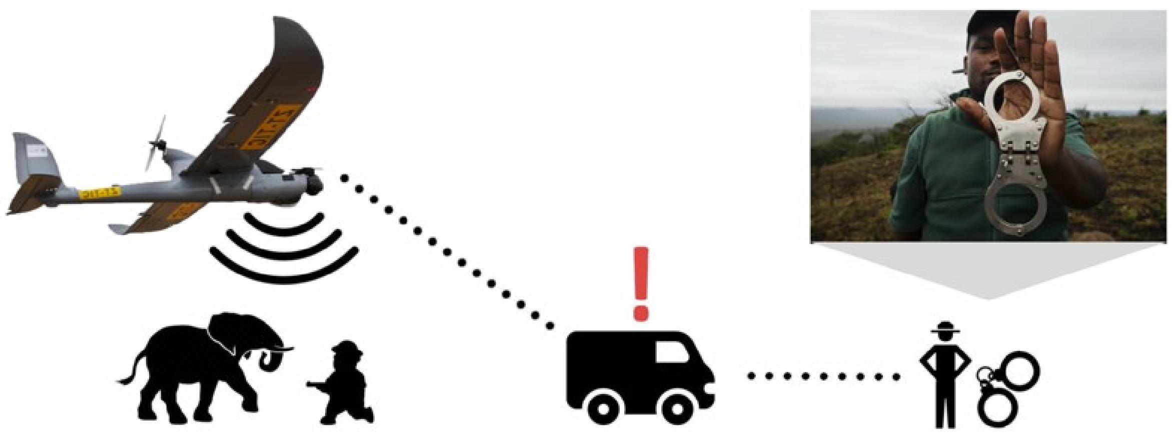

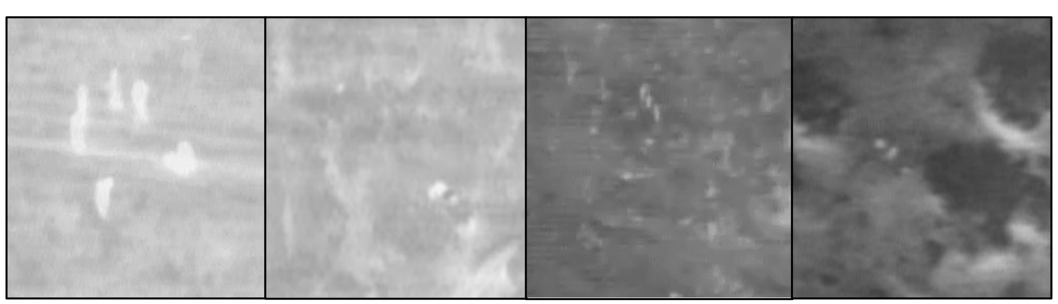

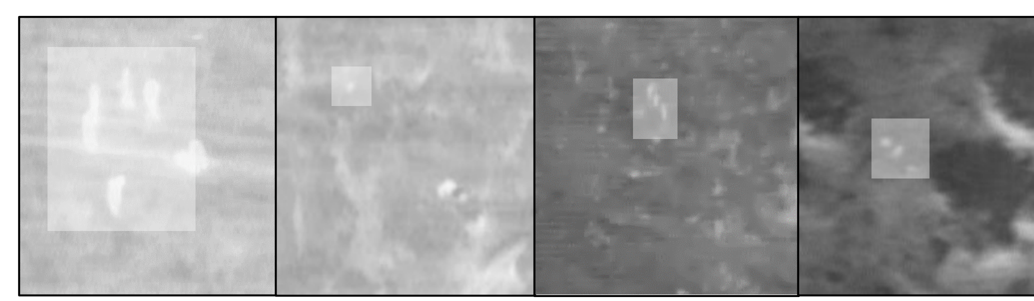

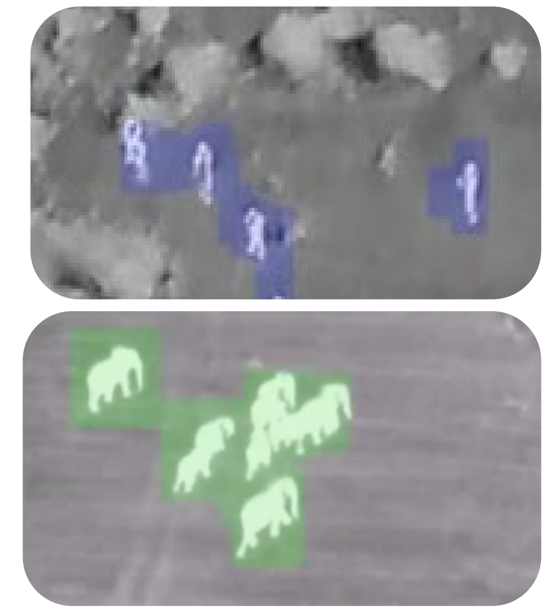

Goal: automatically find poachers

Drone Used to Inform Rangers

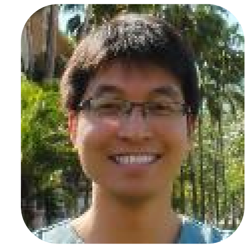
Xu


Bondi

- *Prob(ranger arrives)* = 0.3   [poacher may not be stopped]
- Deceptive signaling to indicate ranger is arriving

Air Shepherd
The Lindbergh Foundation

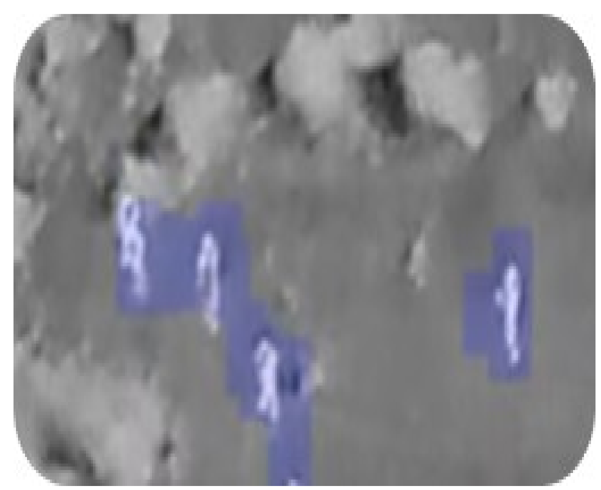

Prob(ranger) = 0.3

Drone Used to Inform Rangers


Xu


Bondi

- *Prob(ranger arrives)* = 0.3 [poacher may not be stopped]
- Deceptive signaling to indicate ranger is arriving

Drone Used to Inform Rangers

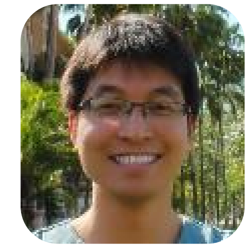
Xu


Bondi

- *Prob(ranger arrives)* = 0.3 [poacher may not be stopped]
- Deceptive signaling to indicate ranger is arriving
- Must be strategic in deceptive signaling

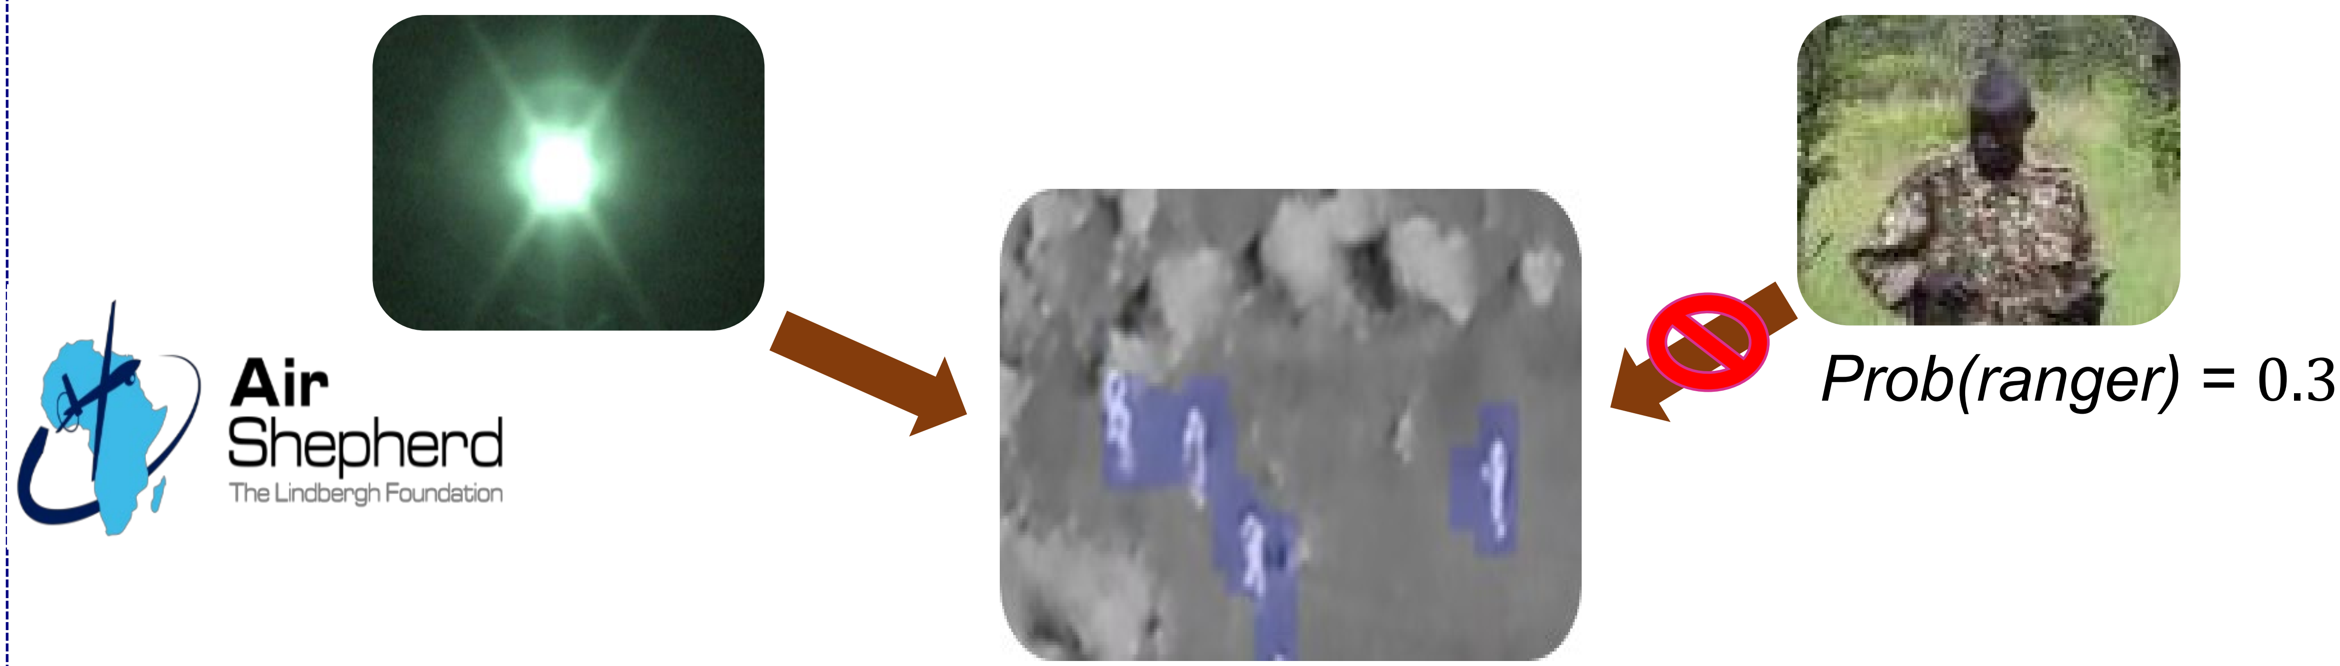

Exploiting Informational Advantage
Defender Knows Pure & Mixed Strategy

(AAAI 2018, AAAI 2020, AAMAS 2021)

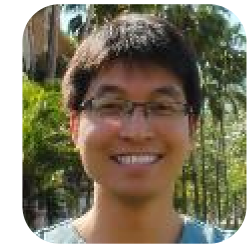

Xu

Bondi

Si-G Model: Stackelberg Security Games with Optimal Deceptive Signaling

- Poacher best interest to “believe signal” even if know 50% defender deception
- Recent work used RL for deception policy generation (AAMAS 2021)

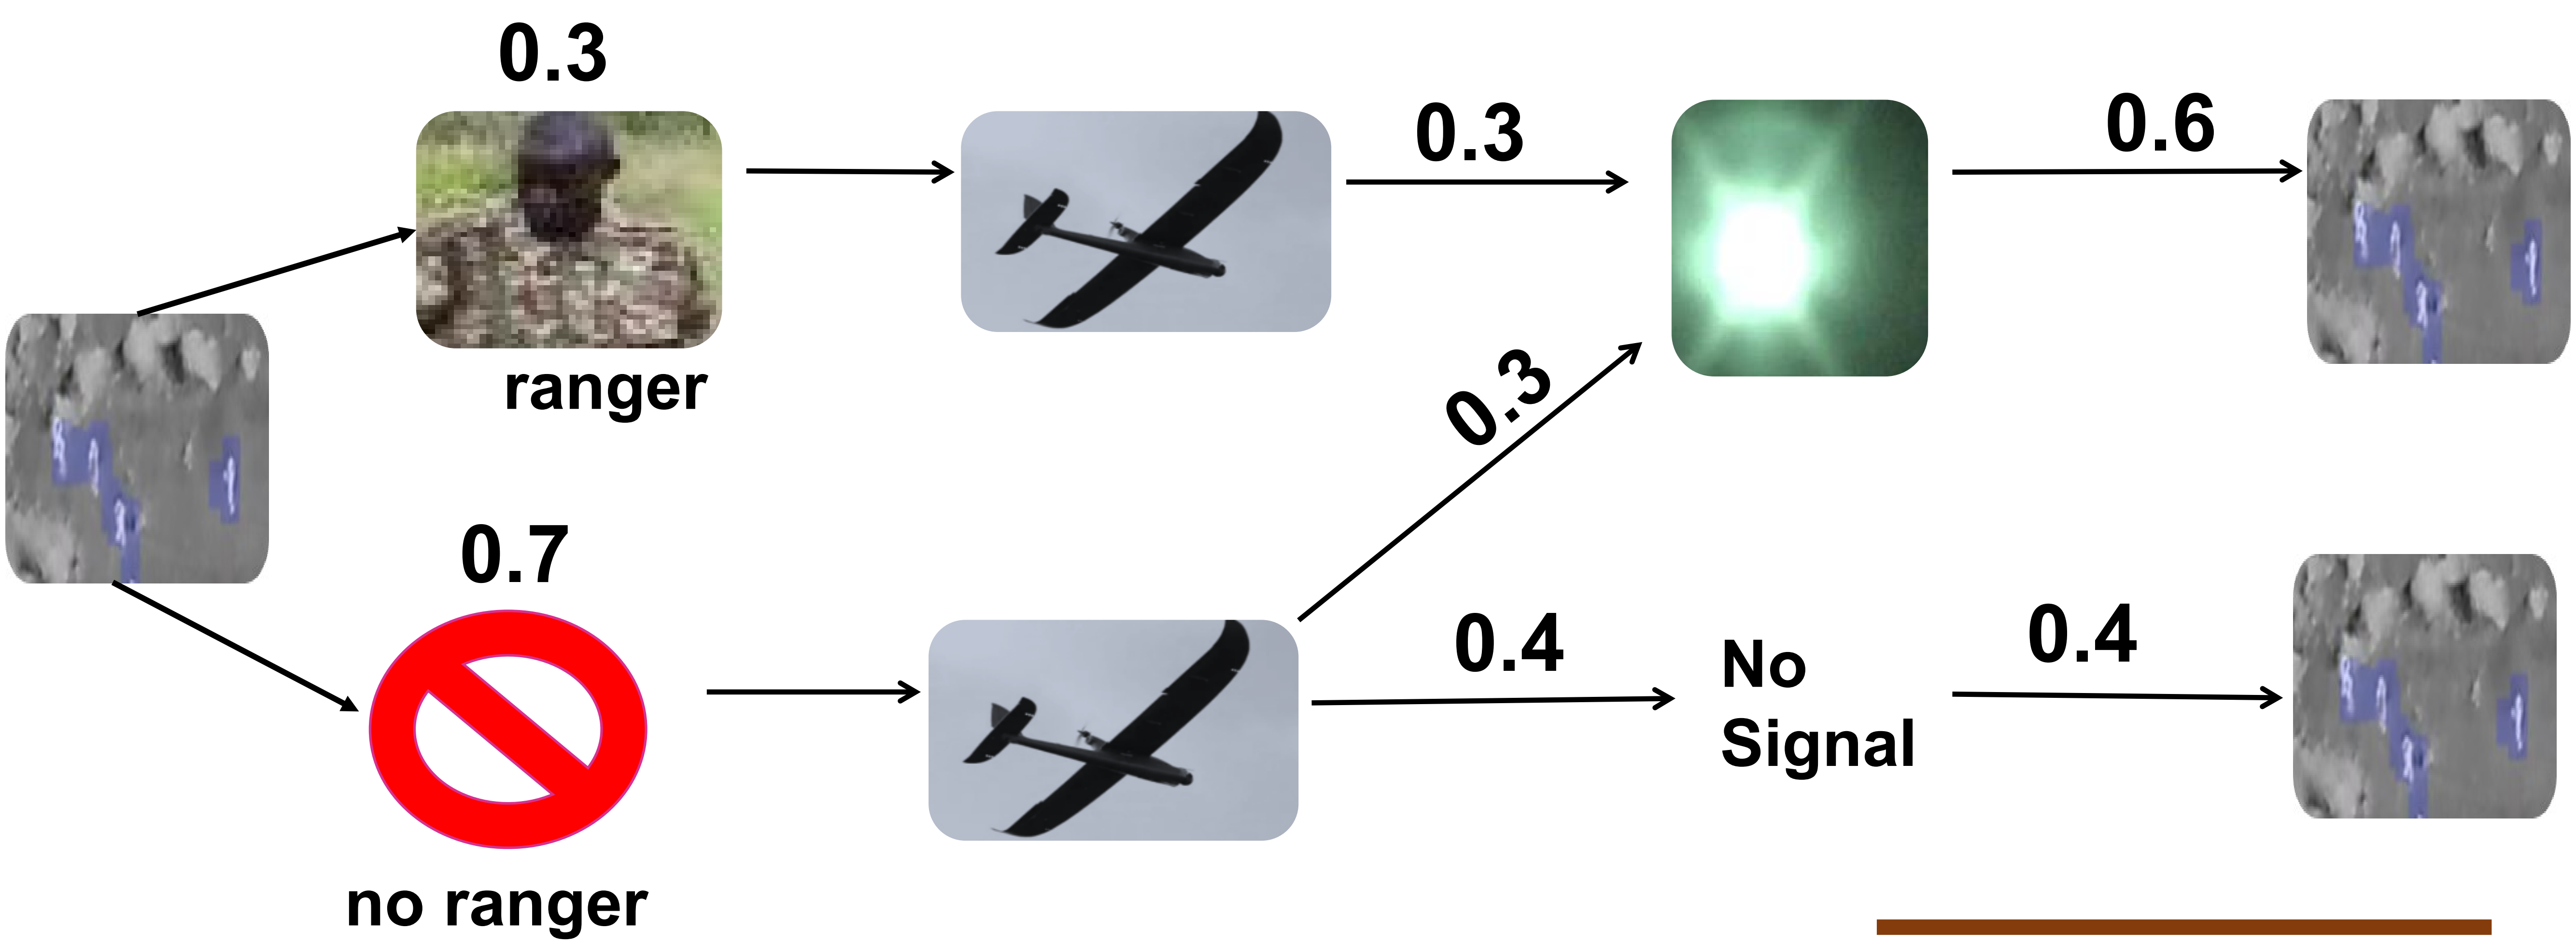

Three new directions
Direction #2: Tailor ML Predictions to Ultimate Objective

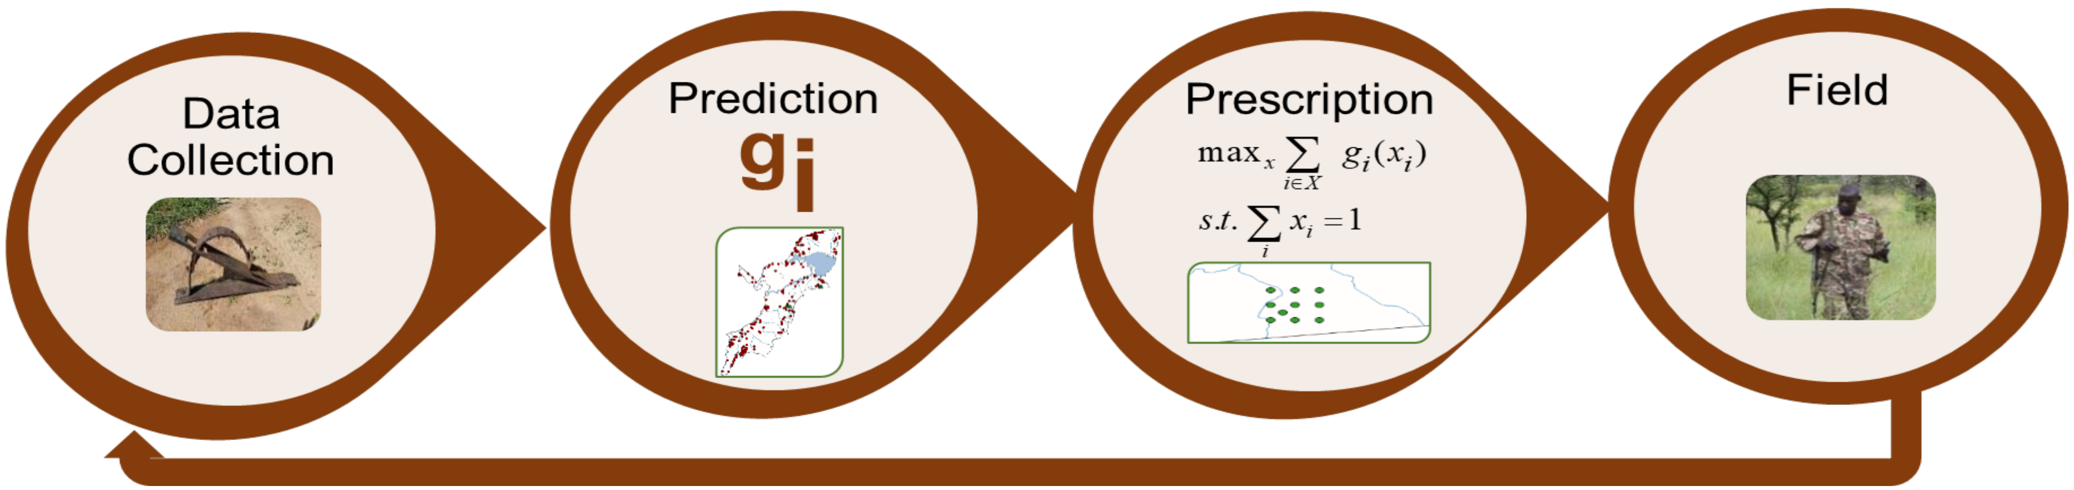

Maximize accuracy in Adversary target values ***estimates***

Prescription: Plan patrol Coverage Game Theory

Minimize $\sum_{\{i \in T\}} q_{\text{empirical}} \log \hat{q}$

Game-Focused Learning:
Modifies Loss Function via Downstream Objective

(AAAI 2019, AAAI 2020, NeurIPS19, NeurIPS20, AAMAS20…)

Perrault

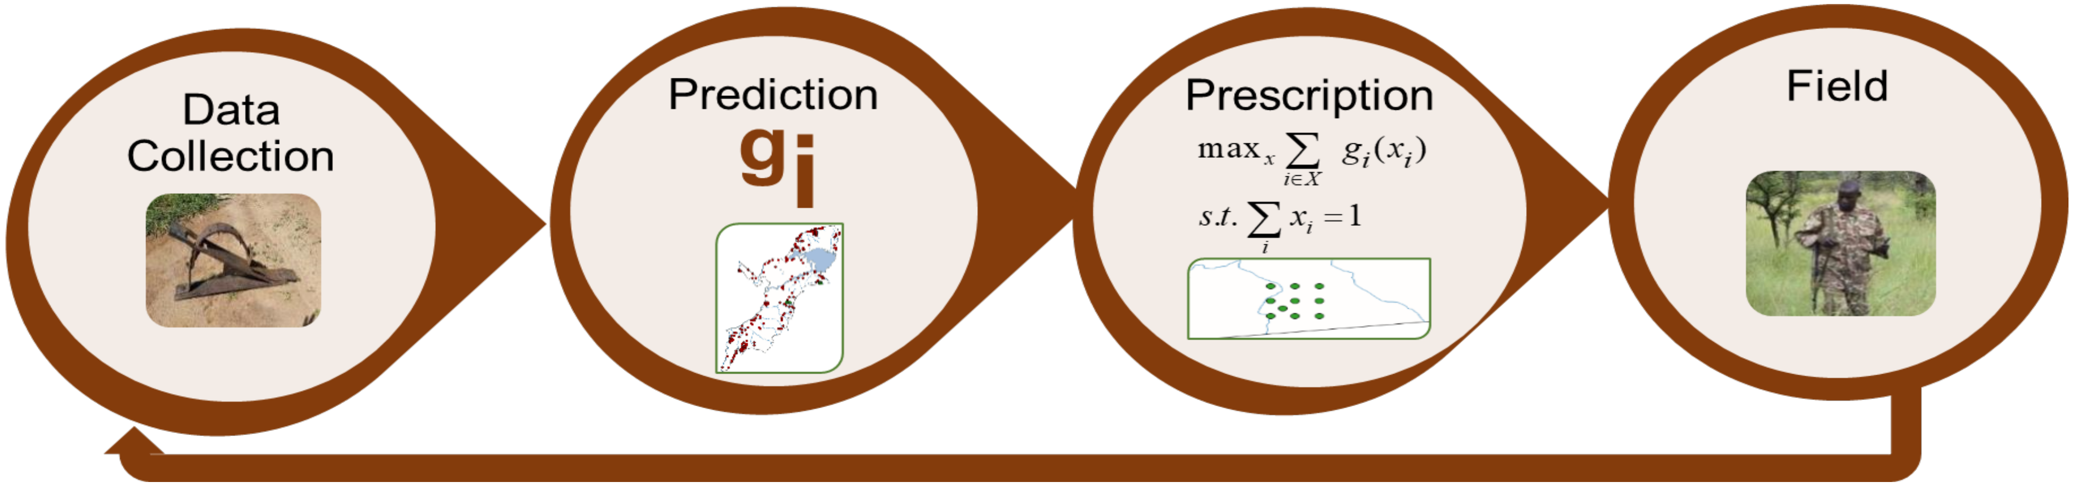

Prediction: Help maximize
Defender's expected utility

Prescription: Plan patrol
Coverage Game Theory

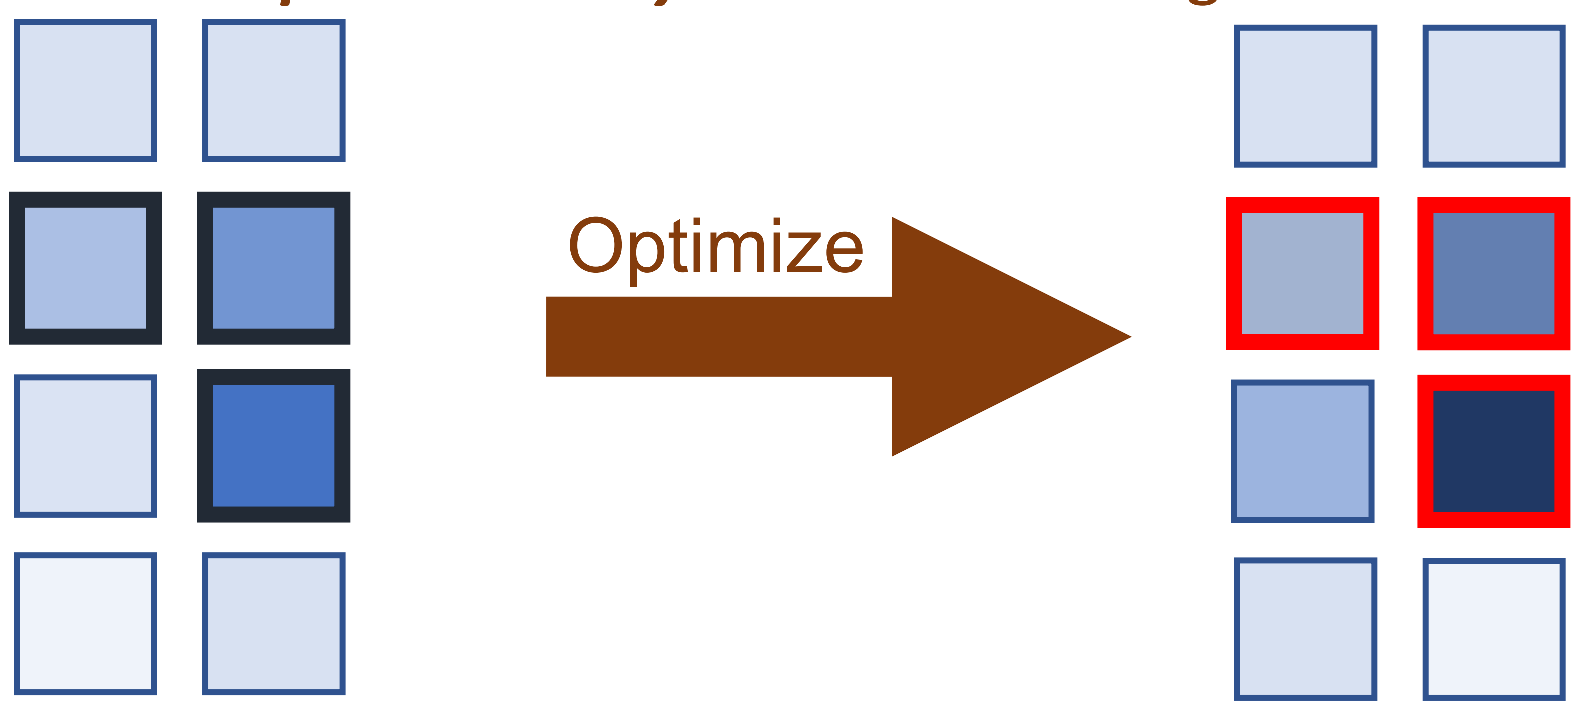

$$\sum (1 - p_i(\hat{q}))\, q_{\text{empirical}}$$

Another View:
Game-Focused Learning: End-to-End Method

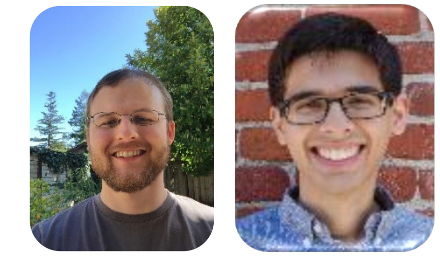

Perrault Wilder

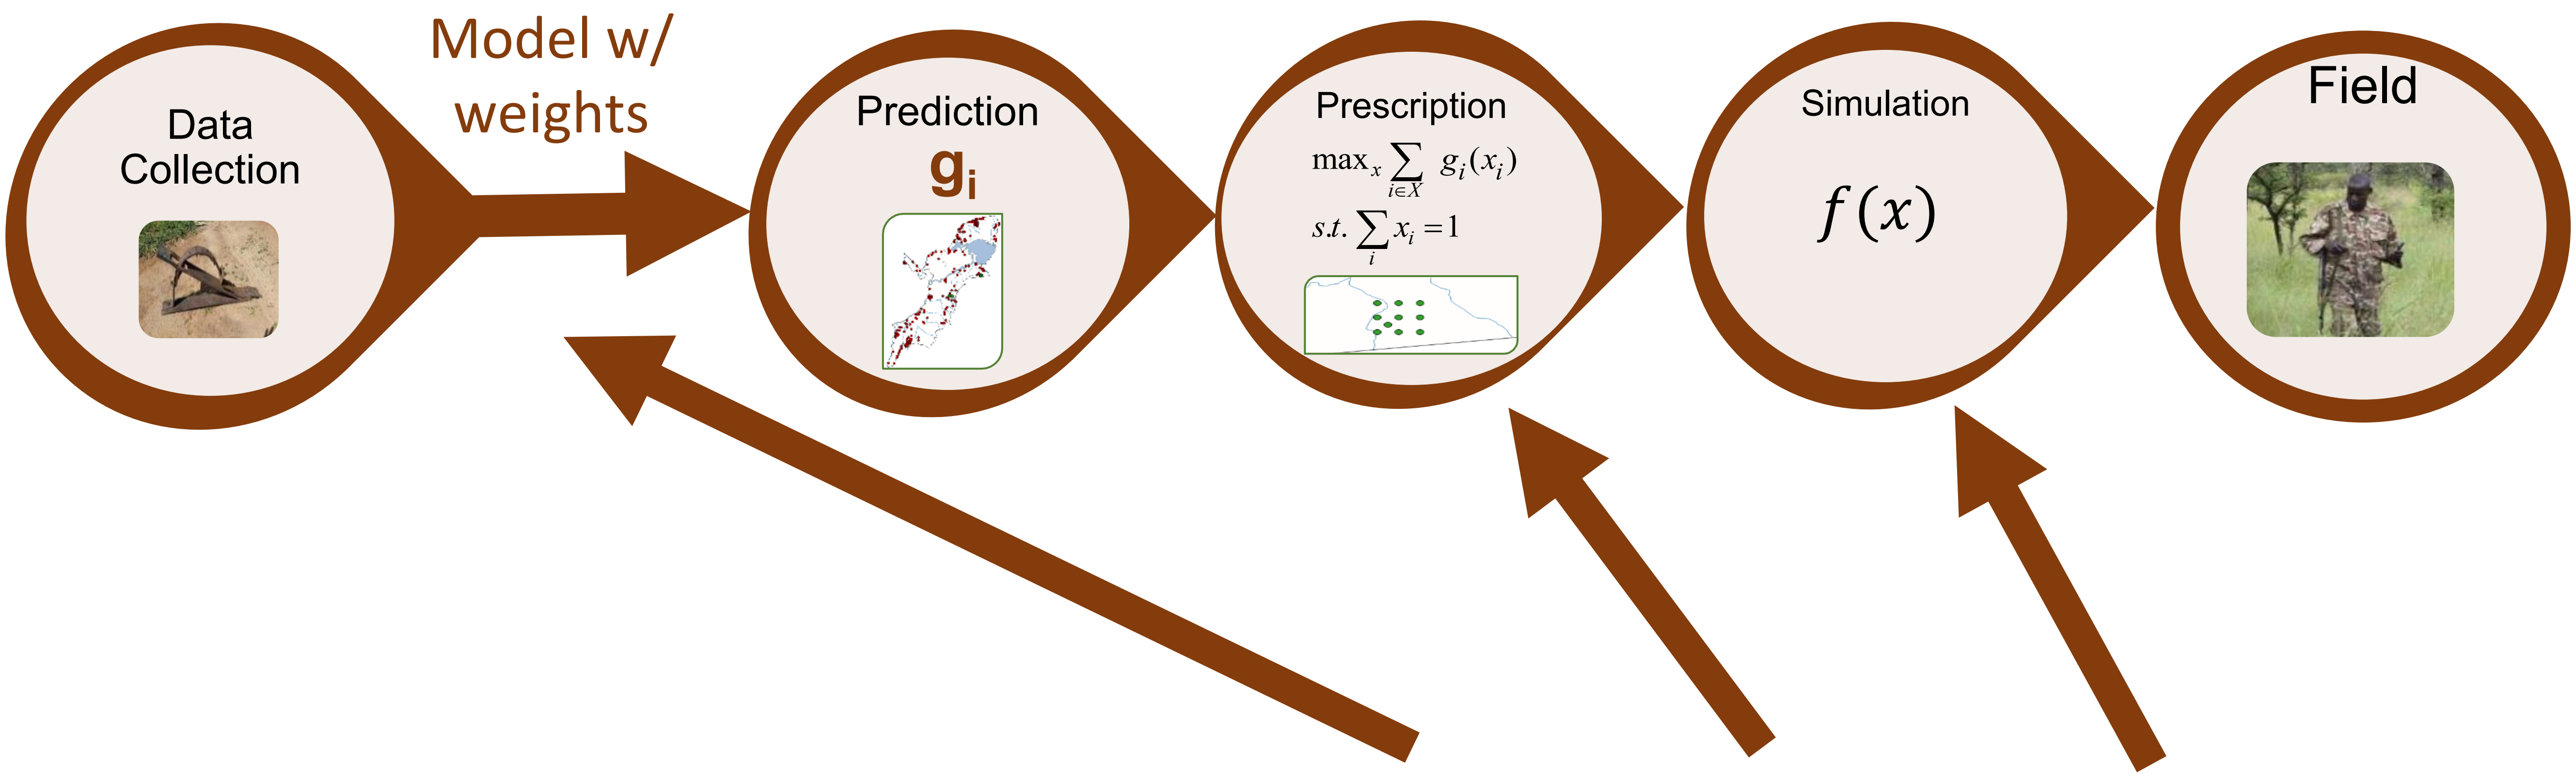

- Game-focused gradient descent:

$$\frac{\partial \text{obj(decision)}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{decision}}{\partial \text{prediction}} \frac{\partial \text{obj(decision)}}{\partial \text{decision}}$$

Game-Focused Learning:
Simulations Murchison Falls National Park

Perrault

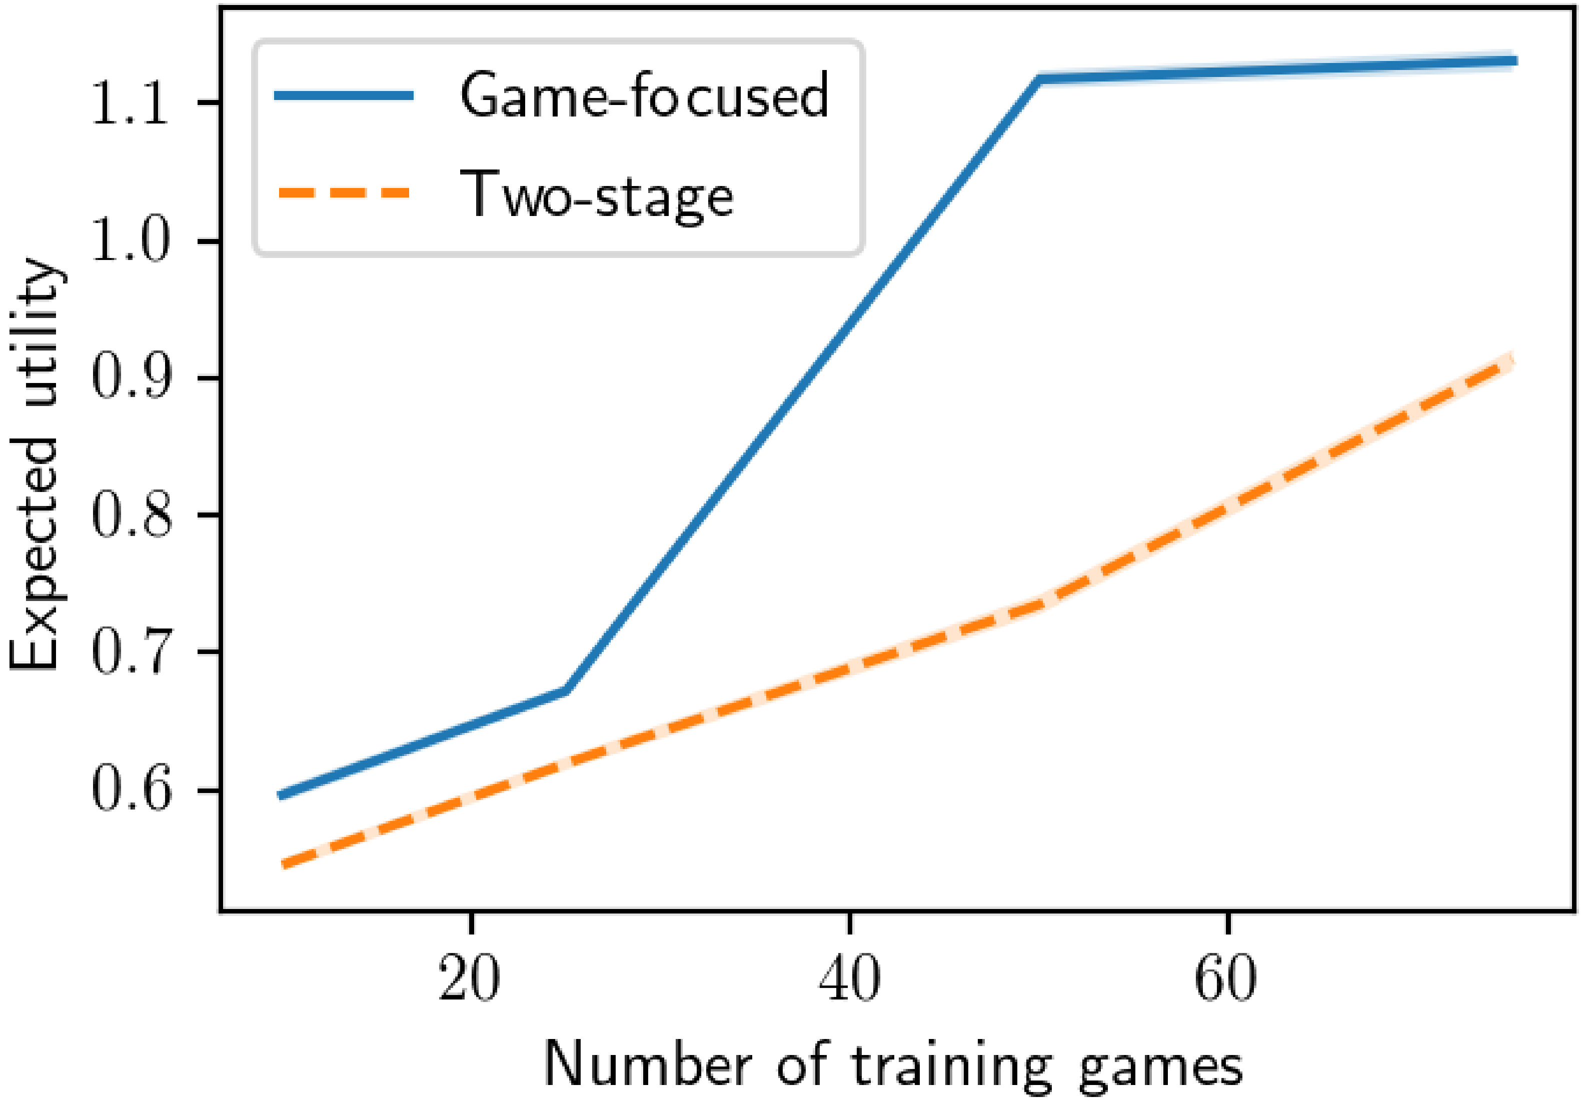

Direction #3: Data Scarce Parks

Xu

exploitation

Data-rich parks: build predictive models to plan patrols

Data-scarce parks: conduct patrols to detect illegal activity and collect data to improve the predictive model

exploration

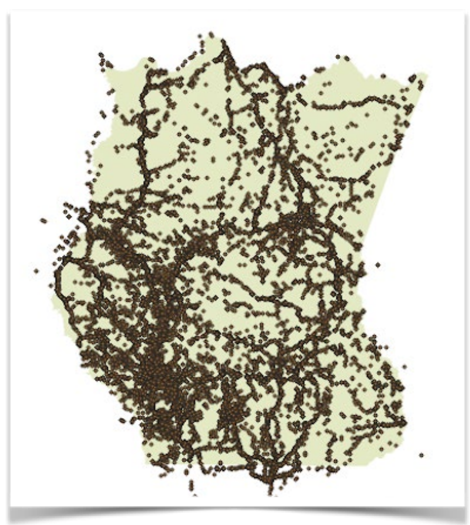

Srepok, Cambodia
43,269 patrol observations
2013 – 2018

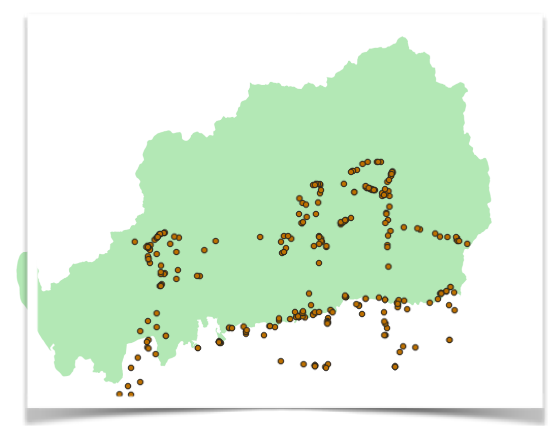

Royal Belum, Malaysia
824 patrol observations
June – August 2018

LIZARD: Multiarmed Bandit
LIpschitZ Arms with Reward Decomposability (AAAI 2021)

Xu

Theorem: Given N targets, Lipschitz constant L, and time horizon T, the regret bound of the LIZARD algorithm is $Reg(T) \leq O\left(L^{\frac{4}{3}} N T^{\frac{2}{3}} (\log T)^{\frac{1}{3}}\right)$:

- *Input:* N Targets with features, T Time, stochastic poacher places snares at targets
- *Output:* Specify patrol effort per target $\leq$ budget B
- Reduce regret wrt OPT, optimal patrol effort, for capturing snares

Lizard exploits decomposability, smoothness, monotonicity

Future: AI for Social Impact (AI4SG or AI4SI)


Achieving social impact & AI innovation go hand in hand

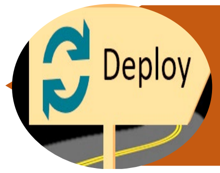

Data to deployment: Not just improving algorithms, new AI4SI evaluation


Important to step out of the lab and into the field


Embrace interdisciplinary research -- social work, conservation


Lack of data is the norm, a feature; part of the project strategy


Empower non-profits to use AI tools; avoid being gatekeepers to AI4SI tech

Key Collaborators on Papers Referenced

(In the order papers referenced)

- Eric Rice (USC)
- Nicole Immorlica (MSR)
- Yair Zick (UMASS, Amherst)
- Balaraman Ravindran (IIT-Madras)
- Amit Sharma (MSR)
- Maia Majumder (Harvard)
- Michael Mina (Harvard)
- Daniel Larremore (Colorado)
- Andy Plumptre (Cambridge)
- Rohit Singh (WWF)
- Phebe Vayanos (USC)
- Bistra Dilkina (USC)

Collaborate to realize AI's tremendous potential to
Improving society & fighting social injustice

@MilindTambe_AI

Job Alert

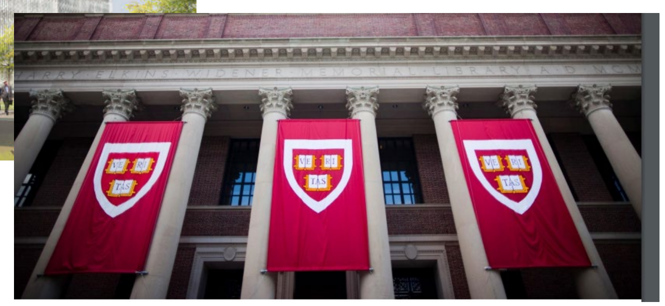

CRCS Center for Research on Computation and Society

at Harvard School of Engineering and Applied Sciences

Postdoc positions: Dec 1 deadline

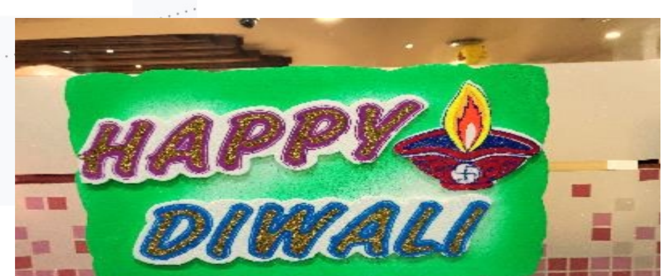

Google Research

Research scientist, visiting scientists, postdocs, predocs: See website

- The END

New Directions in Restless Bandits

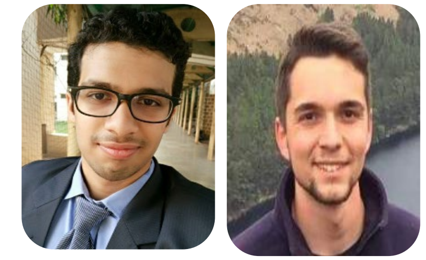

Mate Killian Biswas

Fast algorithms for extending to:

- *Multiple action types*

  (AAMAS 2021a)

- *Risk aware restless bandits*

  (AAMAS 2021b)

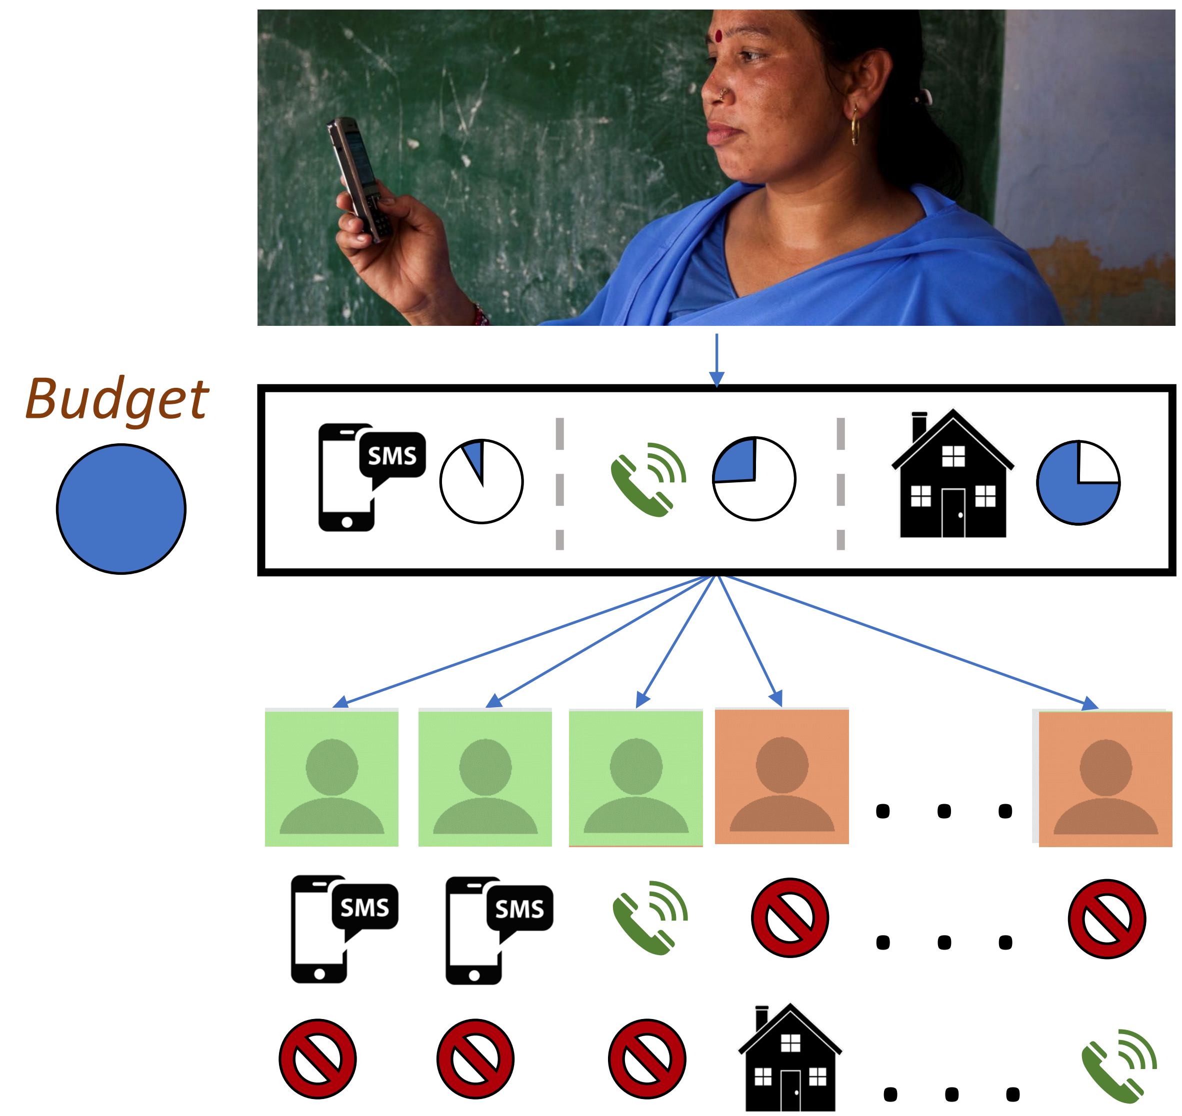

Online learning

- *Learning policies via Index Q-Learning*

  (AAMAS 2021c)

New Directions in Restless Bandits

Mate Killian Biswas

Fast algorithms for extending to:

- ➢ *Multiple action types*

  (AAMAS 2021a)

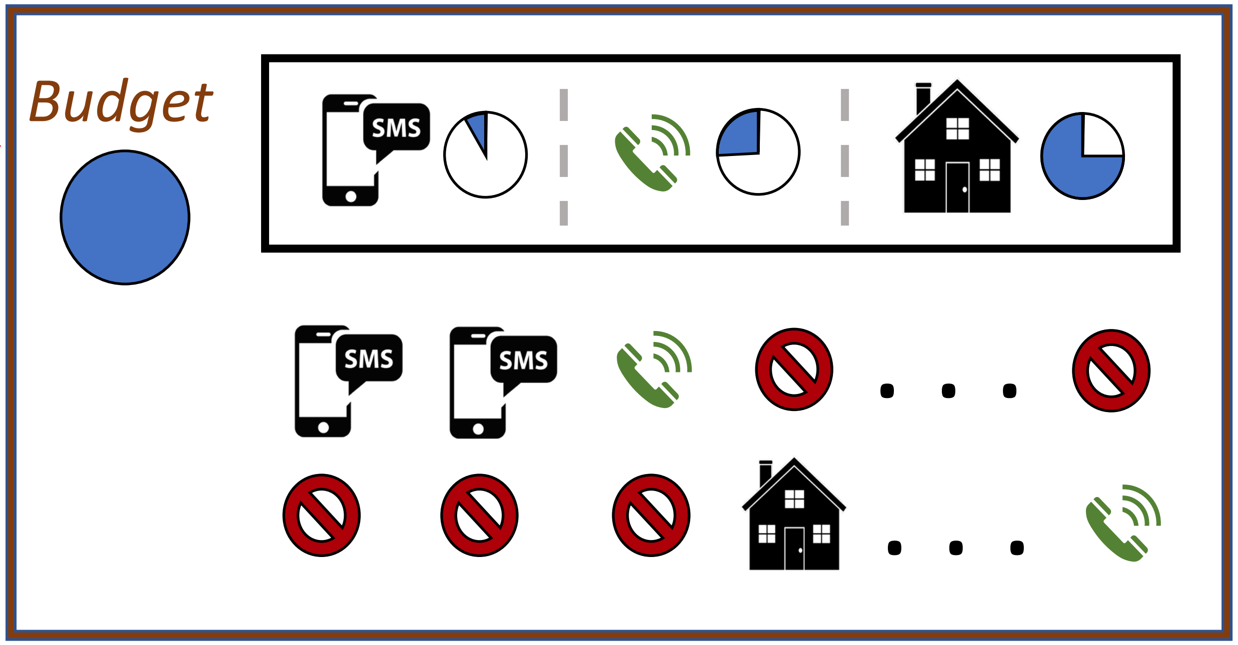

- ➢ *Risk aware restless bandits*

  (AAMAS 2021b)

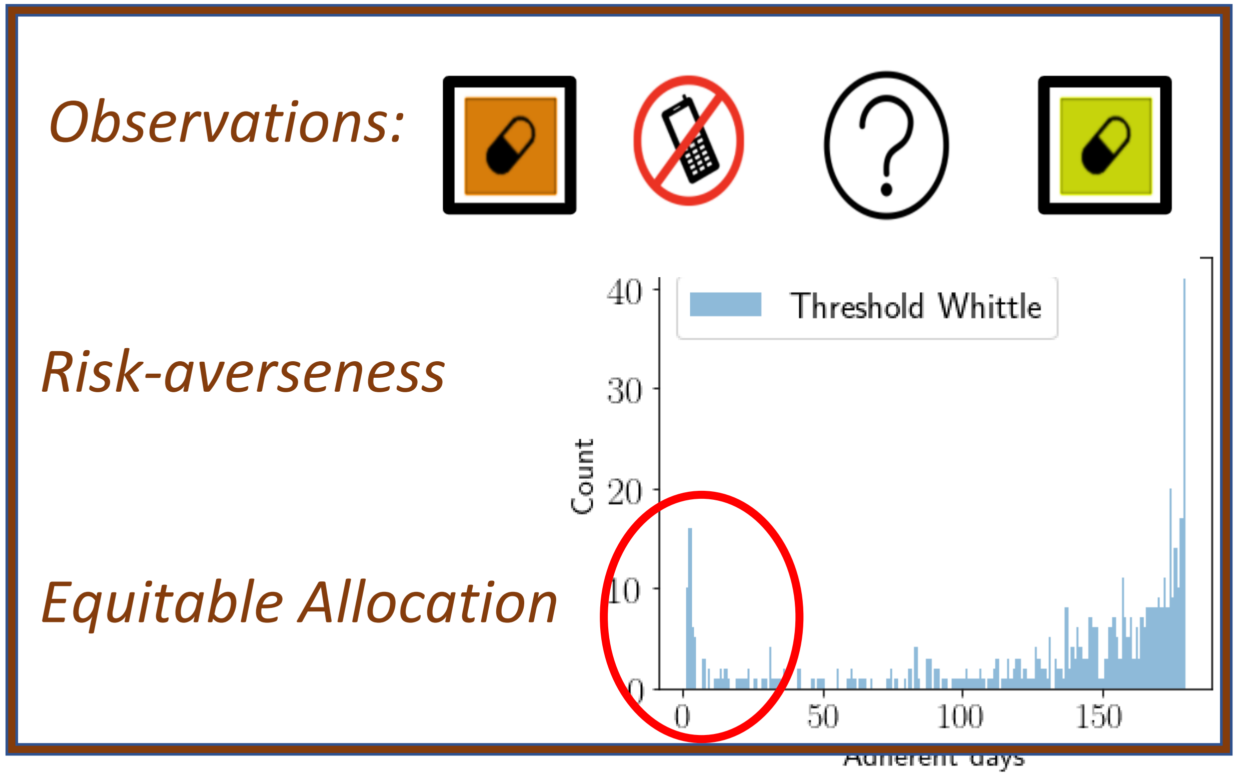

Online learning

- ➢ *Learning policies via Index Q-Learning*

  (AAMAS 2021c)

Another View:

Previous Two-Stage Method: Gradient Descent

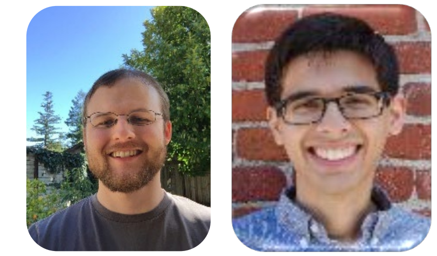

Perrault Wilder

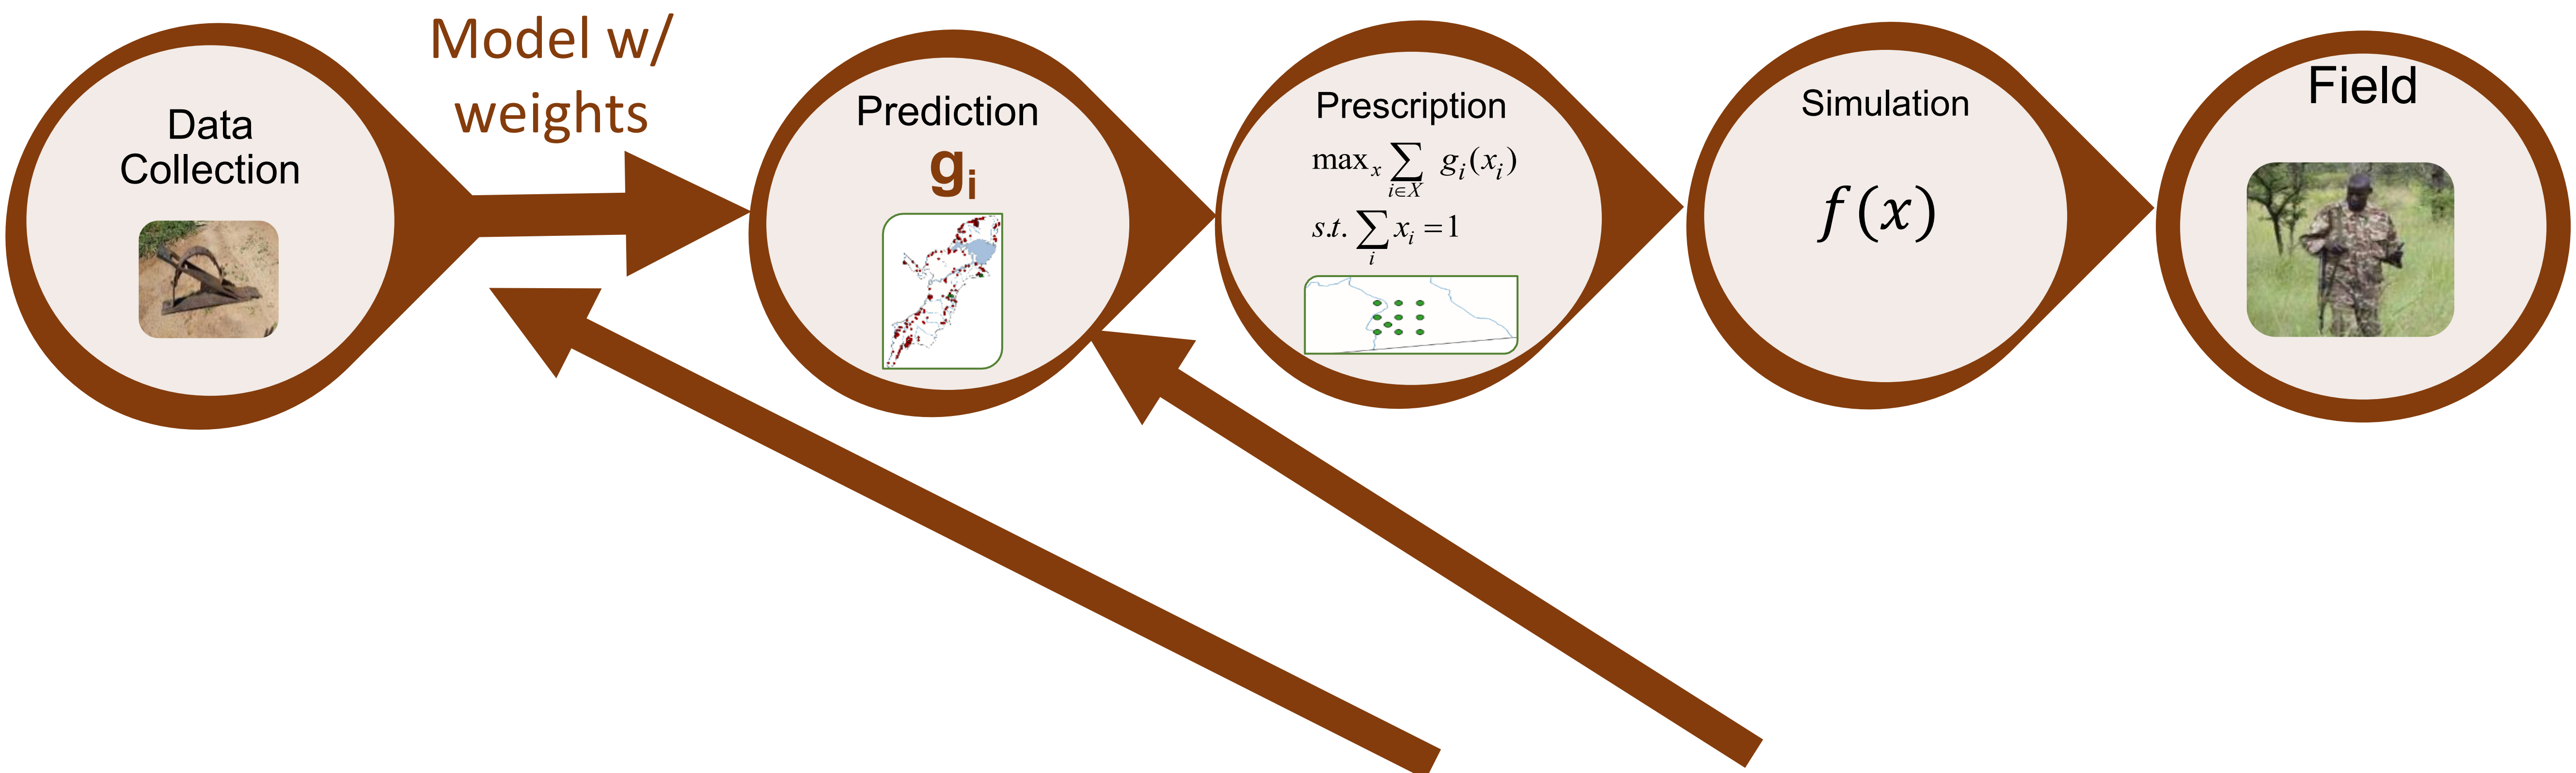

➢ Max accuracy gradient descent:

$$\frac{\partial \text{accuracy}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{accuracy}}{\partial \text{prediction}}$$

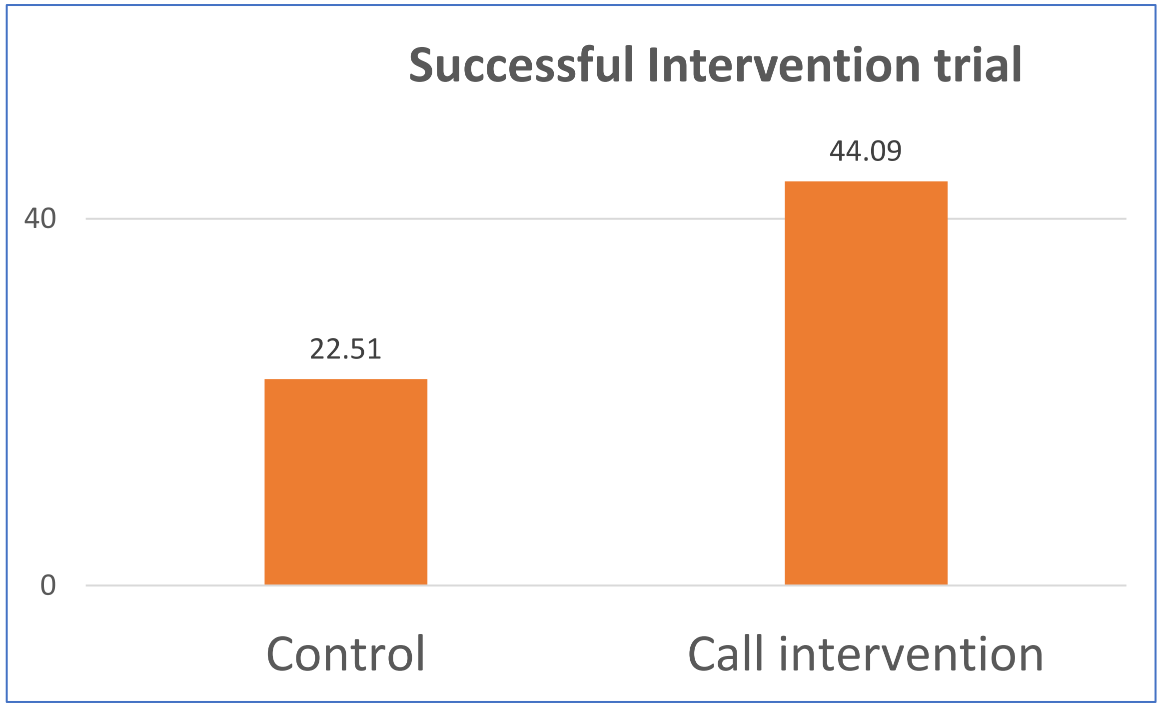
Successful Intervention trial
44.09
40
22.51
0
Control
Call intervention