AI for Public Health & Conservation:
Learning & Planning in the Data to Deployment Pipeline

MILIND TAMBE

Director, Ctr for Research on Computation & Society

Harvard University

Director "AI for Social Good"

Google Research India

@MilindTambe_AI

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

Google Research

AI & Multiagent Systems Research for Social Impact

Public Health

Conservation

Public Safety and Security

Key Research Challenge

Optimize Our Limited Intervention Resources

Optimizing Limited Intervention Resources

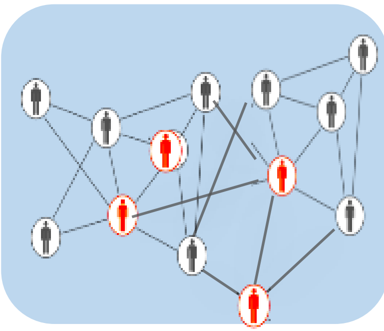

Social Networks & Bandits

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

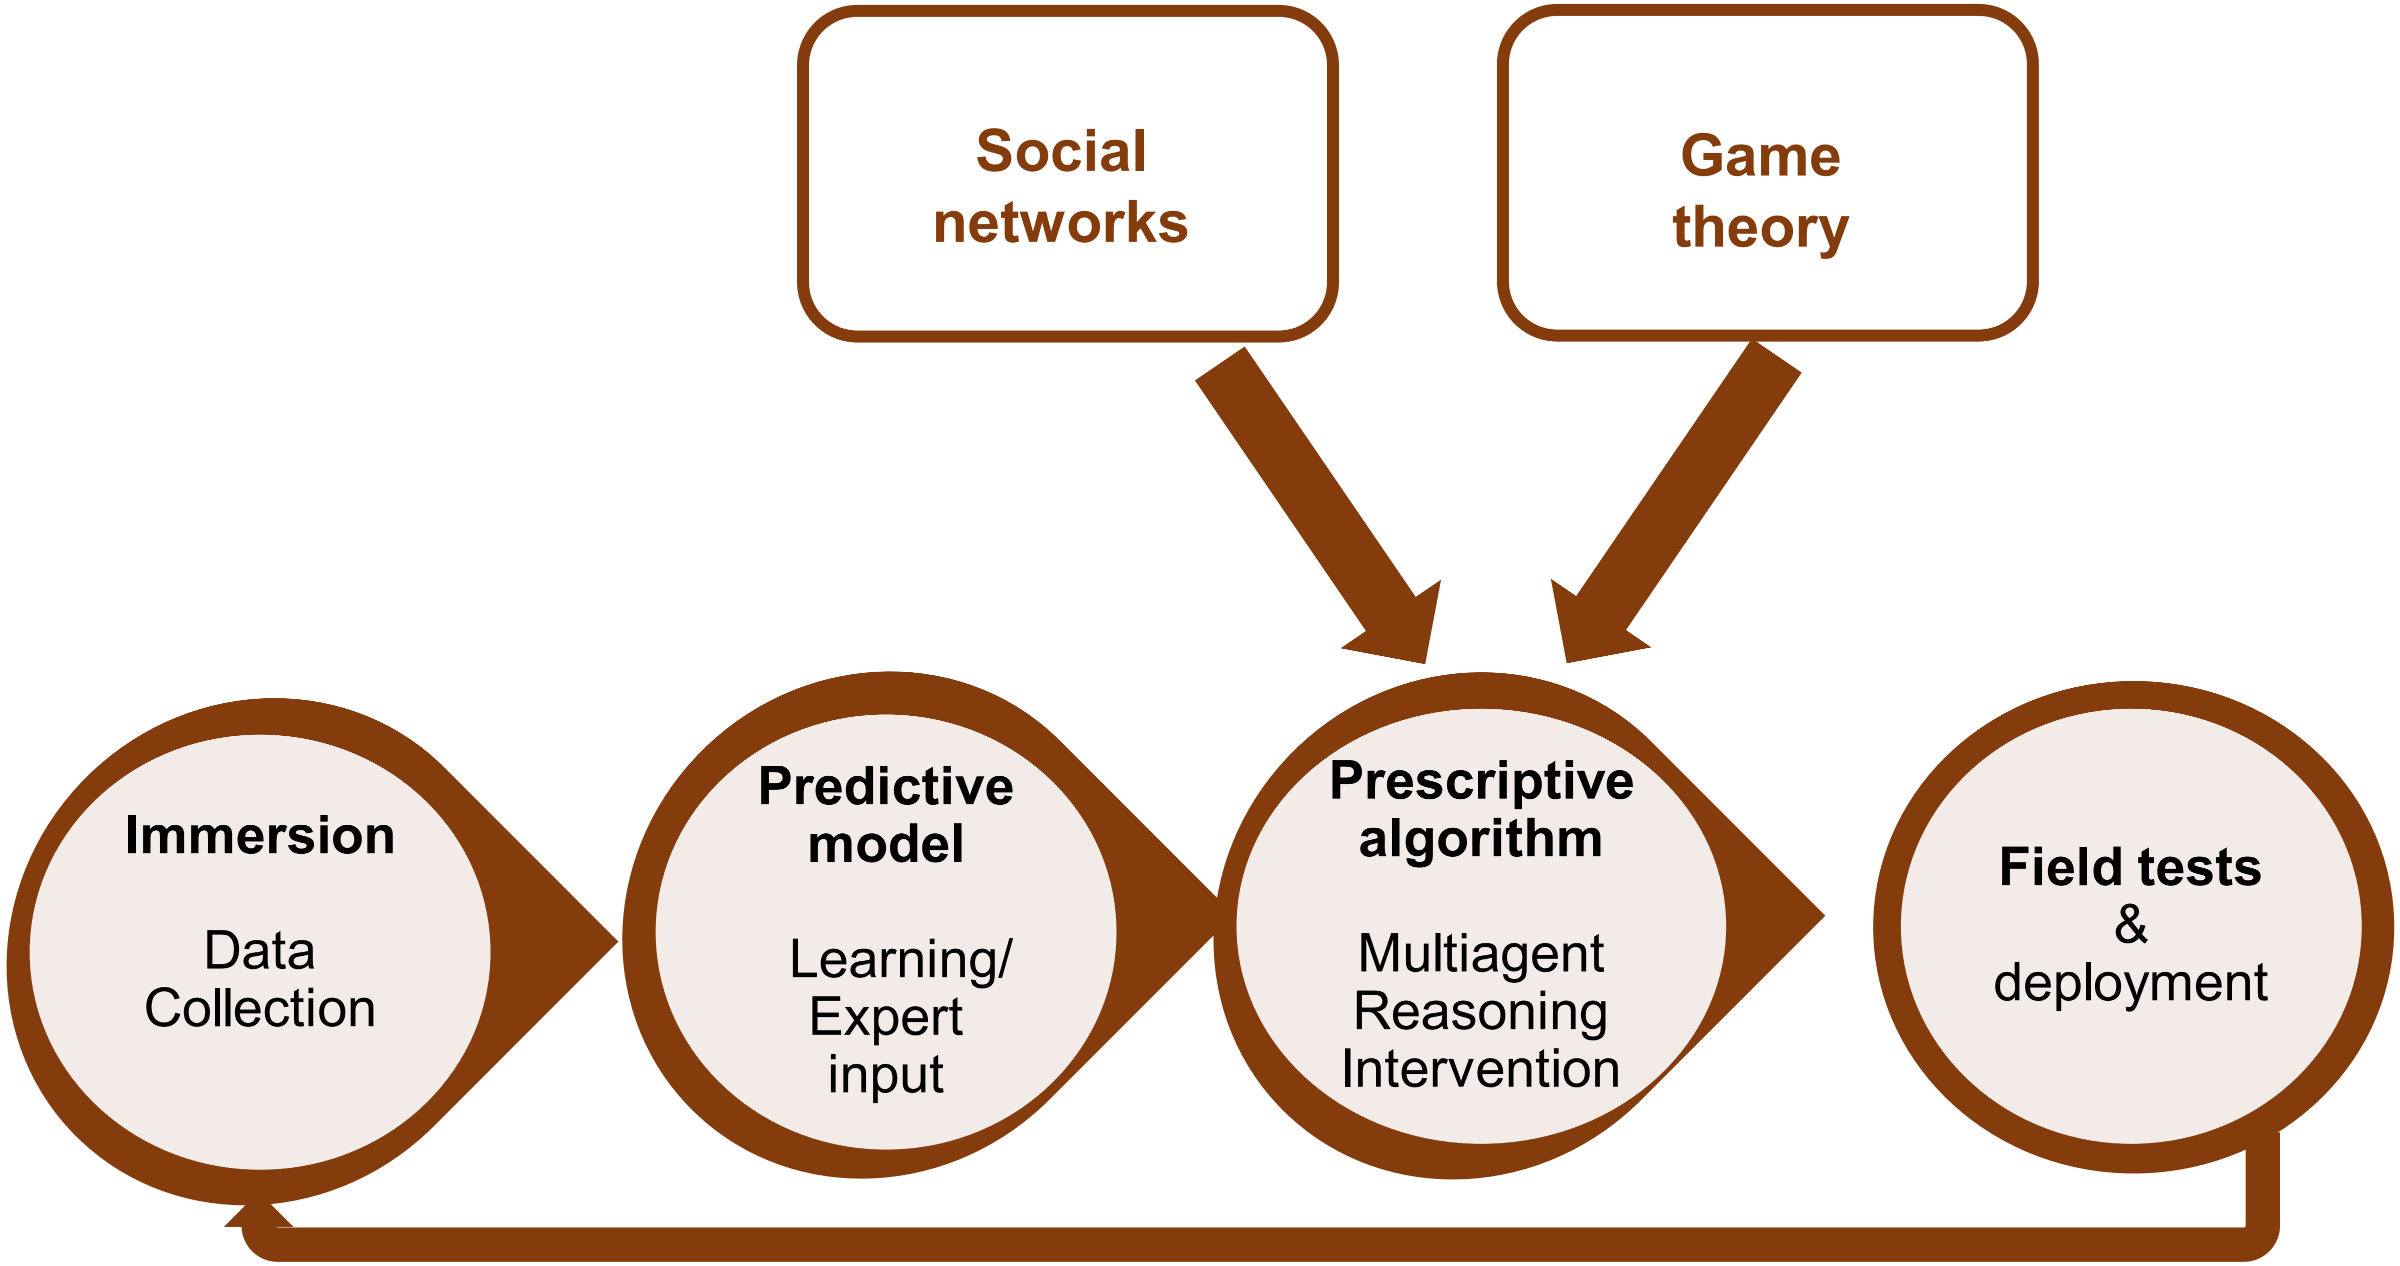

Three Common Themes

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

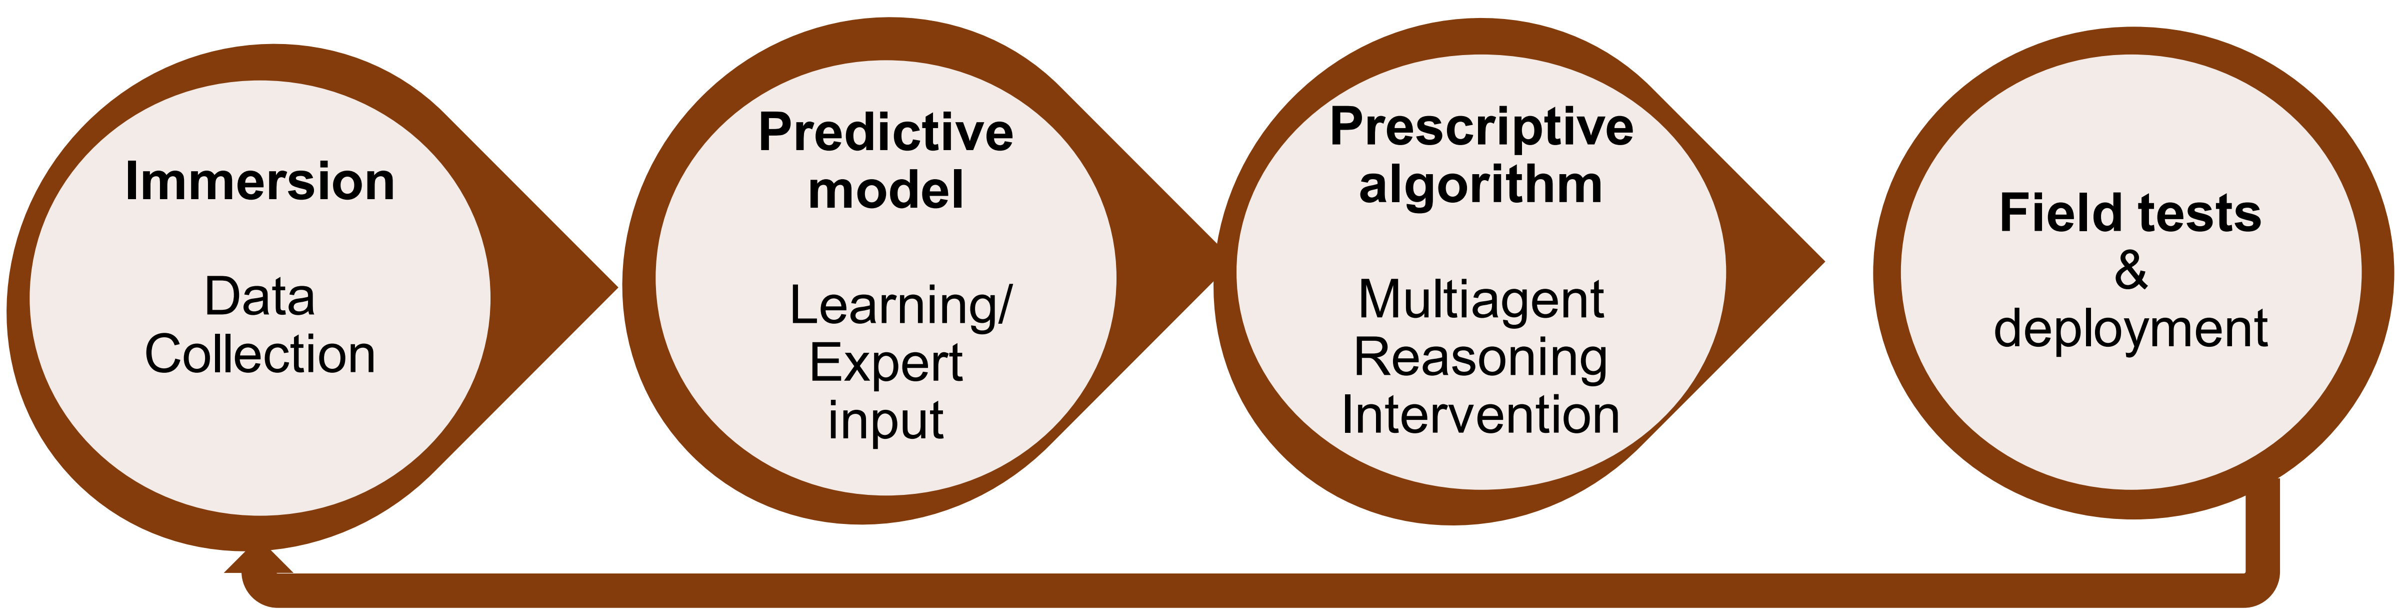

Three Common Themes

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

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 Delivering India's Future

ASHOKA TRUST FOR RESEARCH IN ECOLOGY & THE ENVIRONMENT (atree)

Wildlife Conservation Society

WWF

Wildlife Conservation Trust

SMART

Air Shepherd The Lindbergh Foundation

Swasti THE HEALTH CATALYST

everwell

KhushiBaby

PRATHAM BOOKS

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*

Influence Maximization in Social Networks

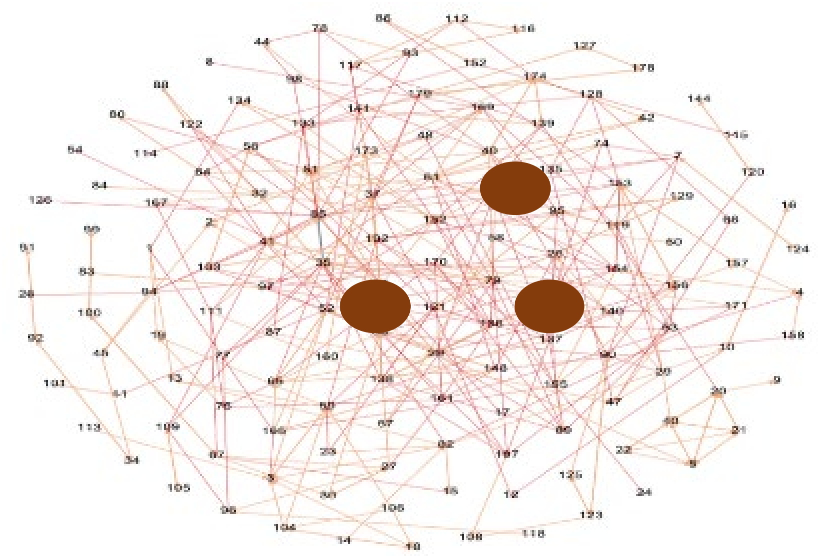

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
  - Assume: Independent cascade model of information spread

- 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

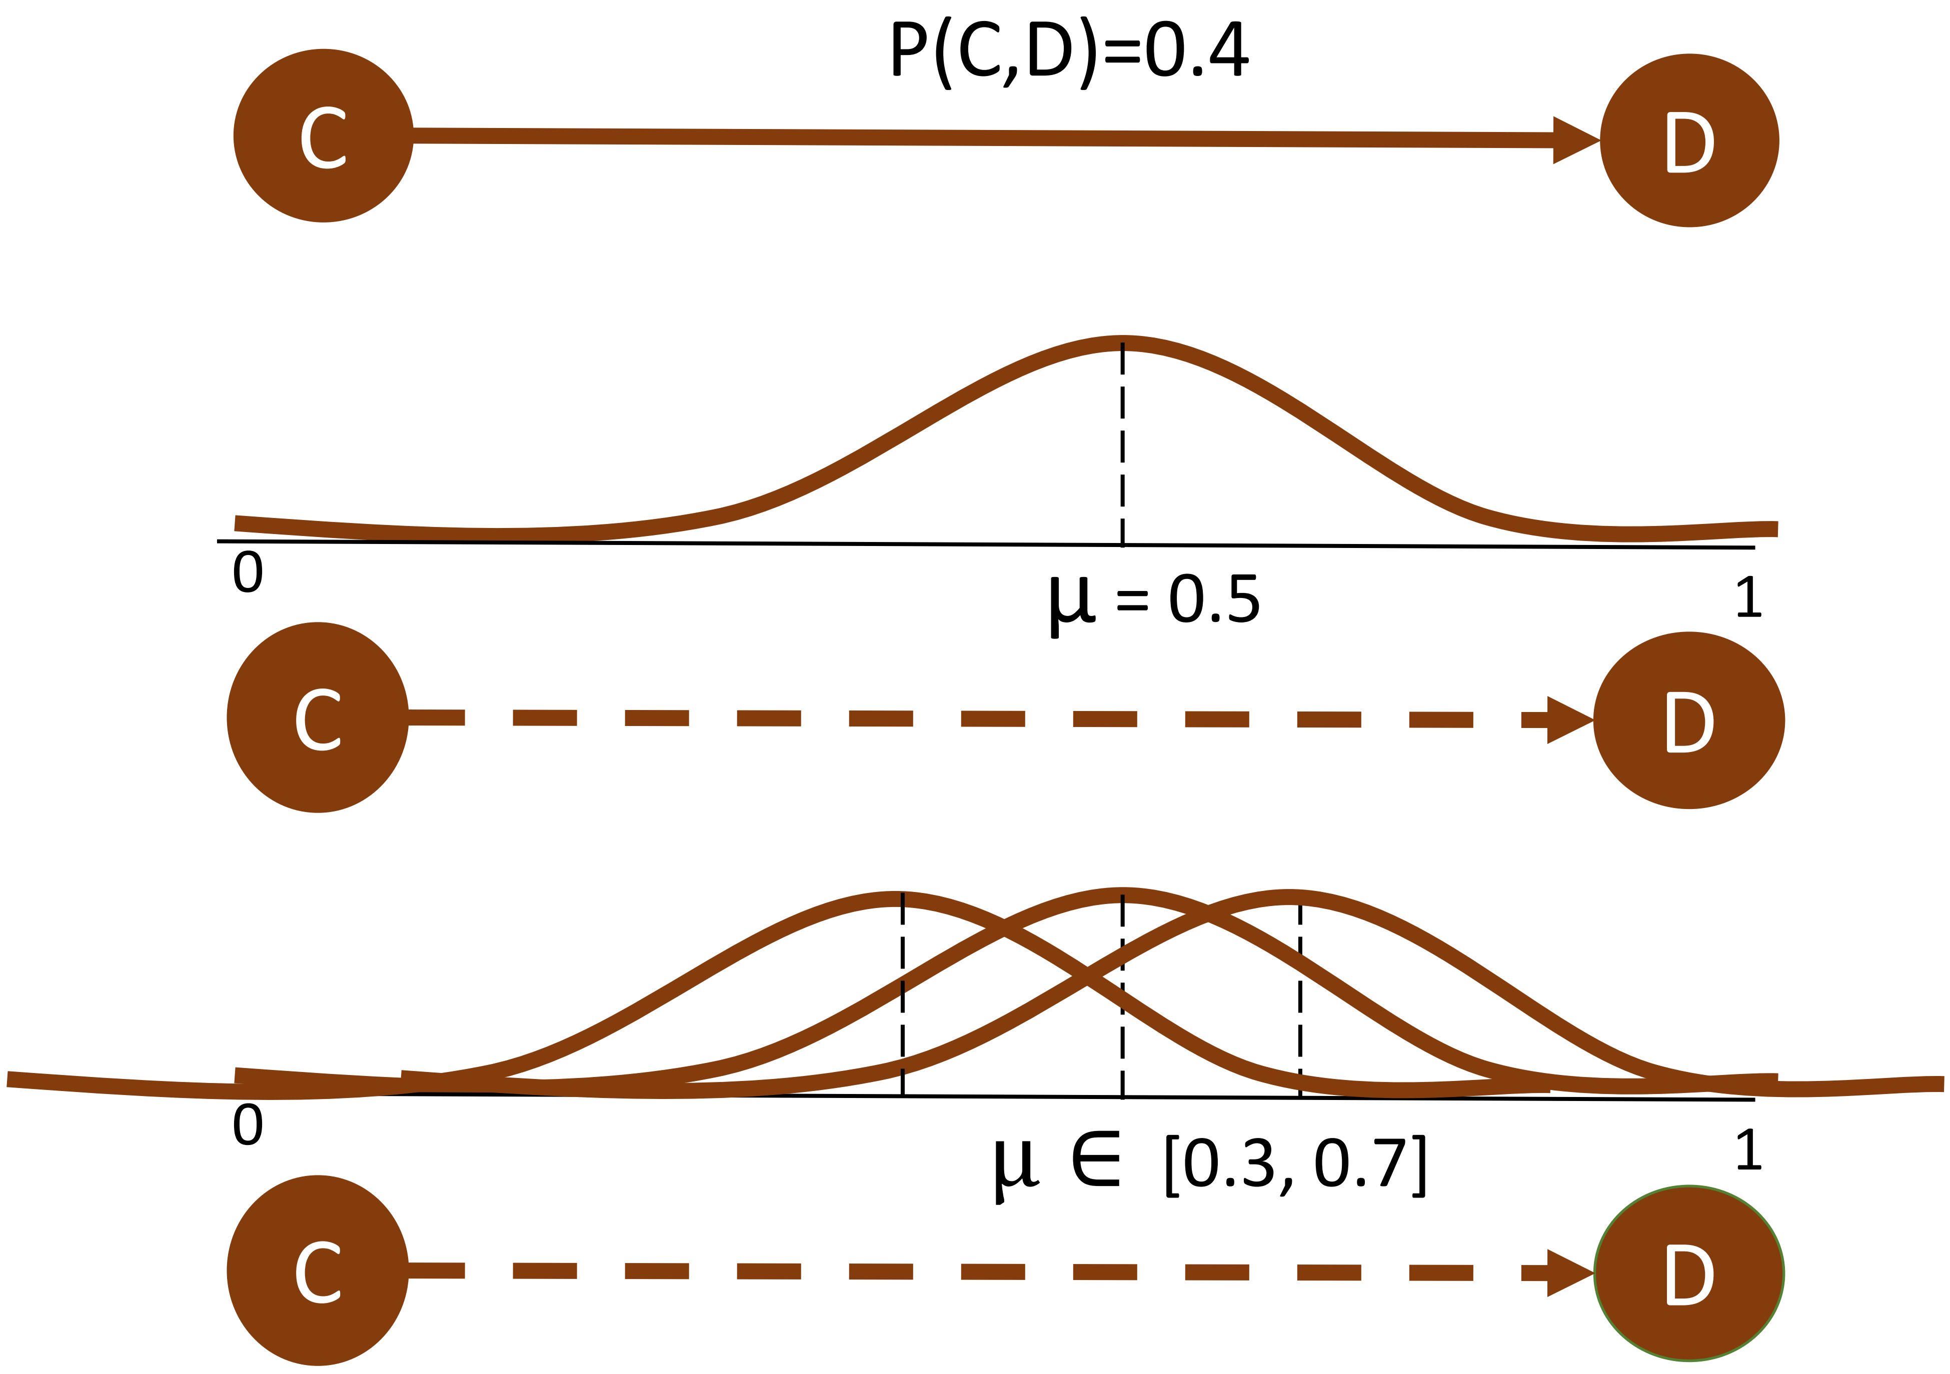

Robust Influence Maximization

(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: POMDPs for Multi-step Policy

(AAMAS 2018a)

Yadav

K = 4
1^{st} time step

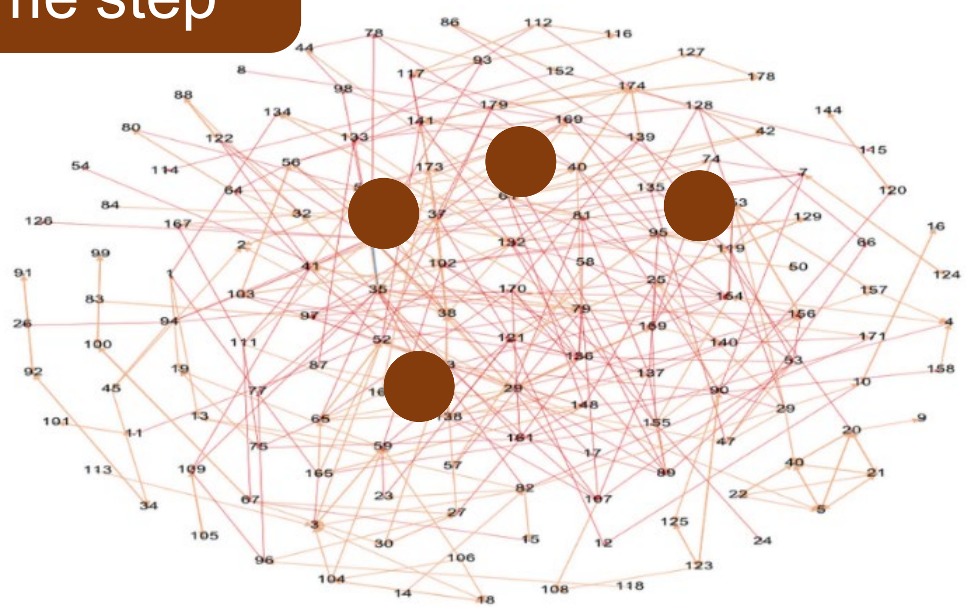

K = 4
2^{nd} time step

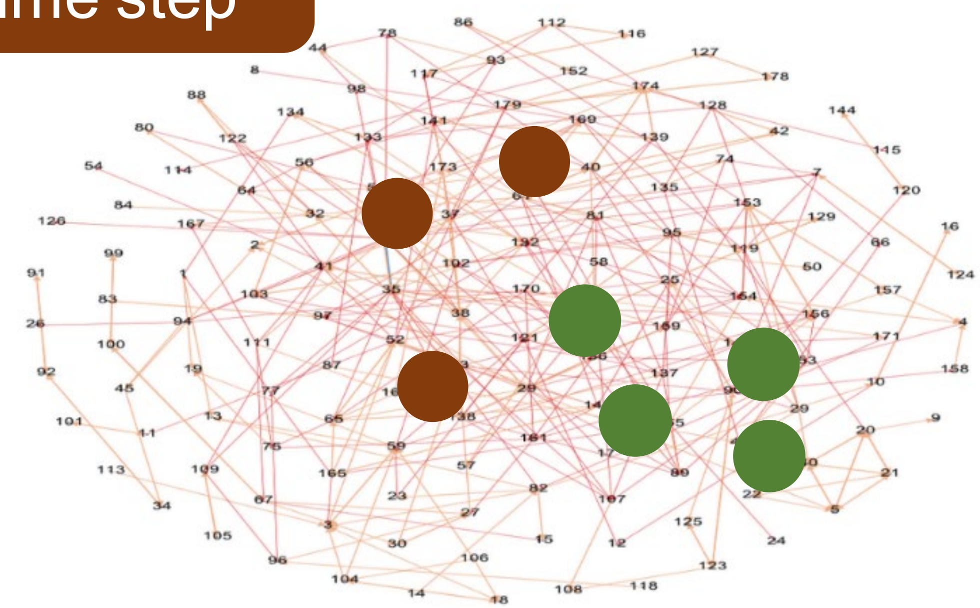

HIDDEN STATE

Action

Choose nodes

POMDP
Policy

Observation: Node presence

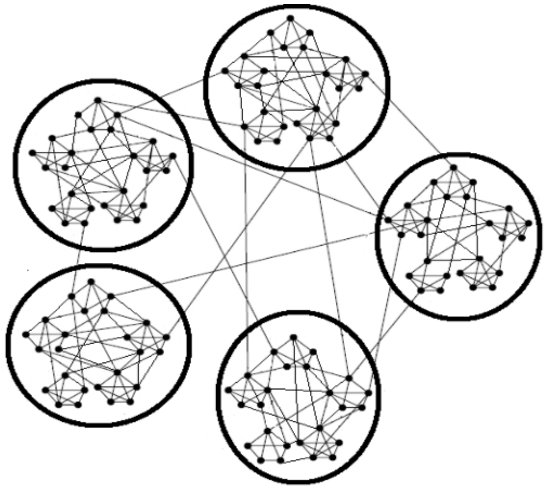

Partition POMDPs:
Exploit community structure

Challenge 3:
Sampling 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**

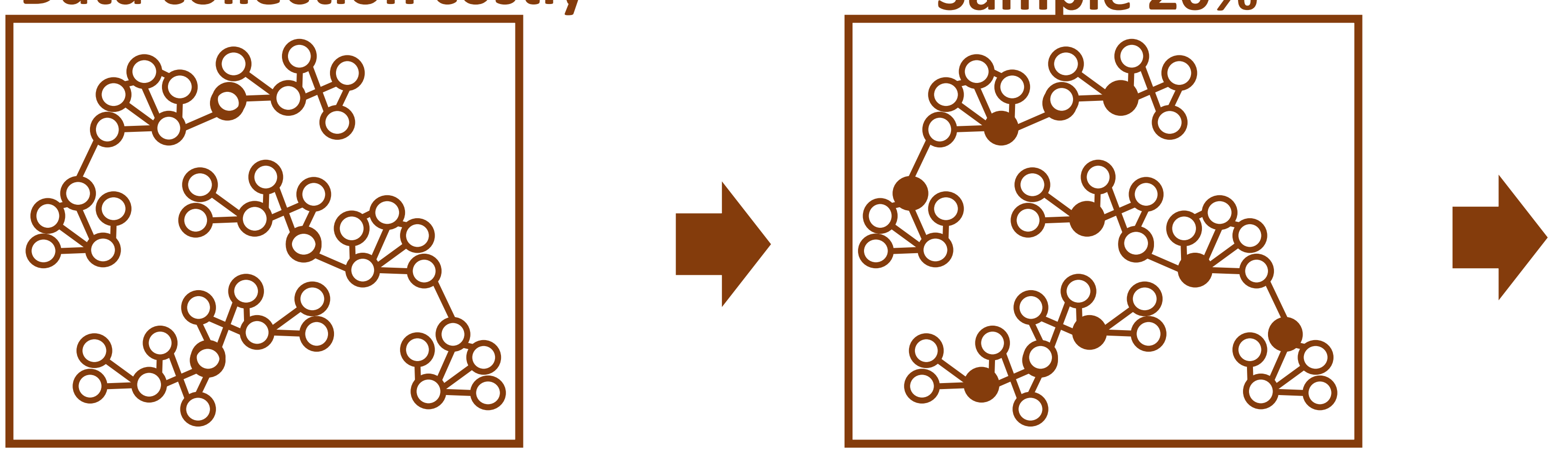

Sample node randomly & estimate size of its community;
Choose seeds from largest K communities

- Input: total number of nodes, n
- Query upto query budget
- 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

“Sampling-HEALER”
Pilot tests with 230 Homeless Youth

(IJCAI 2018)

Yadav

Wilder

Sampling-HEALER

Network Sampling

Robust multi-step policy

Peer leader selection

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

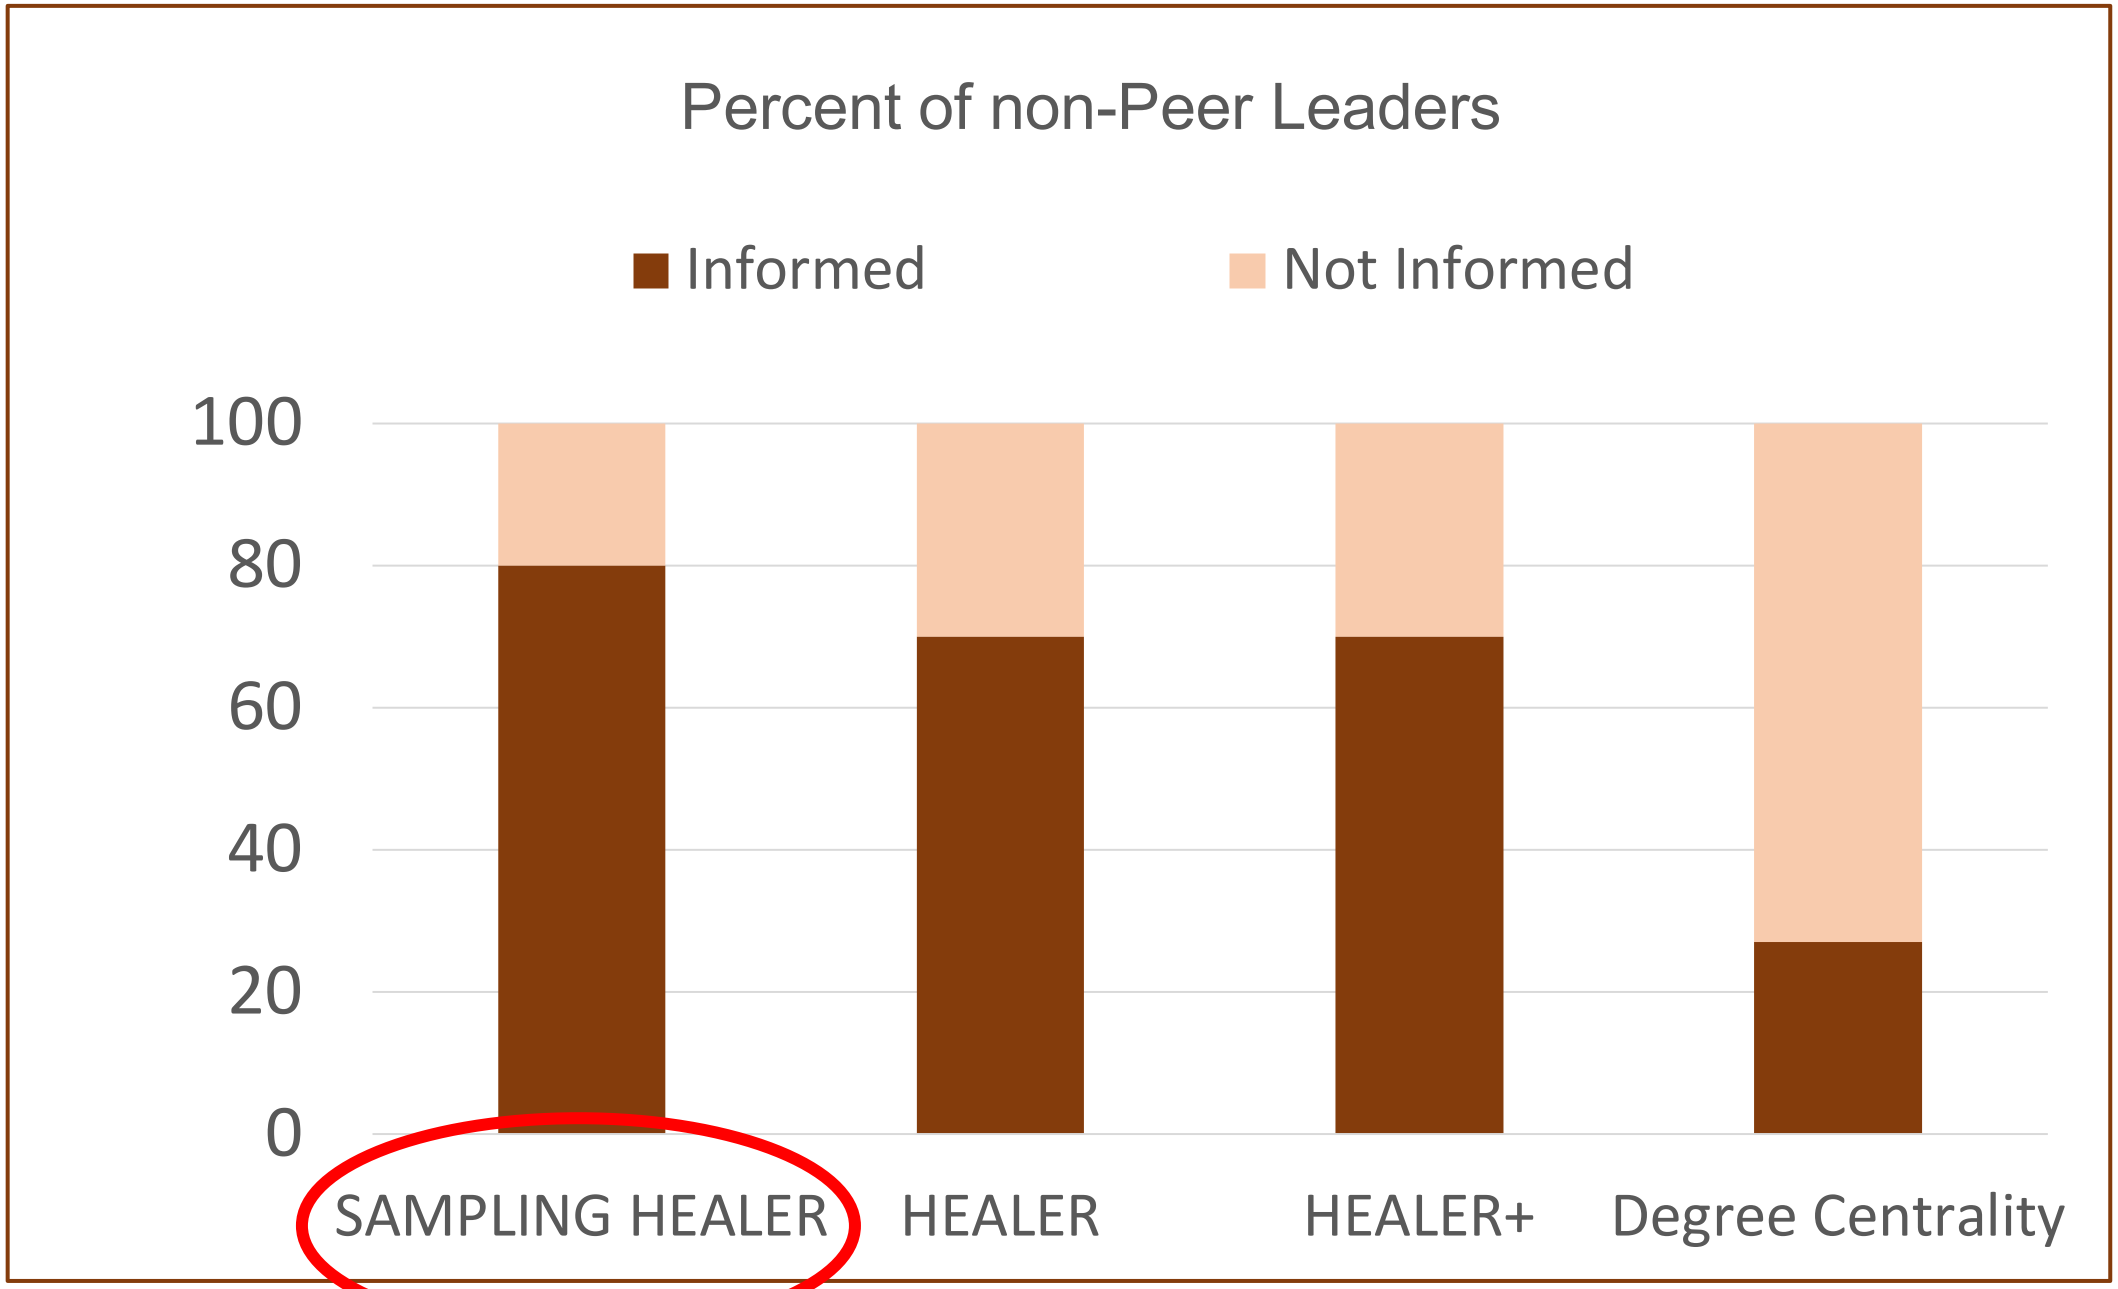

Results of 750 Youth Study [with Prof. Eric Rice]
Actual Change in Behavior?
(Under submission)

First large-scale application of influence maximization for public health

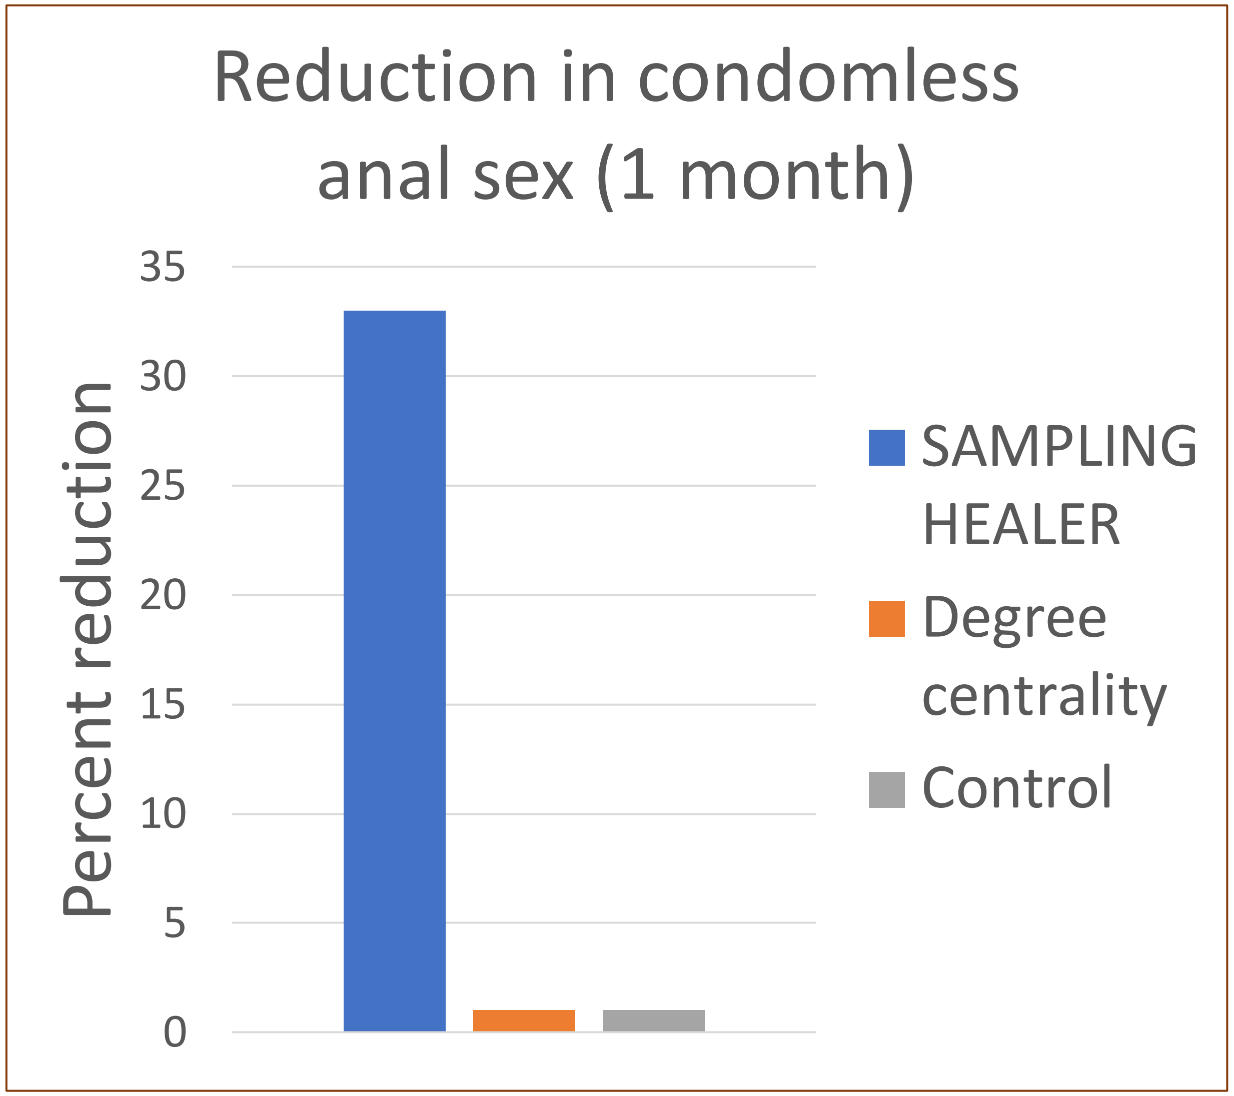

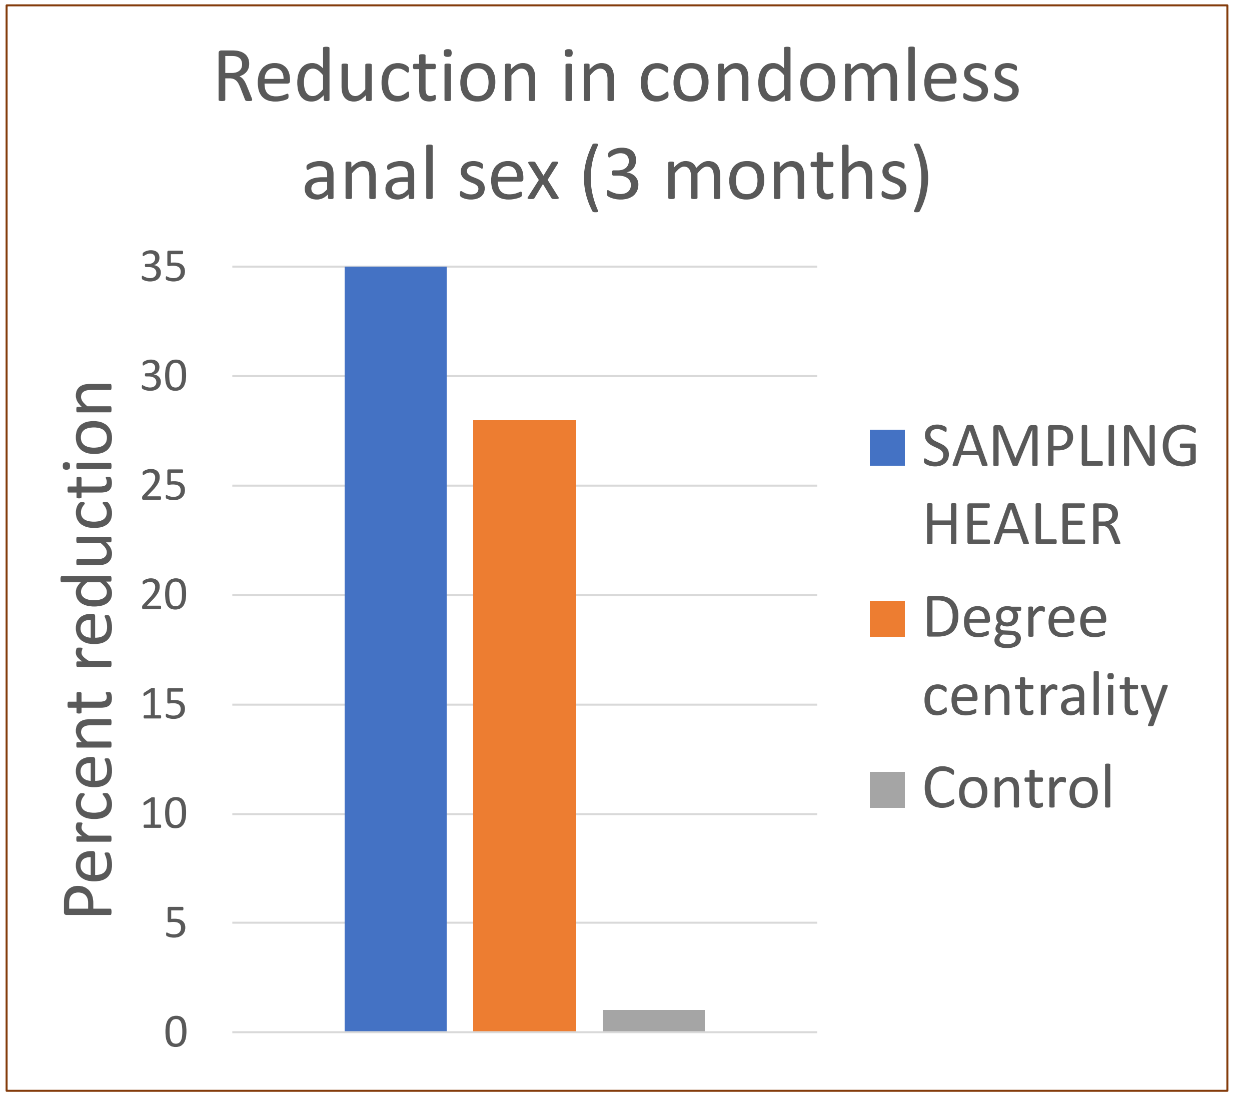

Results of 750 Youth Study [with Prof. Eric Rice]

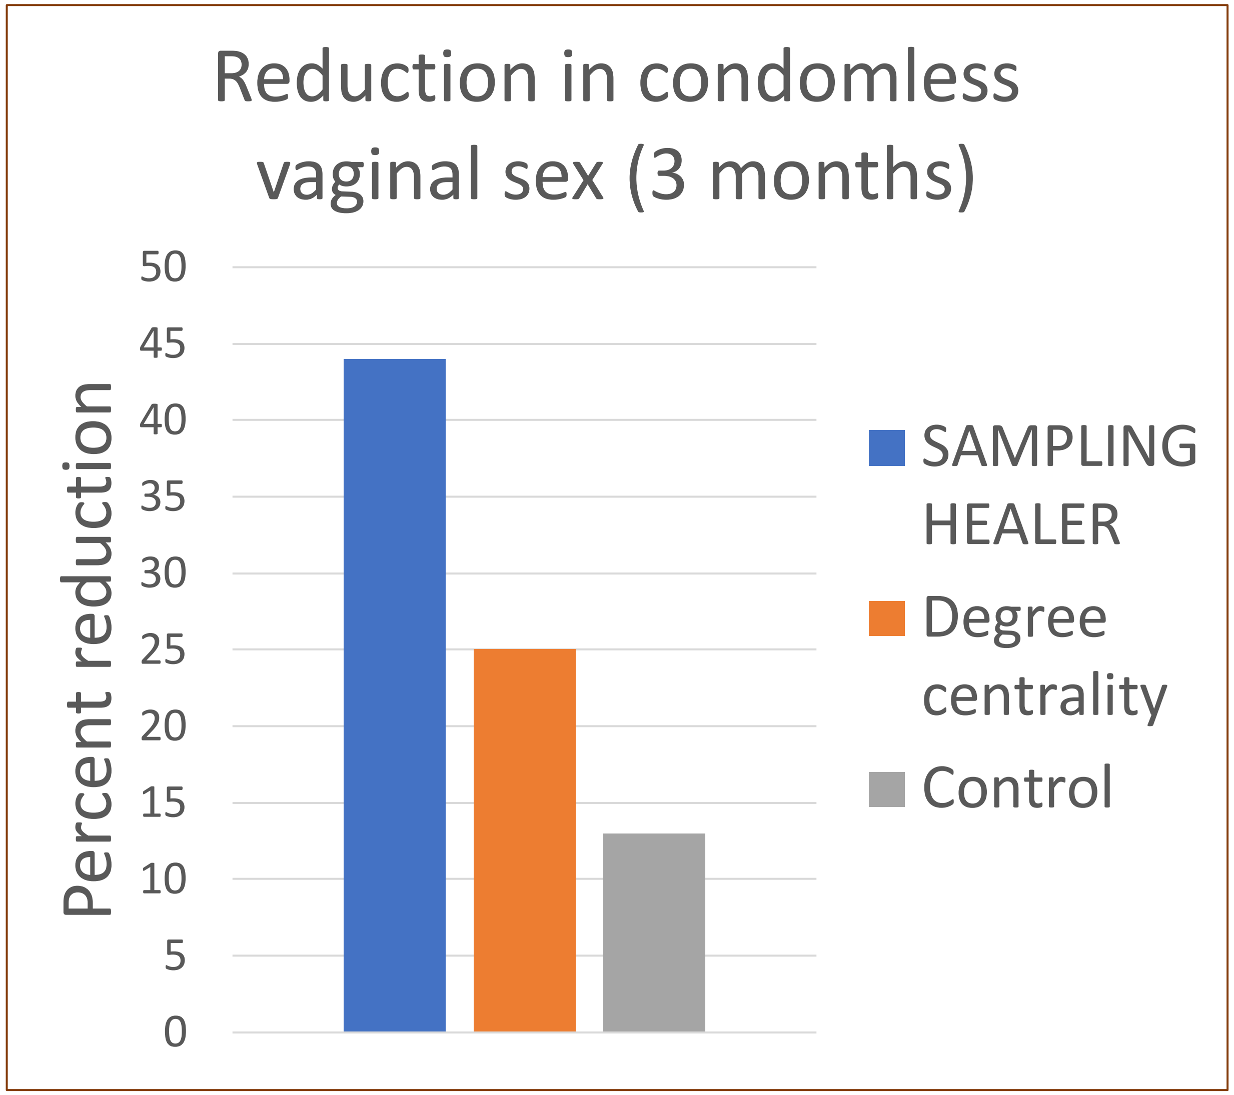

What our collaborators are saying:

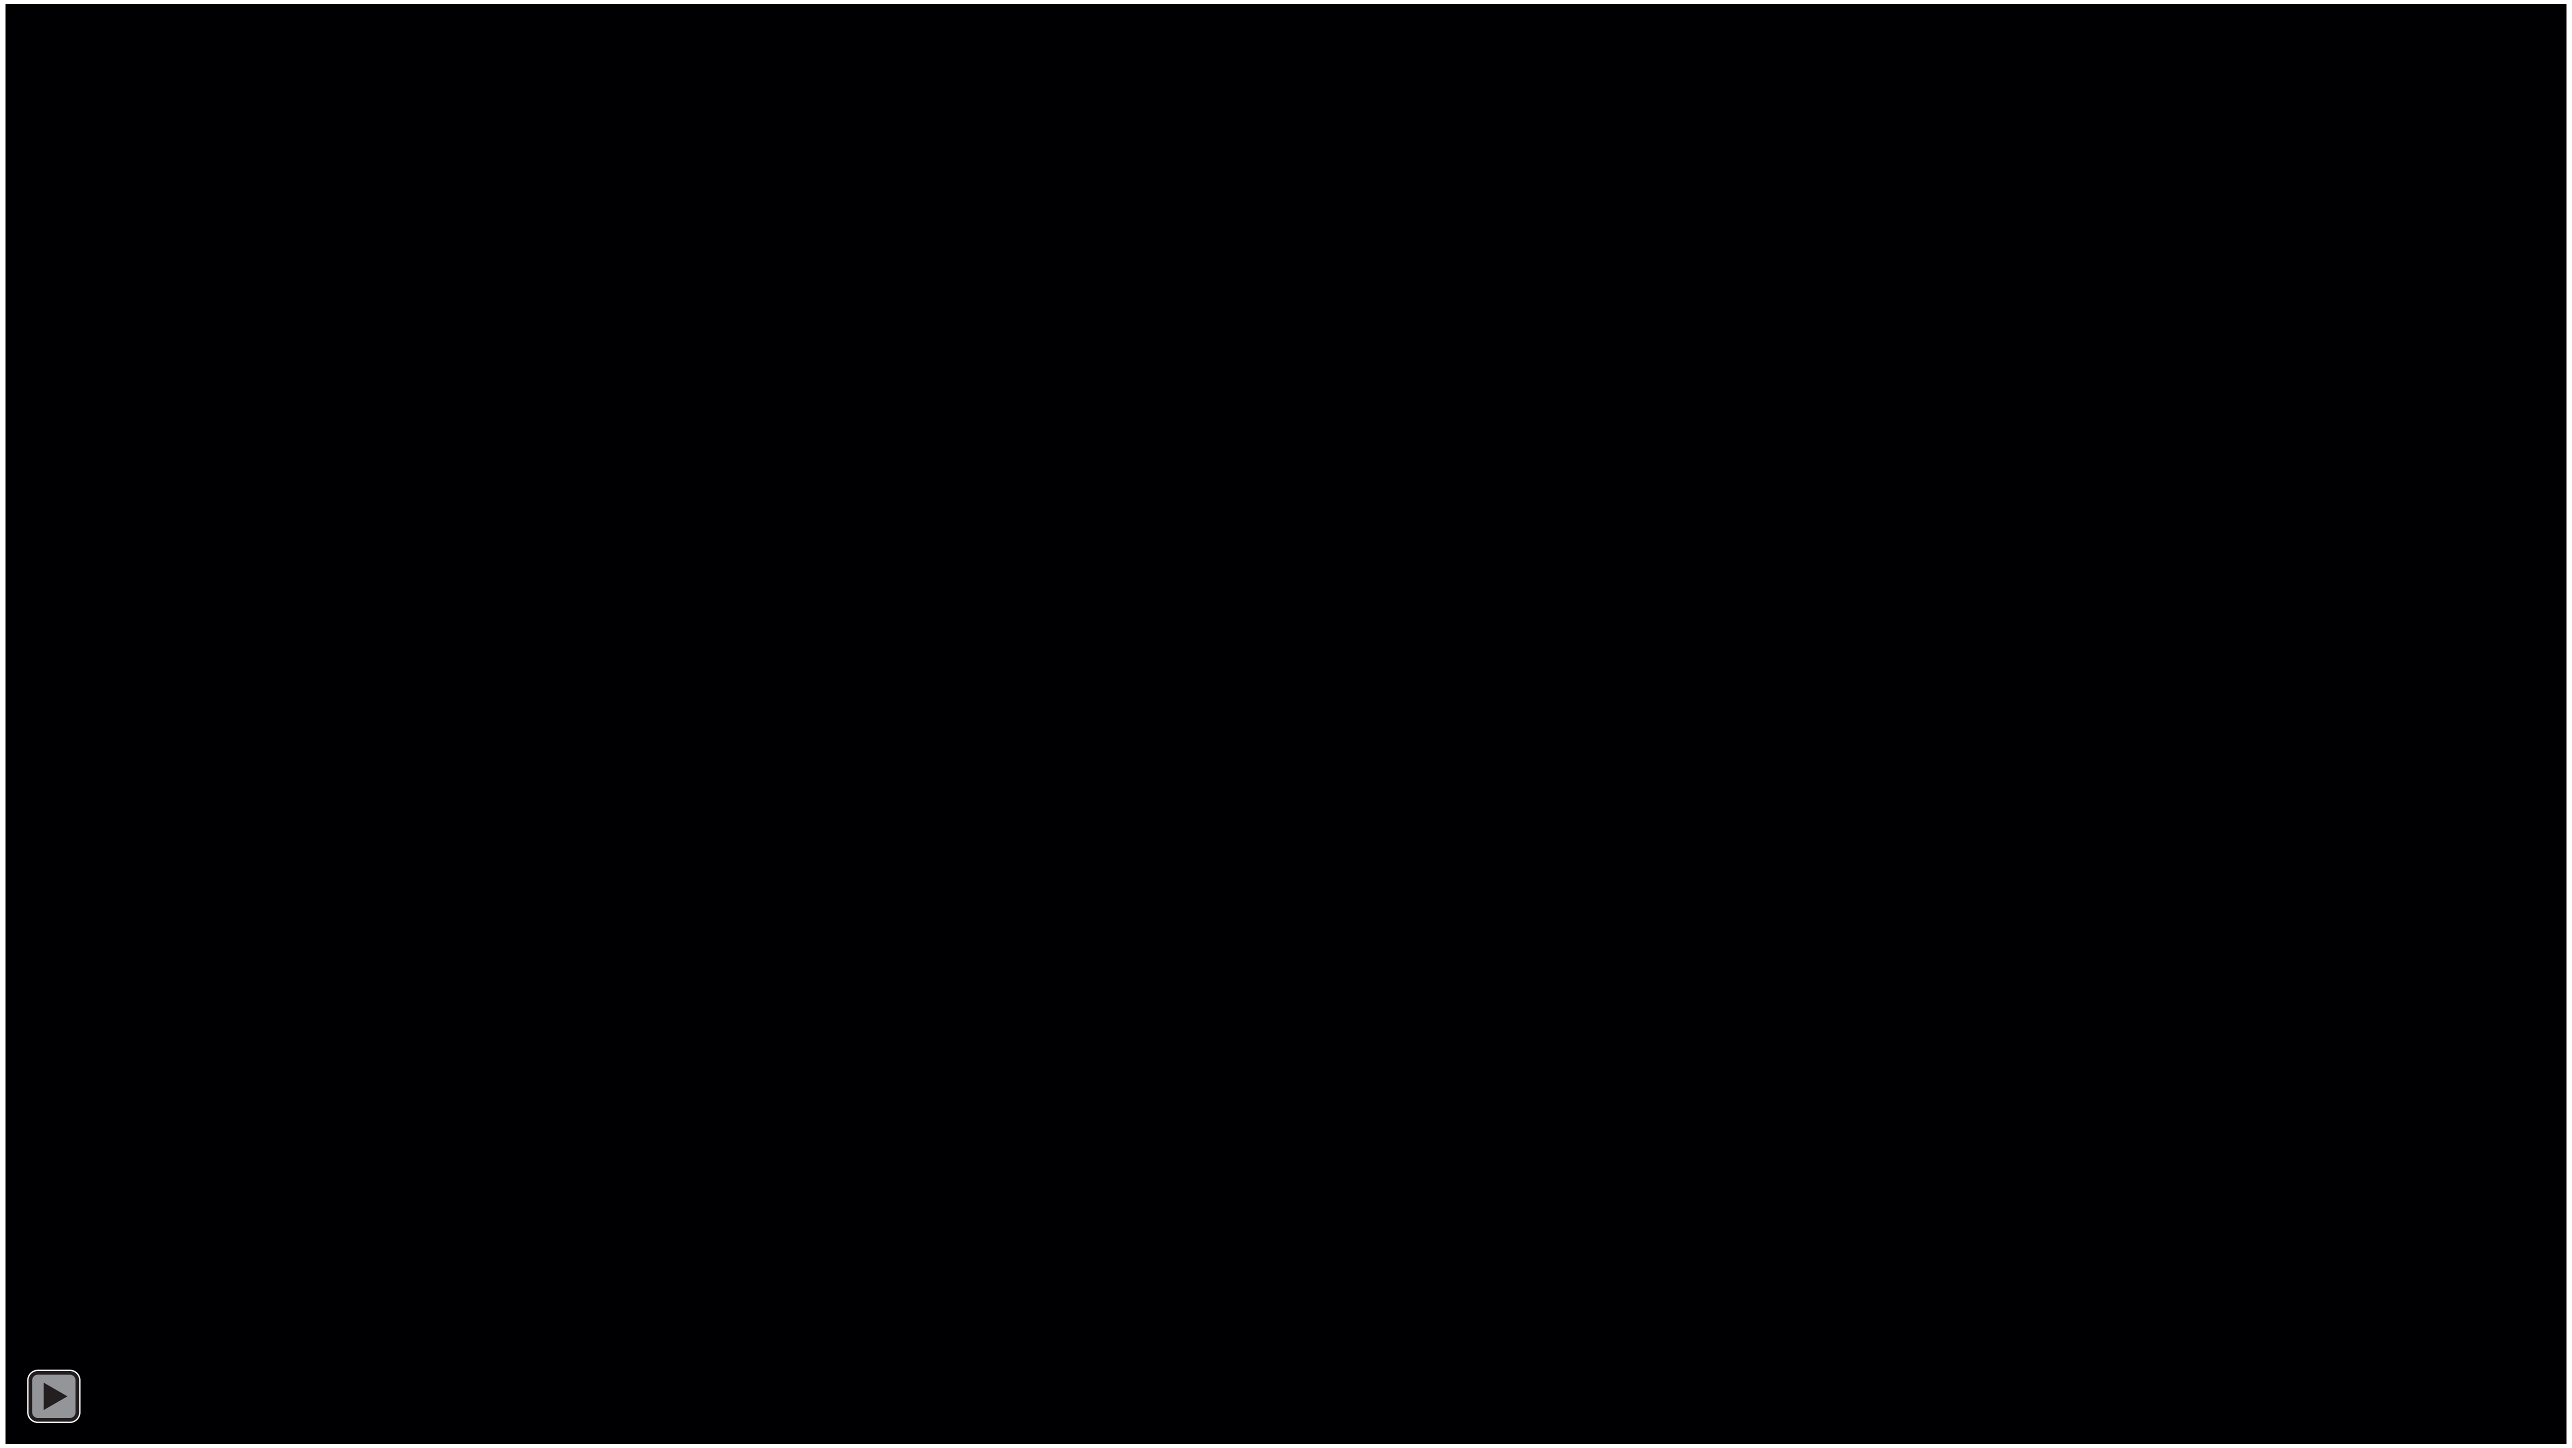

Next steps: Data to Deployment Pipeline Using an RL agent?

(with B. Ravindran & team, AAMAS 2020)

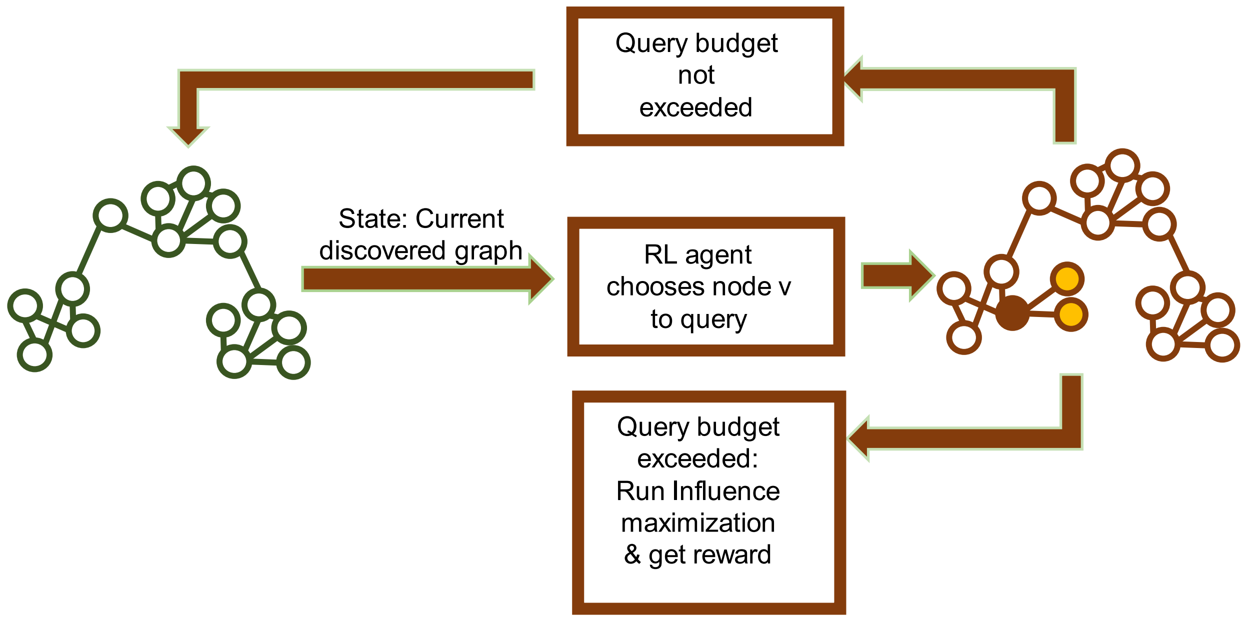

Network Family	Improve %
Rural	23.76
Animal	26.6
Retweet	19.7
Homeless	7.91

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

(NeurIPS 2019, IJCAI 2019)

Rahmattalabi

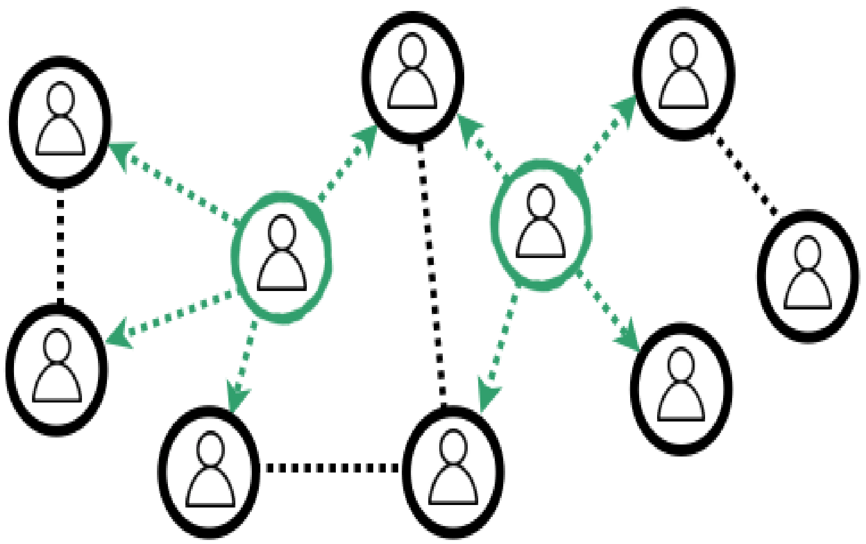

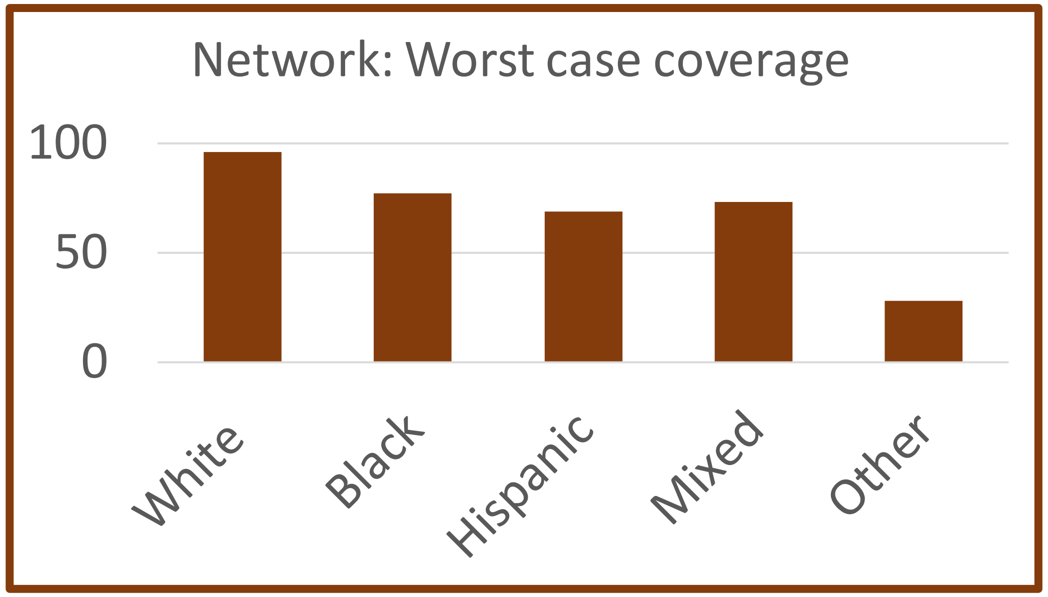

Robust graph covering with gatekeepers, maximize worst case coverage

Disparity in coverage across racial groups

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

γ: 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*

armman

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

Delivering India's Future

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***

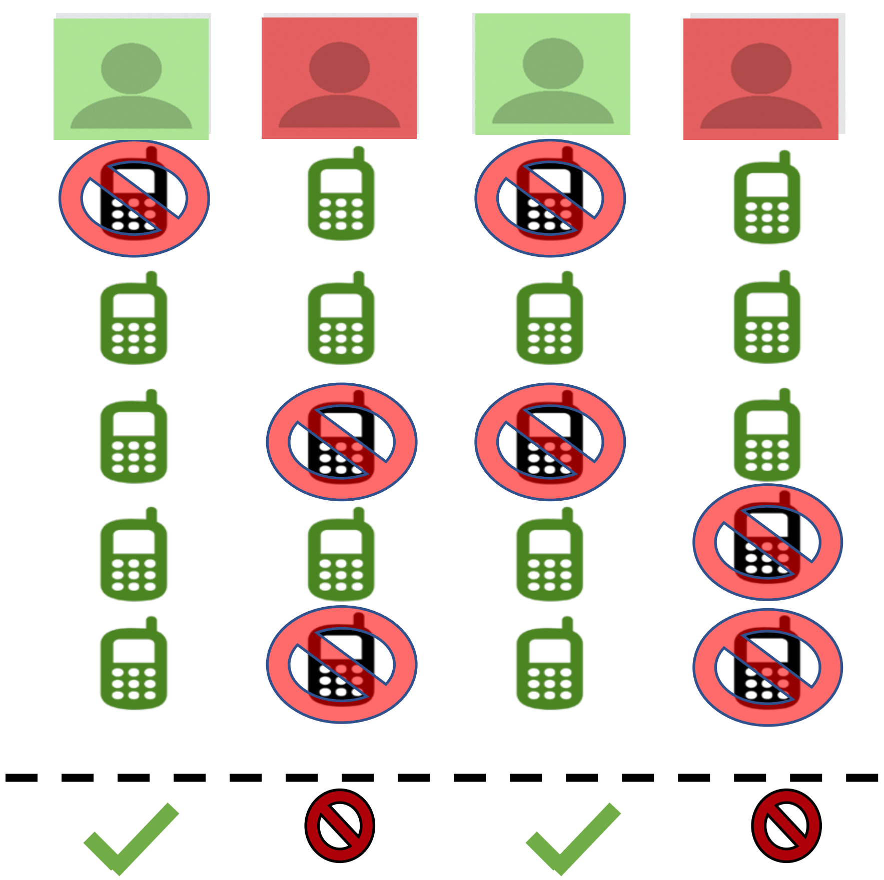

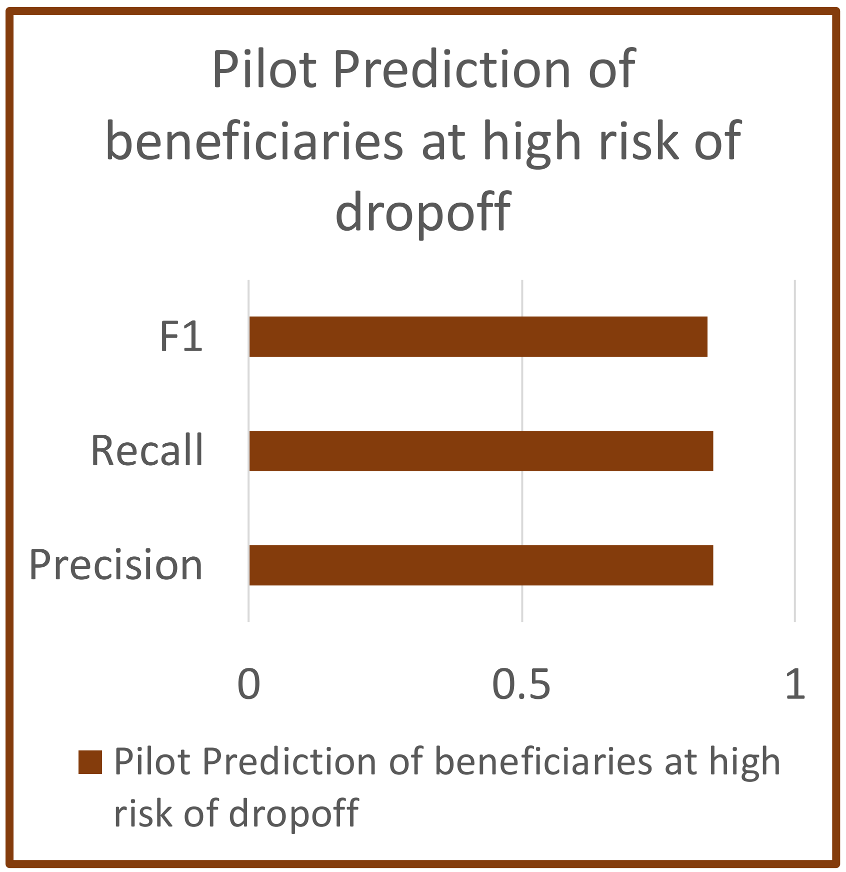

ARMMAN Pilot

- *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

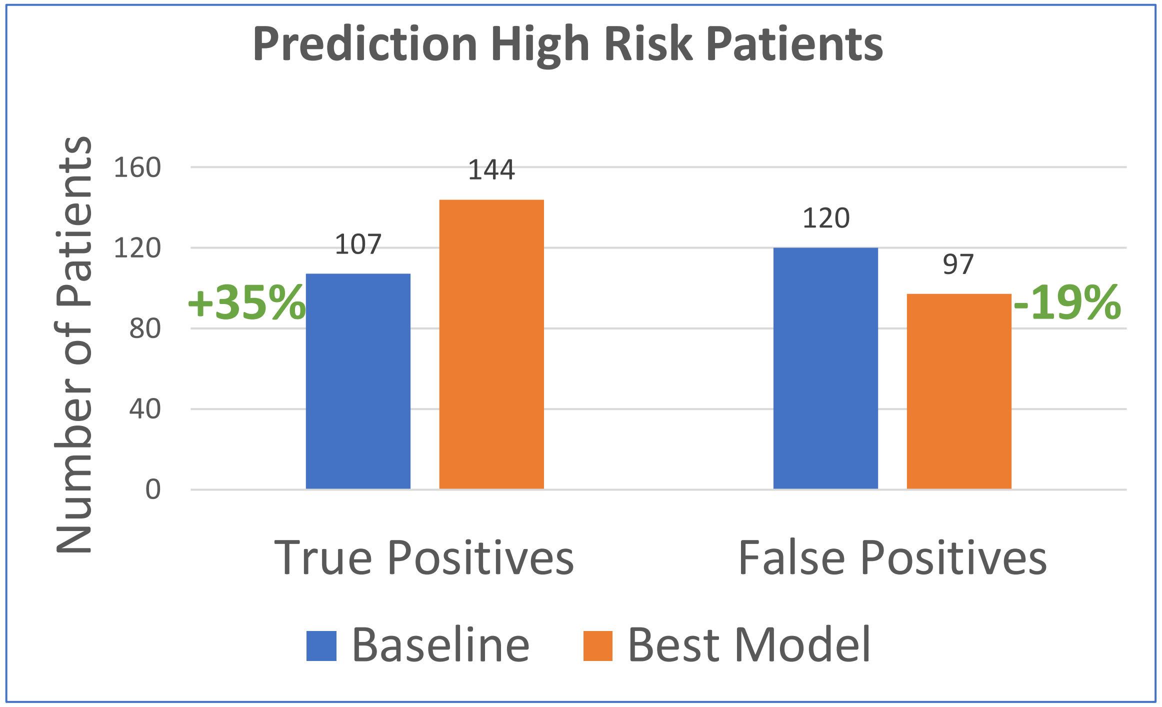

Intervention Reasoning: Active Adherence Monitoring

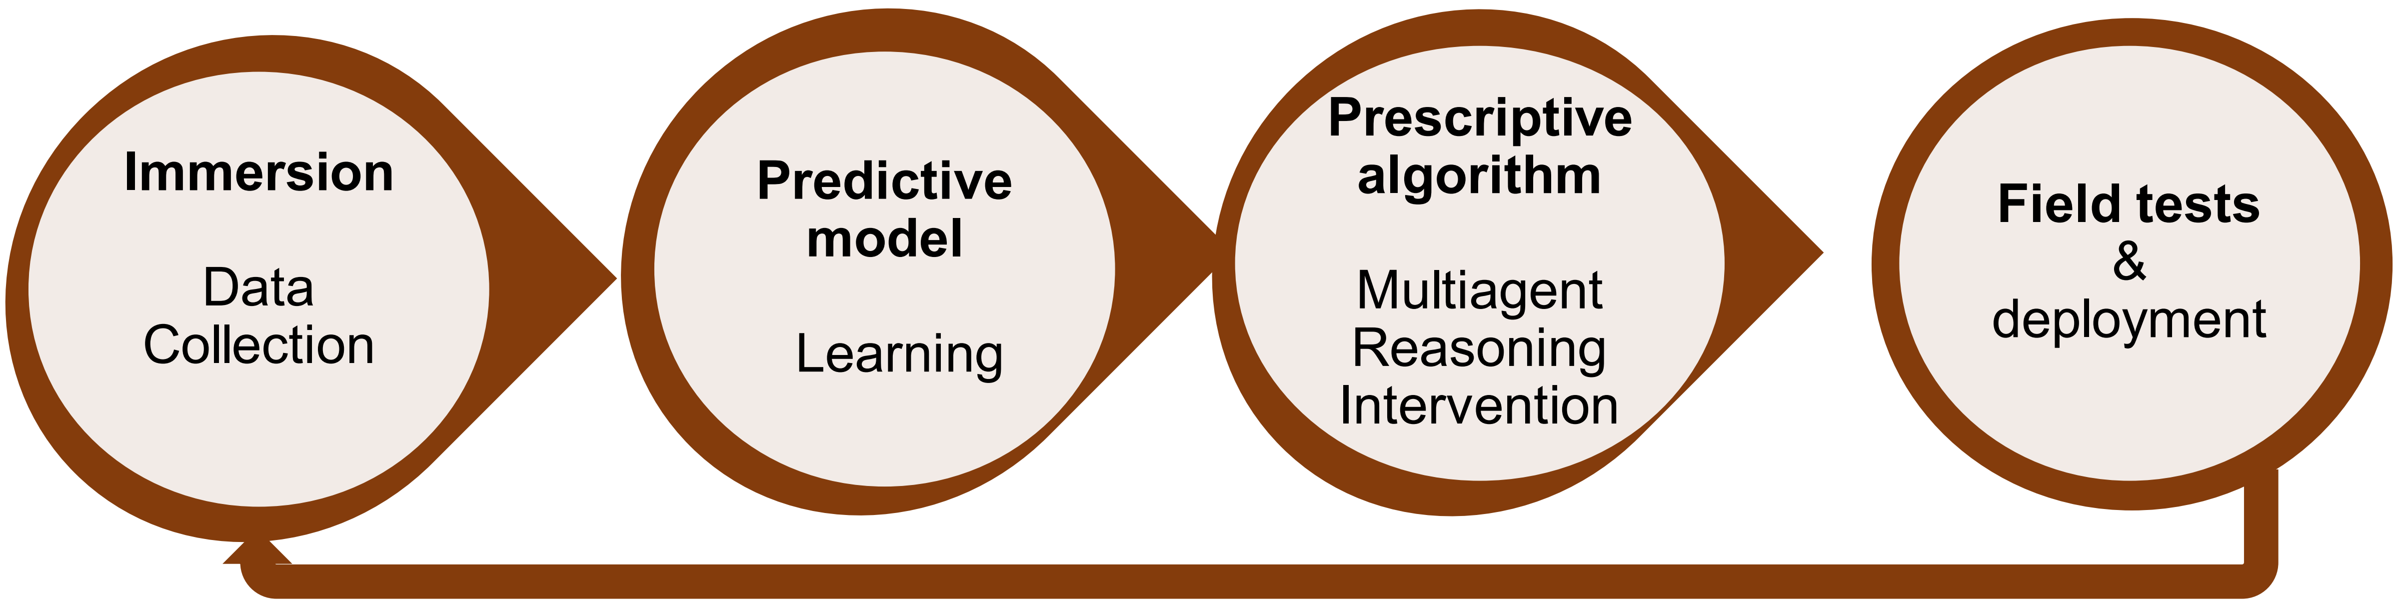

Intervention Scheduling with Scarce Data: Active Adherence Monitoring

(NeurIPS 2020)


Mate


Killian

Health worker intervention

Call patients: Track, improve adherence

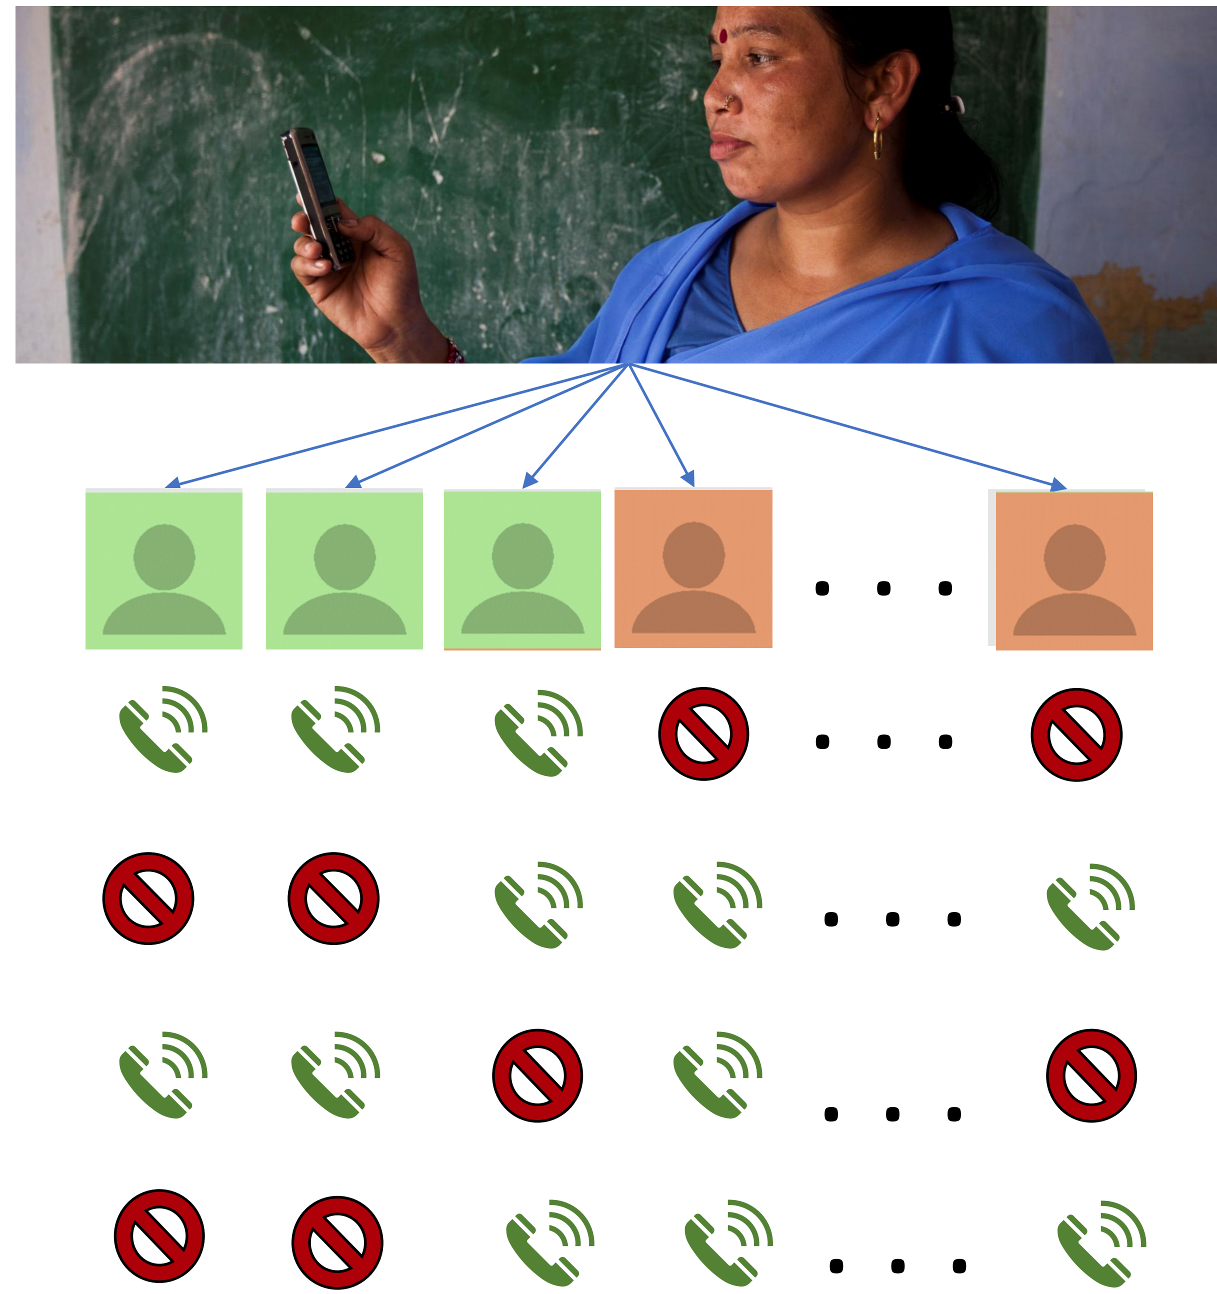

Challenge:

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

Approach:

- *Adherence Restless Bandits*

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

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

(NeurIPS 2020)

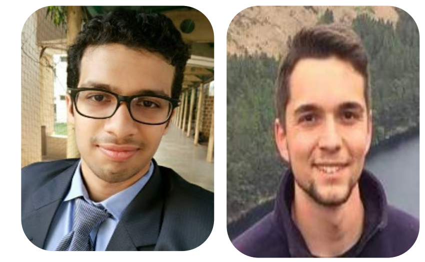

Mate Killian

Restless multiarmed bandits (RMAB)

Adherence RMAB (A-RMAB):

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

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

Patient state may be not observed:

- *Belief state (i.e., probability) of adherence*

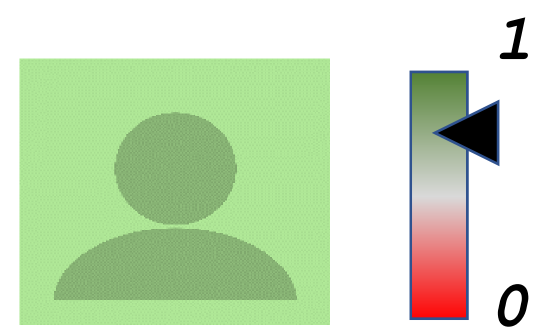

b(adhering,t)

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

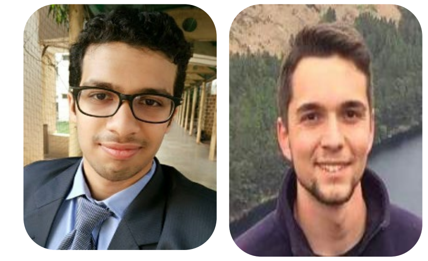

When arm not played

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

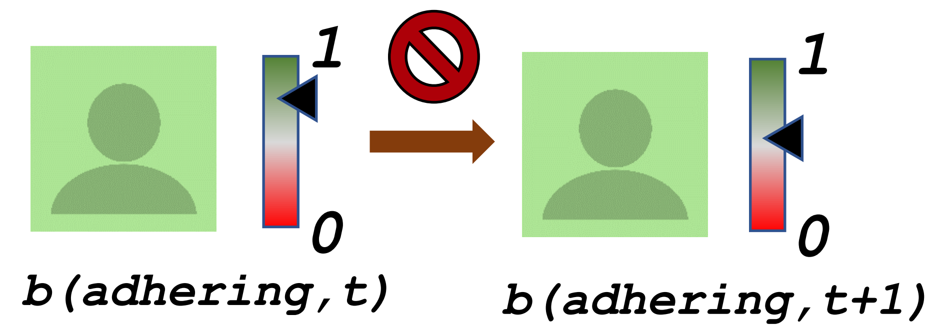

When arm is played

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

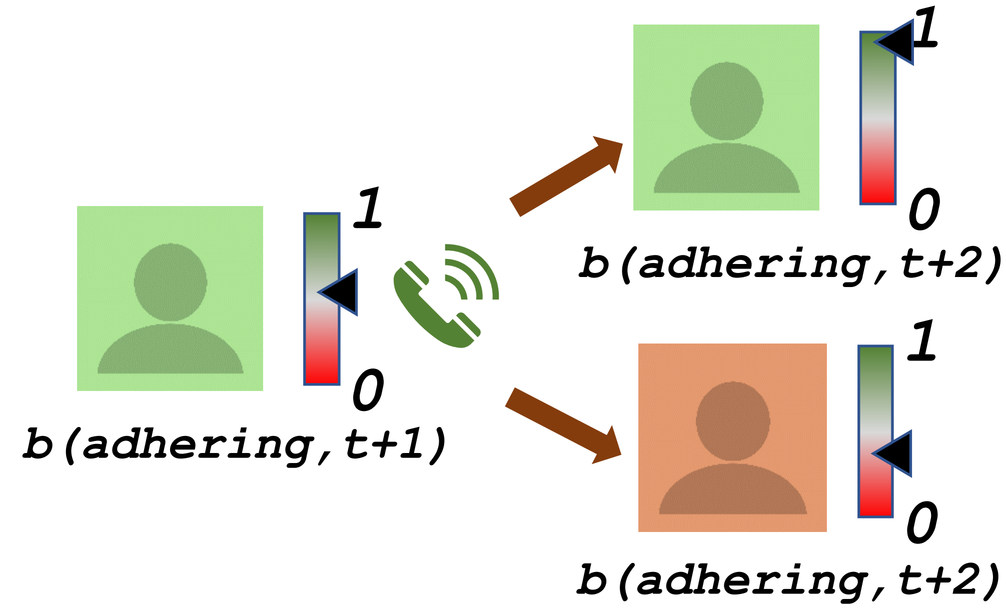

Could convert into a giant POMDP & solve: but inefficient

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

Adherence Restless Bandits(A-RMAB): Whittle Index

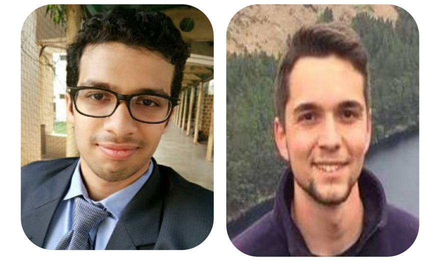
Mate Killian

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

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

- ***Threshold policies: Forward Threshold***

  Call → Belief of adherence below threshold → Call

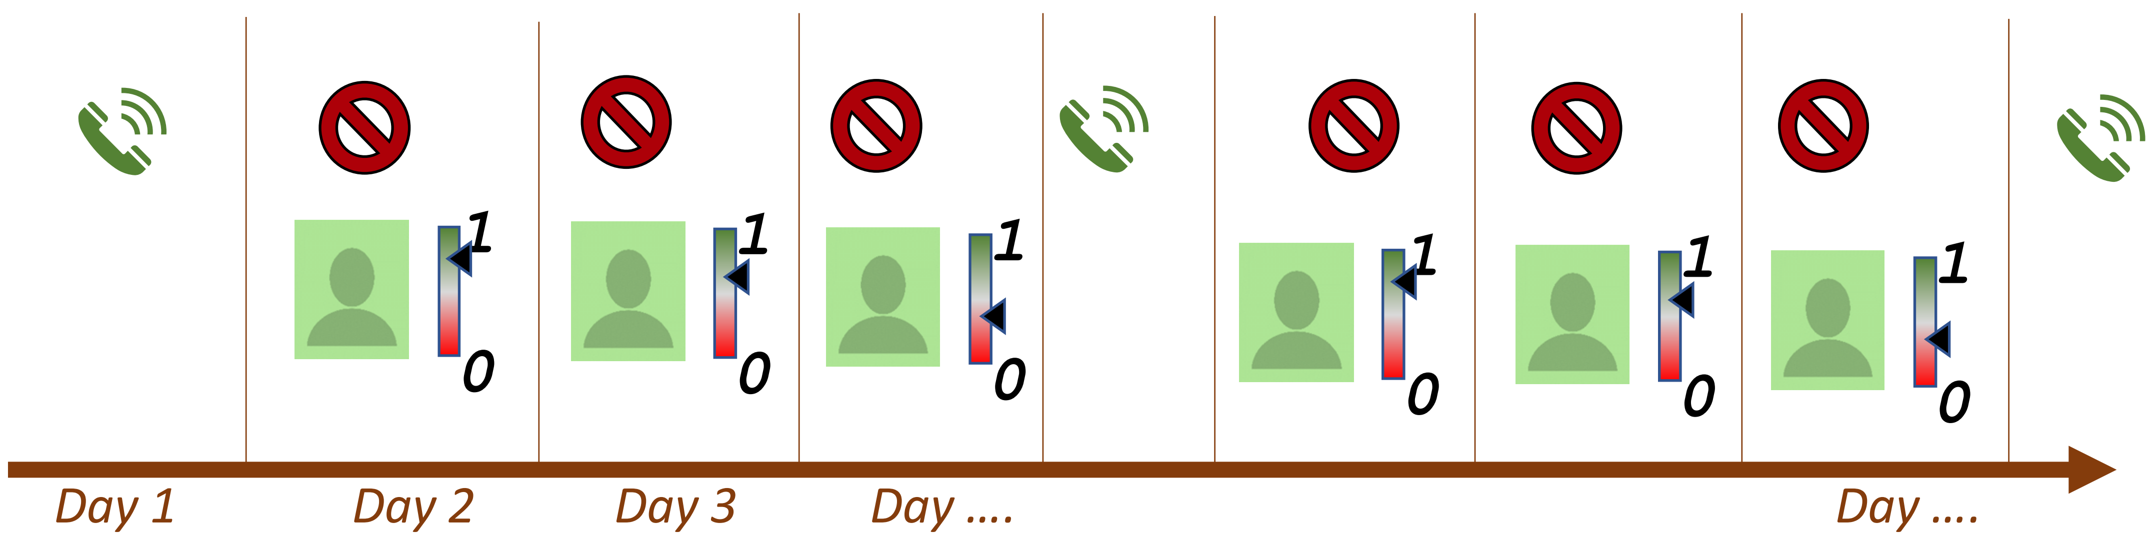

- ***Exploiting threshold policies allow for a fast algorithm***

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

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

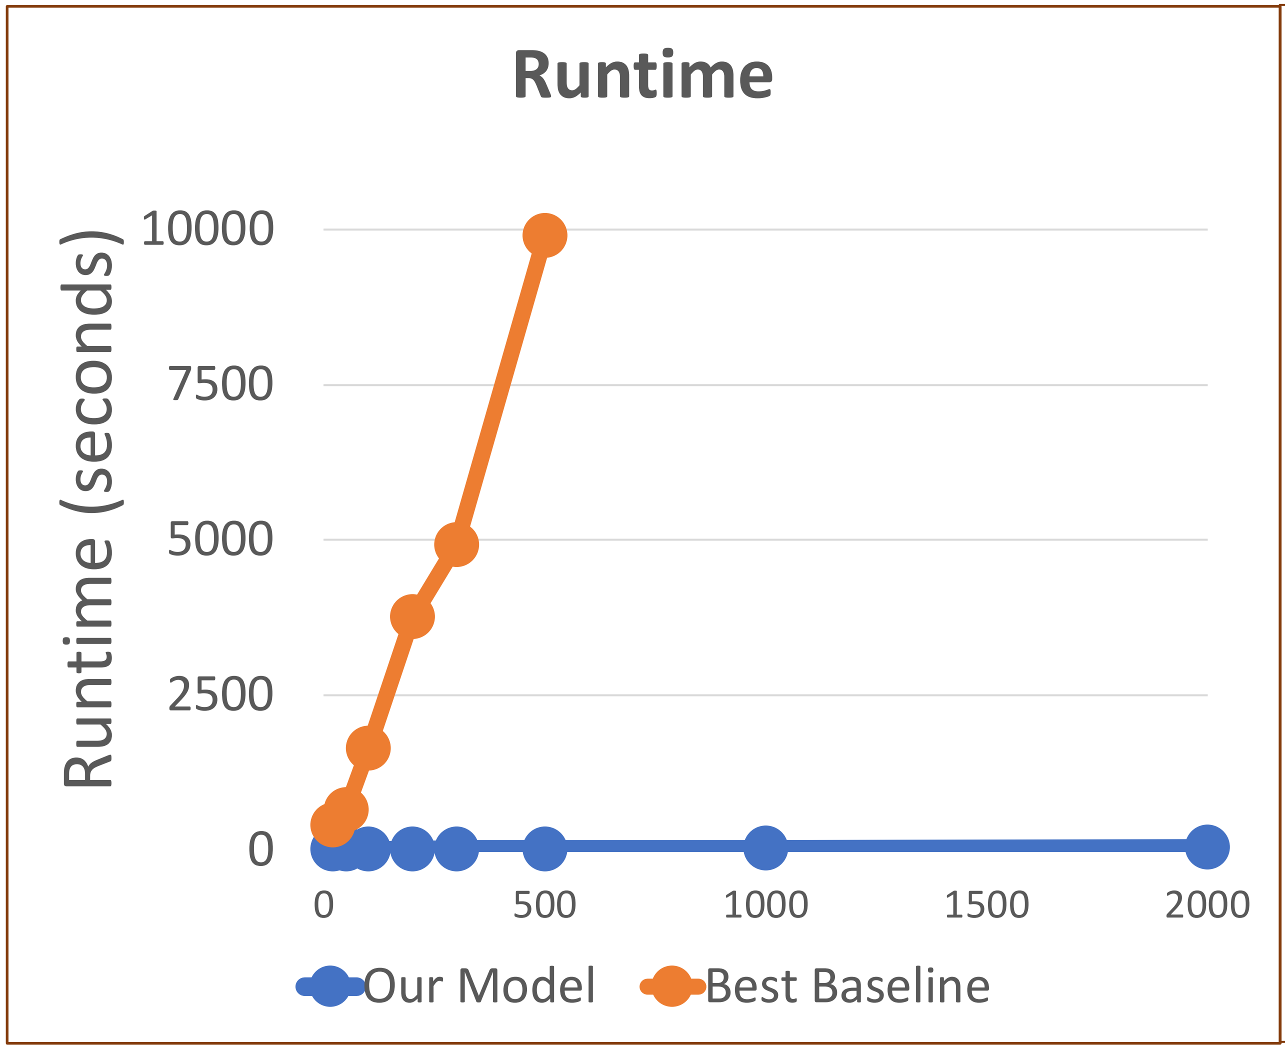

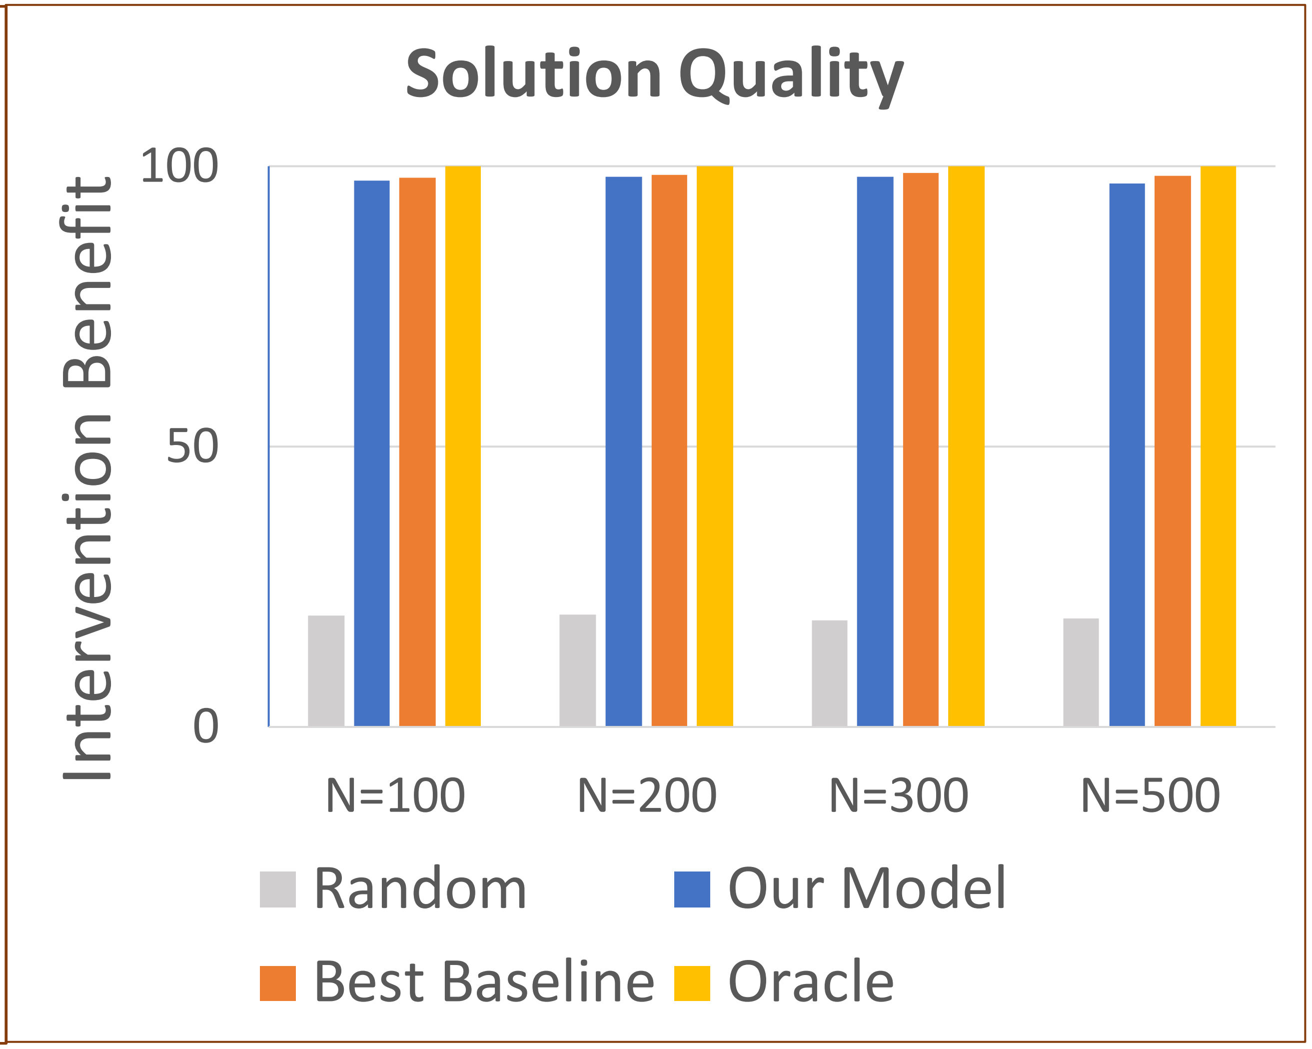

Outline

Public Health

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

Conservation

COVID-19: Agent-based Simulation Model

(Proceedings National Academy of Sciences, 2020)


Wilder

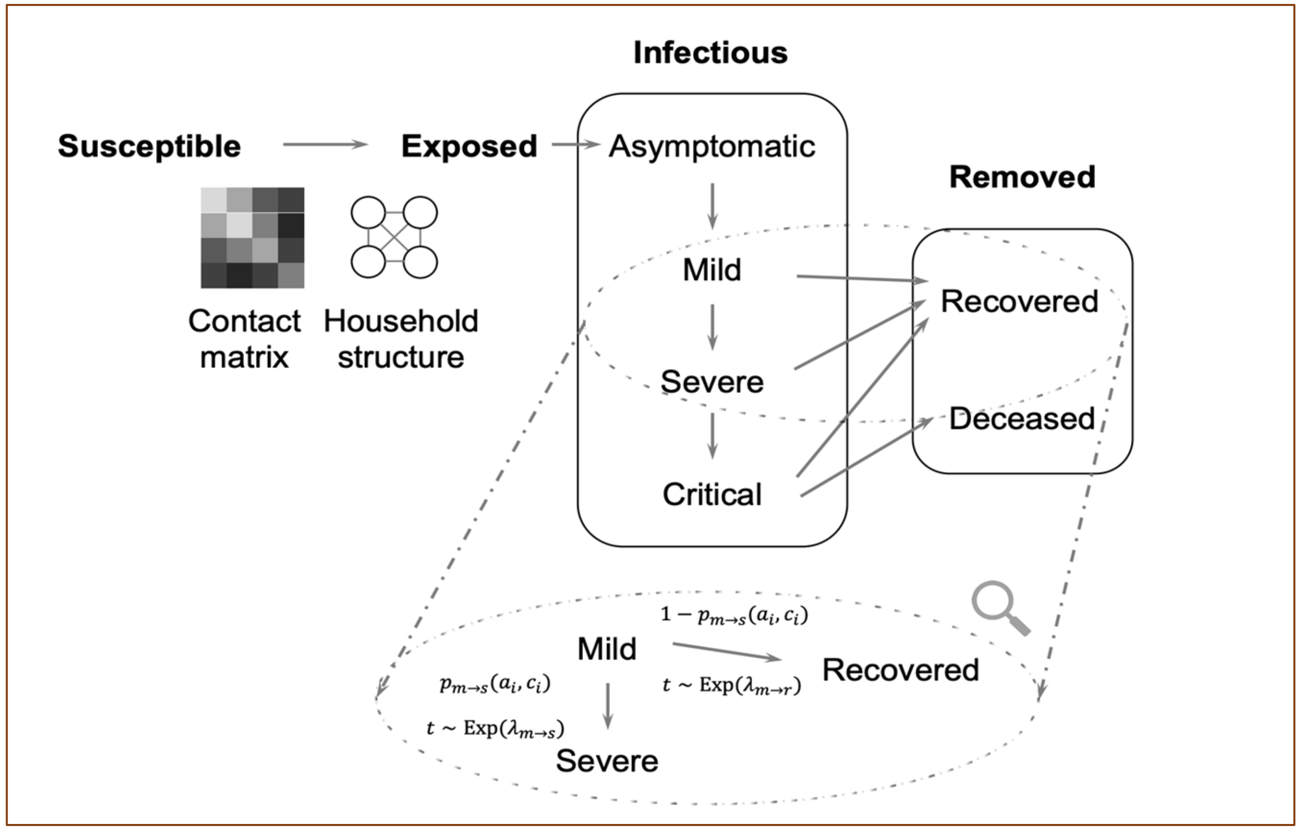

Agent-based model:

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

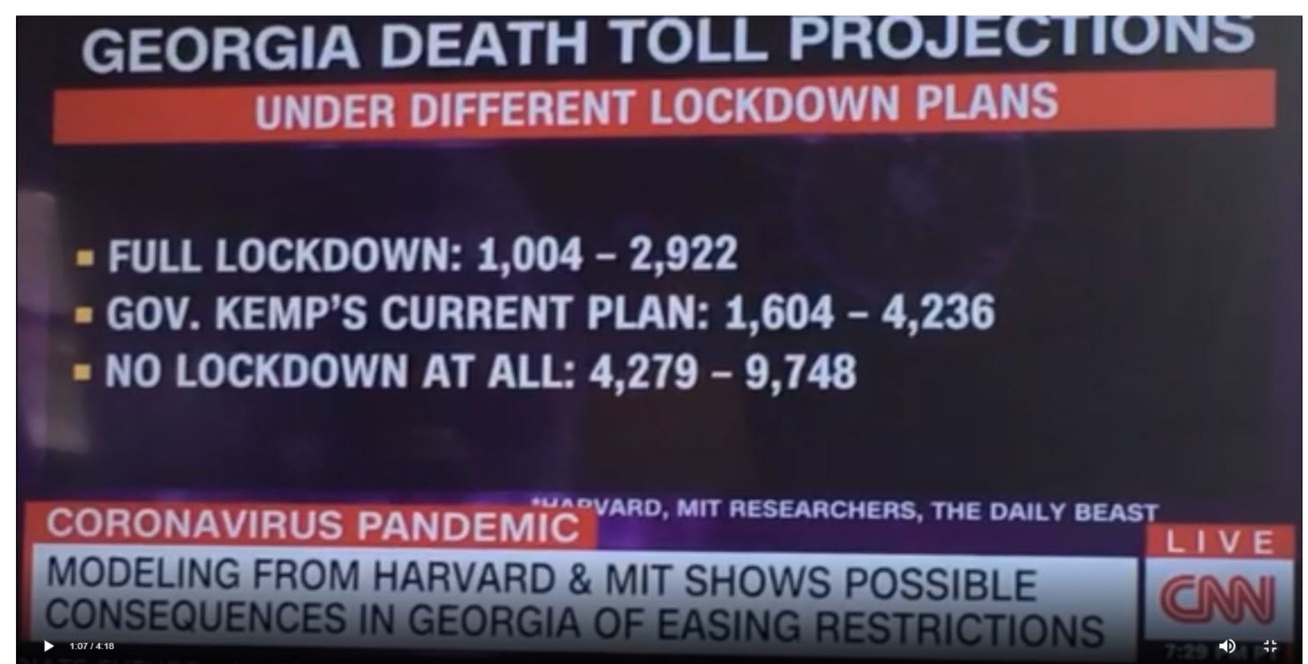

COVID Testing Policy: Accuracy vs Ease
(Science Advances, 2020)


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**

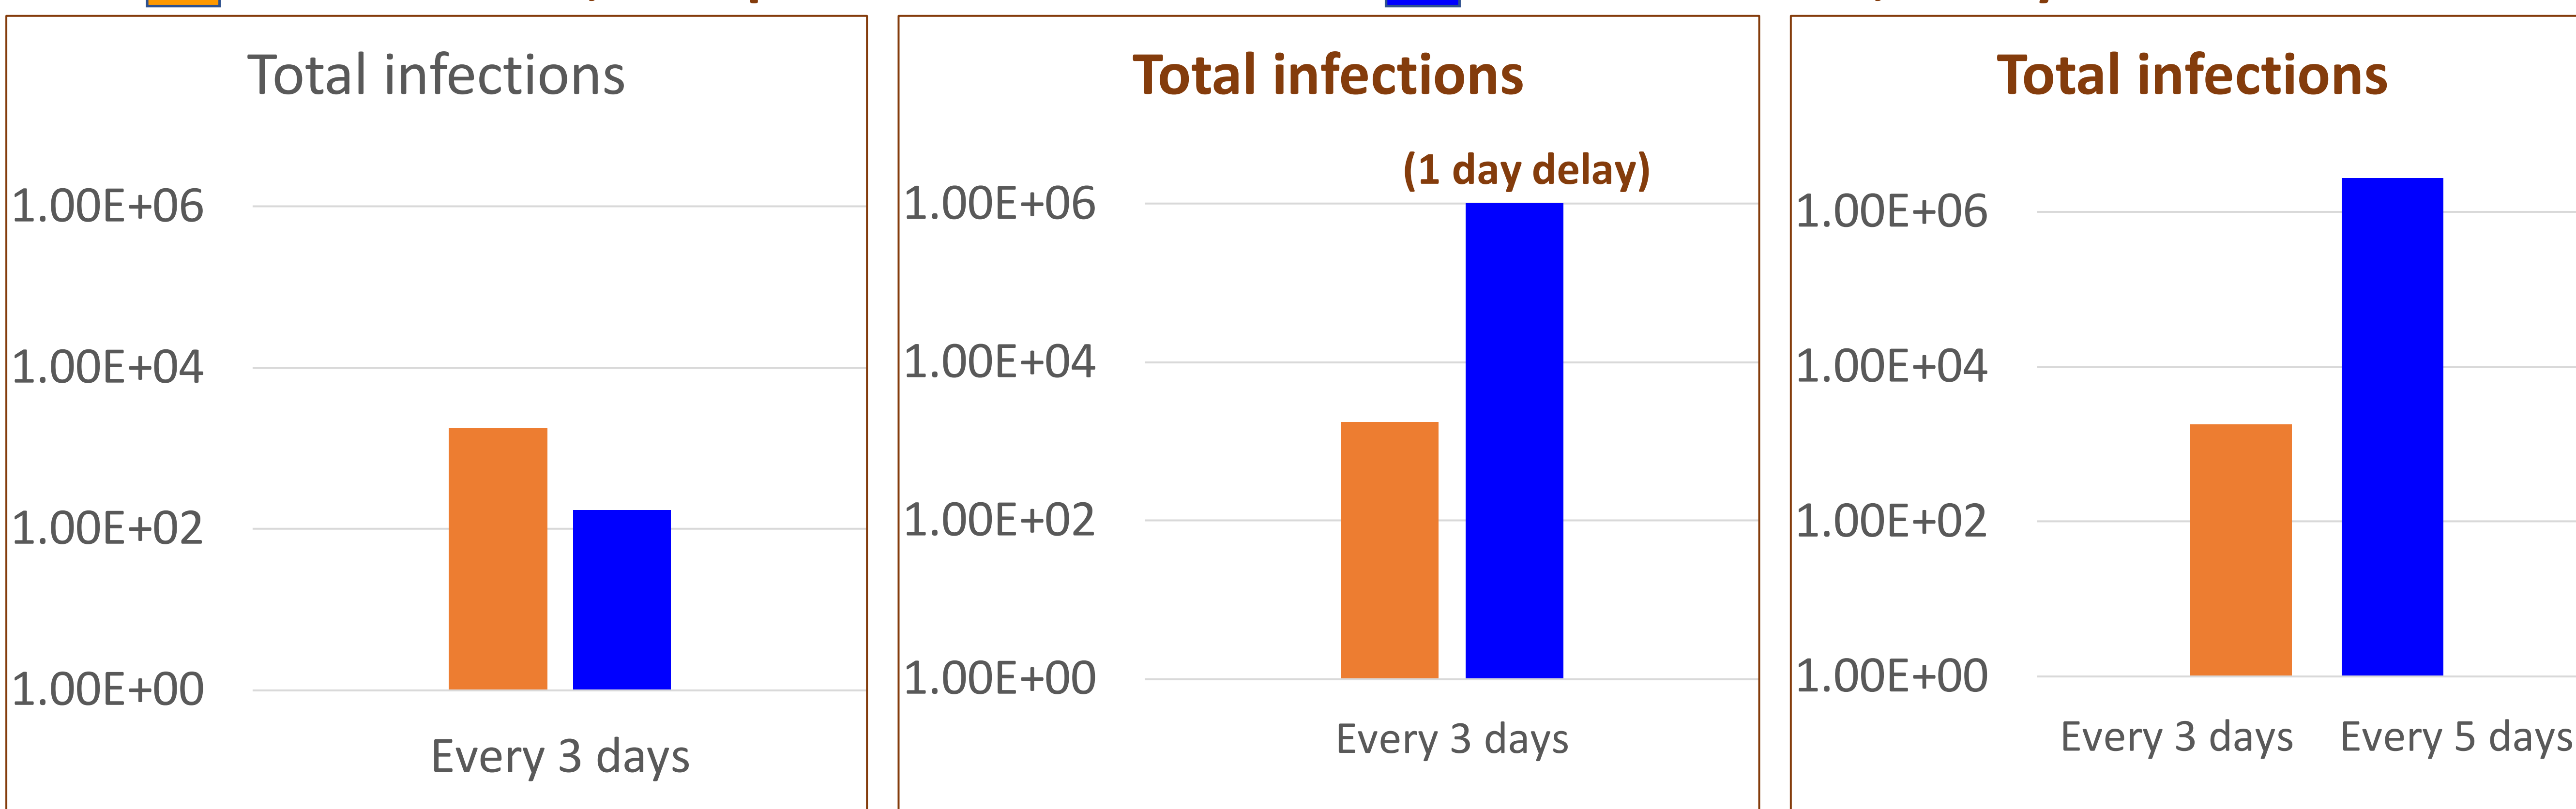

COVID Testing Policy: Impact

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

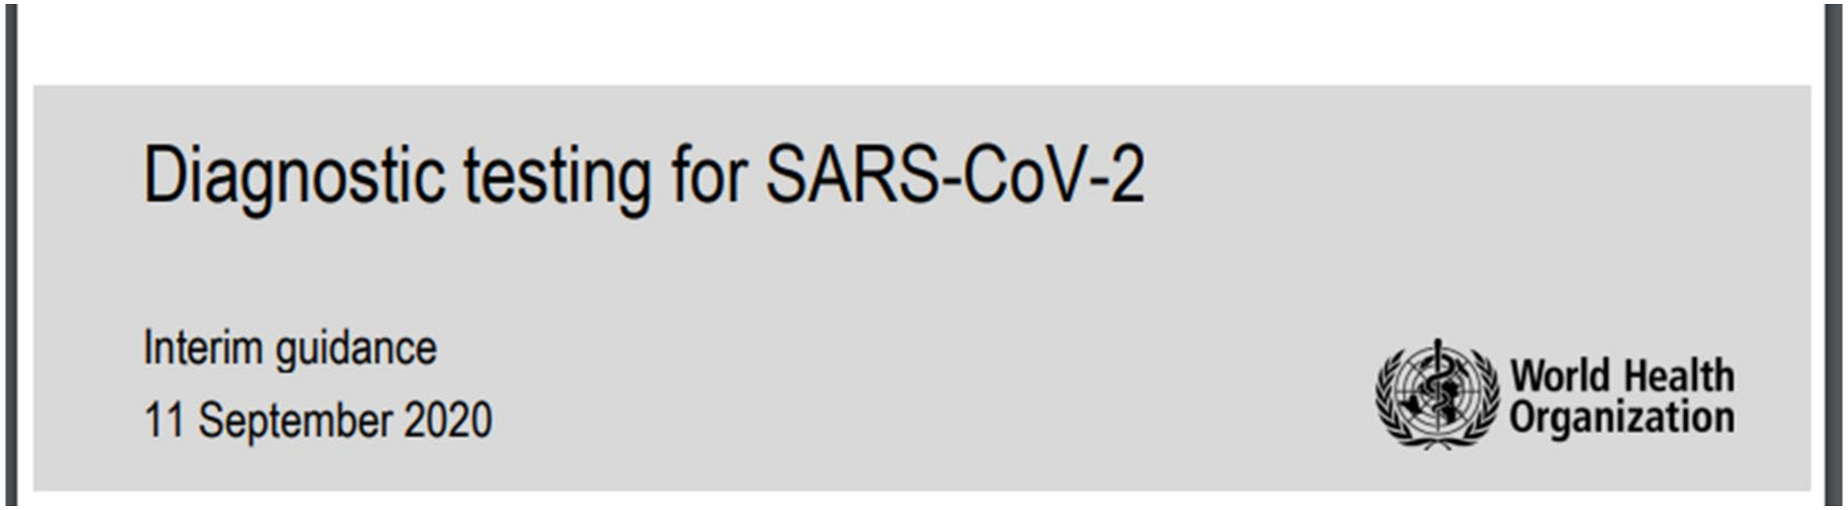
Diagnostic testing for SARS-CoV-2

Interim guidance

11 September 2020

World Health Organization

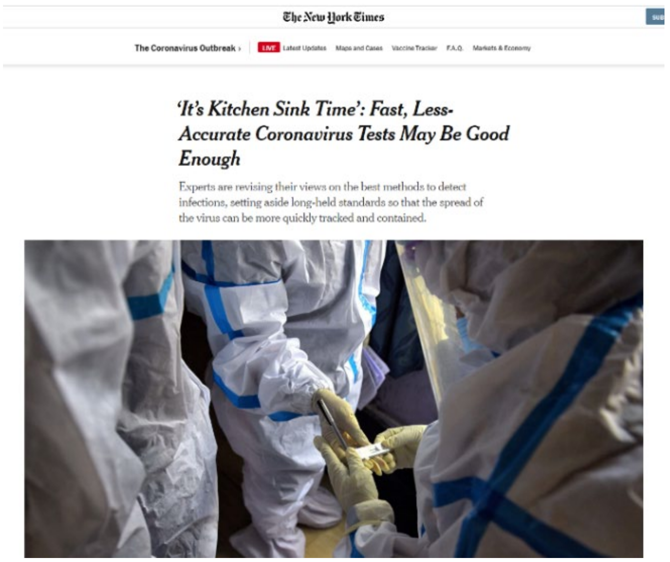
The New York Times

The Coronavirus Outbreak › LIVE Latest Updates Maps and Cases Vaccine Tracker F.A.Q. Markets & Economy

'It's Kitchen Sink Time': Fast, Less-Accurate Coronavirus Tests May Be Good Enough

Experts are revising their views on the best methods to detect infections, setting aside long-held standards so that the spread of the virus can be more quickly tracked and contained.


TIME

AMERICA MUST CHANGE RACISM AND COVID-19 INEQUALITY BLACK, DISABLED AND AT RISK NEWSLETTER

IDEAS • COVID-19

America Needs to Radically Rethink Our COVID-19 Testing Approach

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)

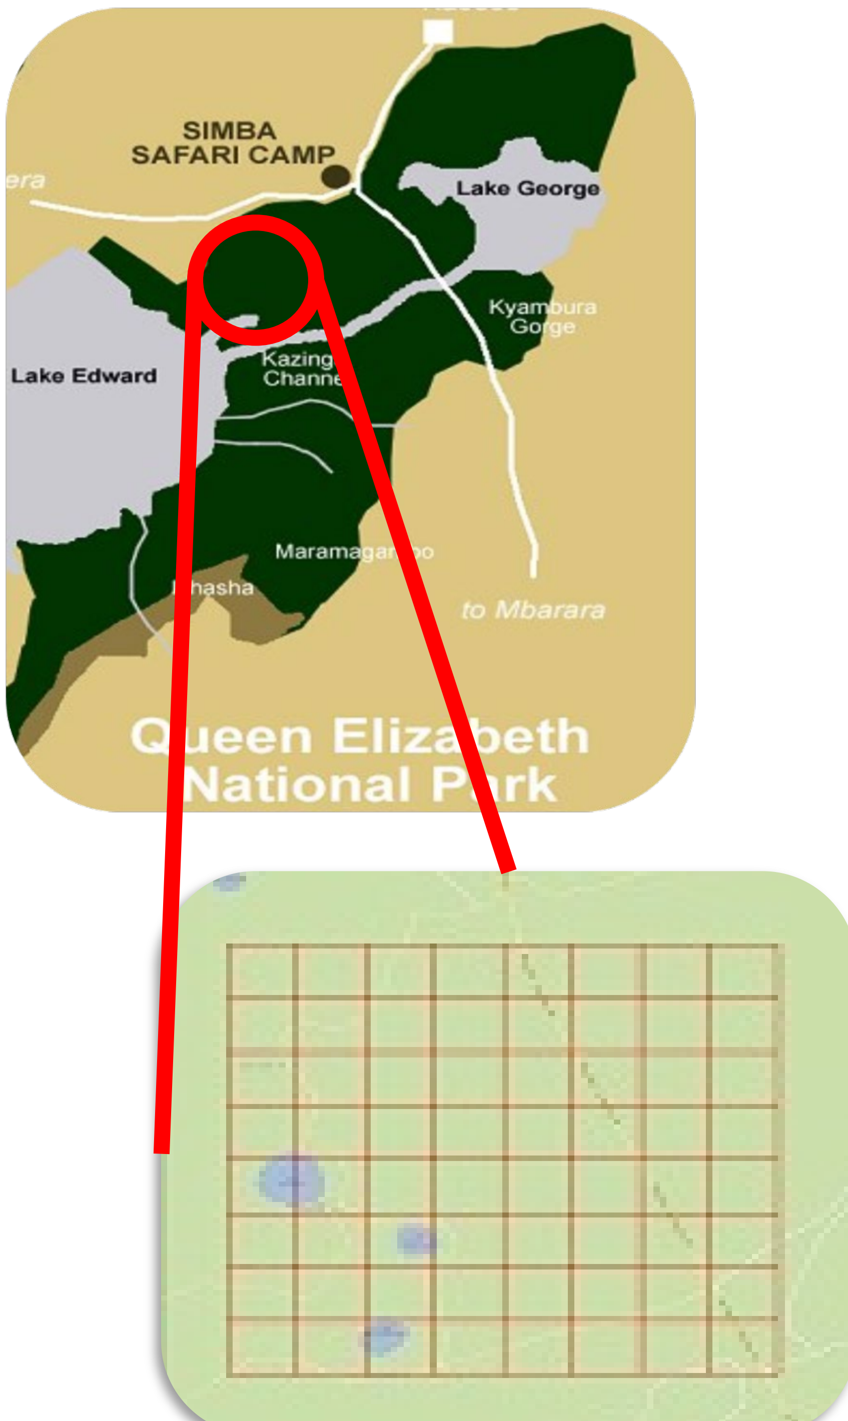

- ***Stackelberg security games (SSG)***
- ***With boundedly rational poachers***
- ***Learn adversary response model at targets "i"***

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

Learning Adversary Response Model: Uncertainty in Observations


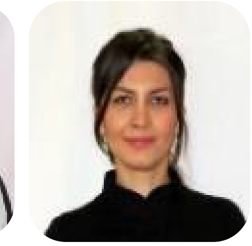

Nguyen Gholami

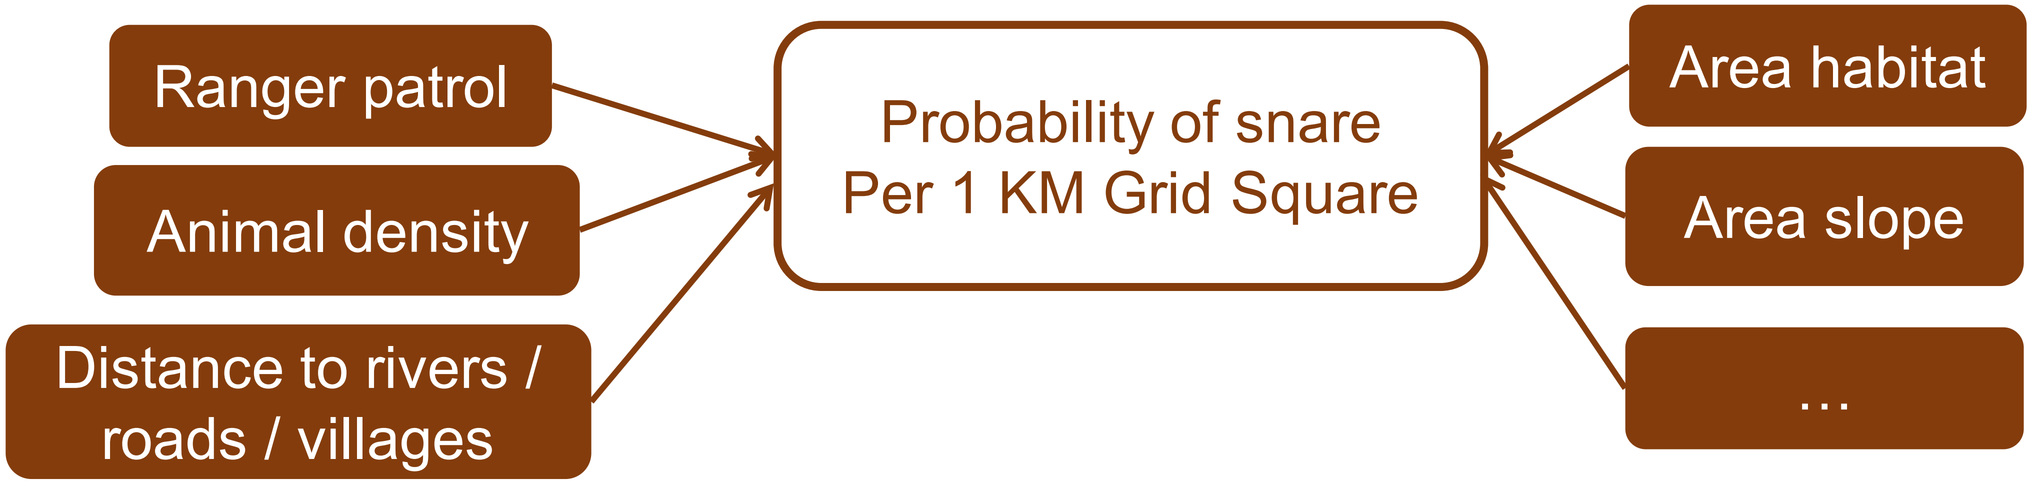

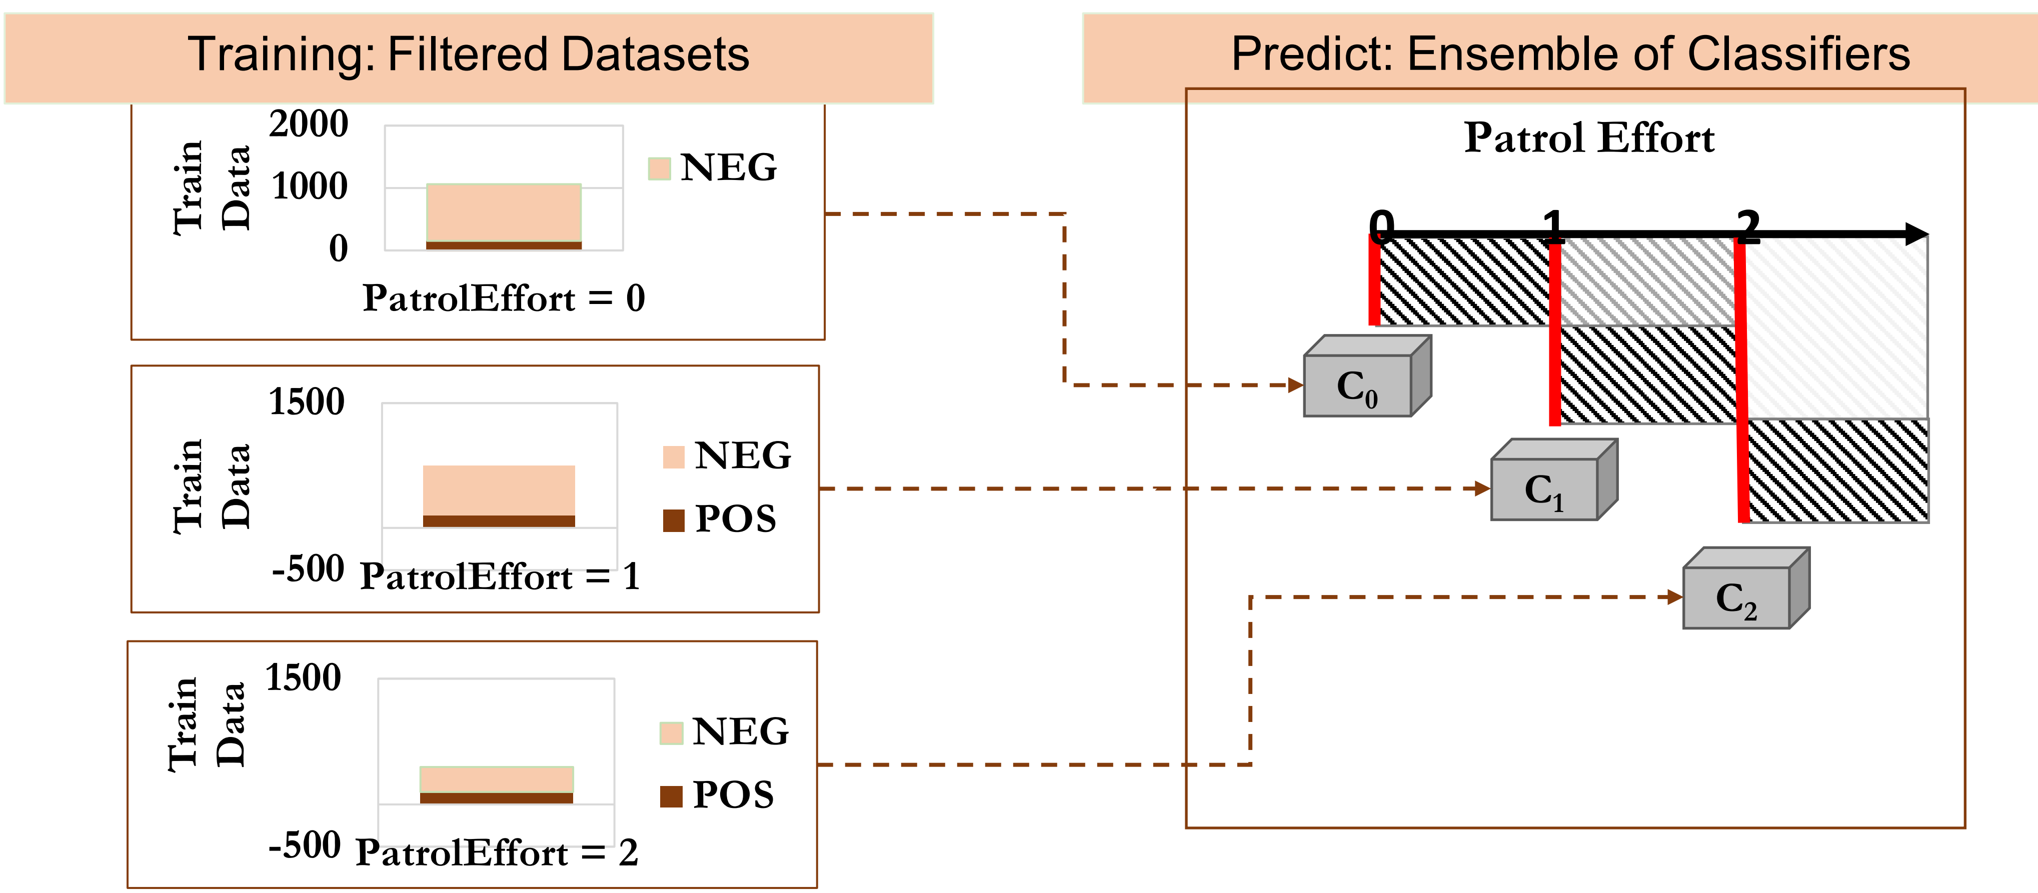

PAWS: First Pilot in the Field

(AAMAS 2017)

Ford

Gholami

- Two 9-sq.km areas, infrequent patrols

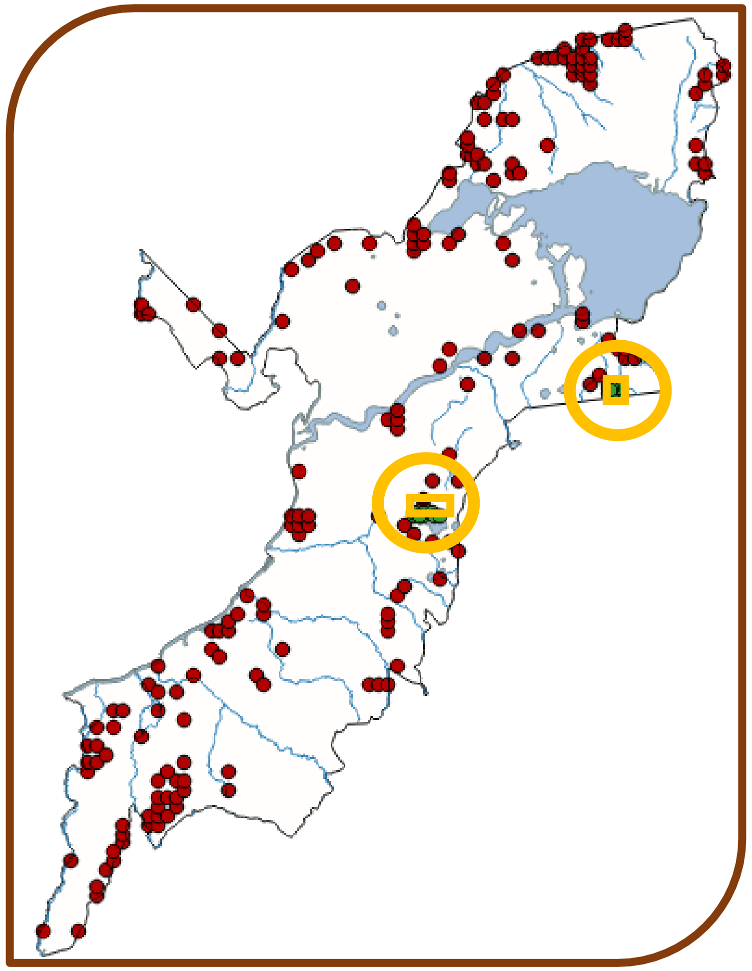

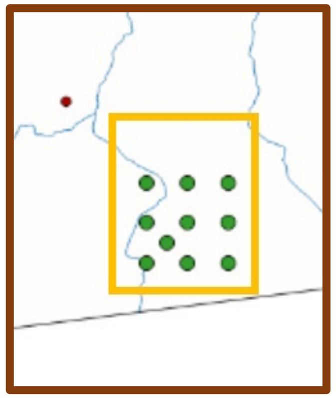

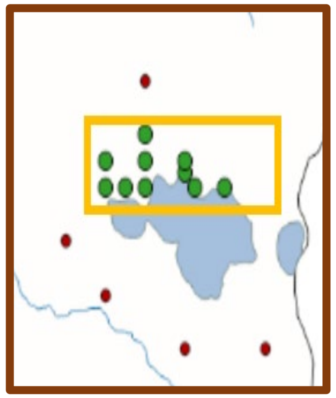

- Poached elephant
- 1 elephant snare roll
- 10 Antelope snares

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

(ECML PKDD 2017, ICDE 2020)


Xu


Gholami

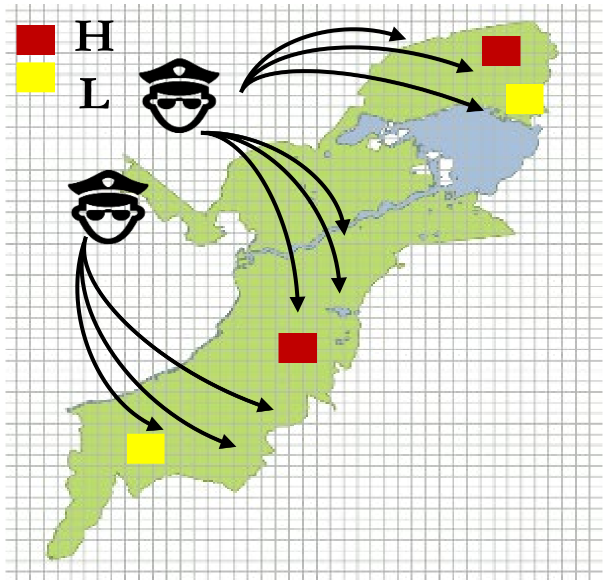

Queen Elizabeth National Park

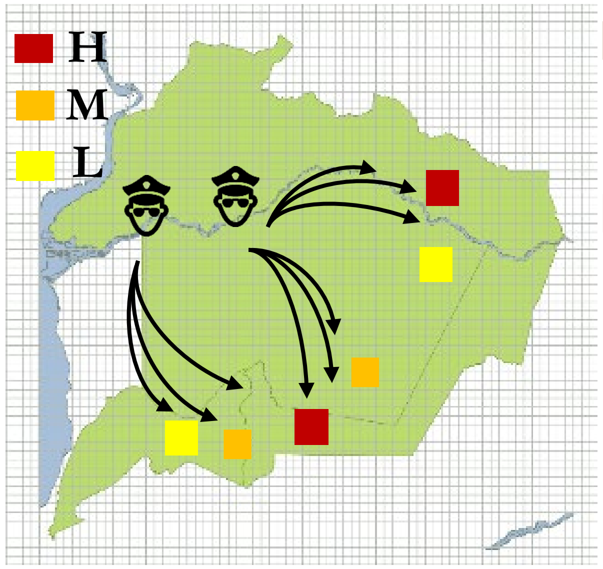

Murchison Falls National Park

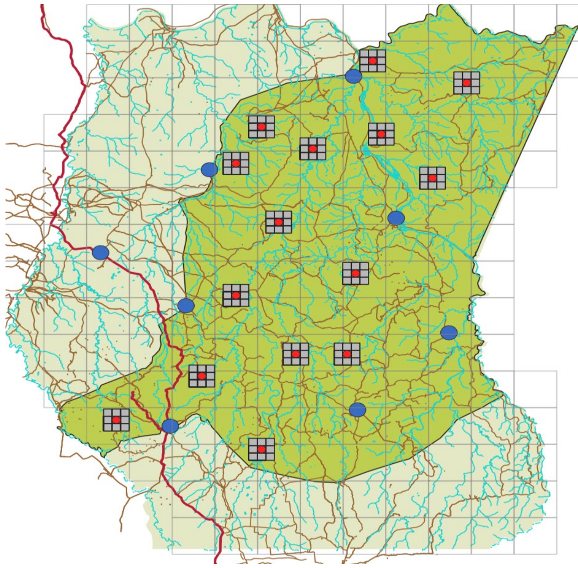

Srepok Wildlife Sanctuary

Snares per patrolled sq. KM

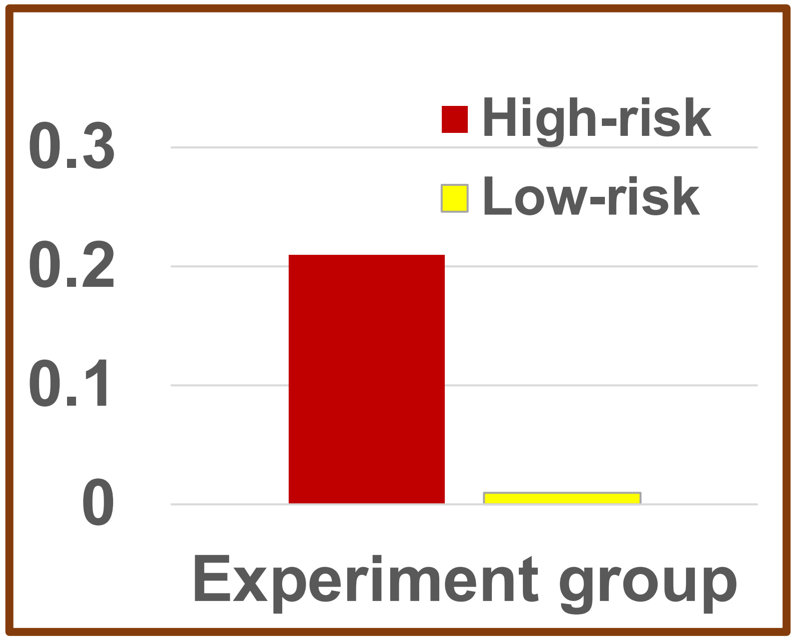

Snares per patrolled sq. KM

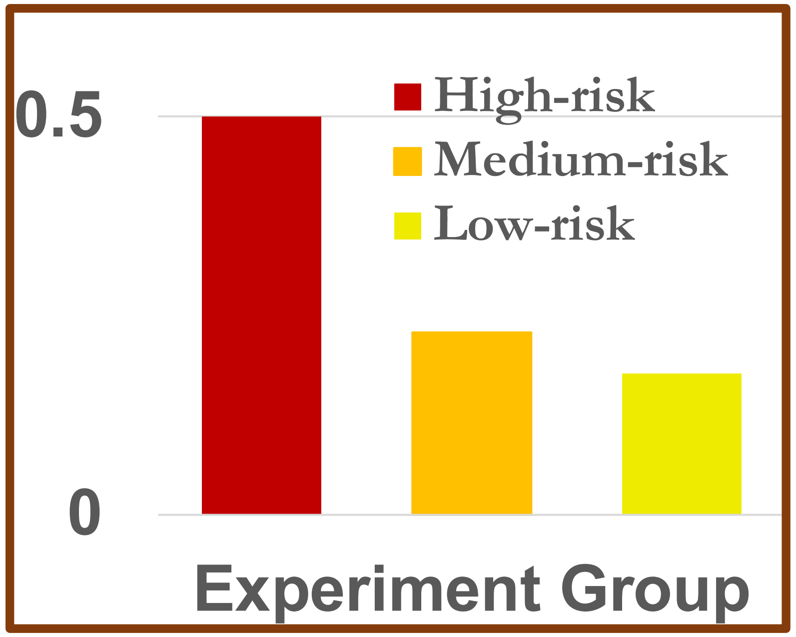

Snares per patrolled sq. KM

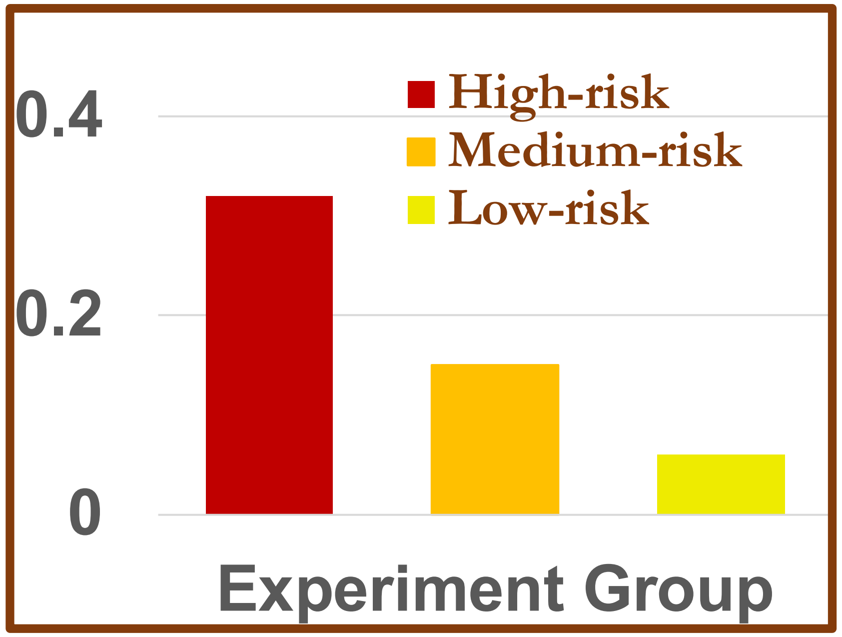

PAWS Real-world Deployment
Cambodia: Srepok Wildlife Sanctuary

(ICDE 2020)

Xu

- *521 snares/month our tests*

vs

- *101 snares/month 2018*

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

Perrault

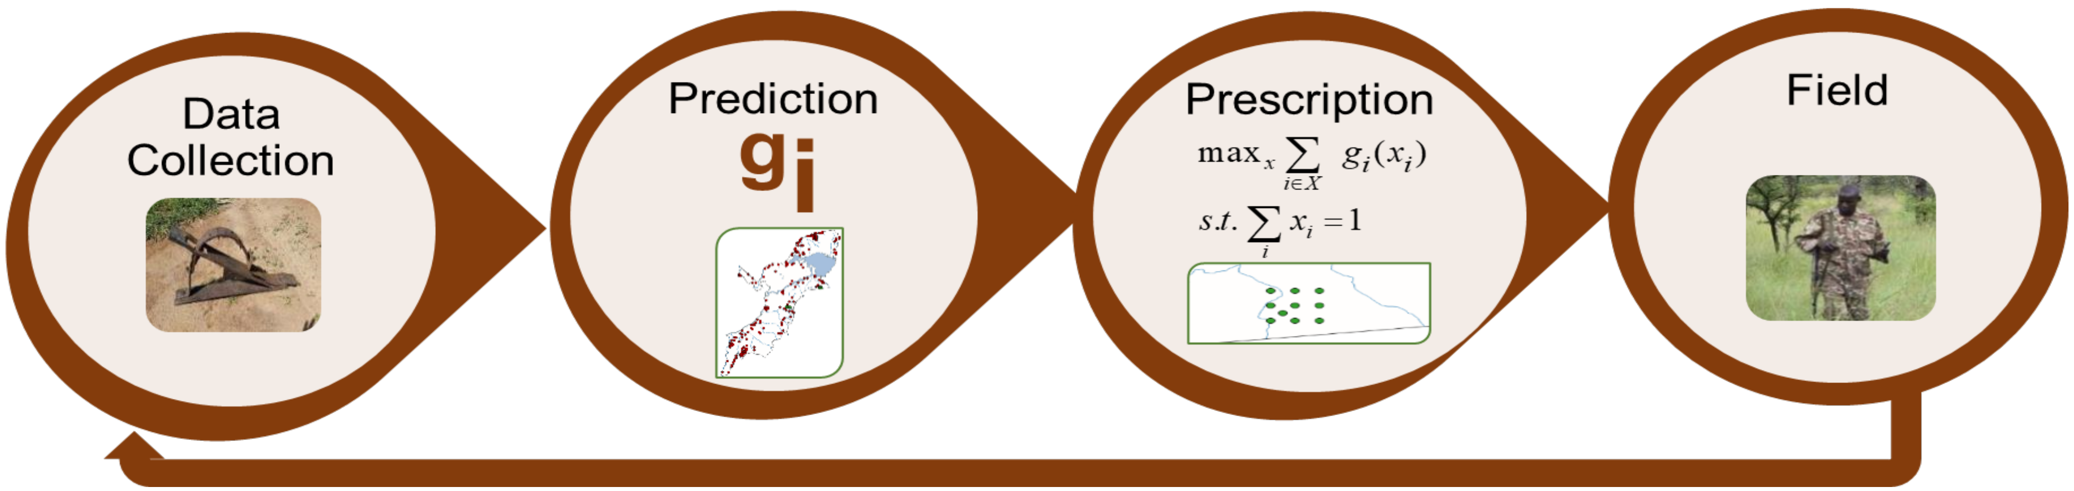

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

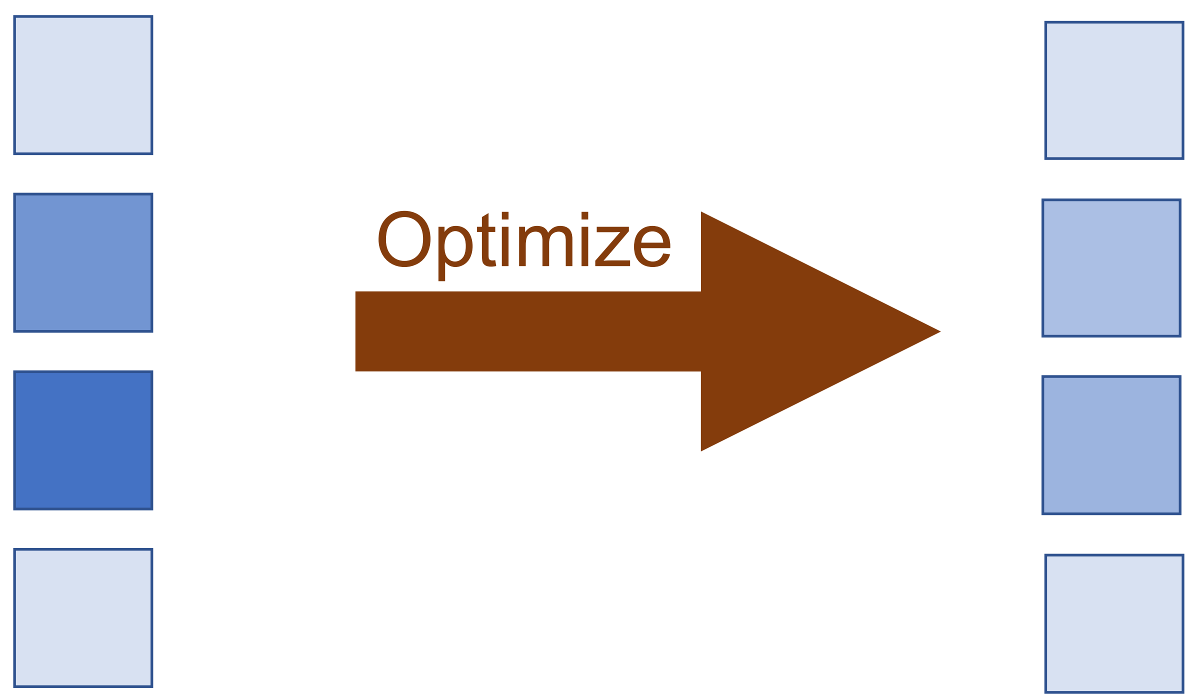

Minimize $\sum_{\{i \in T\}} q_{\text{empirical}} \log \hat{q}$

Game-Focused Learning: End-to-End Method Builds on Decision-focused Learning

(AAAI 2019, AAAI 2020)

Perrault Wilder

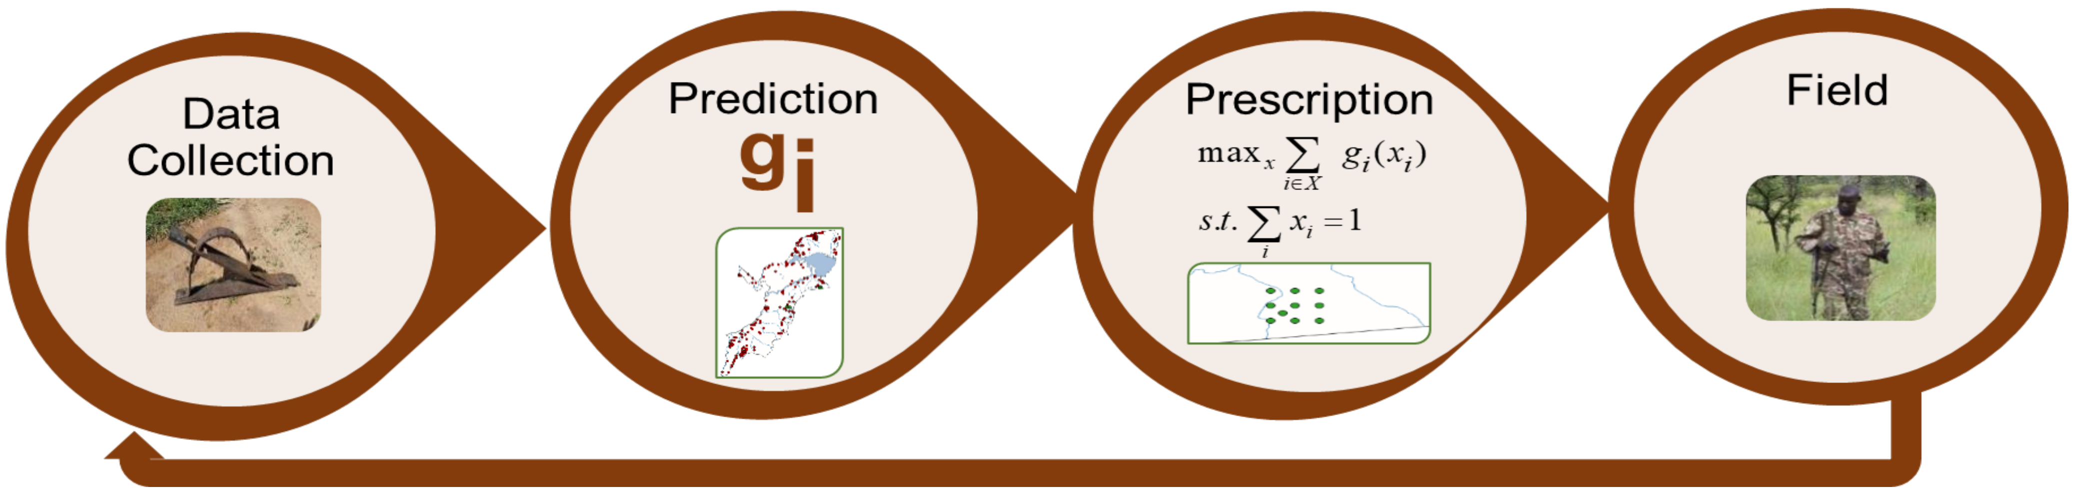

Maximize
defender expected *utility*

Plan patrol
Coverage

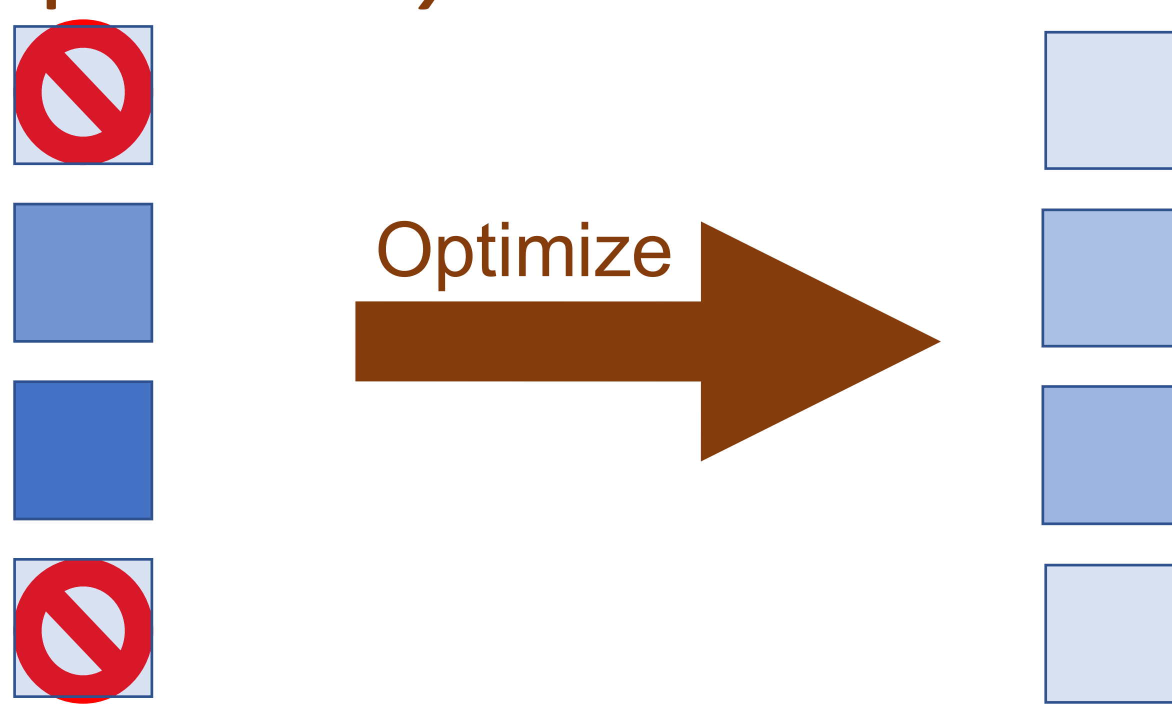

Maximize defender's expected utility $\sum \left(1 - p_i(\hat{q})\right) q_{\text{empirical}}$

Game-Focused Learning: Comparison to Two-Stage

Perrault

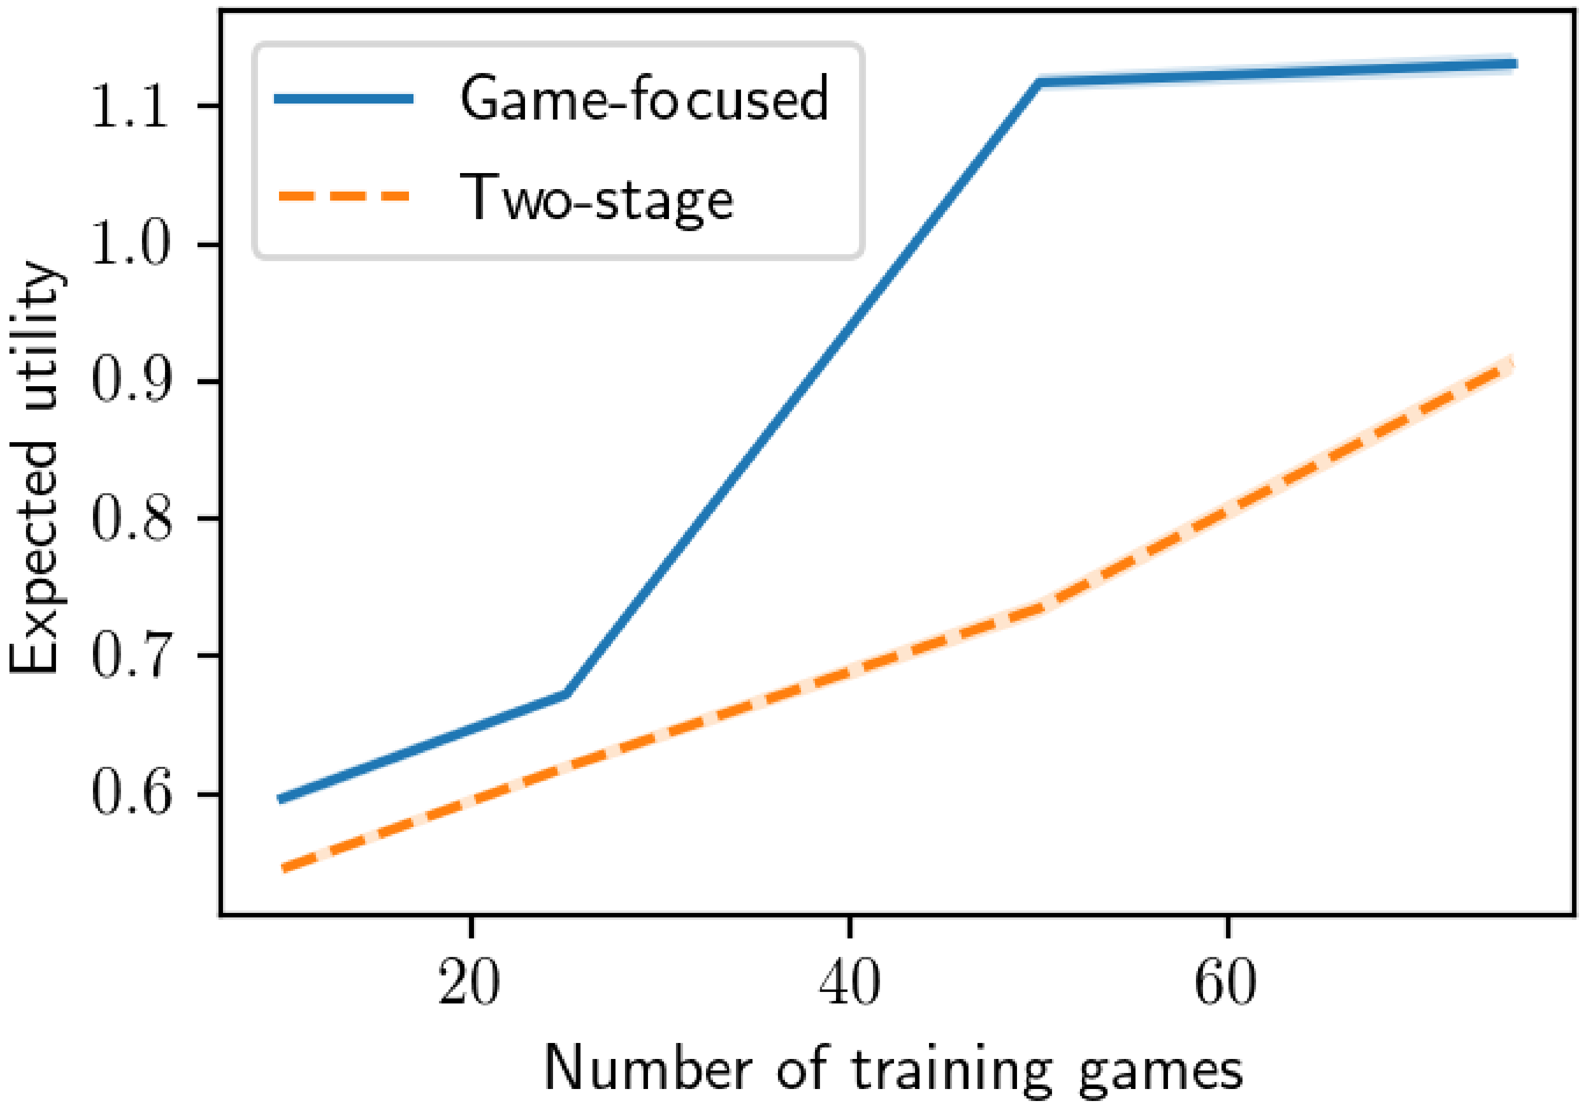

PAWS Goes Global


Xu

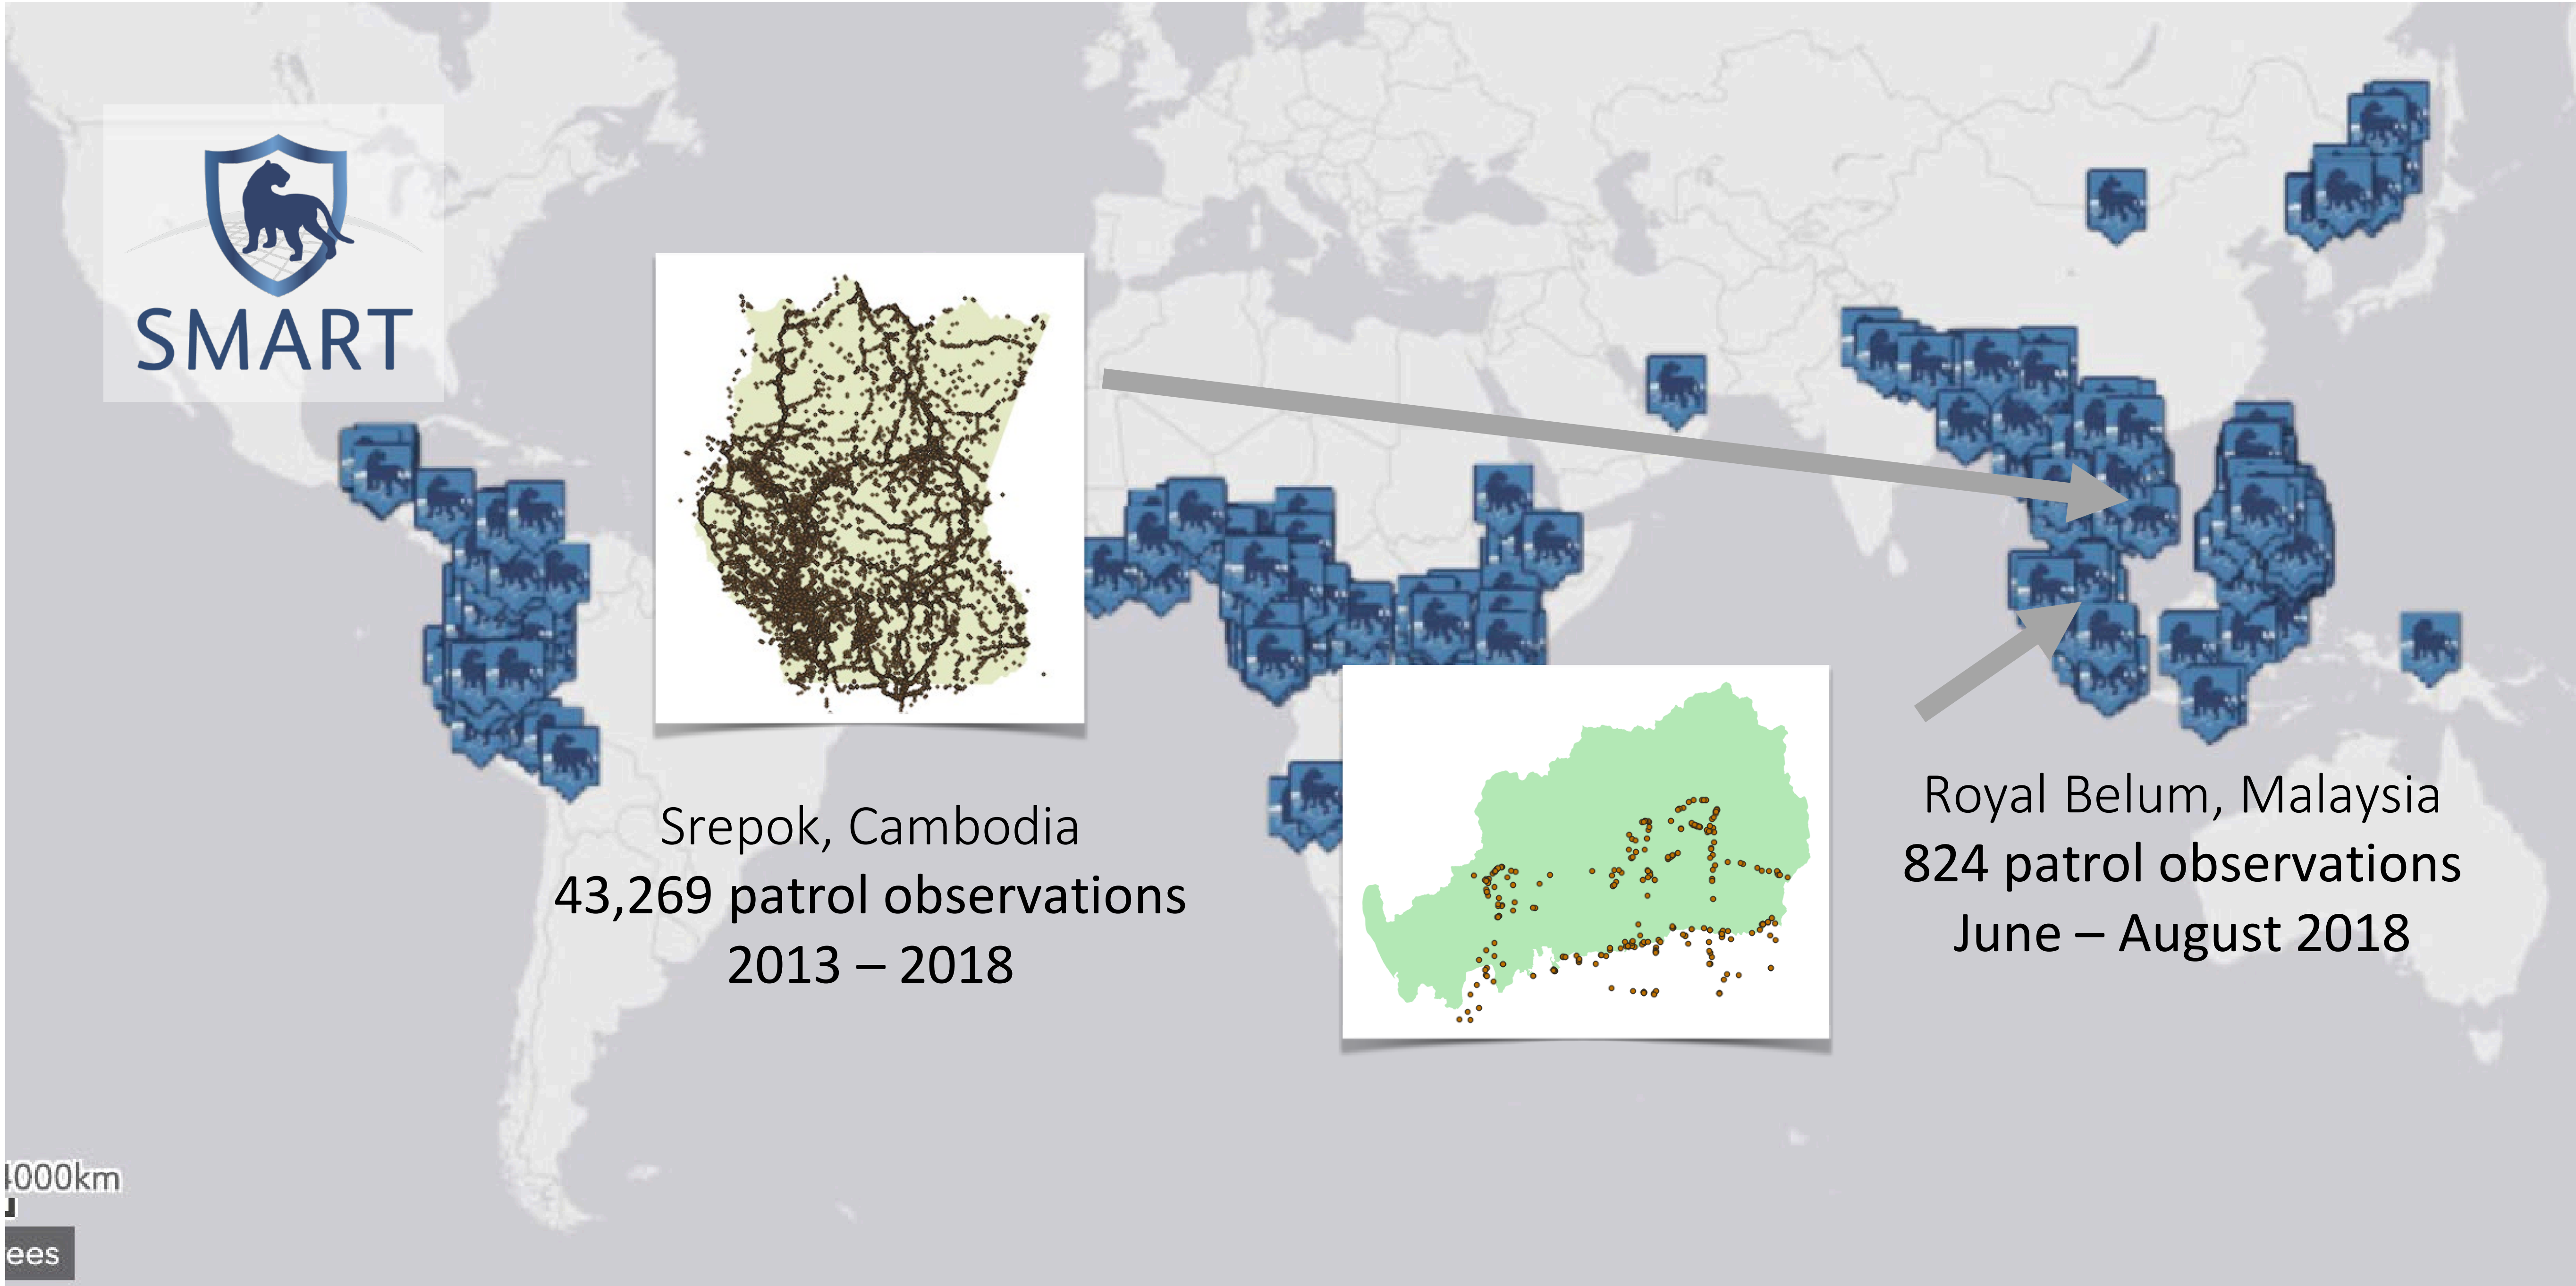

The Dual Mandate

Xu

exploitation

Data-rich parks: build predictive models to plan patrols

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

exploration

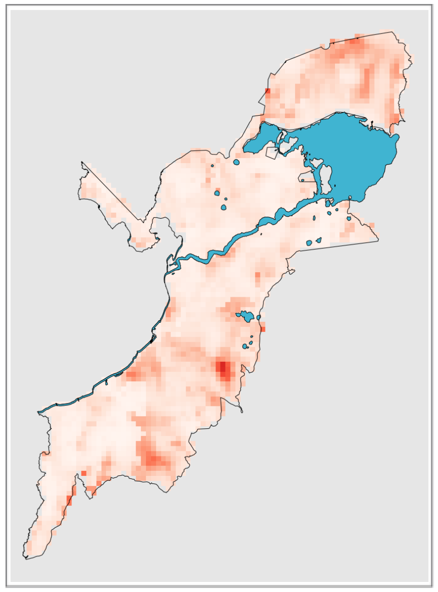

[Gholami et al., AAMAS-18, Xu et al., ICDE-20]

LIZARD: Multiarmed Bandit
LIpschitZ Arms with Reward Decomposability

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 Horizon
- Stochastic adversary, who places snares at targets
- Patrolling algorithm: Specify patrol effort in each target up to budget B
- Reduce regret wrt OPT, optimal patrol effort, for capturing snares

Lizard exploits decomposability, smoothness, monotonicity

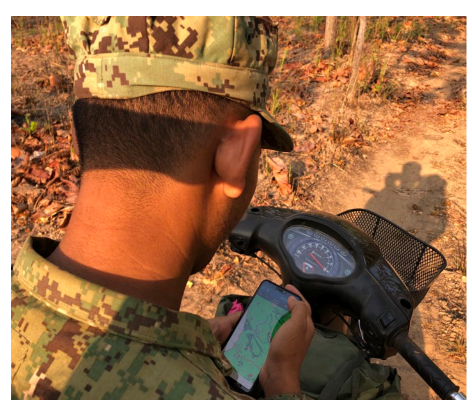

Green Security Games: Integrating Real-Time "SPOT" Information

(IAAI 2018)

Bondi

Air Shepherd
The Lindbergh Foundation

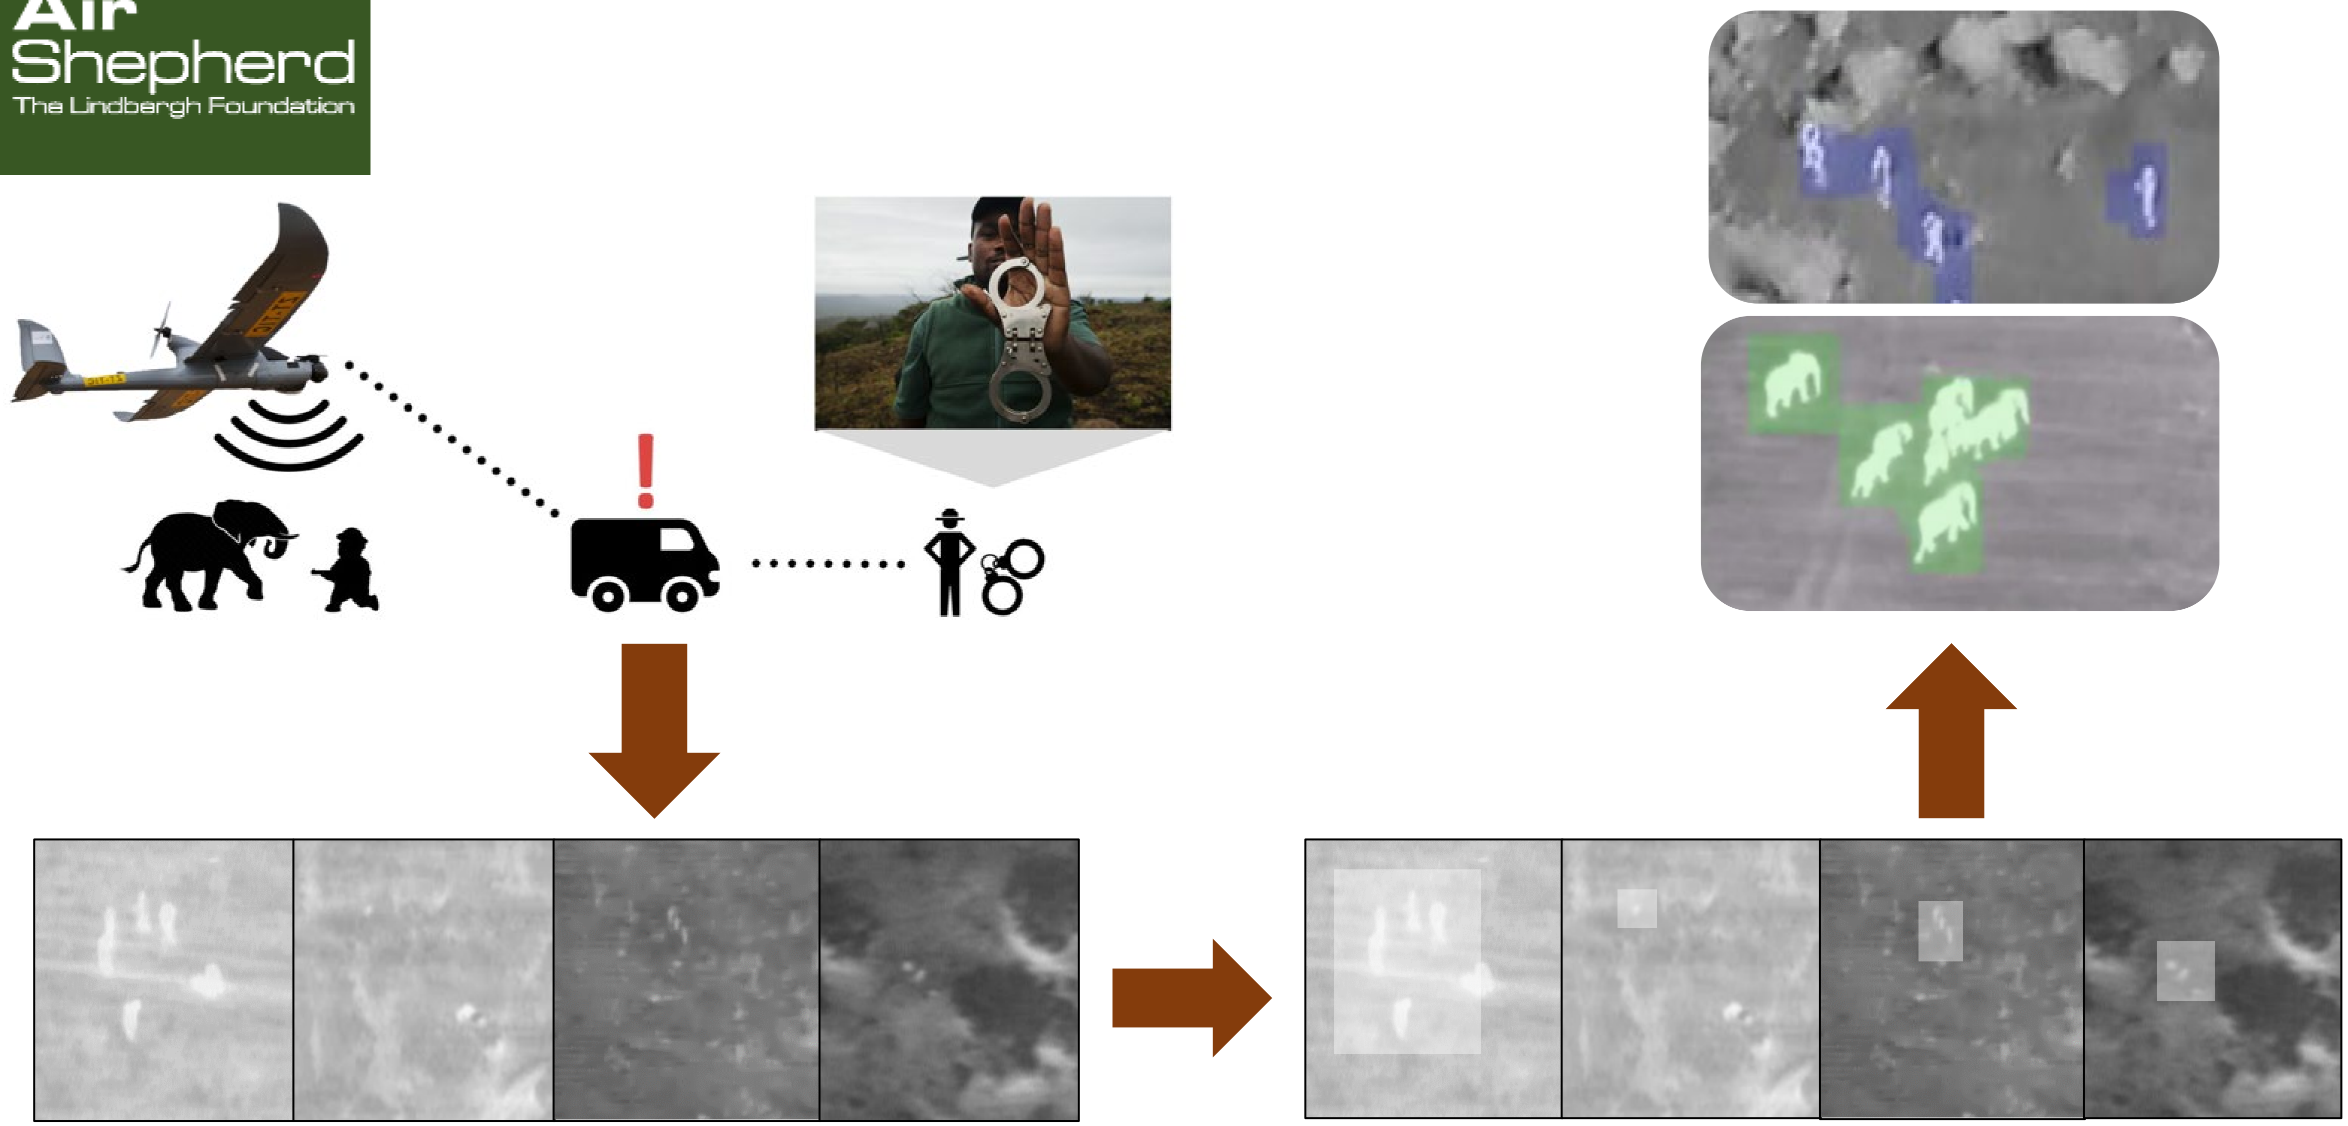

Goal: automatically find poachers

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

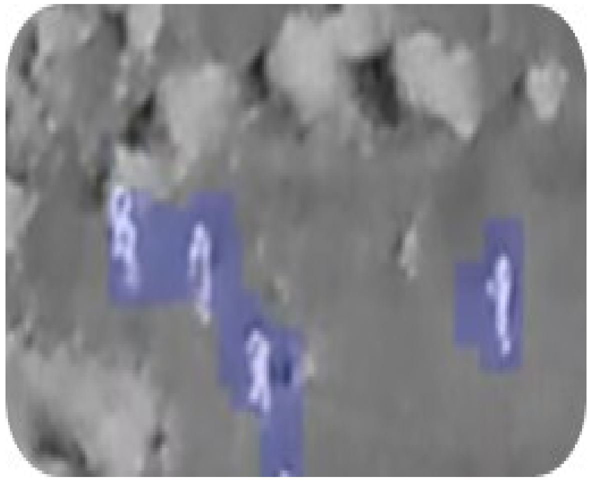

Prob(ranger) = 0.3

Drone Used to Inform Rangers

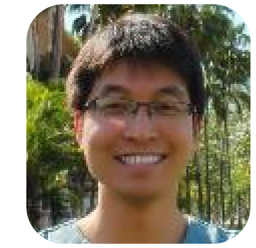

Xu

Bondi

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

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

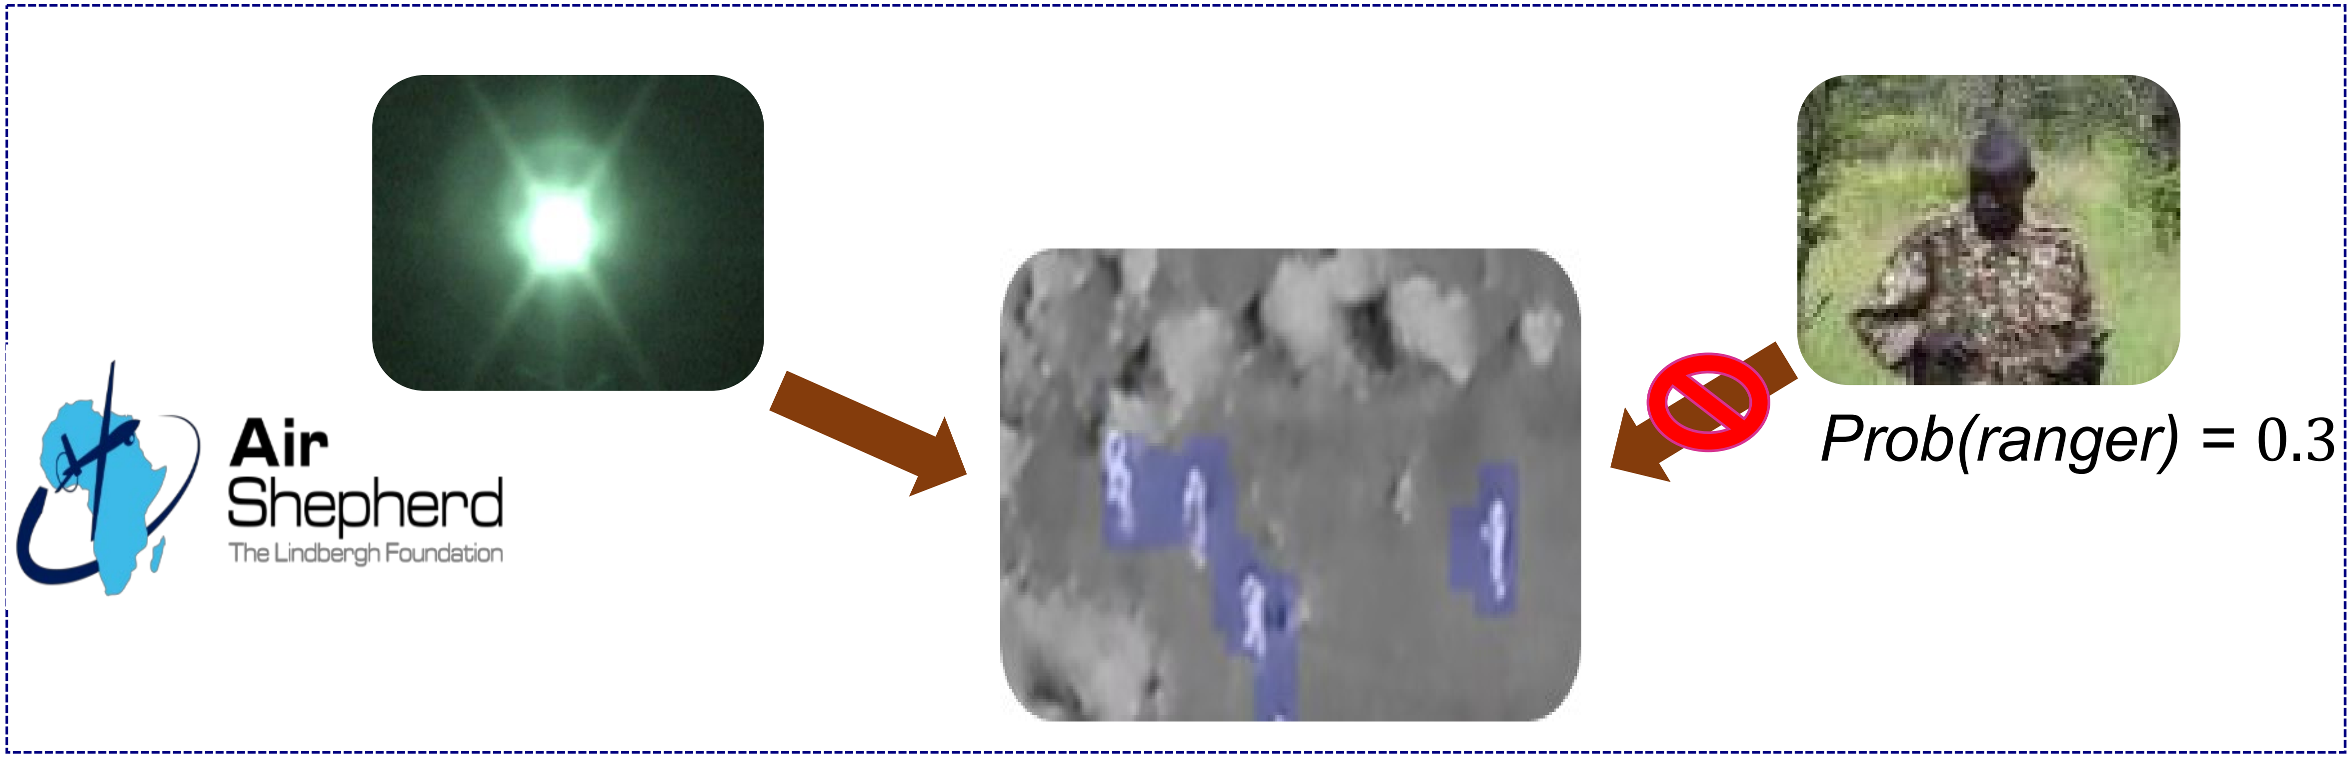

Exploiting Informational Advantage
Defender Knows Pure & Mixed Strategy

(AAAI 2018, AAAI 2020)

Xu

Bondi

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

- Poacher best interest to “believe signal” even if know 50% defender deception

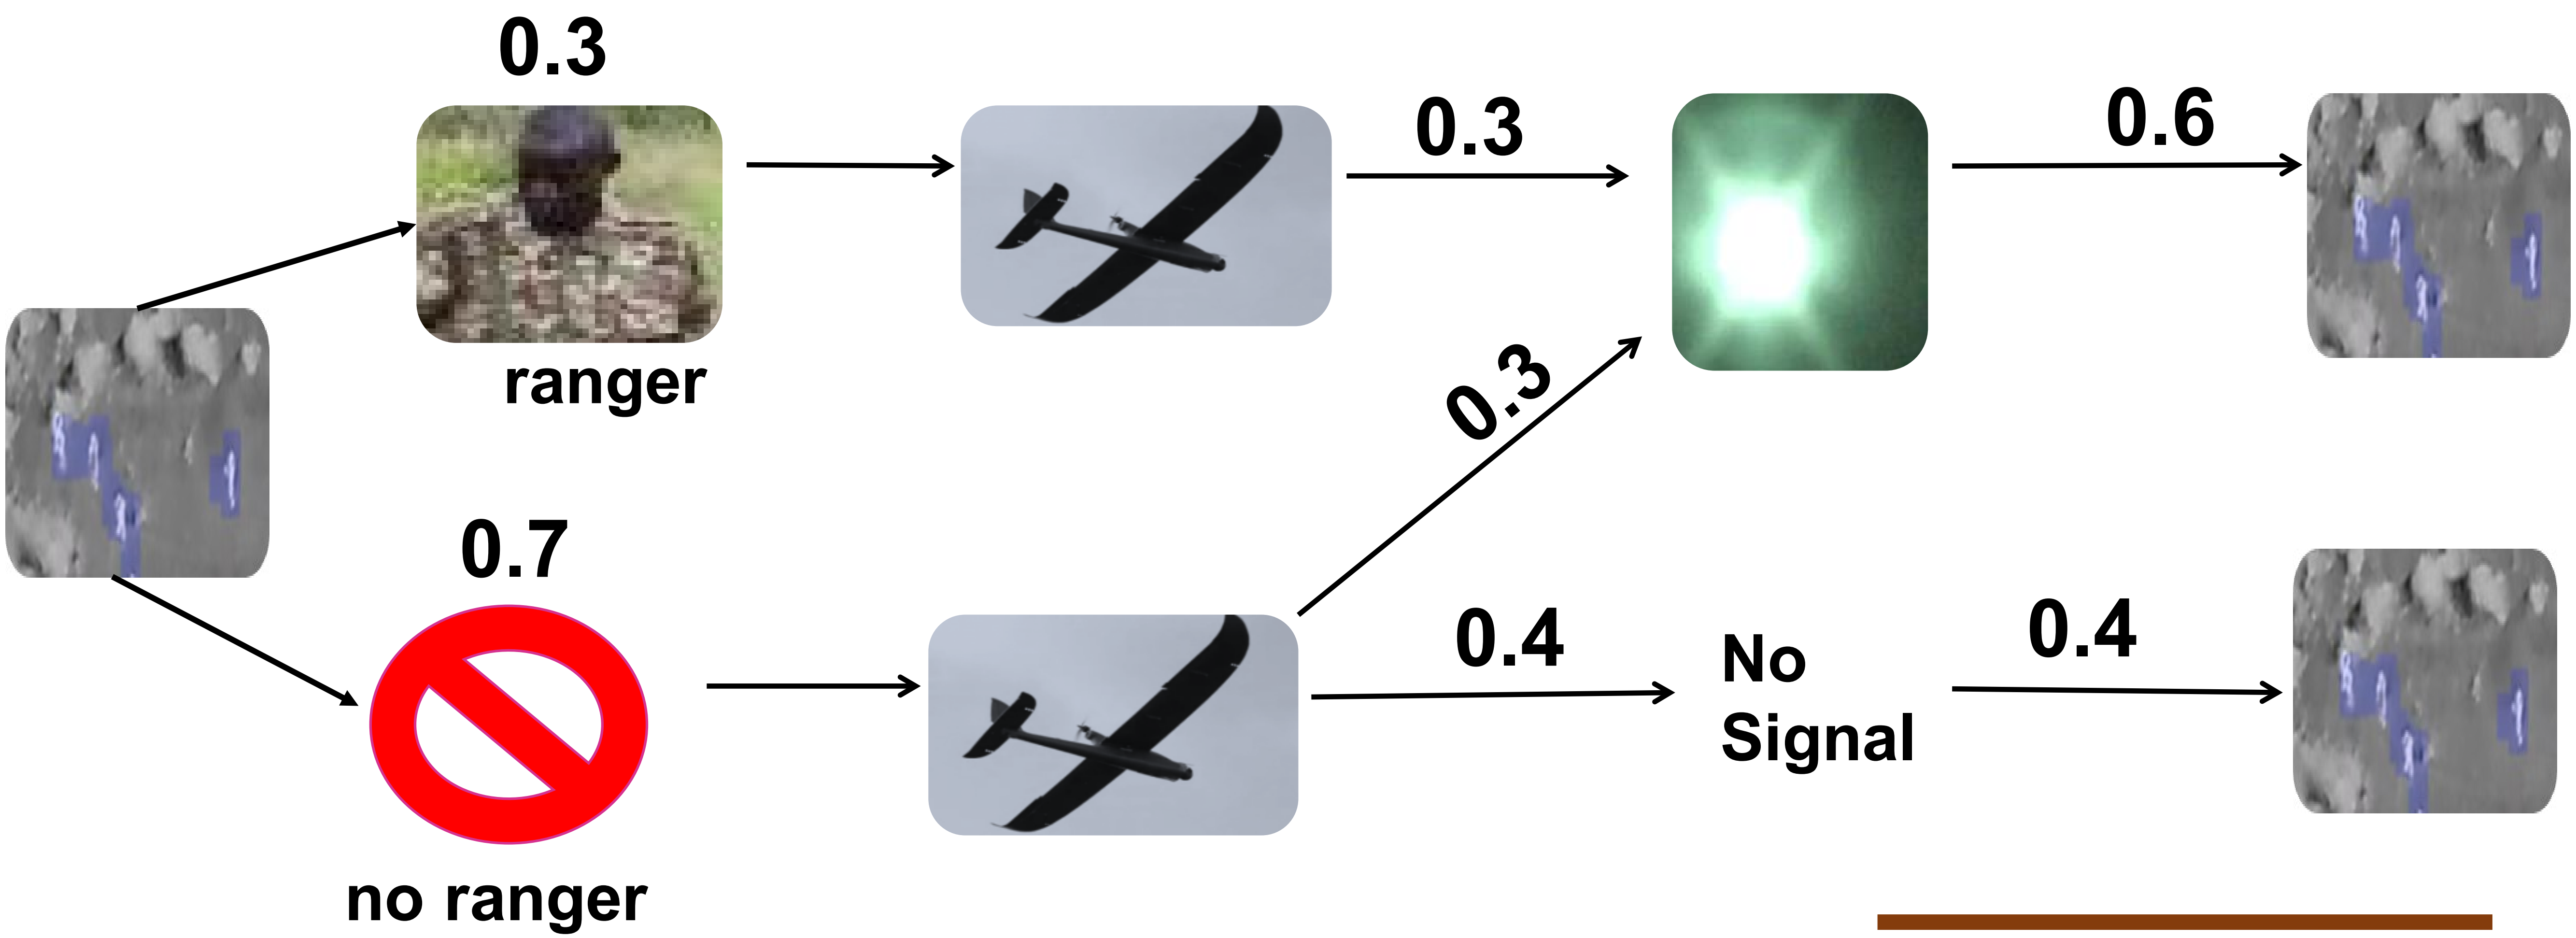

PAWS GOES GLOBAL with SMART platform!!

SMART

WWF

PEACE PARKS FOUNDATION · THE GLOBAL SOLUTION

WCS

AI for Earth

Protect Wildlife
800
National Parks
Around the Globe

Also: Protect Forests, Fisheries…

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

- The END

Another View:
Previous Two-Stage Method: Gradient Descent

Perrault

Wilder

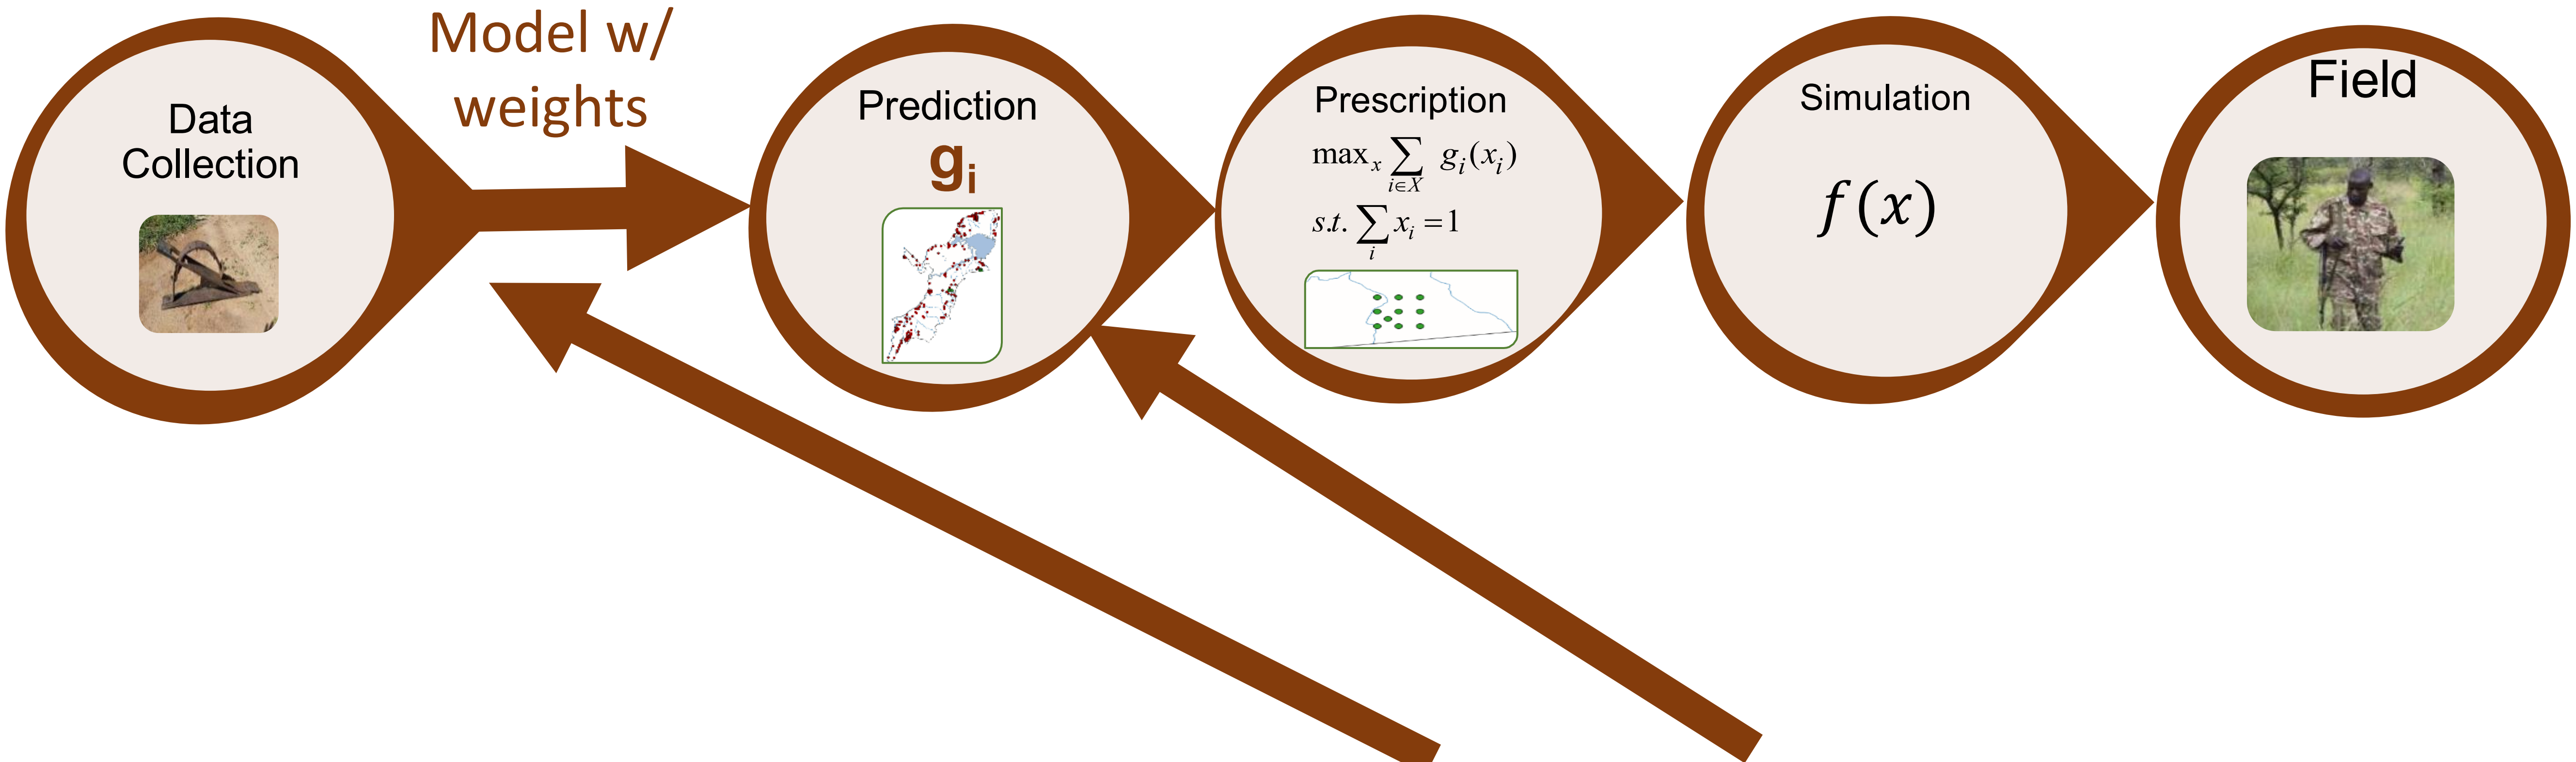

- Max accuracy gradient descent: $\frac{\partial \text{accuracy}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{accuracy}}{\partial \text{prediction}}$

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


Perrault


Wilder

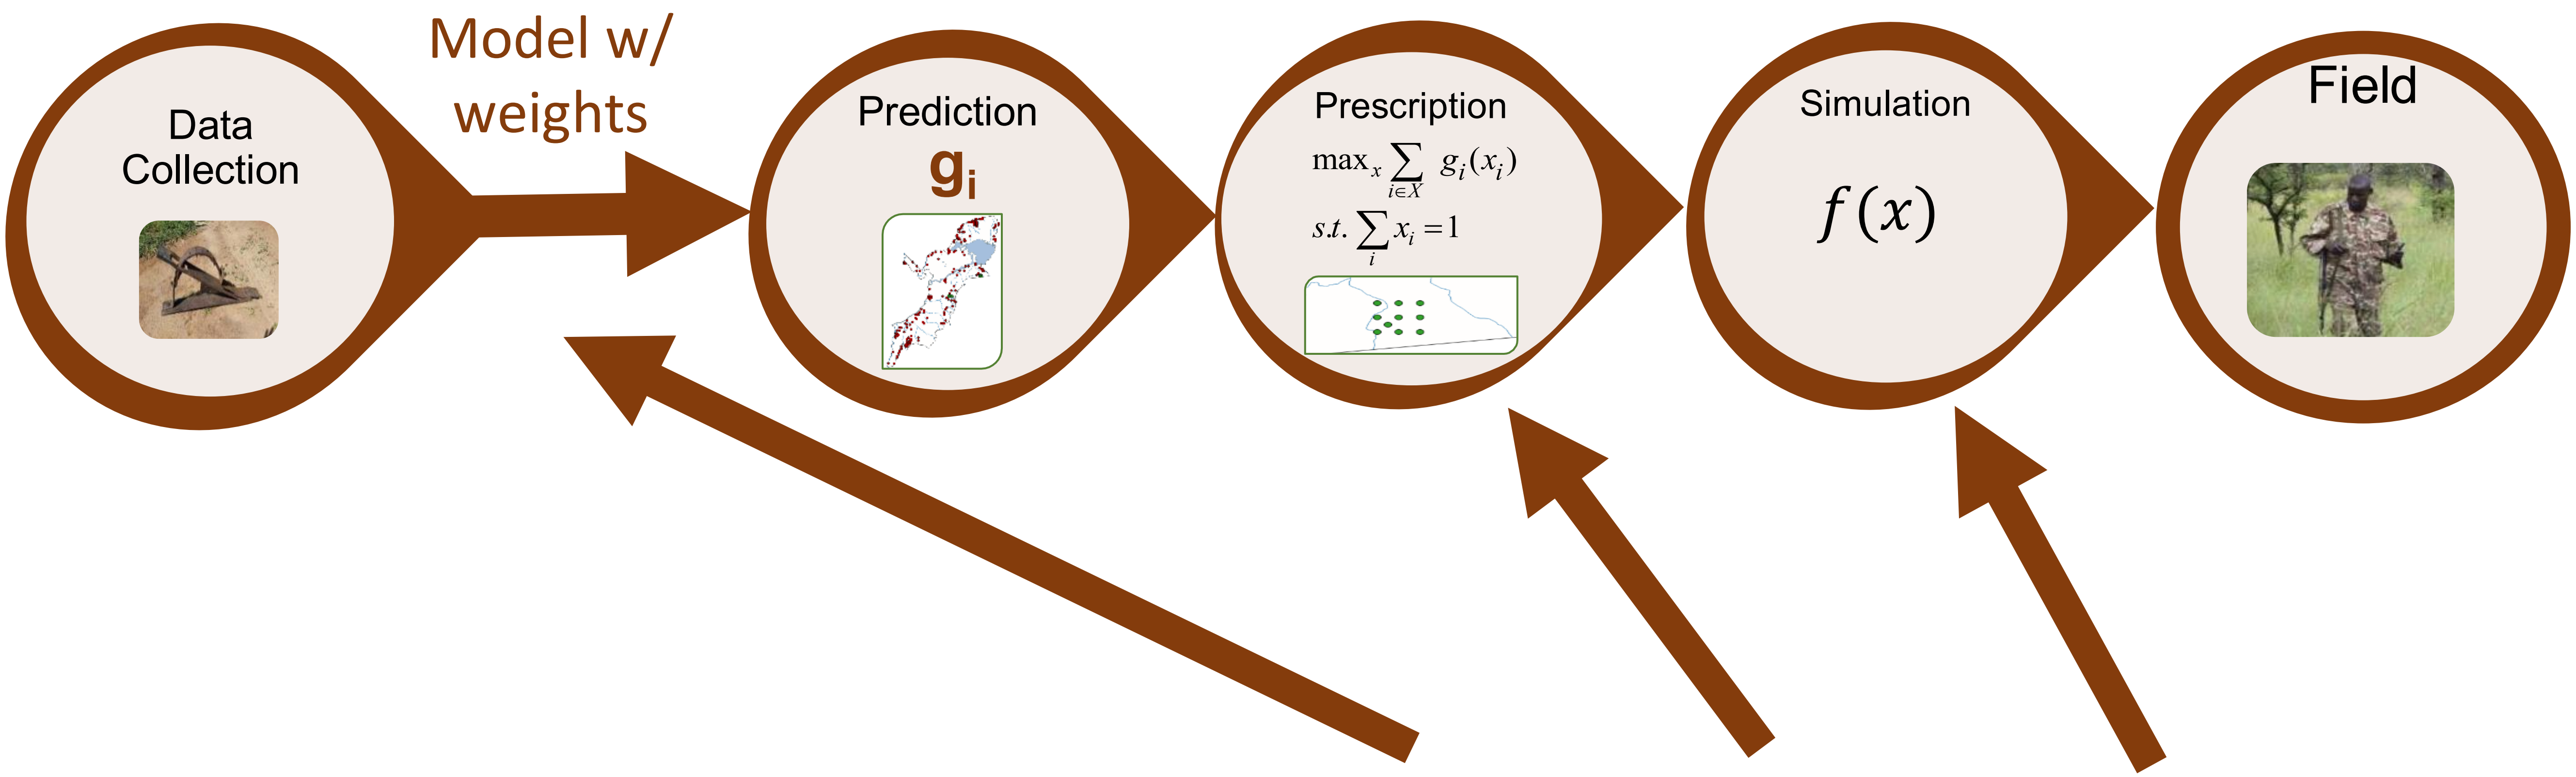

- Game-focused gradient descent: $$\frac{\partial \text{obj(decision)}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{decision}}{\partial \text{prediction}} \frac{\partial \text{obj(decision)}}{\partial \text{decision}}$$