AI for Public Health & Conservation:

Learning & Planning in the Data to Deployment Pipeline

MILIND TAMBE

Director Center for Research on Computation and Society, Harvard University

&

Director "AI for Social Good", Google Research India

@MilindTambe_AI

AI and Multiagent Systems Research for Social Impact

Public Health

Conservation

Public Safety and Security

Viewing Social Problems as Multiagent Systems

Key research challenge across problem areas:

Optimize Our Limited Intervention Resources
when
Interacting with Other Agents

Optimizing Limited Intervention Resources

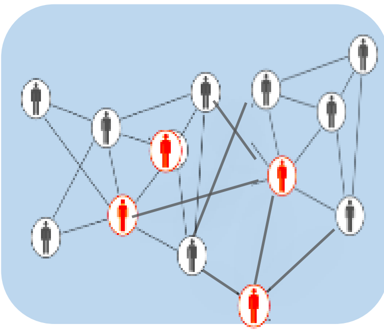

Social Networks & Bandits

Public Health

Green security games

Conservation

Stackelberg security games

Public Safety & Security

Google Research Bangalore
AI for Social Good

armman
Delivering India's Future

KhushiBaby

Swasti
THE HEALTH CATALYST

Public Health

atree
ASHOKA TRUST FOR RESEARCH IN ECOLOGY & THE ENVIRONMENT

Wildlife Conservation Trust

Conservation

Three Common Themes

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

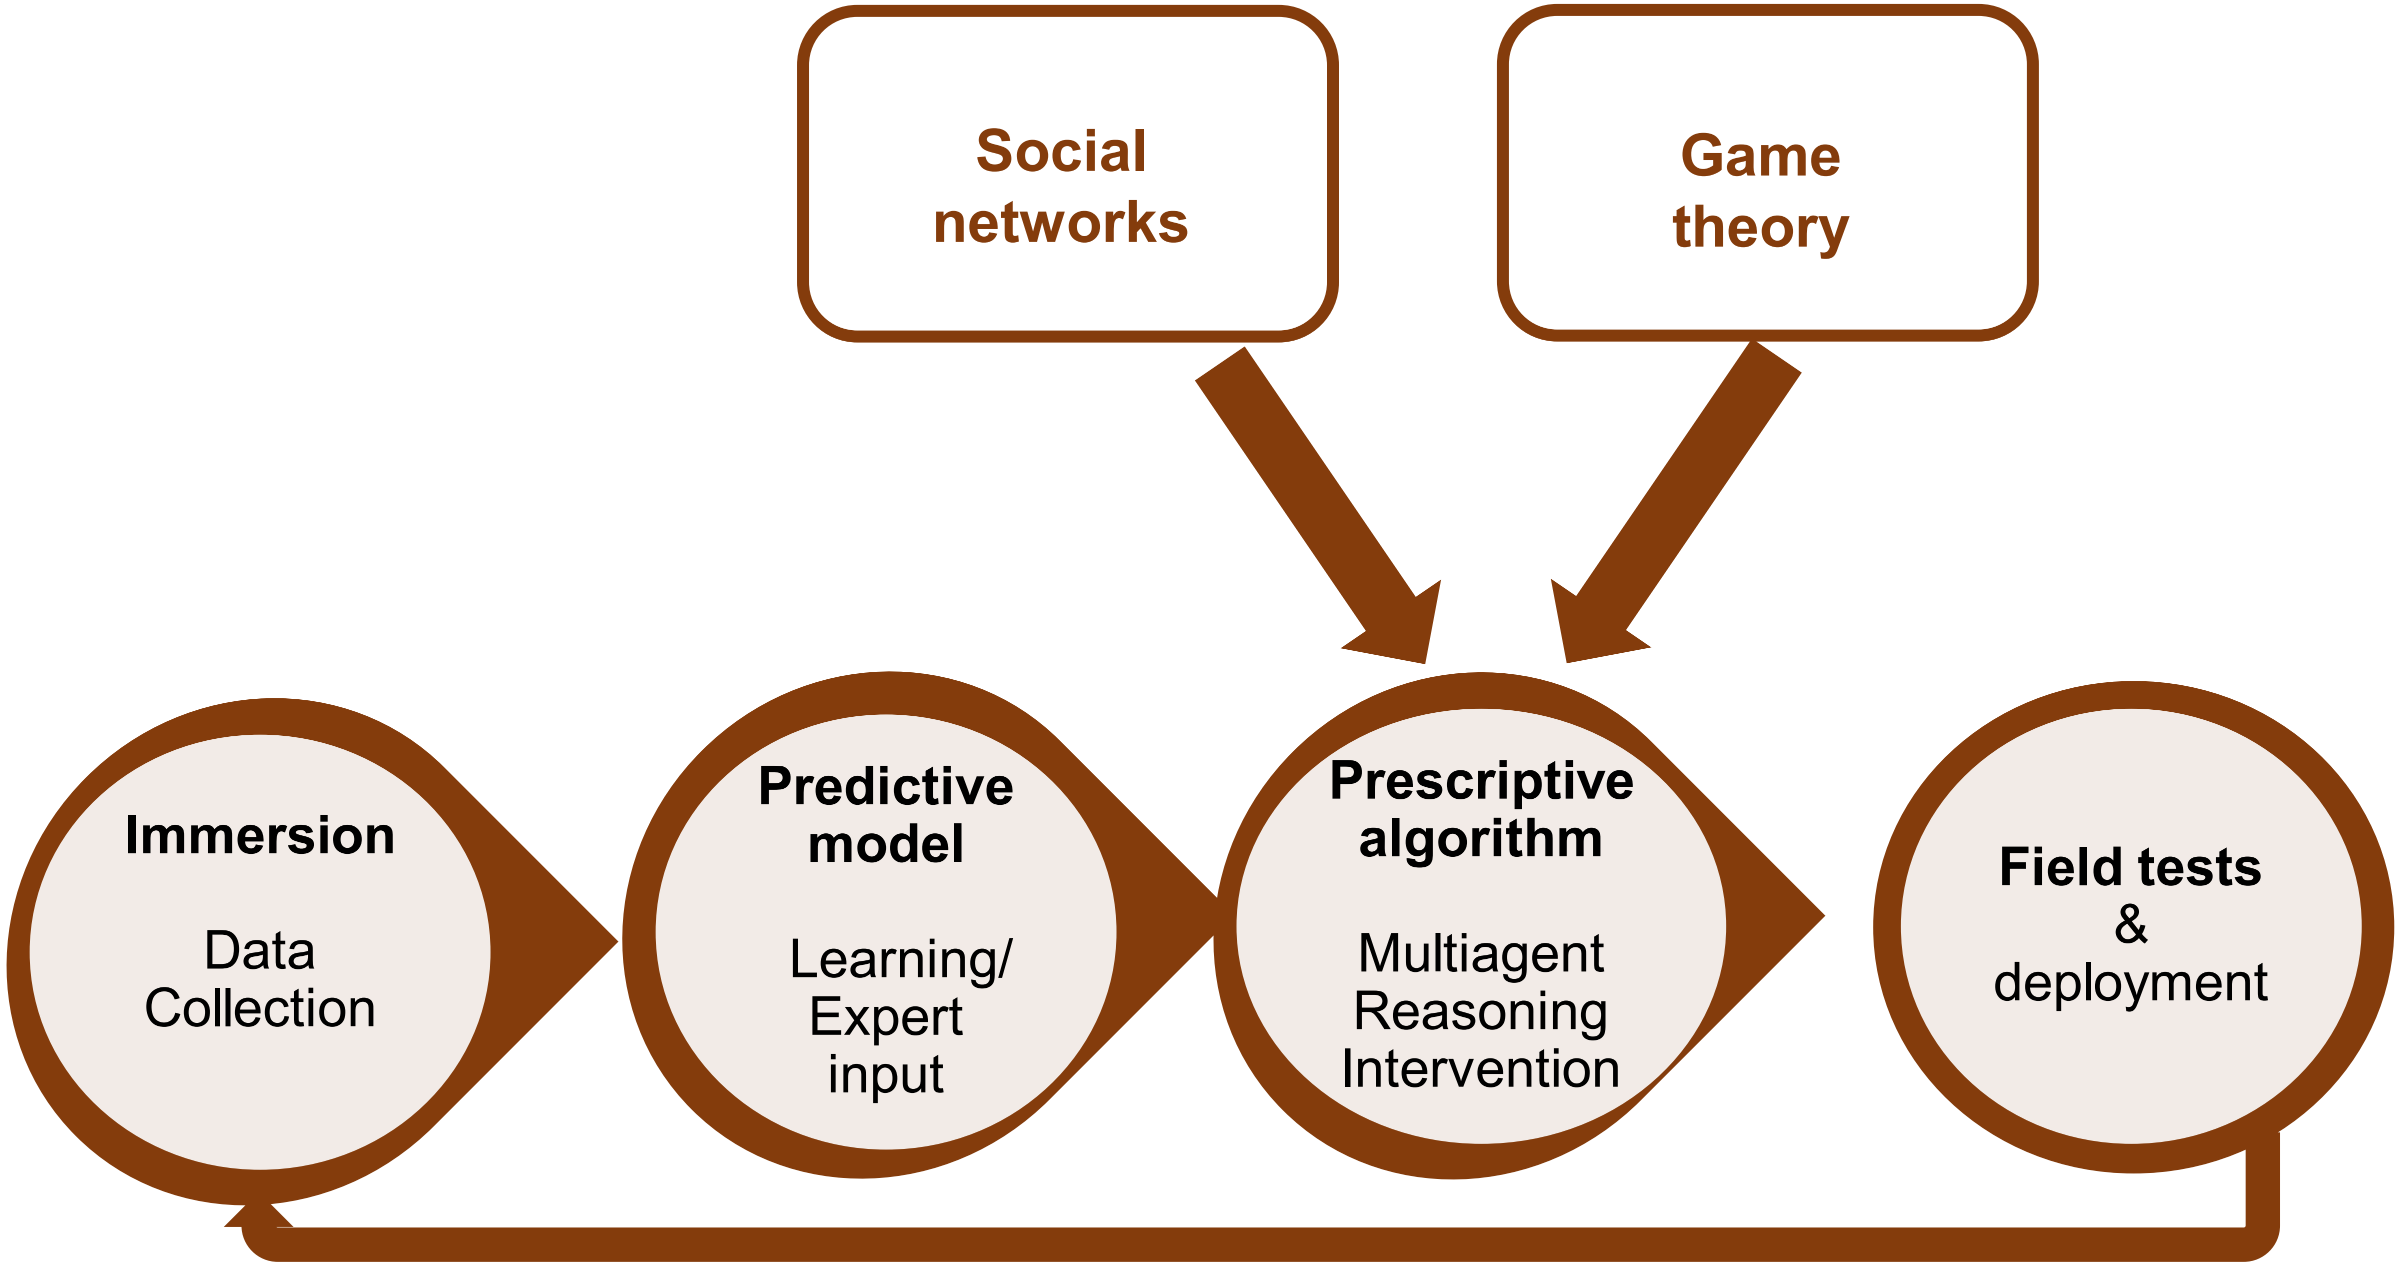

Three Common Themes
Multiagent systems, Data-to-deployment pipeline, Interdisciplinary partnerships

Field test & deployment: Social impact is a key objective

Lack of data is a norm: Must be part of project strategy

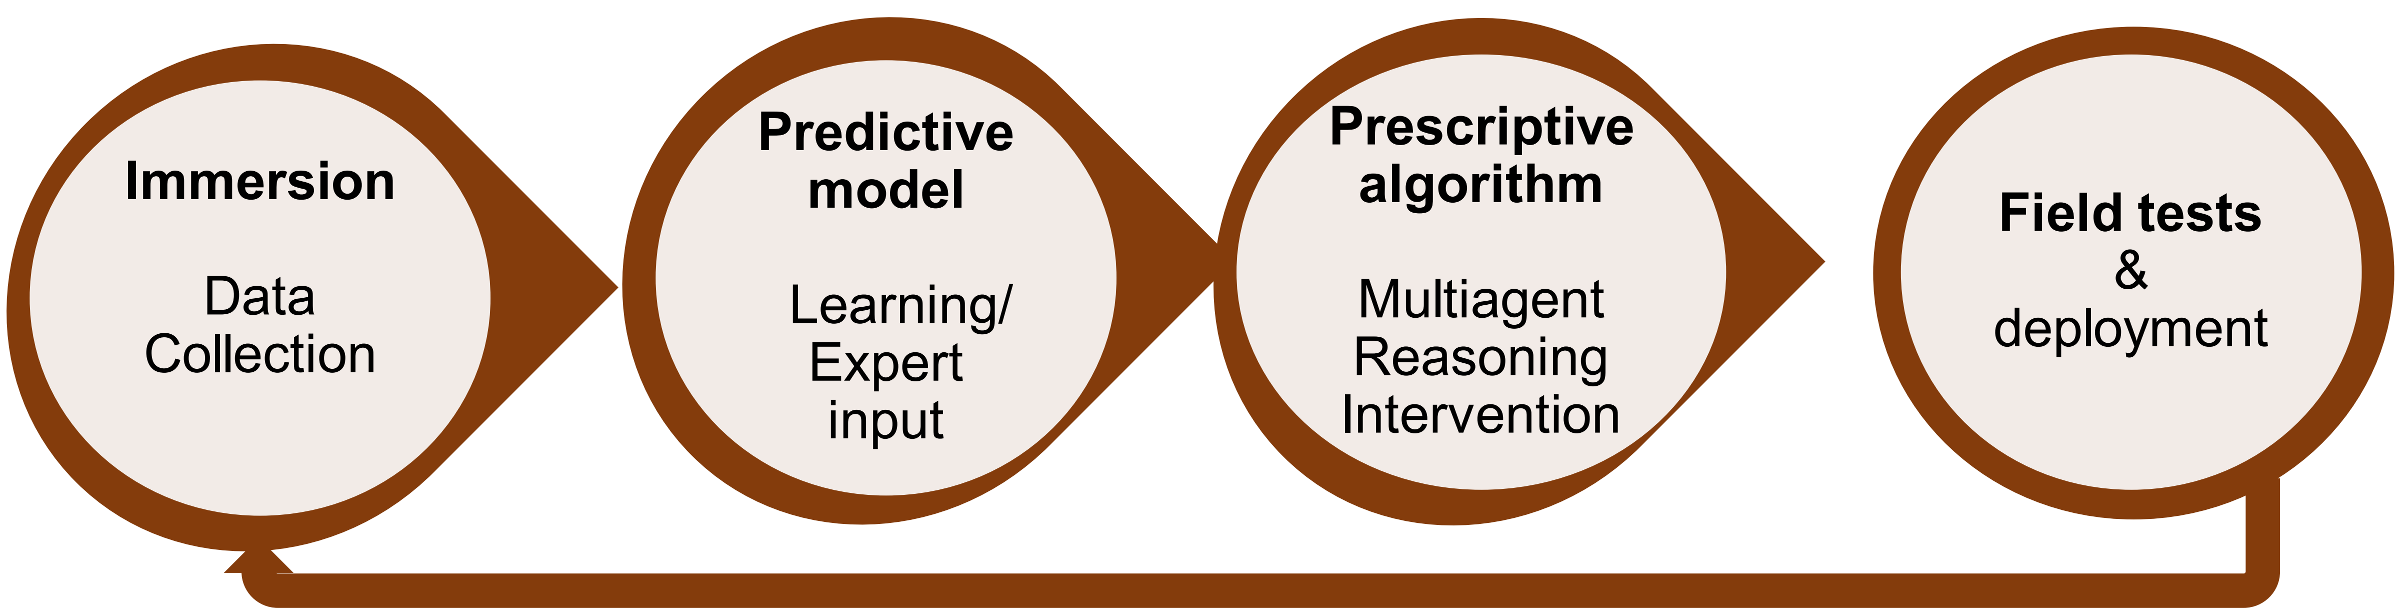

Three Common Themes

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

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

Wildlife Conservation Trust

SMART

Air Shepherd The Lindbergh Foundation

Swasti THE HEALTH CATALYST

everwell

KhushiBaby

PRATHAM BOOKS

United States Coast Guard U.S. Department of Homeland Security

Transportation Security Administration

Three Key Takeaways

EC research on multiagent system impactful for public health/conservation projects

Field deployments bring up new research challenges for EC community

Wealth of new multiagent research challenges via partnerships with NGOs

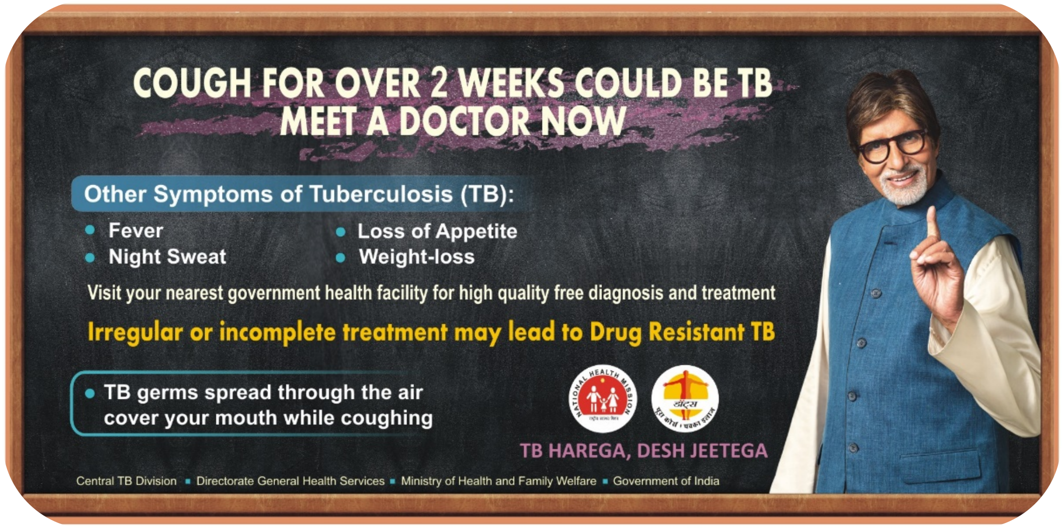

Outline

Public Health

- ➢ *Information dissemination & behavior change: 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

Preventing HIV in homeless youth: Rates of HIV 10 times housed population

- ***Shelters**: Limited number of peer leaders to spread HIV information in social networks*
- *"Real" face-to-face interactions; not Facebook etc*
- *Other applications: HIV prevention (SWASTI), Tuberculosis awareness…*

Influence Maximization in Social Networks

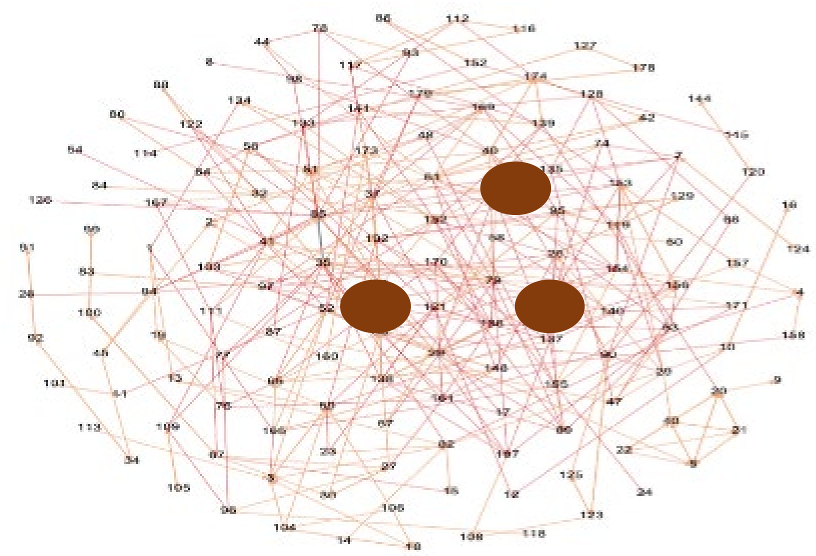

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
  - Assume: Independent cascade model of information spread

- 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 uncover network**

Challenge 1: Uncertainty in Real-world Physical Social Networks

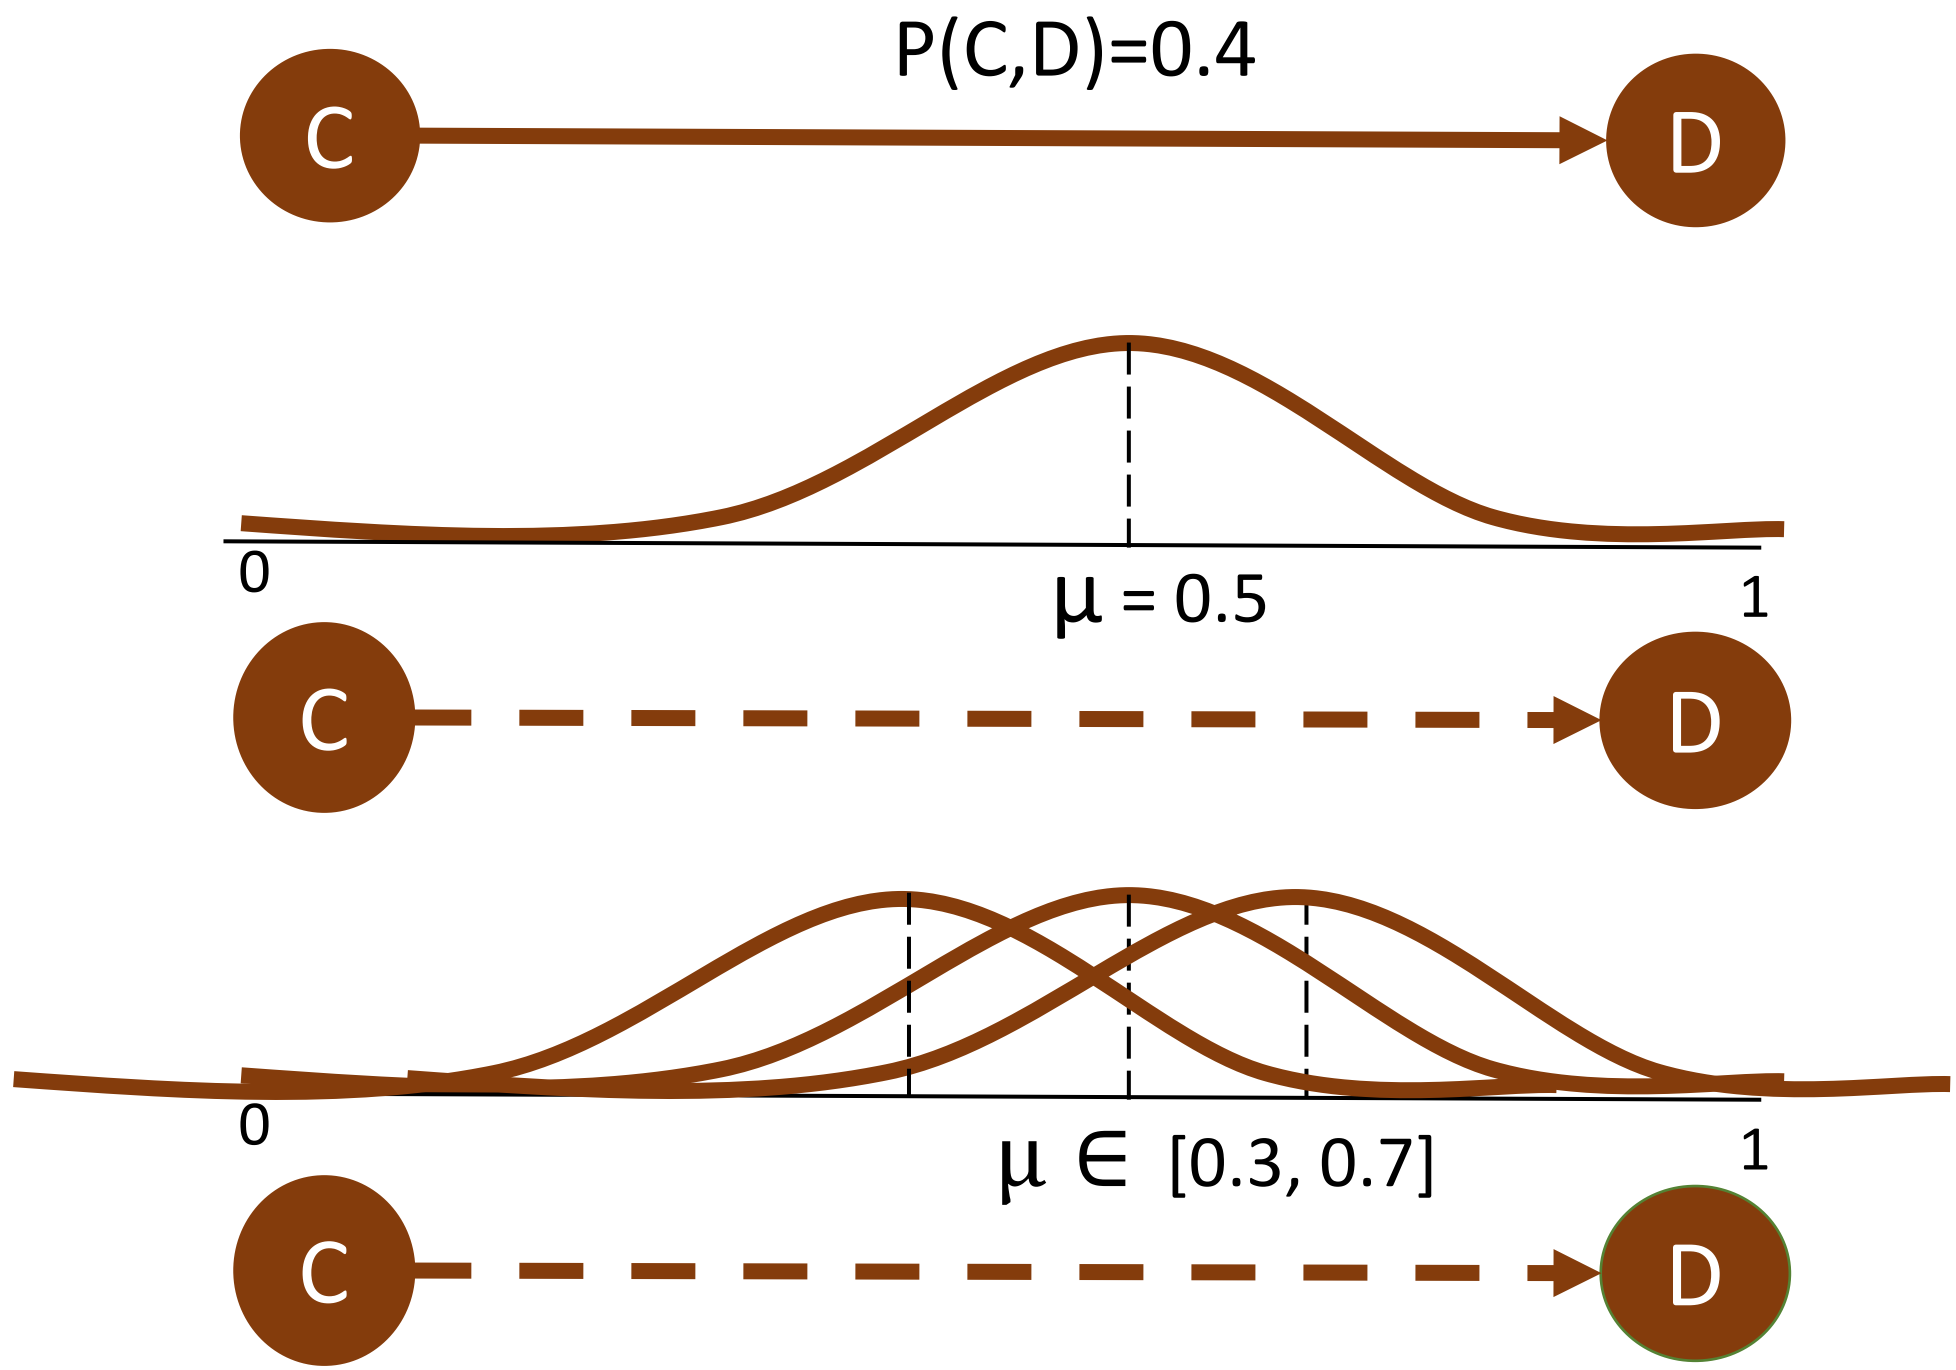

Robust Influence Maximization
(AAMAS 2017)

Wilder

- 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: Multi-step Policy


Yadav


Wilder

K = 4
1st time step

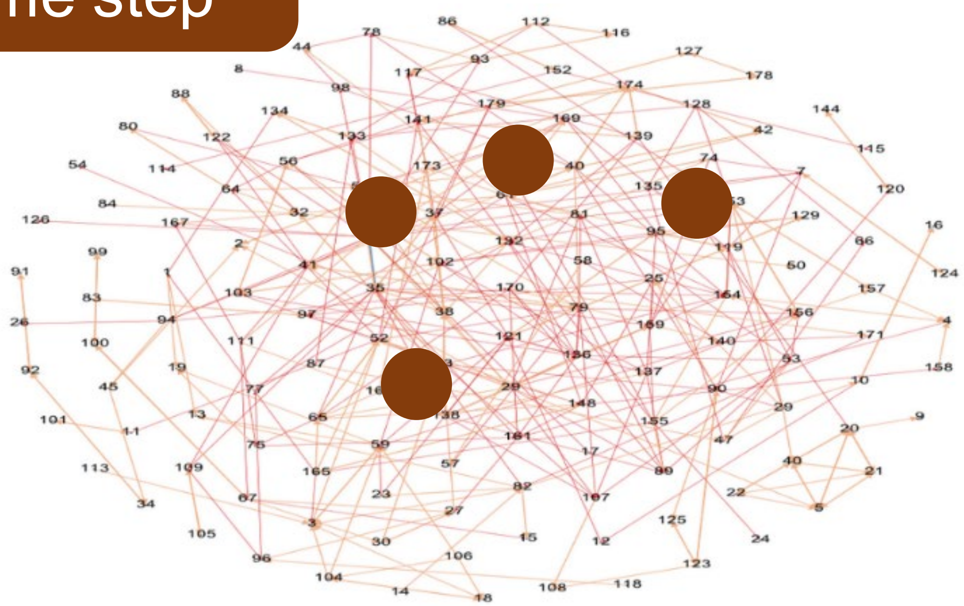

K = 4
2nd time step

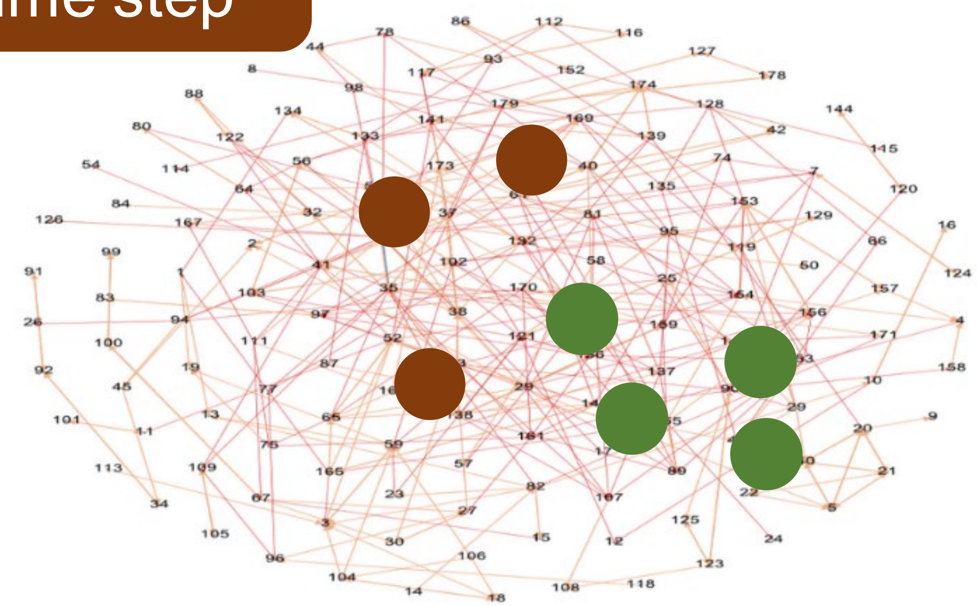

- History: *{(Invited-at-time1, absent-at-time1), (Invited-at-time2, absent-at-time2),…}*
- Provide policy *p: History(t) → Invite peer leaders for t+1*
- State of network is unobservable

POMDPs for Multi-Step Policy for Robust, Dynamic Influence Maximization
(AAMAS 2018a)

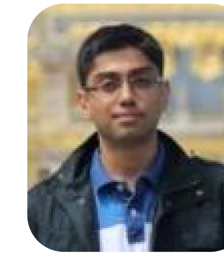

Yadav

HIDDEN STATE

Action

Choose nodes

POMDP Policy

Observation: Node presence

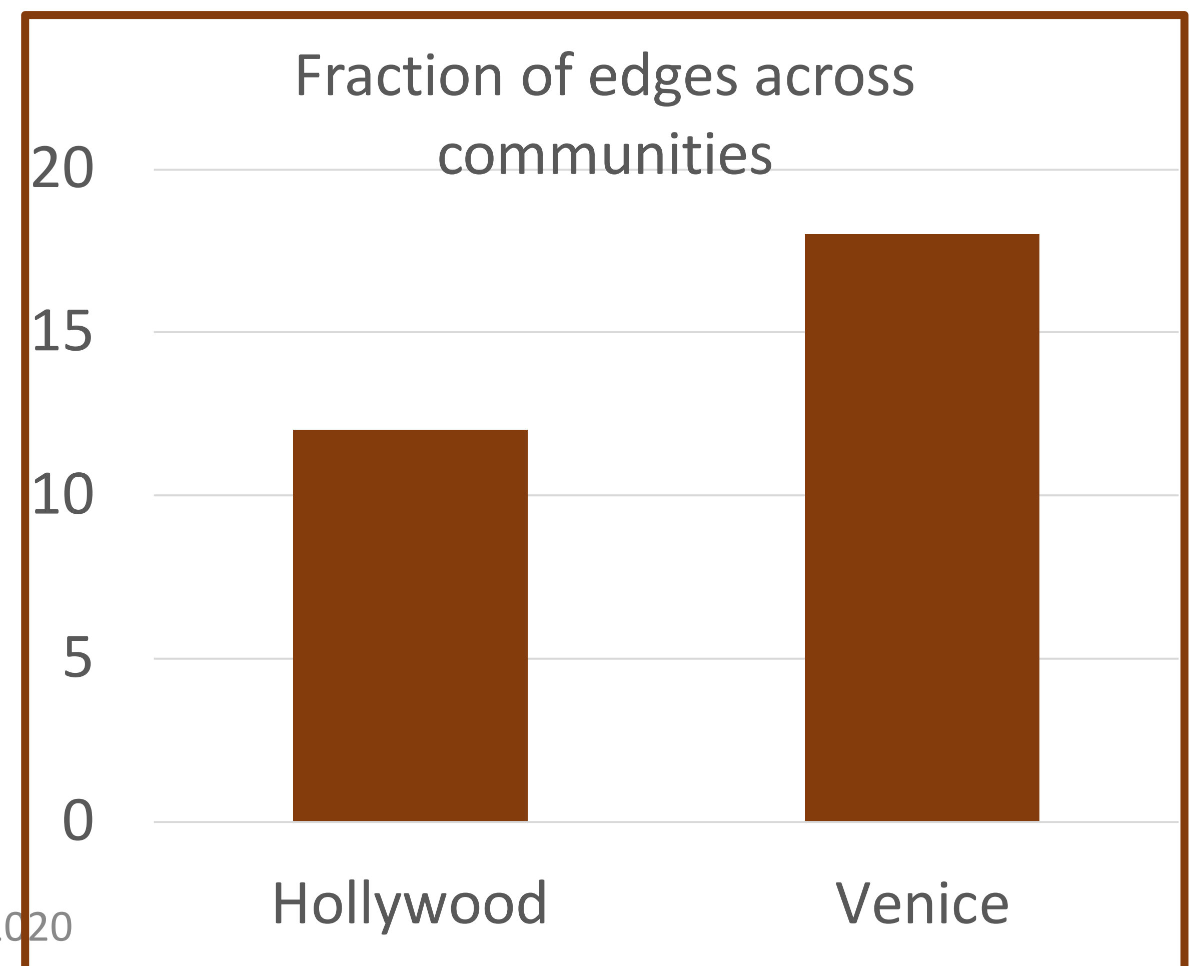

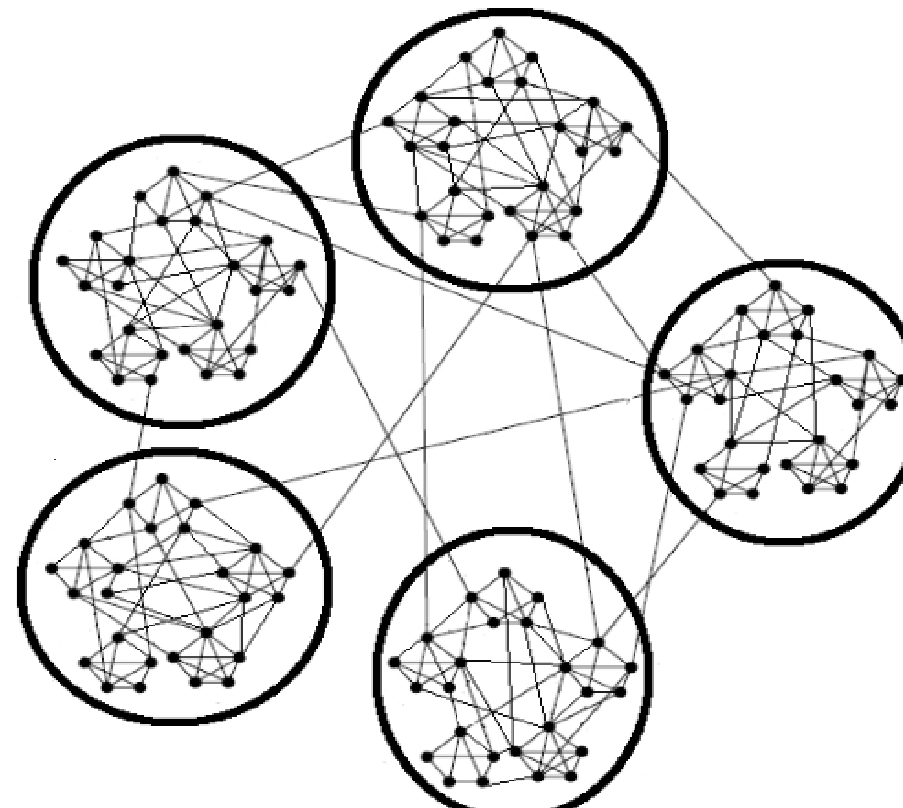

Partition POMDPs: Exploit community structure

Challenge 3:
Sampling to avoid Data Collection Bottleneck
(AAAI 2018)

Wilder

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

Data collection costly → **Sample 20%** → **Sampling Algorithm**

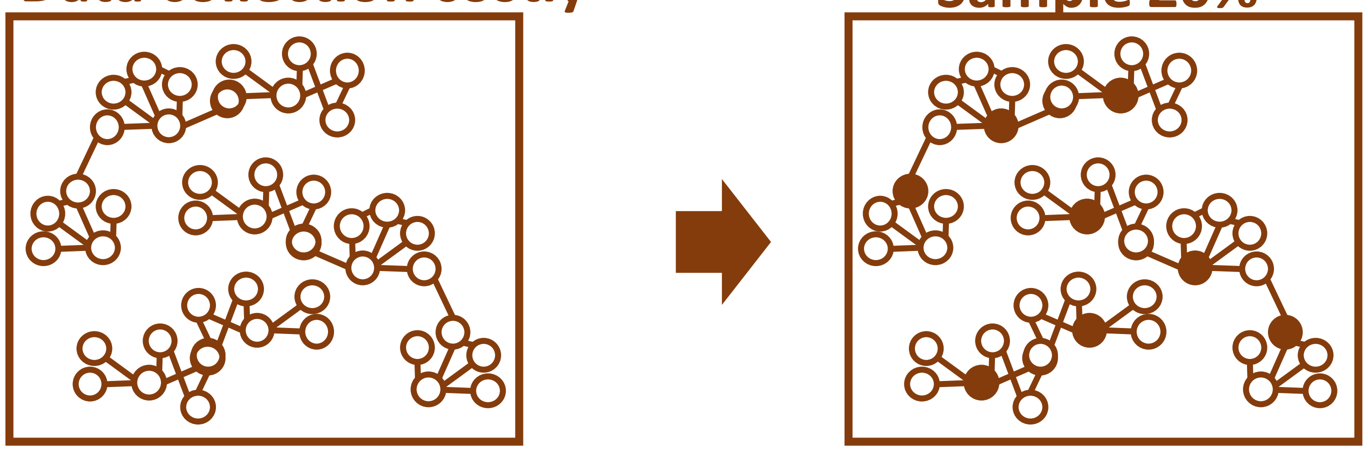

Sampling Algorithm:
Sample node randomly & estimate size of its community;
Choose seeds from largest K communities

- Input: total number of nodes, n
- Sampling algorithm queries upto query budget
- Output K seed nodes; spread influence via independent cascade model
- Compare to OPT, best influence spread by algorithm with full network

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

“Sampling-HEALER”
Pilot tests with 230 Homeless Youth
(IJCAI 2018)


Yadav


Wilder

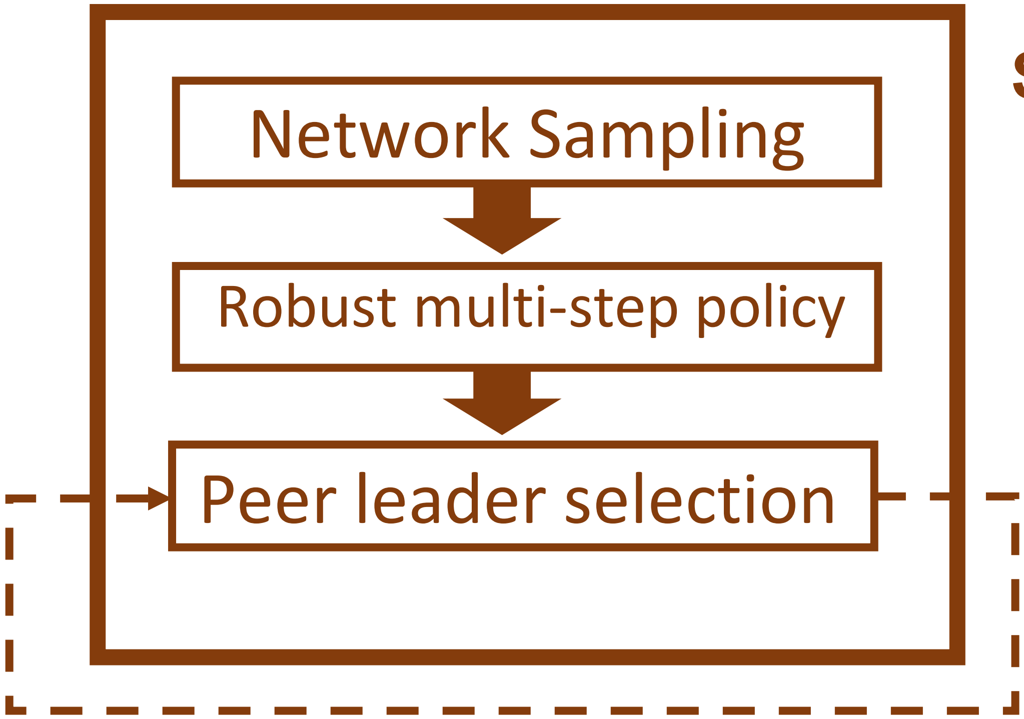

Sampling-HEALER

Observe peer leaders present/absent


12 peer leaders

Sampling HEALER (Sampled Network)	HEALER (Full Network)	HEALER+ (Full Network)	DEGREE CENTRALITY (Full Network)
60 youth	62 youth	56 youth	55 youth

Results: Pilot Studies

(Journal of Society of Social Work & Research 2018)

Yadav

Wilder

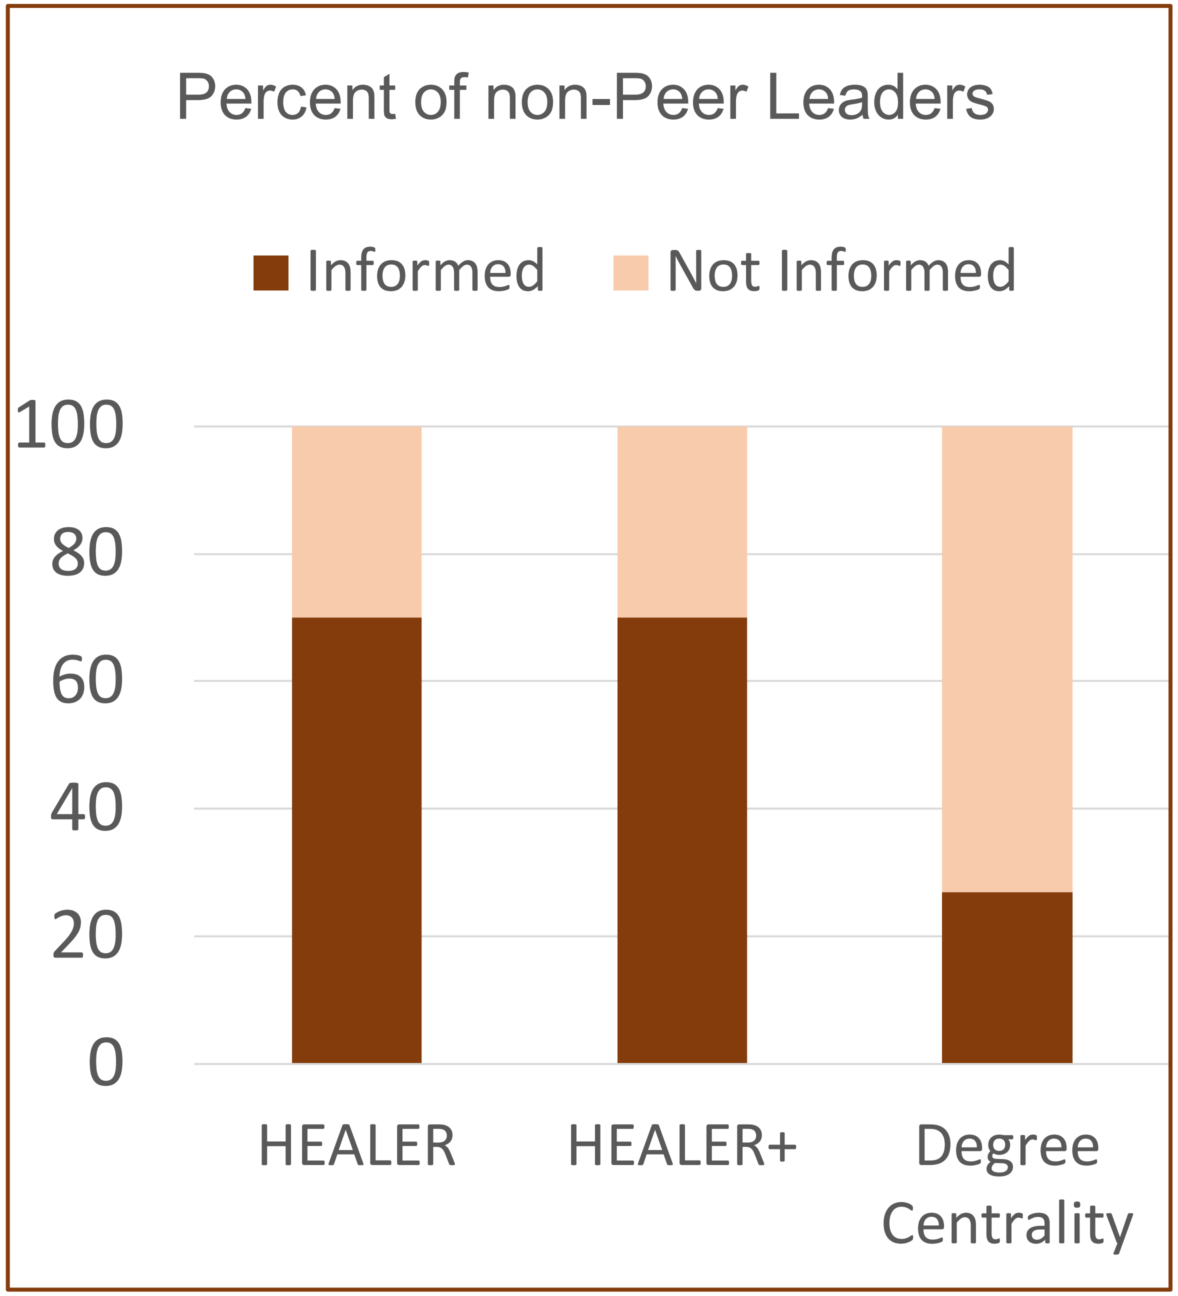

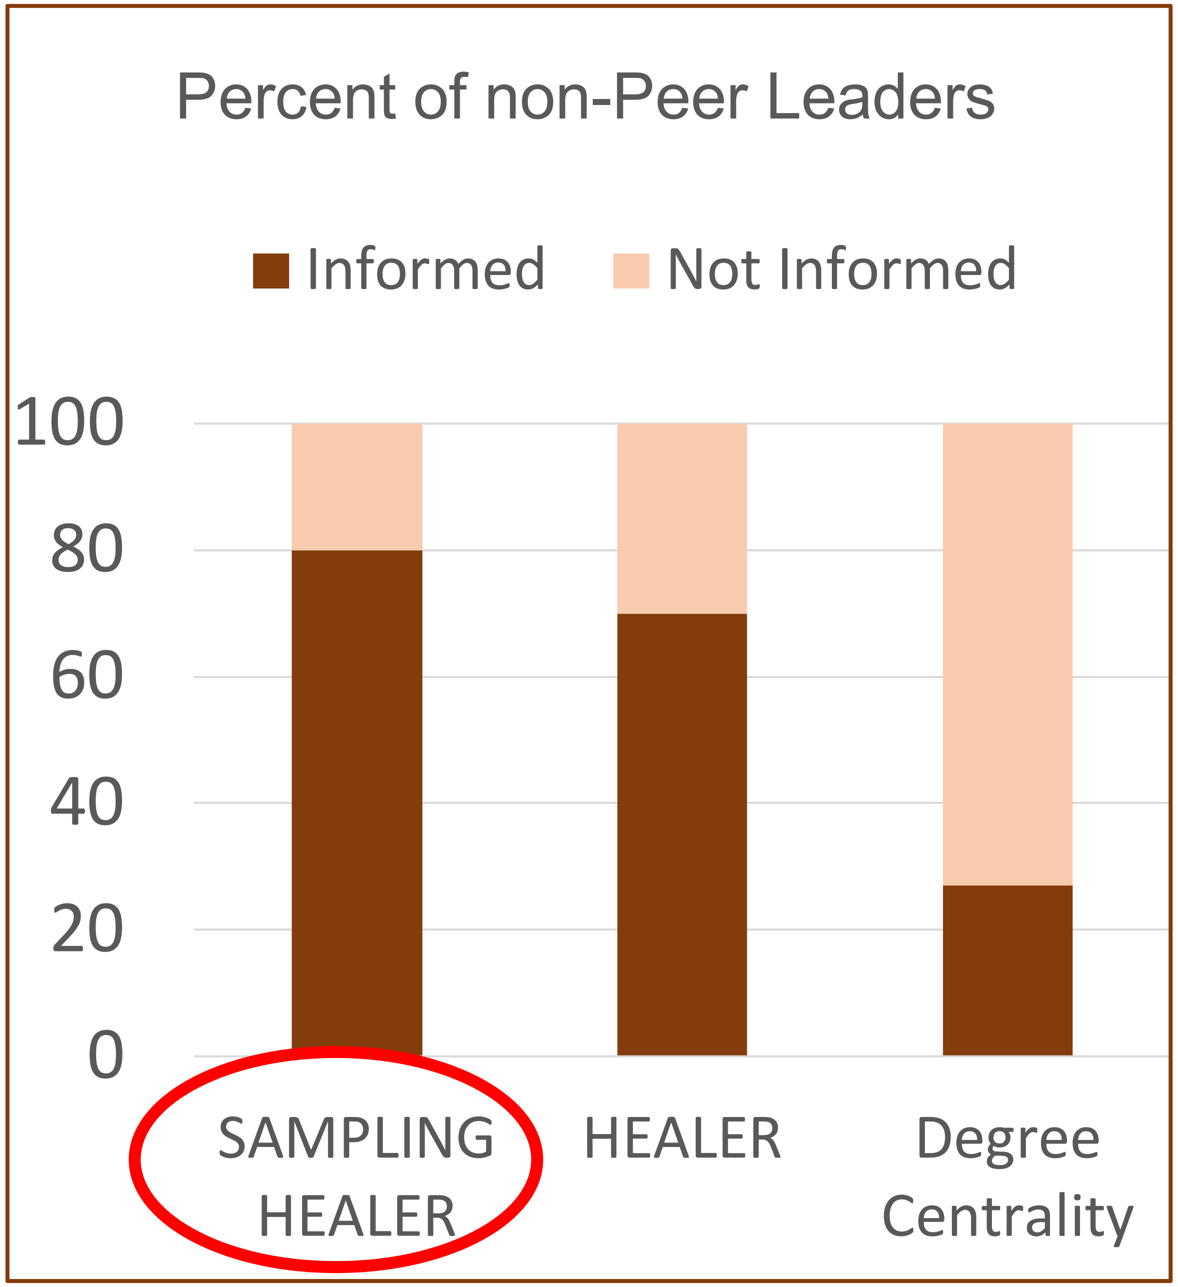

Results of 800 Youth Study [with Prof. Eric Rice]
Actual Change in Behavior?
(Under submission)

First large-scale application of influence maximization for public health

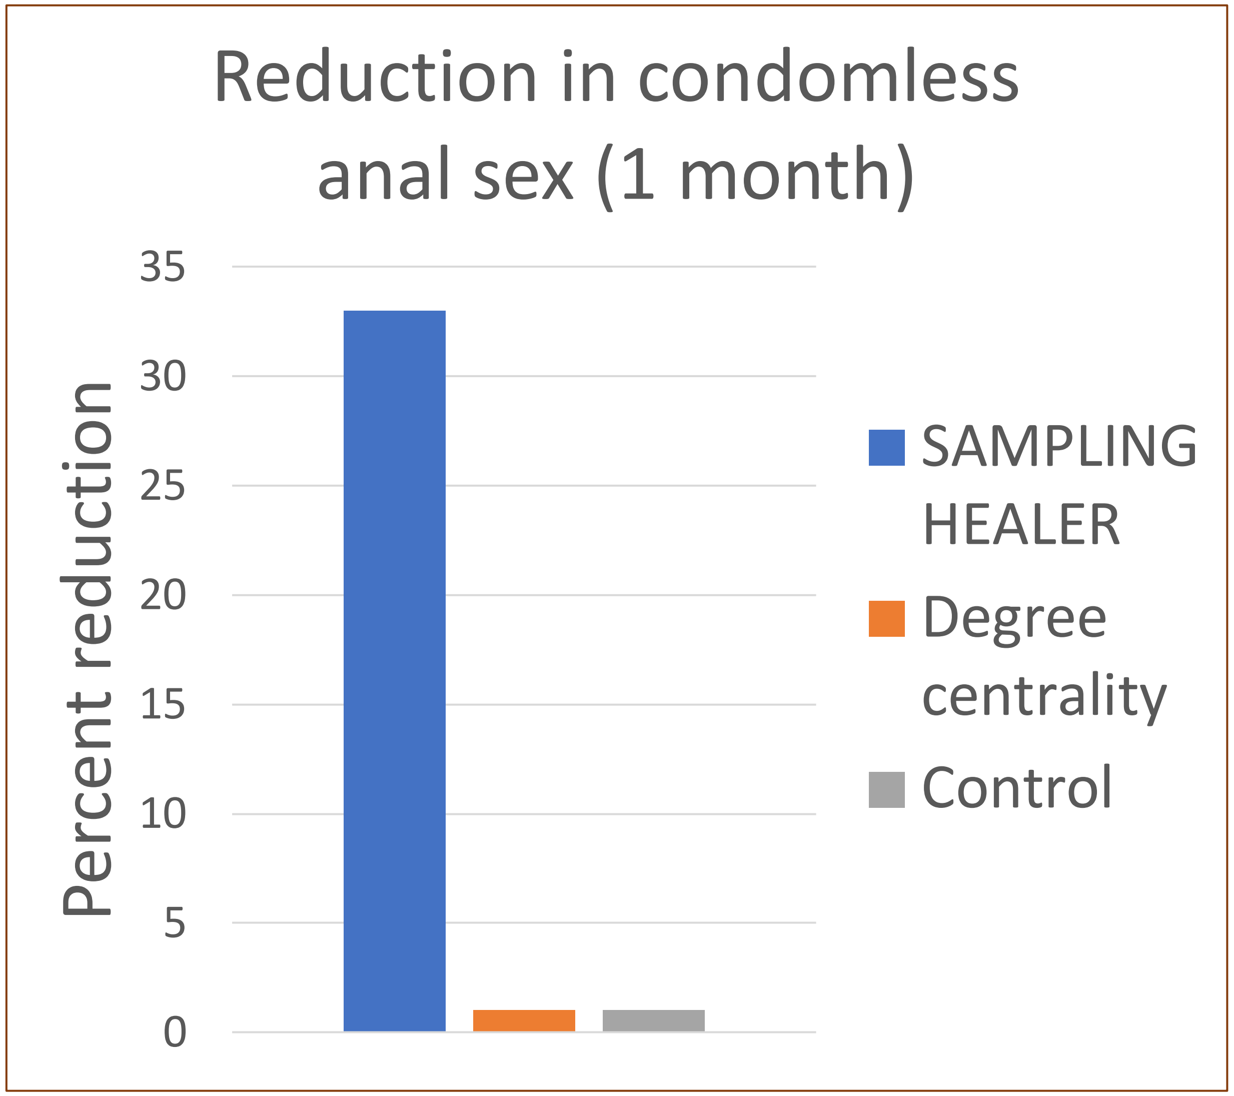

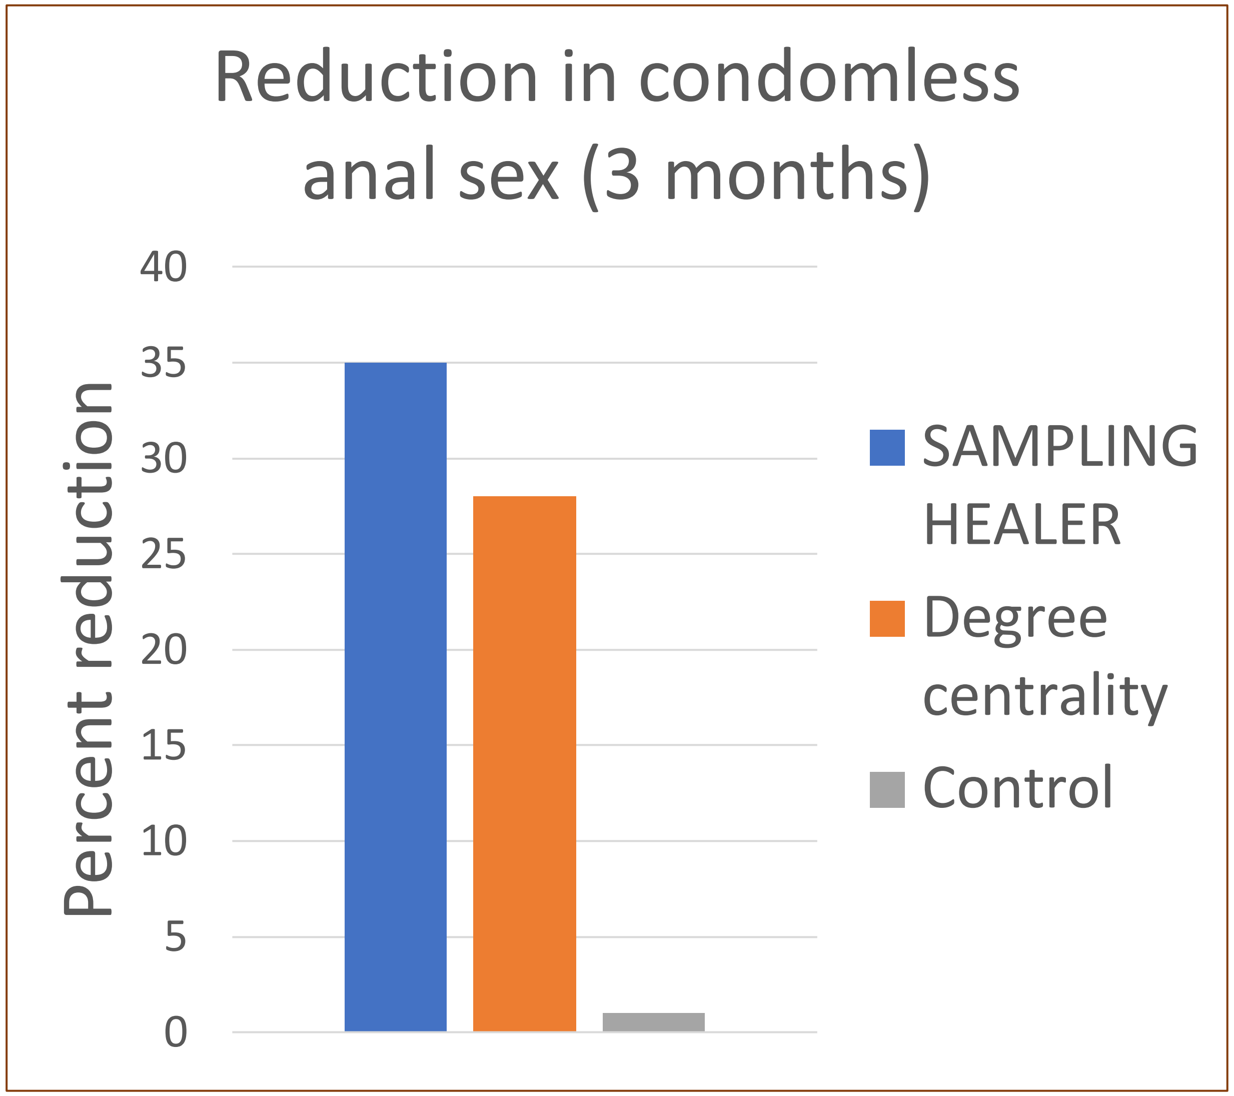

Results of 800 Youth Study [with Prof. Eric Rice]

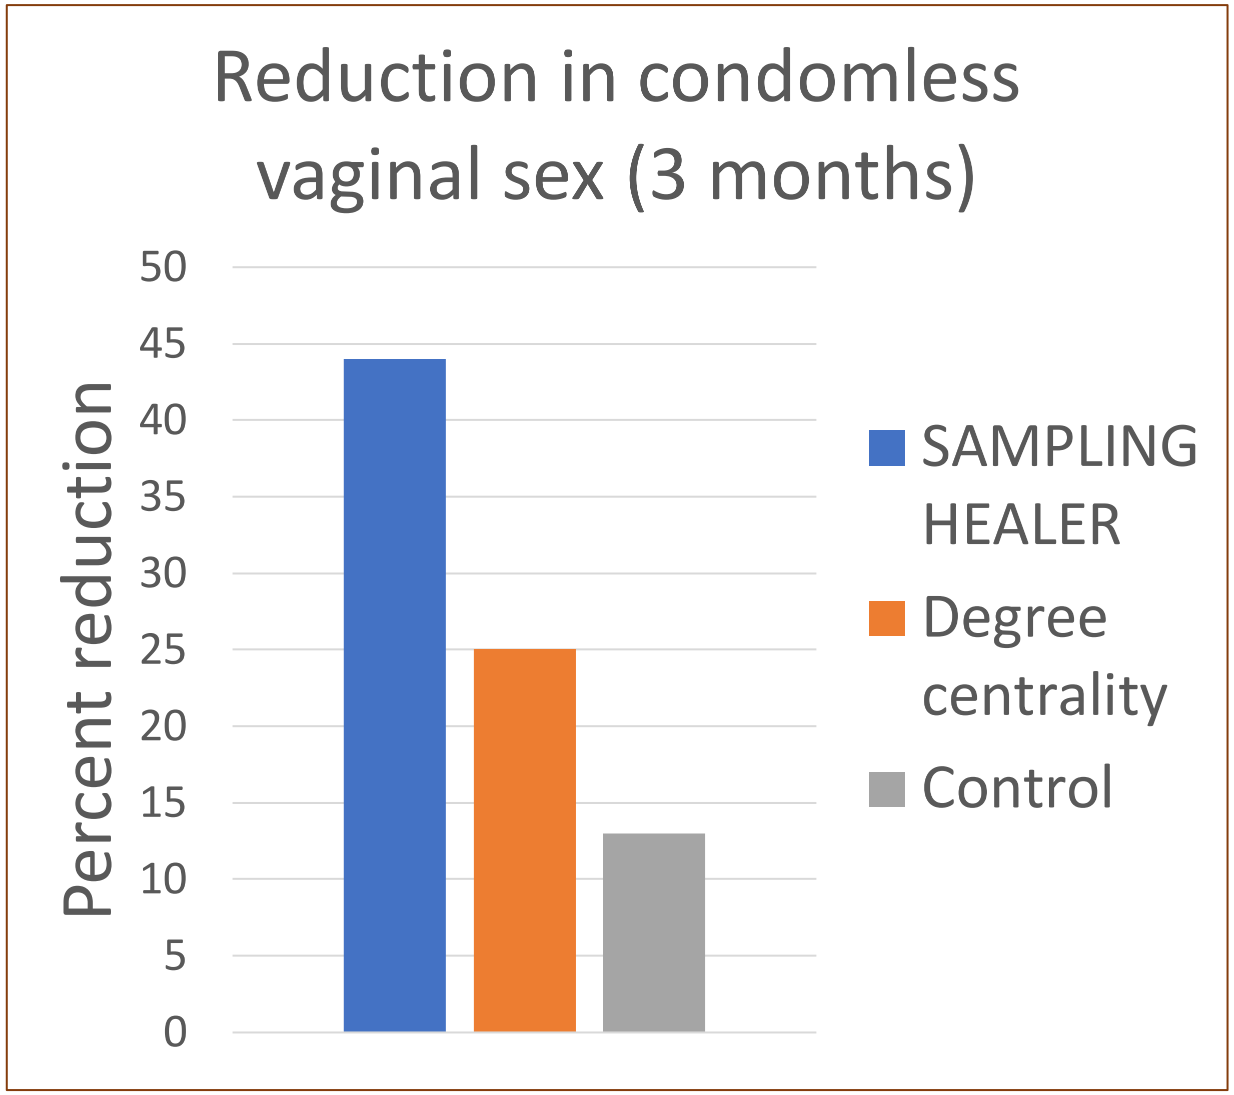

AI Assistant: HEALER

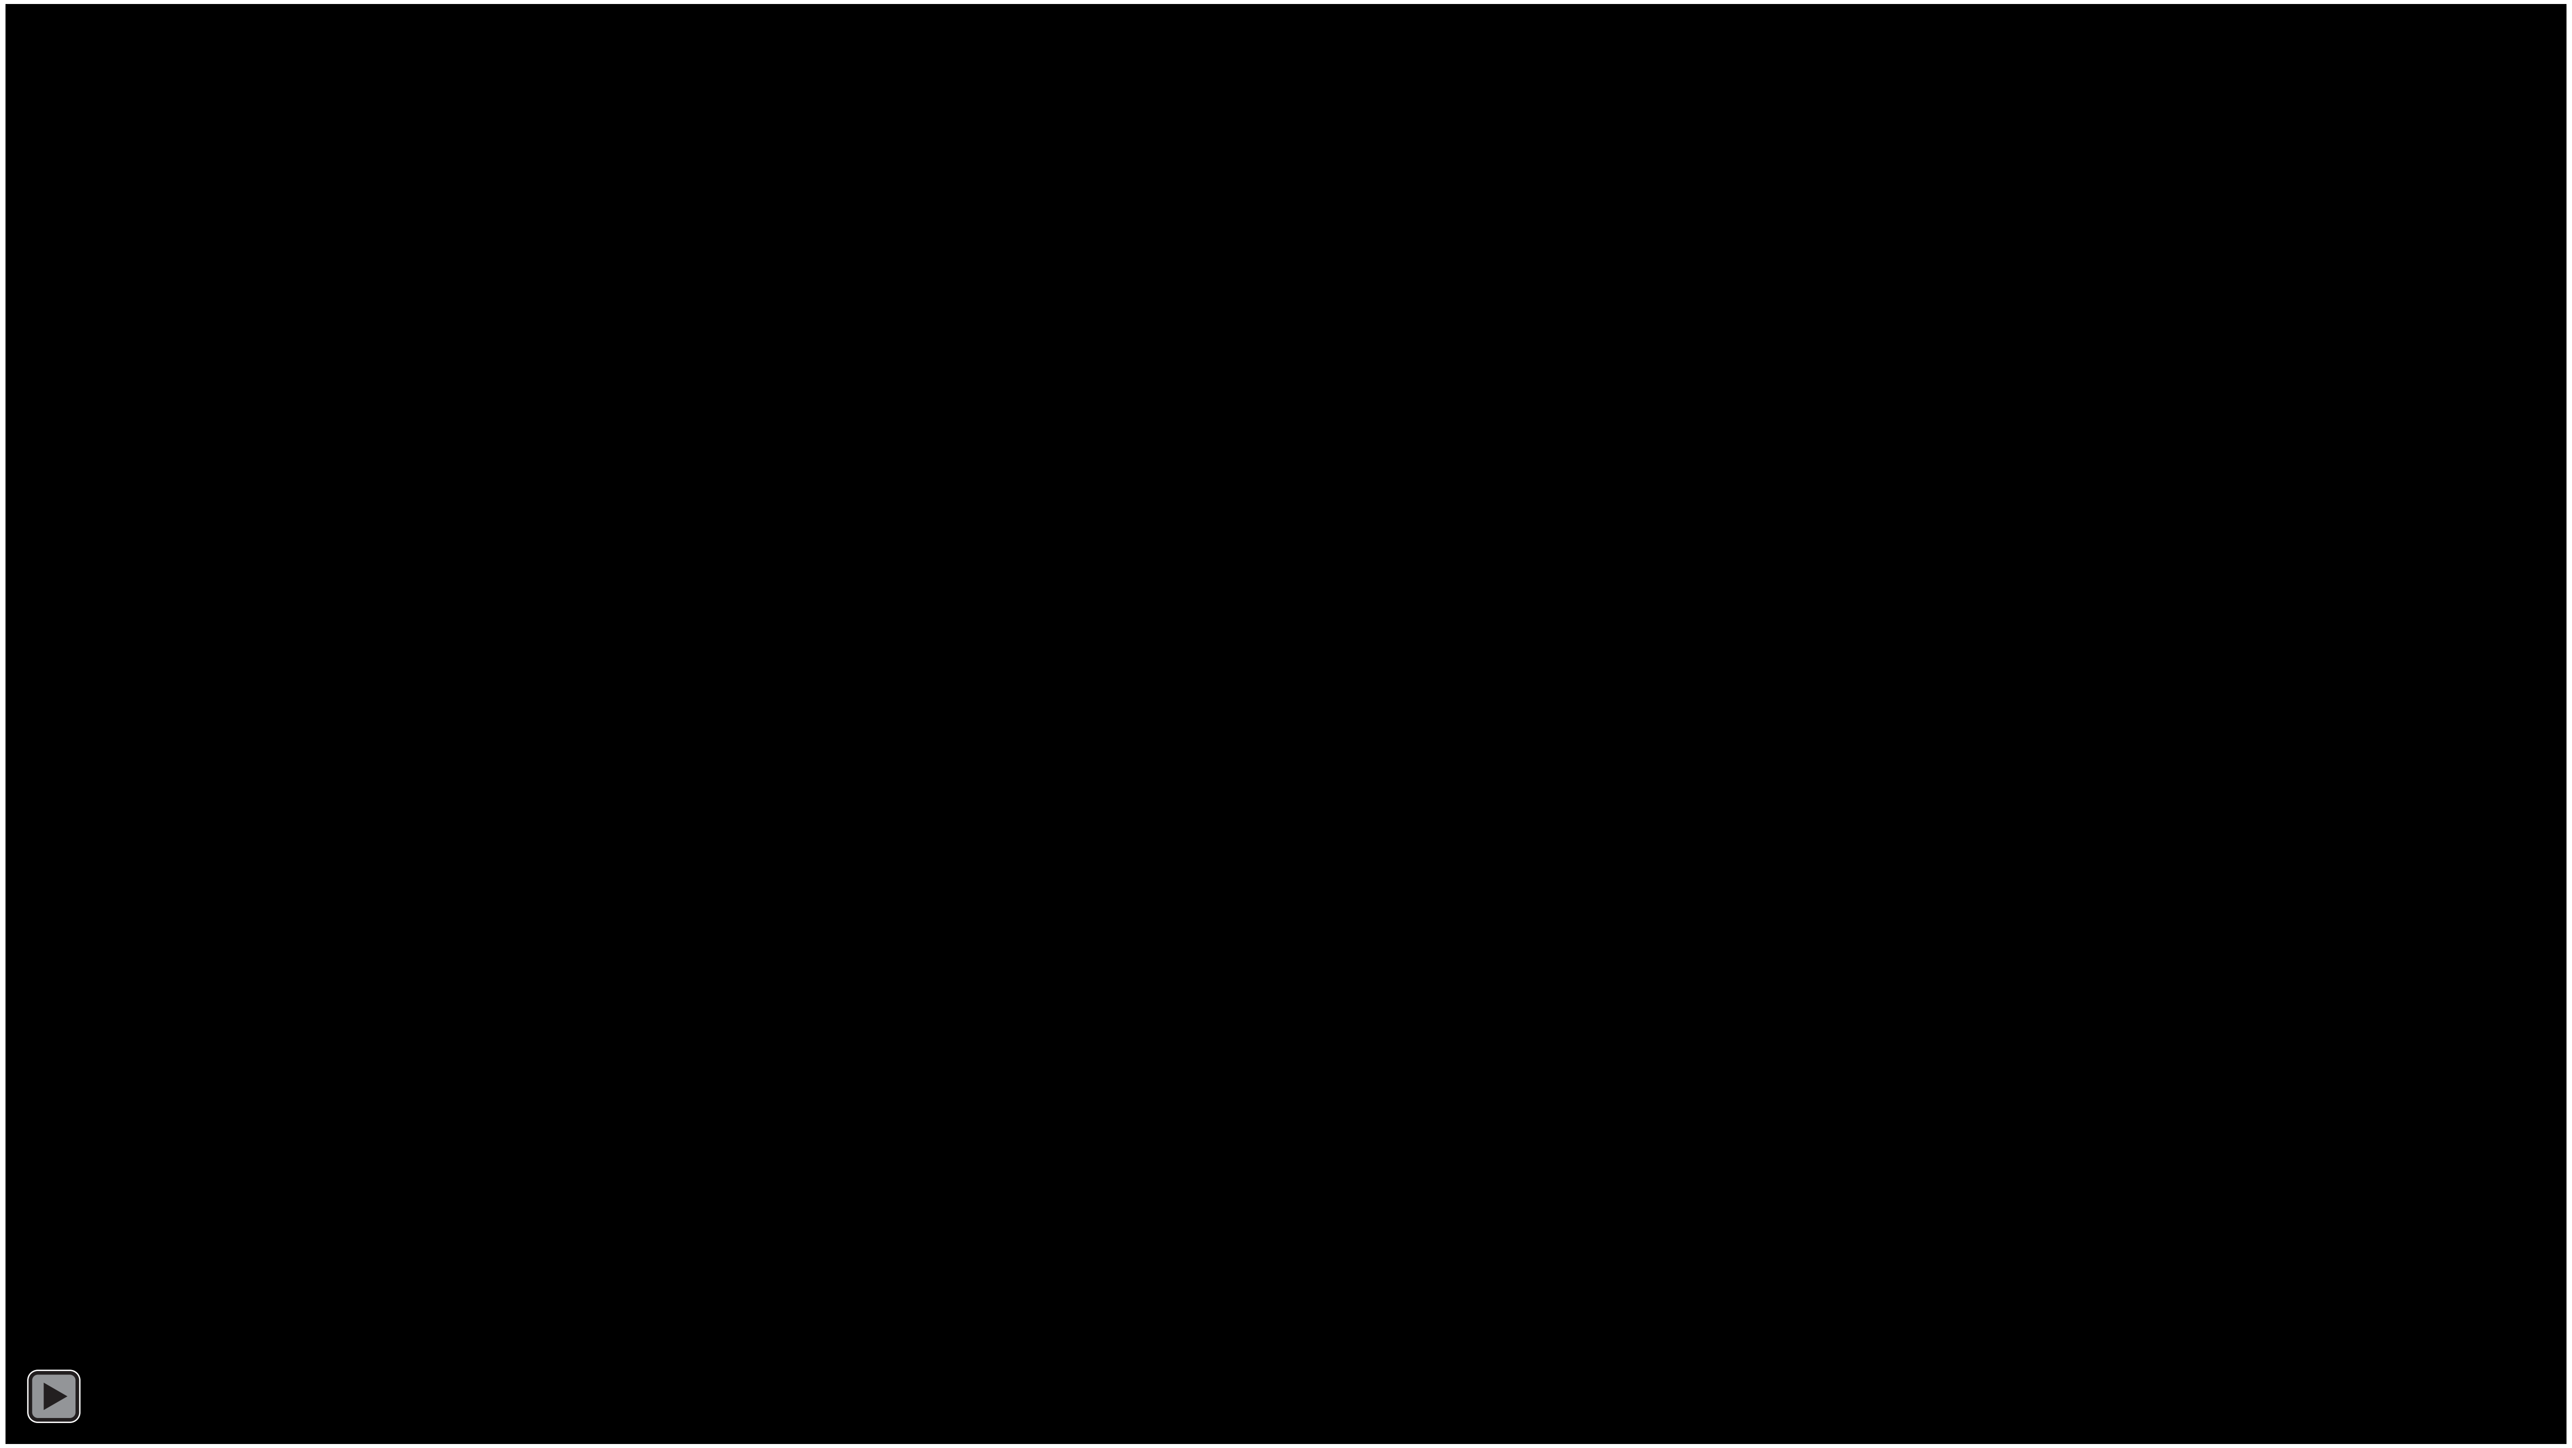

Next steps: Data to Deployment Pipeline Using an RL agent?

(with B. Ravindran & team, AAMAS 2020)

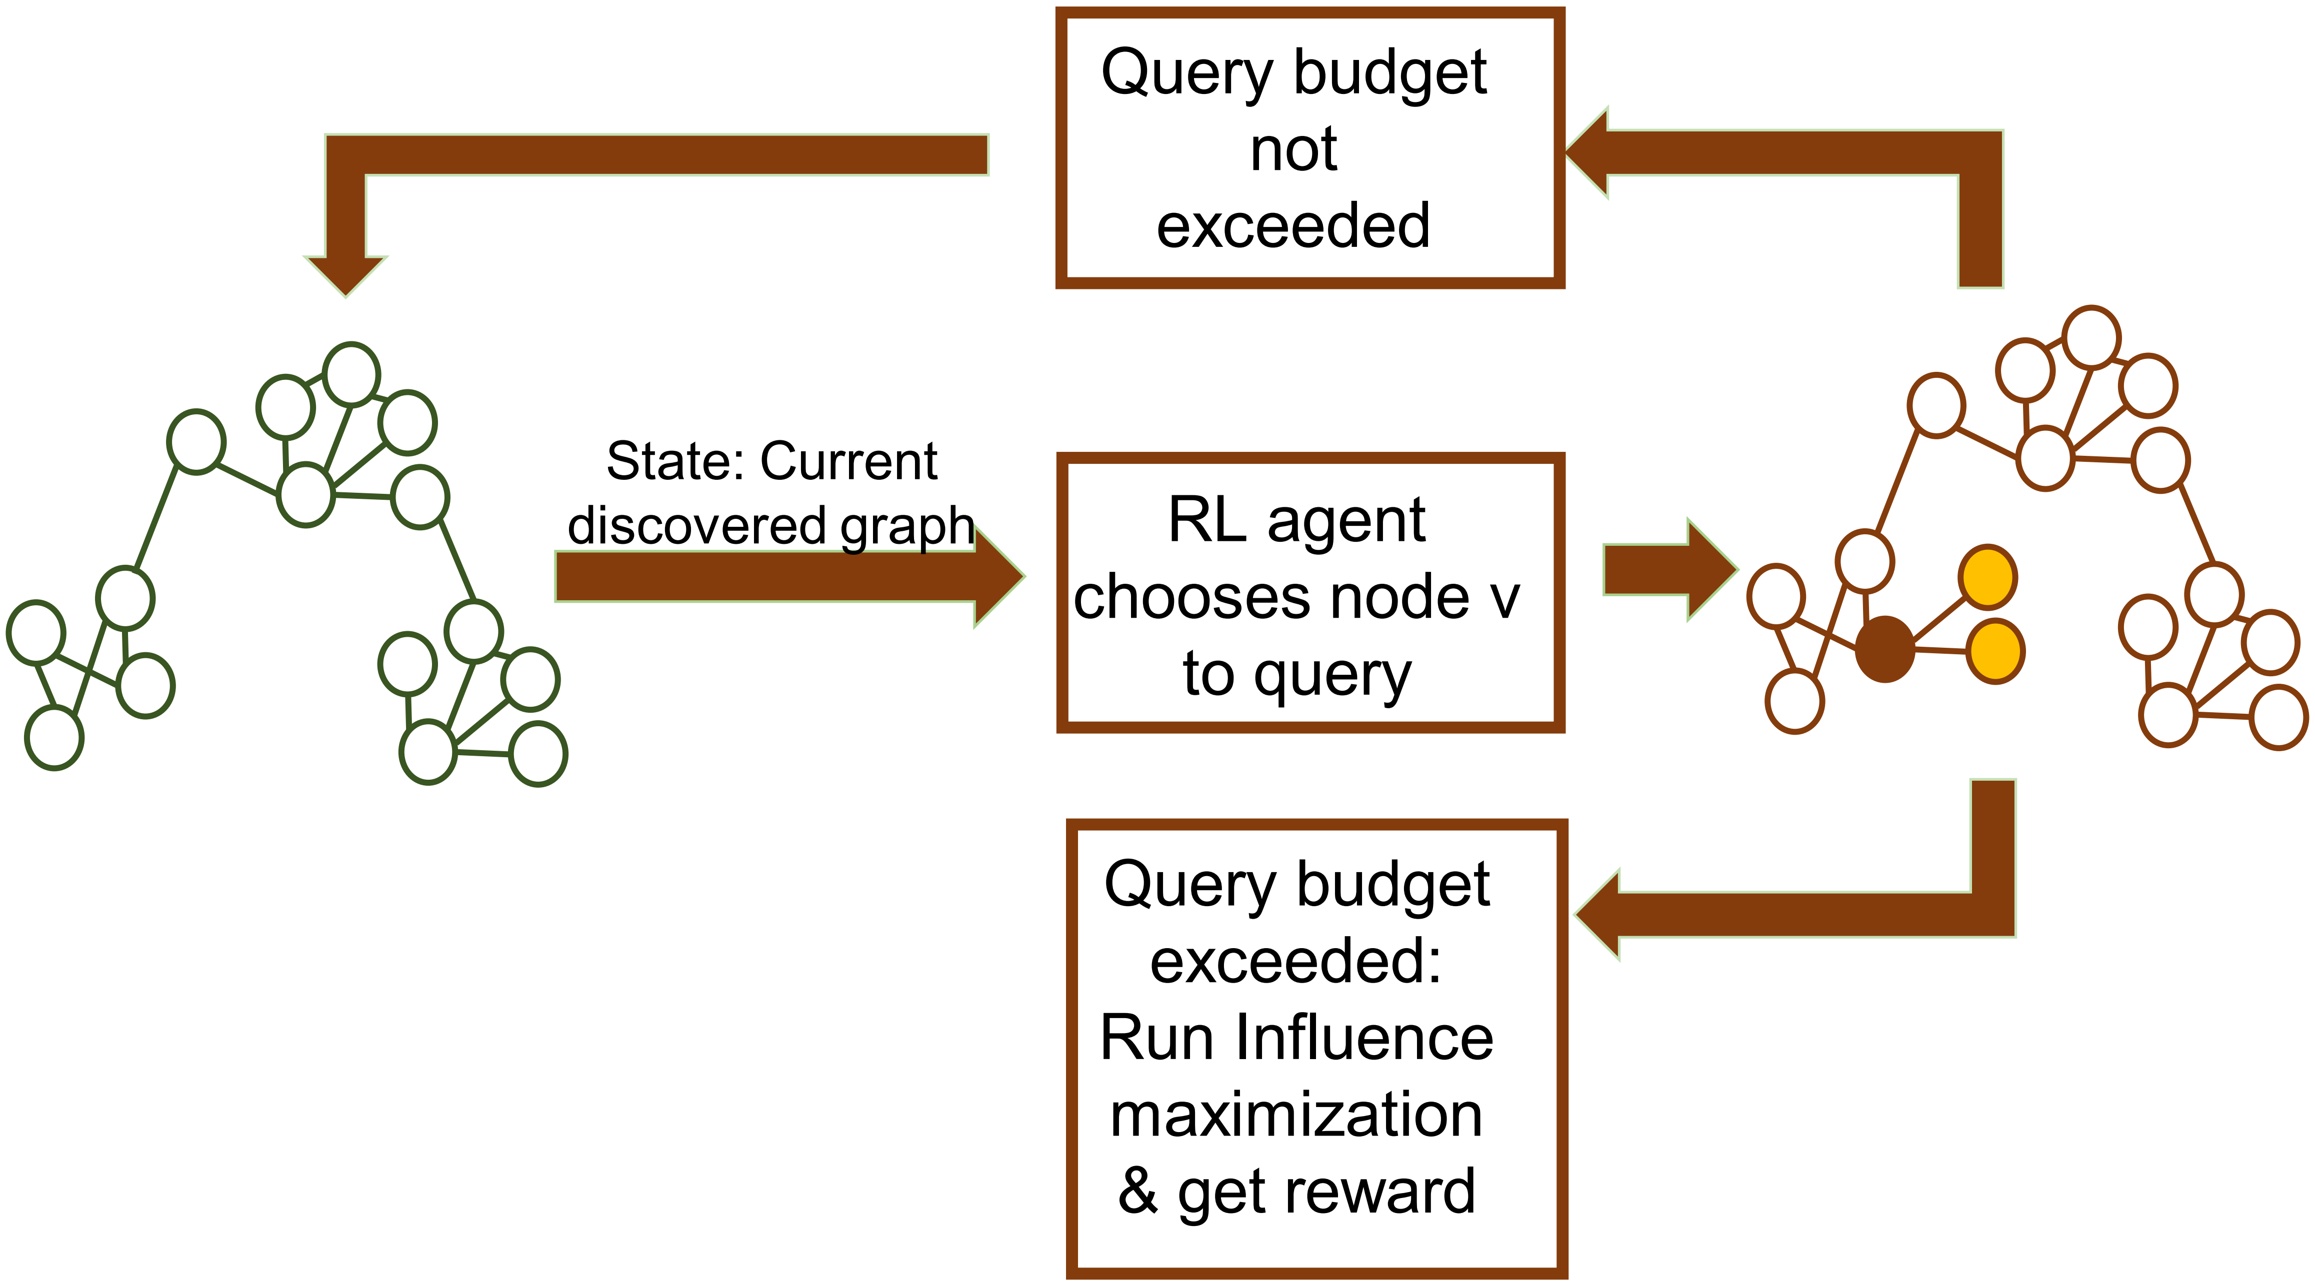

Next steps: Data to Deployment Pipeline Using an RL agent?

(with B. Ravindran & team, AAMAS 2020)

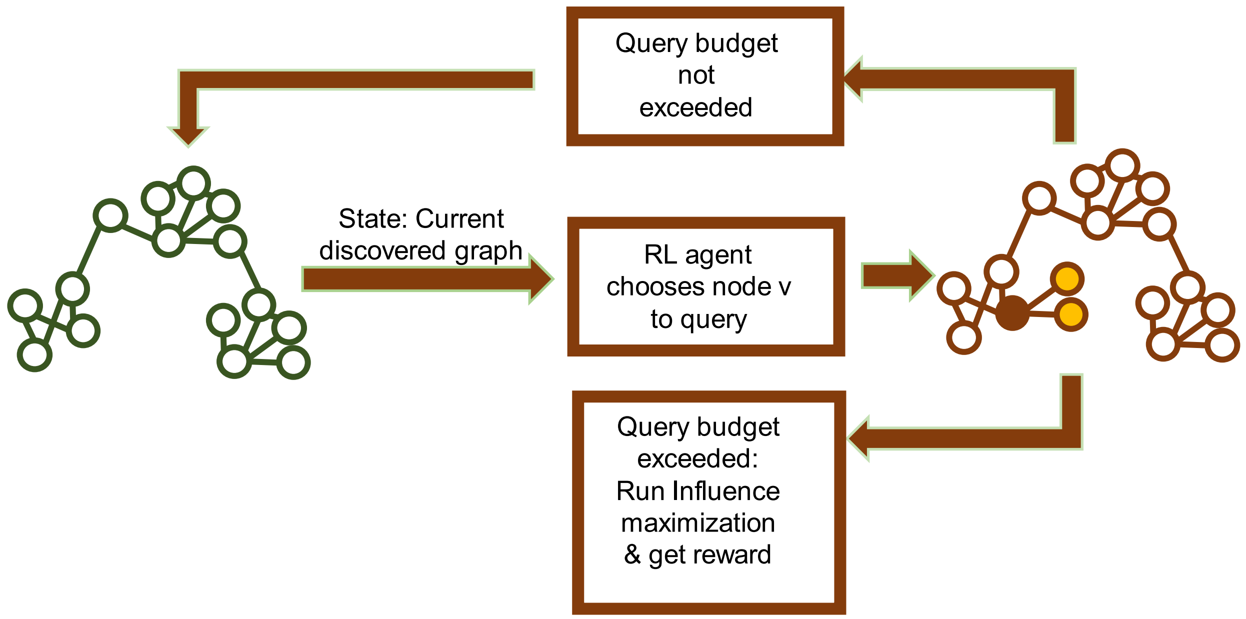

Network Family		Improve %
Rural		23.76
Animal		26.6
Retweet		19.7
Homeless		7.91

Fairness in Reasoning with Social Networks: Suicide Prevention via Gatekeeper Selection

(NeurIPS 2019, IJCAI 2019)

Rahmattalabi

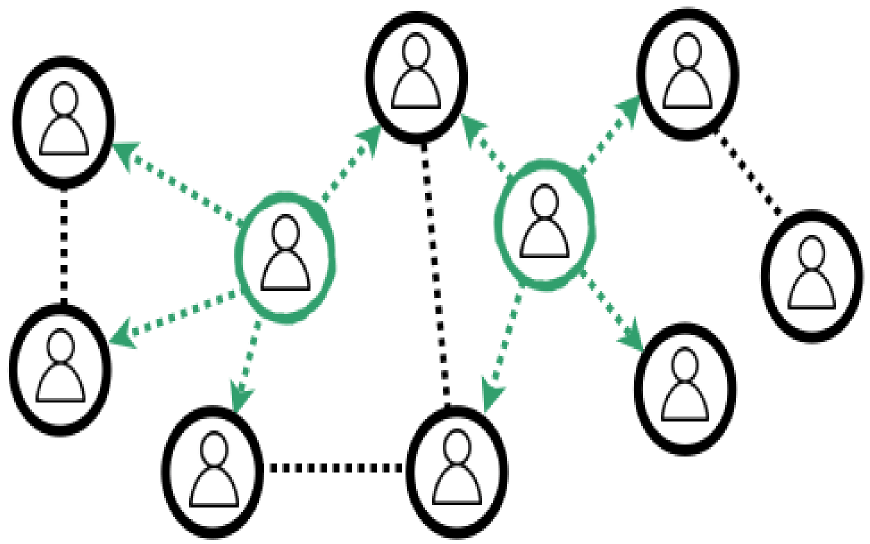

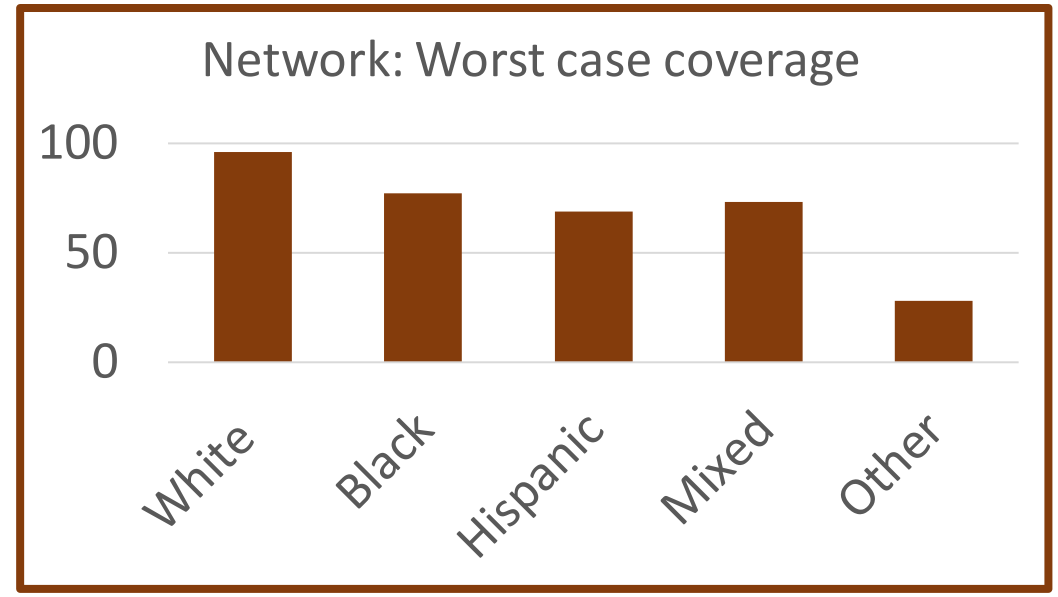

Robust graph covering with gatekeepers, maximize worst case coverage

Disparity in coverage across racial groups

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

Y: Max of minimum utility for any community

Diversity constraints: $u_c(A) \geq U_c$

U_c: *Utility if # gatekeepers allocated proportional to size of community*

Outline

Public Health

- *Information dissemination & behavior change: Social networks*
- *Health program adherence: Passive via ML vs Active via bandits*
- *COVID-19: Agent-based modeling*

Conservation

Health Program Adherence
Maternal & Child Care in India
(Under submission)

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

- *ARMMAN: 18 Million women enrolled, 160000 health workers…*
- *mMitra: Weekly call to new/expectant moms; friendly 3 minute messages about health*
- *mMitra: Significant benefits shown; 2.2 million women enrolled*
- *Unfortunately, significant fraction low-listeners or drop-outs*

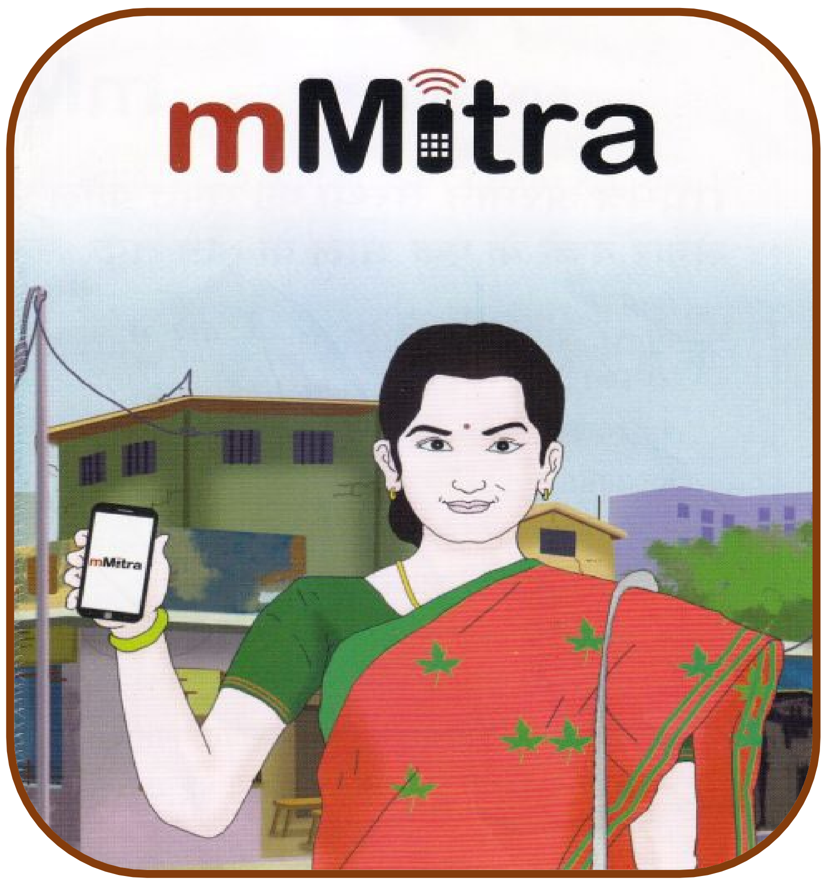

Passive Adherence Monitoring Maternal & Child Care in India

(with B Ravindran IIT Madras)

Predict beneficiaries likely to drop out: Allows ARMMAN to focus intervention

- *Neural networks for prediction: Extensive tests with past data*
- *Results of pilot with 18000 beneficiaries: High precision, recall, accuracy*

- ***Prediction software deployed: helps 300,000 beneficiaries in mMitra***

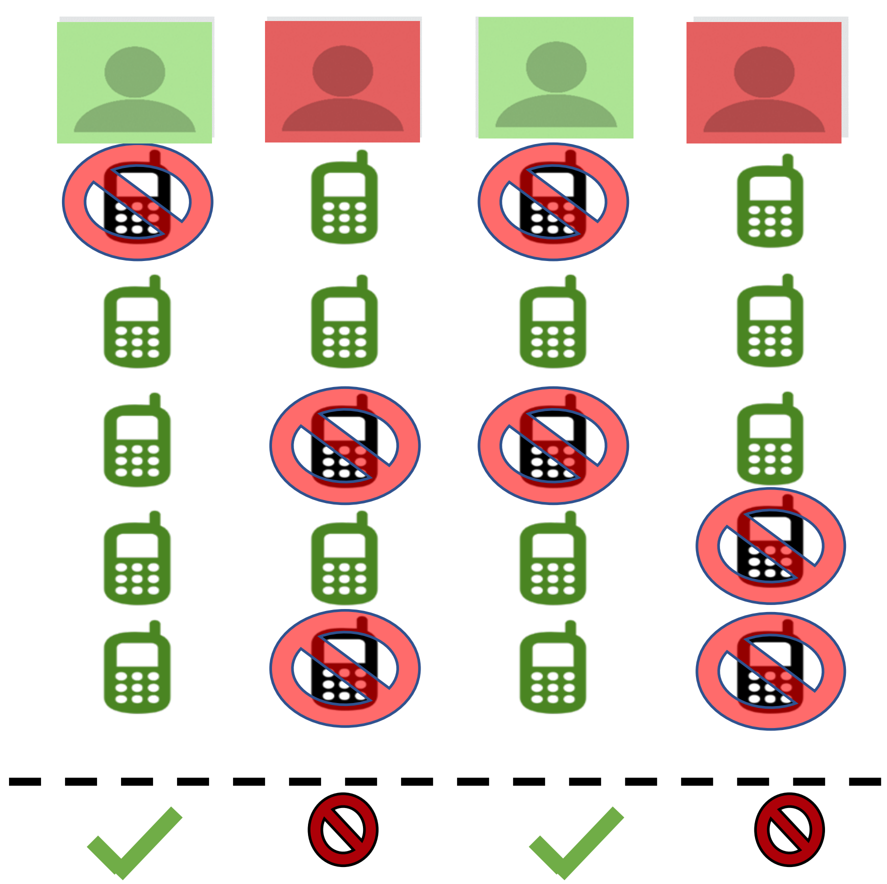

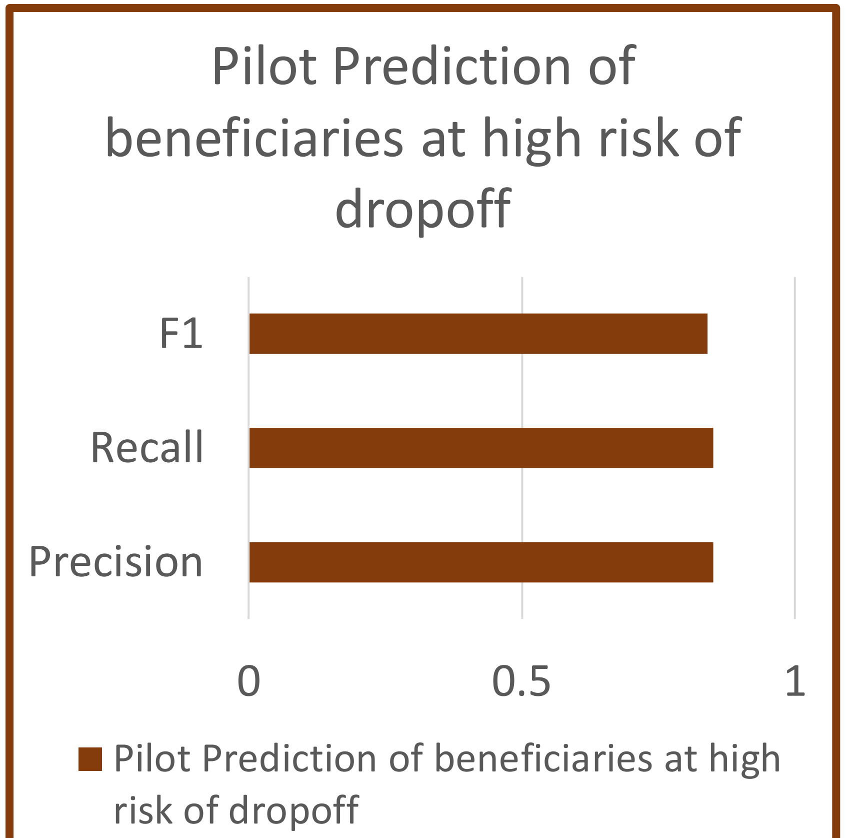

ARMMAN Pilot

- *18000 Beneficiaries*

  Nov & Dec 2019

- *Test: Jan-April 2020*

Passive Adherence Monitoring Preventing Tuberculosis in India

(KDD 2019)

Killian

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

- *Non-adherence to TB Treatment: Digital adherence tracking via daily phone calls*
- *Intervene before patients miss dose*
- *Predict adherence risk from phone call patterns?*
- *Results from Mumbai, India: 15,000 patients, 1.5 million phone calls*

everwell

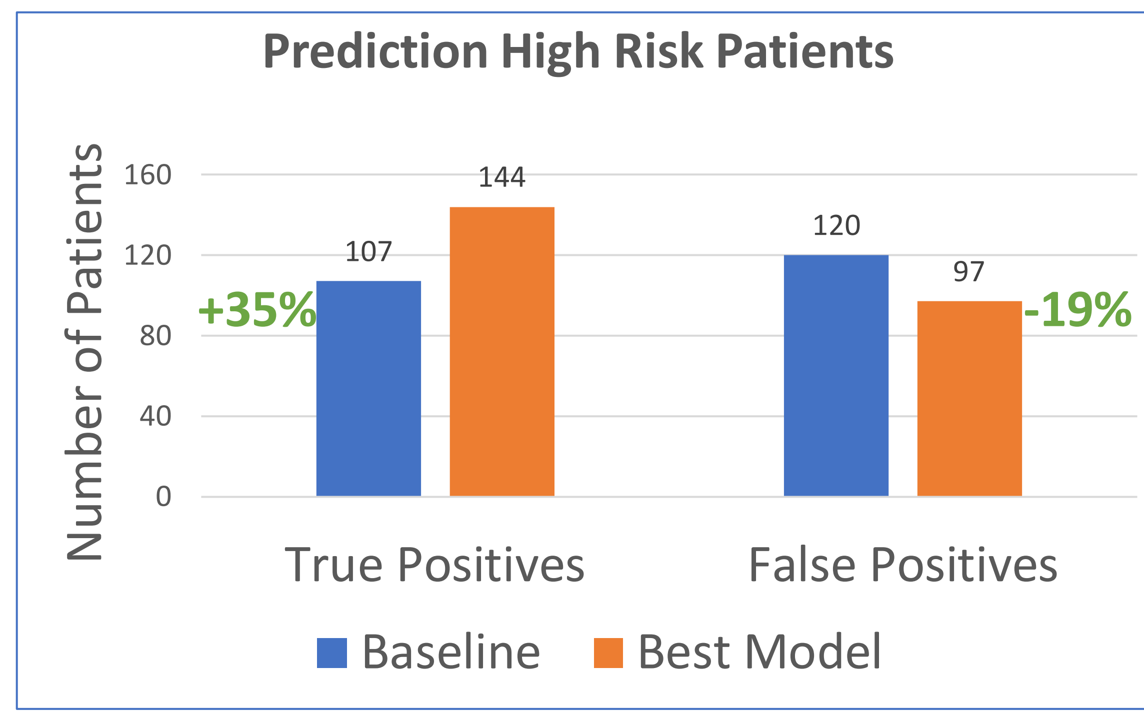

Intervention Reasoning: Active Adherence Monitoring

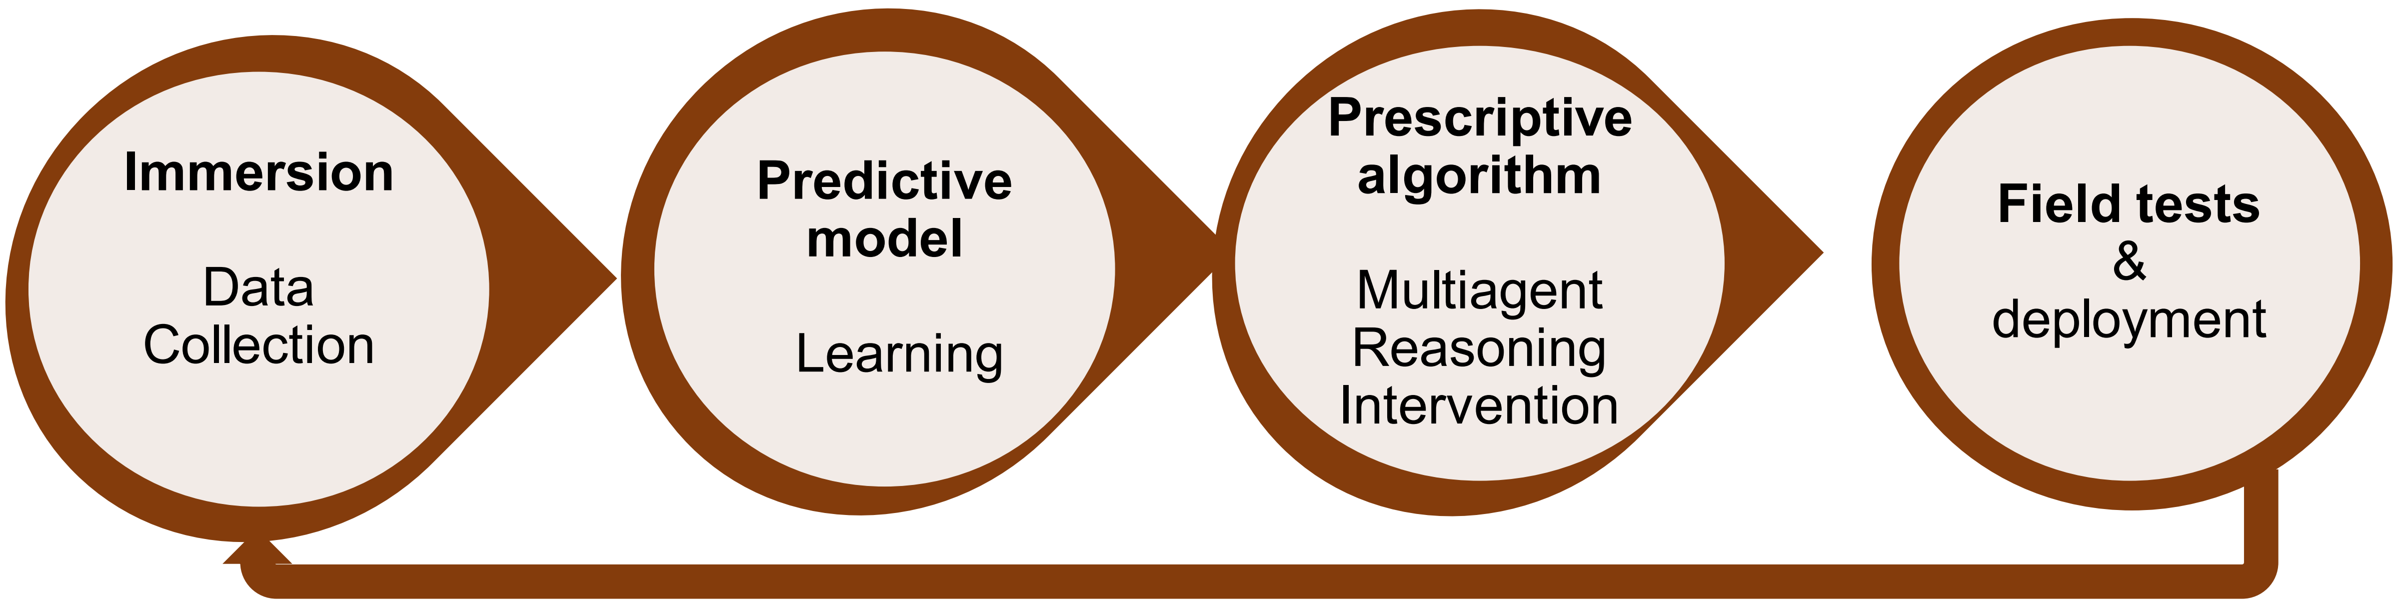

Health worker calls, patients do not call

Intervention Scheduling with Scarce Data: Passive vs Active Adherence Monitoring

(Under submission)


Mate Killian

Health worker

Call patients: Track, improve adherence

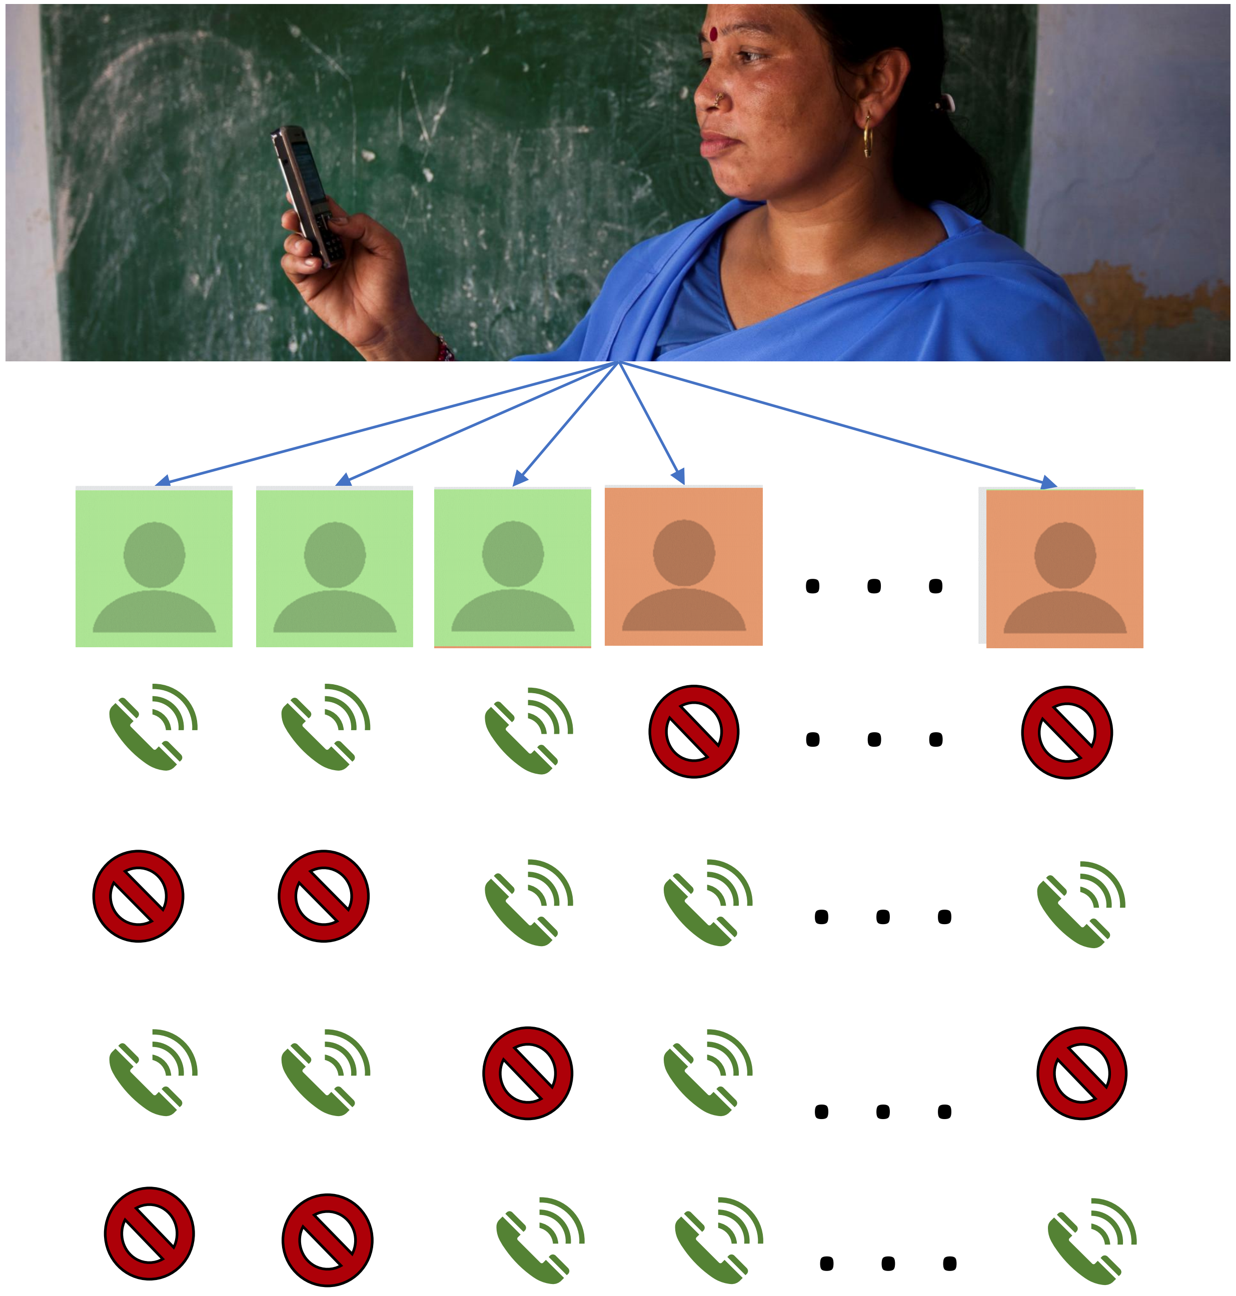

Challenge:

- *Large number of patients (N)*
- *Which 'k' patients to call?*

Approach:

- *Adherence Restless Bandits*

Intervention Scheduling with Scarce Data: Adherence Restless Bandits(A-RMAB)

(Under submission)


Mate


Killian

Restless multiarmed bandits (RMAB)

Adherence RMAB (A-RMAB):

- *Each arm has binary latent state {0, 1}*
- *0= not-adhering; 1= adhering*

	not-adhering	adhering
not-adhering	0.90	0.10
adhering	0.05	0.95

Intervention Scheduling with Scarce Data: Adherence Restless Bandits(A-RMAB)

(Under submission)


Mate Killian

Restless multiarmed bandits (RMAB)

Adherence RMAB (A-RMAB):

- *Each arm has binary latent state {0, 1}*
- *0= not-adhering; 1= adhering*

	not-adhering	adhering
not-adhering	0.90	0.10
adhering	0.05	0.95

When arm not played

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

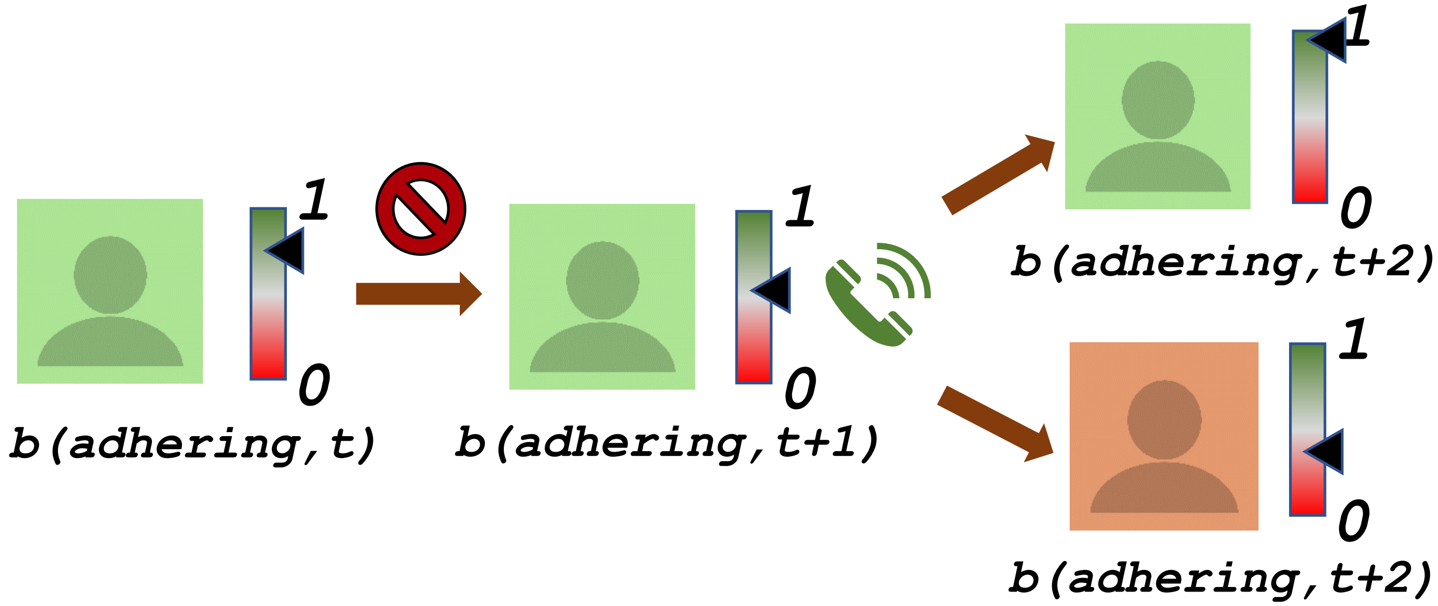

When arm is played

- *Observe current state*
- *Higher chance of adhering next round*

*P(**adhering** | call) > P(**adhering** | no call)*

Intervention Scheduling with Scarce Data: Adherence Restless Bandits(A-RMAB)


Mate


Killian

- ***Performance guarantee requires A-RMAB to be indexable***

Theorem 1: A-RMAB Indexable if threshold policies are optimal.

- ***Threshold policies:*** *Call → Belief of adherence below threshold → Call*

Forward threshold: 

Theorem 2: Forward threshold optimal if intervention effect on "Non-adherent" patients is large.

Empirically, almost all patients are threshold optimal

⟹ Fast algorithm + no sacrifice on performance

Intervention Scheduling with Scarce Data: Adherence Restless Bandits(A-RMAB)


Mate

Killian

- ***Orders of magnitude speedup with no solution quality loss***

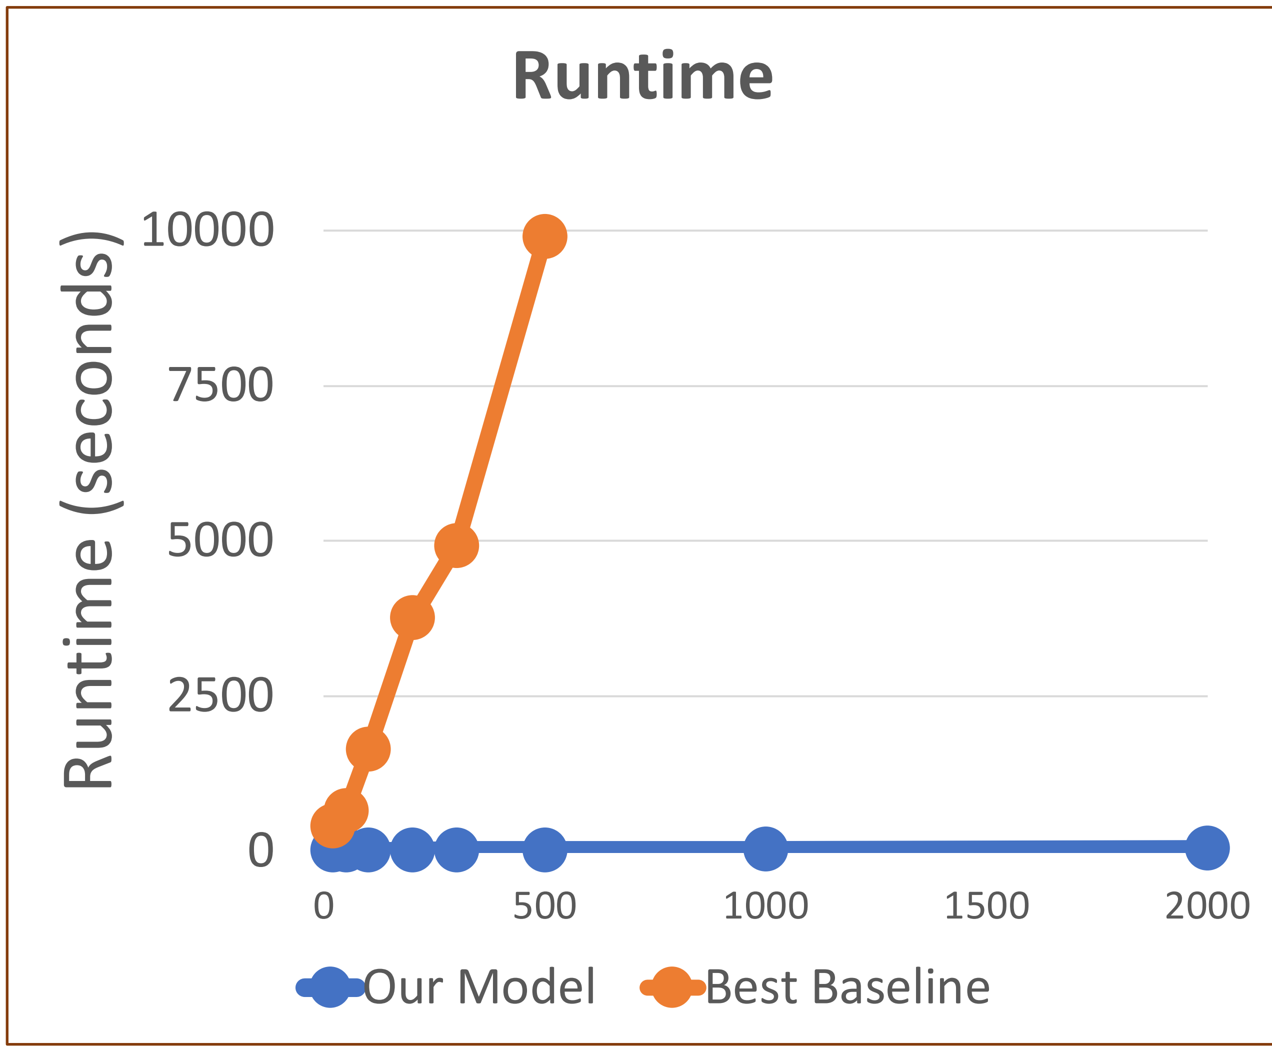

COVID-19: Agent-based Simulation Model
(Under submission)

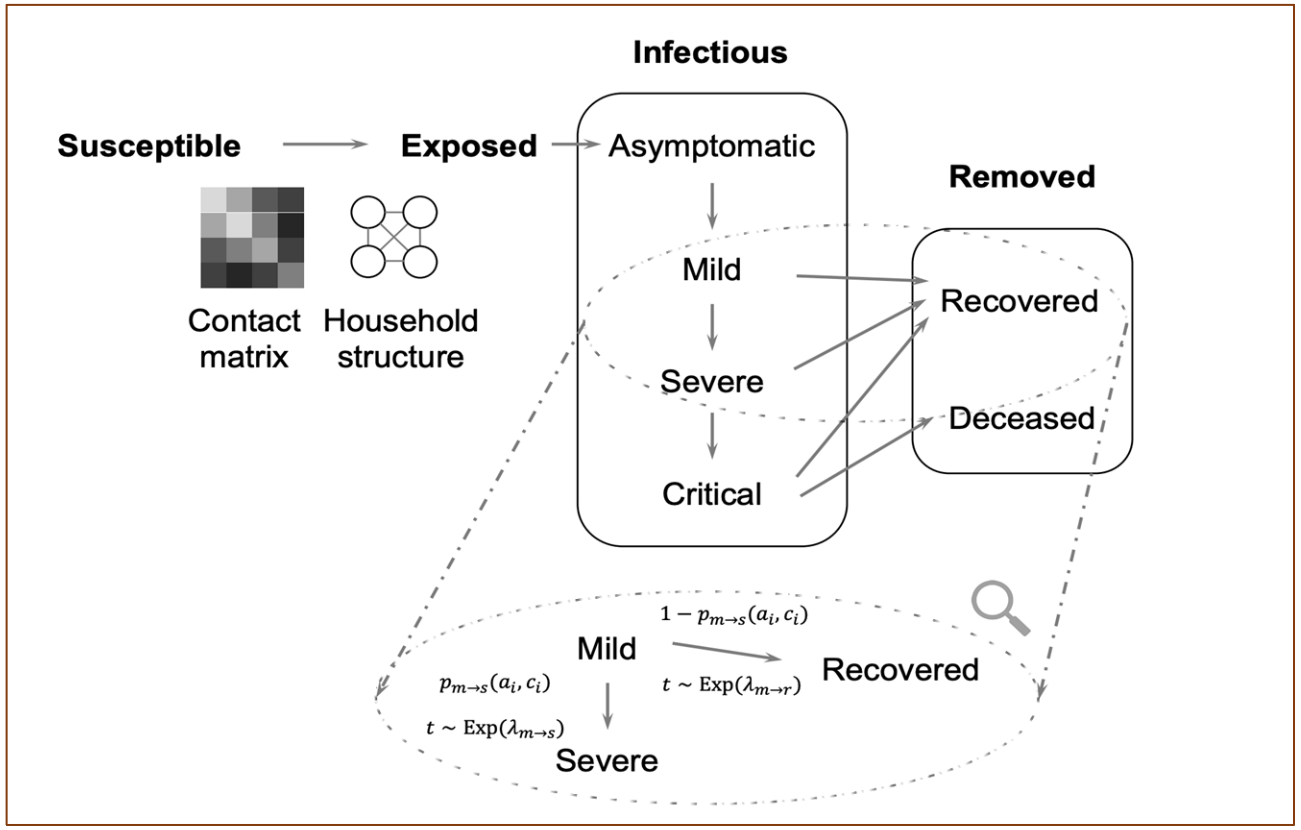

Agent-based model:

- *Families*
- *Co-morbidities*
- *Age*
- *Testing*
- *Contact tracing*

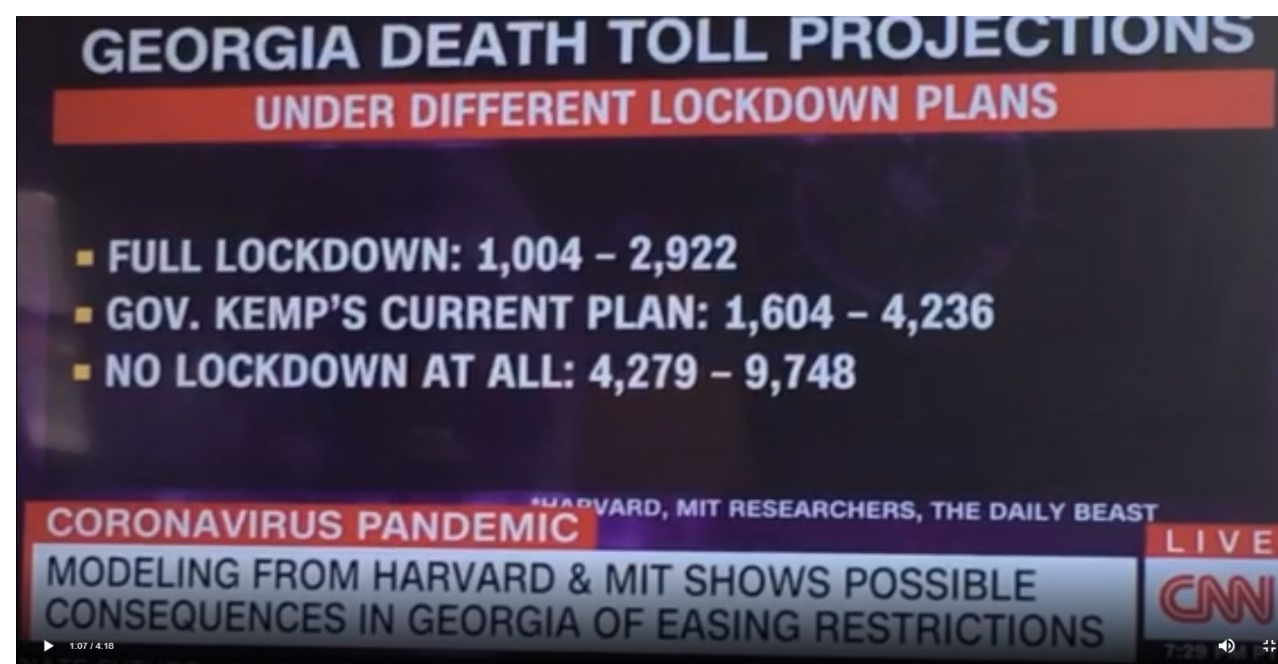

COVID Testing Policy: Accuracy vs Ease
(Under submission)


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
  - RT-LAMP: **10^5**/mL, \$5-30
  - Antigen strip ("Less sensitive"): **10^6**/mL, \$3-5

Test sensitivity is secondary to turnaround time & frequency for COVID-19 surveillance

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

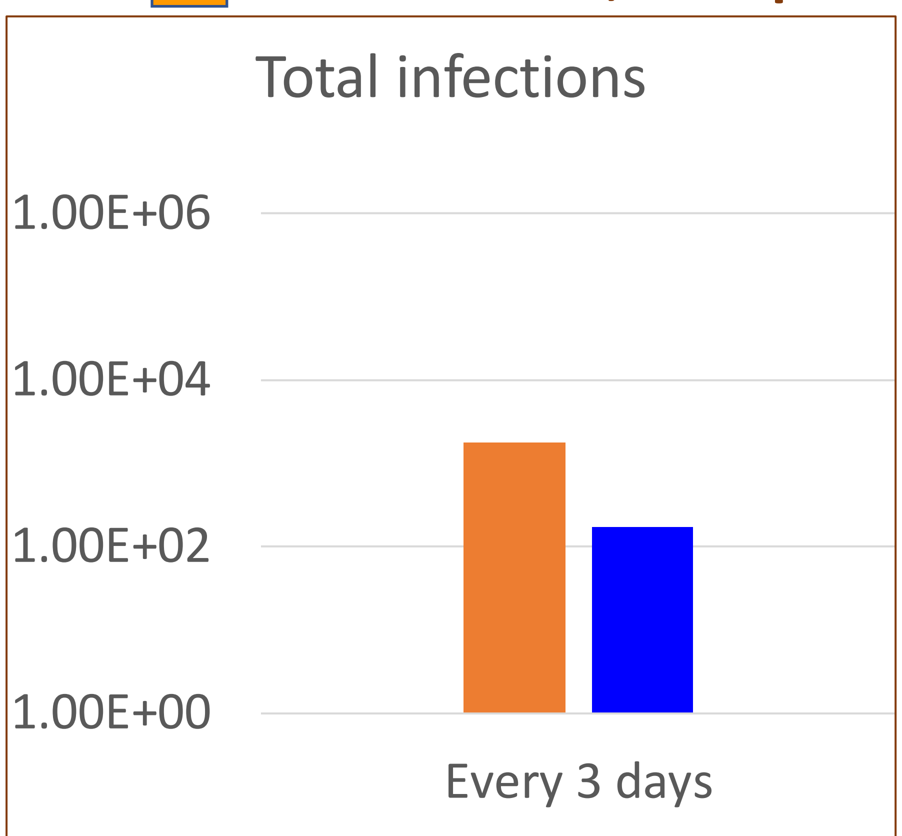

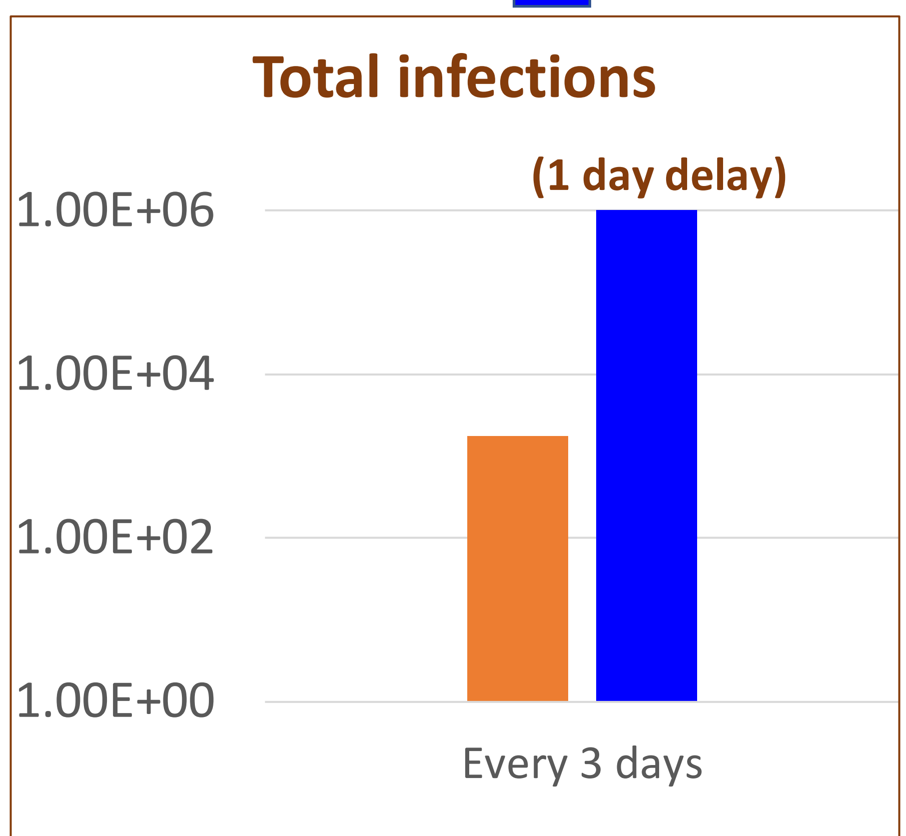

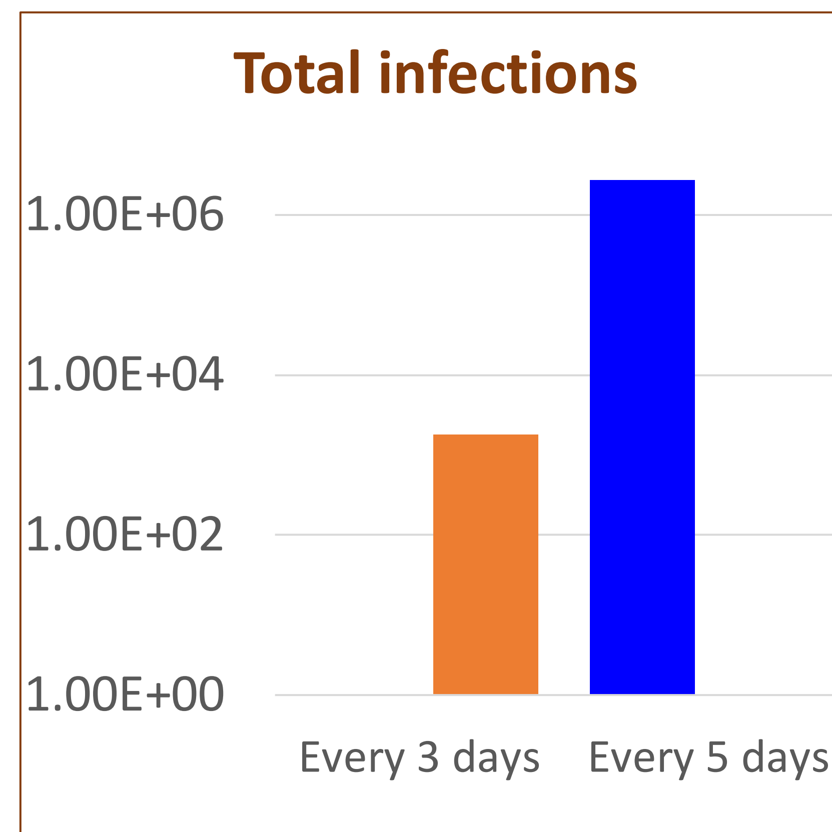

Challenges from field deployments

Influence maximization for health information dissemination

- Performance guarantee? Combined unknown network, robust dynamic policy
- Influence model: Independent cascade vs Activation jump vs…
- Fairness in influence maximization

Restless bandits for health adherence monitoring & intervention planning

- Efficient algorithms under uncertainty of patient observations
- Unknown, evolving patient adherence behavior
- Enable community health workers to interject in intervention planning

Outline

Public Health

Conservation

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

Poaching of Wildlife in Uganda
Limited Intervention (Ranger) Resources to Protect Forests

Snare or Trap

Wire snares

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

Fang

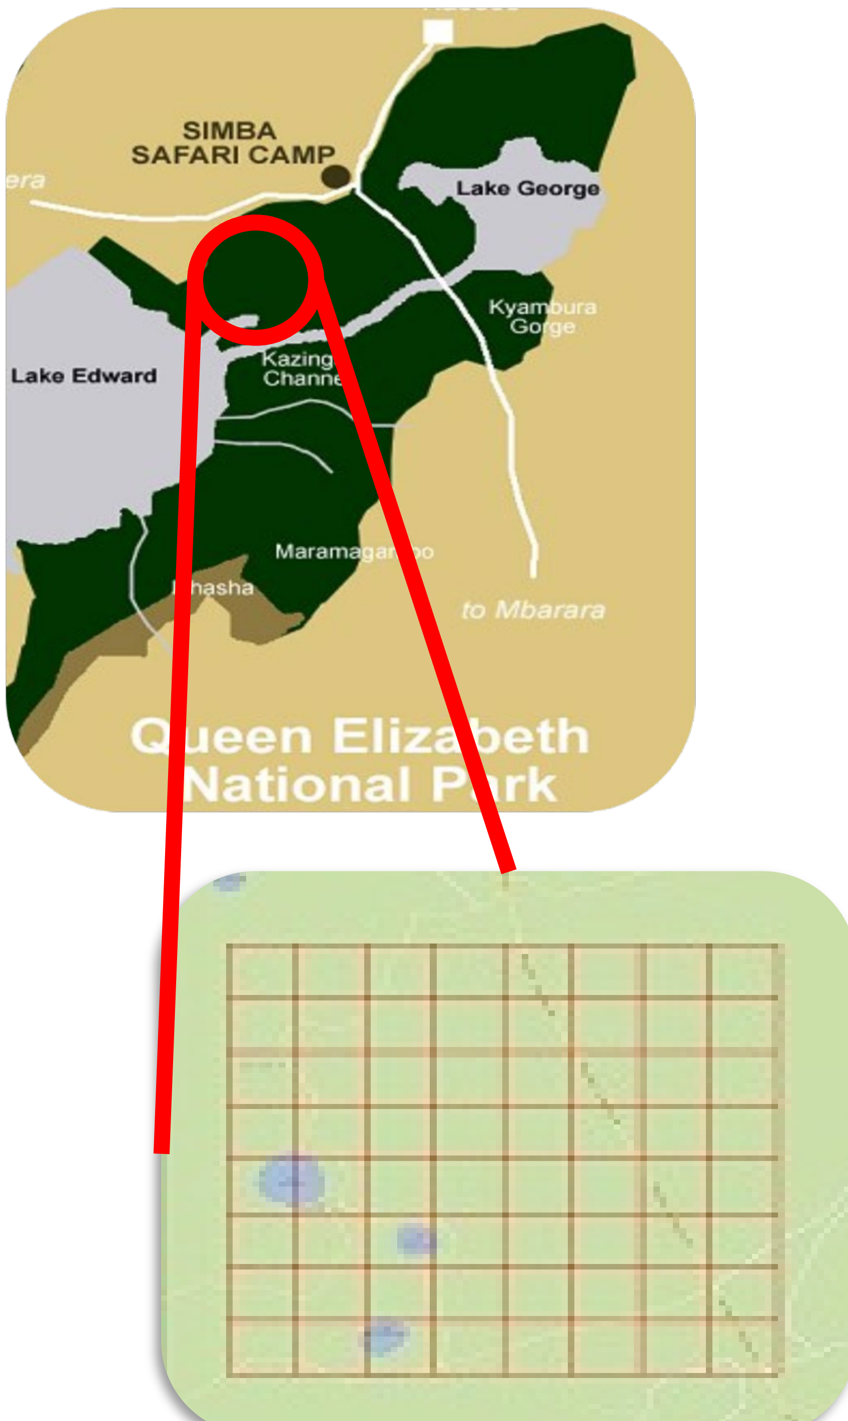

- ***Stackelberg security games (SSG)***
- ***With boundedly rational poachers***
- ***Learn adversary response model at targets "i"***

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

Learning Adversary Response Model: Uncertainty in Observations

Nguyen Gholami

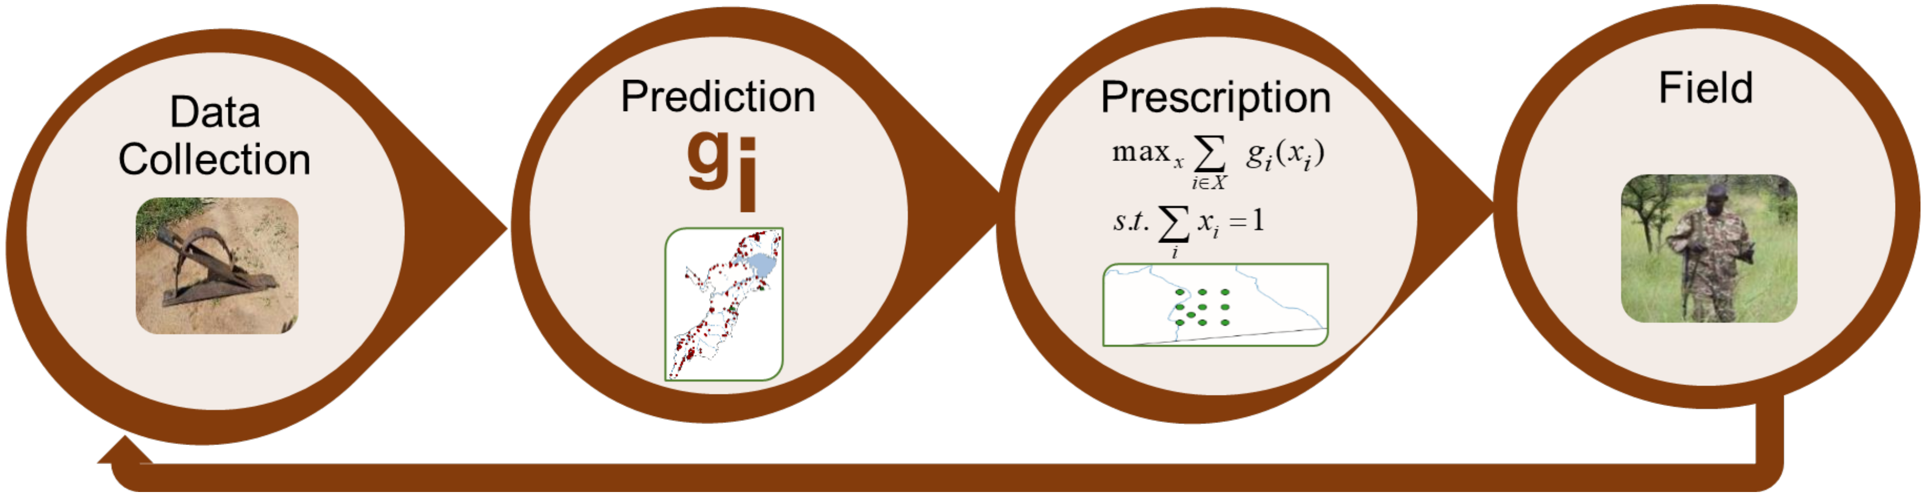

Ranger patrol

Animal density

Distance to rivers / roads / villages

Probability of snare Per 1 KM Grid Square

Area habitat

Area slope

…

Record: No Attack (NEG)

Record: Attack (POS)

Adversary Response Modeling
Imperfect Observation Ensemble Model

Gholami

Training: Filtered Datasets

Predict: Ensemble of Classifiers

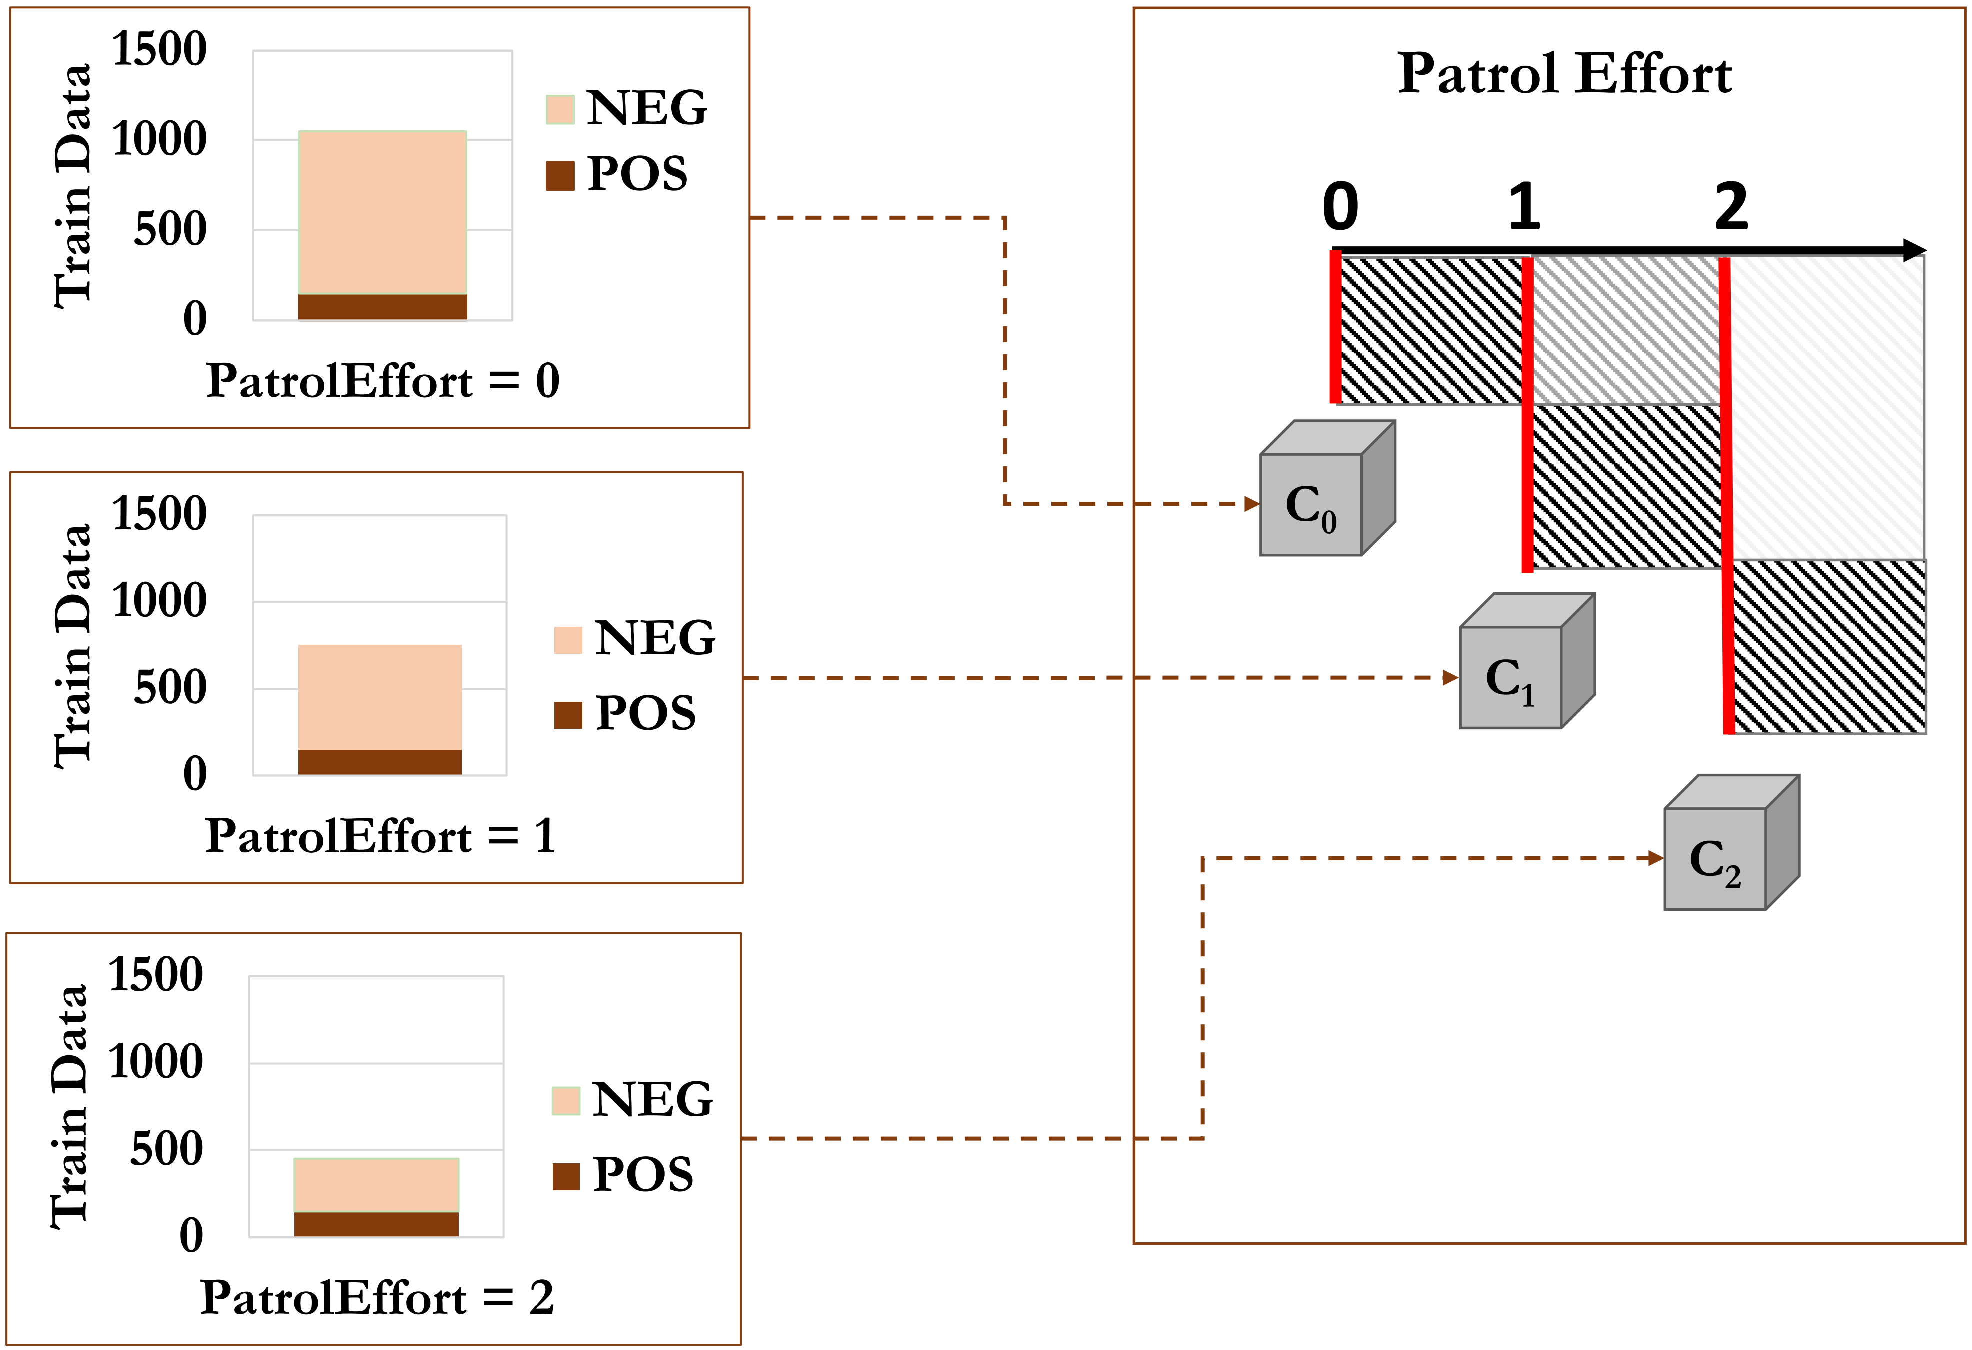

PAWS: Protection Assistant for Wildlife Security
Adversary Model in the Lab

Gholami

Poacher Behavior Prediction

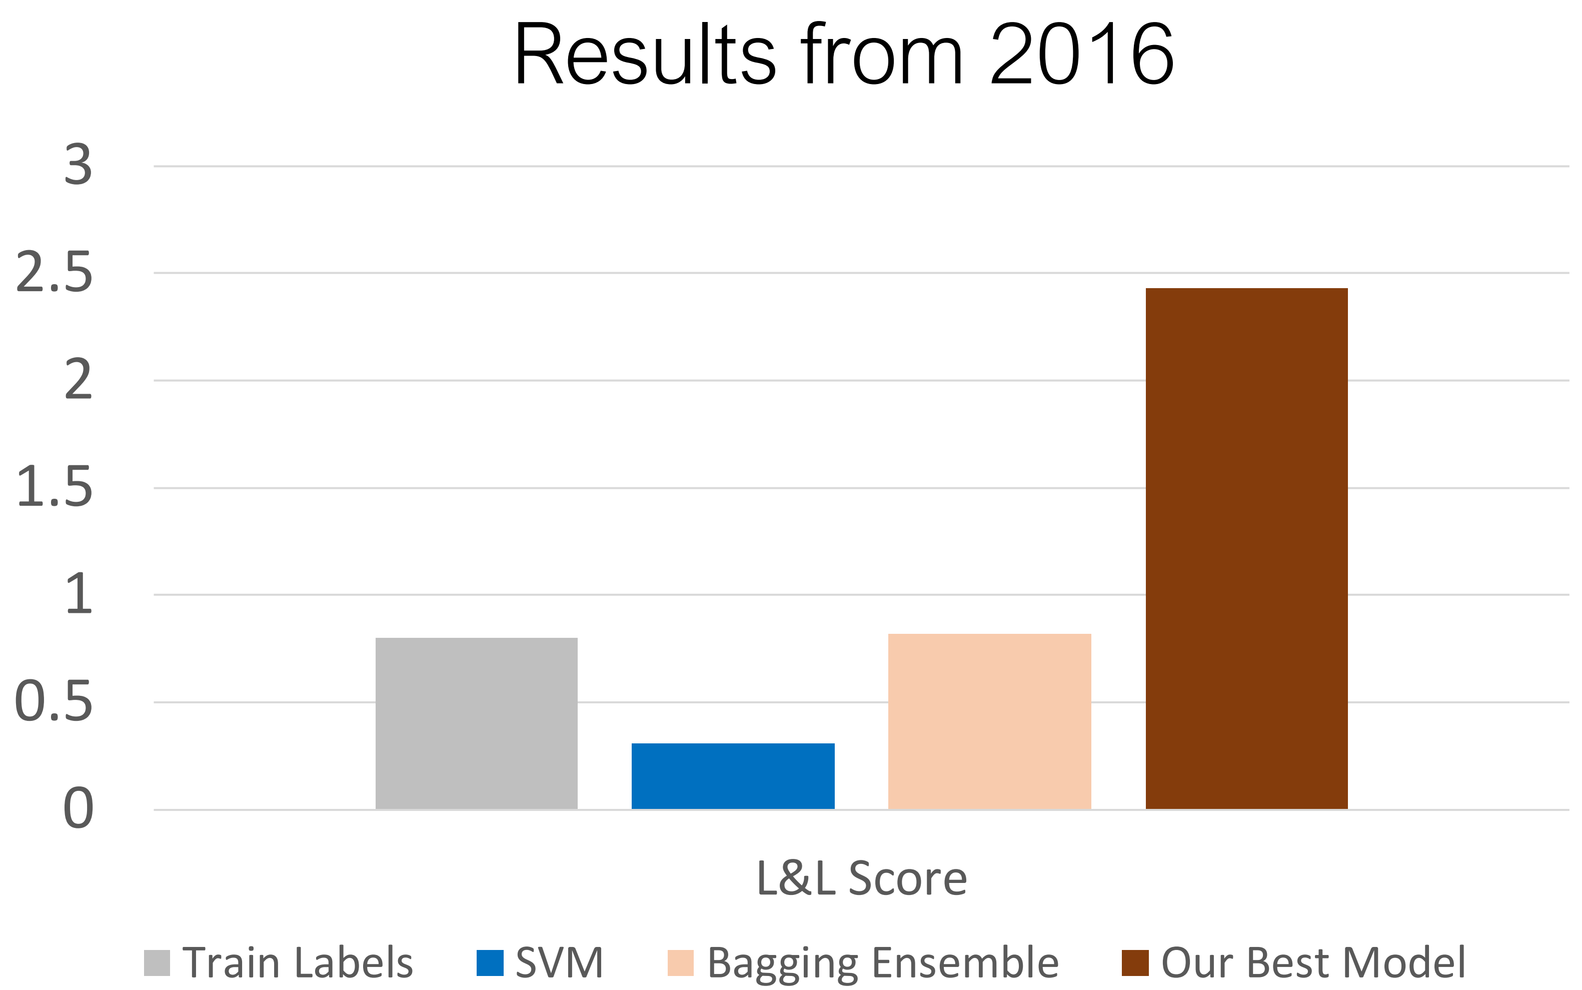

PAWS: First Pilot in the Field

(AAMAS 2017)

Ford

Gholami

- Two 9-sq.km areas, infrequent patrols

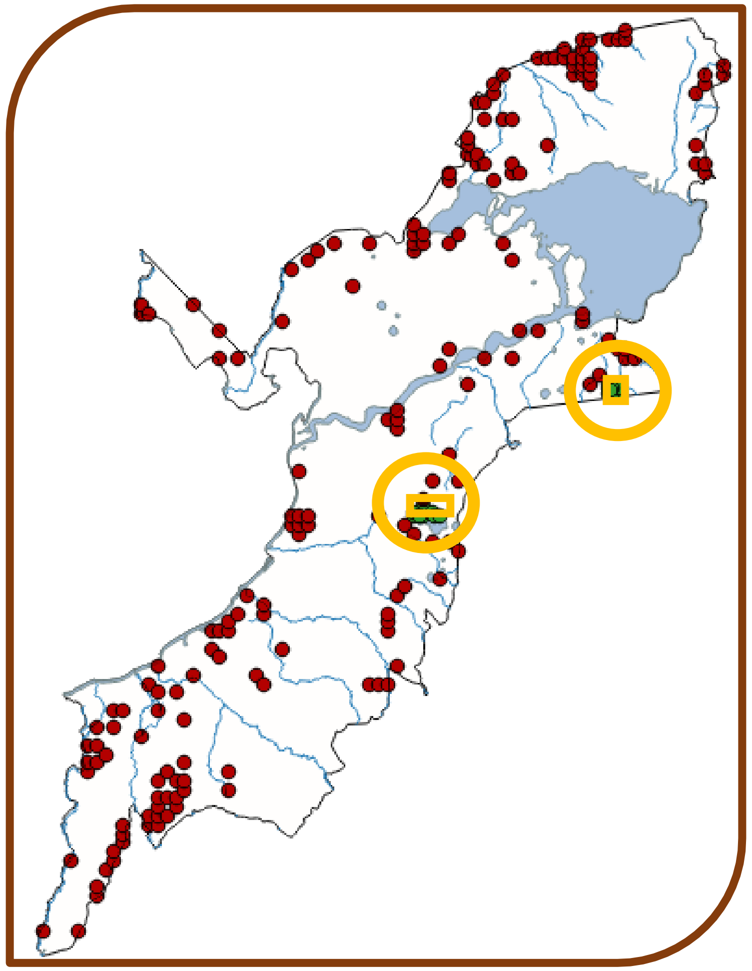

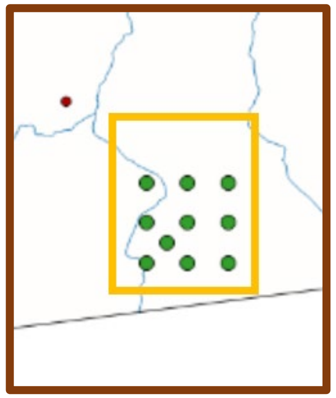

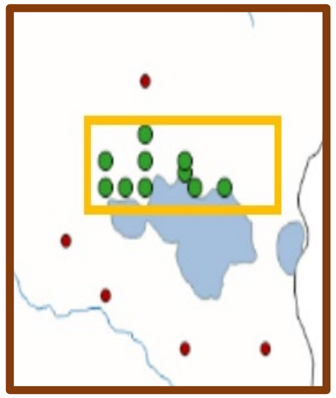

- Poached elephant
- 1 elephant snare roll
- 10 Antelope snares

PAWS Predicted High vs Low Risk Areas: 2 National Parks, 24 areas each, 6 months

(ECML PKDD 2017, ICDE 2020)


Ford


Gholami

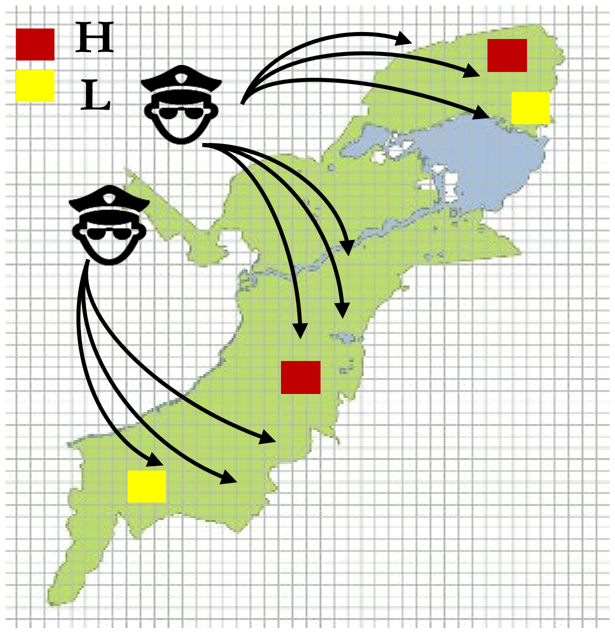

Queen Elizabeth National Park

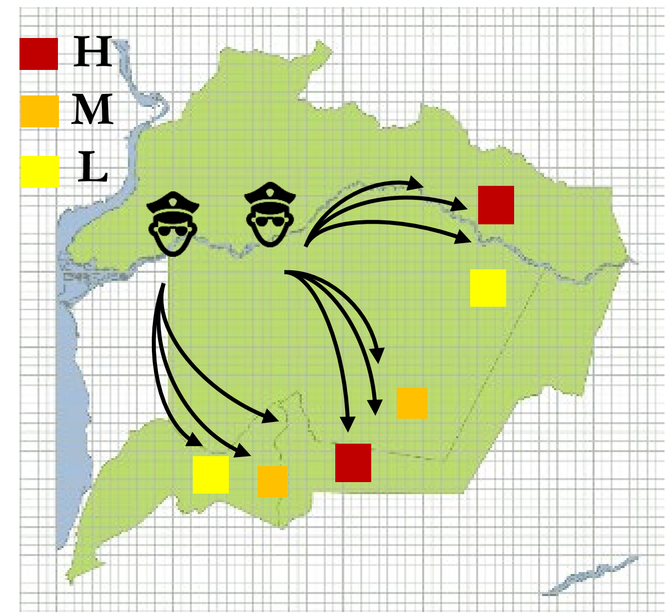

Murchison Falls National Park

Snares per patrolled sq. KM

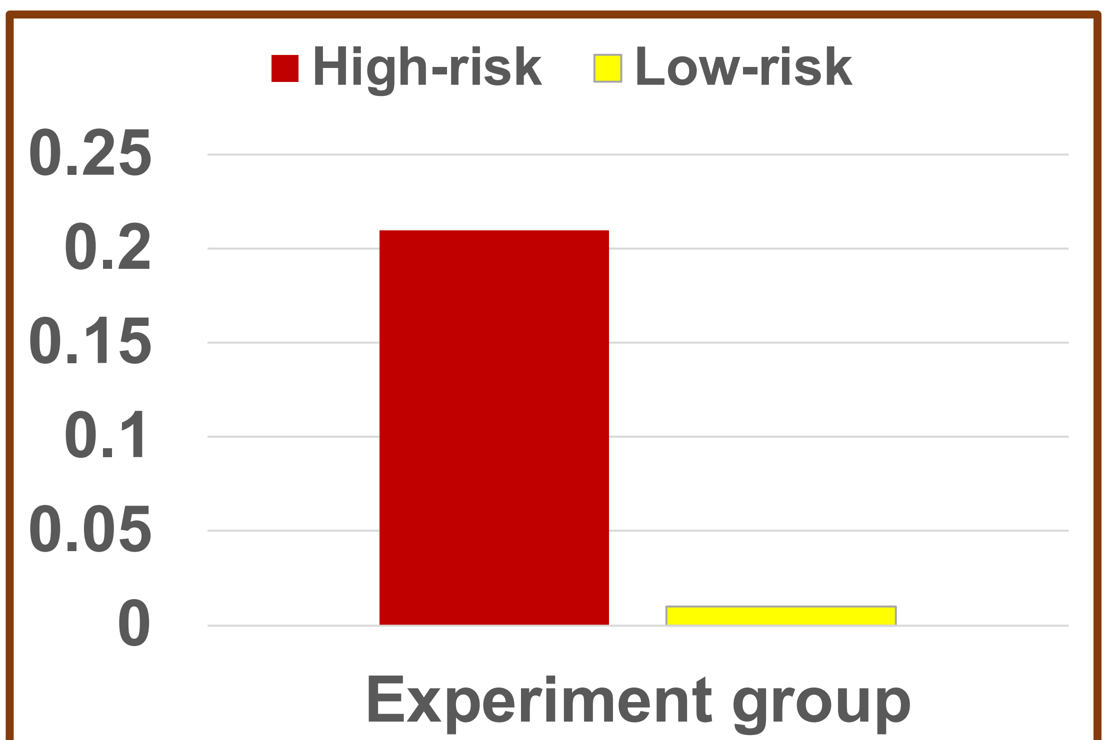

Snares per patrolled sq. KM

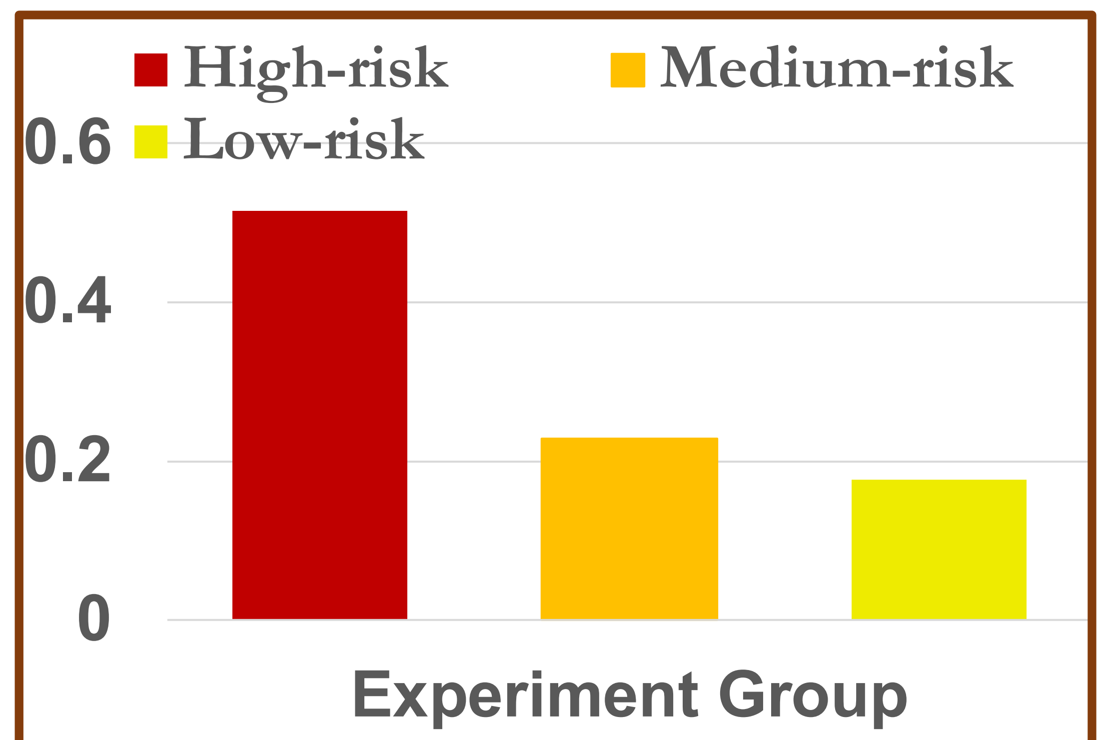

PAWS Real-world Deployment Cambodia: Srepok Wildlife Sanctuary

(ICDE 2020)

Xu

"@Milind: I am Super excited with the results. Let's get this going on other countries too this year."

Rohit Singh, WWF (2019)

Snares per patrolled sq. KM

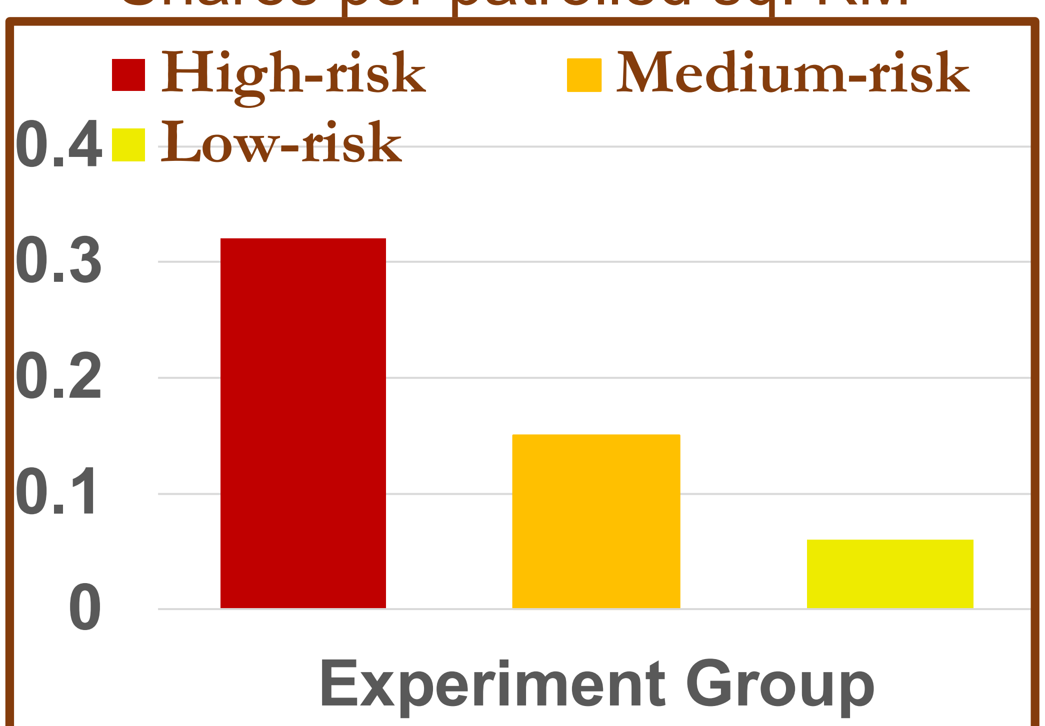

- *521 snares/month our tests*

vs

- *101 snares/month 2018*

Is Adversary observing & Reacting to Patrols?

Evidence from the Field Justifies Stackelberg Assumption

Xu

Perrault

- Logistic regression model

$$a_i + \gamma \cdot \texttt{past\_effort} + \beta \cdot \texttt{current\_effort}$$

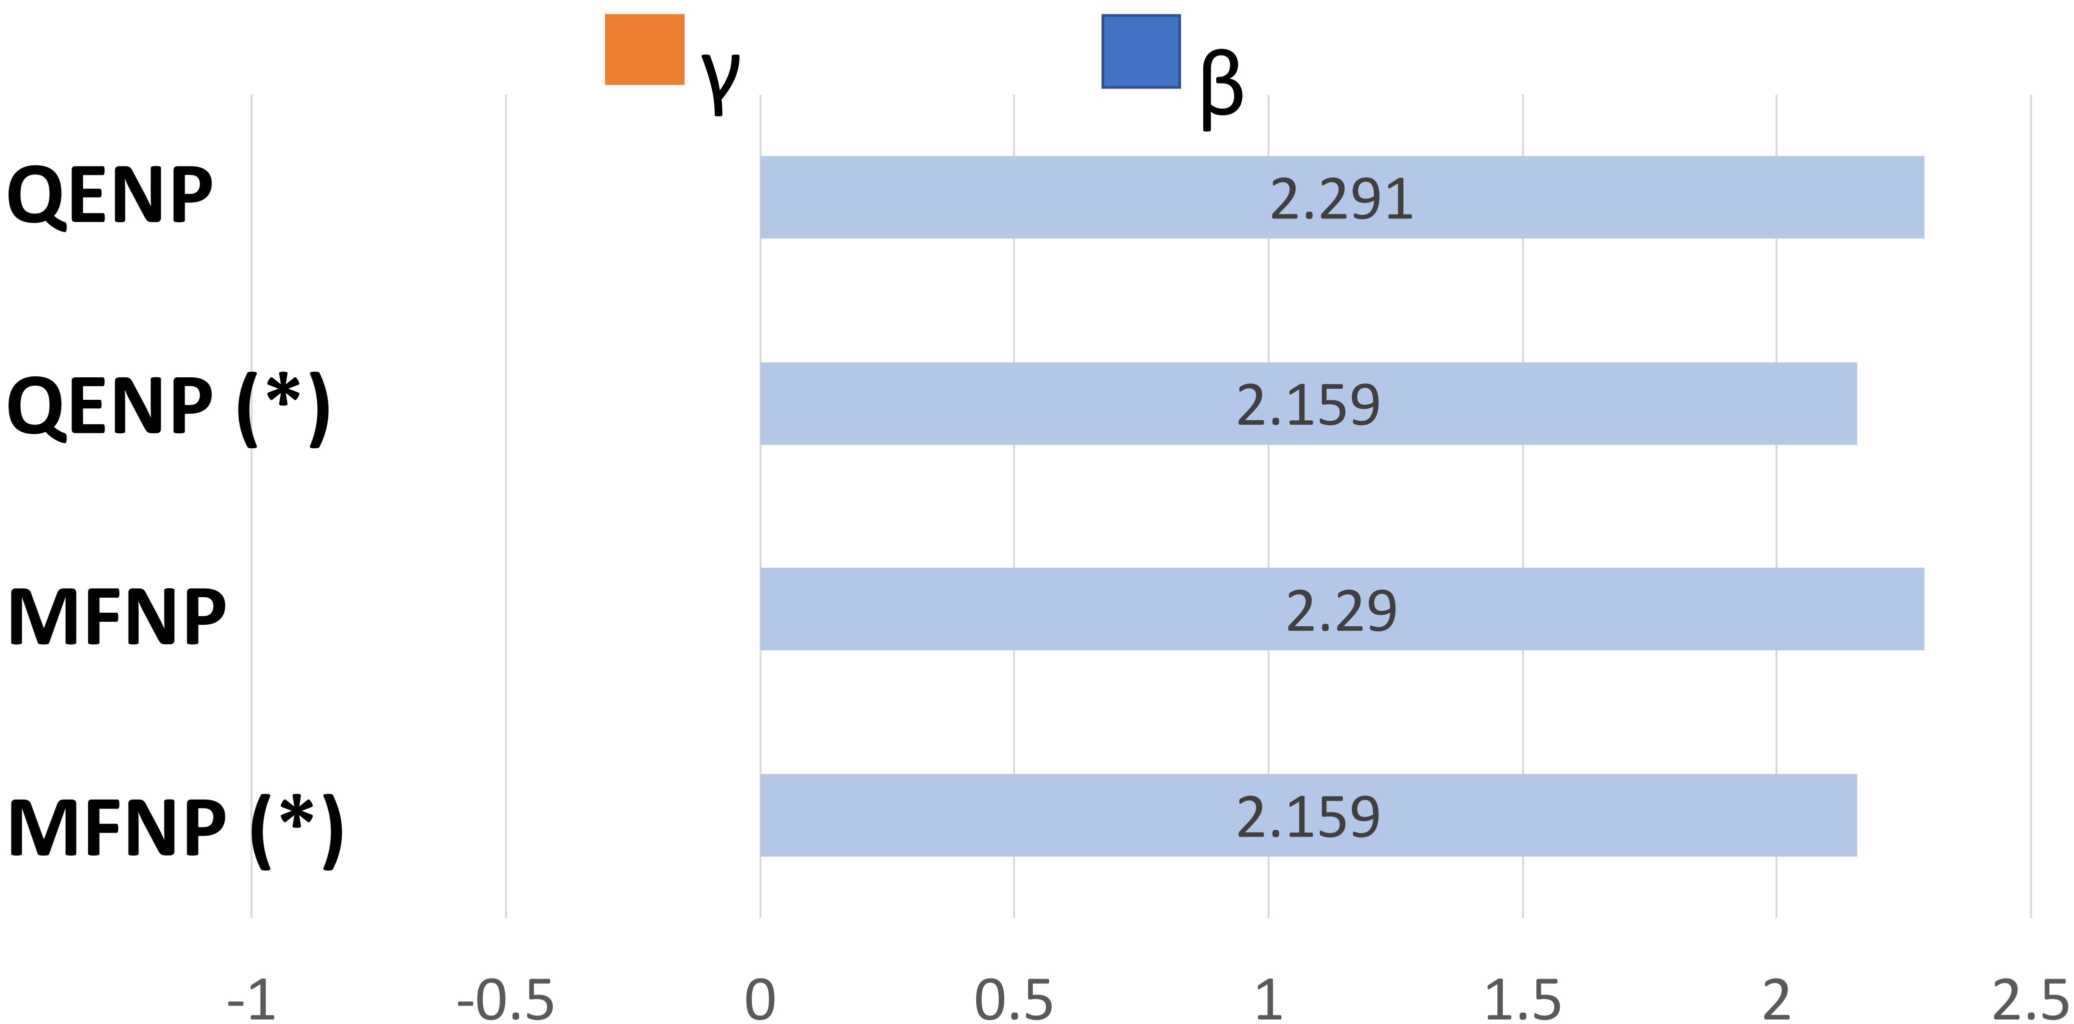

Demonstrating Deterrence:
Evidence from the Field Justifies Stackelberg Assumption

Xu

Perrault

- Is adversary observing & reacting to patrols? Logistic regression model

$$a_i + \gamma \cdot \texttt{past\_effort} + \beta \cdot \texttt{current\_effort}$$

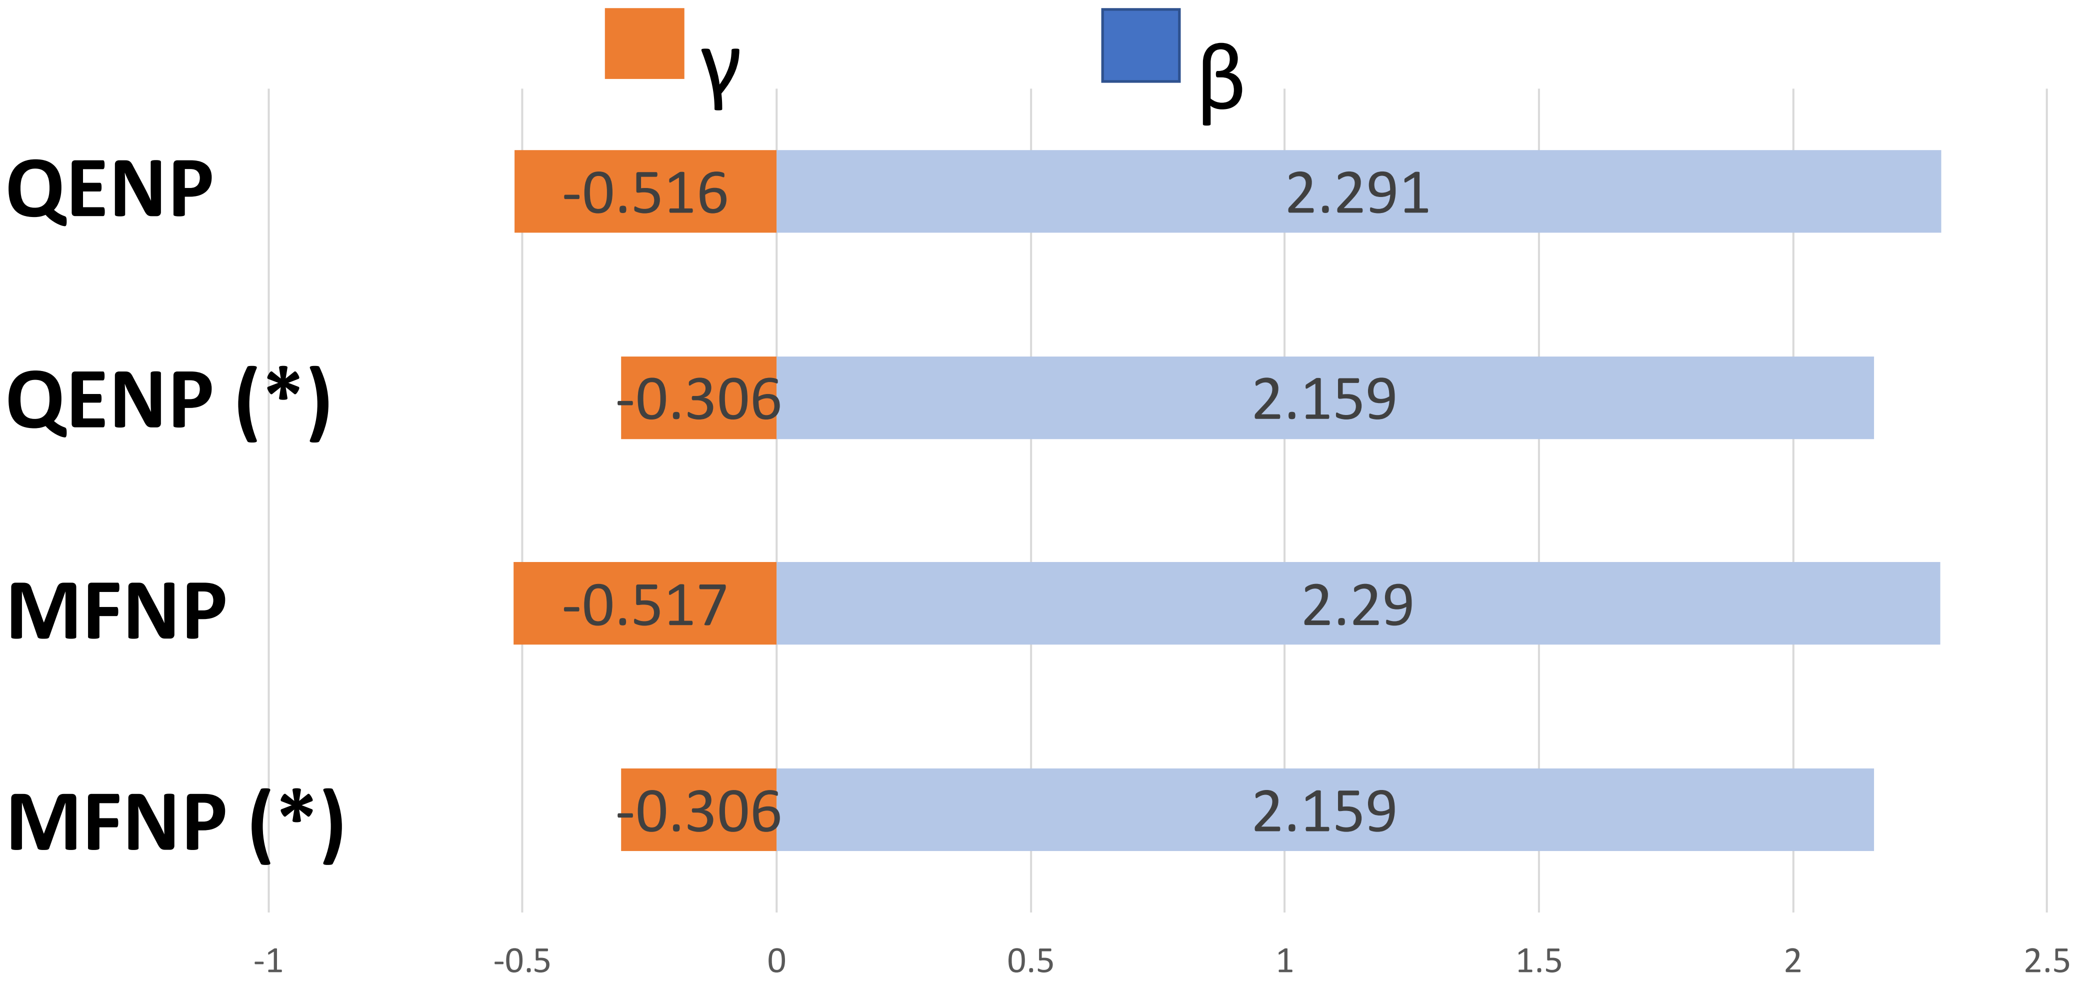

Solving Security Game with Learned Adversary Model

- Solving Stackelberg security game with learned adversary model
  - → Difficulty of generating routes: many constraints on patrols

- *Insufficient data: Errors in planning patrols on targets*
- *Game focused learning to maximize decision quality*
  - → Maximizing learning accuracy ≠ Maximizing decision quality

Previous Stage-by-Stage Method:
Make Prediction as Accurate as Possible Then Plan

Perrault

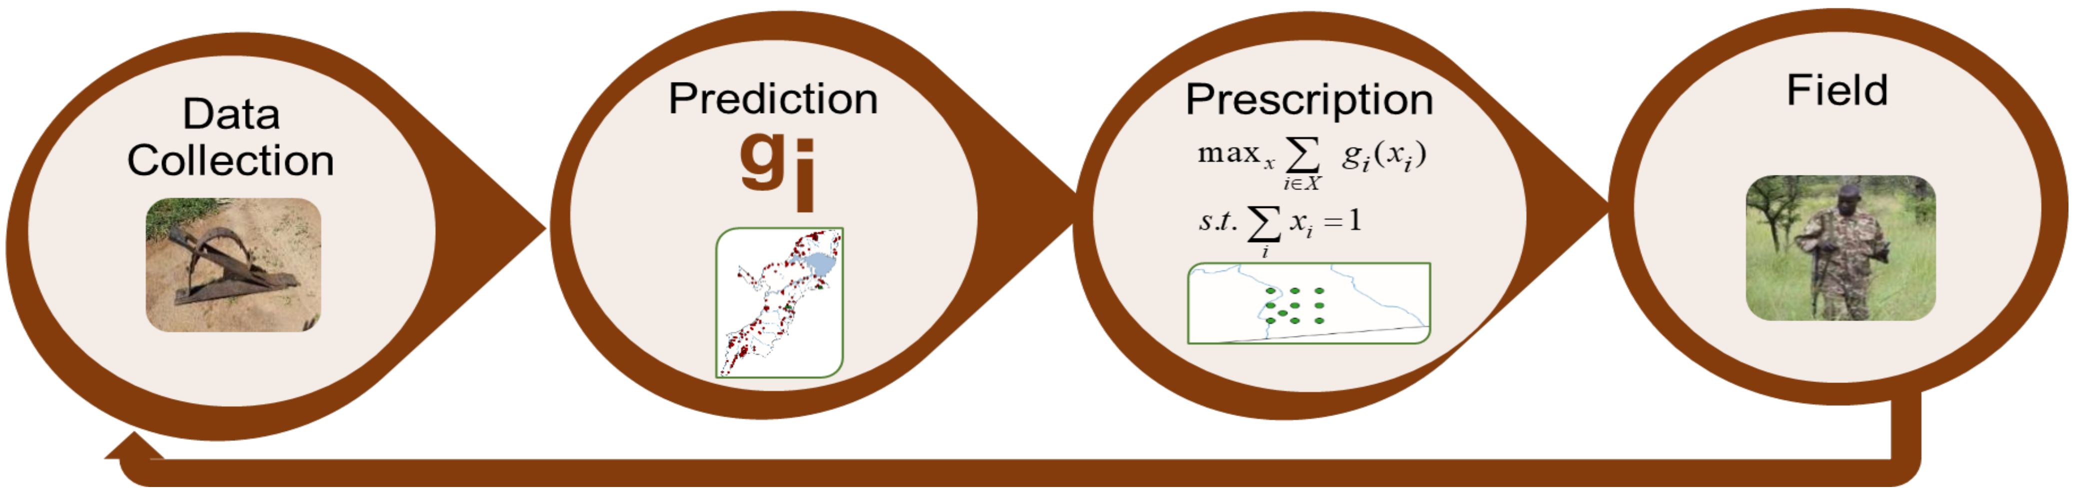

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

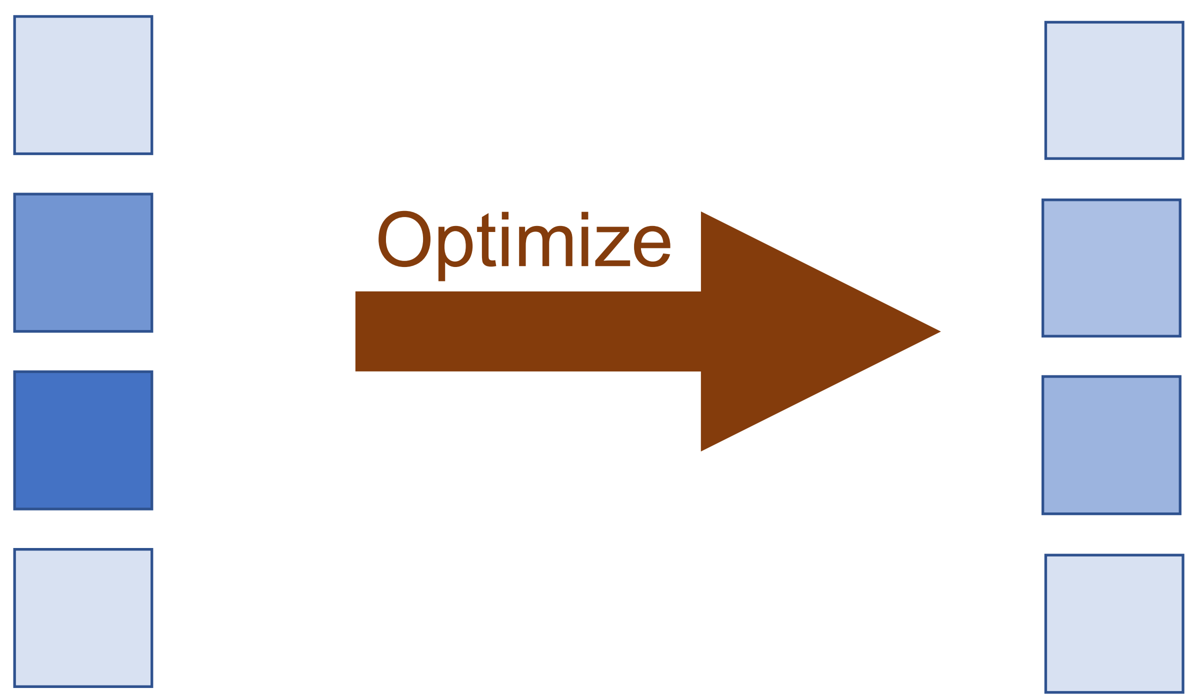

Minimize $\sum_{\{i \in T\}} q_{\text{empirical}} \log \hat{q}$

Stage by Stage Method:
Need to Focus on Important Targets

Perrault

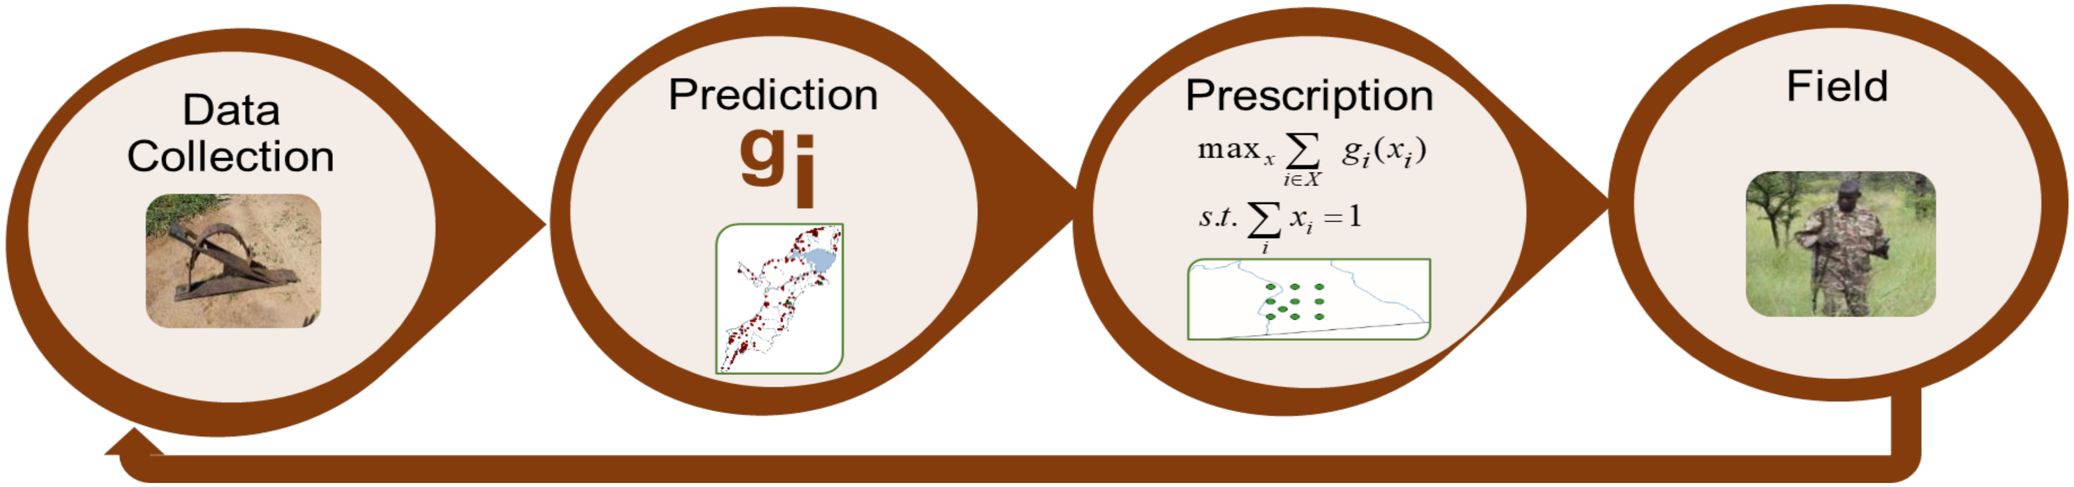

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

Two targets: Large effect Defender EU

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

Game-Focused Learning:
Need to Focus on Important Targets

Perrault

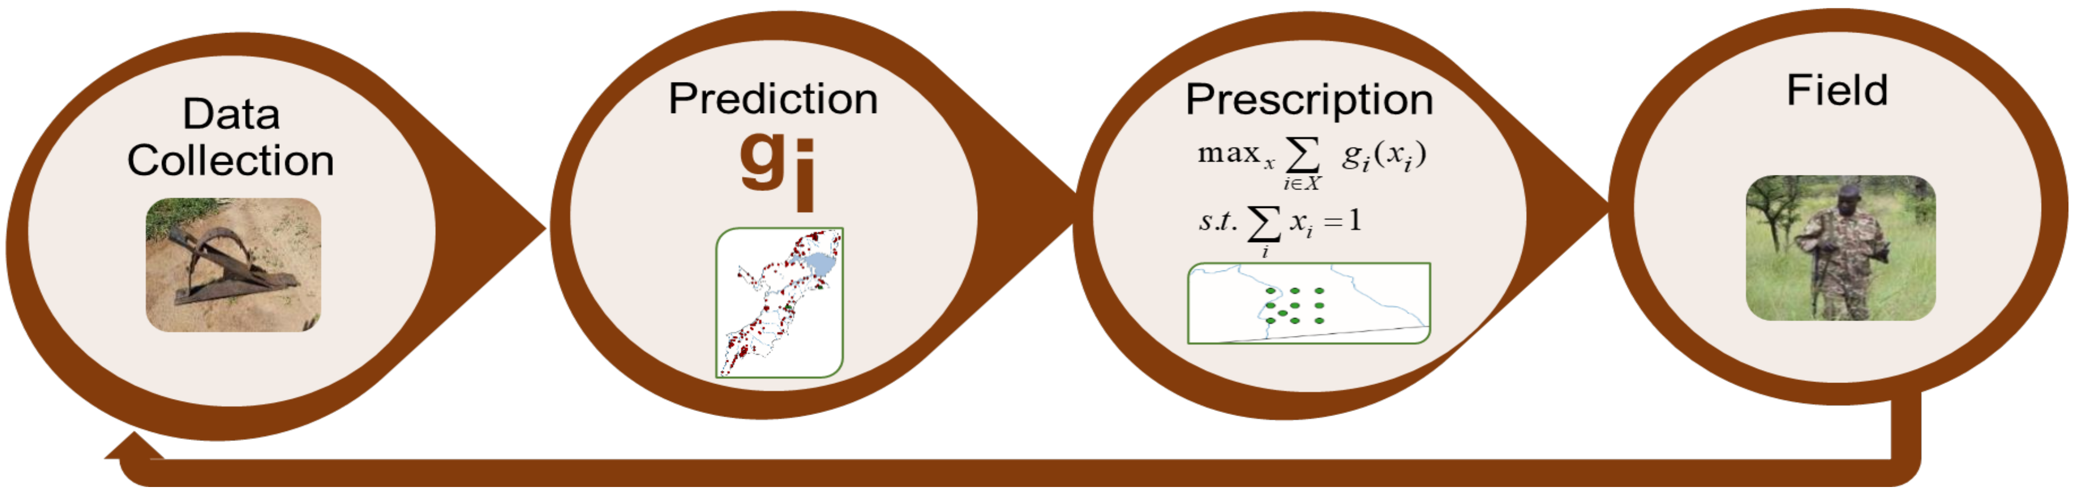

Maximize accuracy only of
Important targets

Plan patrol
Coverage

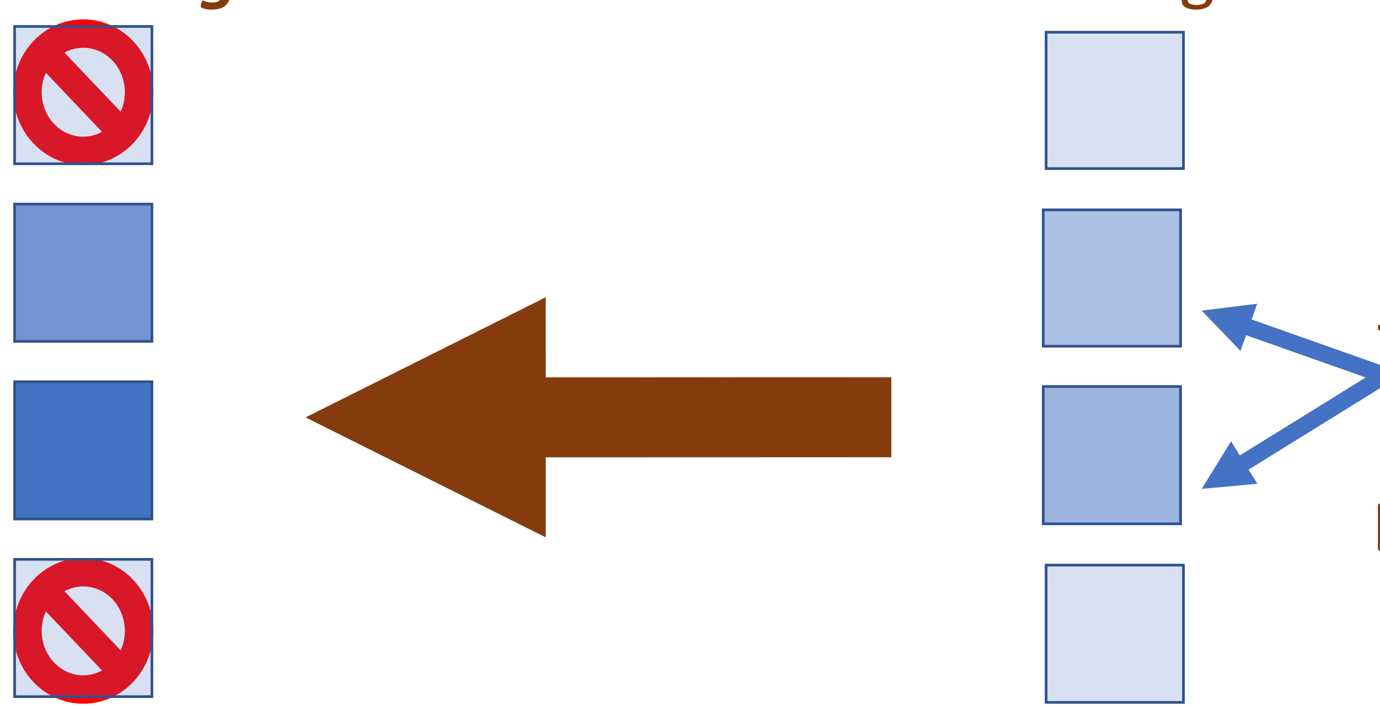

Two targets:
Large effect
Defender EU

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

Game-Focused Learning: End-to-End Method
Builds on Decision-focused Learning
(AAAI 2019, AAAI 2020)


Perrault


Wilder

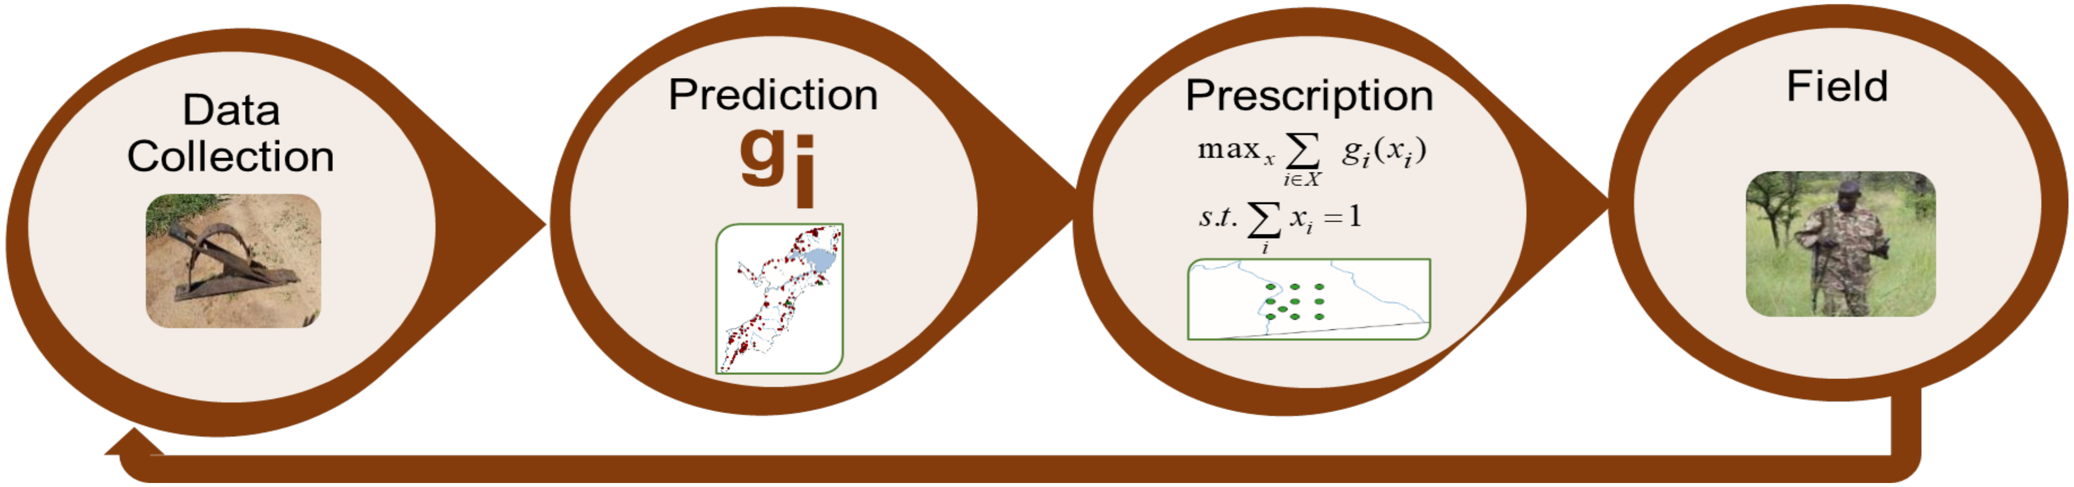

Maximize
defender expected *utility*

Plan patrol
Coverage

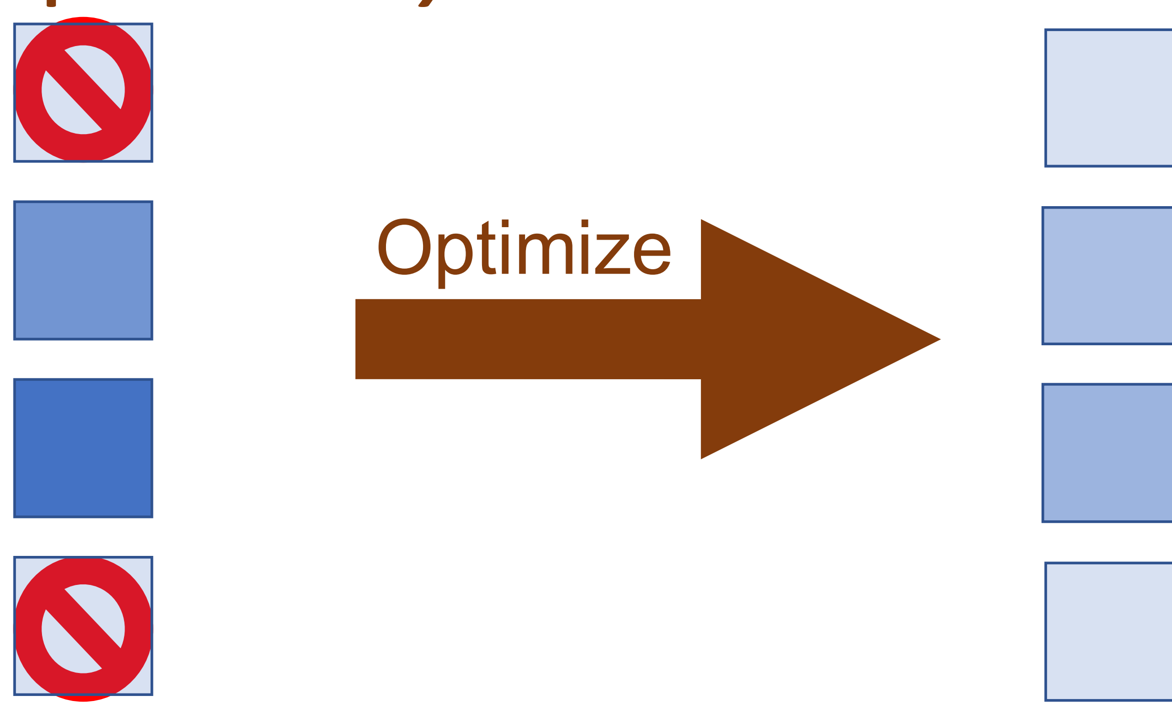

Maximize defender's expected utility $\sum \left(1 - p_i(\hat{q})\right) q_{\text{empirical}}$

Previous Two-Stage Method: Gradient Descent


Perrault


Wilder

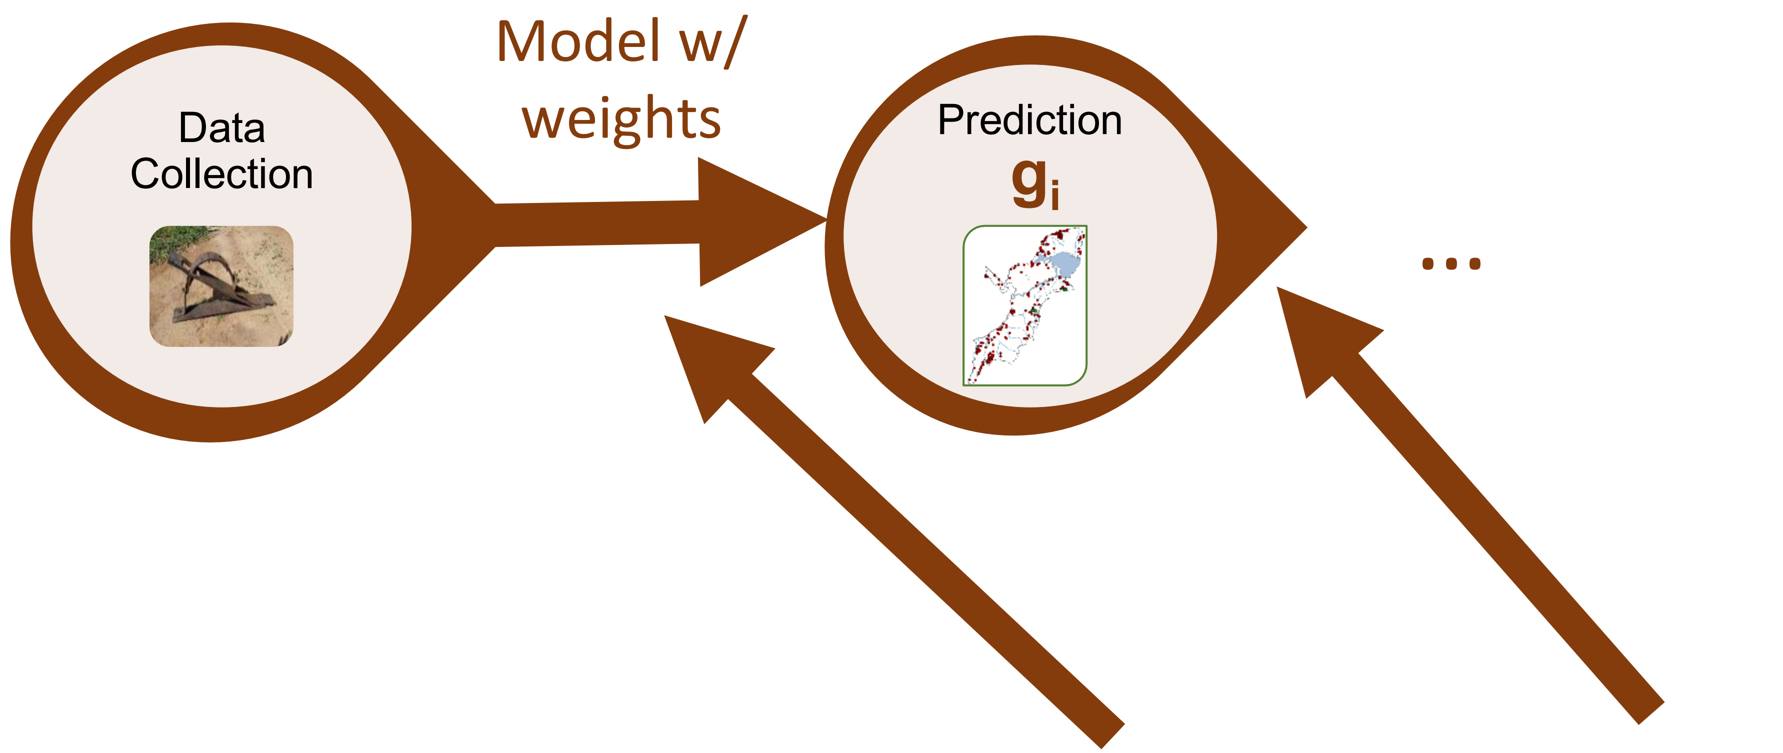

- Max accuracy gradient descent: $$\frac{\partial \text{accuracy}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{accuracy}}{\partial \text{prediction}}$$

Game-Focused Learning: End-to-End Method

Perrault

Wilder

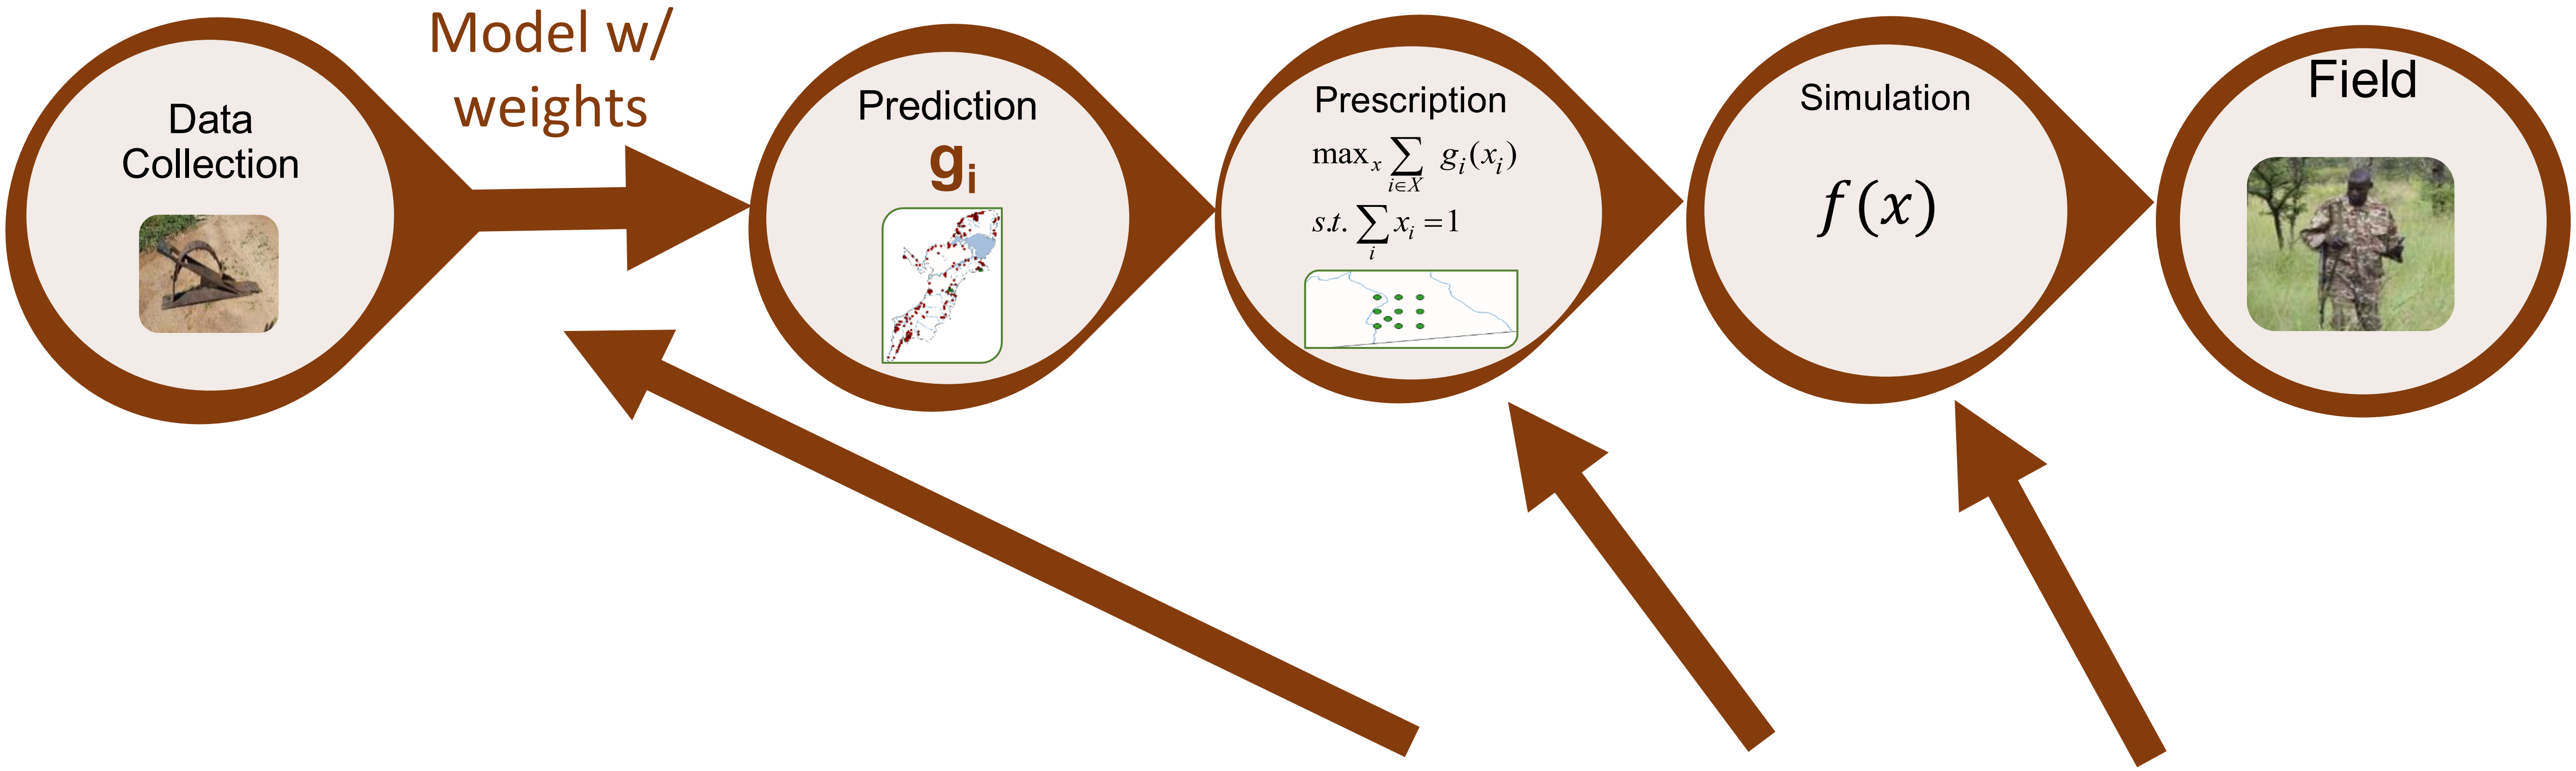

- Game-focused gradient descent: $$\frac{\partial \text{obj(decision)}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{decision}}{\partial \text{prediction}} \frac{\partial \text{obj(decision)}}{\partial \text{decision}}$$

Game-Focused Learning: Comparison to Two-Stage

Perrault

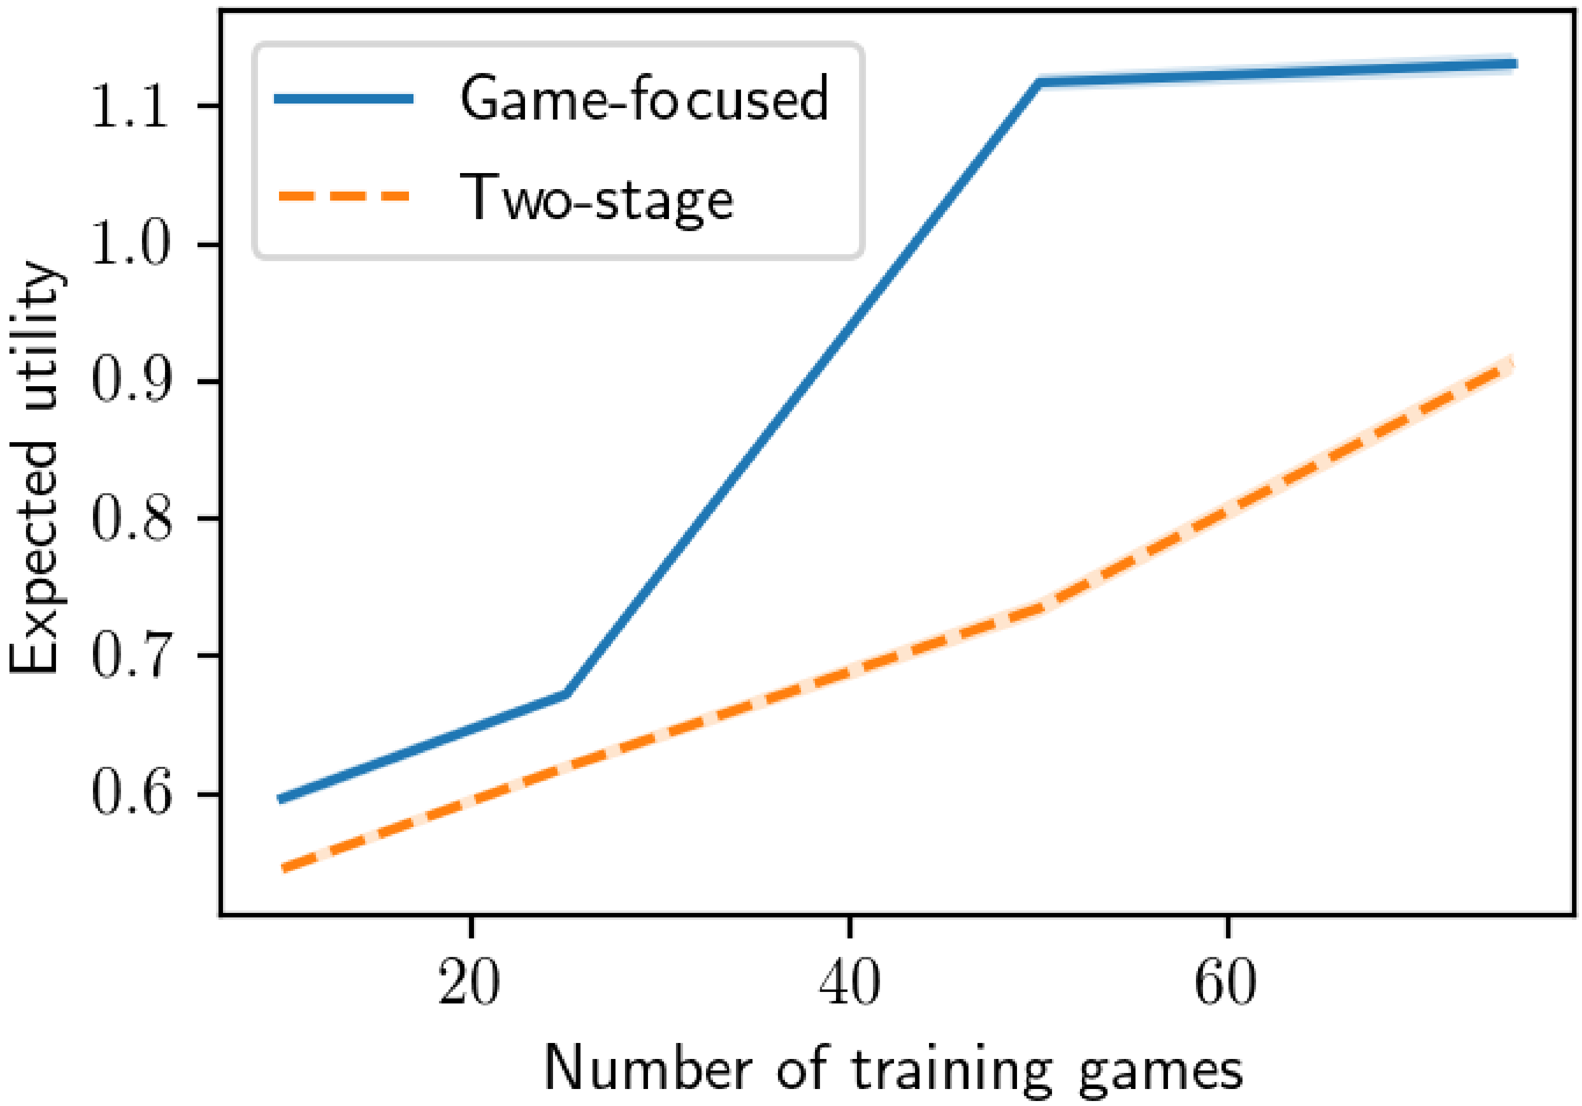

Green Security Games: Integrating Real-Time Information in the Pipeline

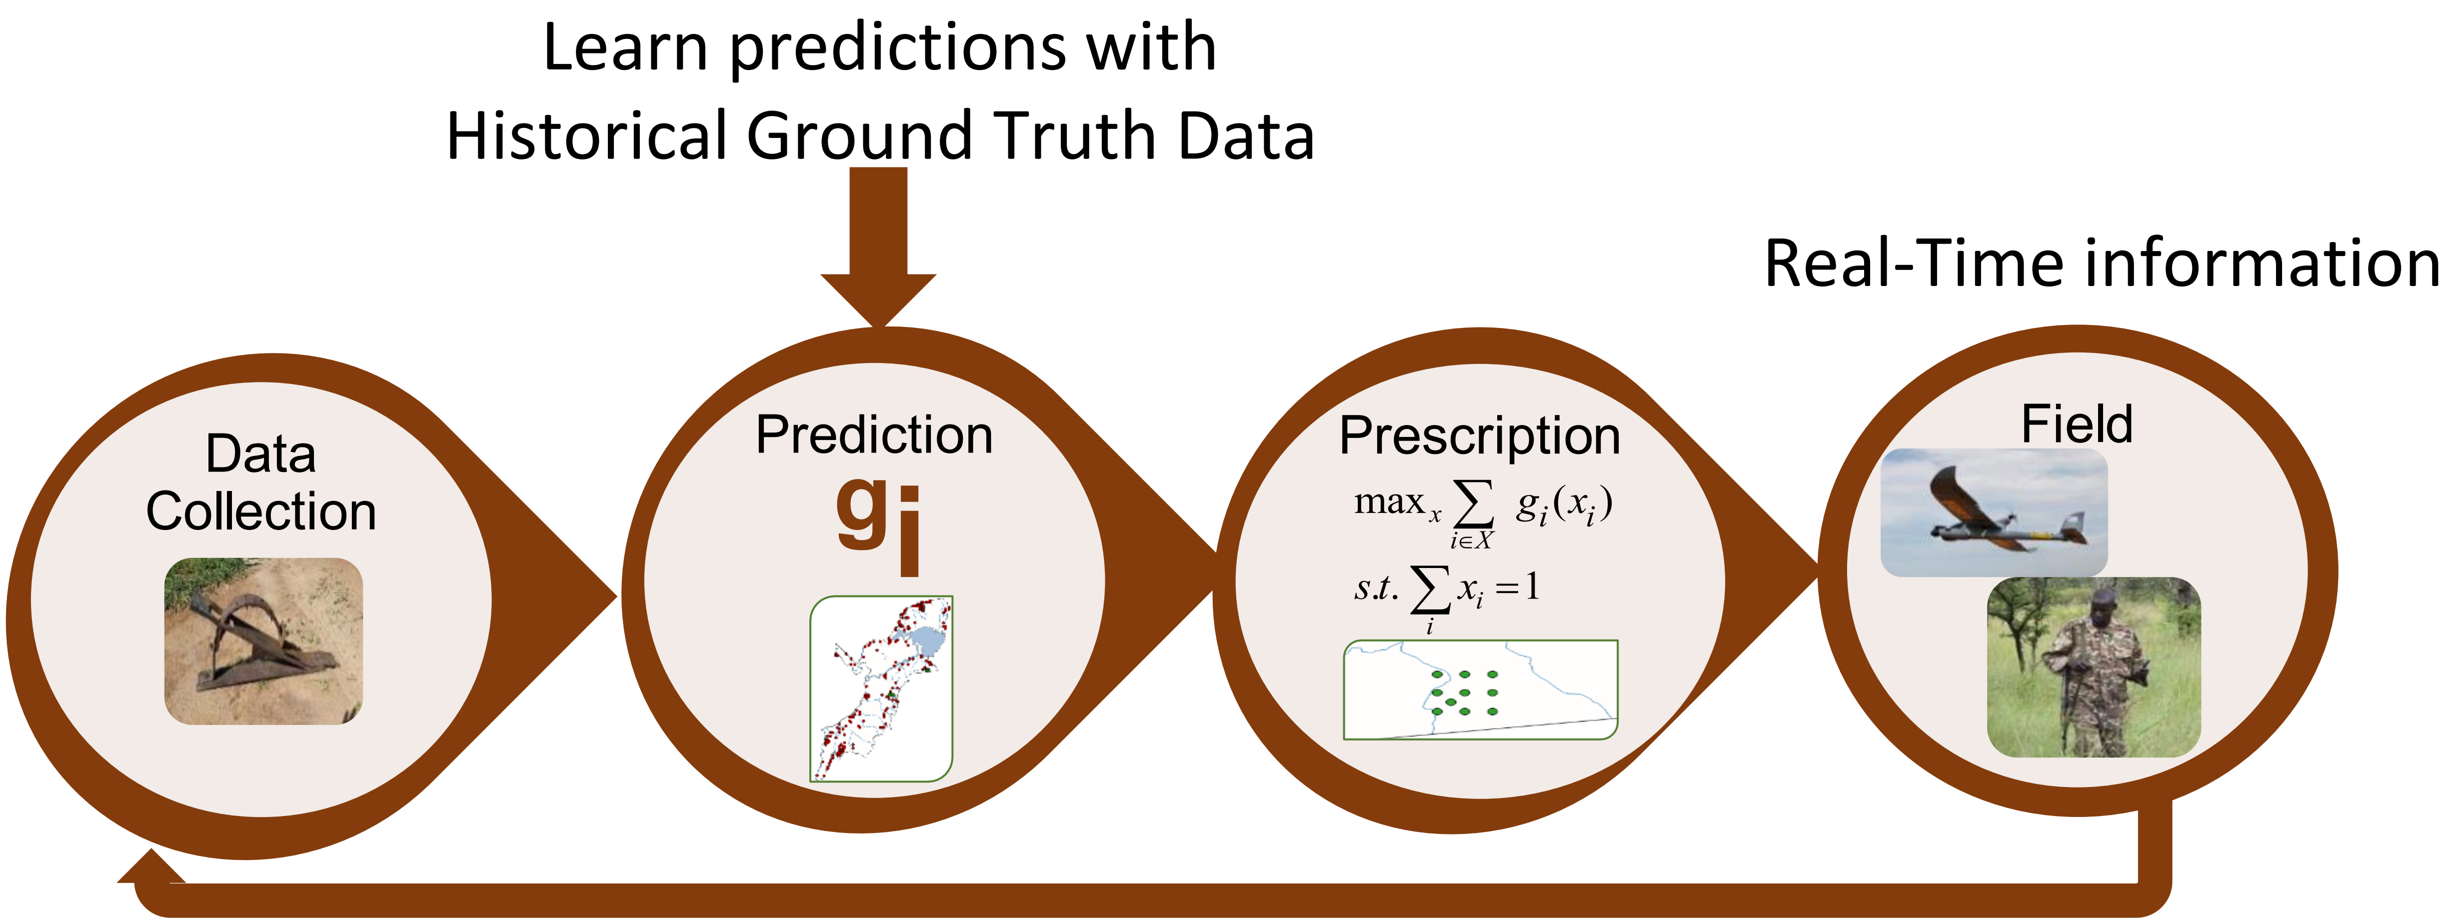

Green Security Games: Integrating Real-Time "SPOT" Information

(IAAI 2018)

Bondi

Air Shepherd
The Lindbergh Foundation

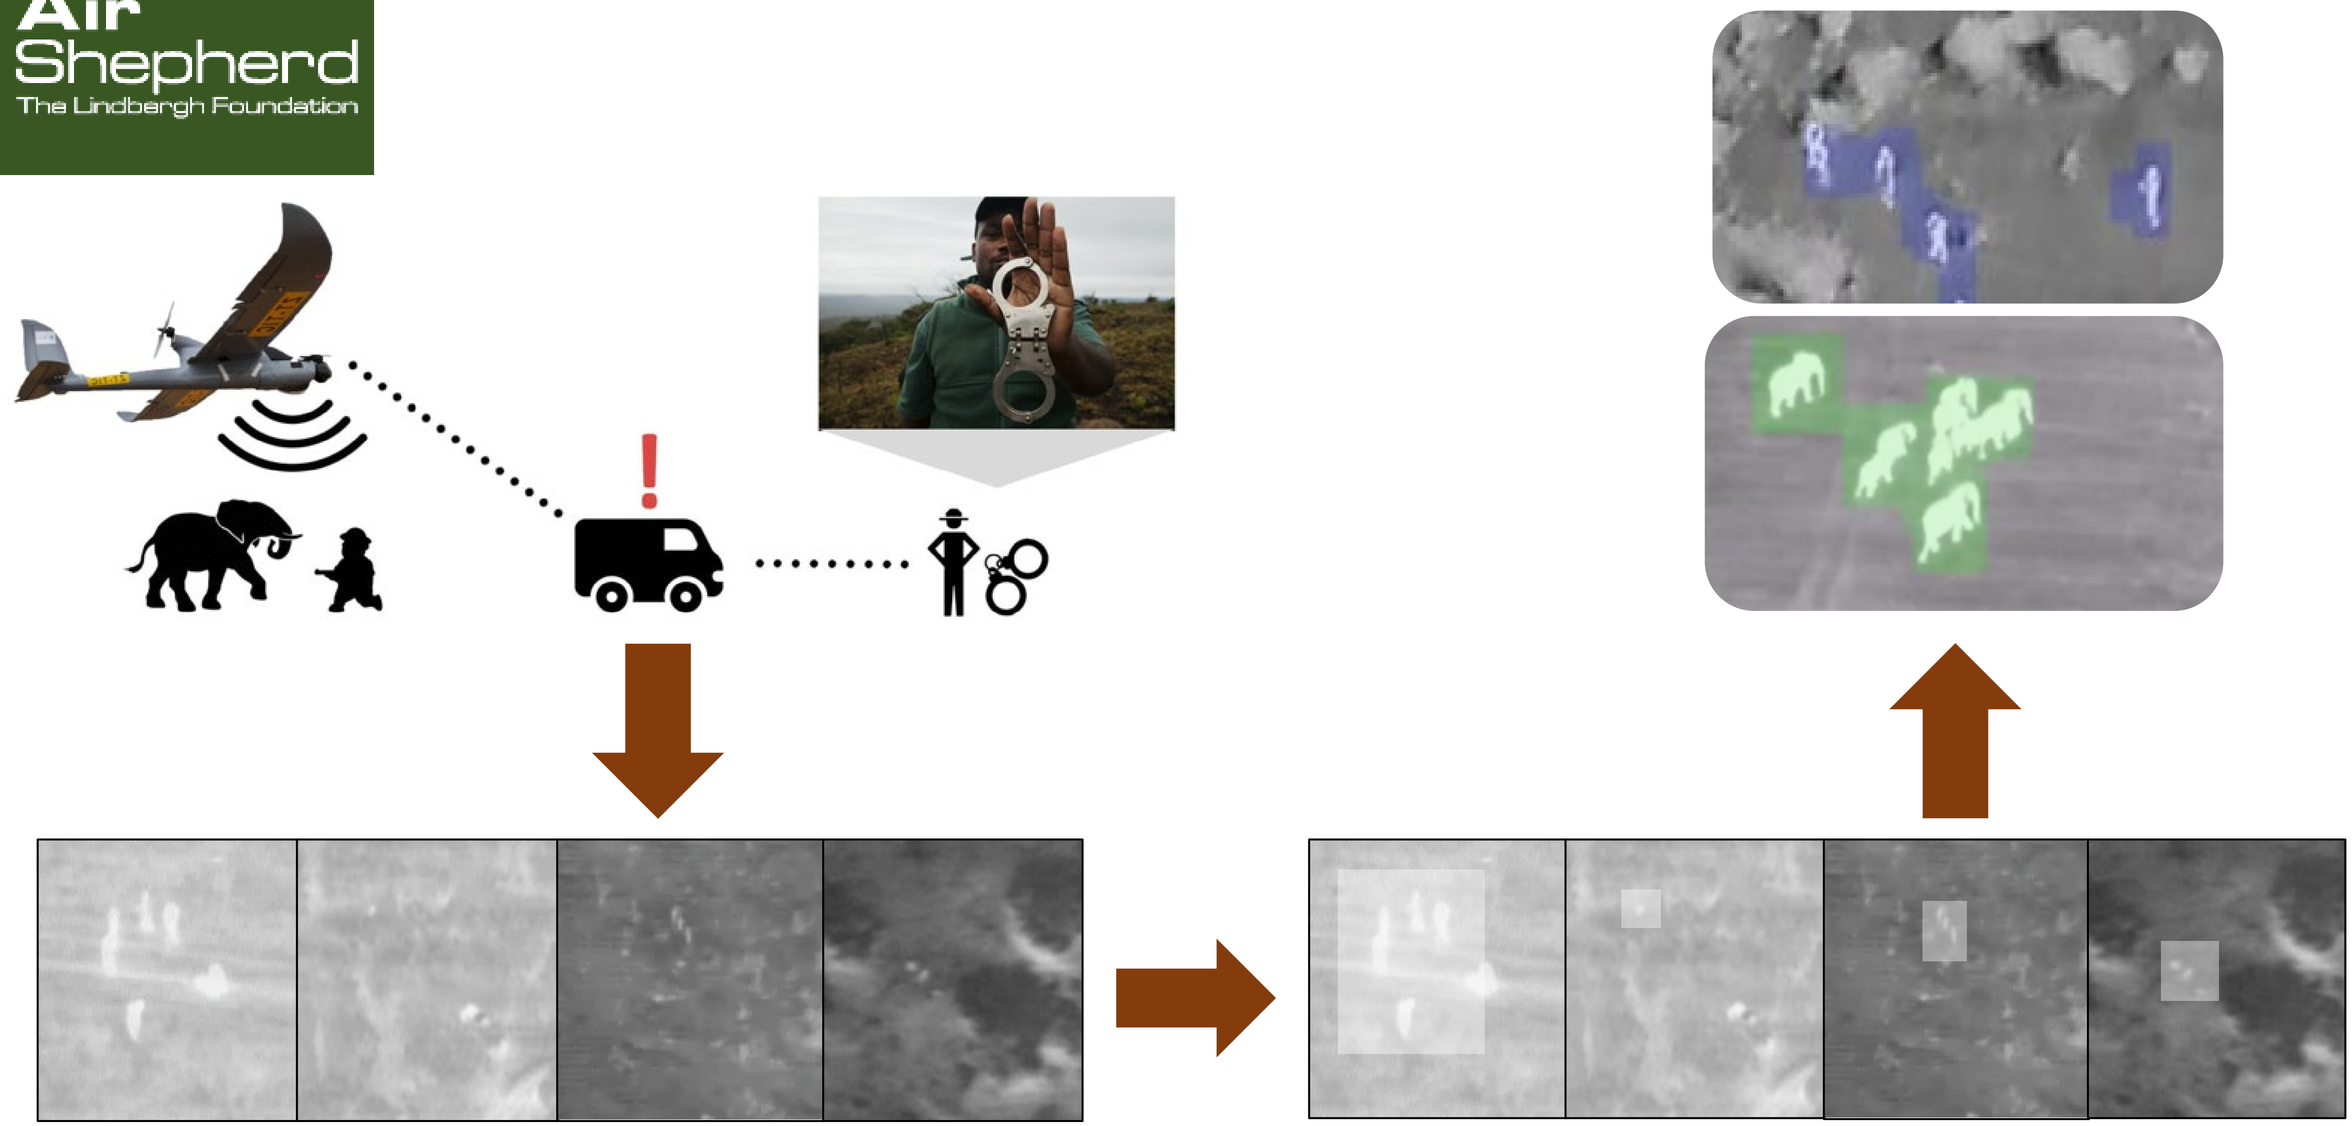

Goal: automatically find poachers

Drone Used to Inform Rangers

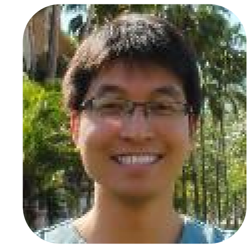

Xu

Bondi

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

Air Shepherd
The Lindbergh Foundation

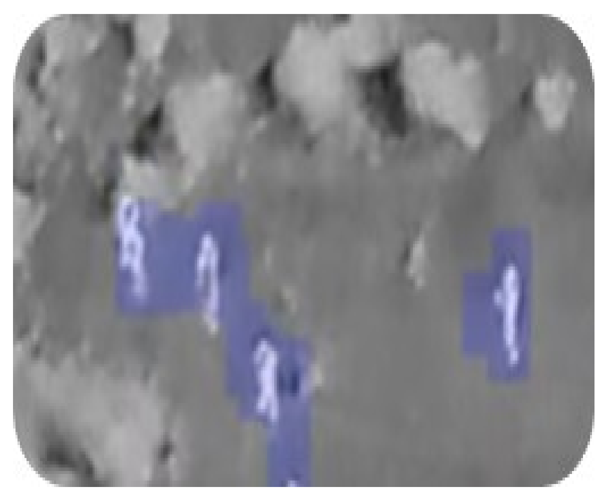

Prob(ranger) = 0.3

Drone Used to Inform Rangers

Xu

Bondi

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

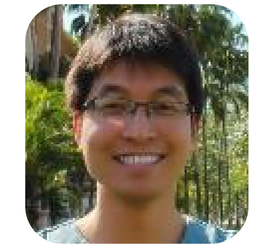
Xu


Bondi

Drone Used to Inform Rangers

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

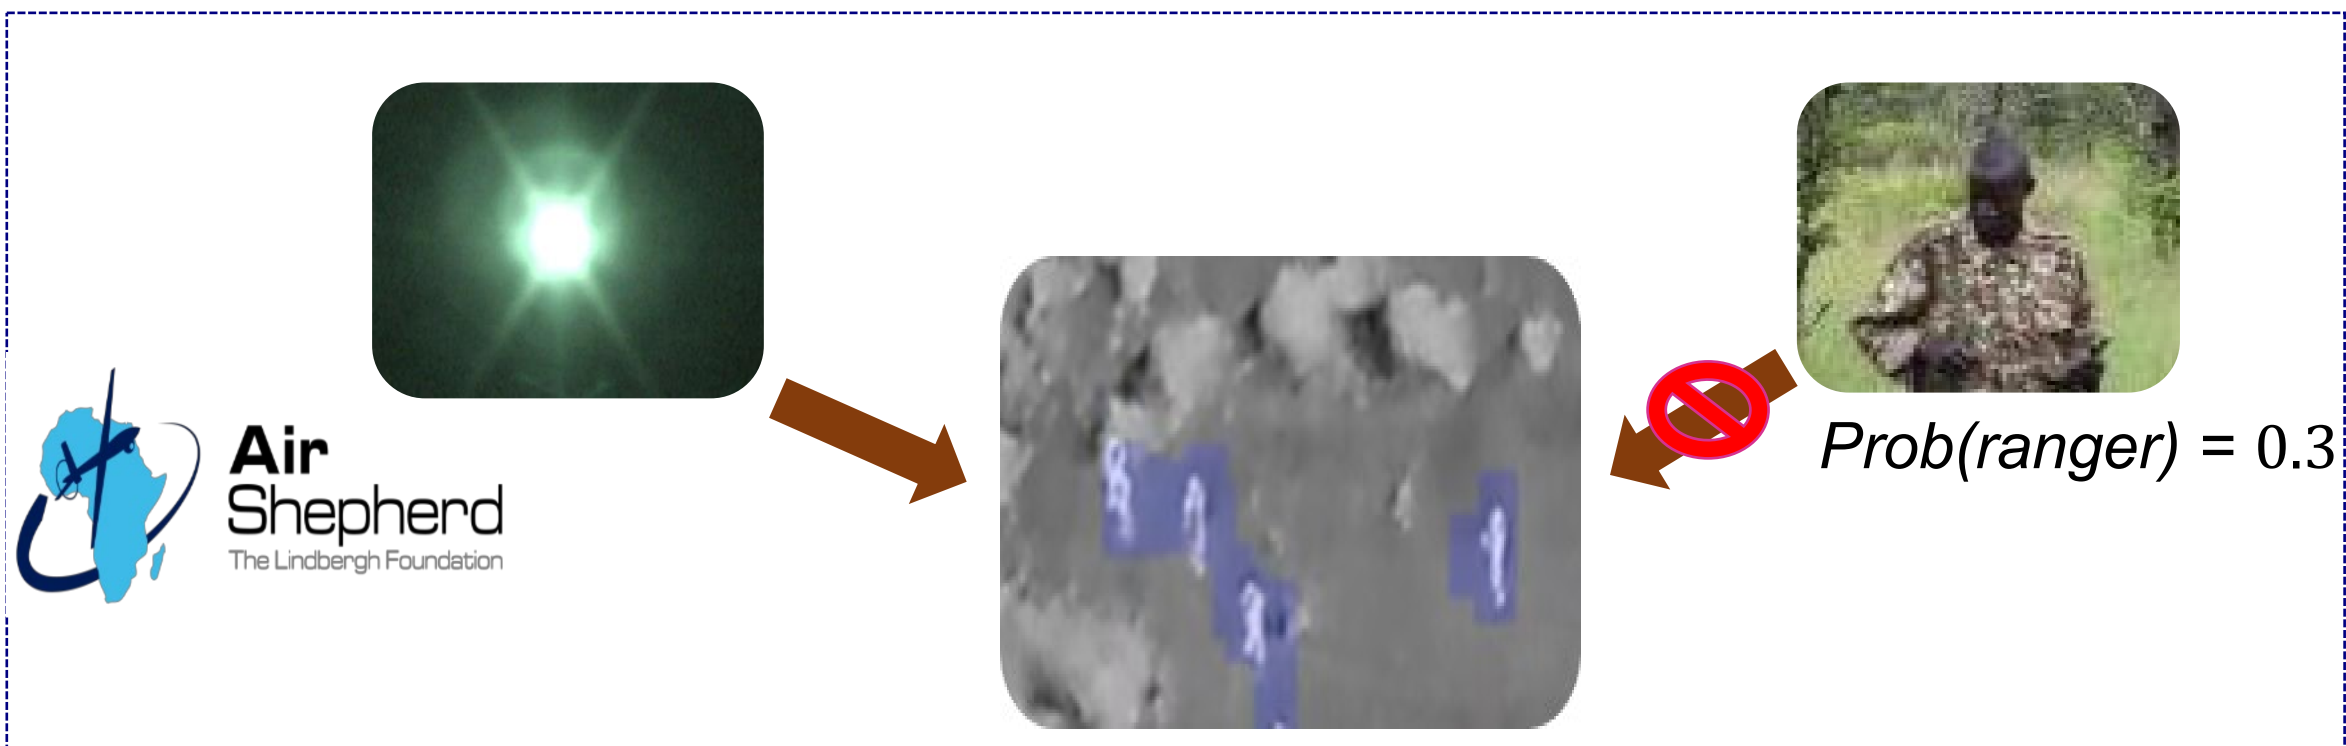

Strategic Signaling: Informational Advantage Defender Knows Pure & Mixed Strategy

(AAAI 2018, AAAI 2020)

Xu

Bondi

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

- Poacher best interest to “believe signal” even if know 50% defender deception
- *Theorem: Computing optimal defender commitment is NP-hard even for zero-sum Si-G*
- Extended to strategic signaling in presence of errors detecting adversaries

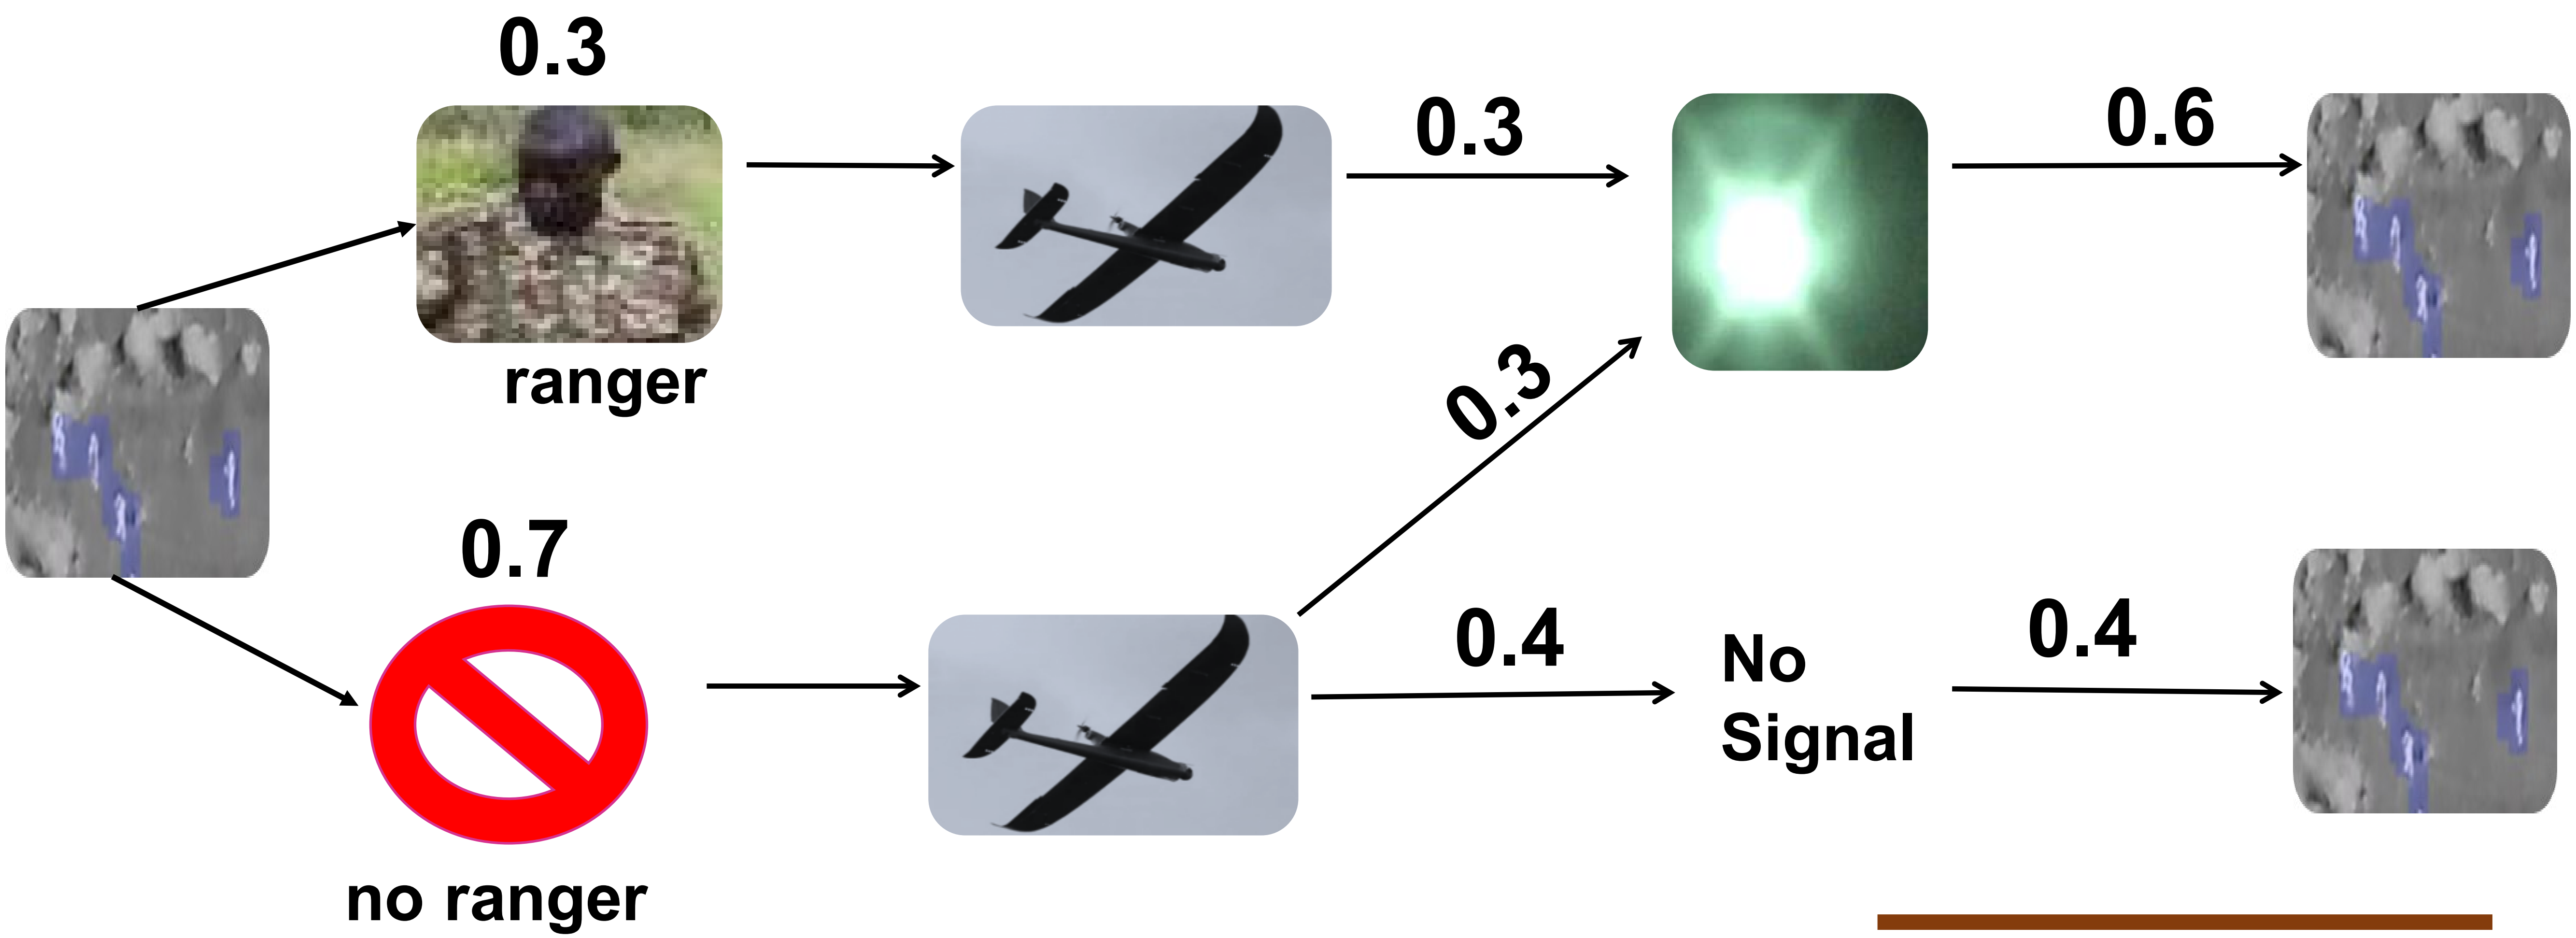

PAWS GOES GLOBAL with SMART platform!!

SMART

WWF

PEACE PARKS FOUNDATION · THE GLOBAL SOLUTION

WCS

AI for Earth

Protect Wildlife
800
National Parks
Around the Globe

Also: Protect Forests, Fisheries…

Challenges from field deployments

Integrated learning of adversary models & game theoretic planning

- ➢ Learning adversary models with limited real-world data
- ➢ Robust game theoretic planning with learned adversary model uncertainty
- ➢ Active gathering of adversary information

Strategic signaling in "green security games" with real-time information

- ➢ Uncertainty in real-time information
- ➢ Multi-agent spatio-temporal coordination in signaling
- ➢ Cost-benefit tradeoff in investment in signaling

Three Key Takeaways

EC research on multiagent system impactful for public health/conservation projects

Field deployments bring up new research challenges for EC community

Wealth of new multiagent research challenges via partnerships with NGOs

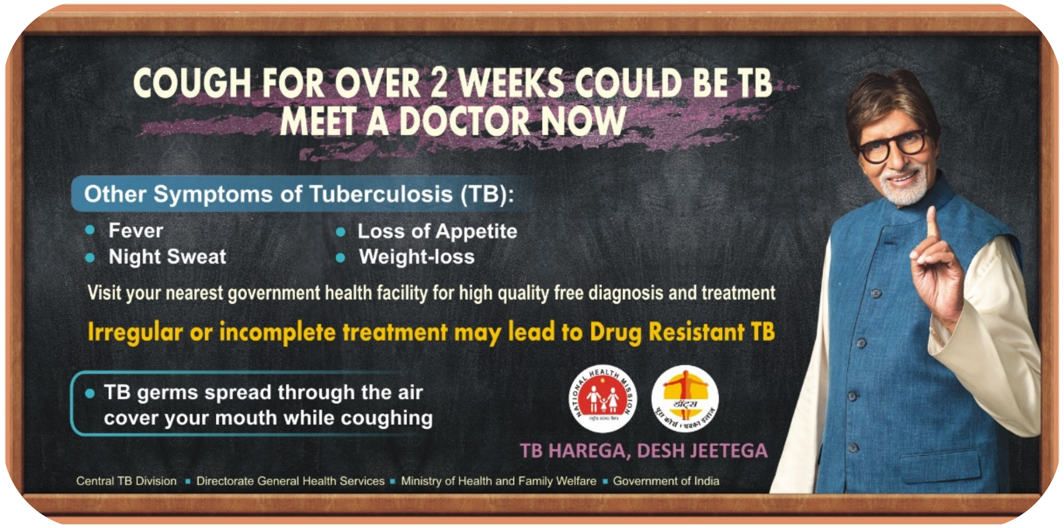

Future: AI Research for Social Impact

It is possible to simultaneously advance AI research & achieve social impact

Data to deployment perspective: Not just improving algorithms

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

AI for Social Impact should be evaluated differently

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