AI for Social Impact:
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

AI and Multiagent Systems Research for Social Impact

Public Safety and Security

Conservation

Public Health

Viewing Social Problems as Multiagent Systems

Key research challenge across problem areas:

Optimize Our Limited Intervention Resources
when
Interacting with Other Agents

Multiagent Systems Reasoning:
Game Theory, Networks

Public Safety and Security
Optimizing Limited Intervention (Security) Resources

Stackelberg Security Games

- Game Theory for security resource optimization
- Real-world: US Coast Guard, US Federal Air Marshals Service…

Conservation/Wildlife Protection
Optimizing Limited Intervention (Ranger) Resources

Green Security Games







- Security games and adversary (poacher) behavior prediction
- <u>Real-world: National parks in Uganda, Malaysia…</u>

Public Health
Optimizing Limited Intervention (Social Worker) Resources

Games against Nature

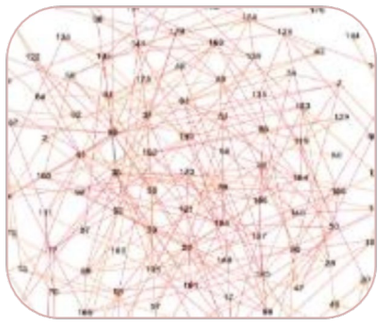

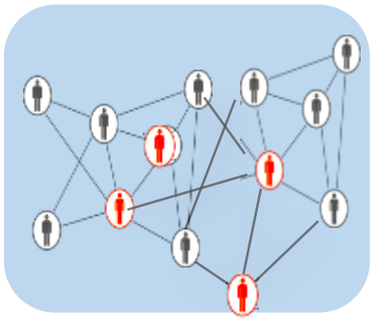

- Social networks to enhance intervention, e.g., HIV information
- <u>Real-world pilot tests: Homeless youth shelters in Los Angeles</u>

Google Research Bangalore (Forthcoming)
Optimizing Limited Intervention Resources

Director, AI for Social Good

Public Health & Welfare

Conservation

Education

- New projects in public health, conservation, education

Common Themes

Interdisciplinary Partnerships, Multiagent Systems, Data-to-deployment pipeline

Wildlife Conservation Society

WWF

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

Transportation Security Administration

LOS ANGELES LGBT CENTER

S.P.Y
safe place for youth

my friend's place

everwell

SMART

Air Shepherd
The Lindbergh Foundation

Common Themes

Interdisciplinary Partnerships, Multiagent Systems, Data-to-deployment pipeline

Immersion

Data Collection

Predictive model

Learning/ Expert input

Prescriptive algorithm

Game theory Intervention

Field tests & deployment

Wildlife Conservation Society

WWF

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

Transportation Security Administration

LOS ANGELES LGBT CENTER

S.P.Y
safe place for youth

my friend's place

everwell

AI for Social Impact
Observations on Area of Research

Field tests & deployments because Social Impact is a key objective!

Encourage AI for Social Impact research: Value entire pipeline (contributions beyond algorithms in data collection, model, impact)

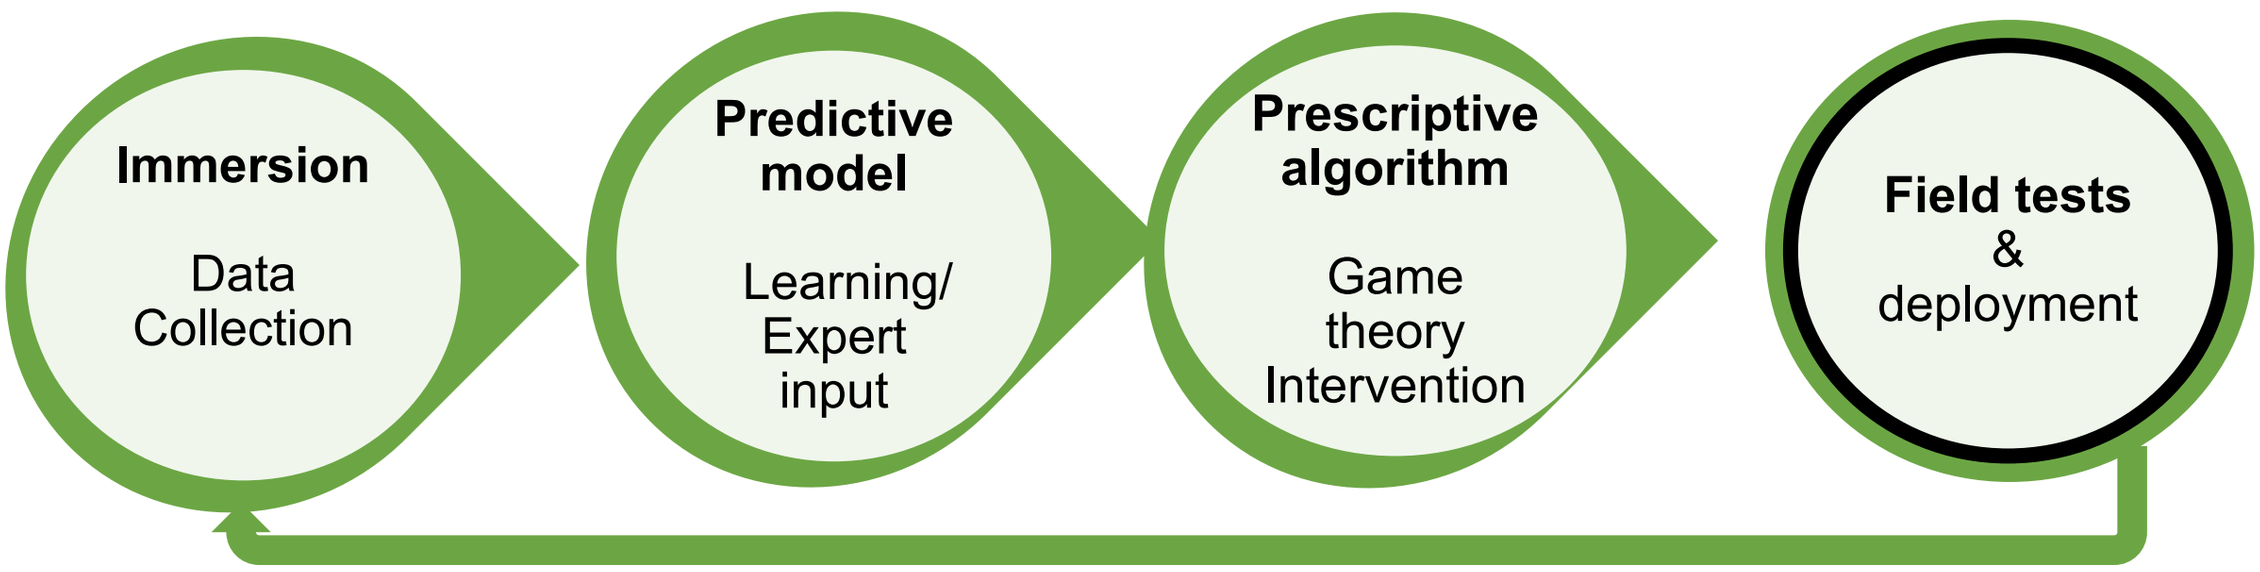

Outline: Overview of Past 12 Years of Research

Public Safety & Security: Stackelberg Security Games

Conservation/Wildlife Protection: Green Security Games

Public Health: Influence maximization/Game against nature

- AAMAS,AAAI,IJCAI
- Real world evaluation
- PhD students & postdocs

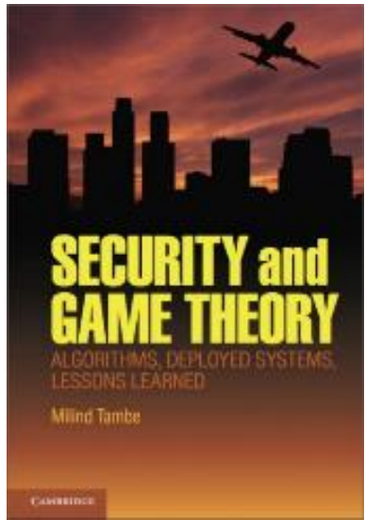

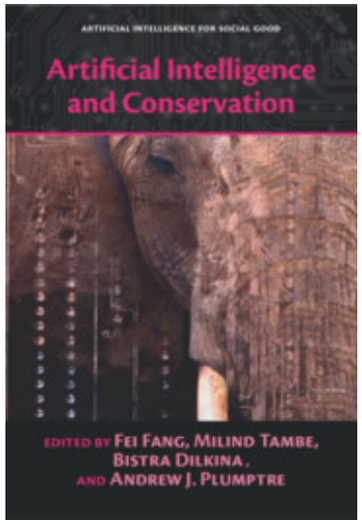

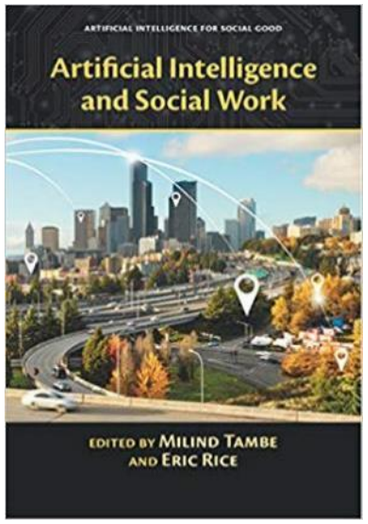

ARMOR Airport Security: LAX(2007)
Game Theory direct use for security resource optimization?

Erroll Southers

LAX Airport, Los Angeles

Game Theory for Security Resource Optimization

New Model: Stackelberg Security Games

Adversary

Defender

	Terminal #1	Terminal #2
Terminal #1	4, -3	-1, 1
Terminal #2	-5, 5	2, -1

Game Theory for Security Resource Optimization

New Model: Stackelberg Security Games

Adversary

Defender

	Terminal #1	Terminal #2
Terminal #1	4, -3	-1, 1
Terminal #2	-5, 5	2, -1

Game Theory for Security Resource Optimization

New Model: Stackelberg Security Games

Stackelberg: Defender commits to randomized strategy, adversary responds

Security game: Played on targets, payoffs based on targets covered or not

Optimization: Not 100% security; increase cost/uncertainty to attackers

Adversary

Defender

	Terminal #1	Terminal #2
Terminal #1	4, -3	-1, 1
Terminal #2	-5, 5	2, -1

ARMOR at LAX
Basic Security Game Operation [2007]


Kiekintveld


Pita

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

Mixed Integer Program

Pr (Canine patrol, 8 AM @Terminals 2,5,6) = 0.17

Canine Team Schedule, July 28								
	Term 1	Term 2	Term 3	Term 4	Term 5	Term 6	Term 7	Term 8
8 AM		Team1			Team3	Team5		
9 AM			Team1	Team2				Team4
…	…	…	…	…	…	…	…	…

Security Game MIP [2007]
Payoffs Estimated From Previous Research


Kiekintveld


Pita

j →

i ↓

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

$$\max \sum_{i \in X} \sum_{j \in Q} R_{ij} \times x_i \times q_j$$

Maximize defender expected utility

$$s.t. \quad \sum_i x_i = 1$$

Defender mixed strategy

$$\sum_{j \in Q} q_j = 1$$

Adversary response

$$0 \le (a - \sum_{i \in X} C_{ij} x_i) \le (1 - q_j) M$$

Adversary best response

SECURITY GAME PAYOFFS [2007]
Previous Research Provides Payoffs in Security Games

	Target #1	Target #2	Target #3
Defender #1	2, -1	-3, 4	-3, 4
Defender #2	-3, 3	3, -2	….
Defender #3	….	….	….

Maximize defender expected utility

$$\max \sum_{i \in X} \sum_{j \in Q} R_{ij} \times x_i \times q_j$$

+ Handling Uncertainty

ARMOR:
Optimizing Security Resource Allocation [2007]

First application: Computational game theory for operational security

January 2009

- January 3rd — *Loaded 9/mm pistol*
- January 9th — *16-handguns, 1000 rounds of ammo*
- January 10th — *Two unloaded shotguns*
- January 12th — *Loaded 22/cal rifle*
- January 17th — *Loaded 9/mm pistol*
- January 22nd — *Unloaded 9/mm pistol*

ARMOR AIRPORT SECURITY: LAX [2008]
Congressional Subcommittee Hearings

Commendations
City of Los Angeles

Erroll Southers testimony
Congressional subcommittee

Newsweek

ARMOR…throws a digital cloak of invisibility….

IRIS
Federal Air Marshals Service [2009]


Kiekintveld


Jain

1000 flights, 20 air marshals: 10^{41} combinations

	Attack 1	Attack 2	Attack …	Attack 1000
1, 2, 3 ..	5,-10	4,-8	…	-20,9
1, 2, 4 ..	5,-10	4,-8	…	-20,9
1, 3, 5 ..	5,-10	-9,5	…	-20,9
…	← 10^{41} rows			

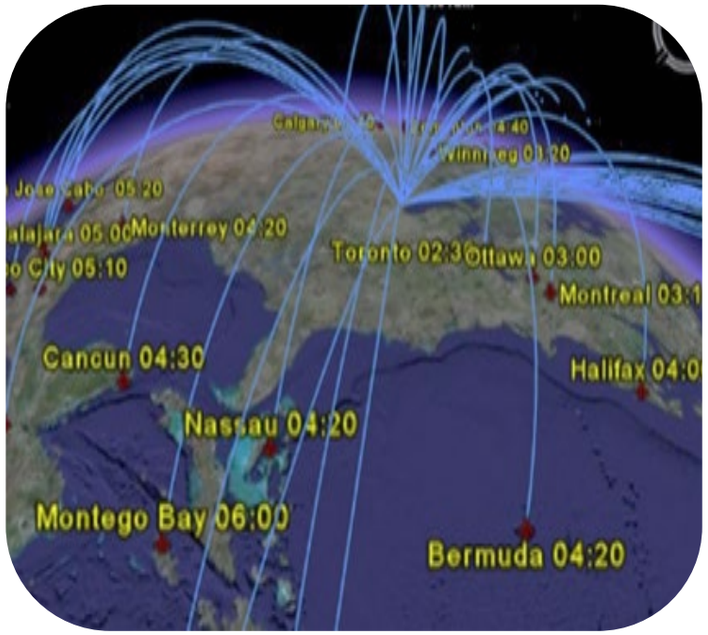

Master

	Attack 1	Attack 2	…	Attack 6
1,2,4	5,-10	4,-8	…	-20,9

	Attack 1	Attack 2	…	Attack 6
1,2,4	5,-10	4,-8	…	-20,9
3,7,8	-8,10	-8,10	…	-8, 10

	Attack 1	Attack 2	…	Attack 6
1,2,4	5,-10	4,-8	…	-20,9
3,7,8	-8,10	-8,10	…	-8, 10
…	…	…	…	…

Slave (LP Duality Theory)
Best new pure strategy

Slave (LP Duality Theory)
Next best new pure strategy

PROTECT: Port and Ferry Protection Patrols [2011] Using Marginals for Scale up


Shieh


An

Boston

Los Angeles

New York

Meritorious Team Commendation from Commandant (US Coast Guard)

Solving Problems: Overall Research Framework
End-to-End, Data to Deployment Pipeline

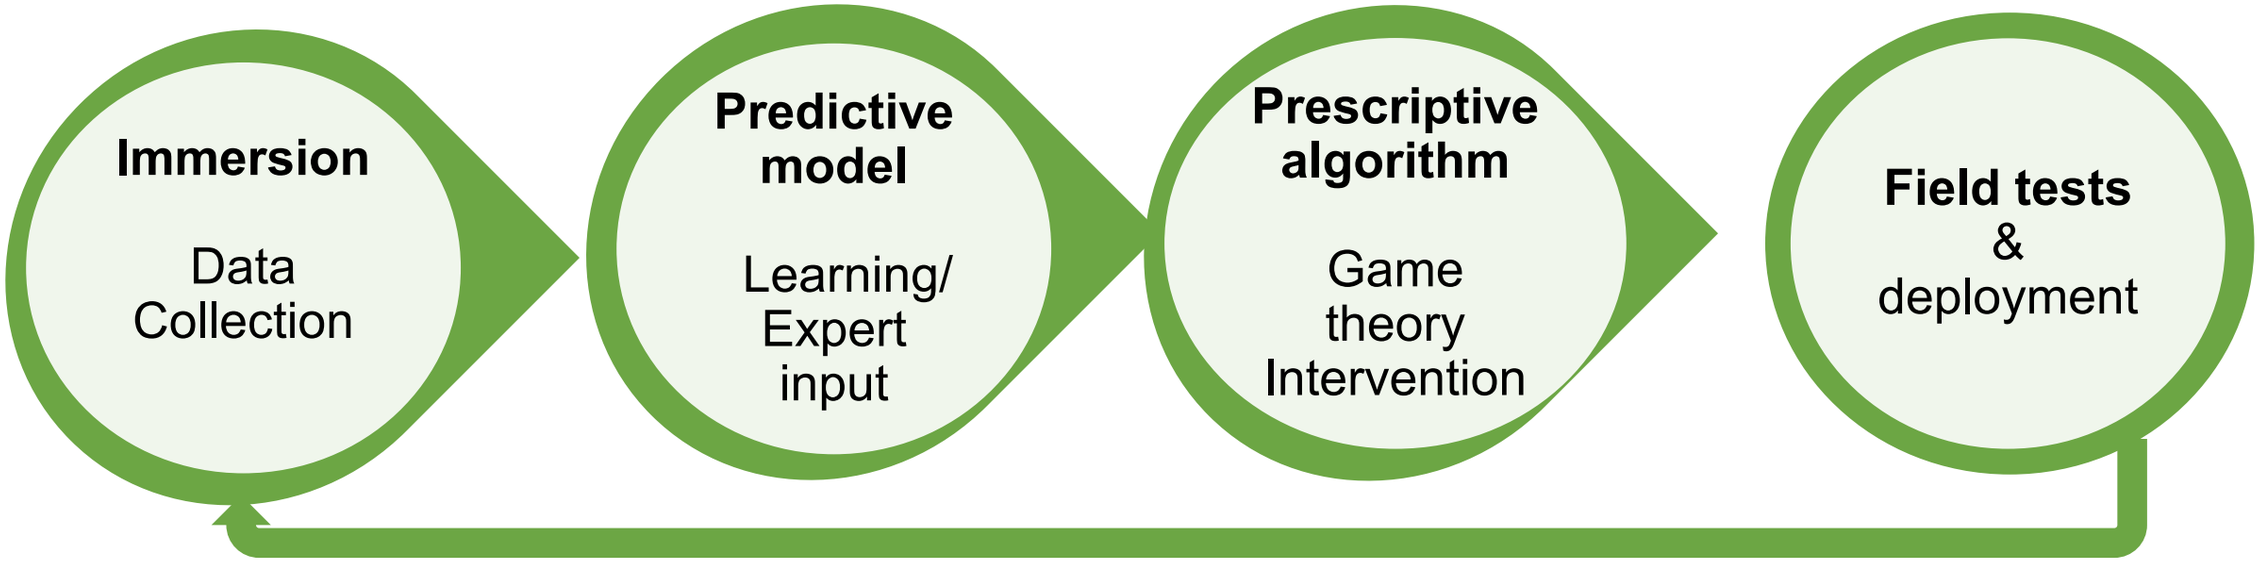

Global Presence of Security using Game Theory

Significant Real-World Evaluation Effort

Security Games superior in
Optimizing Limited Security Resources
Vs

Human Schedulers/“simple random”

Field Evaluation of Schedule Quality

Improved Patrol Unpredictability & Coverage for Less Effort

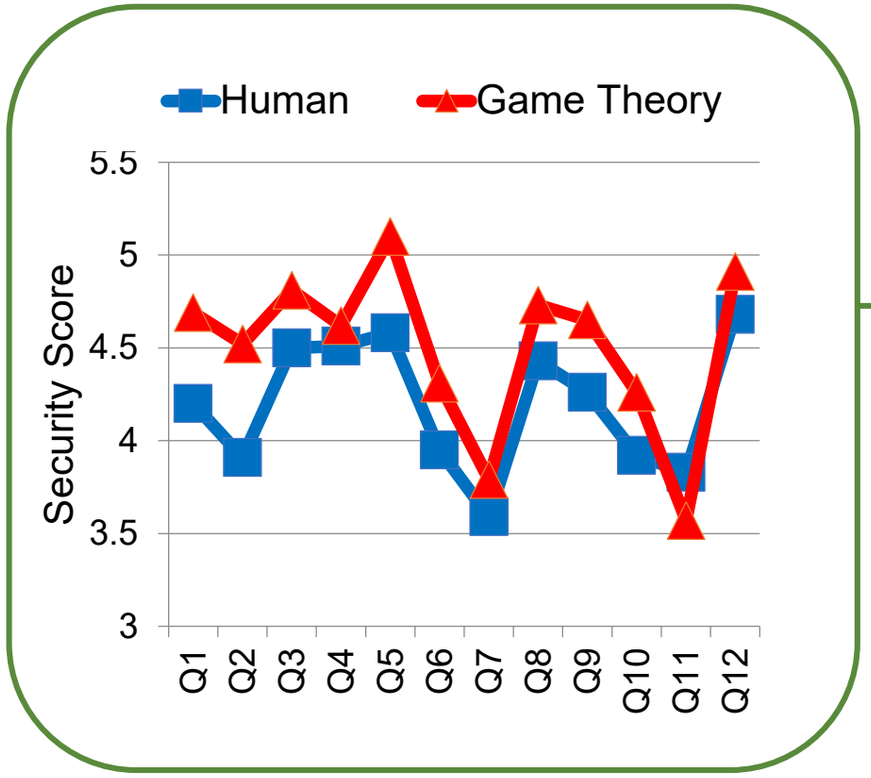

Train patrols: Game theory outperformed expert humans schedule 90 officers

Field Tests Against Adversaries

Computational Game Theory in the Field

Controlled

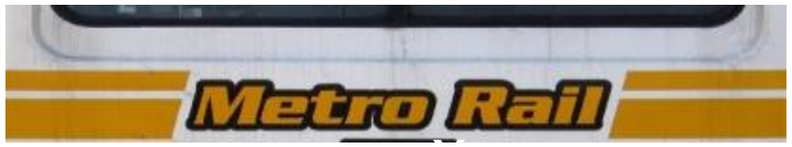

- 21 days of patrol, identical conditions
- Game theory vs Baseline+Expert

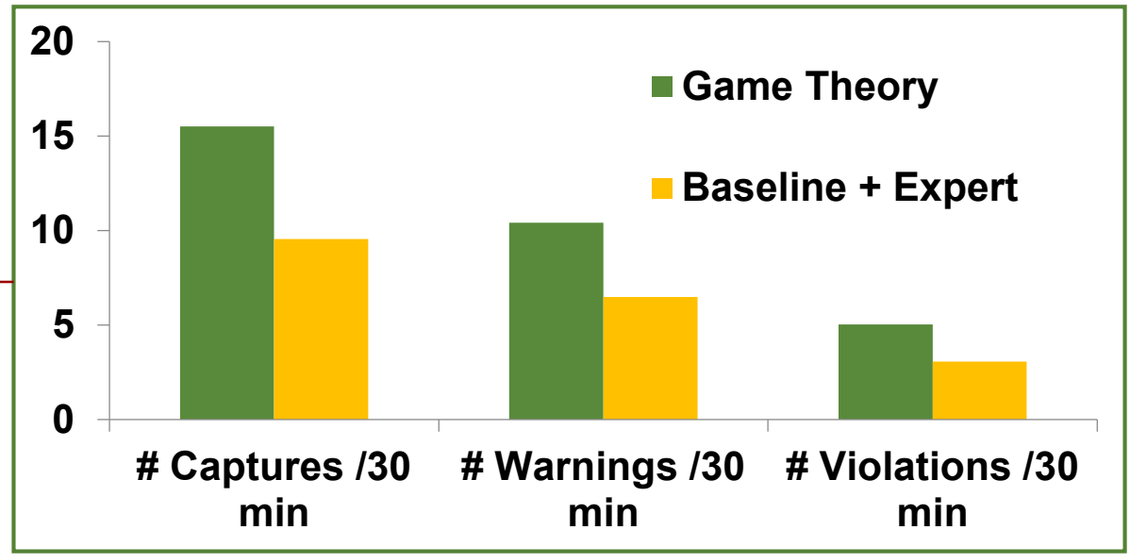

Not Controlled

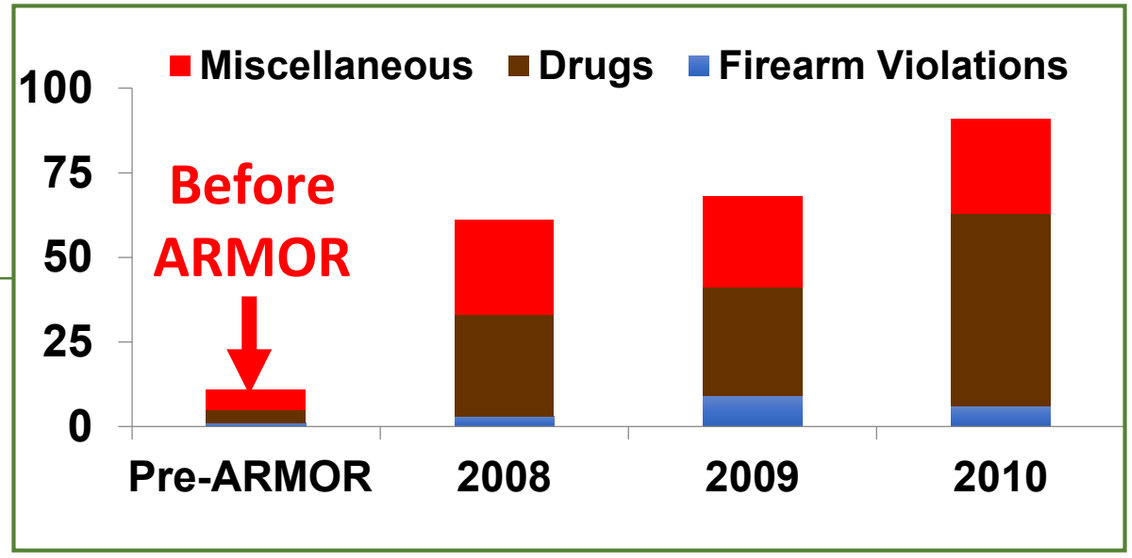

Outline

Public Safety & Security: Stackelberg Security Games

Conservation/Wildlife Protection: Green Security Games

Dr Andy Plumptre
Conservation Biology

Public Health: Influence maximization/Game against nature

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

Snare or Trap

Wire snares

Green Security Games[2015]
Limited Ranger Resources to Protect Forests

Fang

Adversary not fully strategic; multiple “bounded rational” poachers

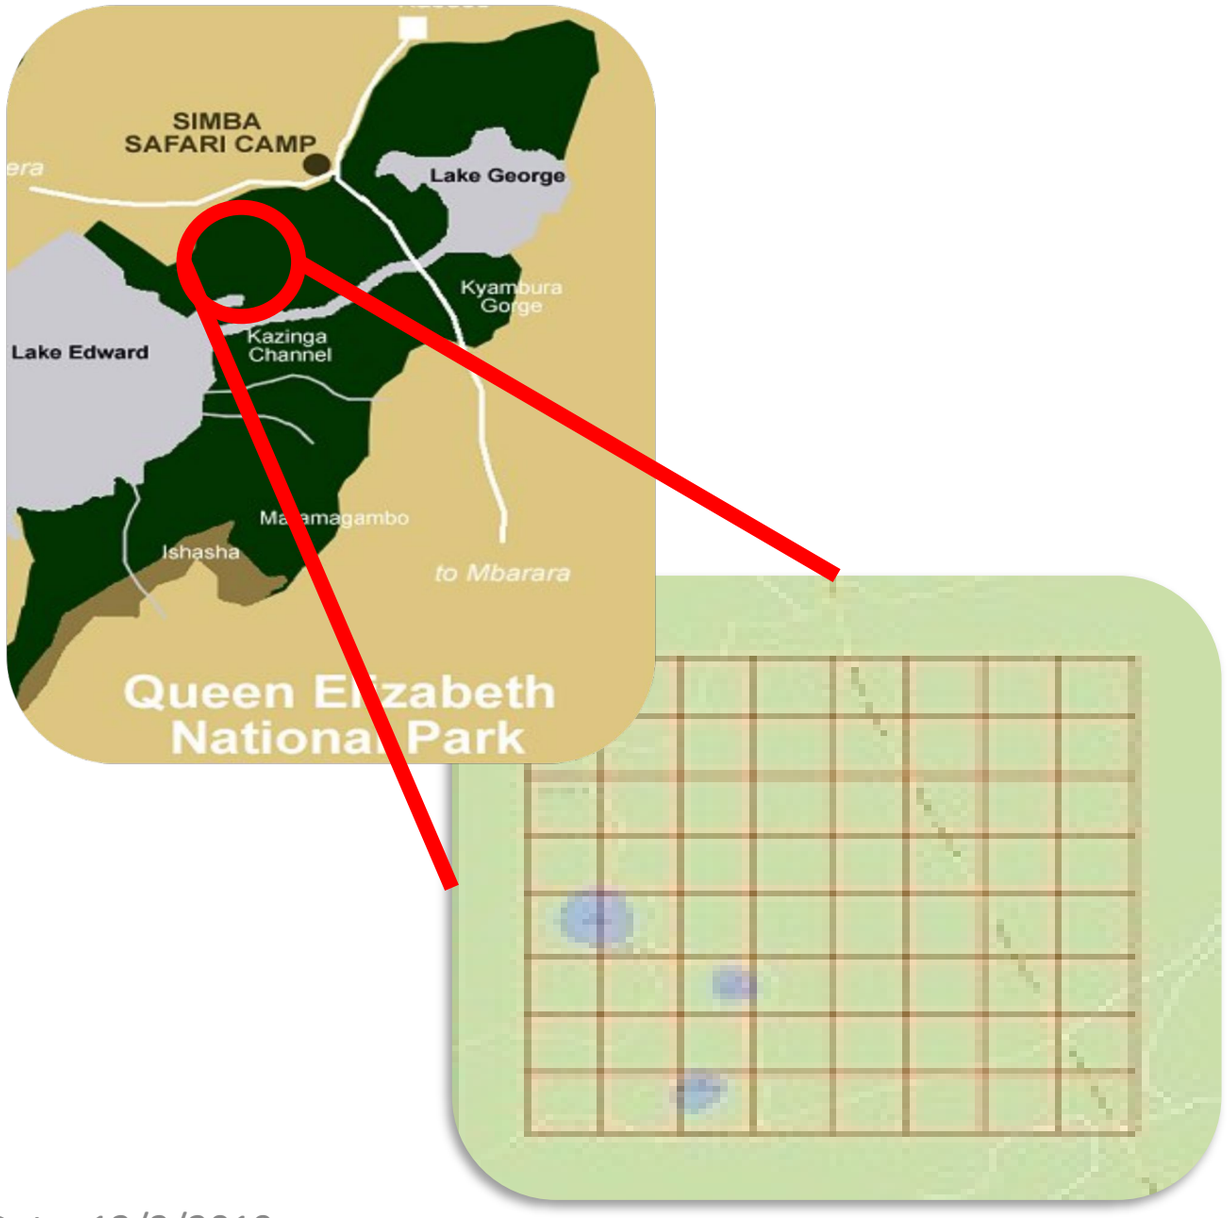

$$\max_{x,q} \sum_{i \in X} \sum_{j \in Q} R_{ij} x_i q_j$$

Max defender utility

$$s.t. \sum_i x_i = 1$$

Defender mixed strategy

$$0 \le (a - \sum_{i \in X} x_i) \le (1 - q_j) M$$

Green Security Games [2015]
Game Theory + Machine Learning Poacher Behavior

Xu

Learn adversary bounded rational response: At each grid **location i**

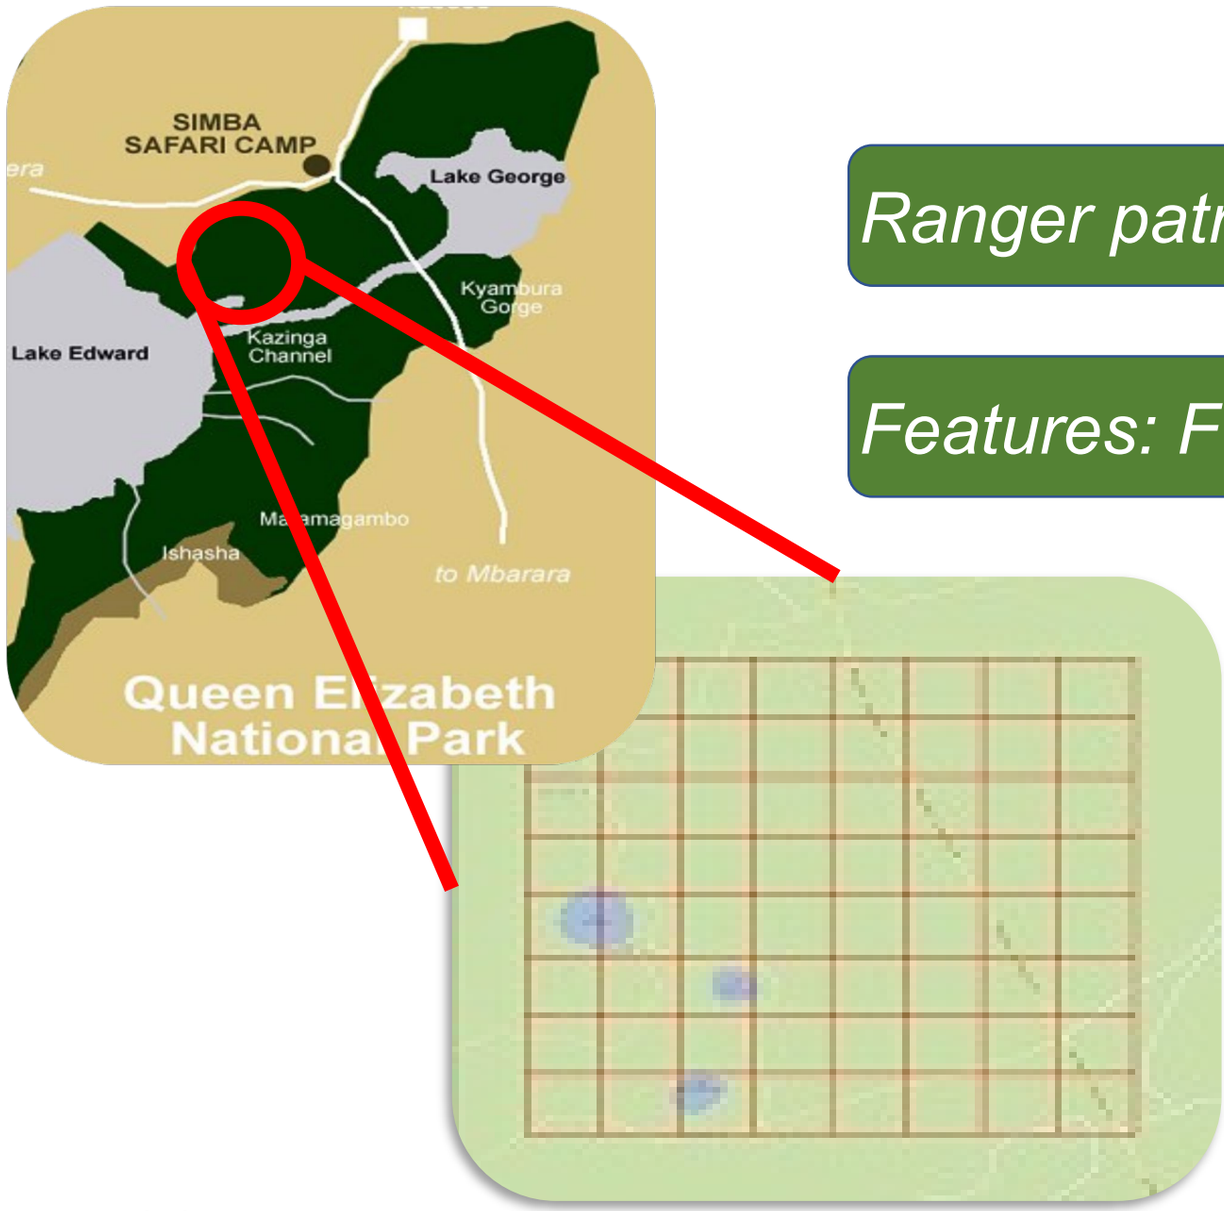

Ranger patrols: $X_{(i)}$

Features: $F_{(i)}$

Machine Learning

g_i

Probability of finding snare in cell i

$$\max_x \sum_{i \in X} g_i(x_i)$$

Max defender utility

$$s.t. \sum_i x_i = 1$$

Defender mixed strategy

Learning Adversary Model
12 Years of Past Poaching Data

Nguyen

g_i

Ranger patrol

Animal density

Distance to rivers / roads / villages

Probability of snare Per 1 KM Grid Square

Area habitat

Area slope

...

Learning Adversary Model Uncertainty in Observations

Nguyen

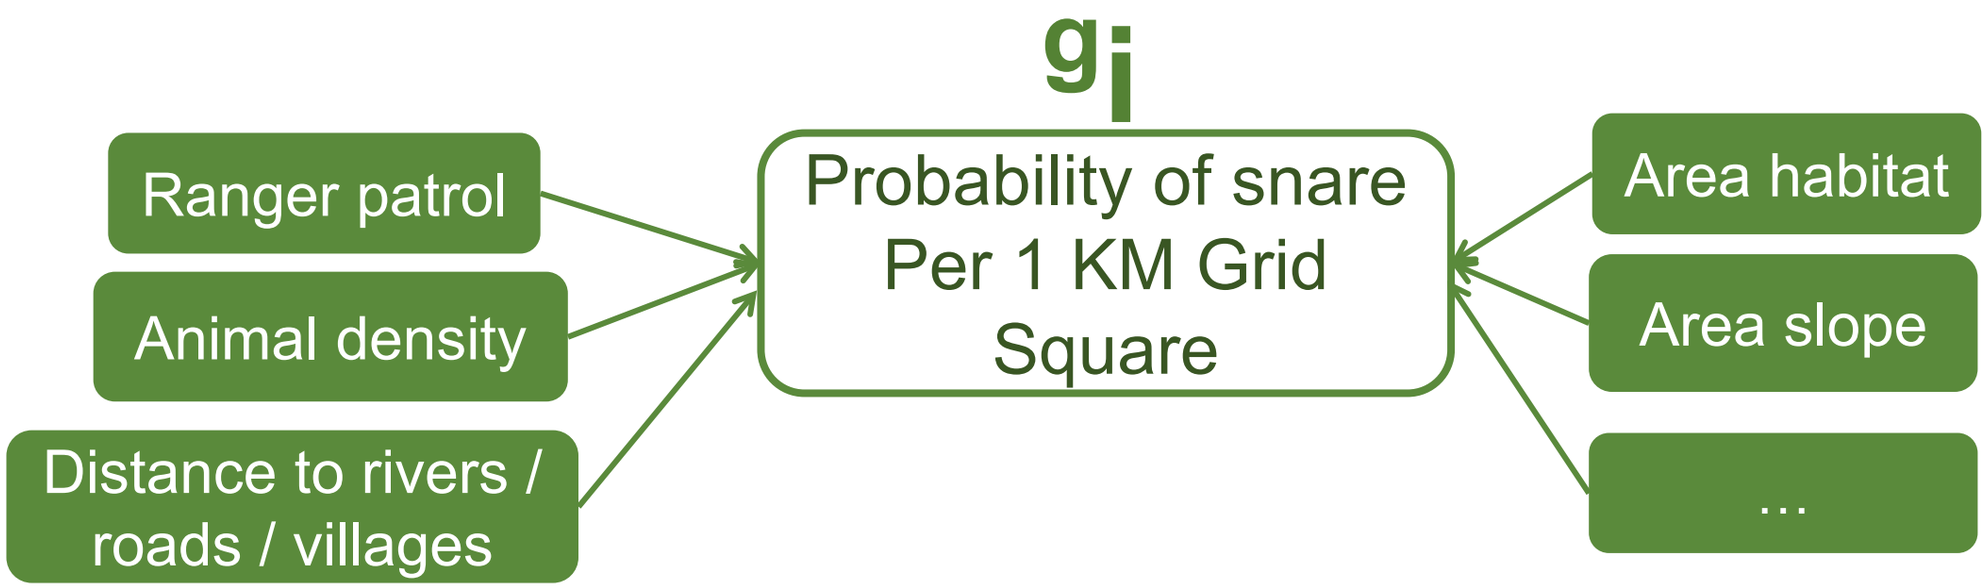

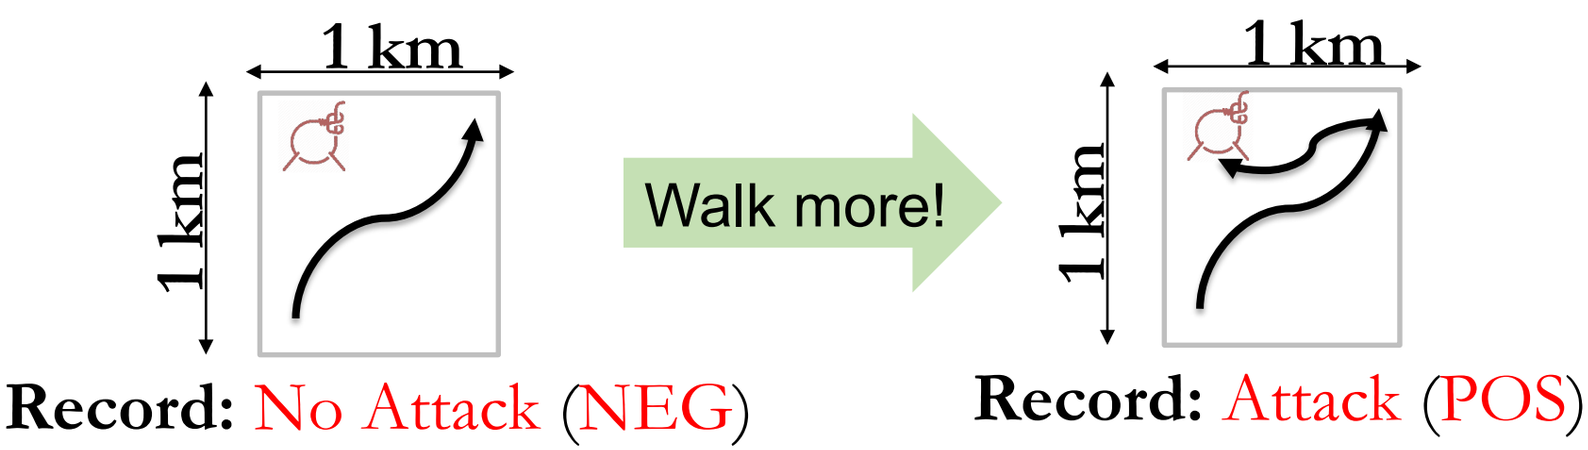

Adversary Modeling [2016]
Imperfect Crime Observation-aware Ensemble Model

Gholami

Training: Filtered Datasets

Predict: Ensemble of Classifiers

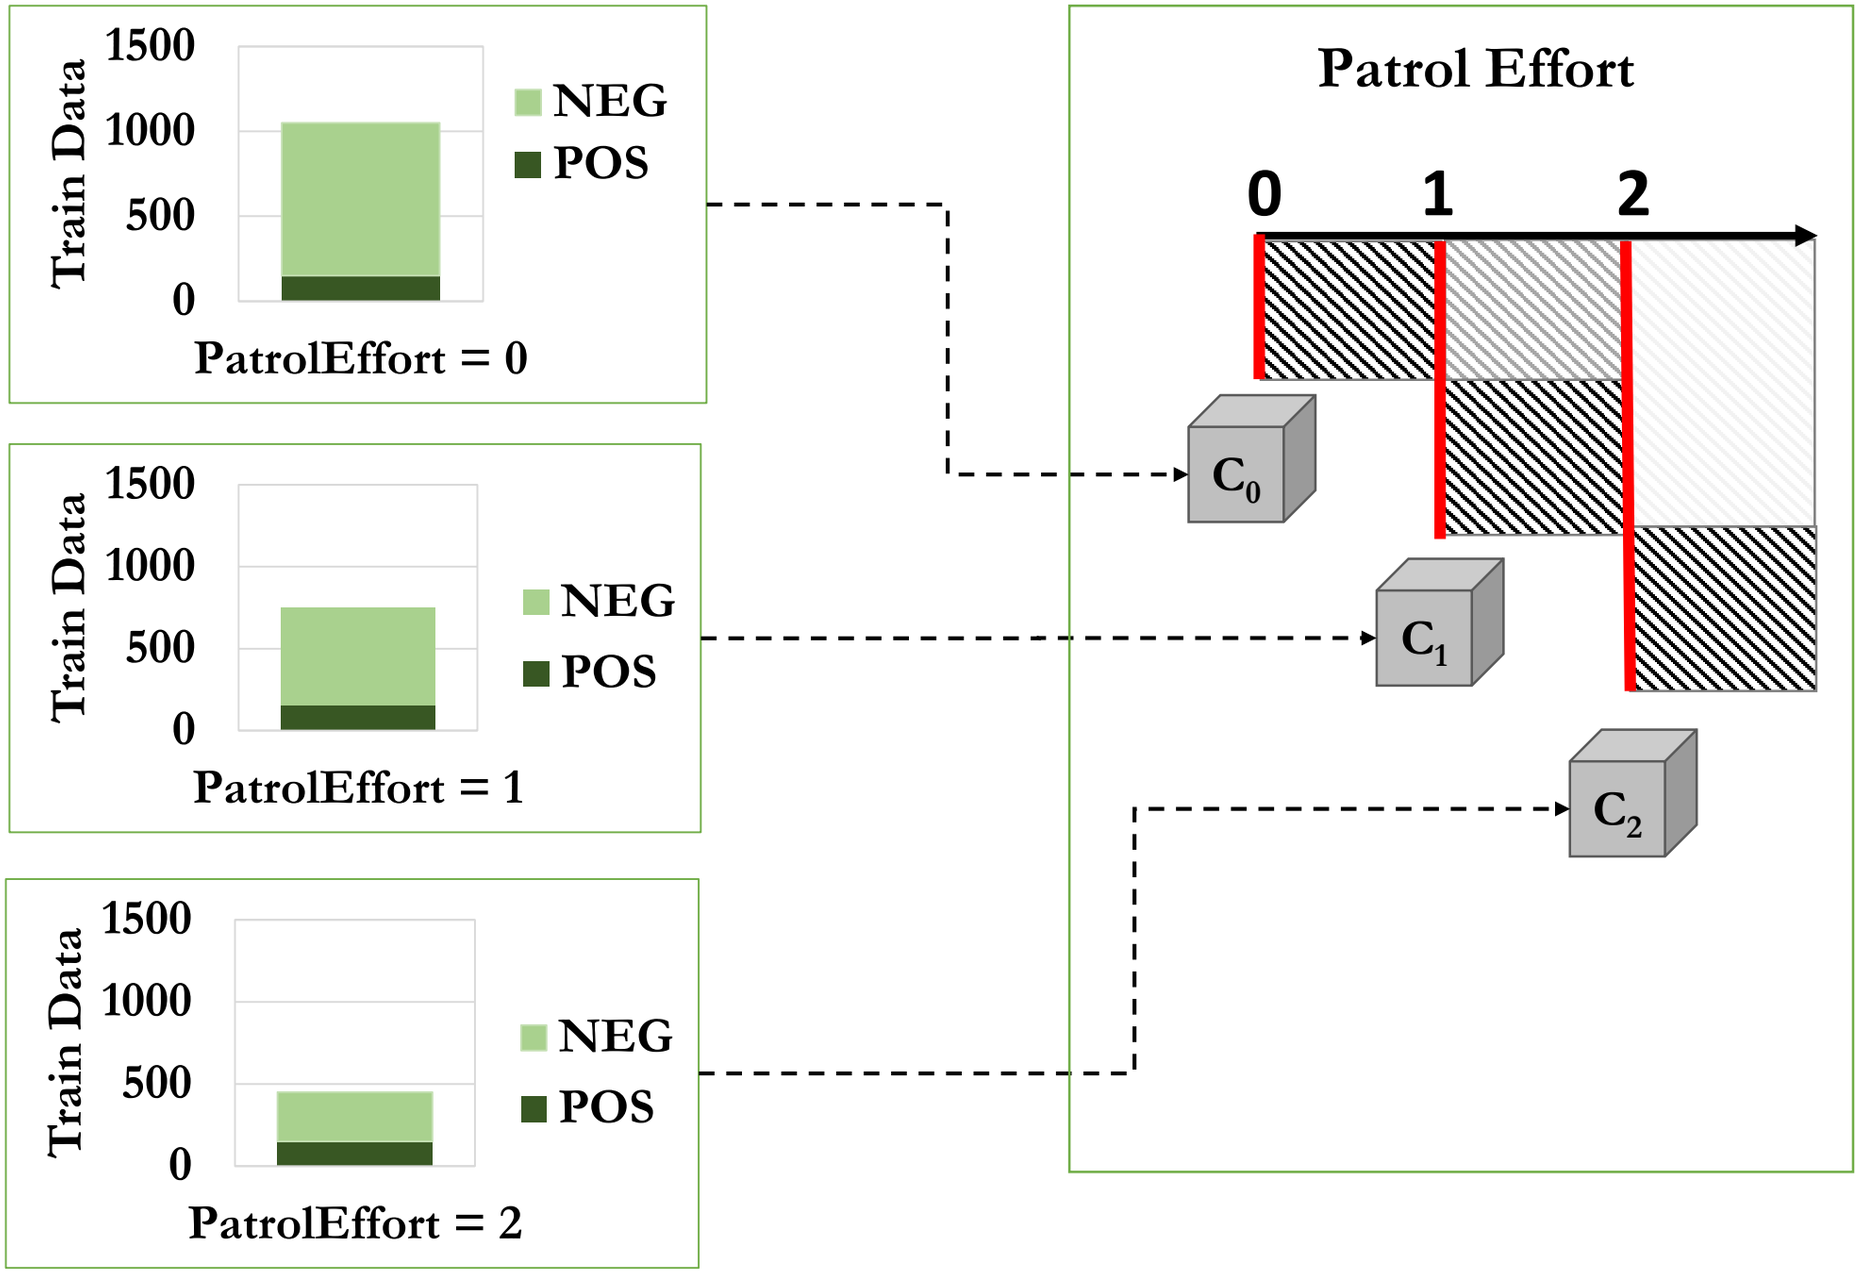

PAWS: Protection Assistant for Wildlife Security
Poacher Attack Prediction in the Lab

Gholami

Poacher Behavior Prediction

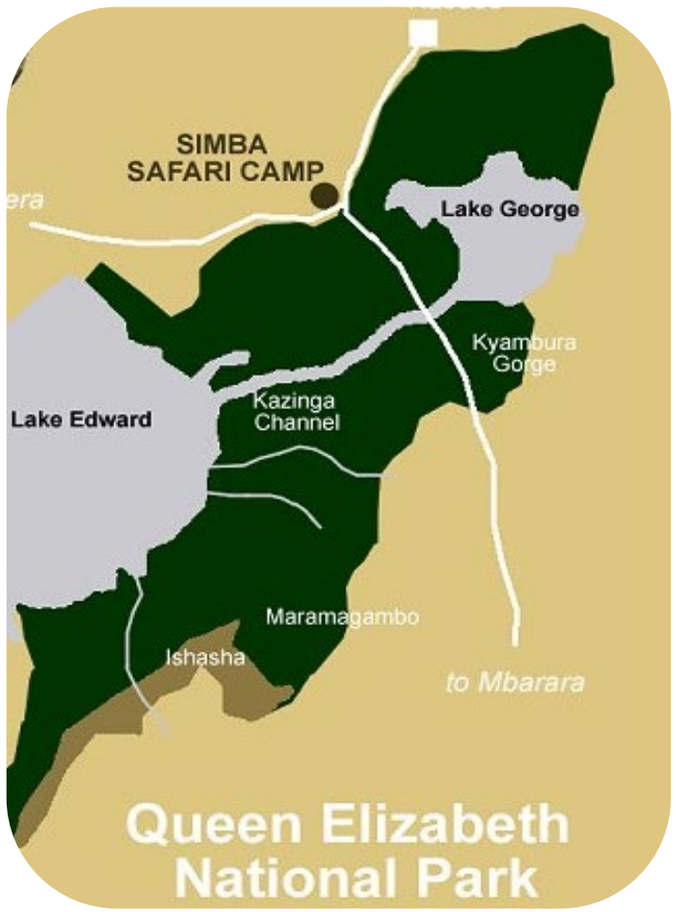

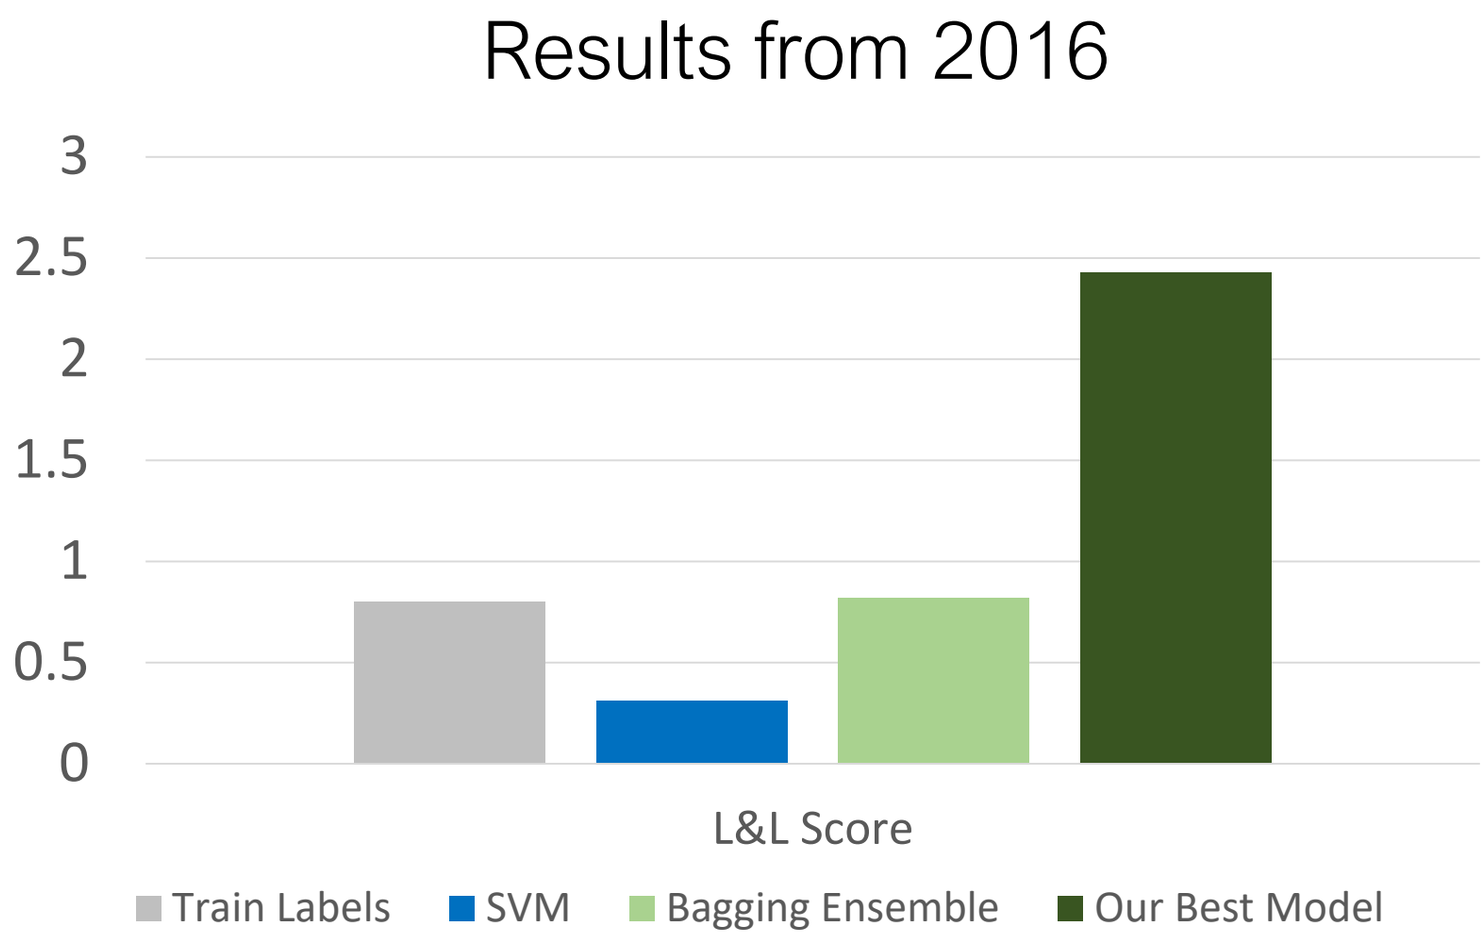

PAWS:
Real-world Deployment 2016: First Trial


Ford


Gholami

- Two 9-sq. km patrol areas
  - Where there were infrequent patrols
  - Where no previous hot spots

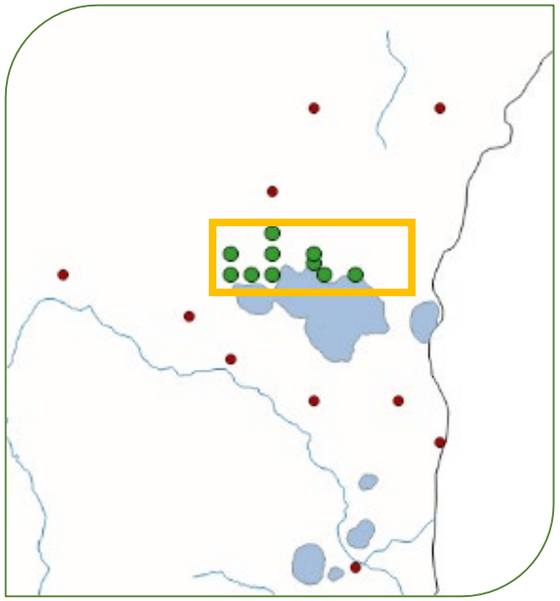

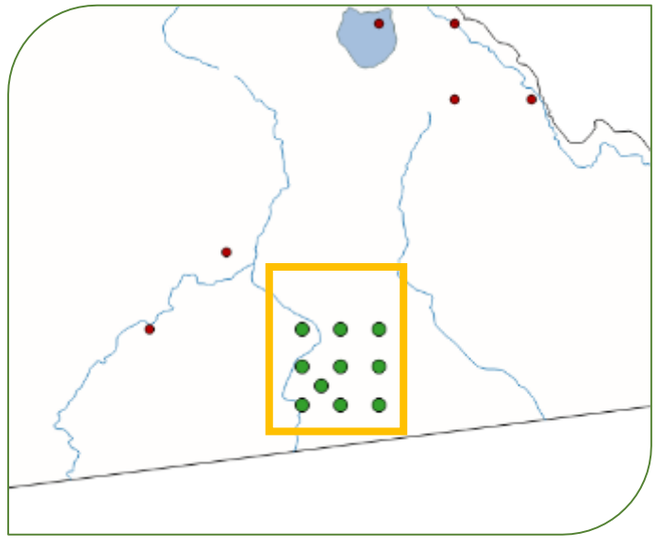

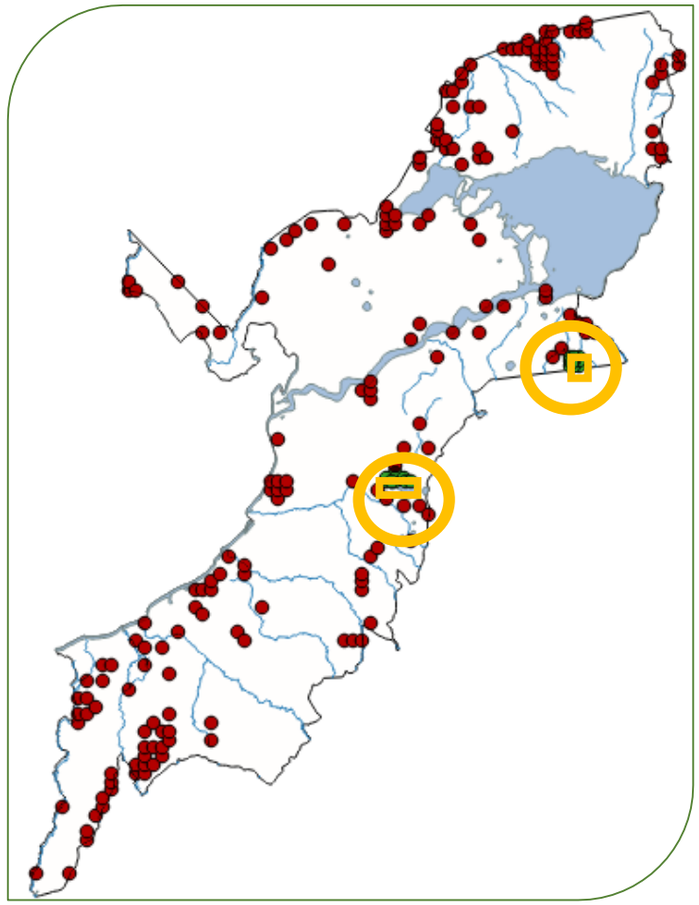

PAWS Real-world Deployment
Two Hot Spots Predicted


Ford


Gholami

- Poached Animals: Poached elephant
- Snaring: 1 elephant snare roll
- Snaring: 10 Antelope snares

Historical Base Hit Rate	Our Hit Rate
Average: 0.73	3

PAWS Predicted High vs Low Risk Areas: 2 National Parks, 24 areas each, 6 months [2017]

Gholami

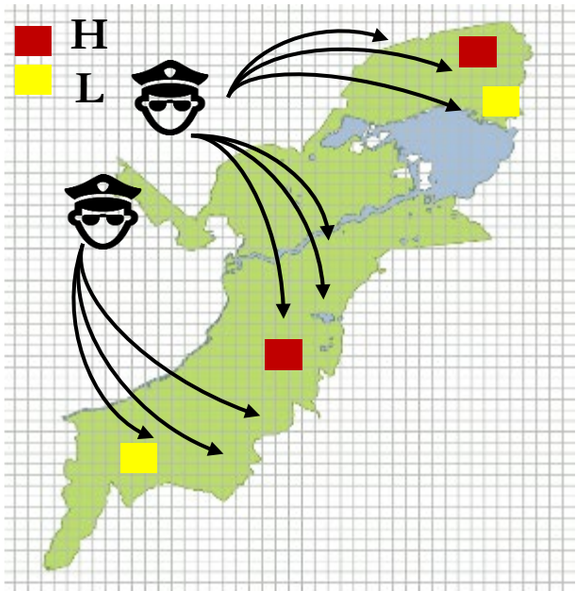

Queen Elizabeth National Park

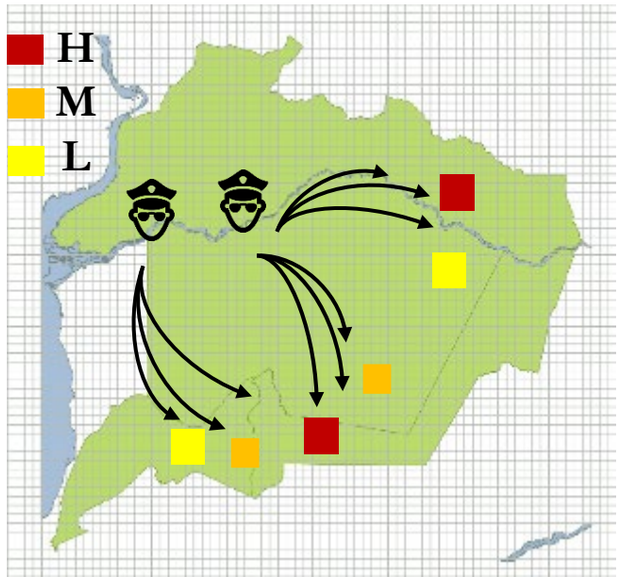

Murchison Falls National Park

Snares per patrolled sq. KM

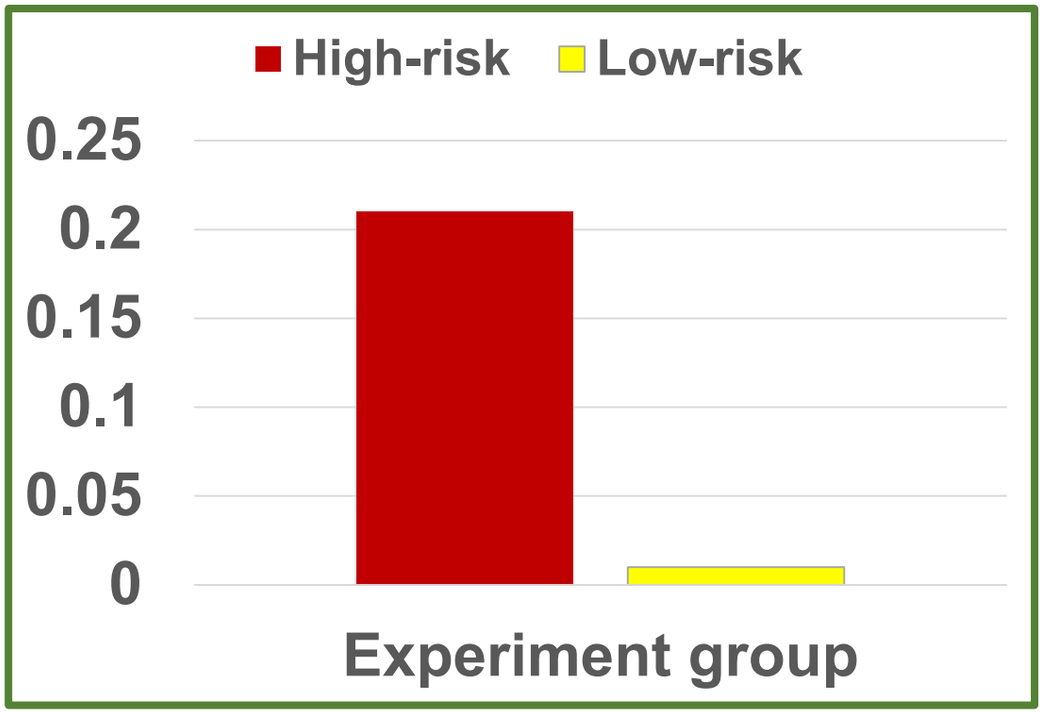

Snares per patrolled sq. KM

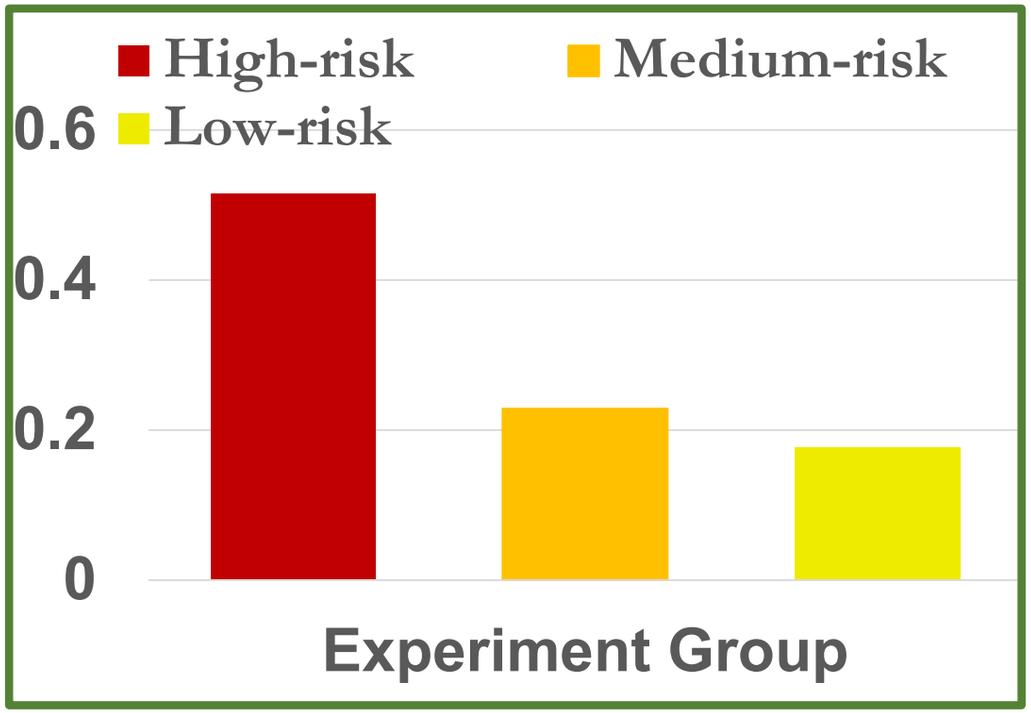

PAWS Real-world Deployment
Cambodia: Srepok Wildlife Sanctuary [2018-2019]

Xu

PAWS Real-world Deployment
Trials in Cambodia: Srepok National Park [2018-2019]

Xu

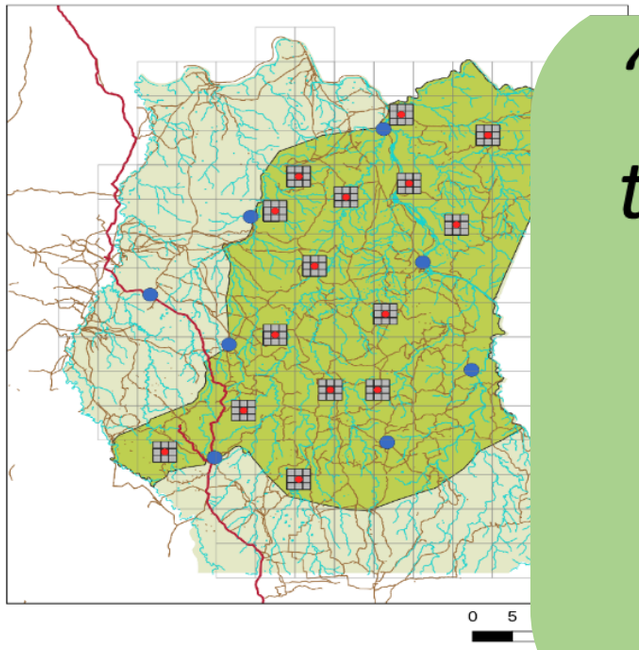

"@Milind: I am Super excited with the results. Let's get this going on other countries too this year."

Rohit Singh, WWF (2019)

- 521 snares/month our tests

vs

- 101 snares/month 2018

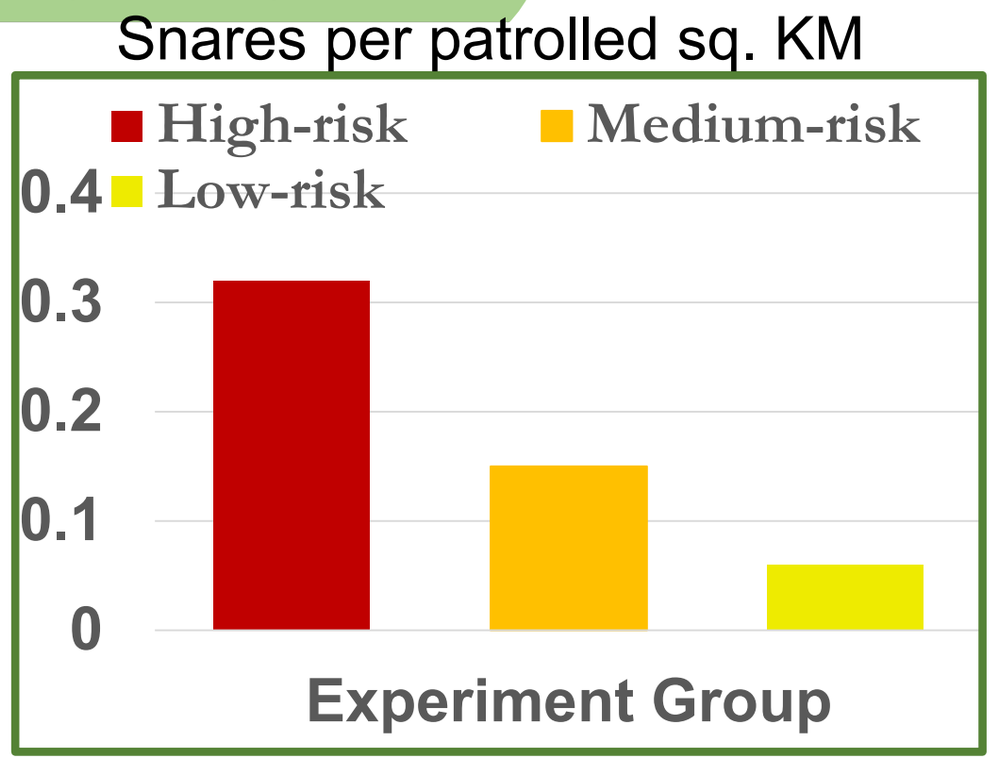

Green Security Games:
Around the Globe with SMART partnership [2019]

Protect Wildlife
800
National Parks
Around the Globe

Also: Protect Forests, Fisheries…

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

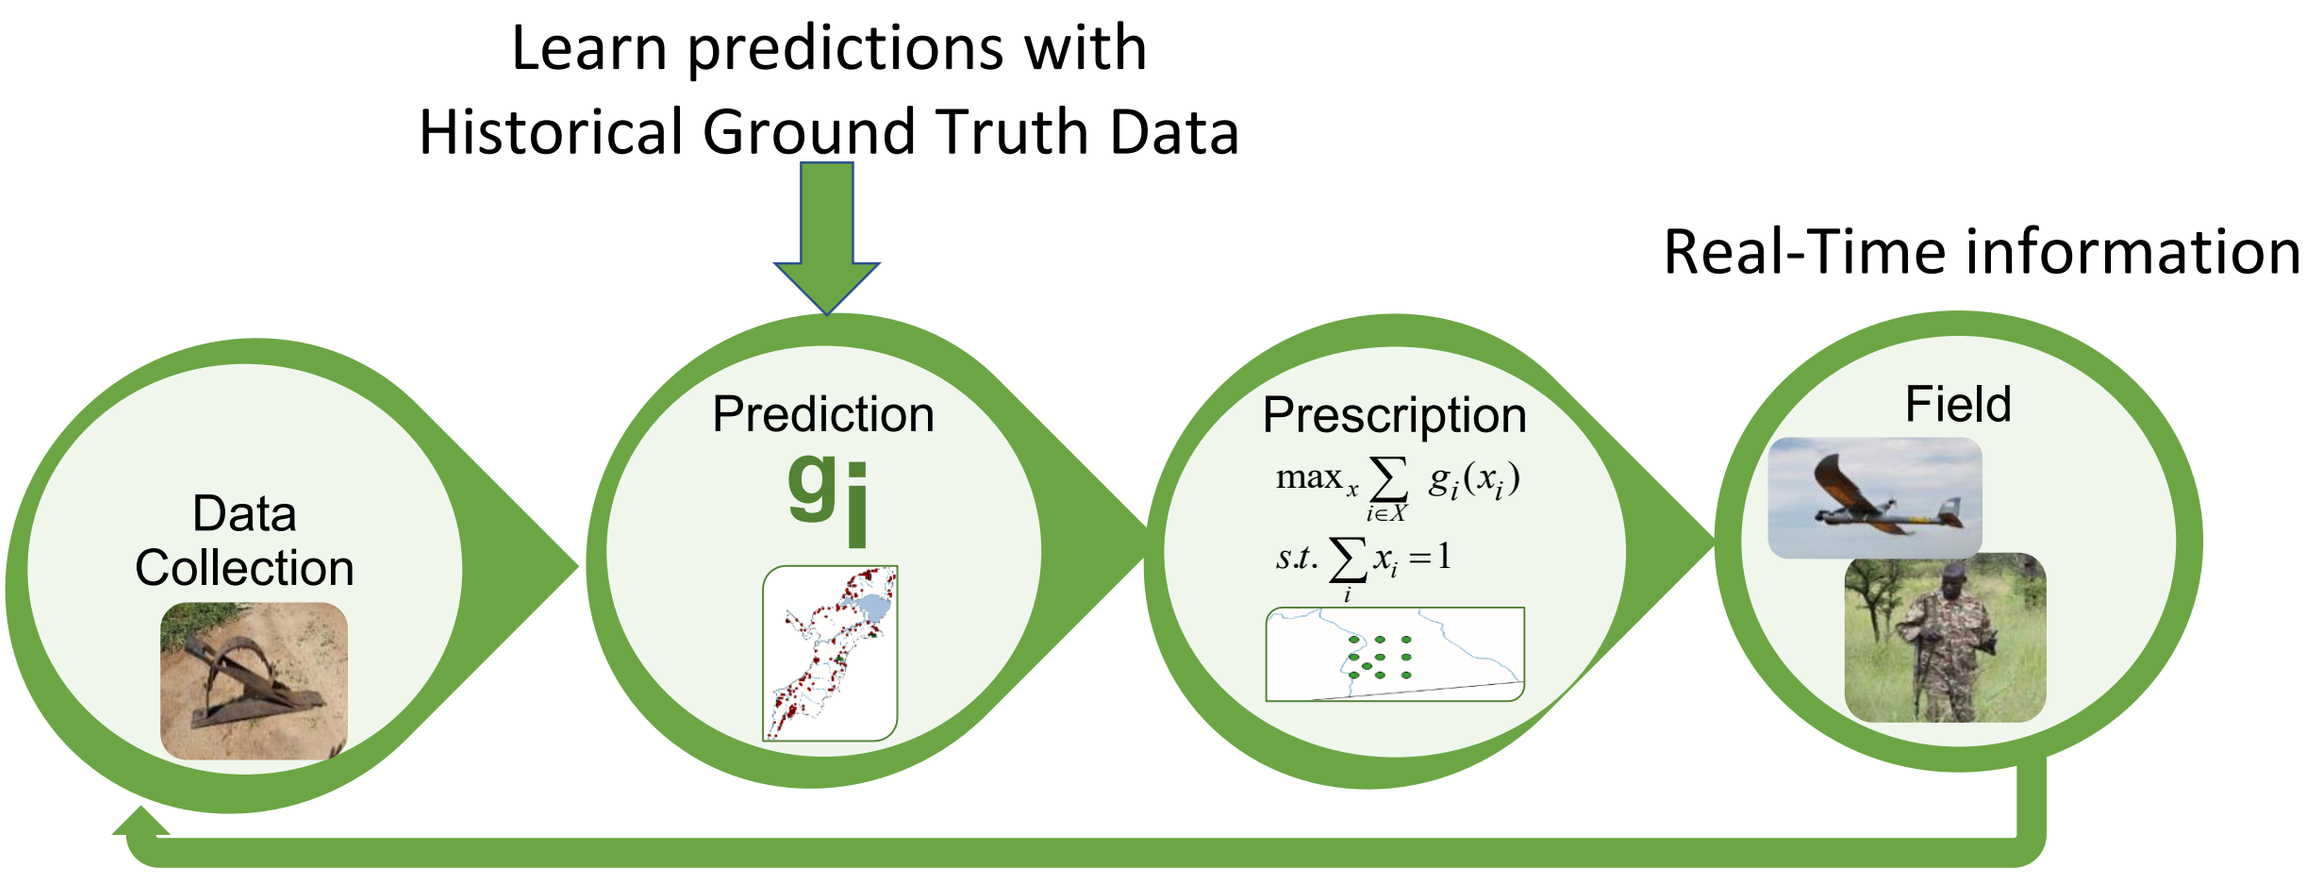

Green Security Games: Integrating Real-Time "SPOT" Information [2018]

Bondi

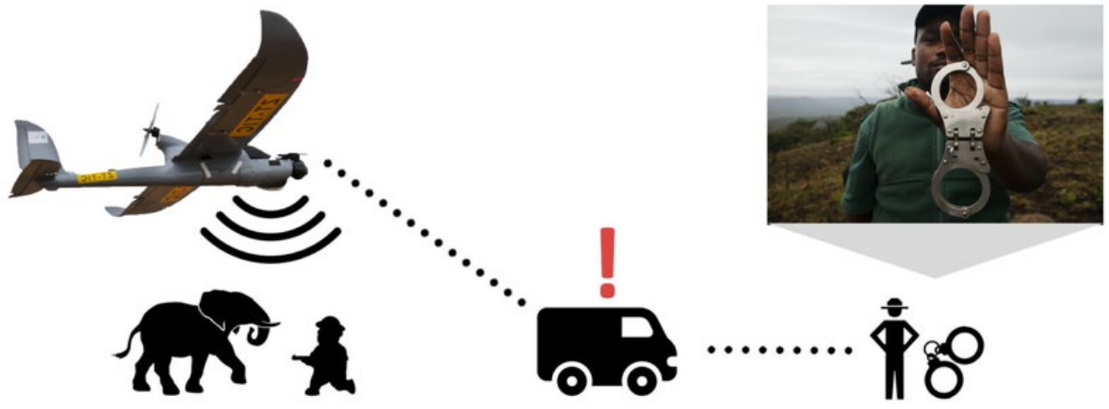

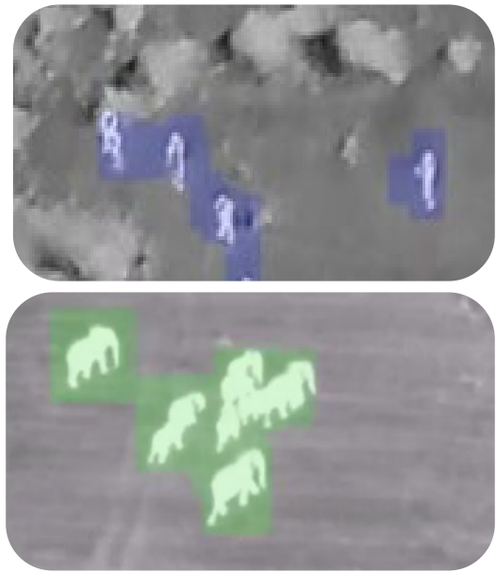

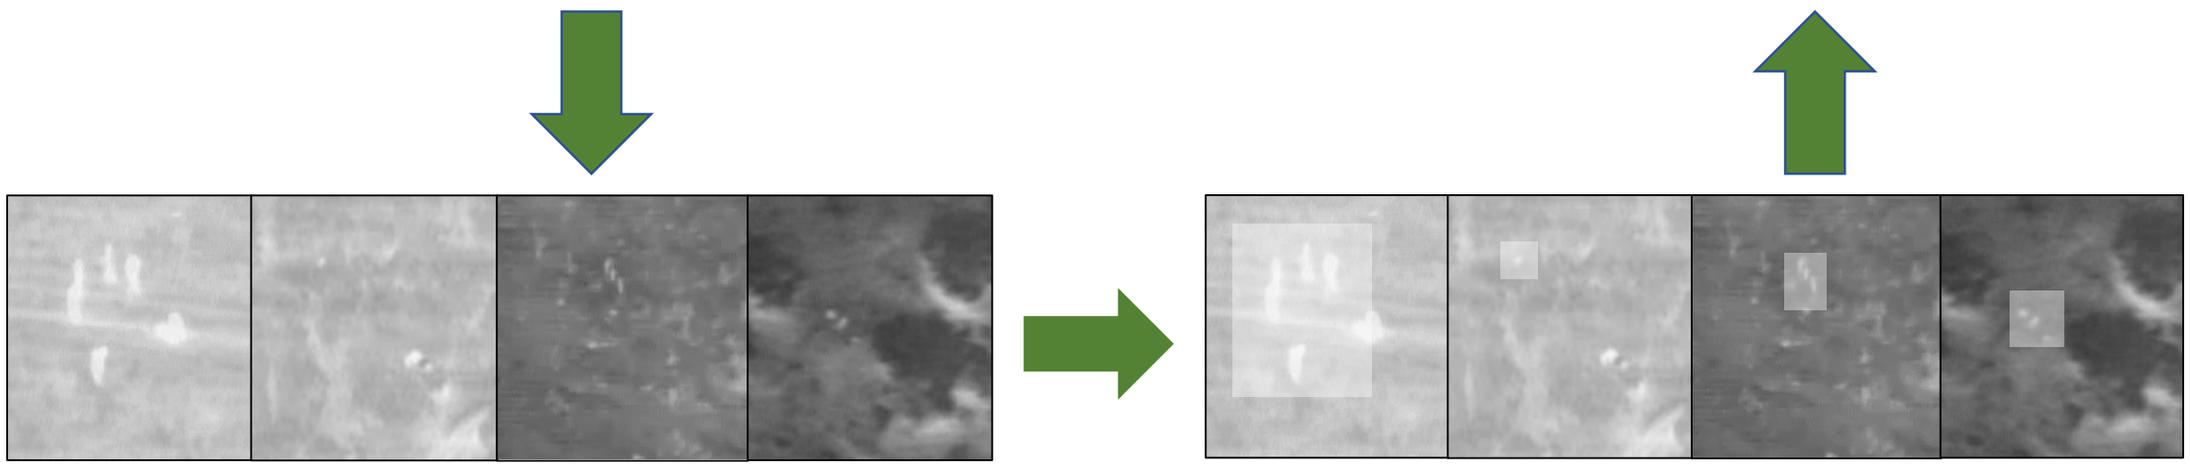

Goal: automatically find poachers

Drone Used to Inform Rangers [2019]


Xu


Bondi

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

Air Shepherd
The Lindbergh Foundation

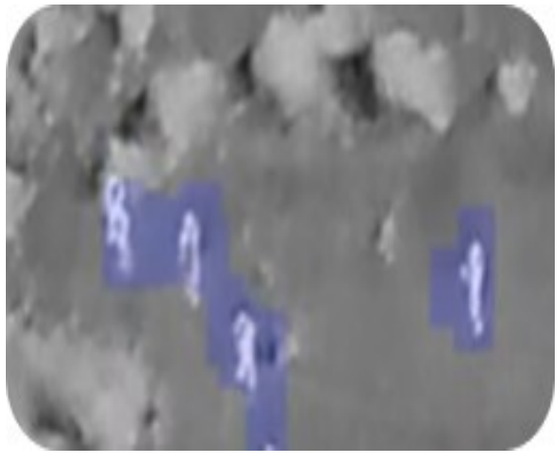

Prob(ranger) = 0.3

Drone Used to Inform Rangers [2019]


Xu


Bondi

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

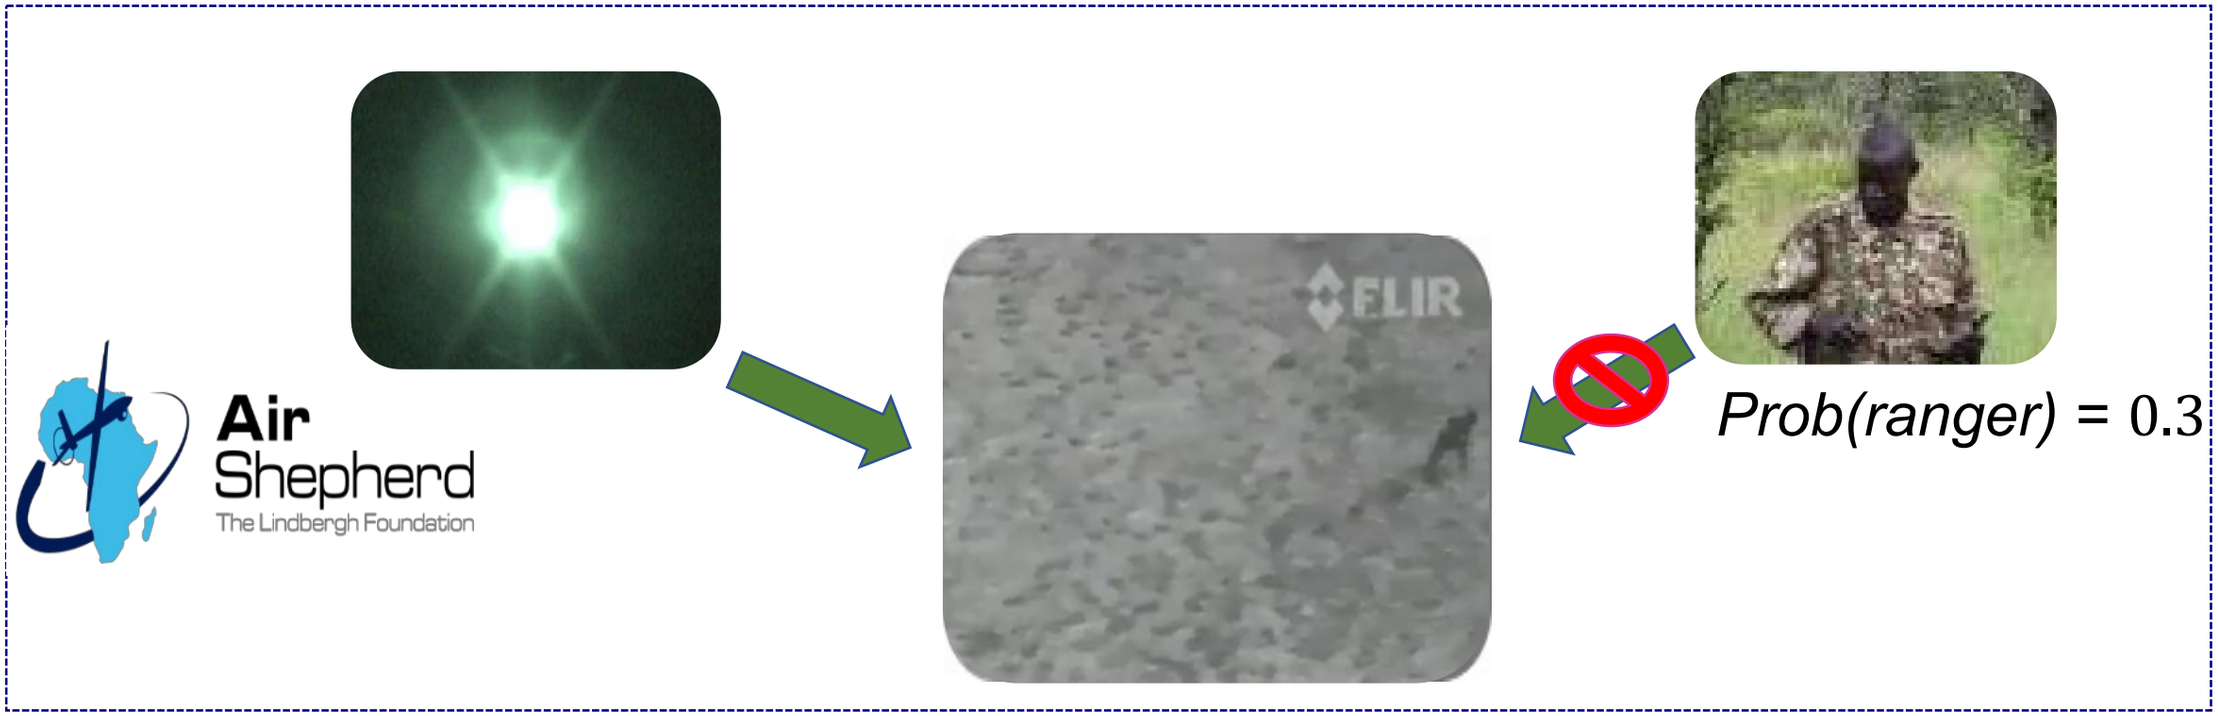

Drone Used to Inform Rangers [2019]


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

Air Shepherd
The Lindbergh Foundation

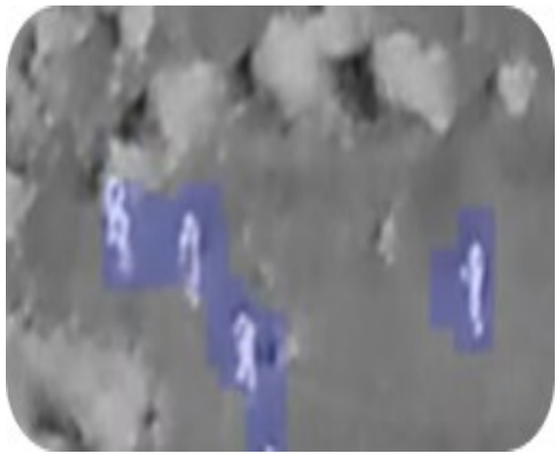

Prob(ranger) = 0.3

Strategic Signaling: Informational Advantage Defender Knows Pure & Mixed Strategy

Xu

New Model: Stackelberg Security Games with Optimal Deceptive Signaling

- Poacher best interest to "believe signal" even if know 50% time defender is lying
- Theorem: Signaling reduces complexity of equilibrium computation

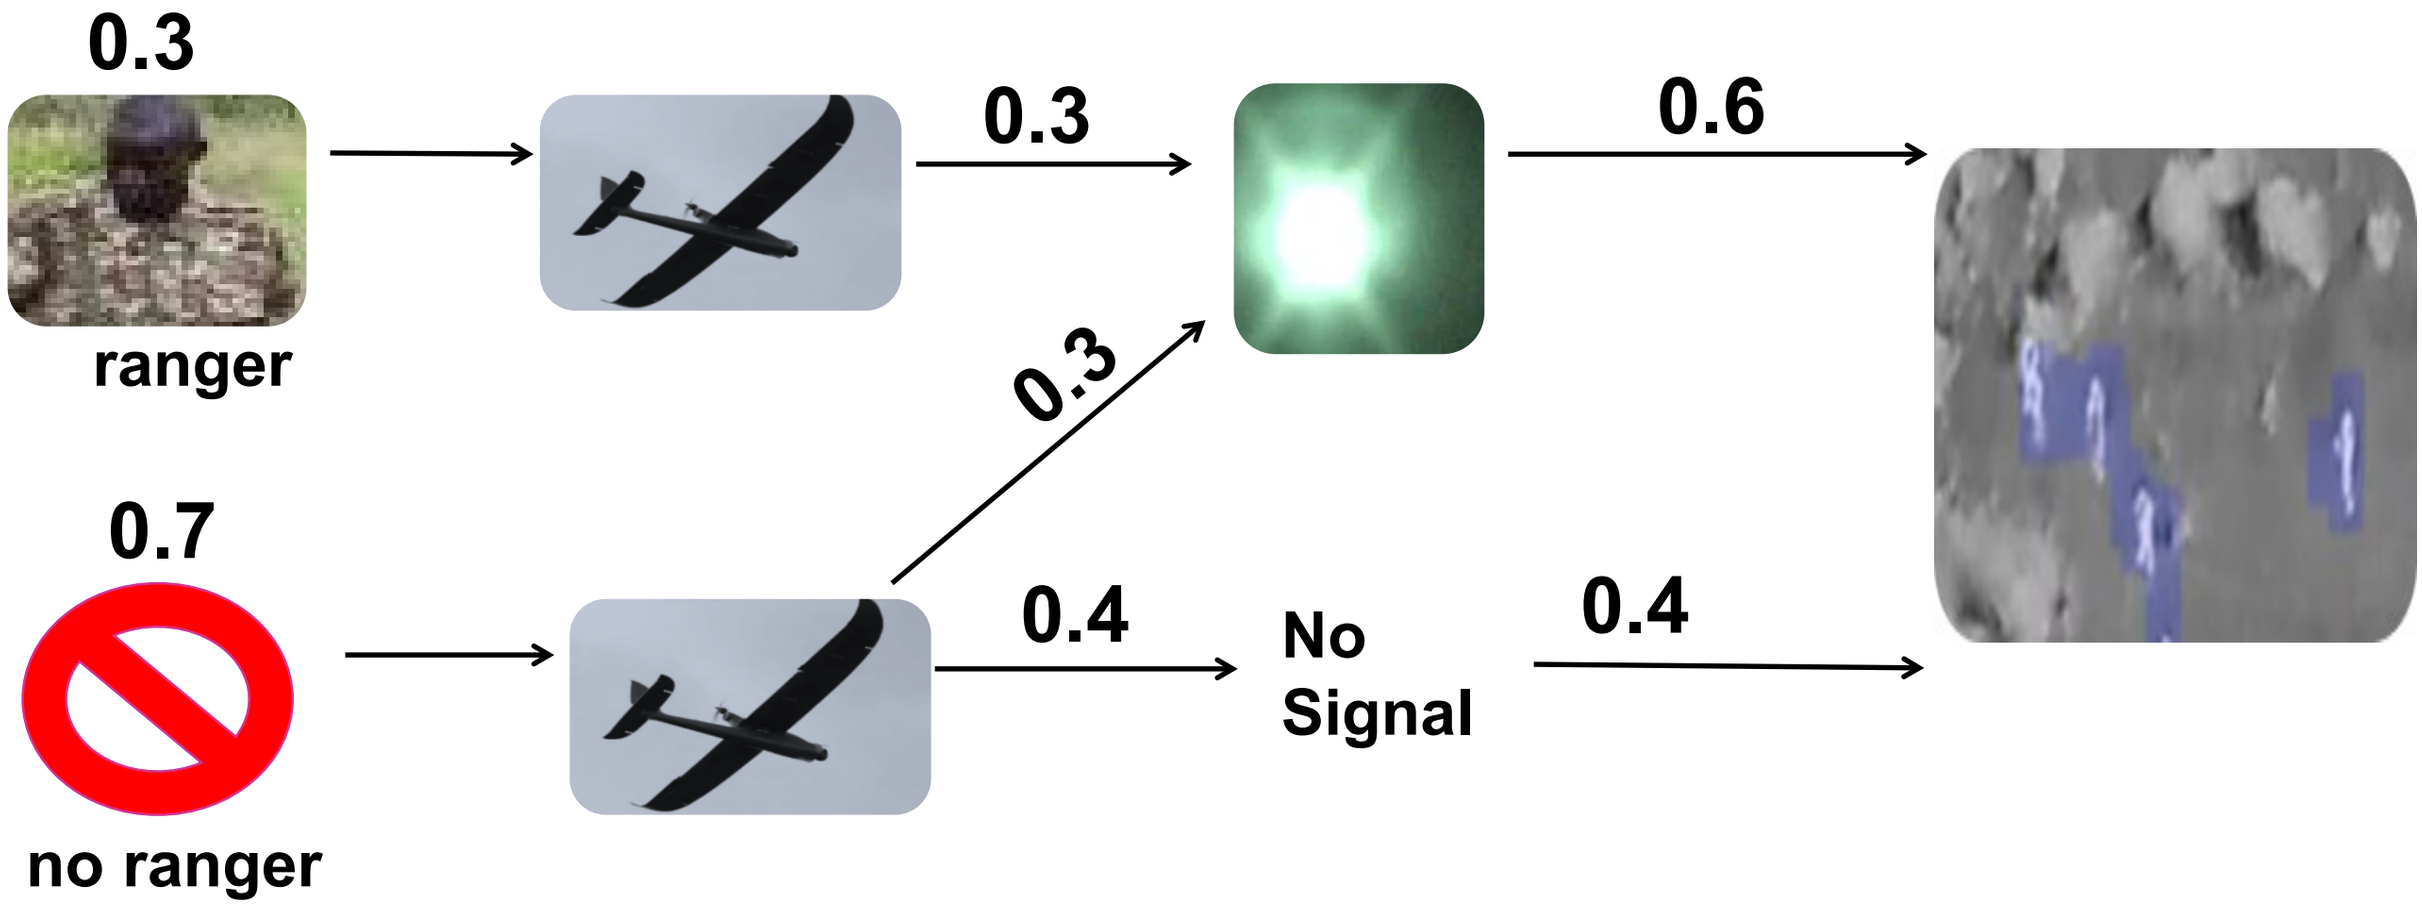

Strategic Signaling: Handling Detection Error
Exploit Informational Asymmetry to Mitigate Impact


Bondi

Strategic signaling in presence of error in detecting adversaries

- Defender knows adversary detected or not; Attacker knows Prob (detection error)
- Signal (at random) when no adversary detection

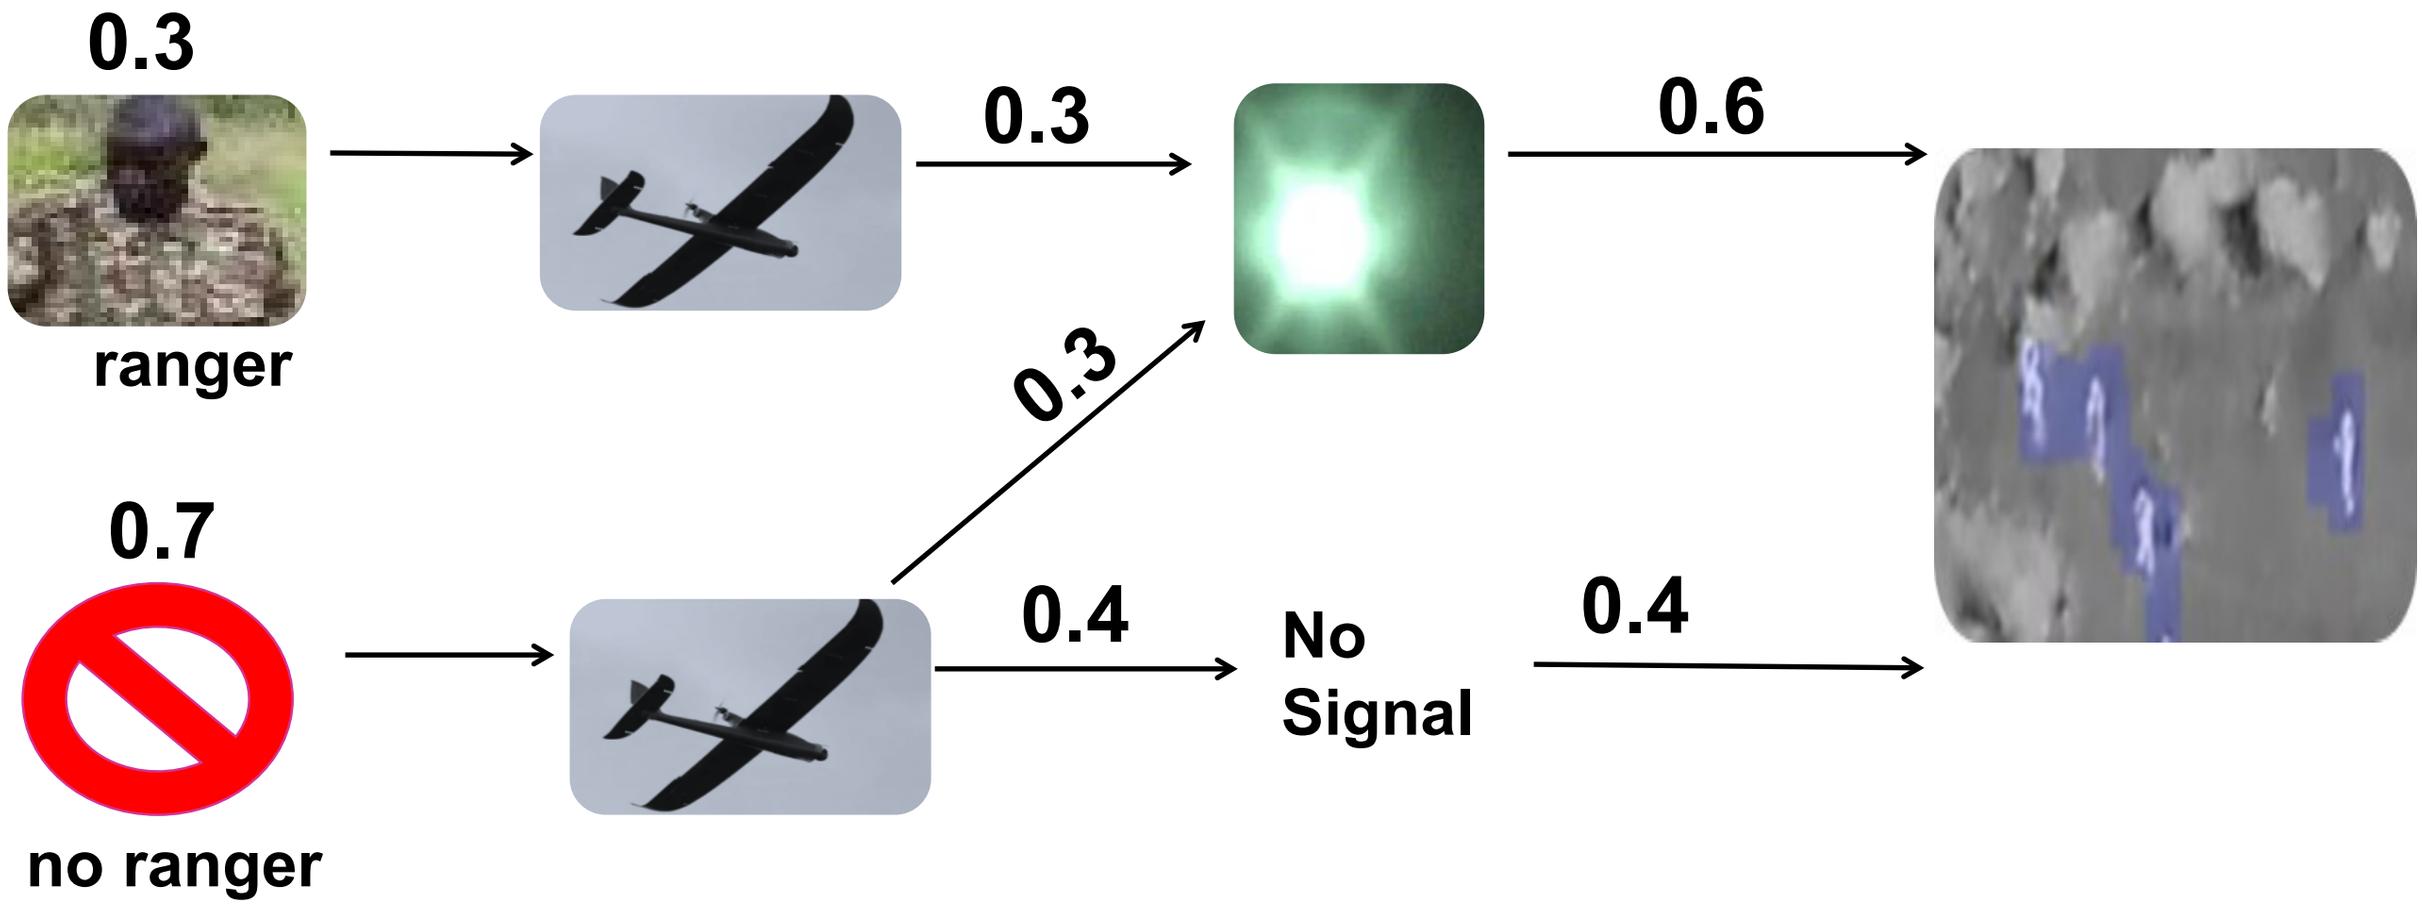

Green Security Games Pipeline: Not Enough Data in Some Parks

Wilder

- Not enough data
  - → May not learn accurate enough adversary model
  - → May lead to errors in planning patrols on targets

- *Game focused learning*
  - → Maximizing learning accuracy ≠ Maximizing decision quality
  - → Learn to maximize decision quality

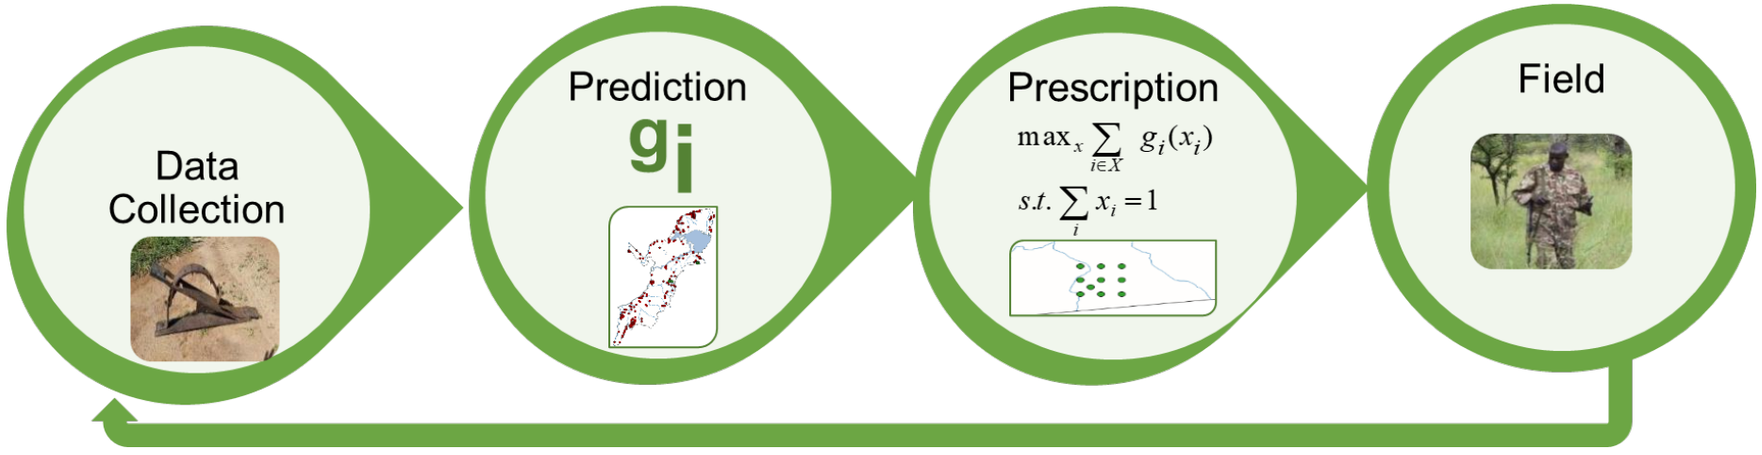

Previous Stage-by-Stage Method:

Make Prediction as Accurate as Possible Then Plan

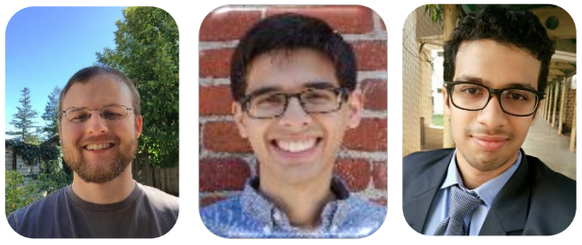

Perrault Wilder Mate

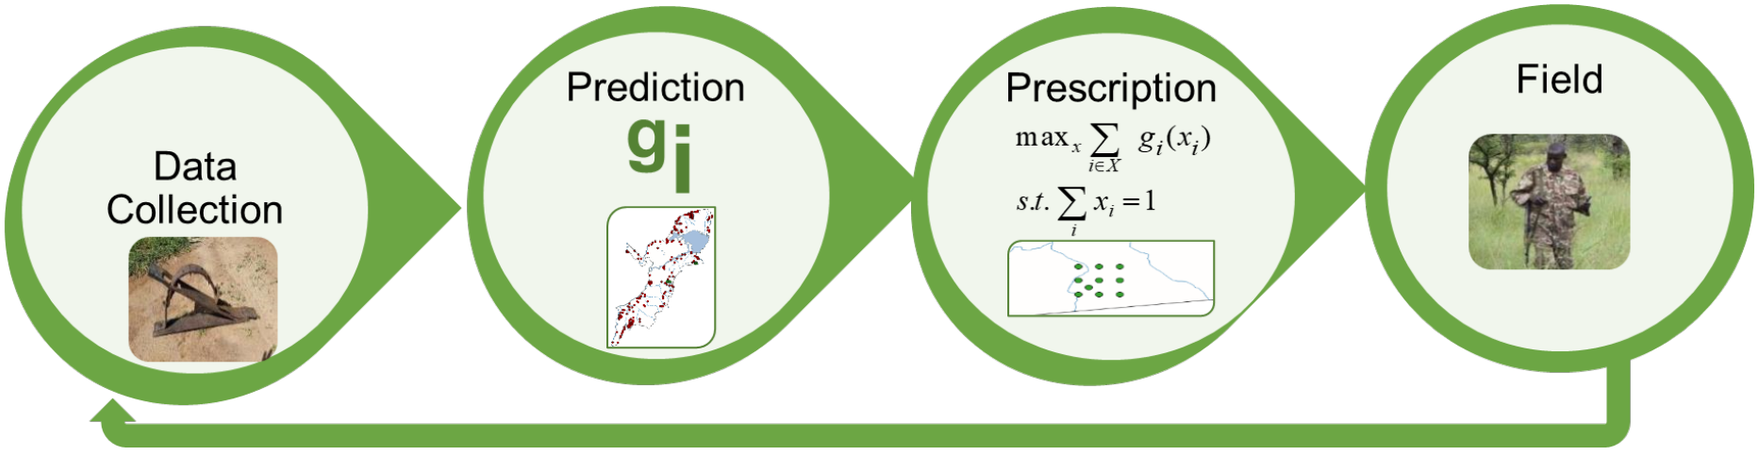

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

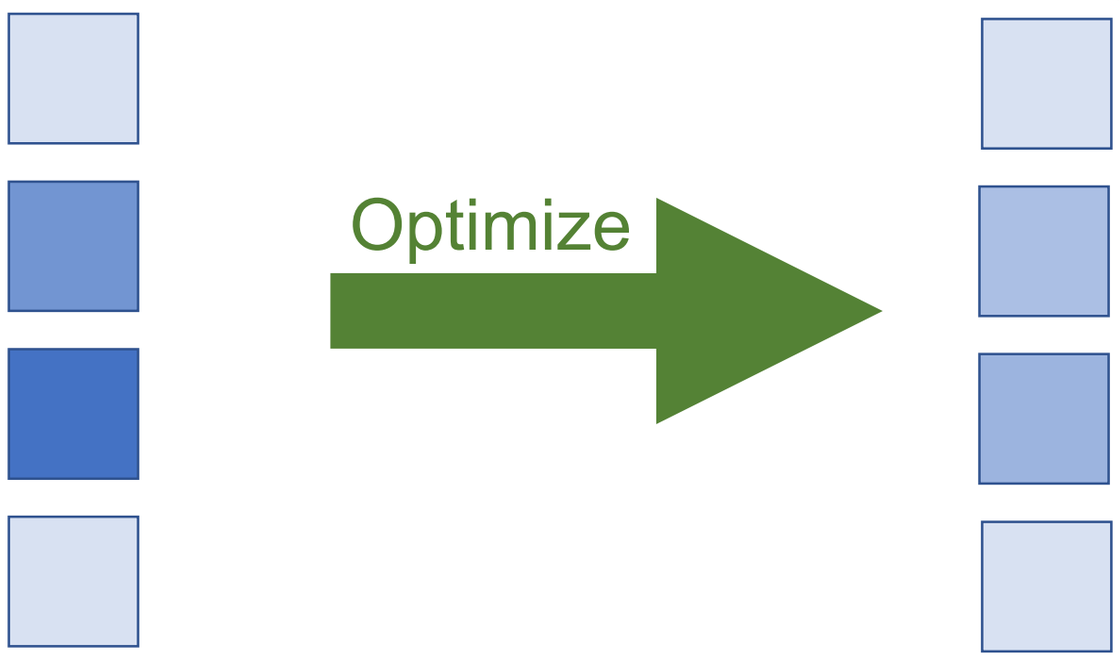

Minimize $\sum_{\{i \in T\}} q_{\text{empirical}} \log \hat{q}$

Stage by Stage Method:
Need to Focus on Important Targets

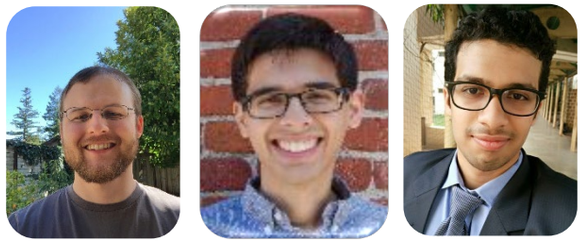

Perrault Wilder Mate

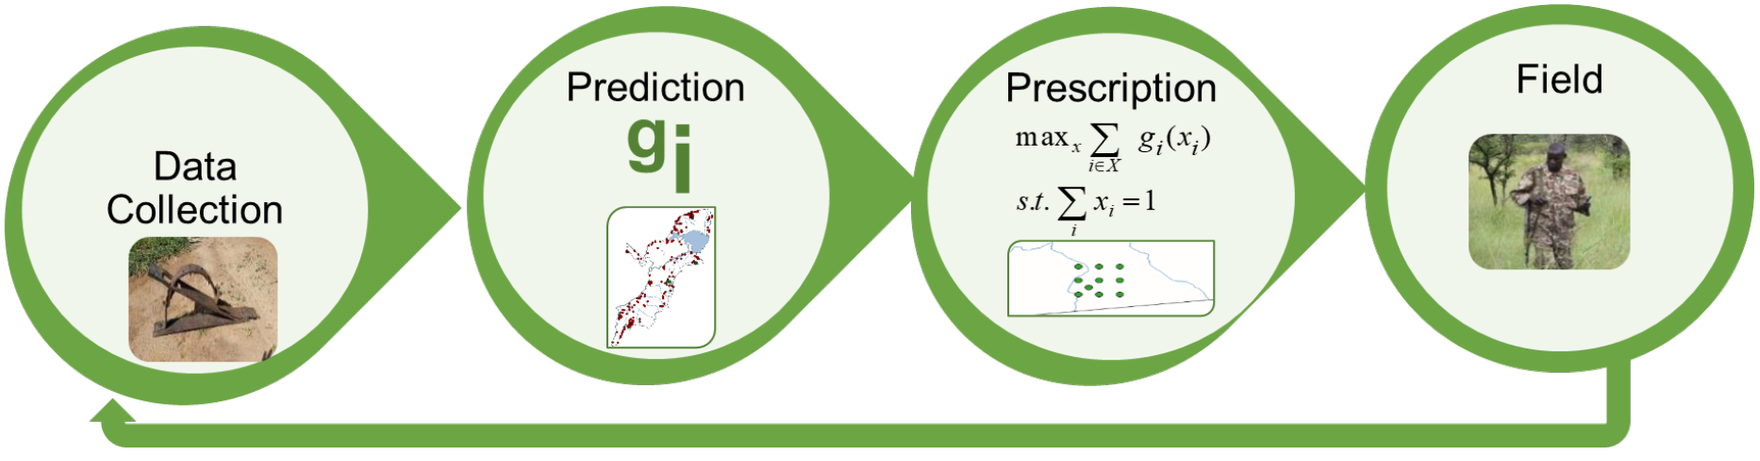

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

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

Game-Focused Learning:
Need to Focus on Important Targets



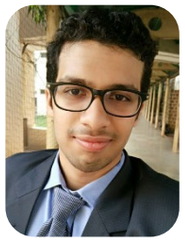

Perrault Wilder Mate

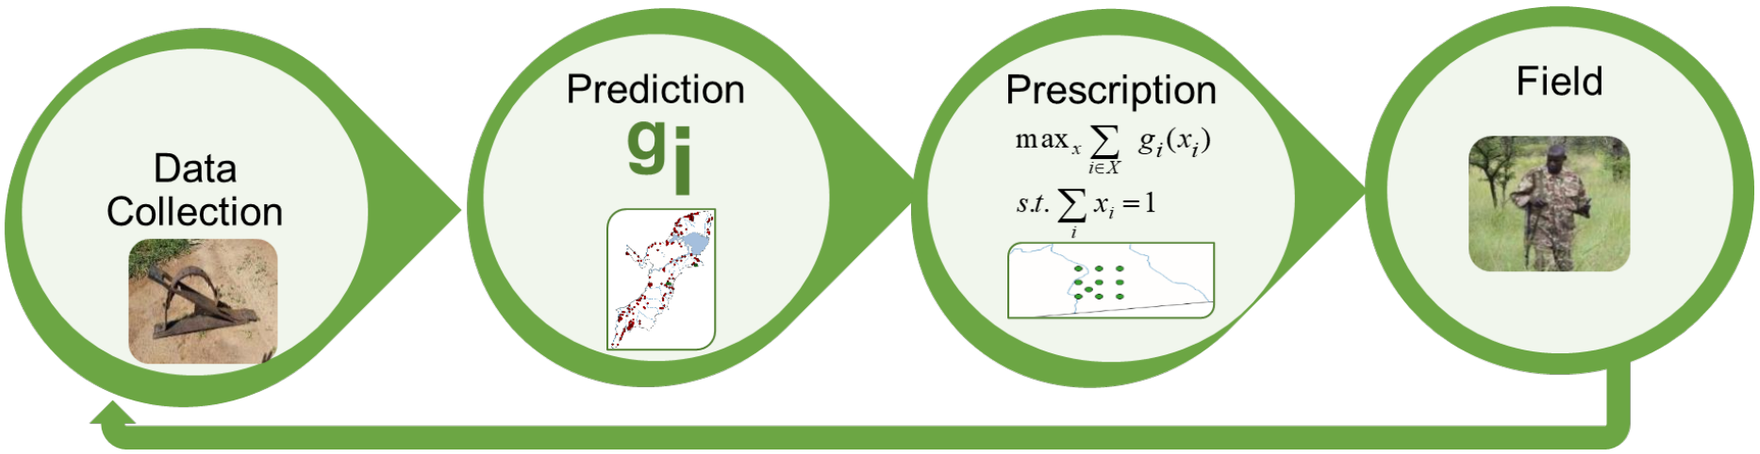

Maximize accuracy only of ***Important targets***

Plan patrol Coverage

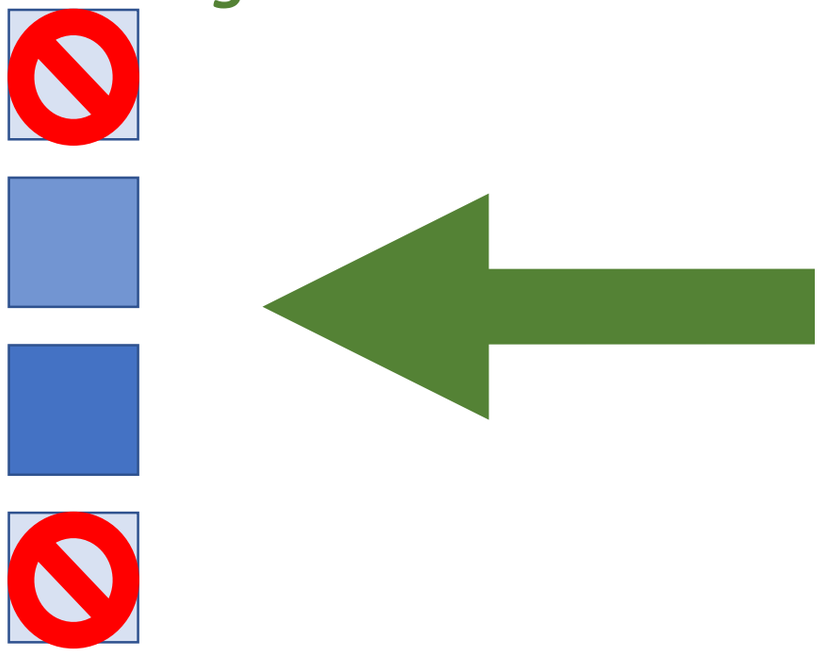

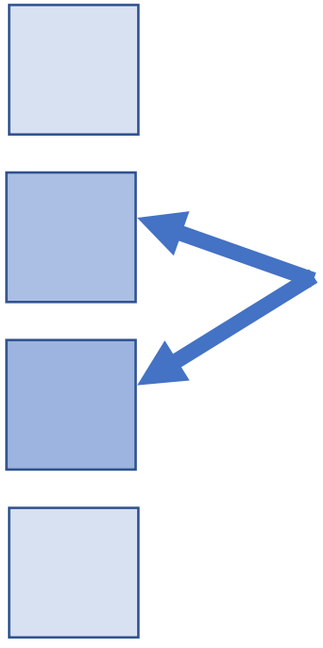

Two targets:
Large effect
Defender EU

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

Game-Focused Learning:
End-to-End Method

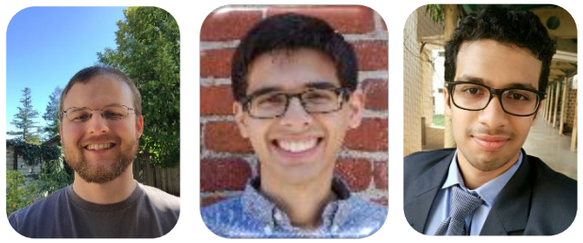

Perrault Wilder Mate

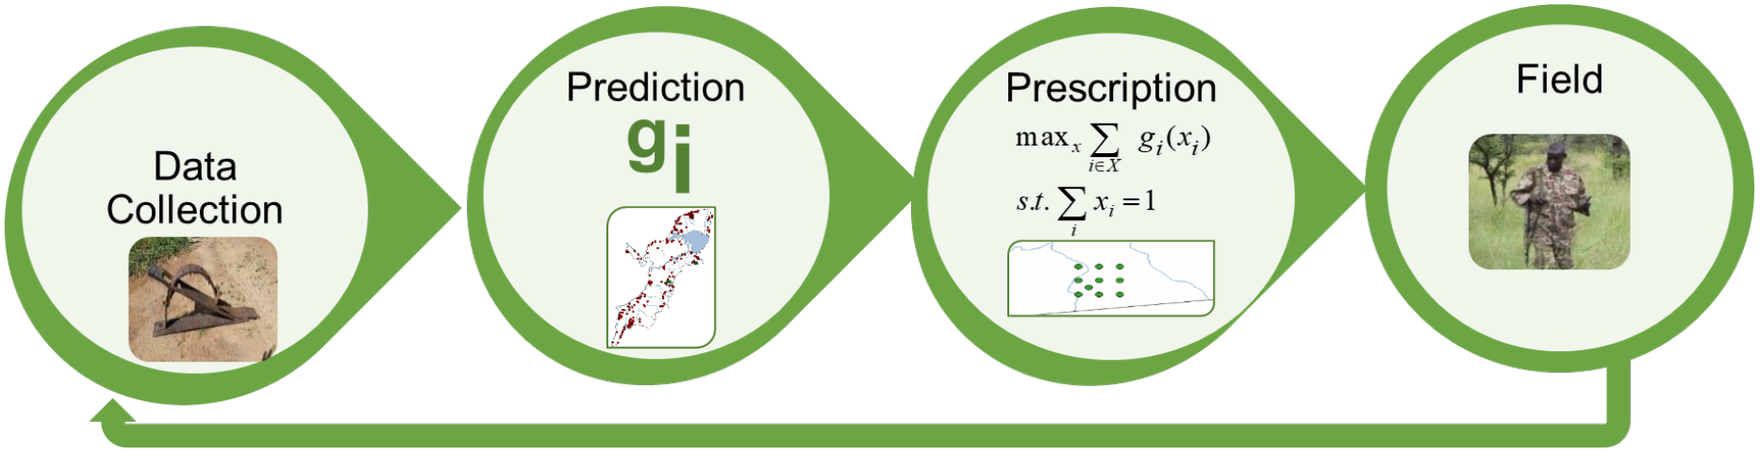

Maximize
defender expected *utility*

Plan patrol
Coverage

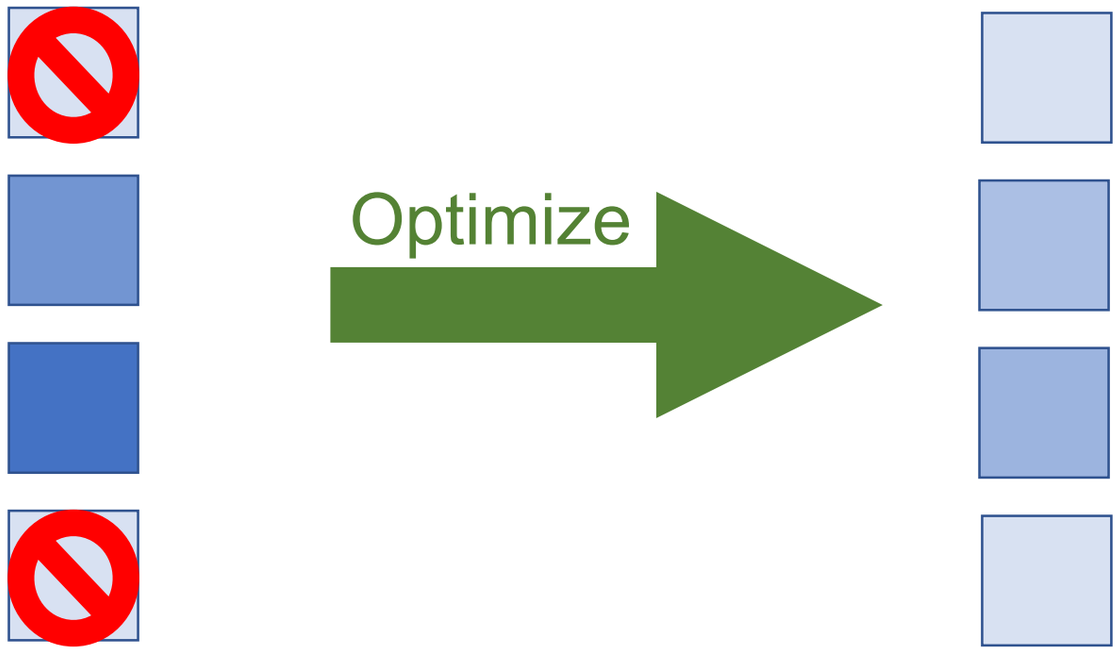

Maximize defender's expected utility

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

Previous Two-Stage Method: Gradient Descent

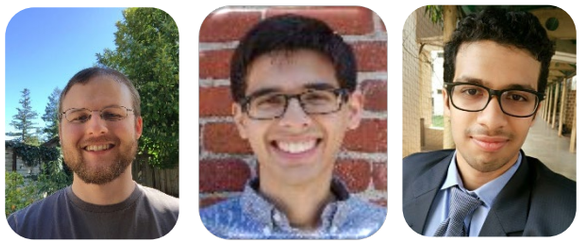

Perrault Wilder Mate

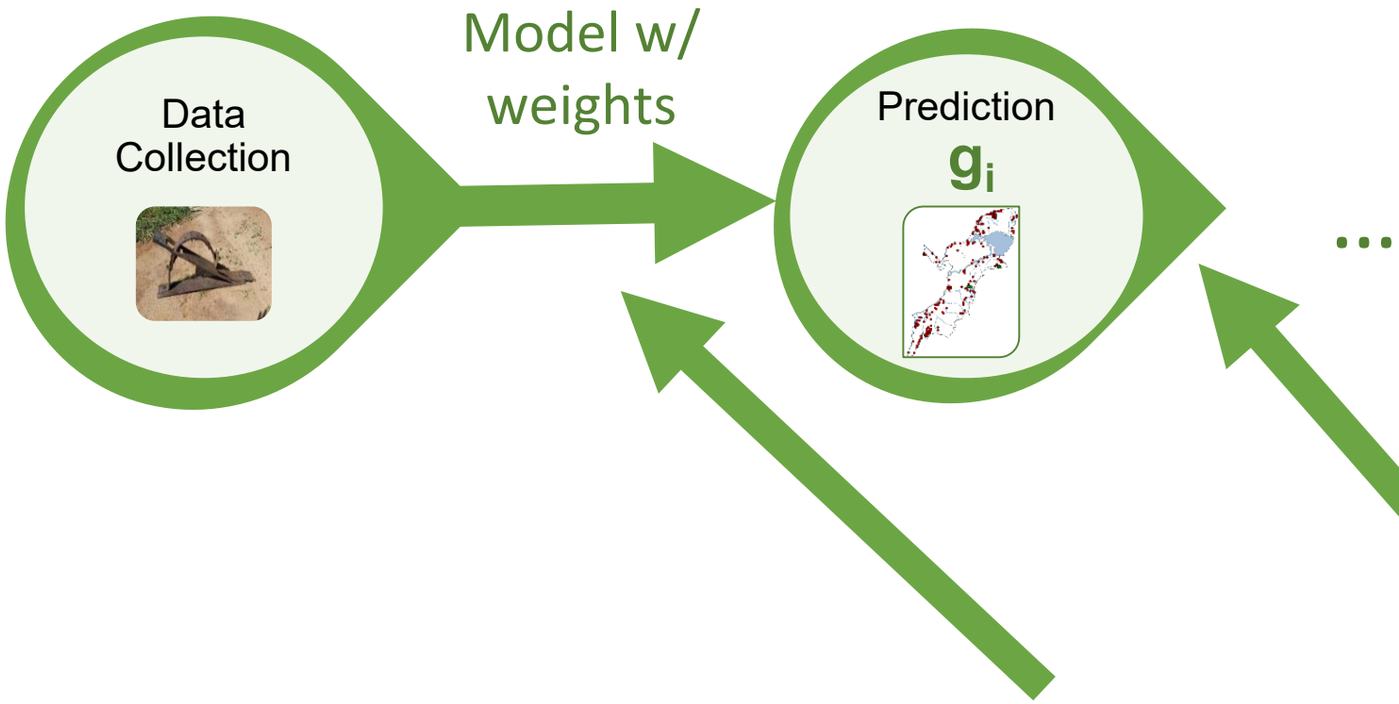

- Max accuracy gradient descent:

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

Game-Focused Learning: End-to-End Method


Perrault


Wilder

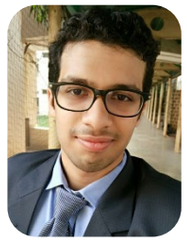
Mate

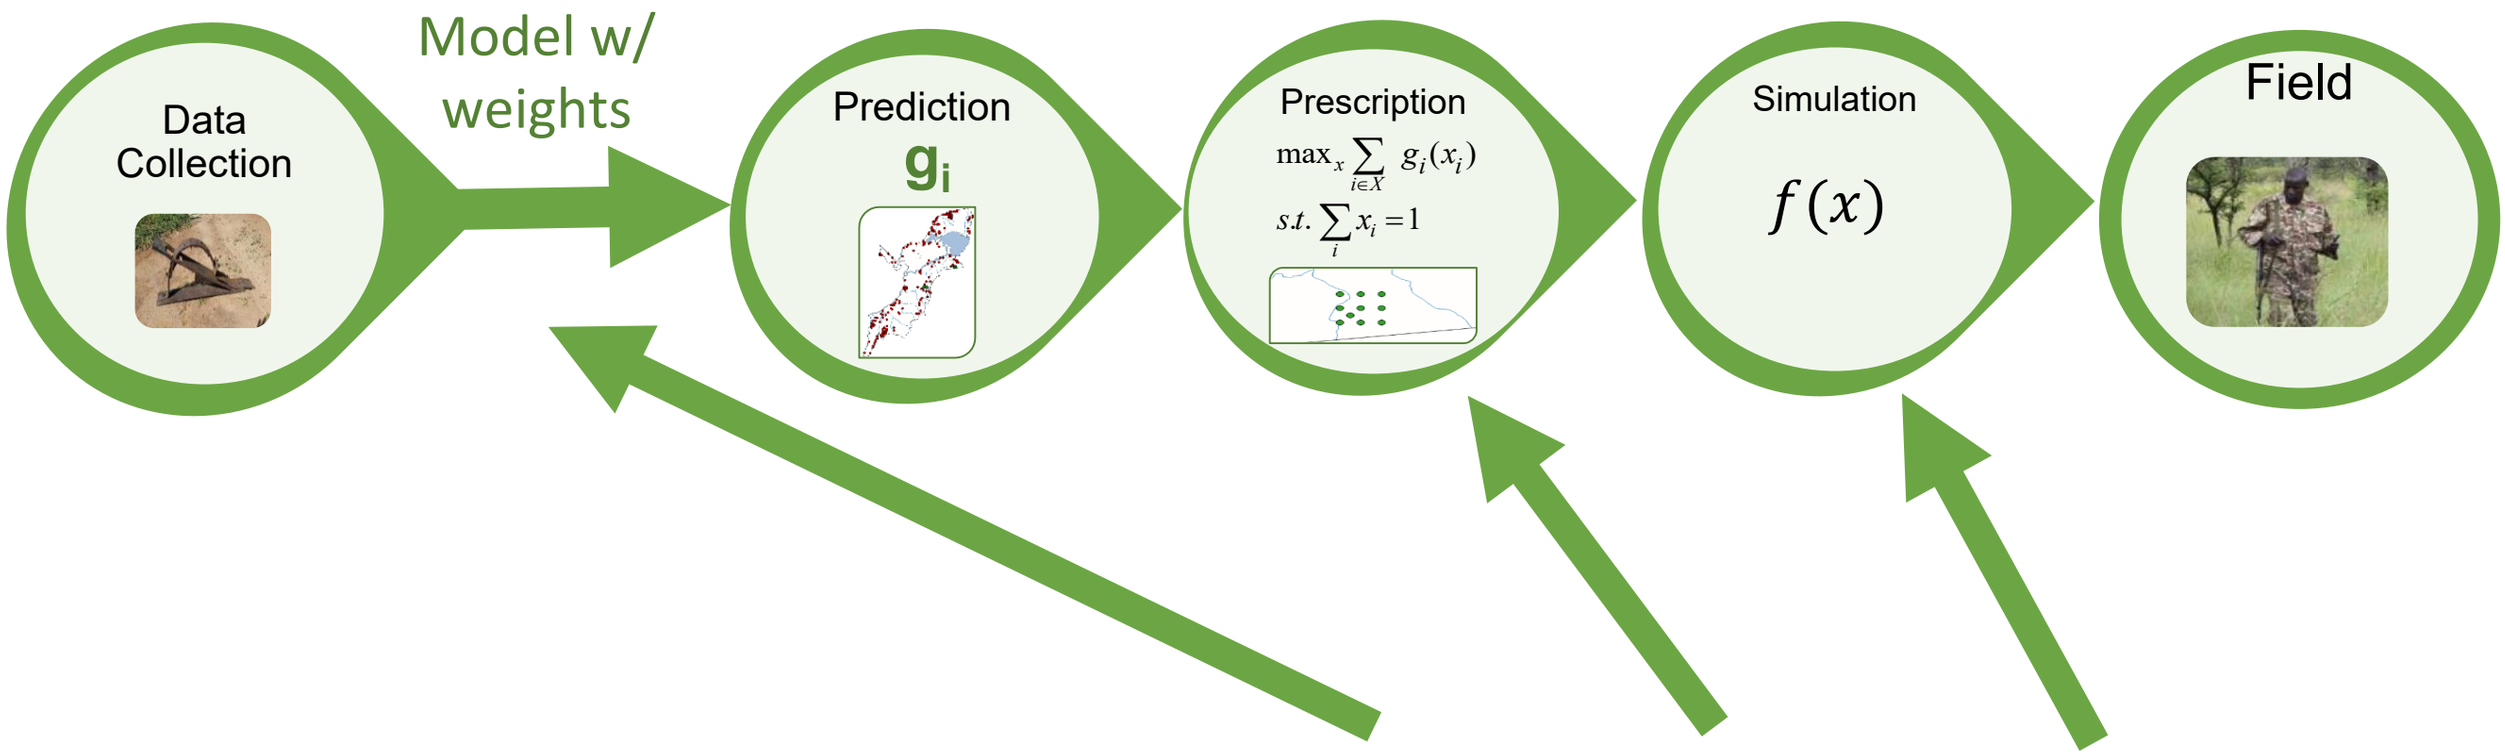

- Game-focused gradient descent:

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

Game-Focused Learning:
End-to-End Method

- Focusing learning on important targets increases defender utility

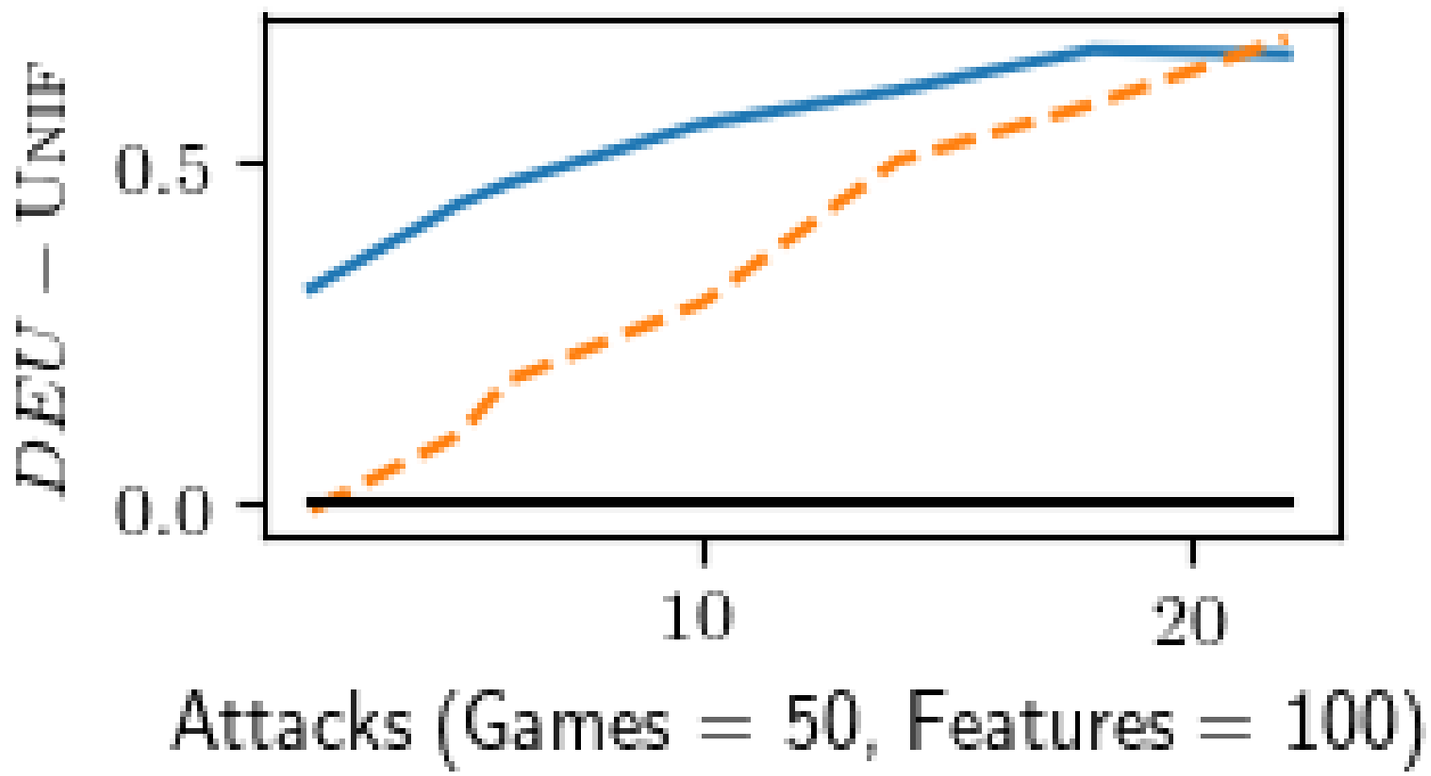

Outline

Public Safety & Security: Stackelberg Security Games

Conservation/Wildlife Protection: Green Security Games

Public Health: Game against nature

Prof Eric Rice
Social Work

Public Health
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" social networks gathered from observations in the field; not facebook etc*

Influence Maximization Background

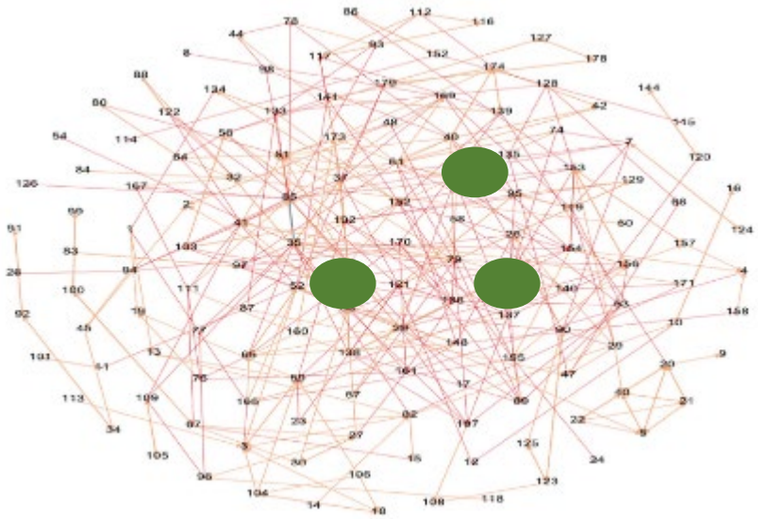

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
- Objective:
  - Maximize expected number of influenced nodes
- *Assumption: Independent cascade model of information spread*

Independent Cascade Model and Real-world Physical Social Networks

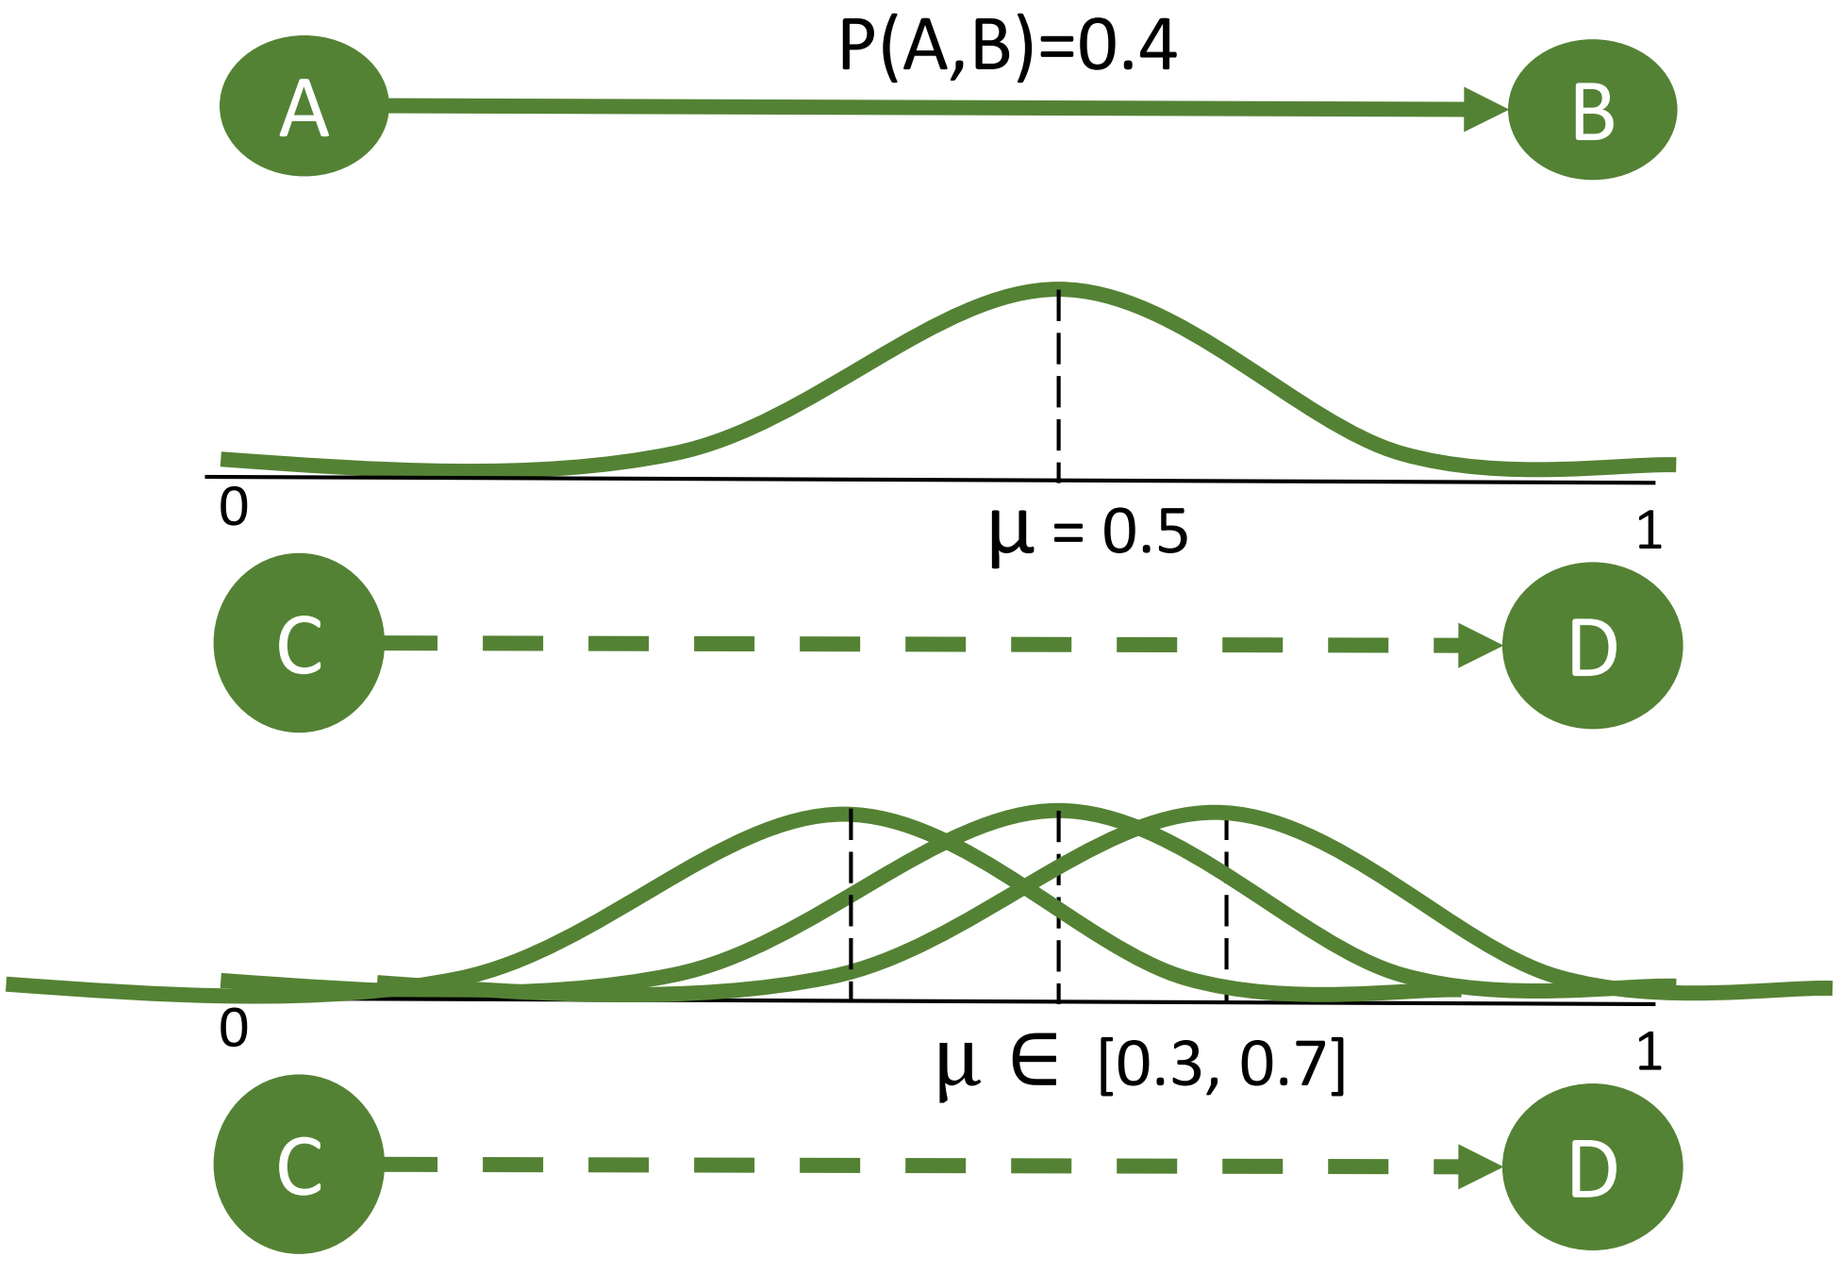

Robust, Dynamic Influence Maximization


Wilder

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

Algorithm
Chooses policy, i.e.,
Chooses Peer leaders

vs

Nature
Chooses parameters
μ,σ

- Payoffs: (performance of algorithm)/OPT

HEALER Algorithm [2017]
Robust, Dynamic Influence Maximization


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: Multi-step Policy


Yadav


Wilder

	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

K = 4
1st time step

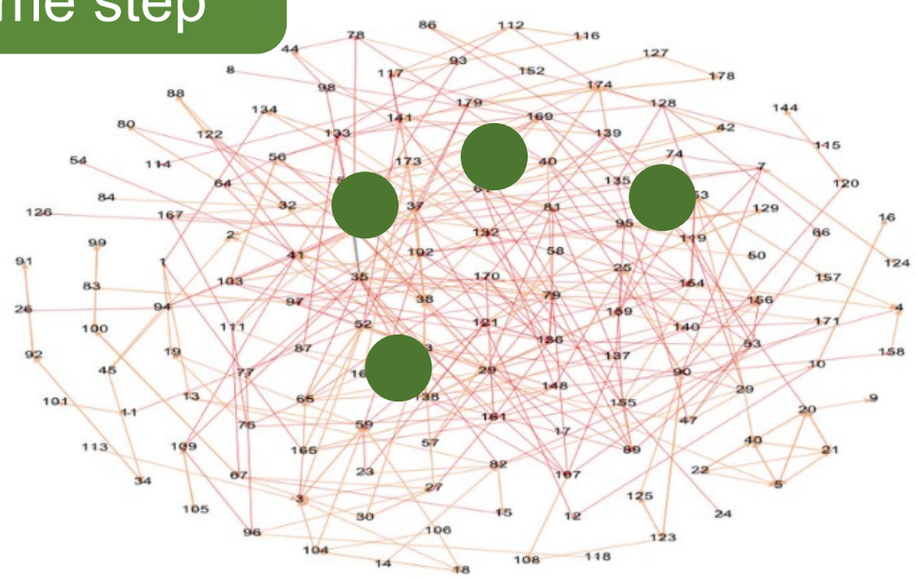

K = 4
2nd time step

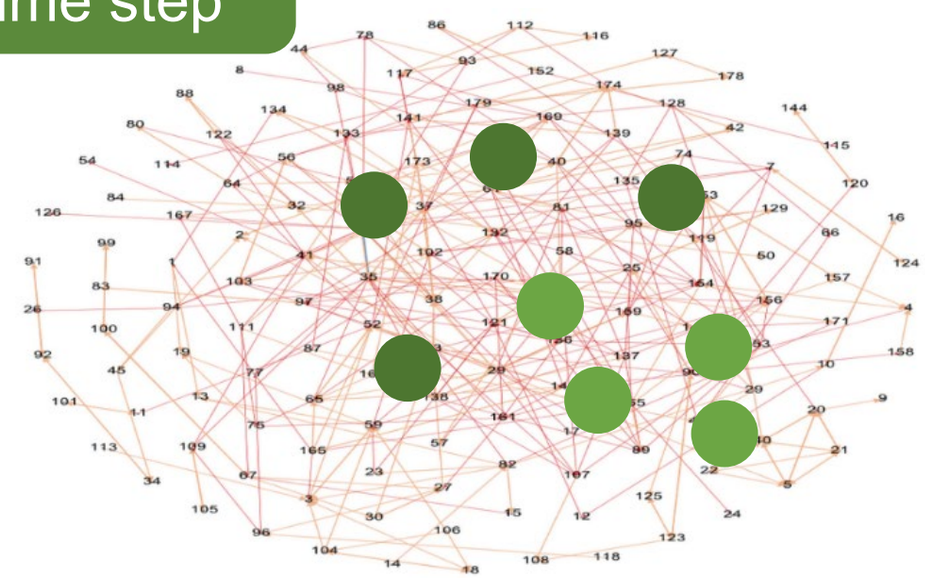

HEALER: POMDP Model for Multi-Step Policy Robust, Dynamic Influence Maximization

Yadav

	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

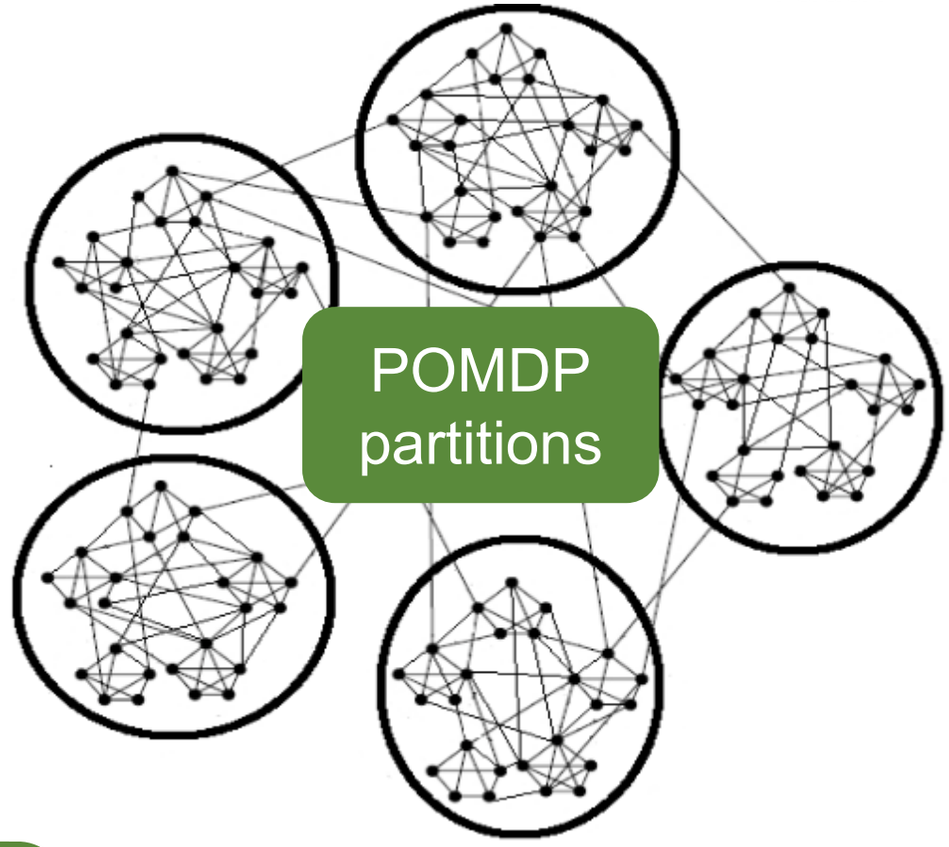

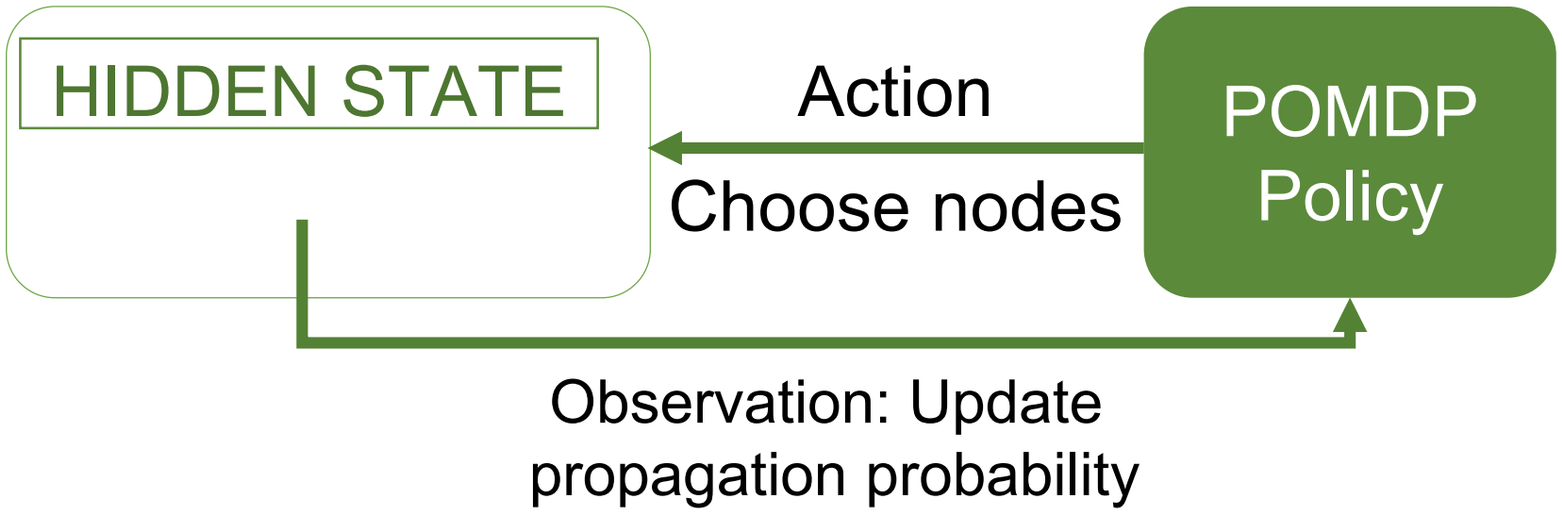

Pilot Tests with HEALER with 170 Homeless Youth [2017]


Yadav


Wilder

Recruited youths:

HEALER	HEALER++	DEGREE CENTRALITY
62	56	55

12 peer leaders

Results: Pilot Studies [2017]


Yadav


Wilder

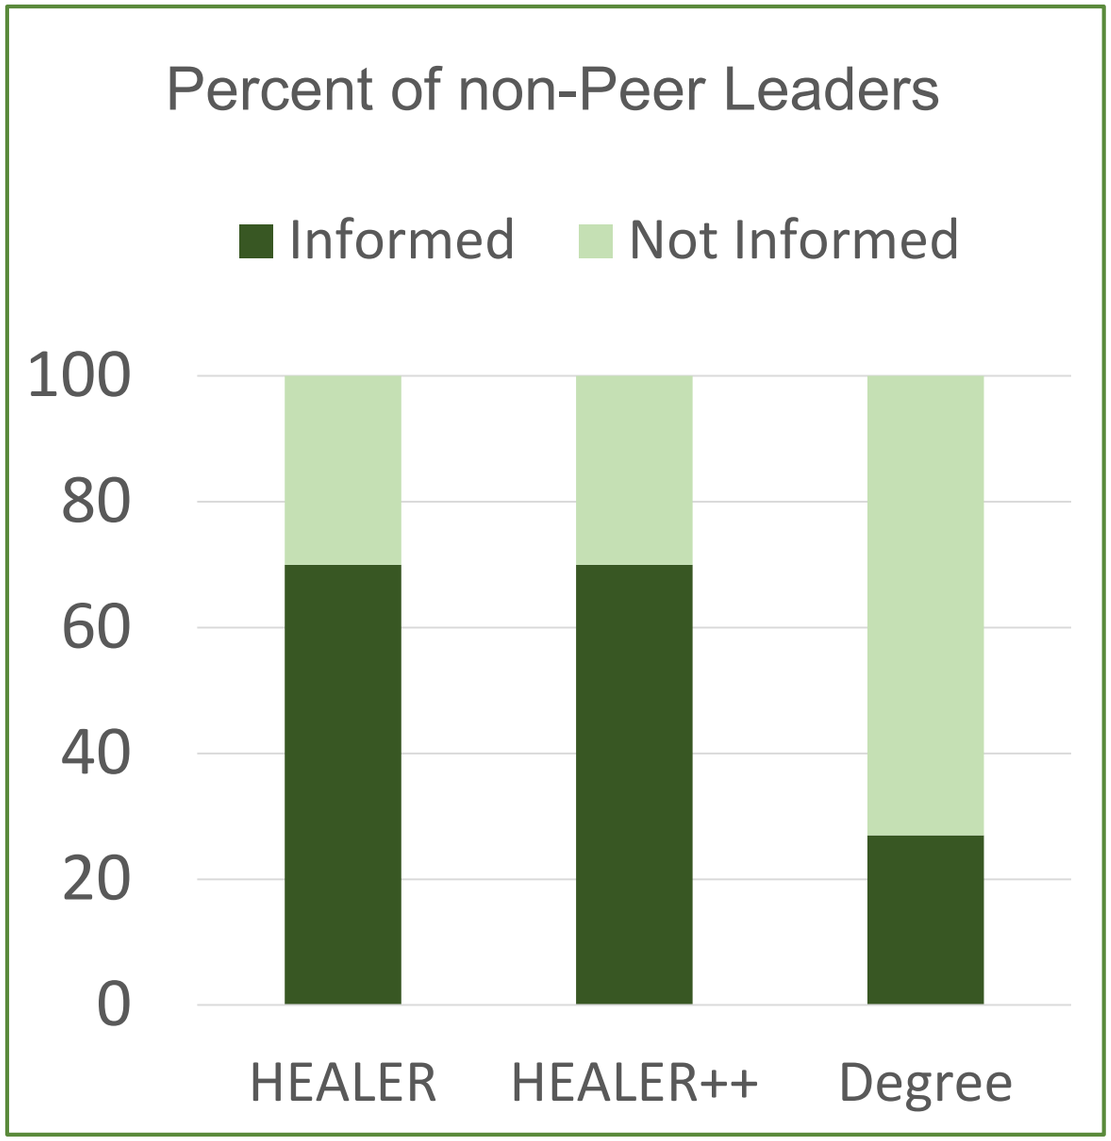

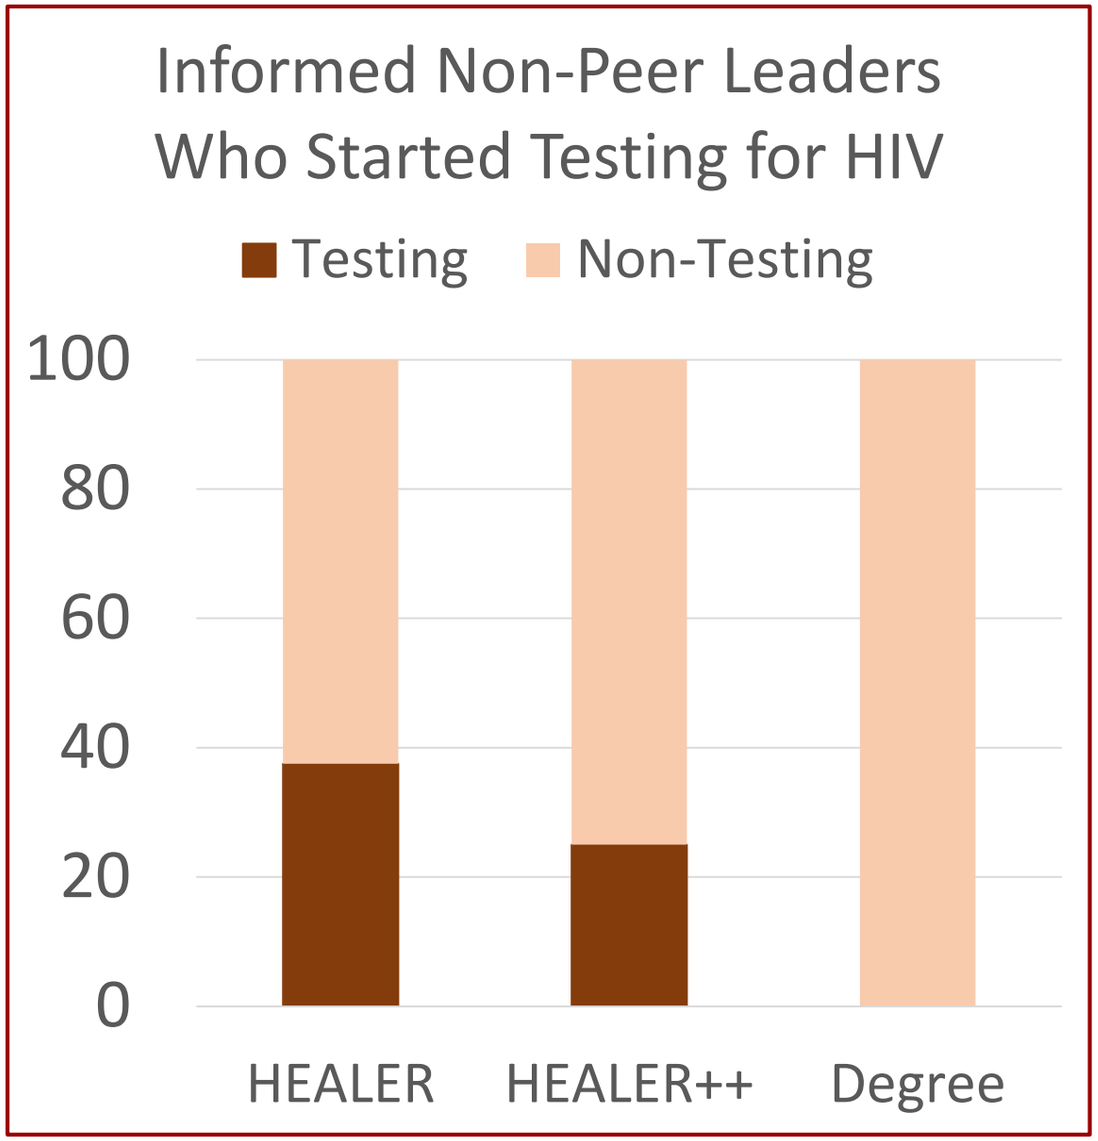

More details: Journal of Society of Social Work & Research (Nov 2018)

Data to Deployment Pipeline:
Network Sampling to avoid Data Collection Bottleneck


Wilder

Data collection costly

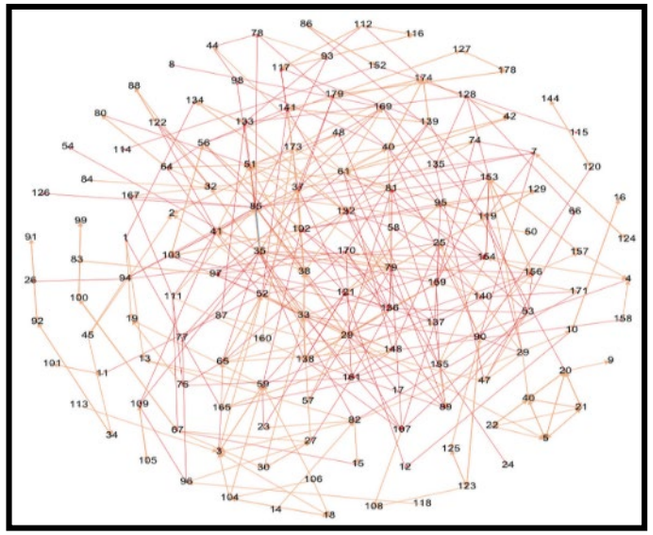

Sample 18%

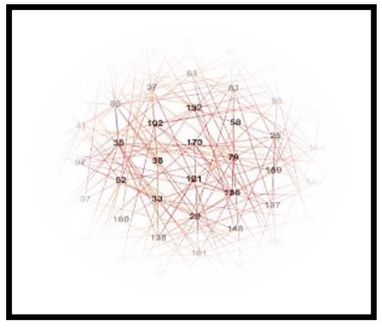

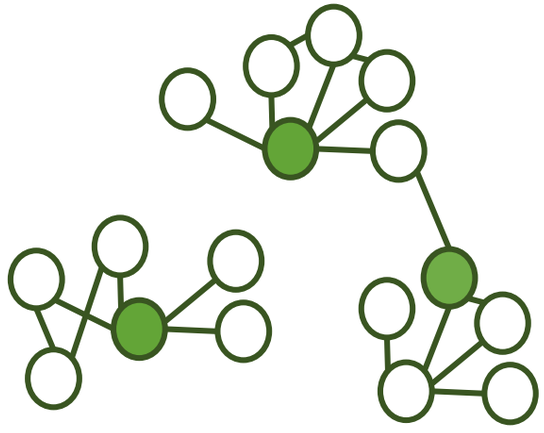

Sampling from largest communities

New experiment With 60 homeless youth

12 peer leaders

Results: Pilot Studies with New Sampling Algorithm [2018]


Wilder

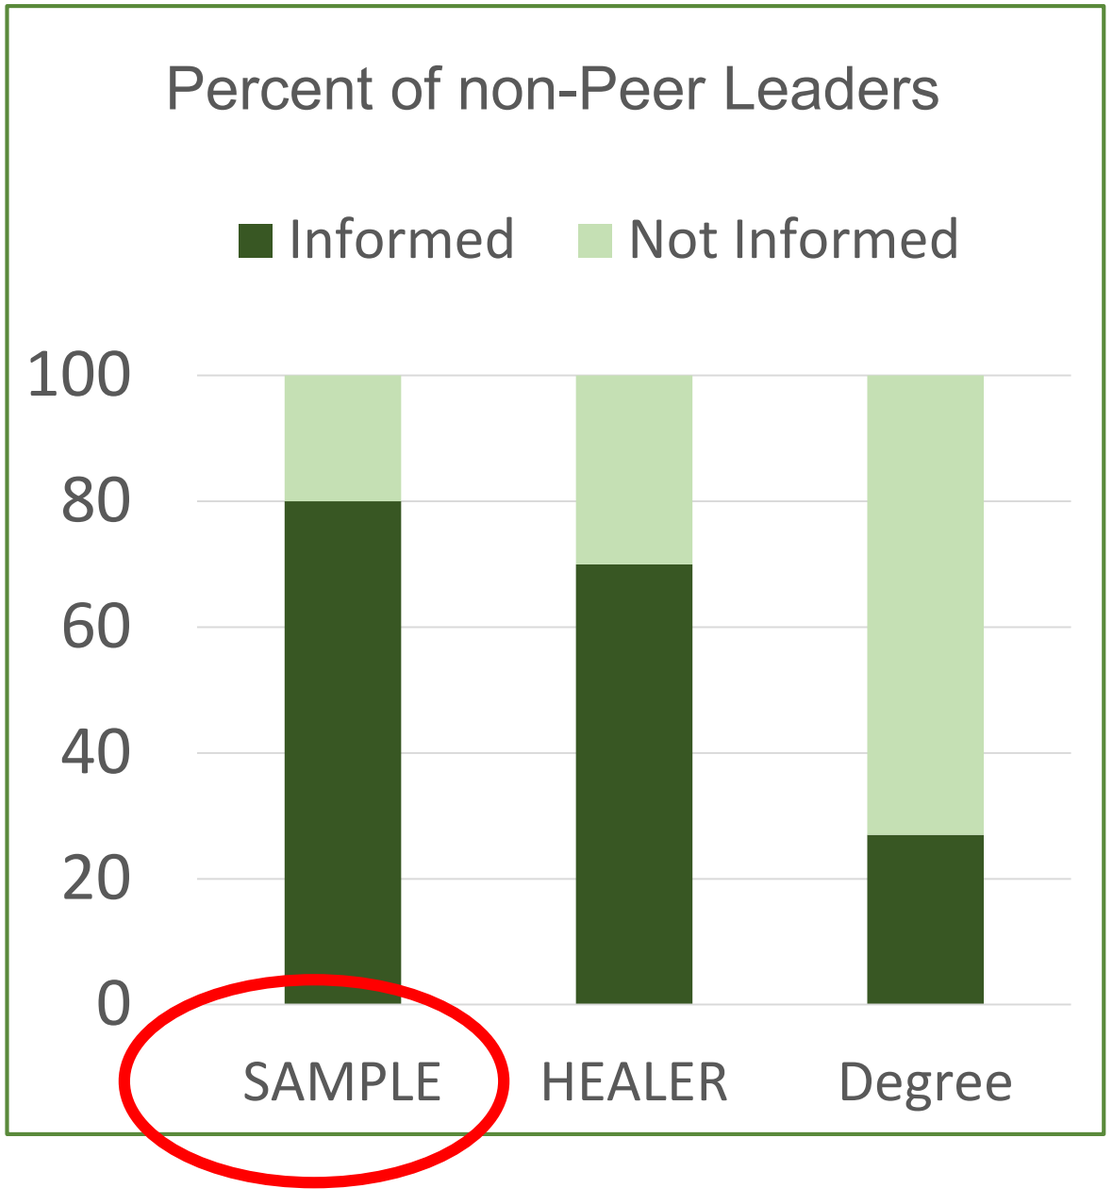

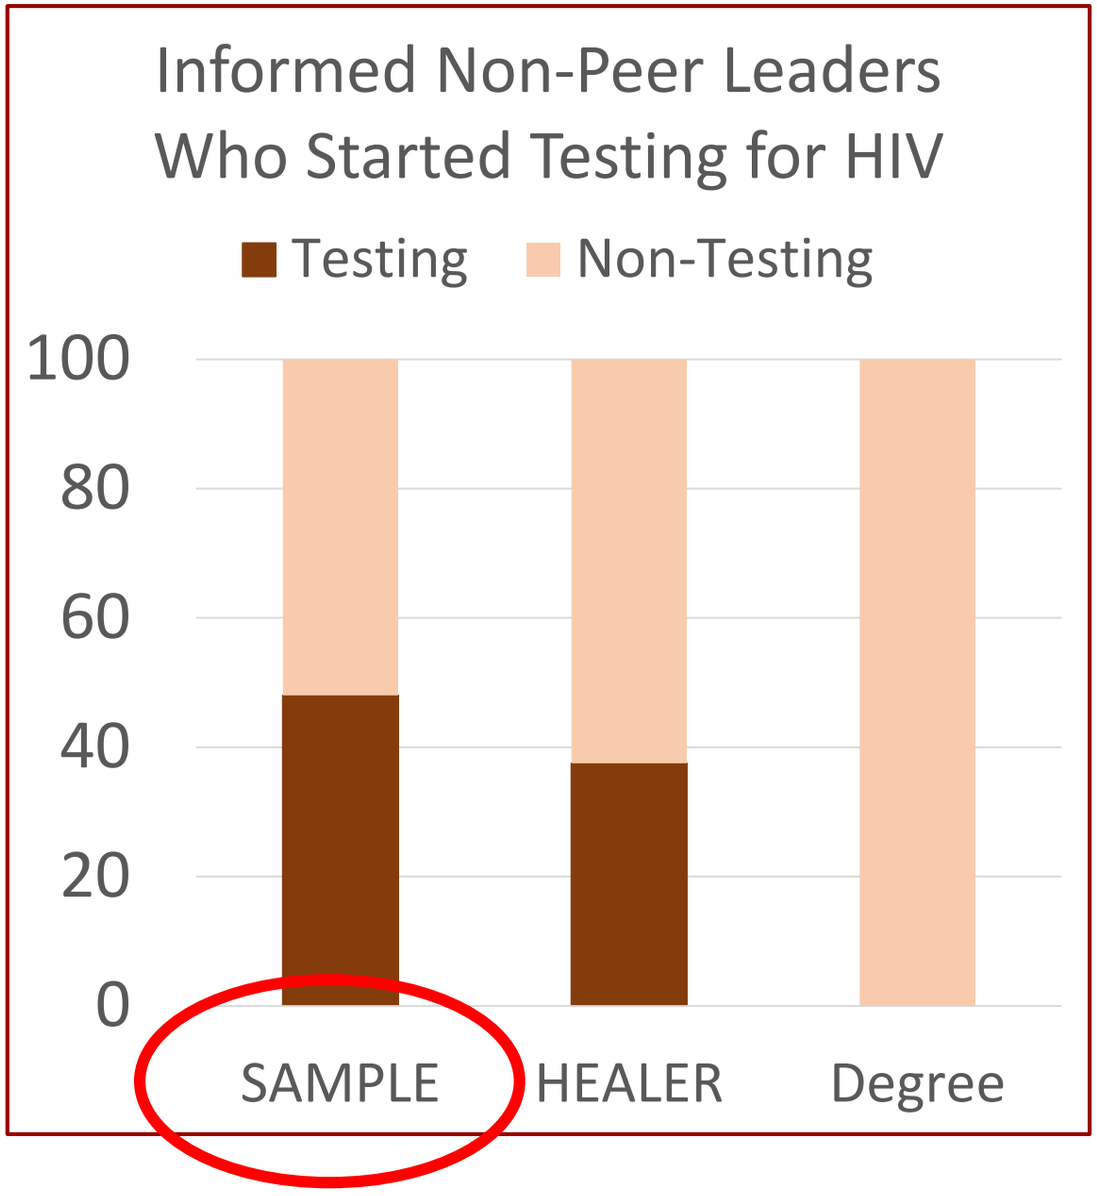

AI Assistant: HEALER

Continuing Research on HIV prevention [2019]

- Completing 900 youth study at three homeless shelters

Public Health: Optimizing Limited Social Worker Resources Preventing Tuberculosis in India [2019]

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

- *Patient in low resource communities: Non-adherence to TB Treatment*
- *Digital adherence tracking: Patients call phone #s on pill packs; many countries*
- *Predict adherence risk from phone call patterns? Intervene before patients miss dose*

TB Treatment Adherence but Limited Resources: Intervening Selectively before patients miss doses

Killian

Data Collect

Phone logs

Predict high risk patients

RF or LSTM

Prescription

Constraint Top K

Field

- *15K patients, 1.5M calls*

everwell

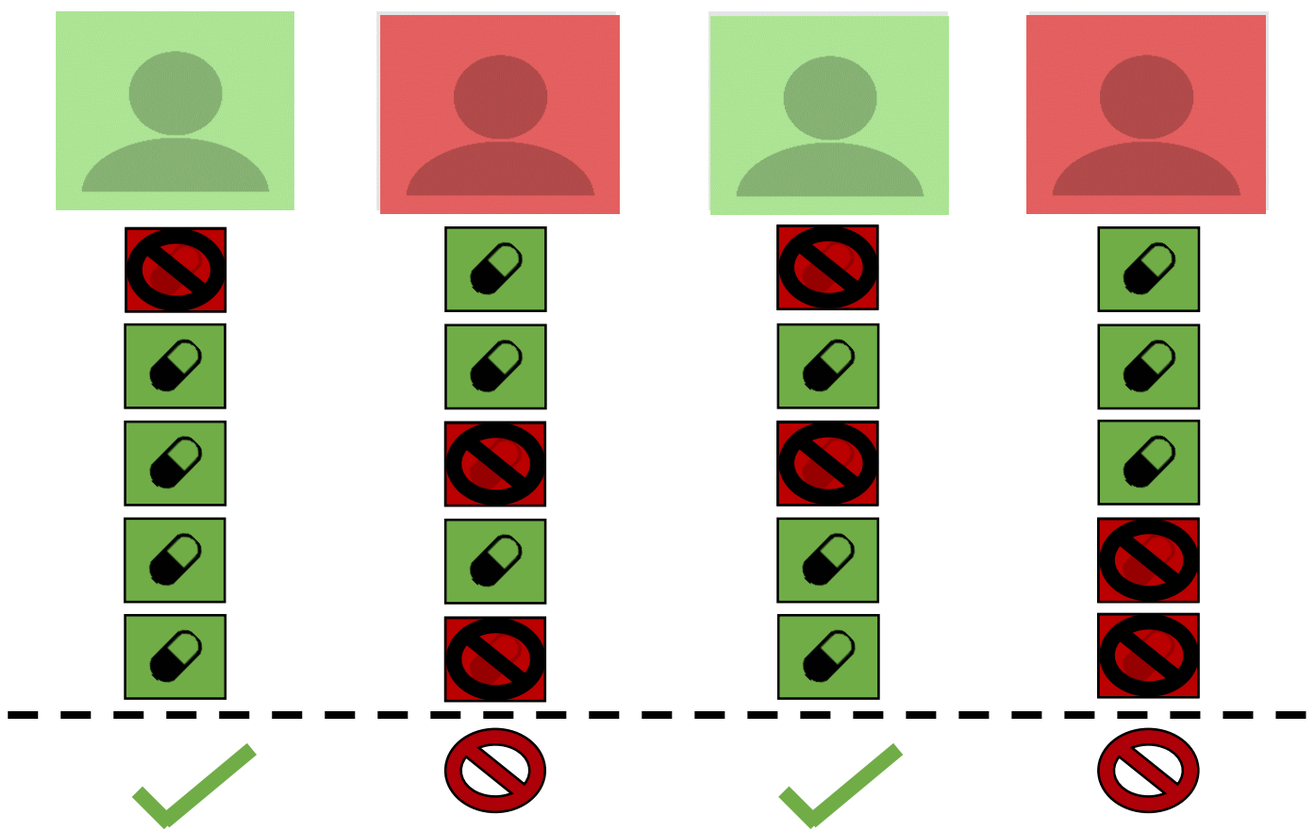

Increasing TB Treatment Adherence: Intervening before patients miss doses [2019]

Killian

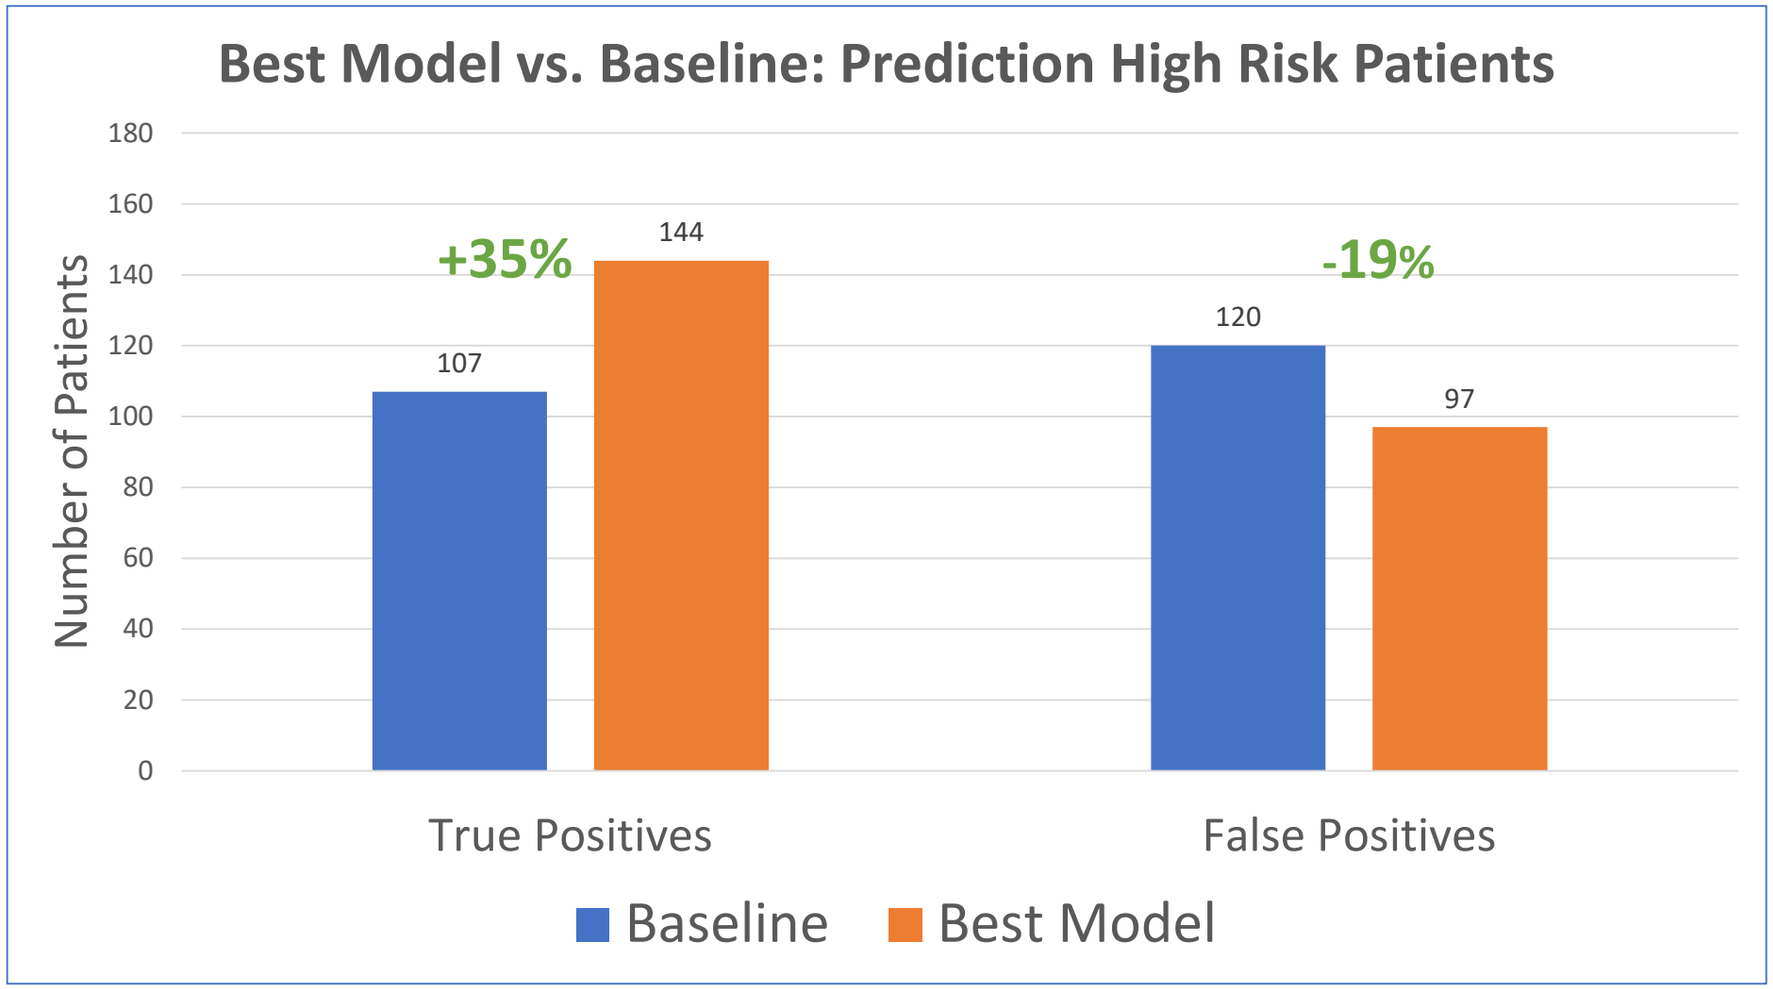

Data from

State of Maharashtra

India

Improving TB interventions
Decision-Focused vs Stage by Stage Methods


Wilder

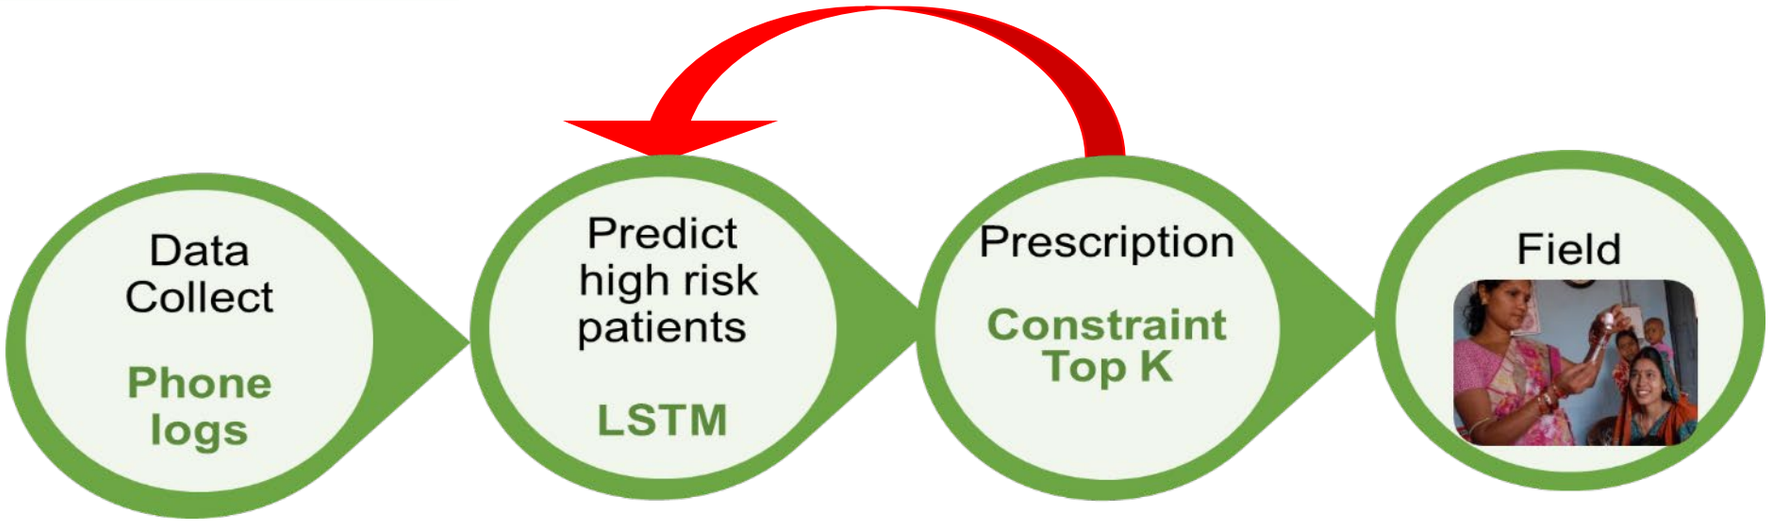

Decision focused learning improves TB interventions

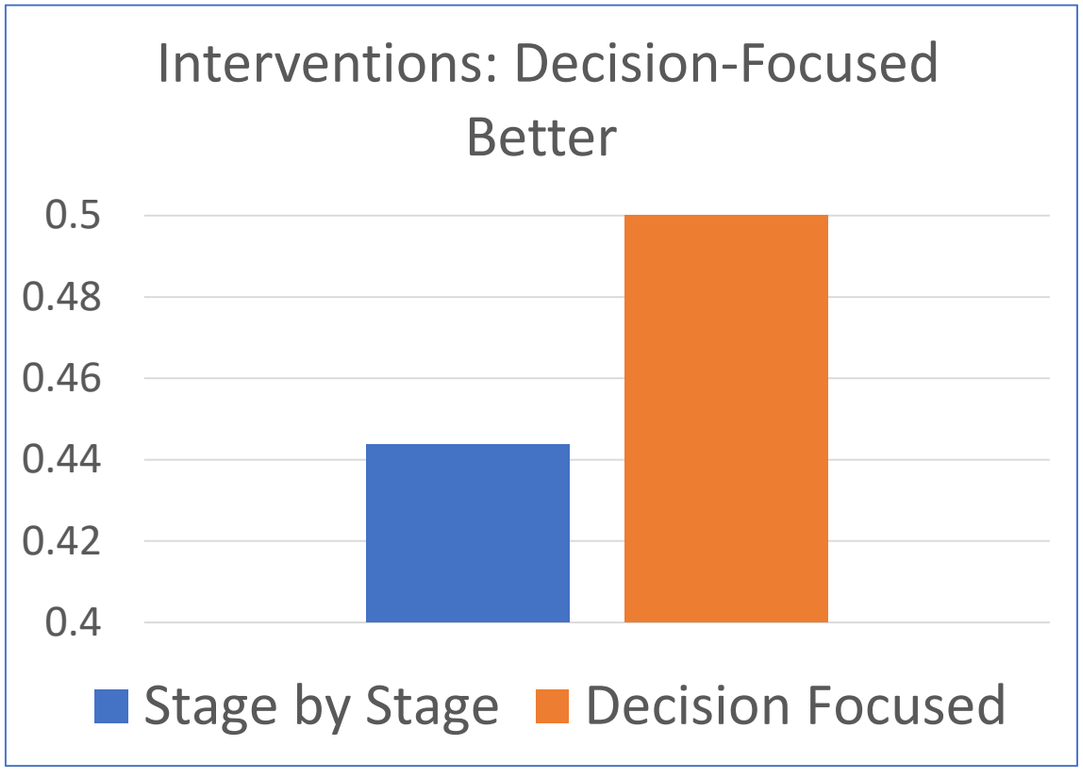

Integrating with Everwell's Platform

Killian

everwell

This work has a lot of potential to save lives.

Bill Thies

Co-founder, Everwell Health Solutions

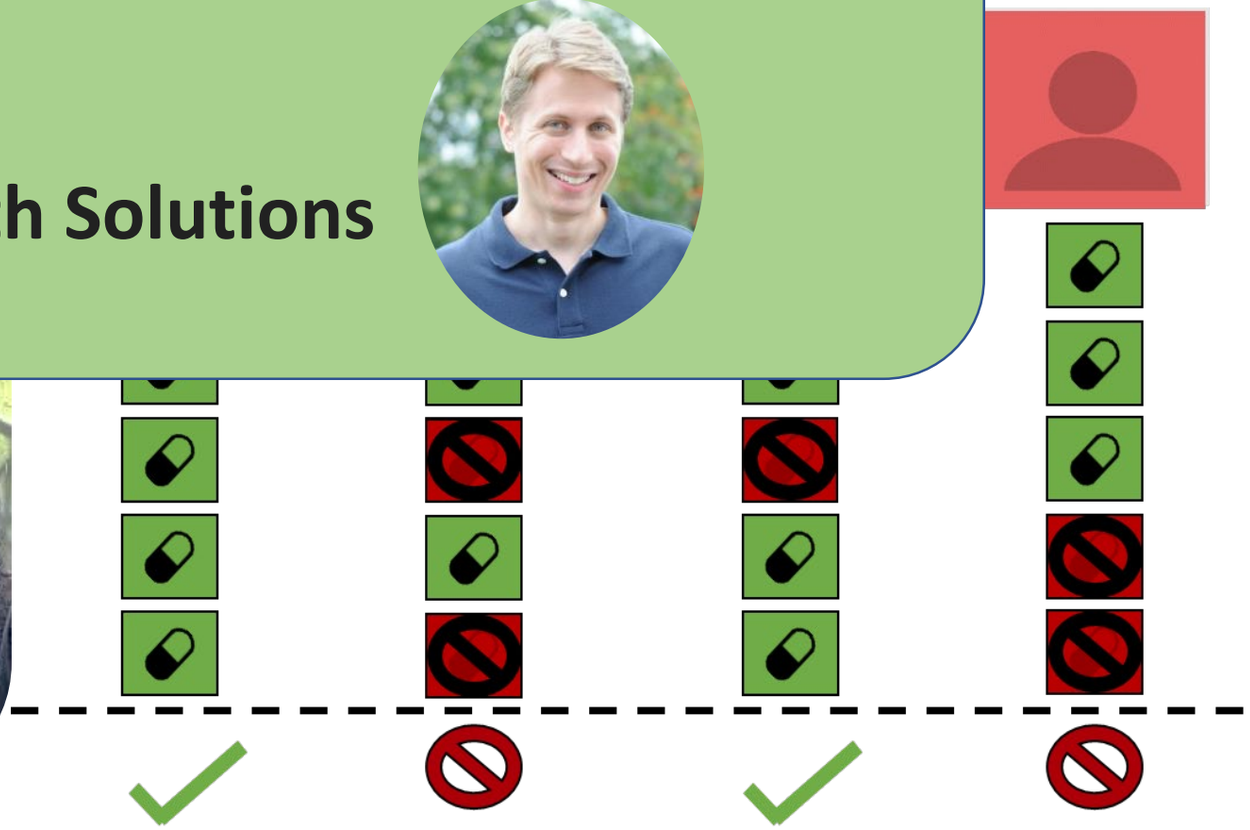

Suicide Prevention in Marginalized Populations: Choose Gatekeepers in social networks

Rahmattalabi

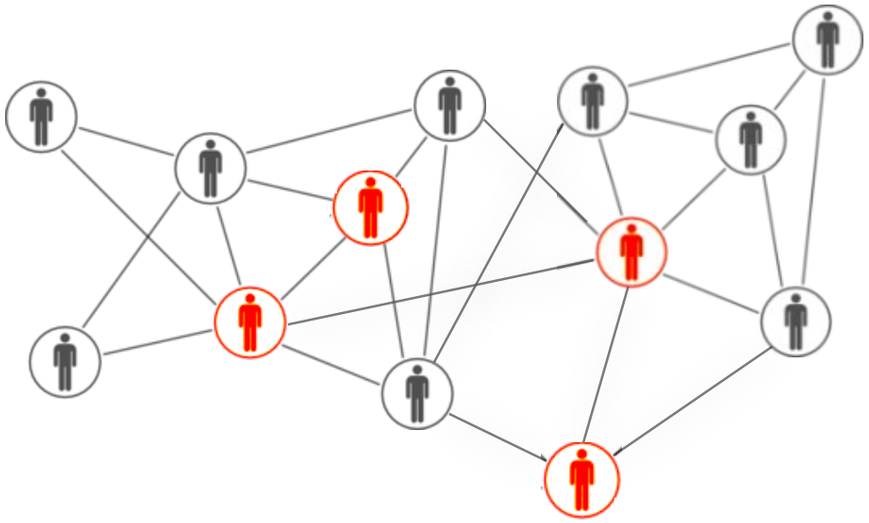

- Worst case parameters: a zero-sum game against nature
- Fairness of coverage

Algorithm
Chooses K gatekeepers

vs

Nature
Chooses some gatekeepers to not participate

Childhood Obesity Prevention via Network Optimization


Wilder


Ou

- *Childhood obesity: Diabetes, stroke and heart disease*
- *Early intervention with mothers: Change diet/activity using social networks*
- *Competitive influences in networks: Add/remove edges for behavior change*

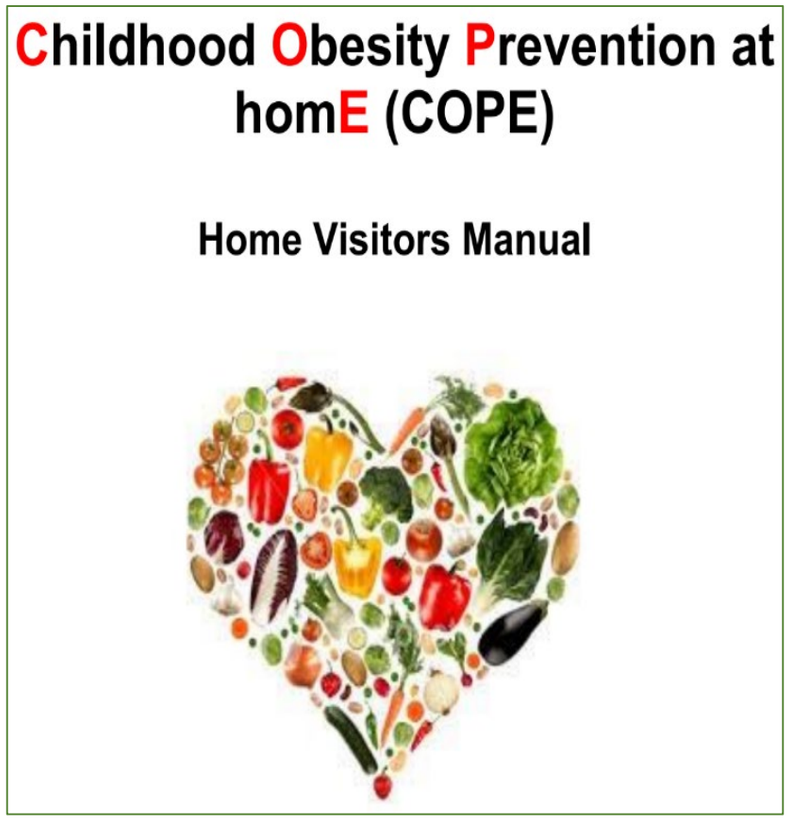

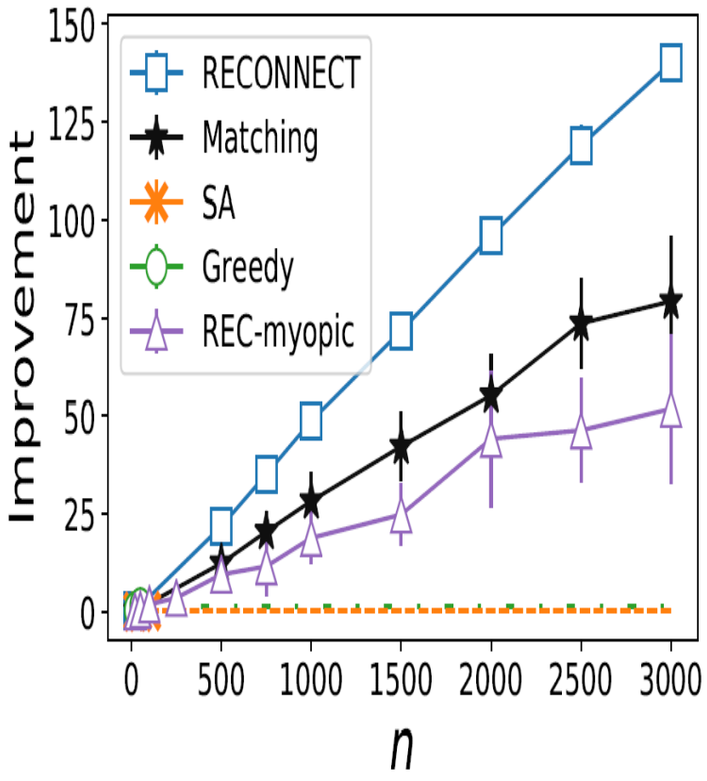

Summary
AI & Multiagent Systems for Social Impact

Cross-cutting challenge: How to optimize limited intervention resources

- Public safety & security, conservation, public health

Unifying themes

- *Multiagent systems reasoning*
- *Data to deployment*

Research contributions:

- *Models, algorithms*: Stackelberg Security Games, game-focused learning
- *Beyond models and algorithms…*

Future: Multiagent Systems and AI Research for Social Good

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

Thank you!

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

@MilindTambe_AI

Game-Focused Learning: End-to-End Method

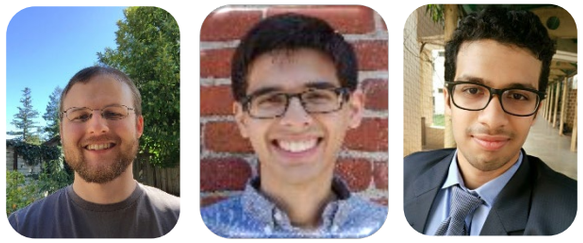

Perrault Wilder Mate

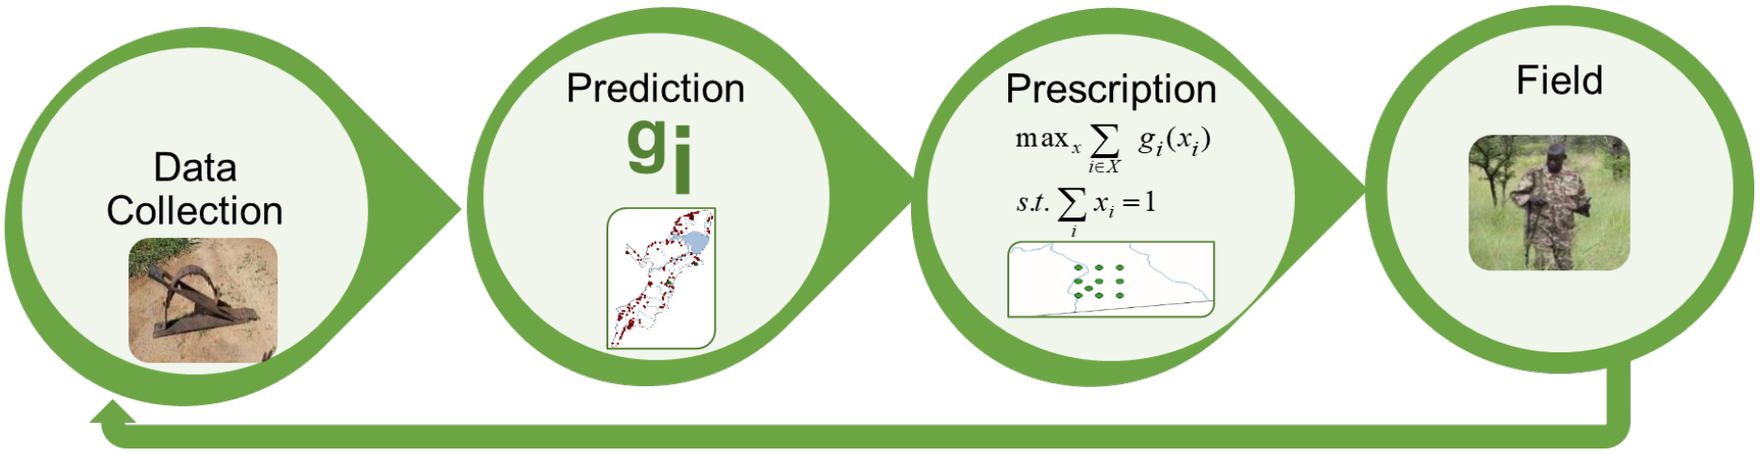

- Game-focused learning approach
  - → Goal of learning: *defender's utility* as high as possible
  - → Not accurate predictions; but predictions for best decision

Game-Focused Learning:
Reduces Errors on Important Targets

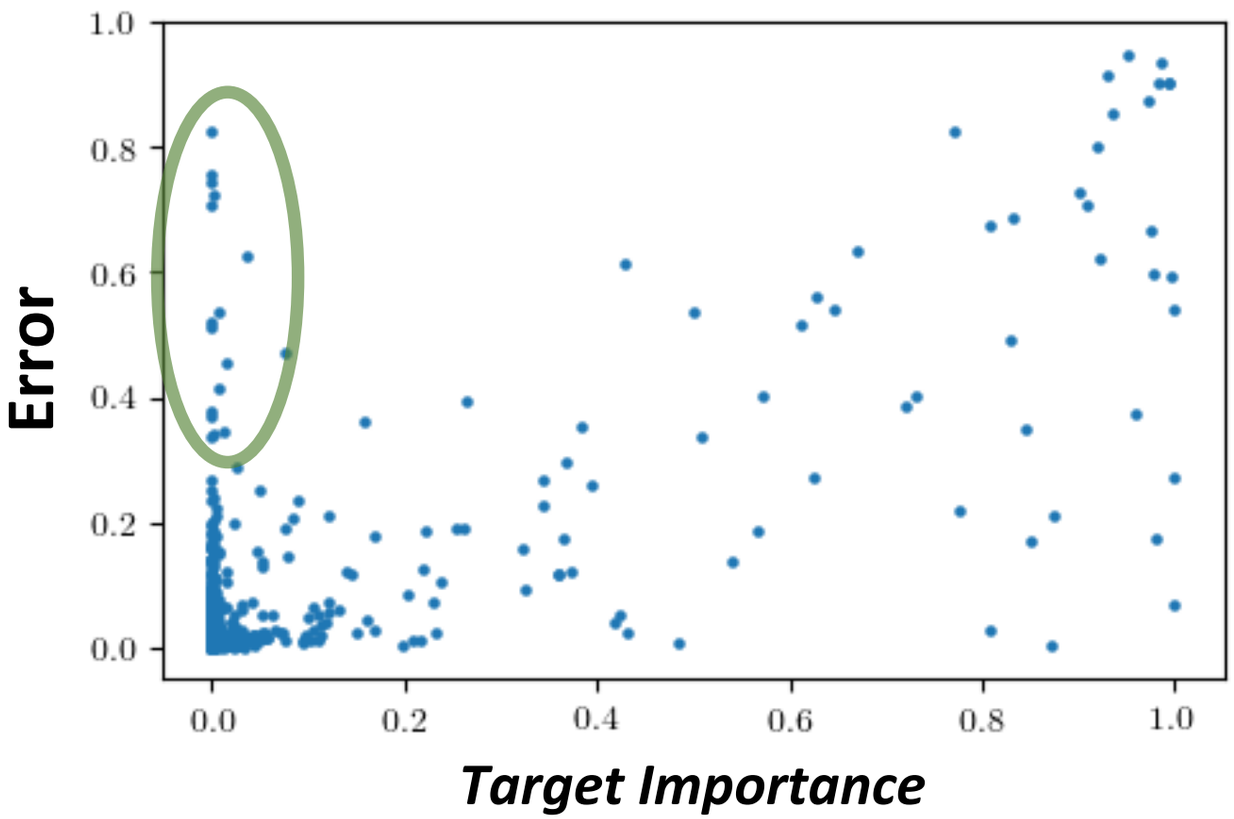

Game-focused

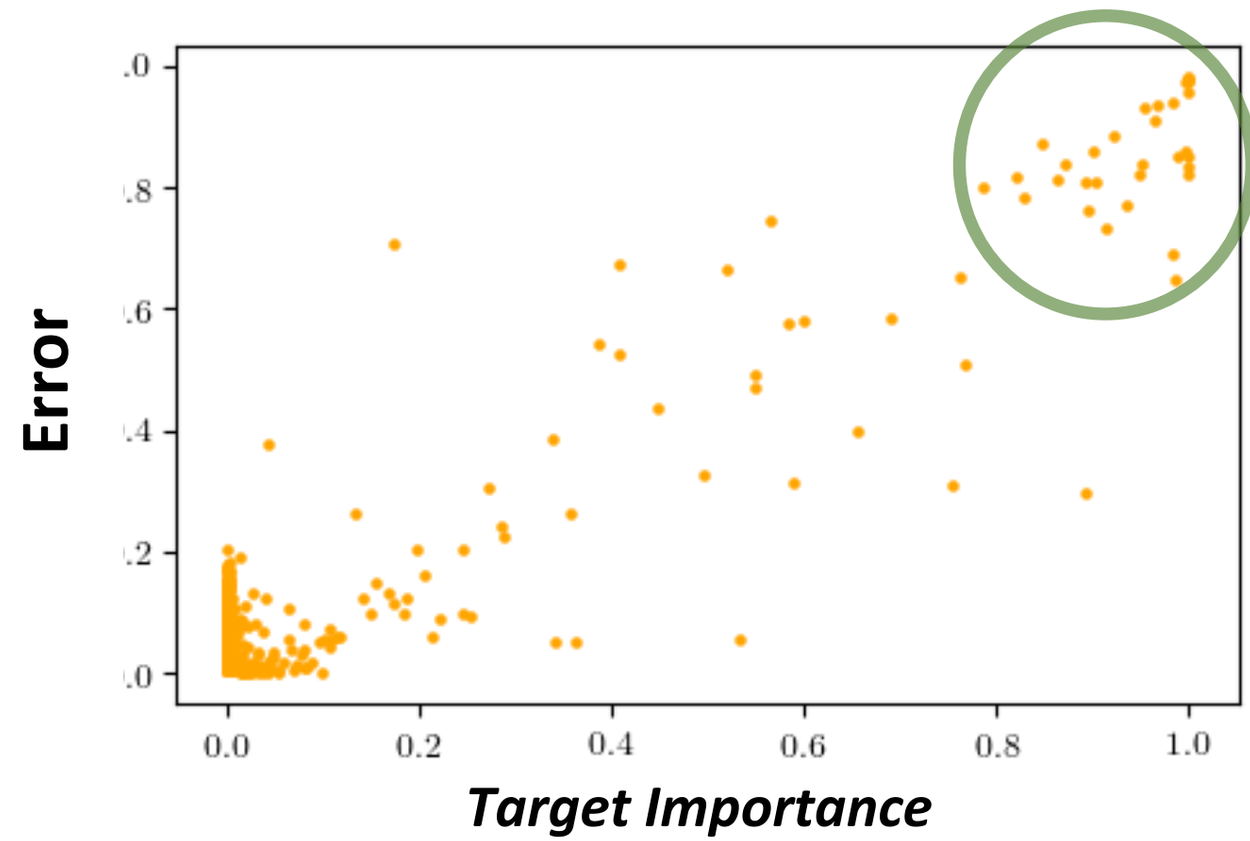

Two-stage