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

Optimizing Limited Intervention Resources


Stackelberg security games

Public Safety & Security


Green security games

Conservation

Sampled social networks


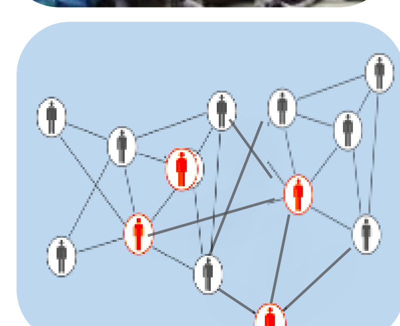

Public Health

Google Research Bangalore
Director, AI for Social Good

AI for Social Good workshop

armman
Delivering India's Future

KhushiBaby

Swasti
THE HEALTH CATALYST

Public Health

Education

atree
ASHOKA TRUST FOR RESEARCH IN ECOLOGY & THE ENVIRONMENT

Wildlife Conservation Trust

Conservation

Three Common Themes

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

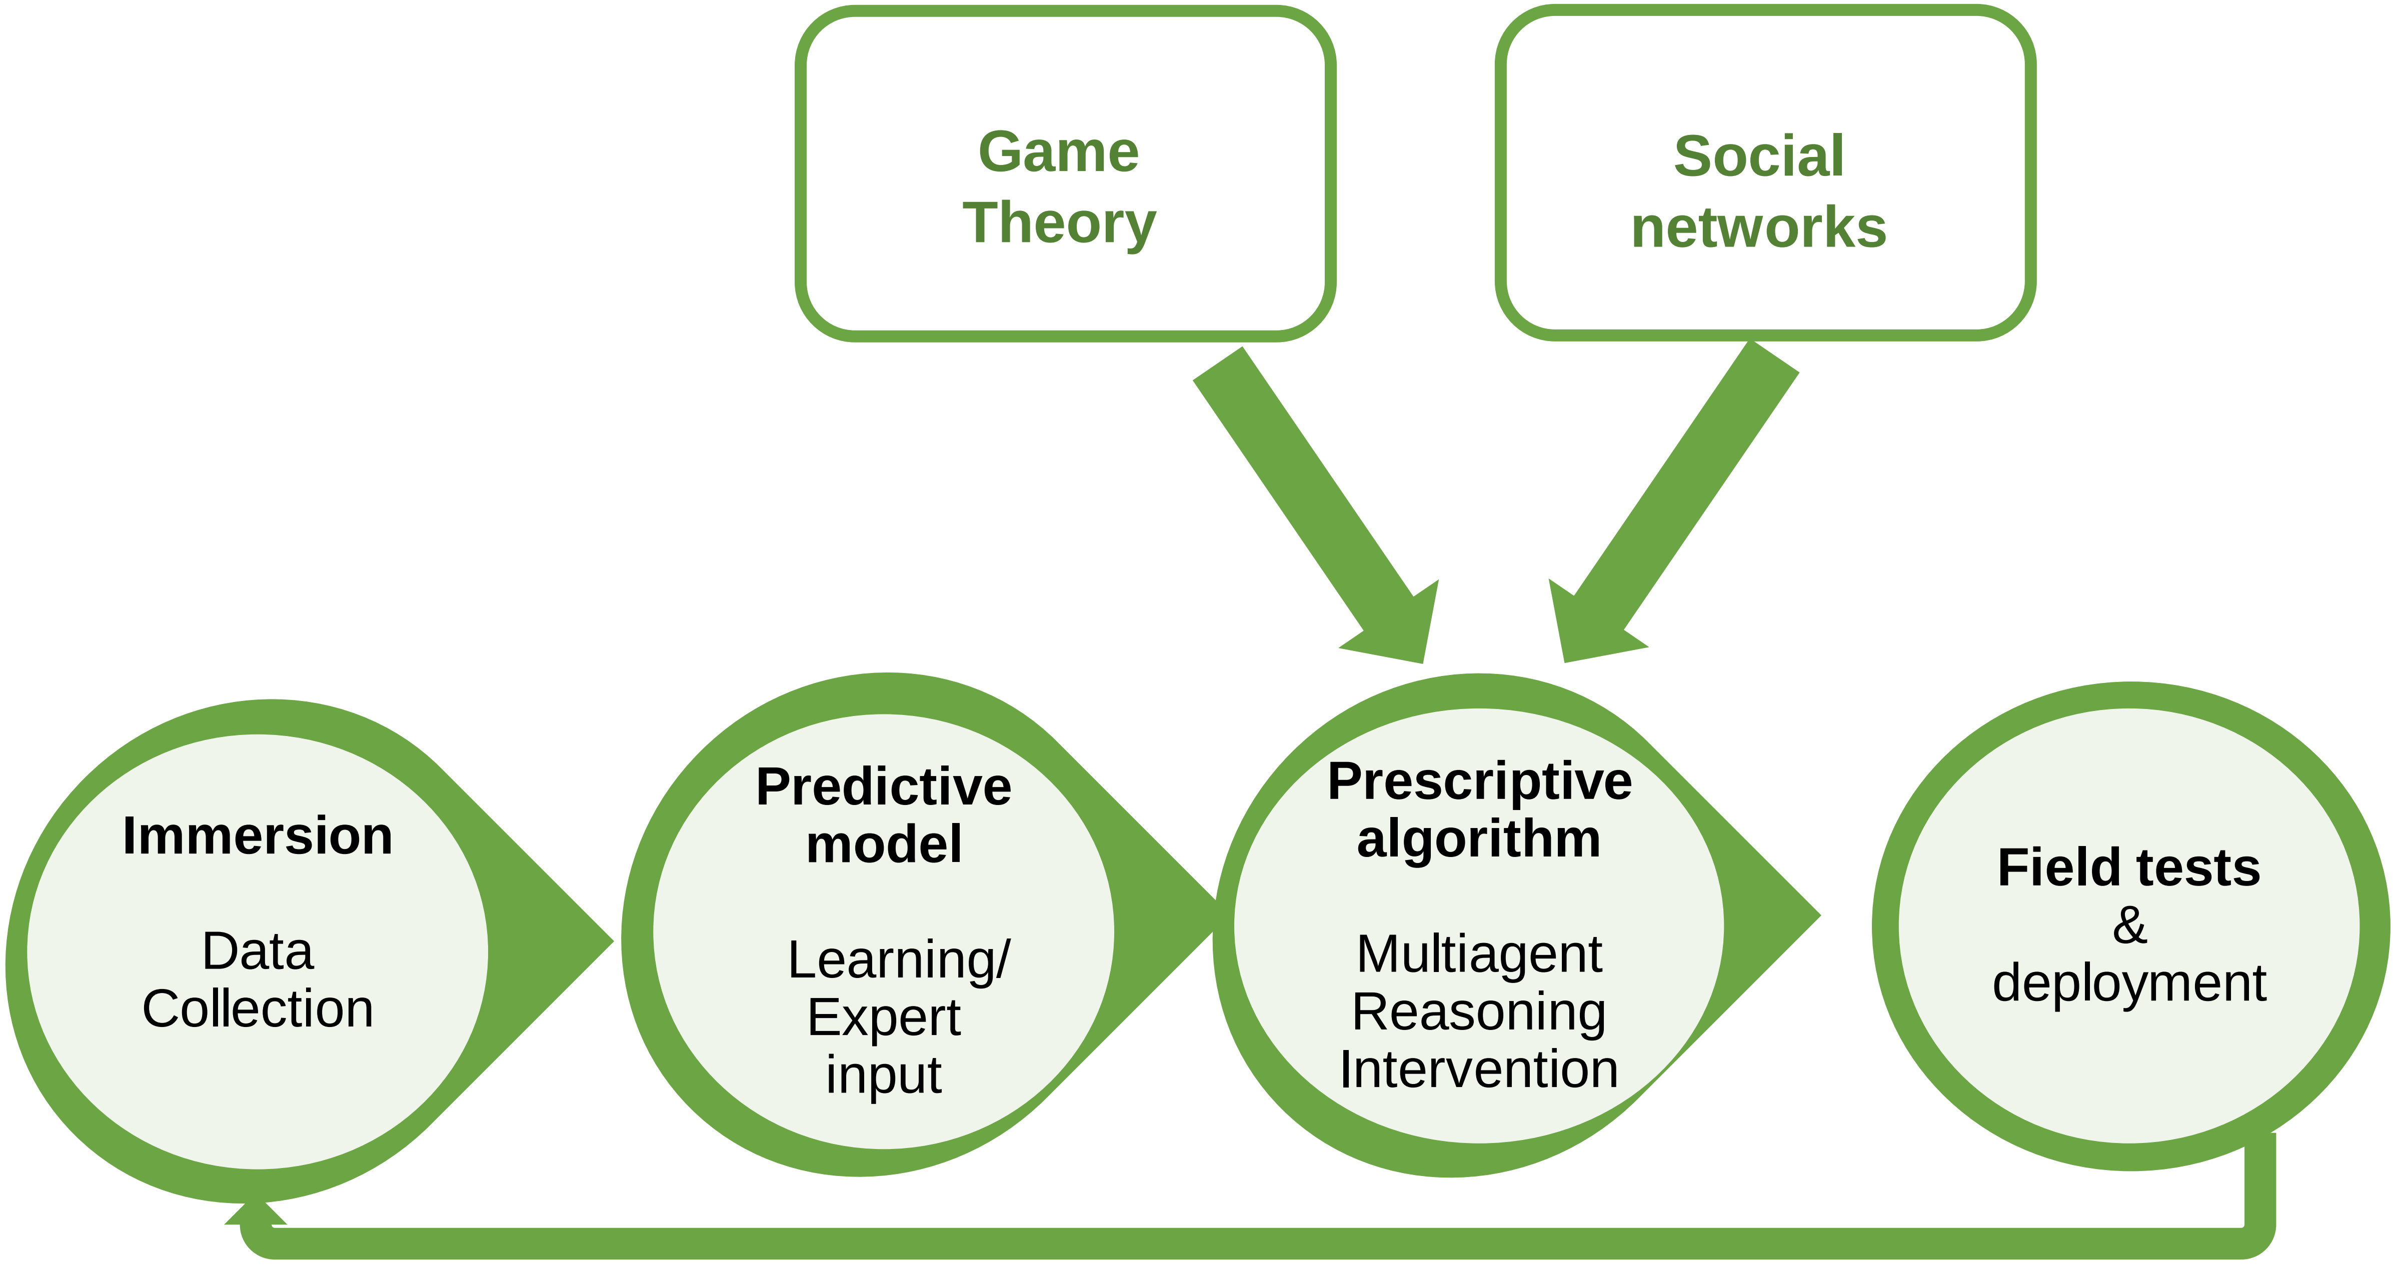

Three Common Themes

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

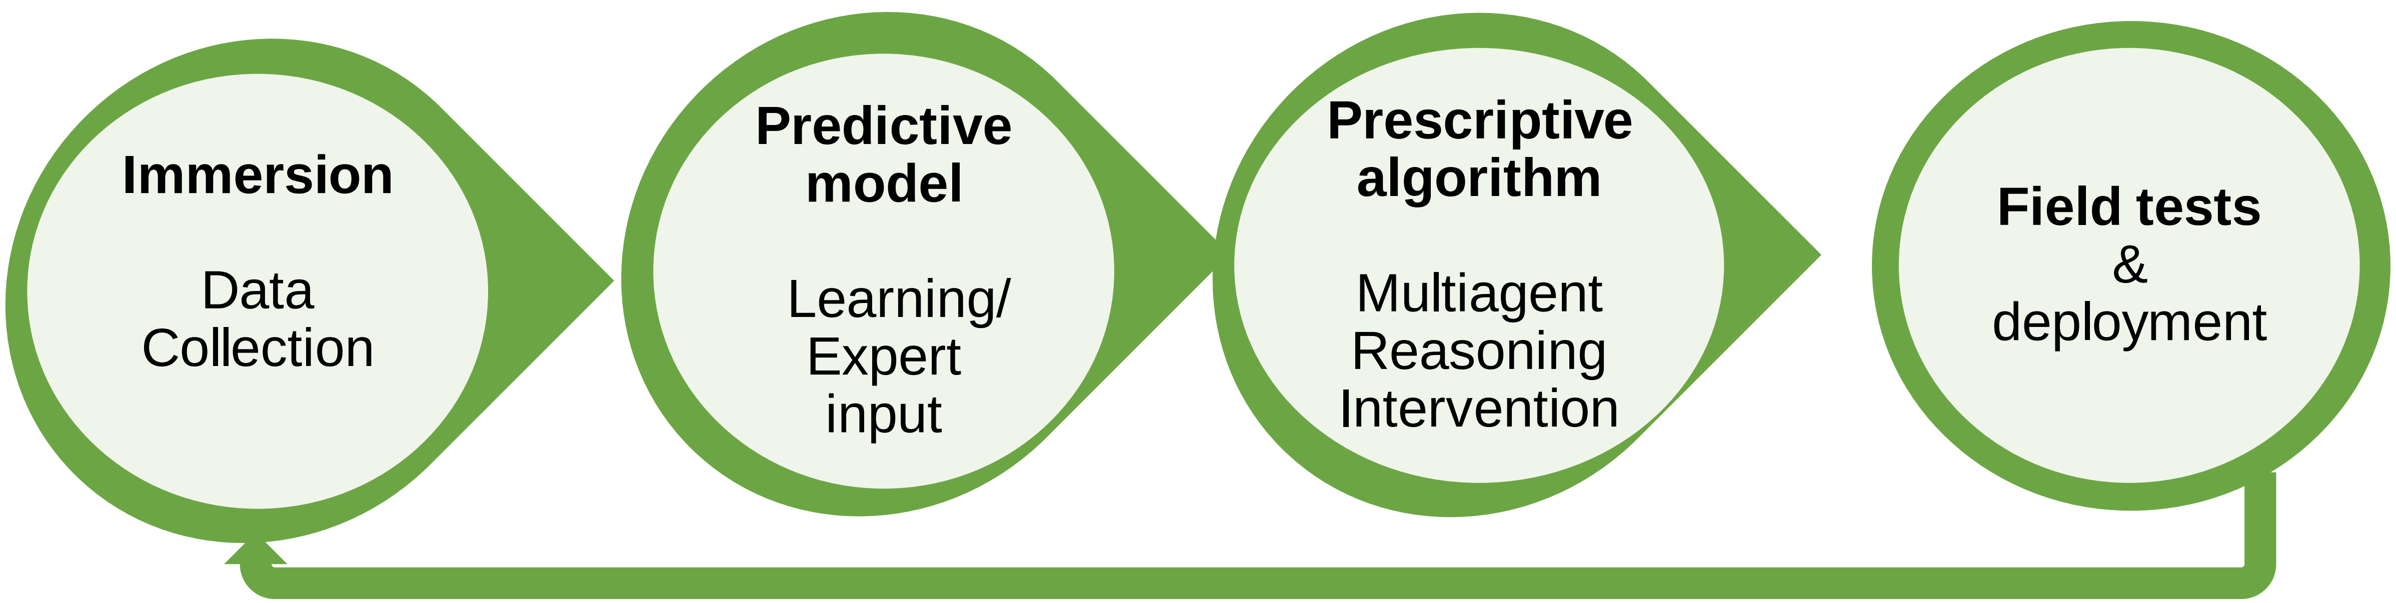

Three Common Themes

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

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

armman
Delivering India's Future

Swasti
THE HEALTH CATALYST

Wildlife Conservation Trust

KhushiBaby

atree
ASHOKA TRUST FOR RESEARCH IN ECOLOGY & THE ENVIRONMENT

PRATHAM BOOKS

Outline: Overview of Past 14 Years of Research

Public Safety & Security: Stackelberg Security Games (brief)

Conservation/Wildlife Protection: Green Security Games

Public Health: Influence maximization & social networks

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

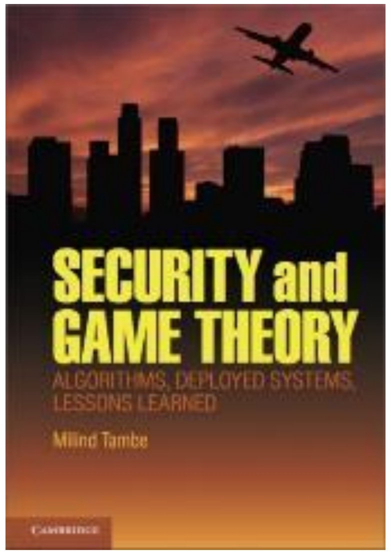

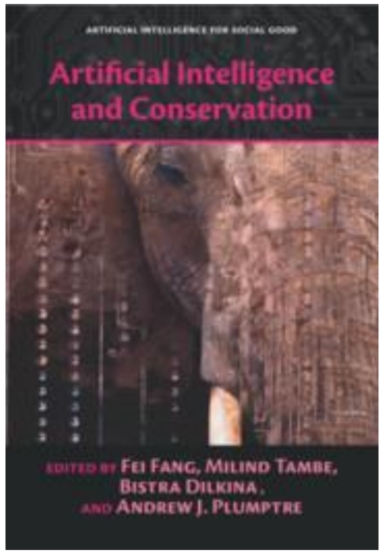

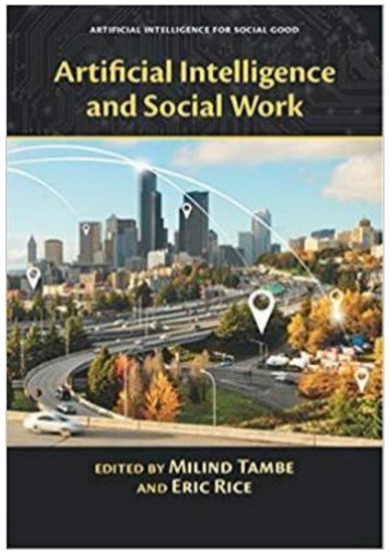

ARMOR Airport Security: LAX(2007)
Game Theory direct use for security resource optimization?

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	$\boxed{-5, 5}$	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 calculated losses

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

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*

Massive Scale Security Games

Large Number of Combinations: Guards to Targets

1000 targets, 20 guards: 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

Slave (LP Duality Theory)
Best new pure strategy

	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)
Next best new pure strategy

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

Deployed Security Games Systems…

ARMOR 2007

IRIS 2009

PROTECT 2011

Erroll Southers testimony
Congressional subcommittee

TSA testimony
Congressional subcommittee

US Coast Guard testimony
Congressional subcommittee

Reviewer 2 is not impressed!

Significant Real-World Evaluation Effort

Security Games superior in
Optimizing Limited Security Resources
Vs

Human Schedulers/“simple random”

Field Tests Against Adversaries

Computational Game Theory in the Field

Controlled

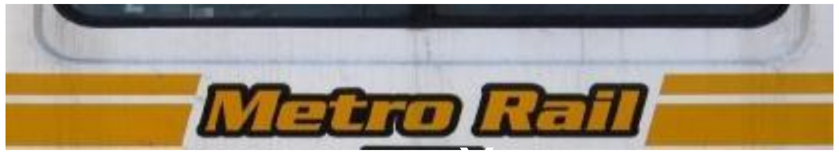

- 21 days of patrol, identical conditions
- Game theory vs Baseline+Expert

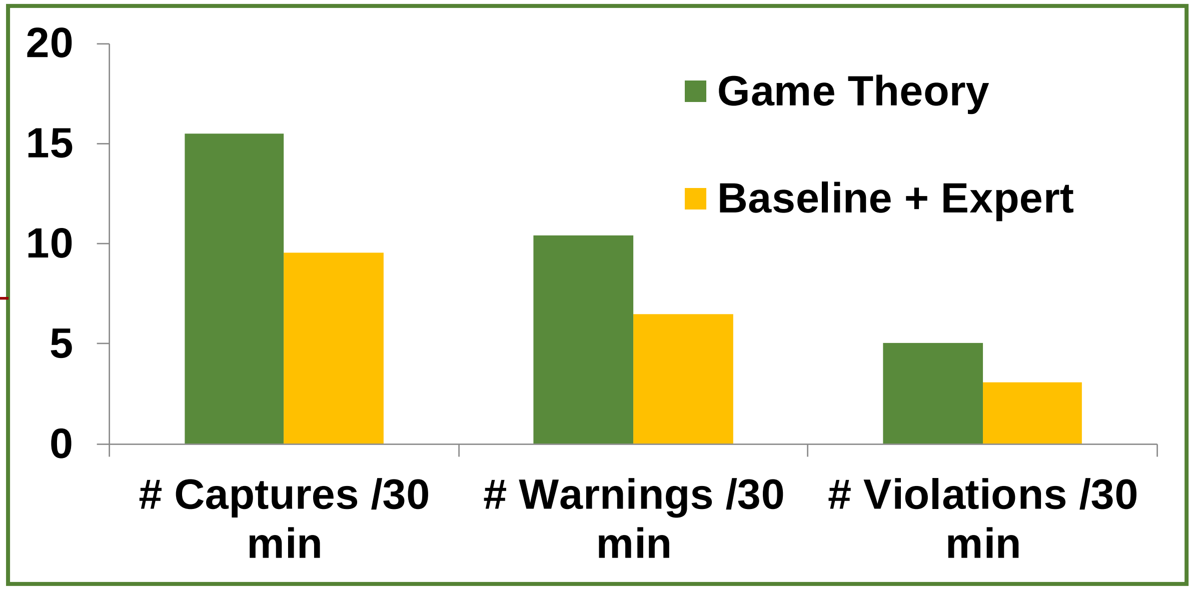

Not Controlled

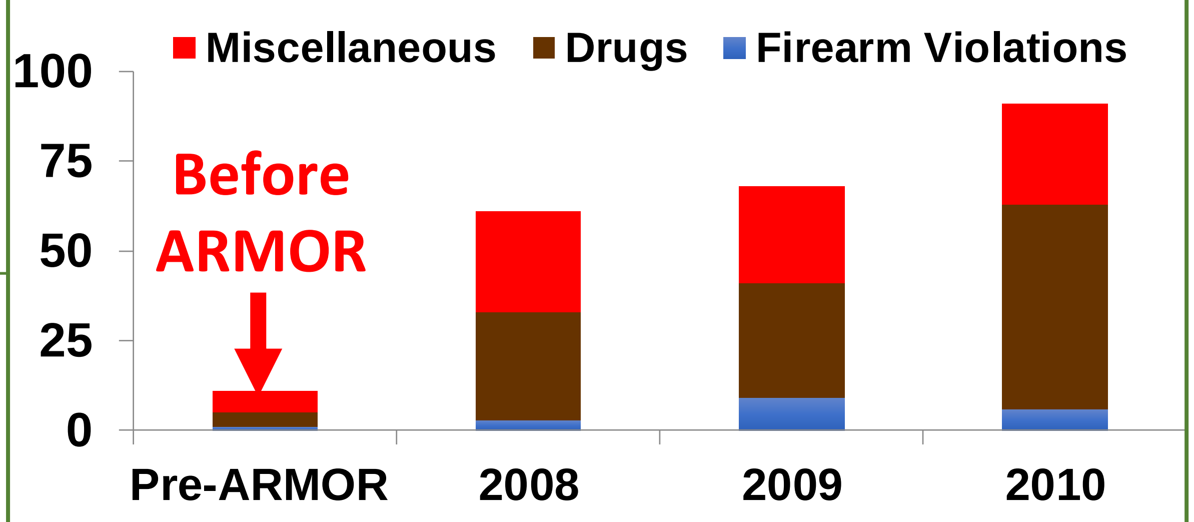

Outline

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

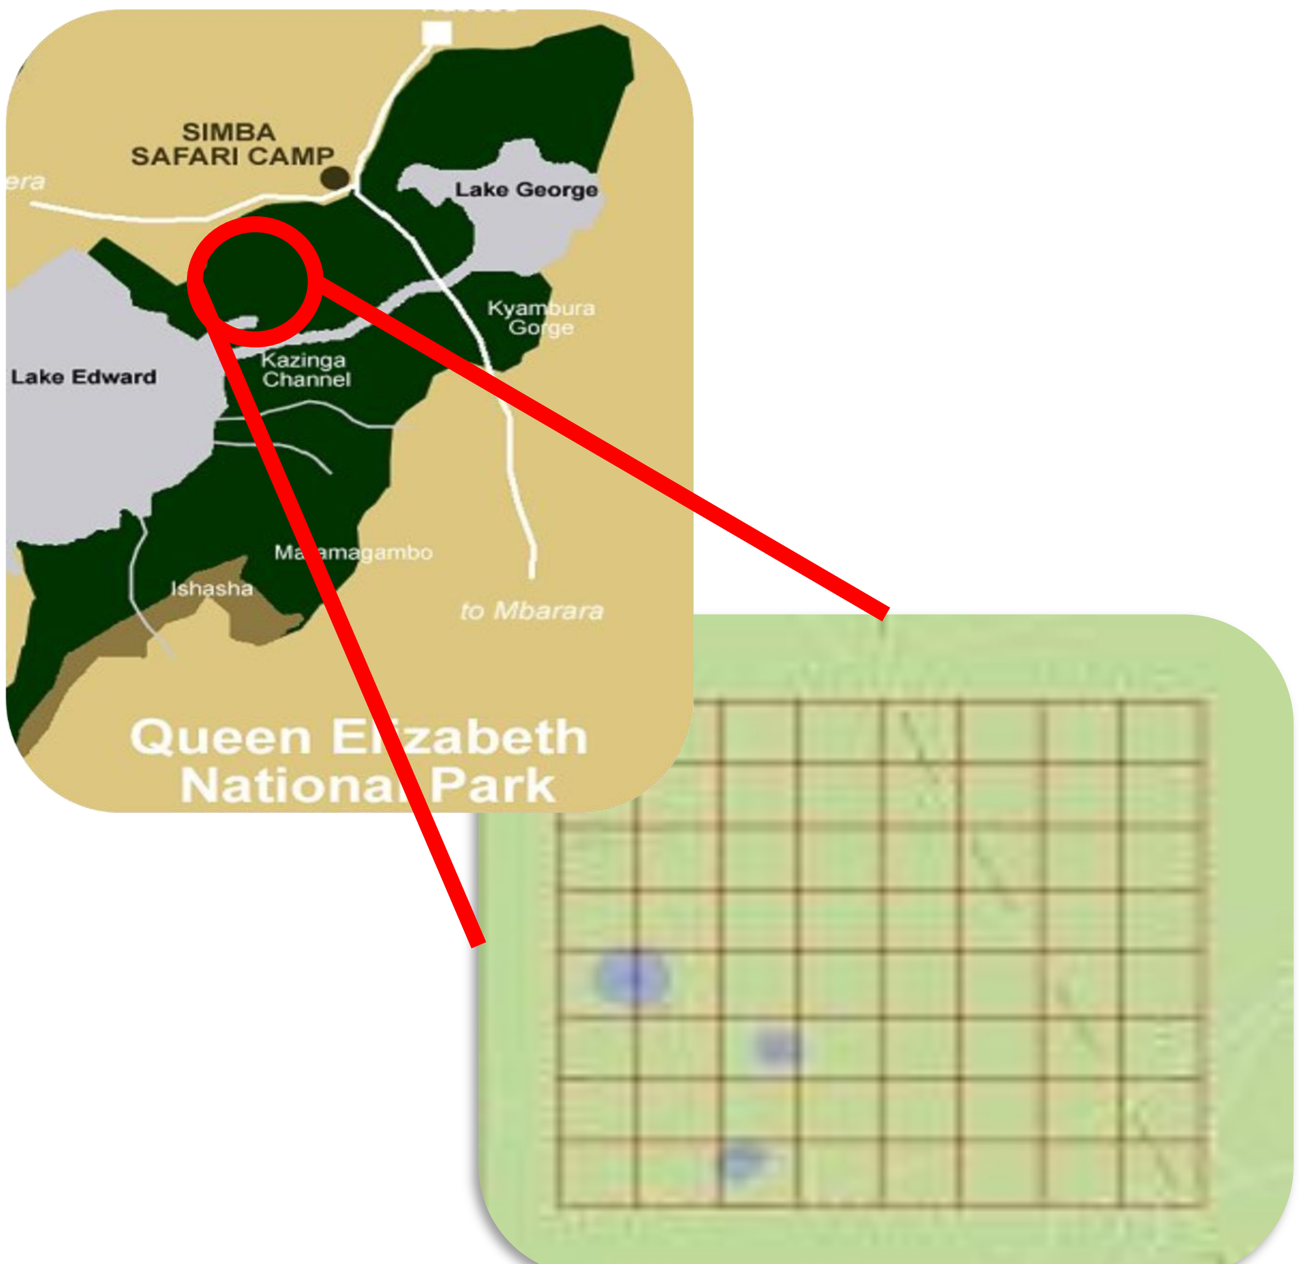

$$\max_{x,q} \sum_{i\in X} \sum_{j\in Q} R_{ij} x_i q_j$$

Max defender utility

$$s.t. \sum_{i} x_i = 1$$

Defender mixed strategy

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

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

Xu

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

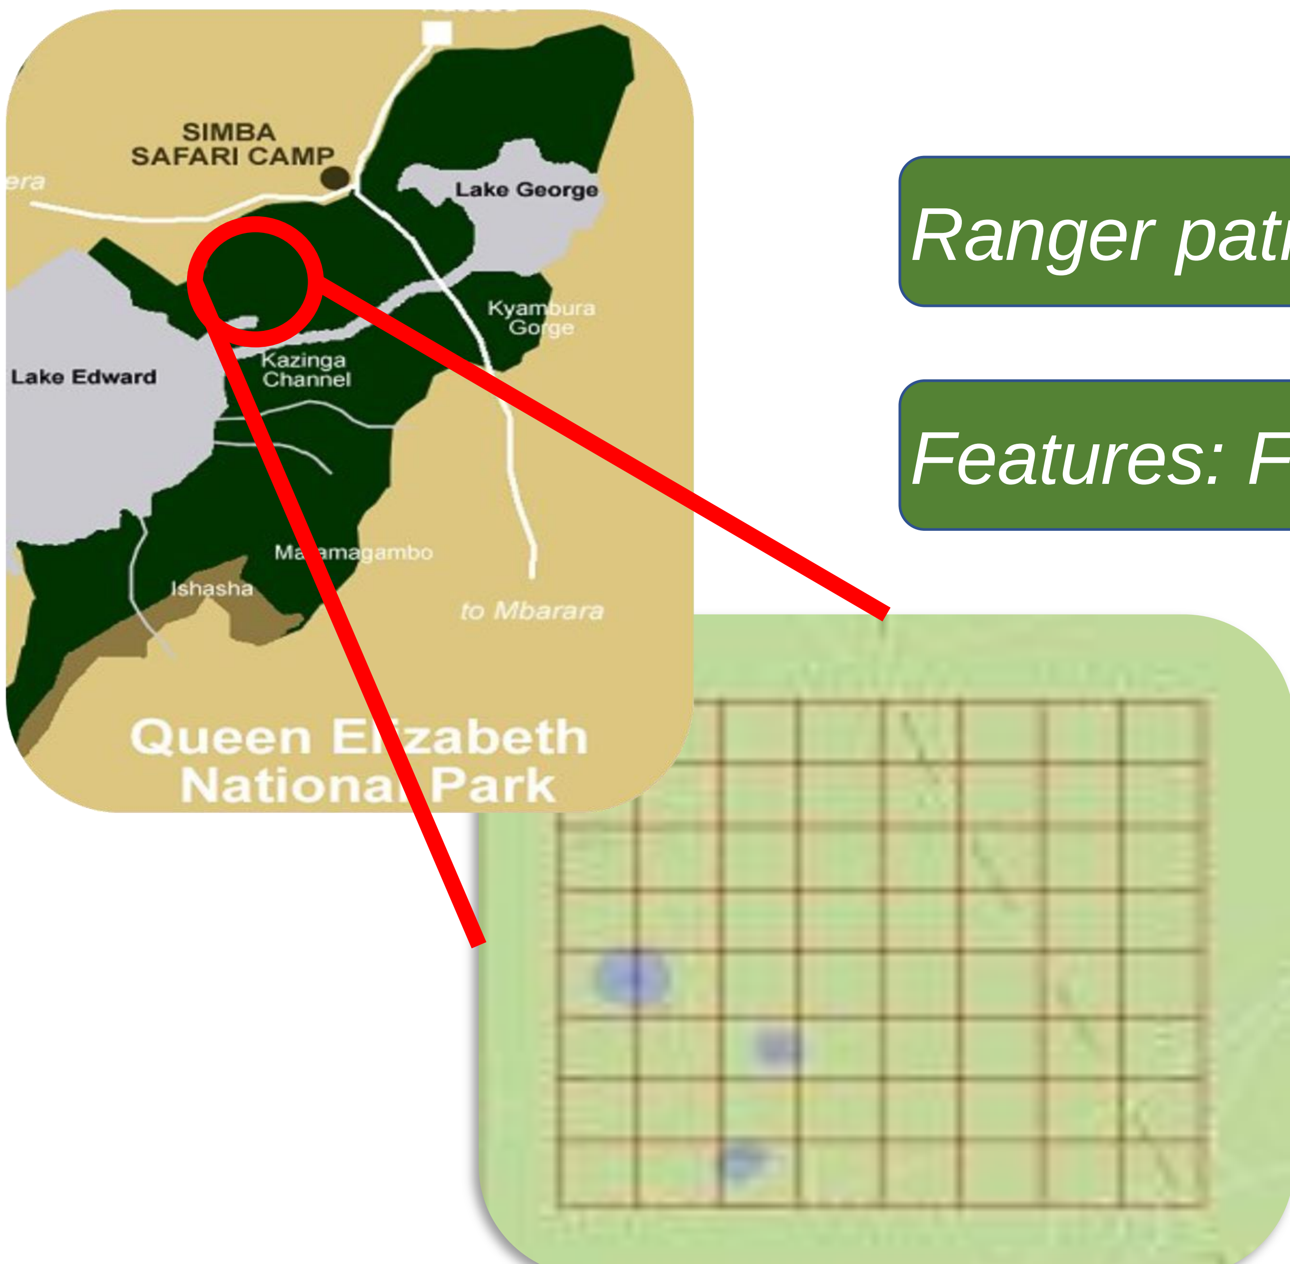

Ranger patrols: $X_{(i)}$

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

Date: 7/17/2020

Learning Adversary Model
12 Years of Past Poaching Data

Nguyen

g_j

Probability of snare Per 1 KM Grid Square

Ranger patrol

Animal density

Distance to rivers / roads / villages

Area habitat

Area slope

…

Learning Adversary Model
Uncertainty in Observations

Nguyen

g_i

Probability of snare Per 1 KM Grid Square

- Ranger patrol
- Animal density
- Distance to rivers / roads / villages
- Area habitat
- Area slope
- …

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

Gholami

Training: Filtered Datasets

Predict: Ensemble of Classifiers

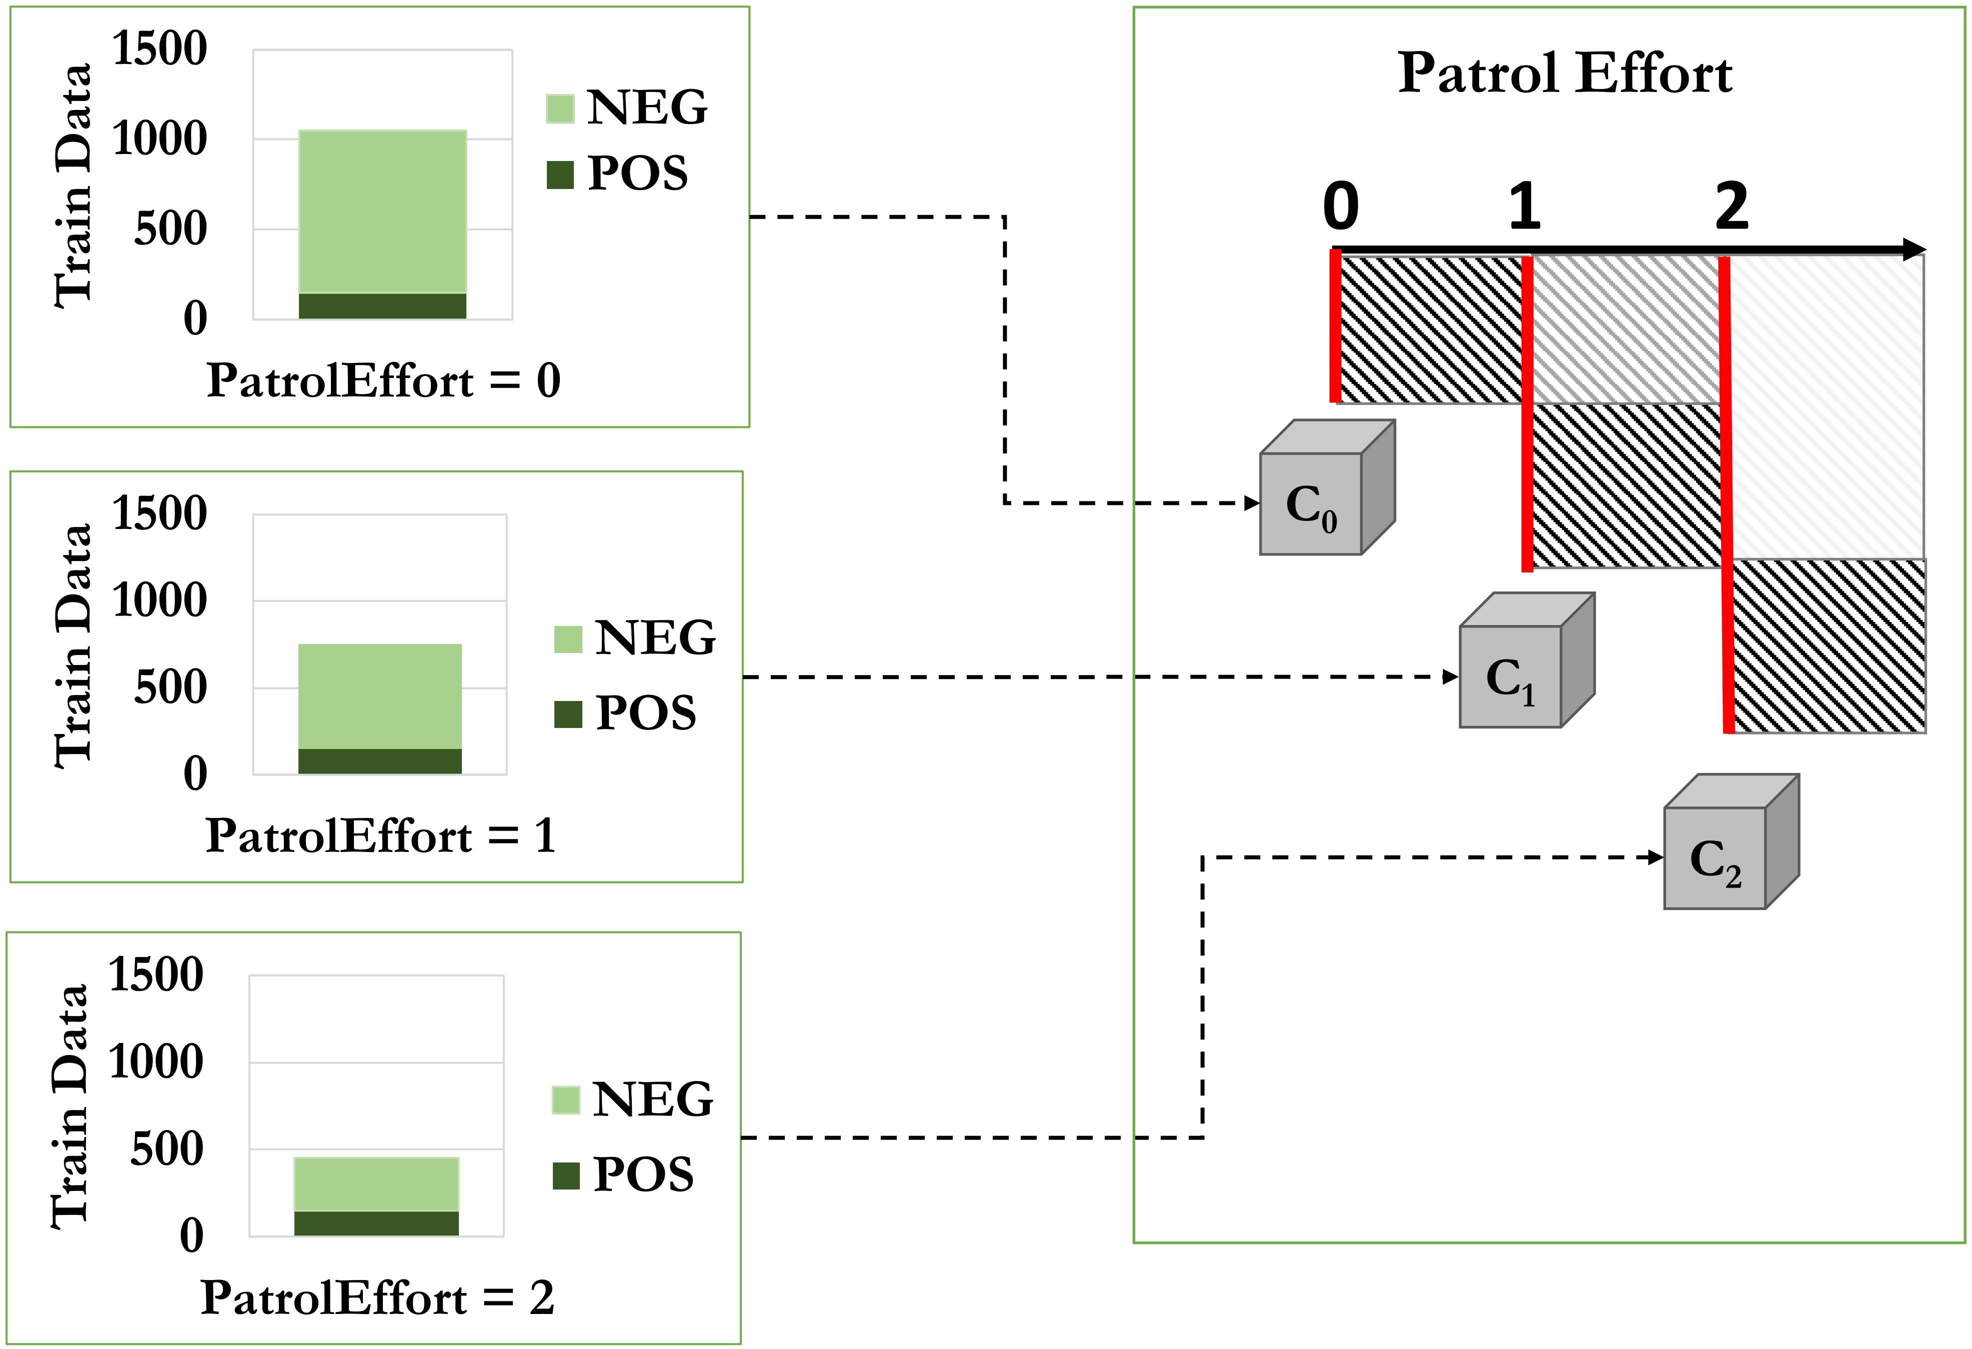

PAWS: Protection Assistant for Wildlife Security
Poacher Attack Prediction in the Lab

Gholami

Poacher Behavior Prediction

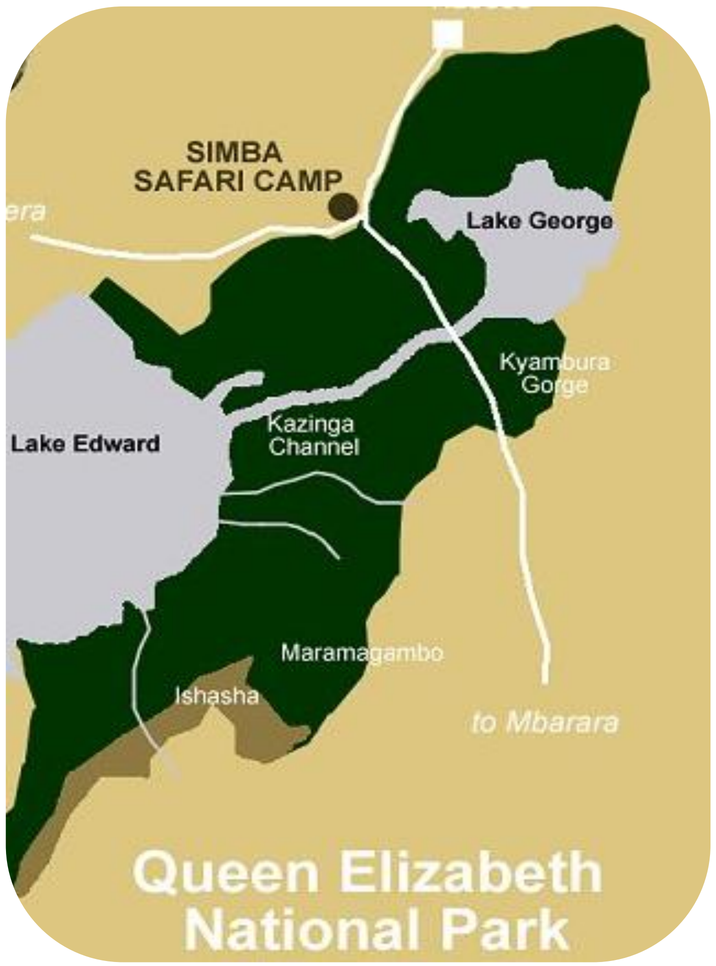

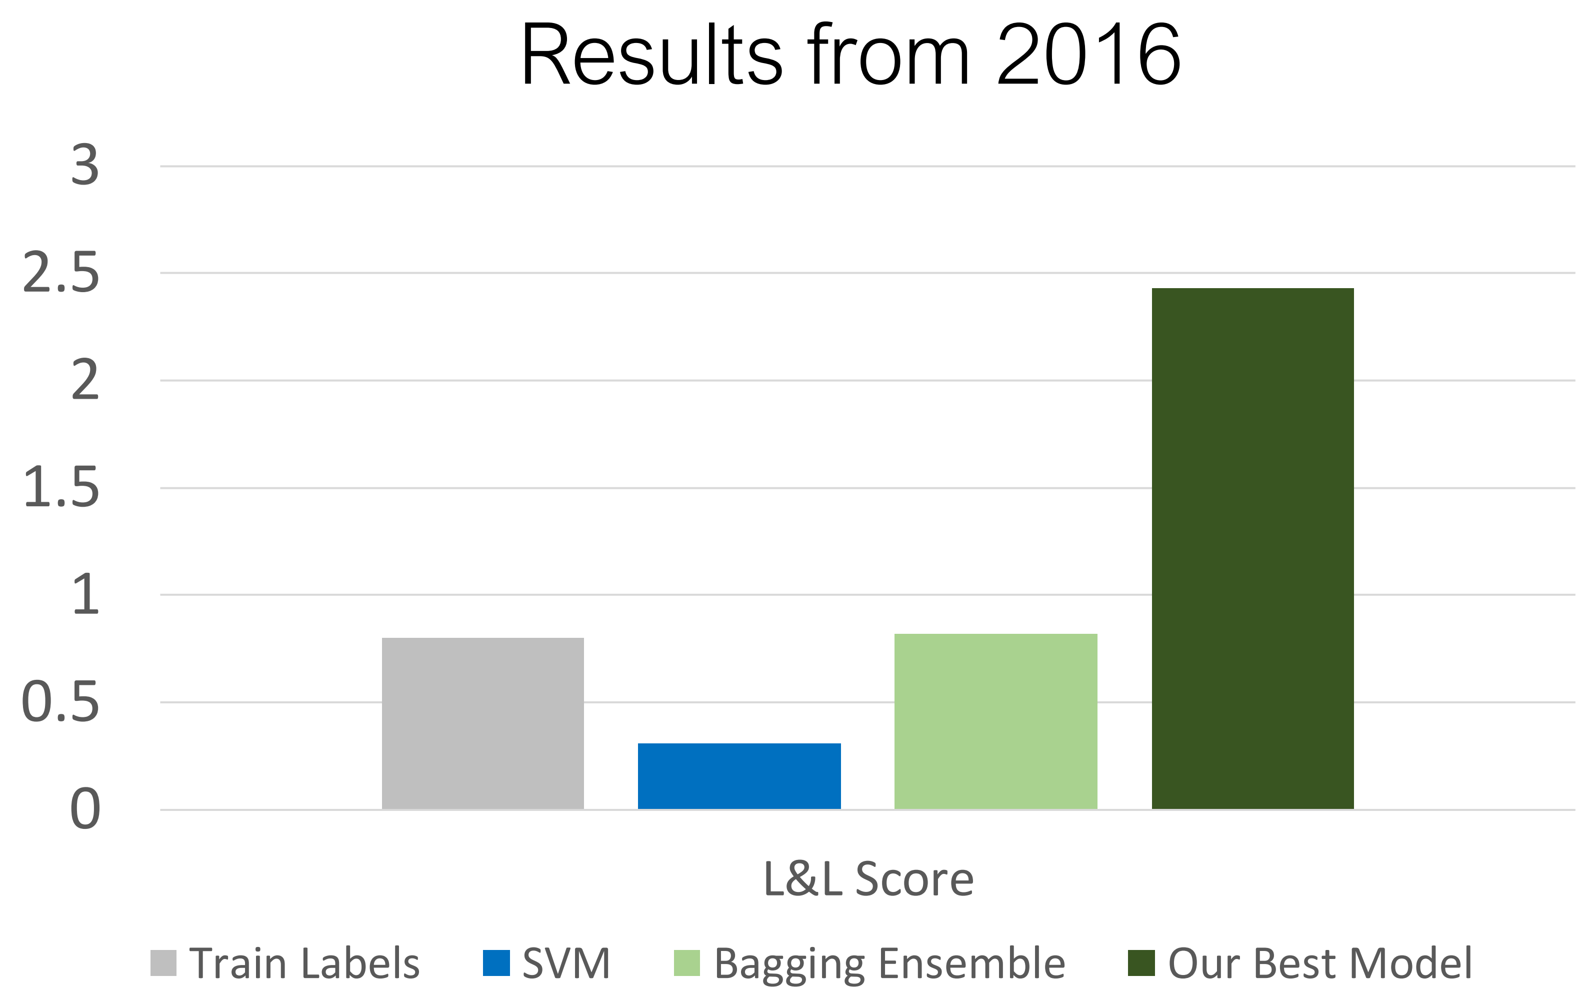

PAWS:
Real-world Deployment 2016: First Trial


Ford


Gholami

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

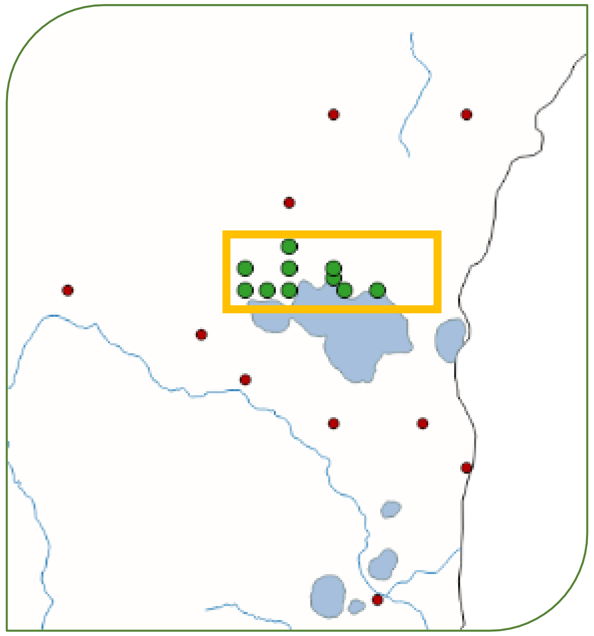

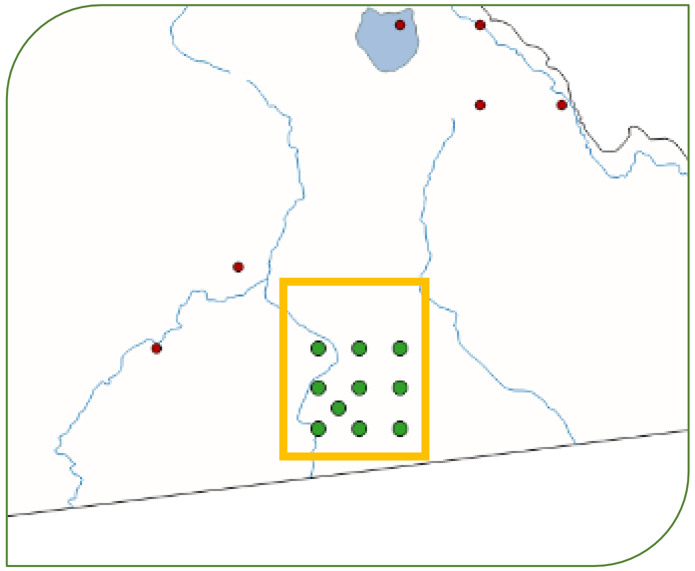

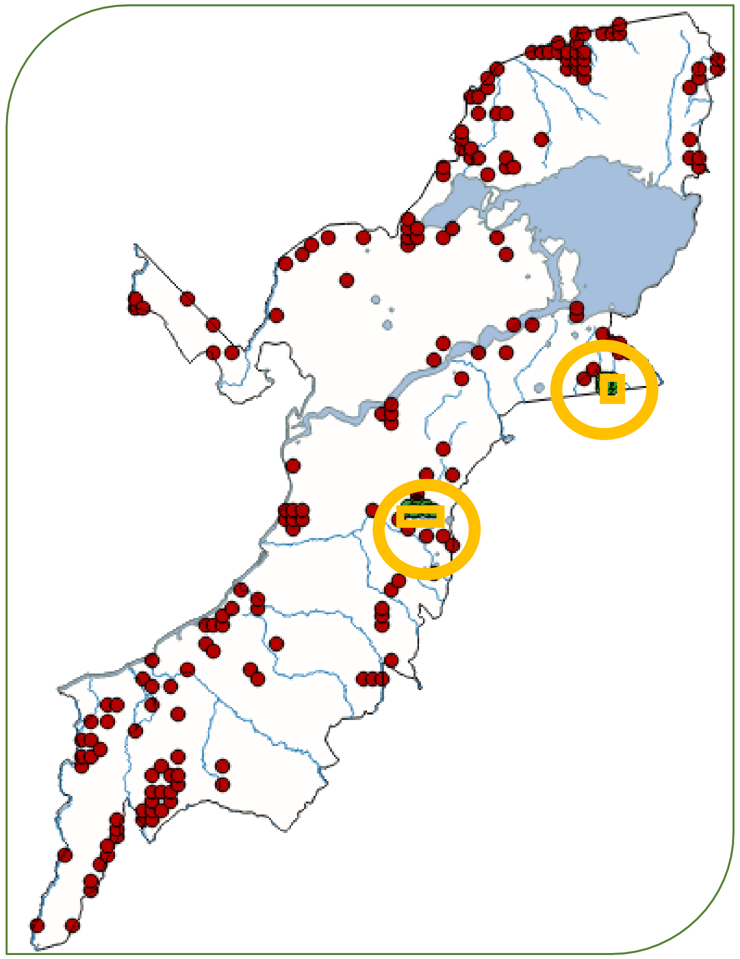

PAWS Real-world Deployment
Two Hot Spots Predicted

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

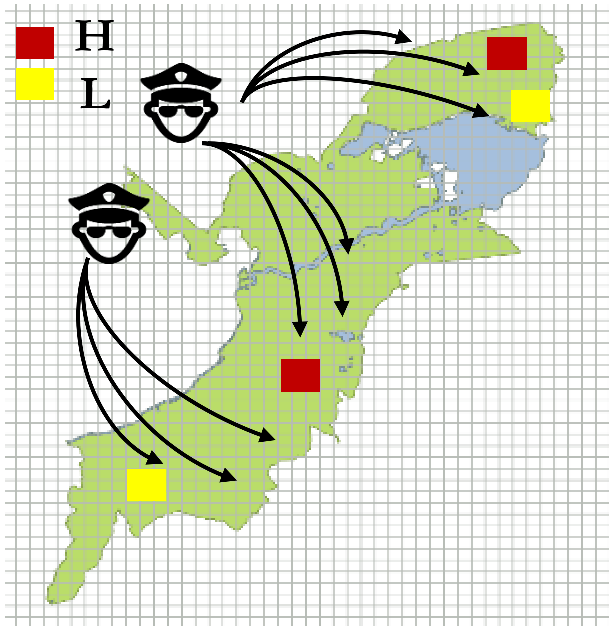

Queen Elizabeth National Park

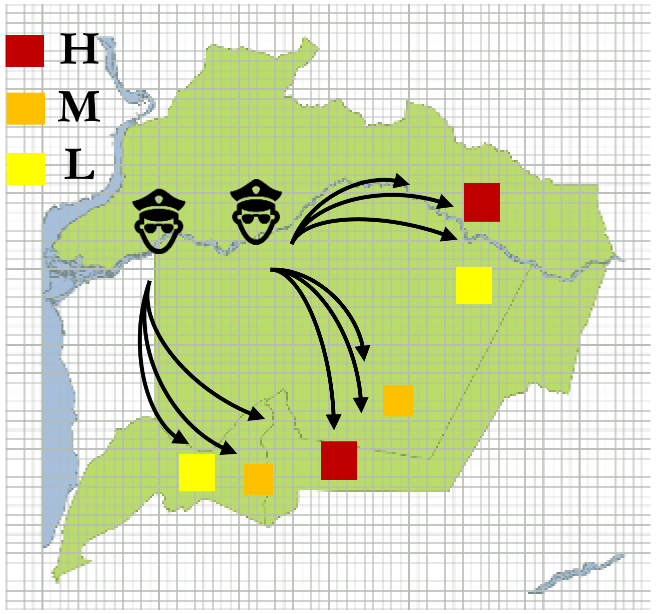

Murchison Falls National Park

Snares per patrolled sq. KM

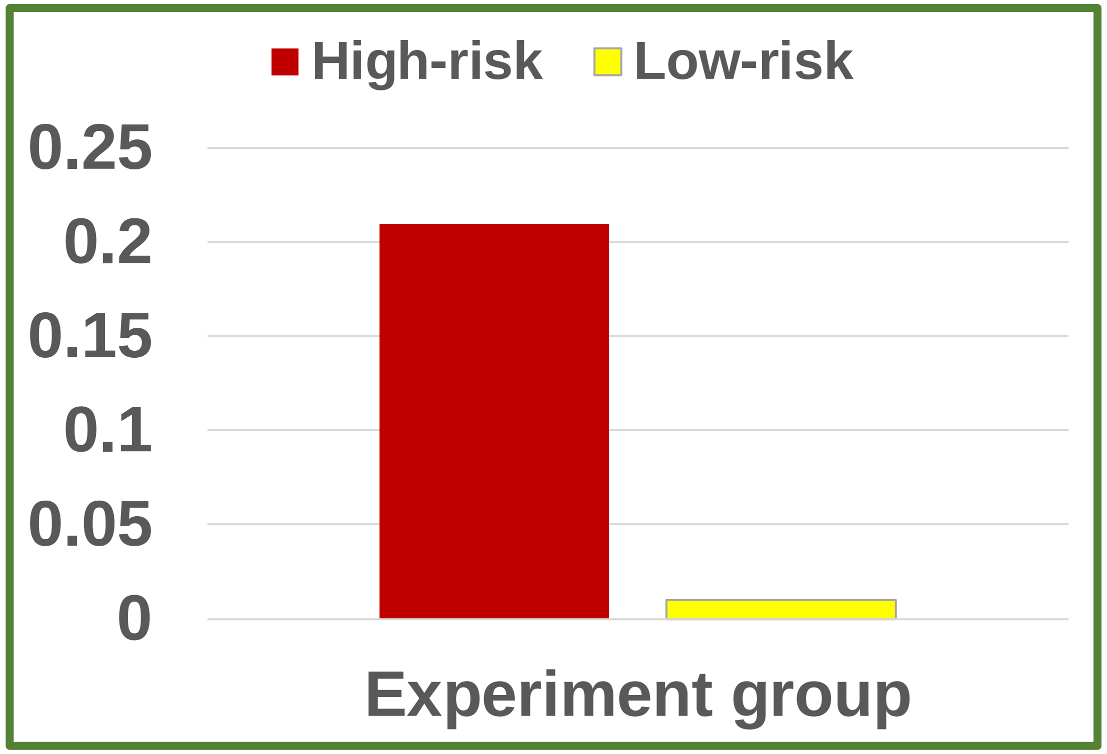

Snares per patrolled sq. KM

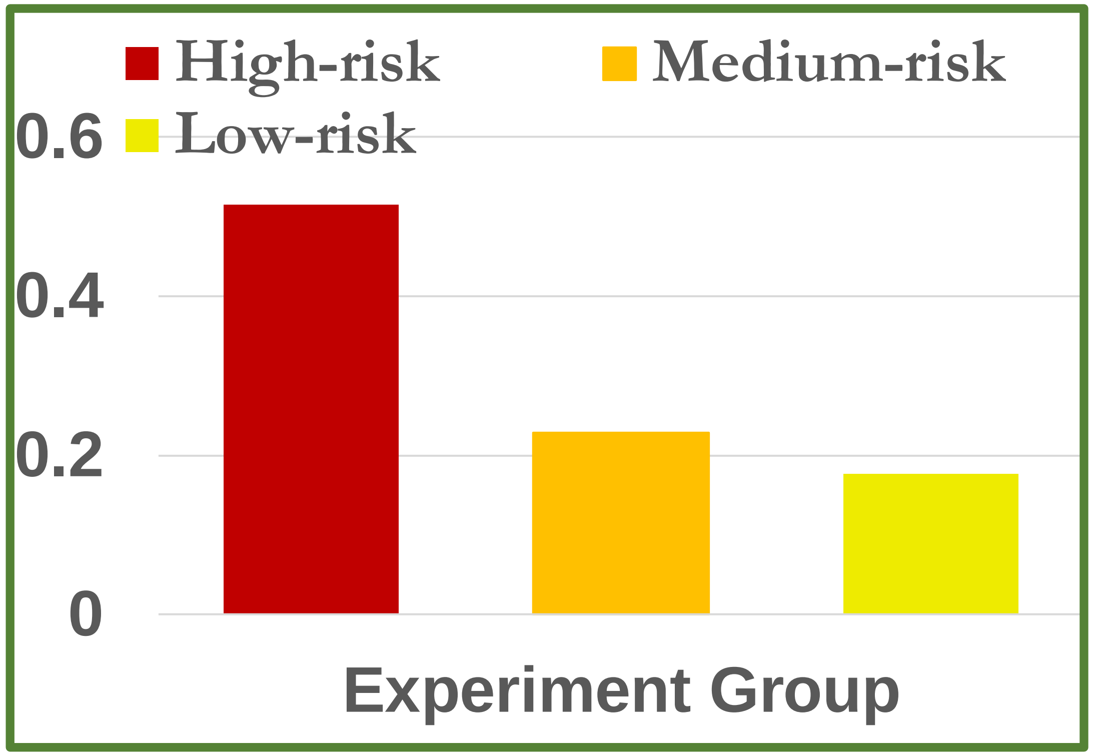

PAWS Real-world Deployment
Cambodia: Srepok Wildlife Sanctuary [2018-2019]

Xu

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

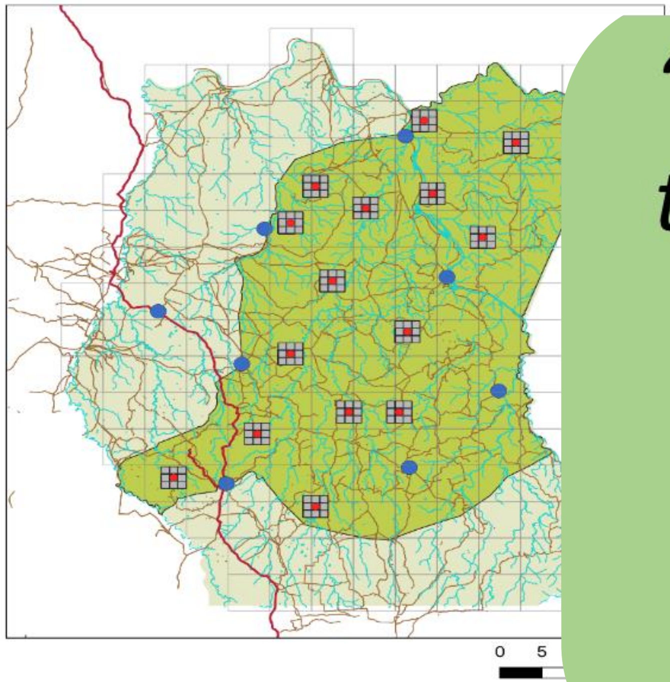

"@Milind: I am Super excited with the results. Let's get this going on other countries too this year."

Rohit Singh, WWF (2019)

- 521 snares/month our tests

vs

- 101 snares/month 2018

Snares per patrolled sq. KM

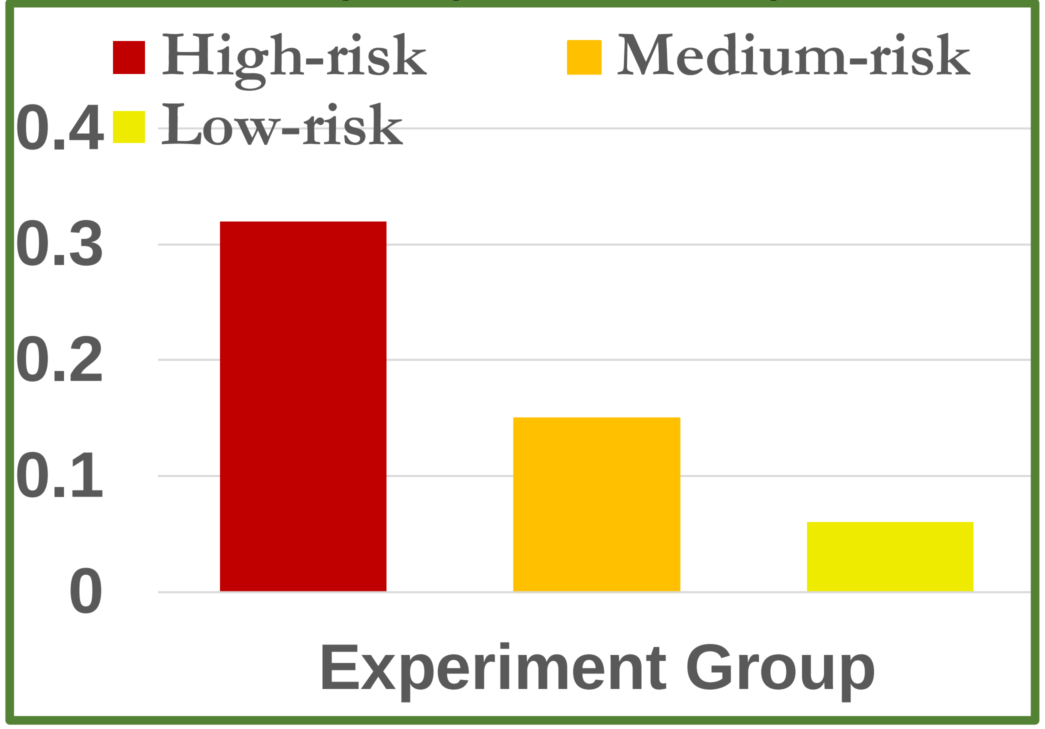

Green Security Games:
Around the Globe with SMART partnership [2019]

SMART

WWF

PEACE PARKS FOUNDATION · THE GLOBAL SOLUTION

WCS

AI for Earth

Protect Wildlife
800
National Parks
Around the Globe

Also: Protect Forests, Fisheries…

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

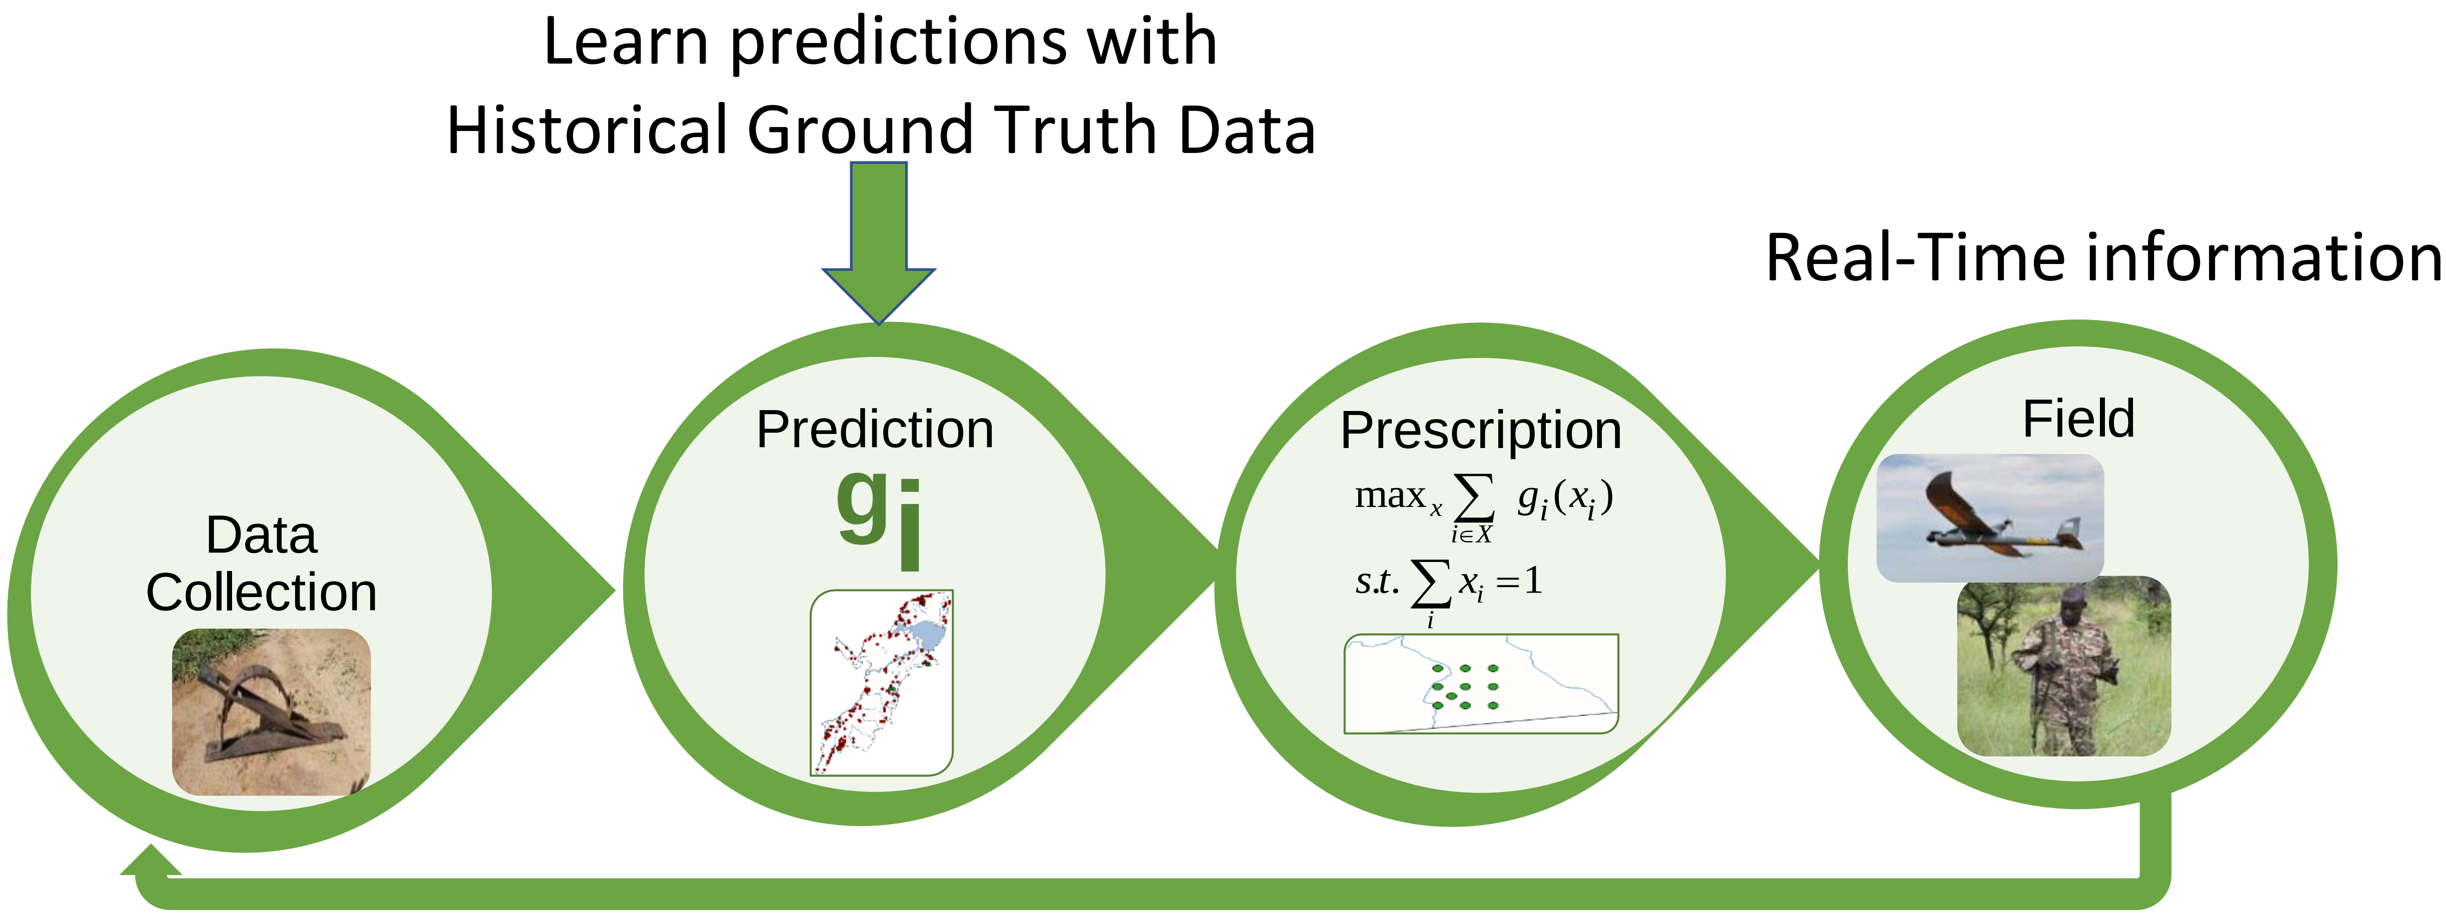

Green Security Games:
Integrating Real-Time “SPOT” Information [2018]


Bondi

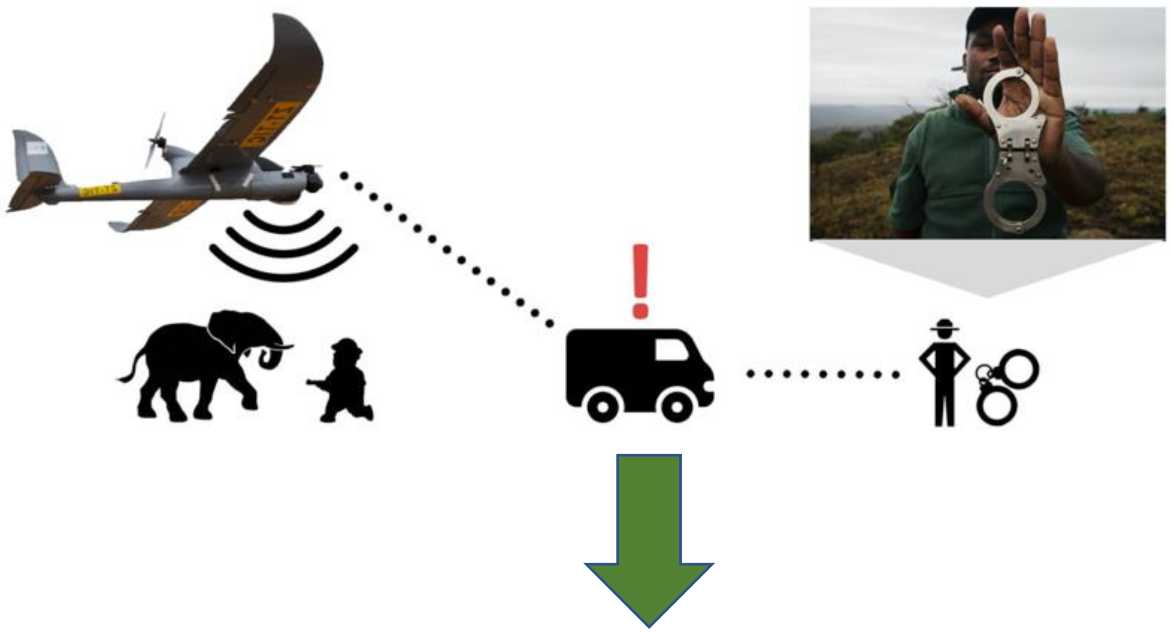

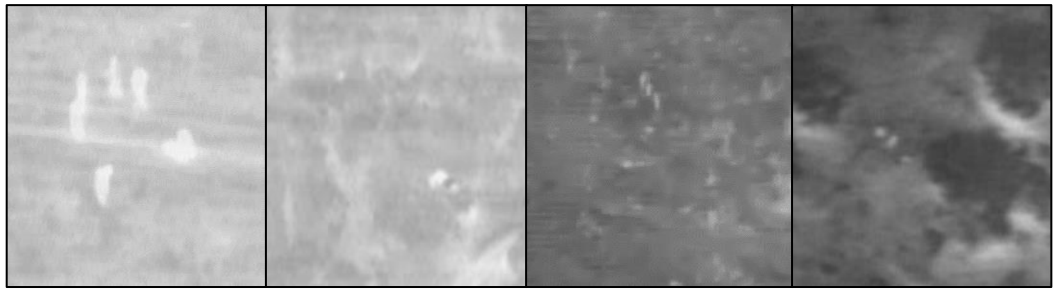

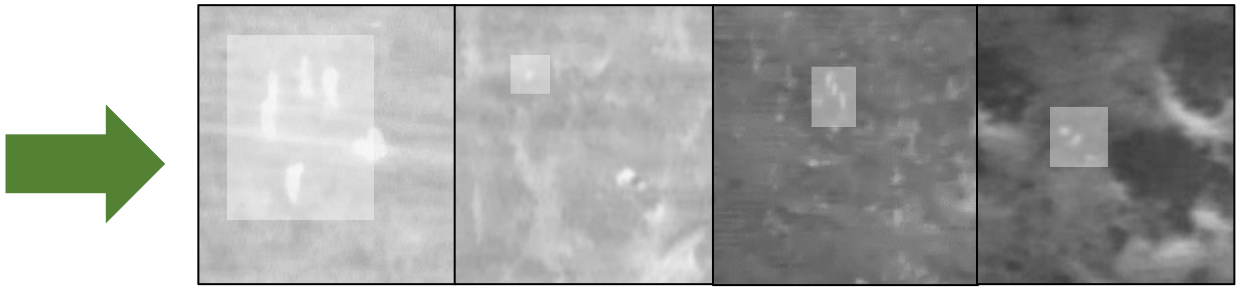

Goal: automatically find poachers

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

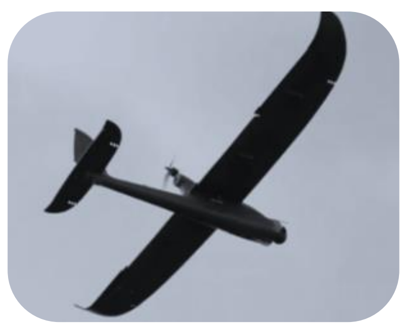

Air Shepherd
The Lindbergh Foundation

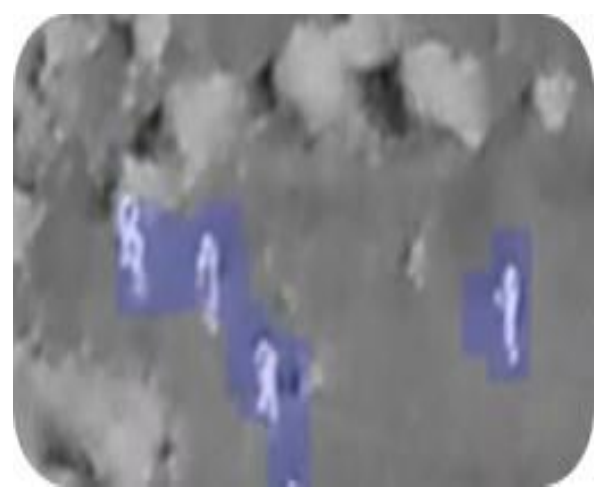

Prob(ranger) = 0.3

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

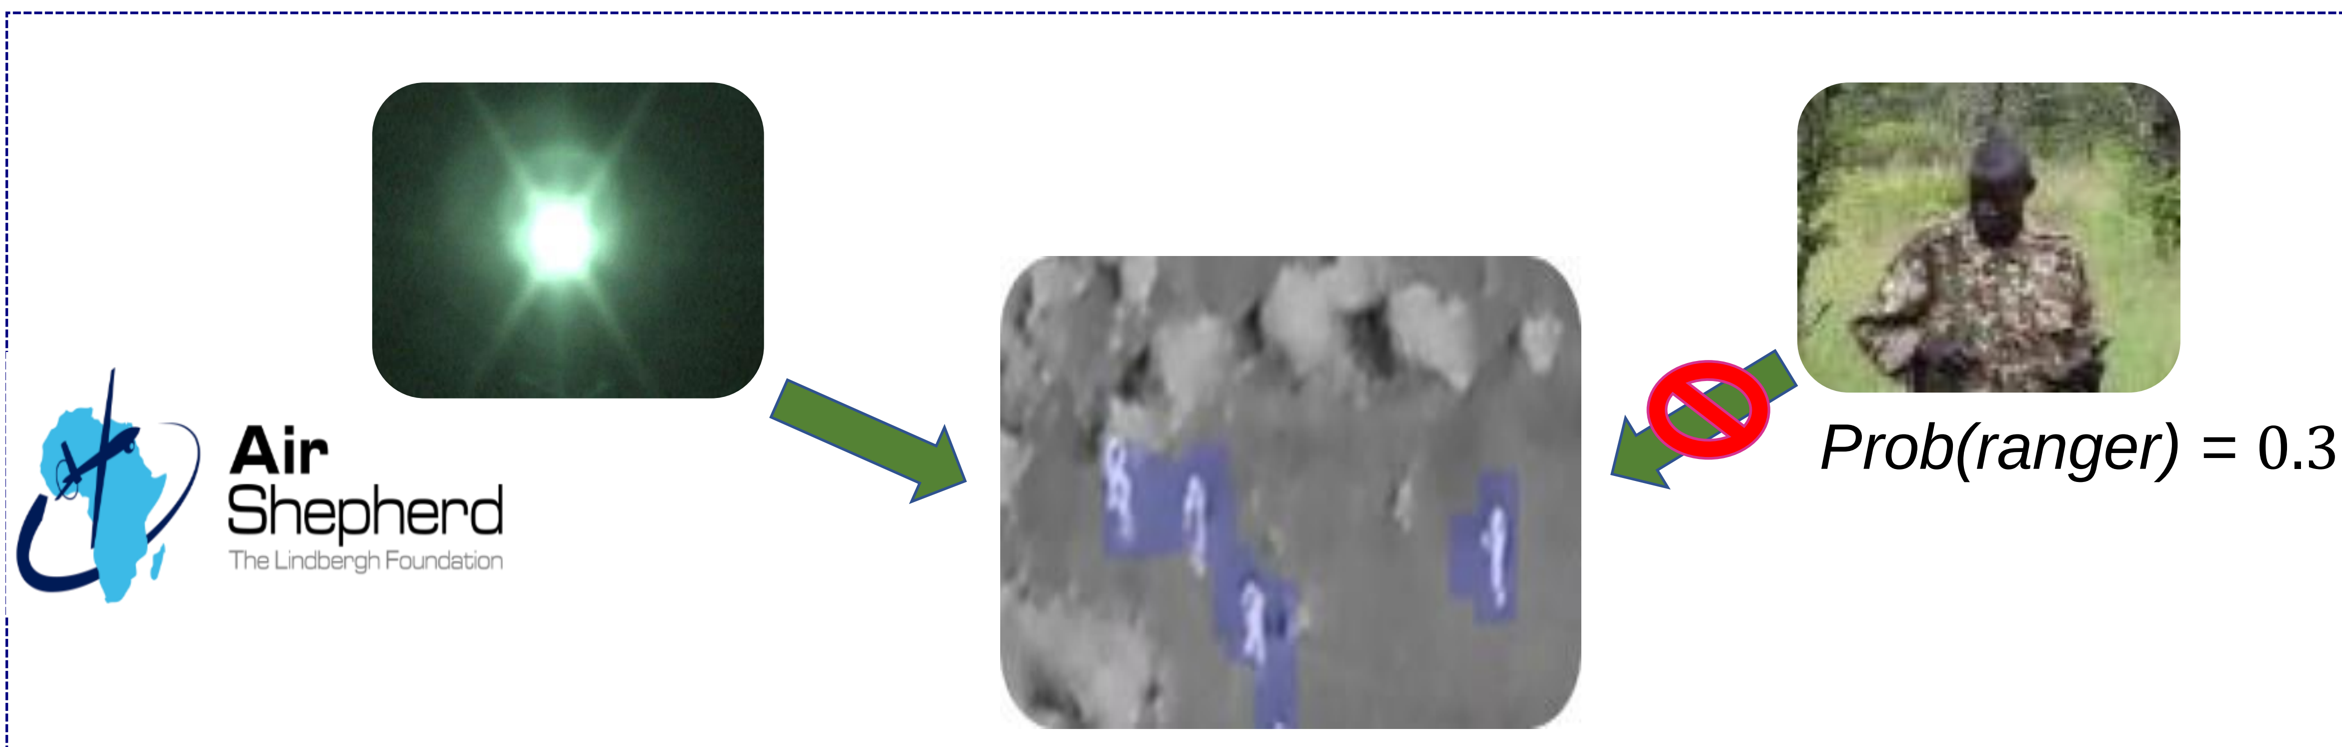

Strategic Signaling: Informational Advantage
Defender Knows Pure & Mixed Strategy

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

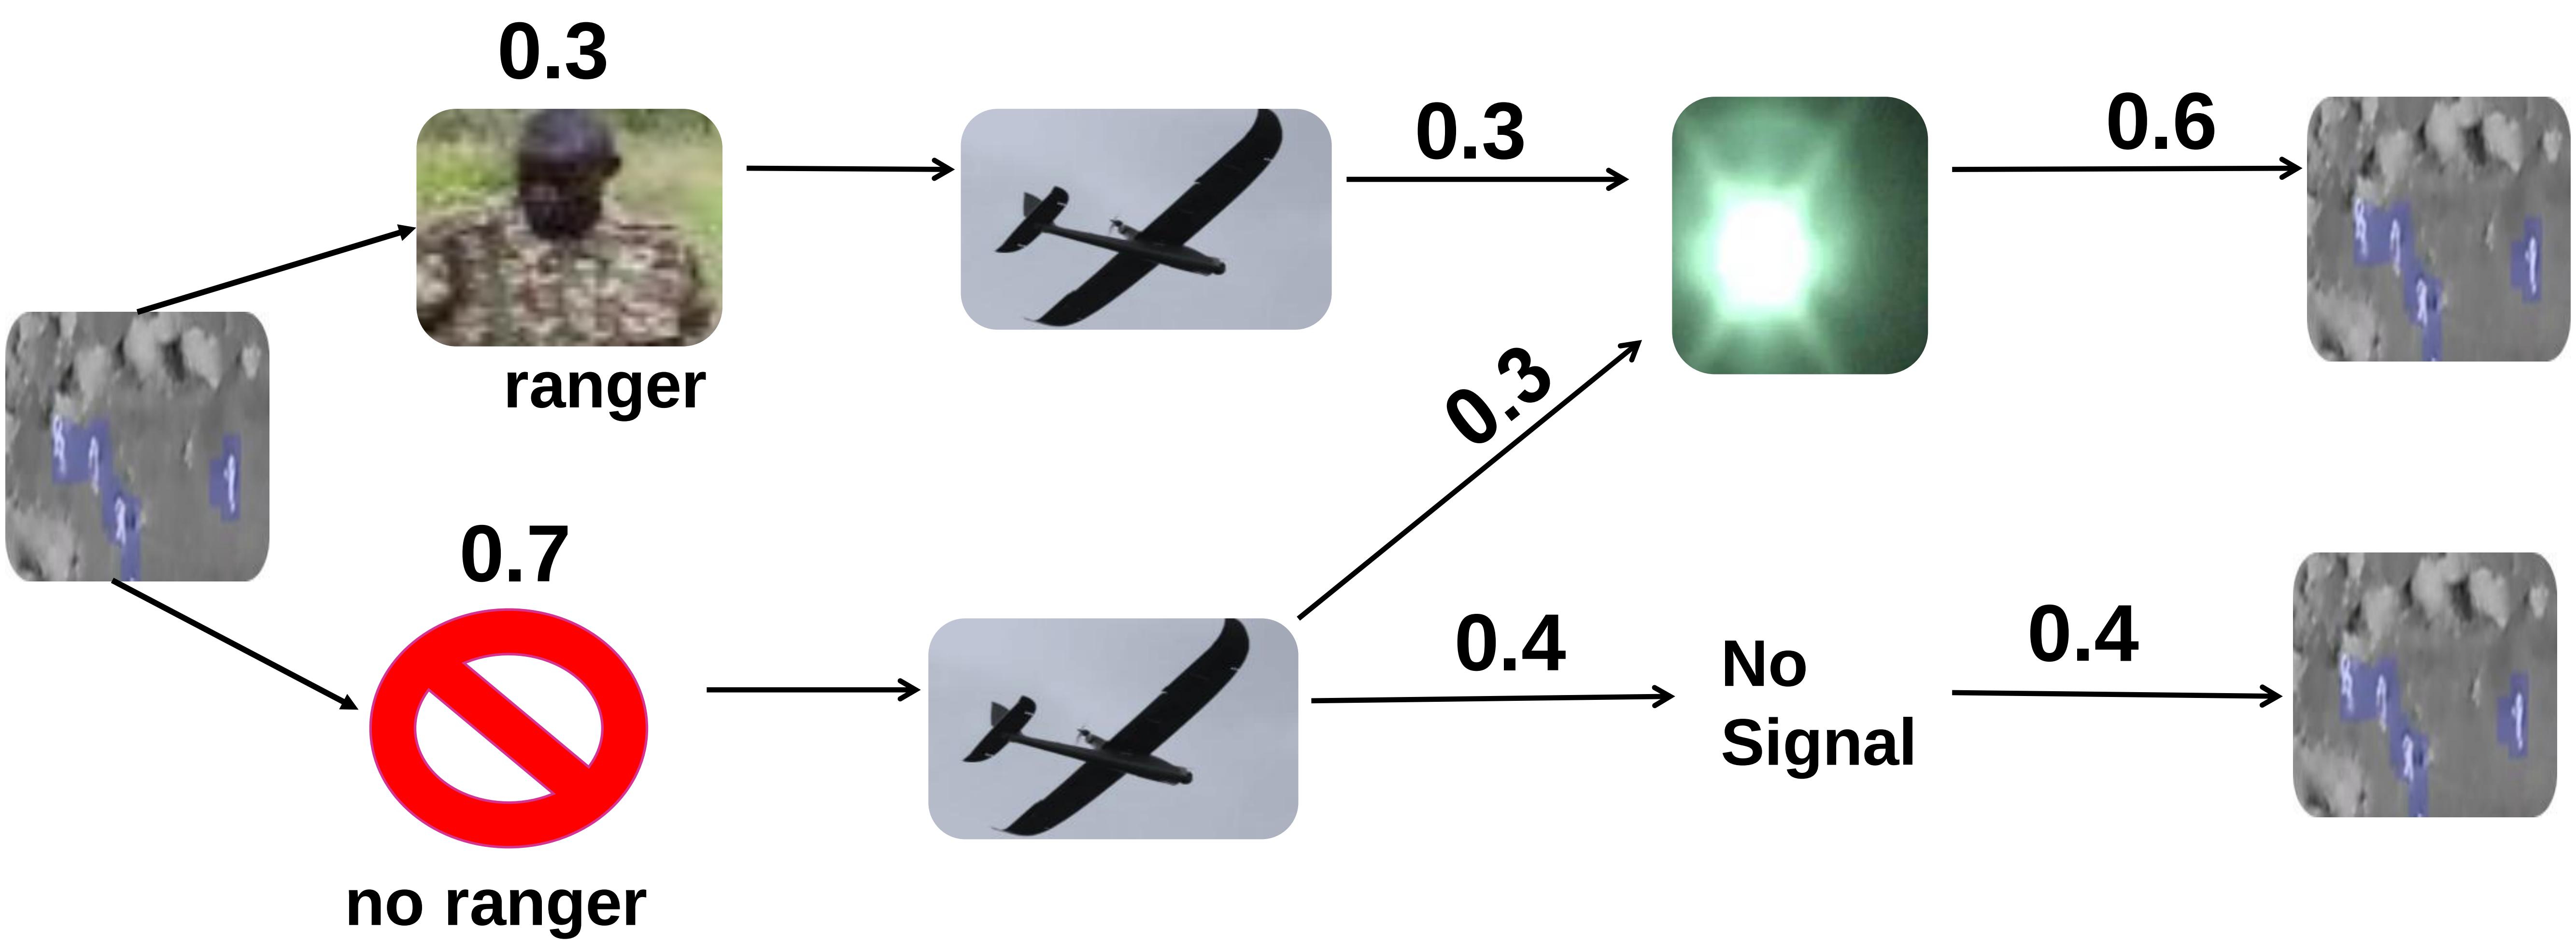

Strategic Signaling: Handling Detection Error
Exploit Informational Asymmetry to Mitigate Impact

Strategic signaling in presence of error in detecting adversaries

➢ Signal selectively when no adversary detection

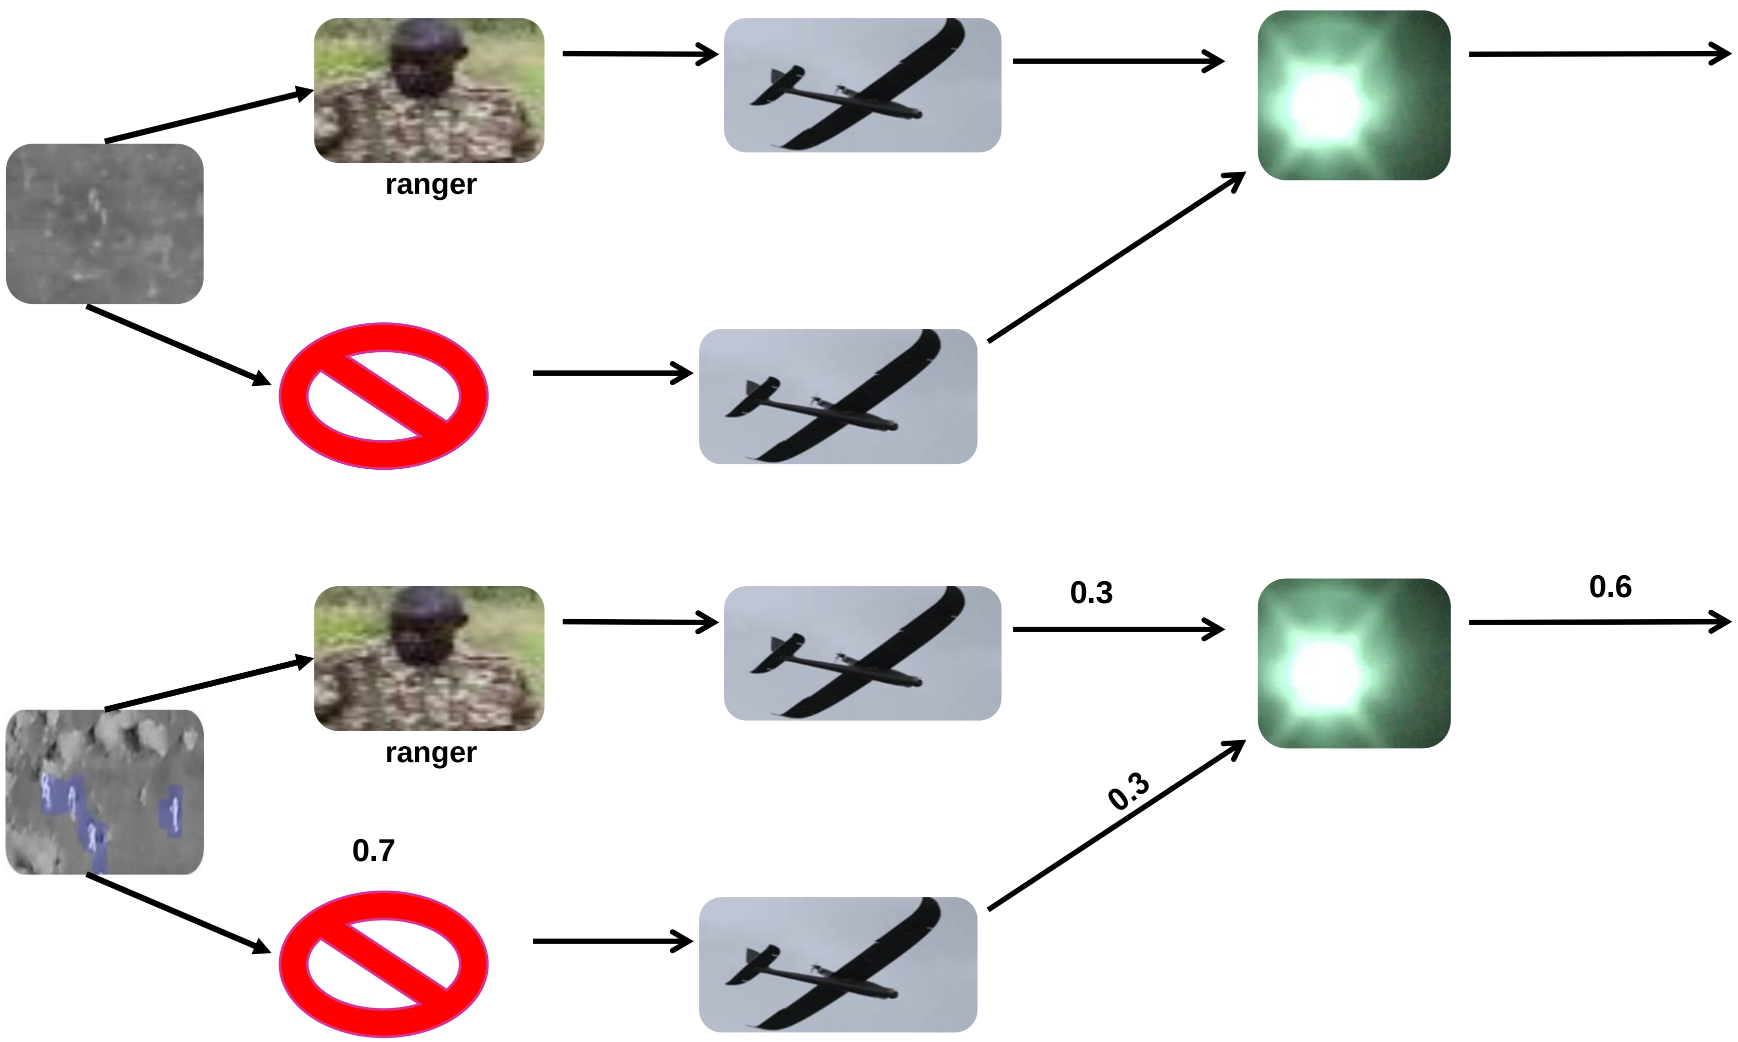

Green Security Games Pipeline: Not Enough Data in Some Parks

Perrault

- 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

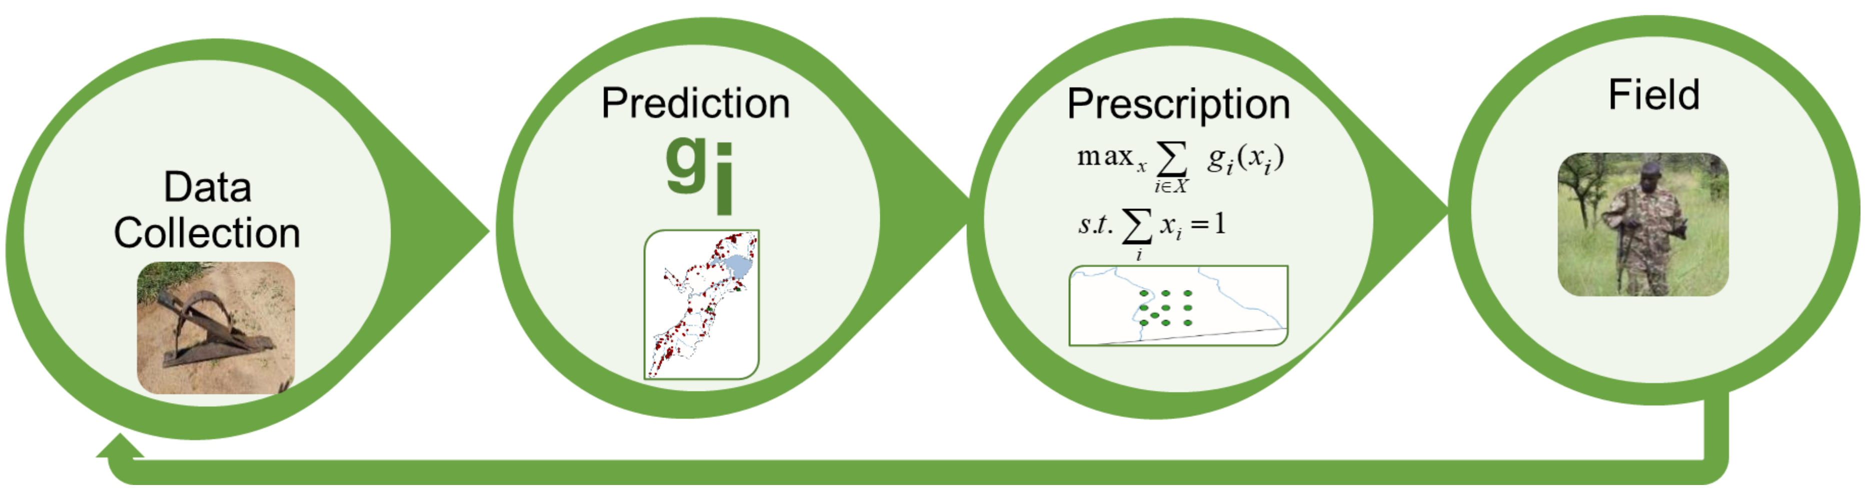

Previous Stage-by-Stage Method:
Make Prediction as Accurate as Possible Then Plan



Perrault Wilder Mate

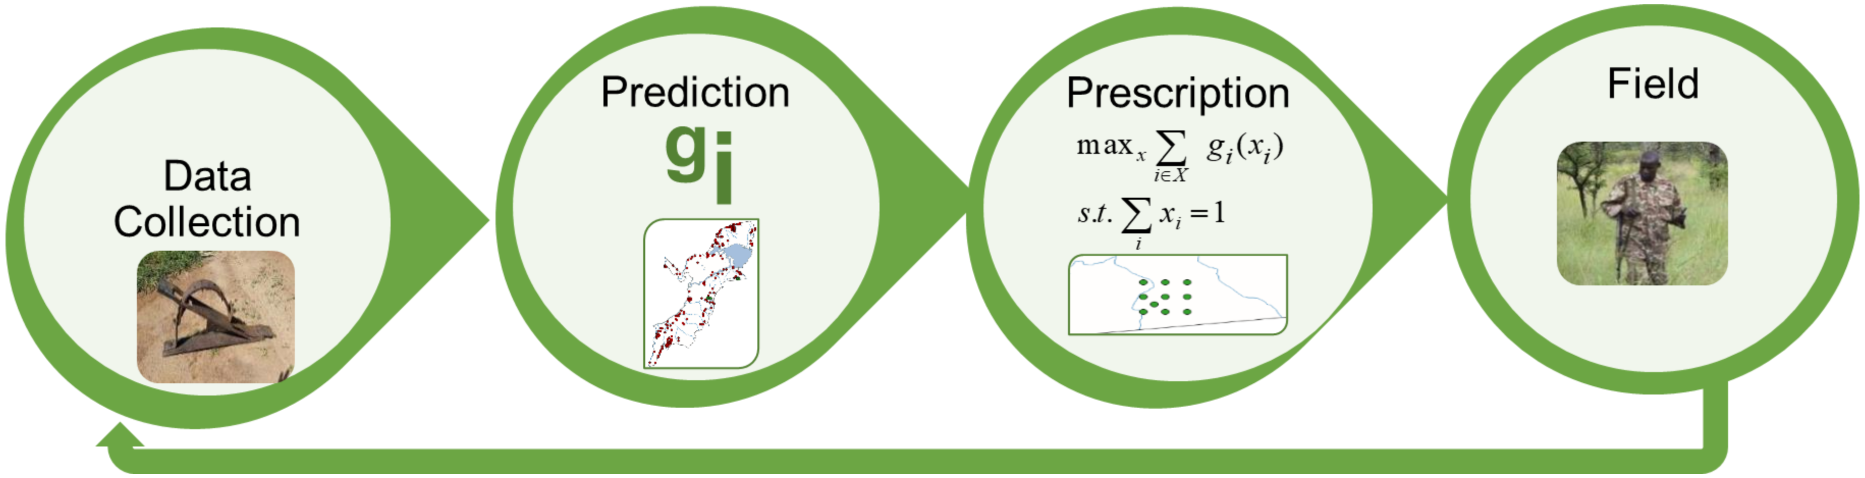

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

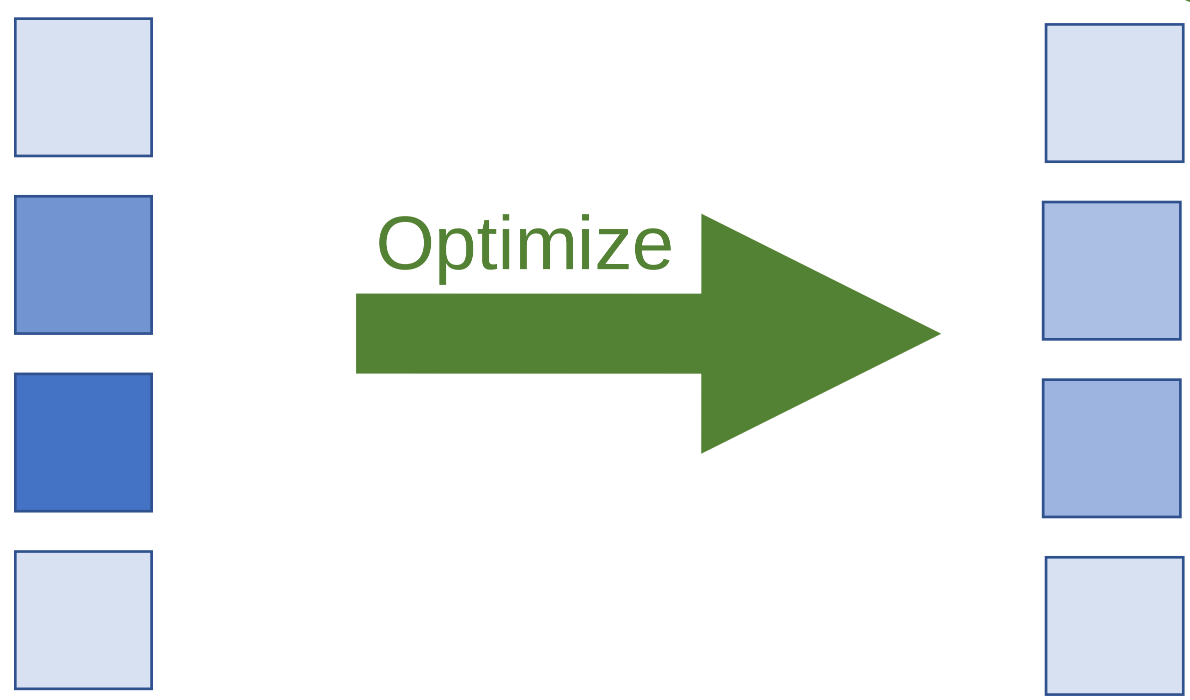

Minimize $\sum_{\{i \in T\}} q_{\text{empirical}} \log \hat{q}$

Stage by Stage Method:
Need to Focus on Important Targets

Perrault Wilder Mate

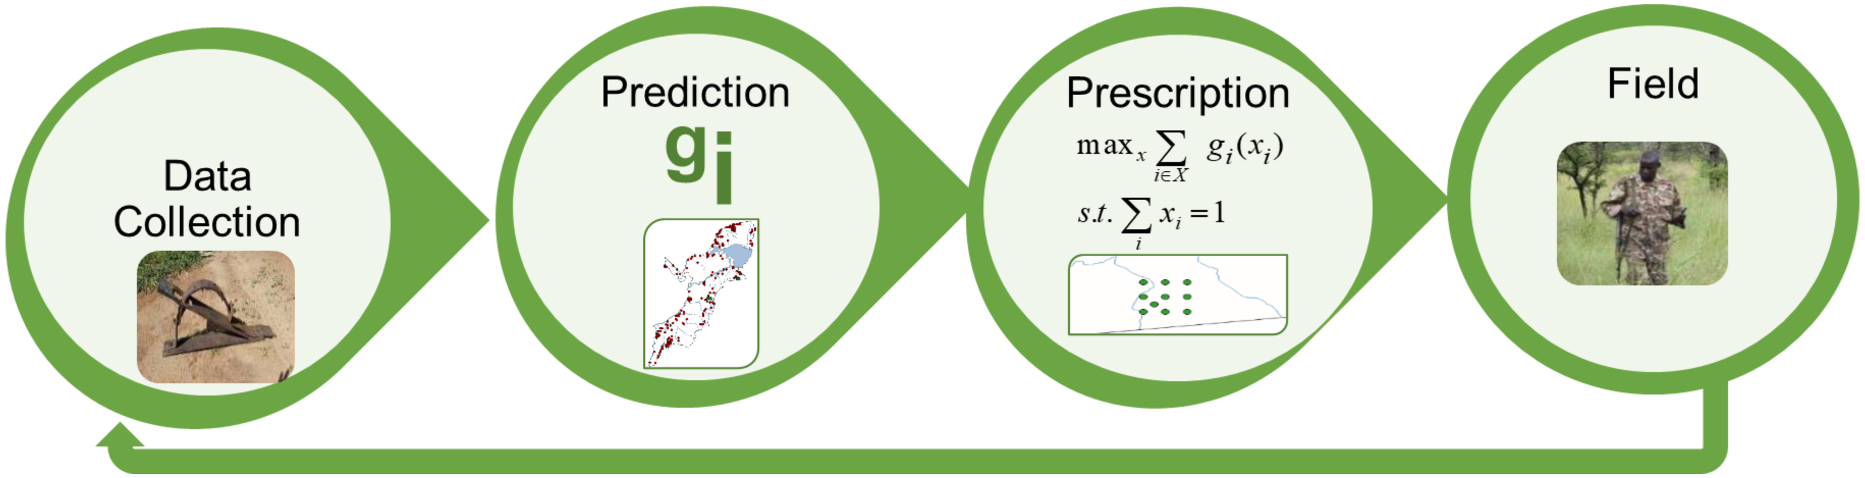

Maximize accuracy in Adversary target values ***estimates***

Plan patrol Coverage

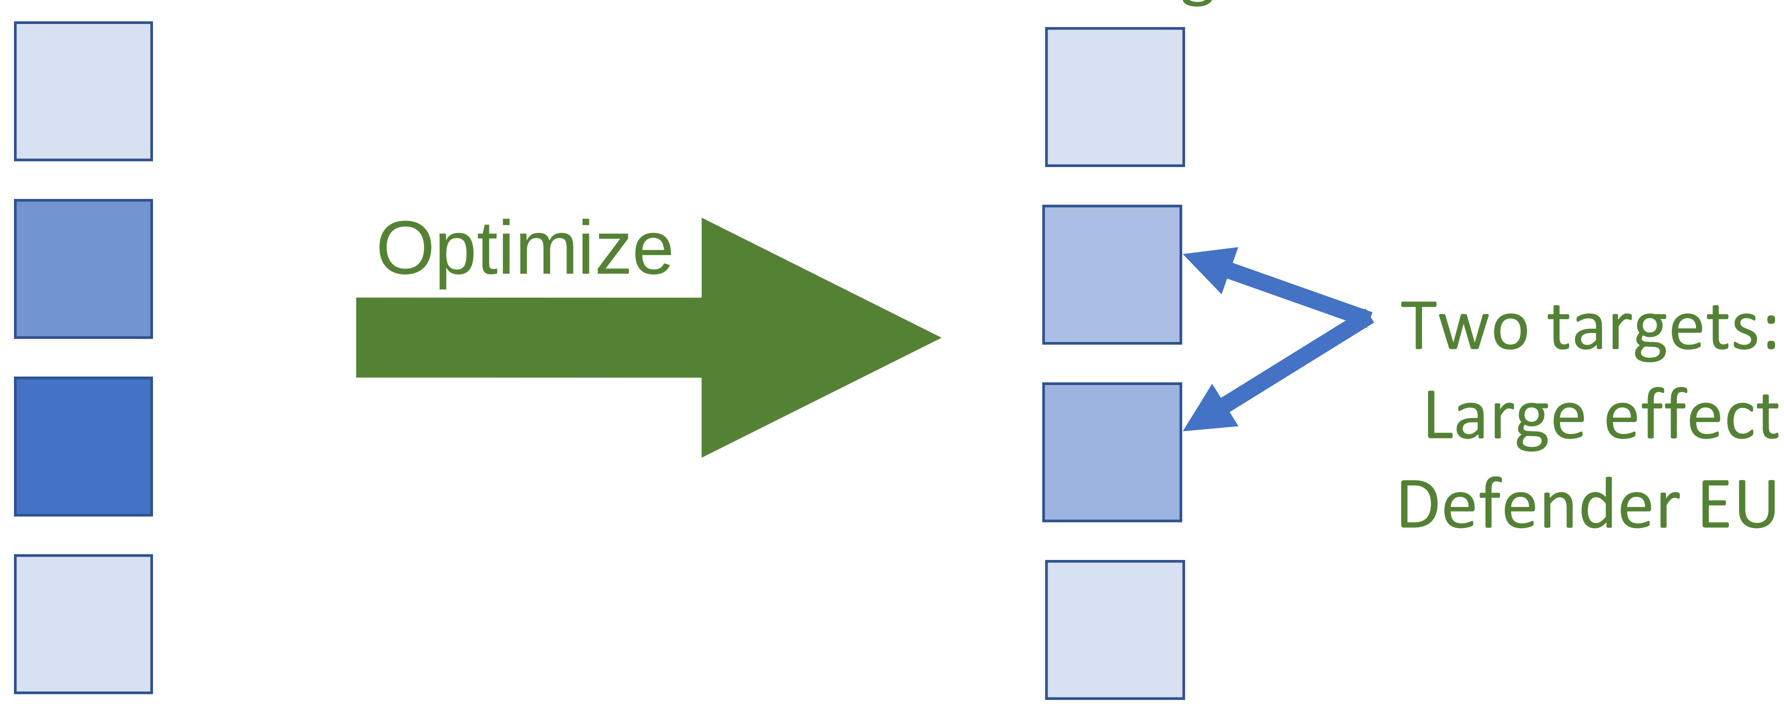

Minimize $\sum_{\{i \in T\}} q_{\text{empirical}} \log \hat{q}$

Game-Focused Learning:
Need to Focus on Important Targets

Perrault Wilder Mate

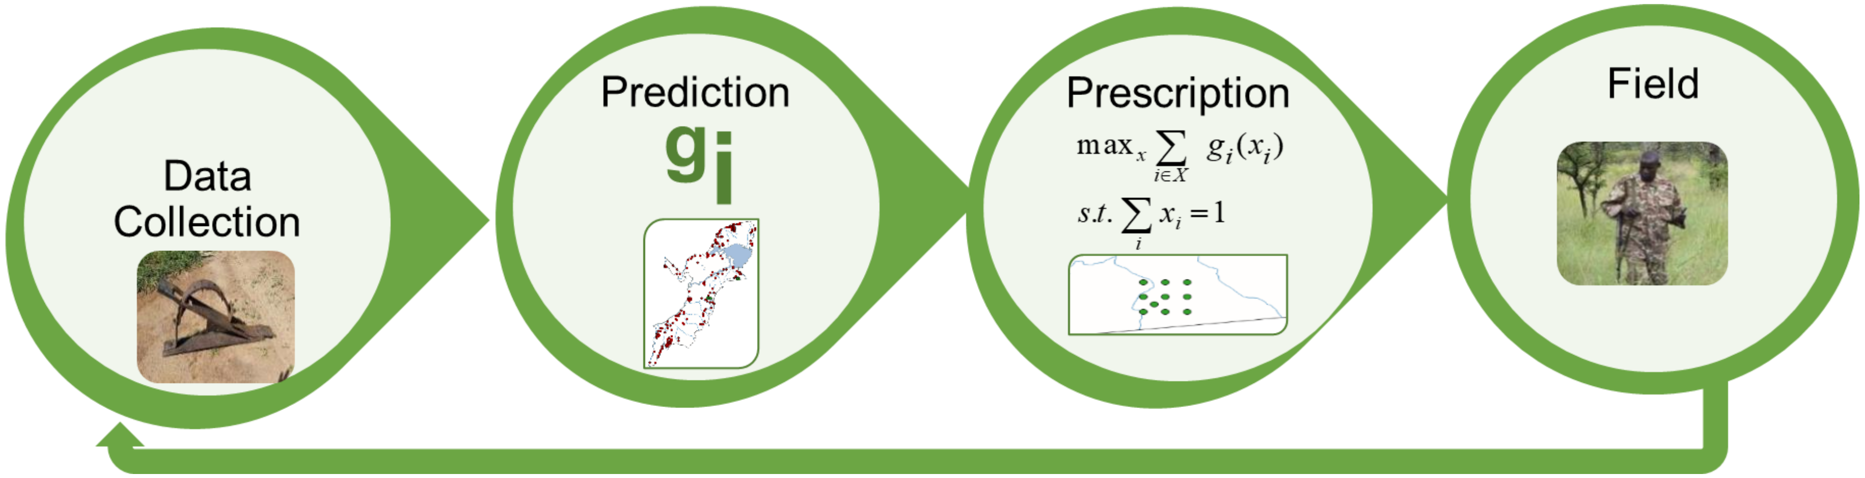

Maximize accuracy only of ***Important targets***

Plan patrol Coverage

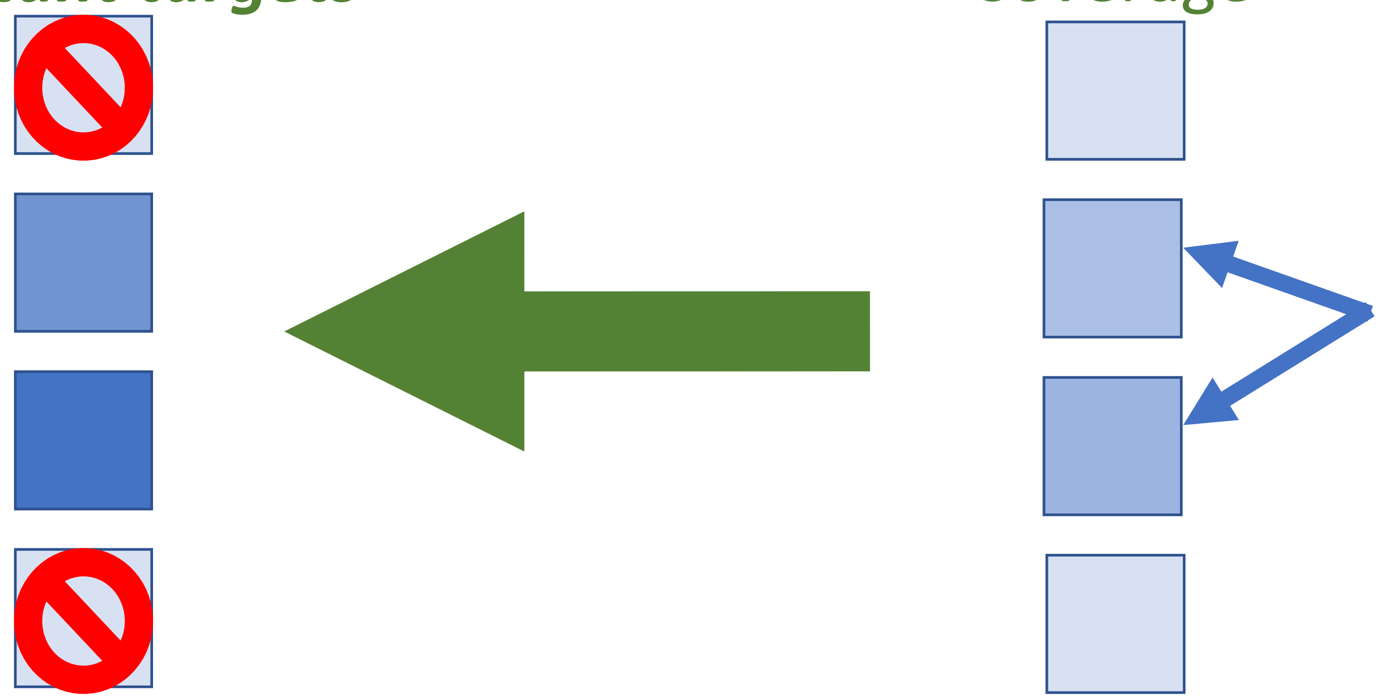

Two targets: Large effect Defender EU

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

Game-Focused Learning:
End-to-End Method

Perrault Wilder Mate

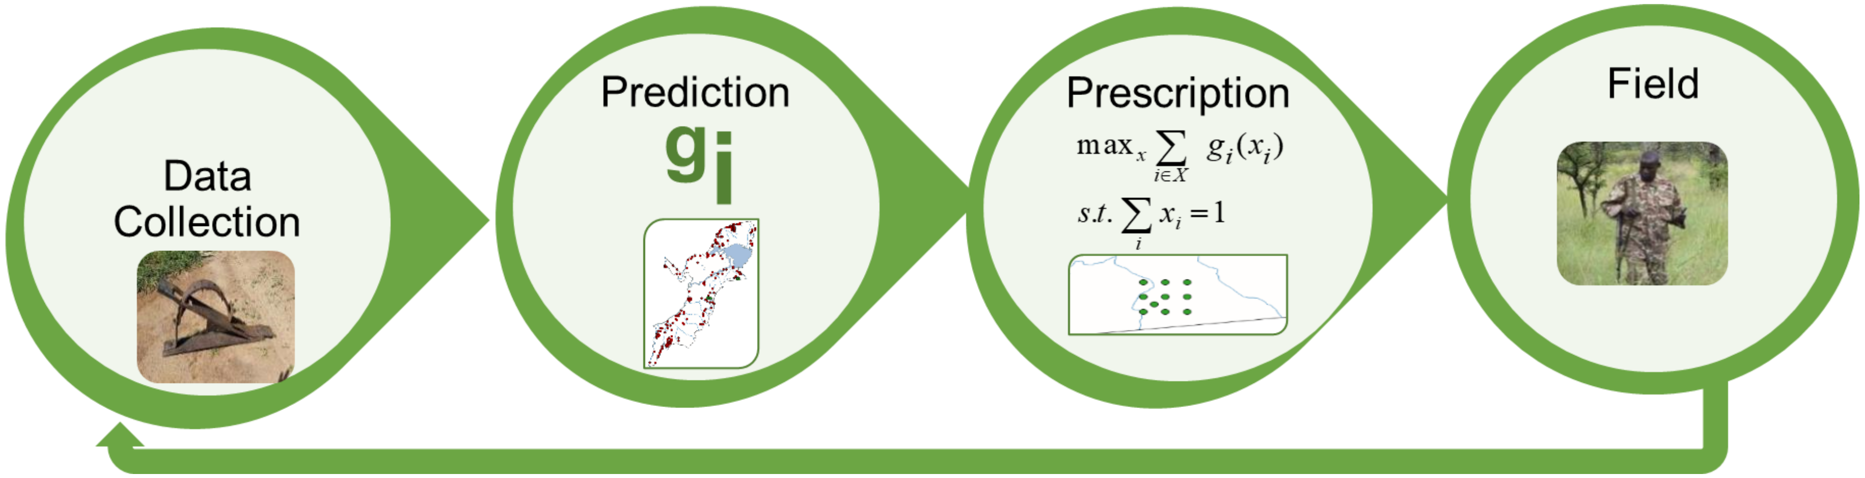

Maximize
defender expected *utility*

Plan patrol
Coverage

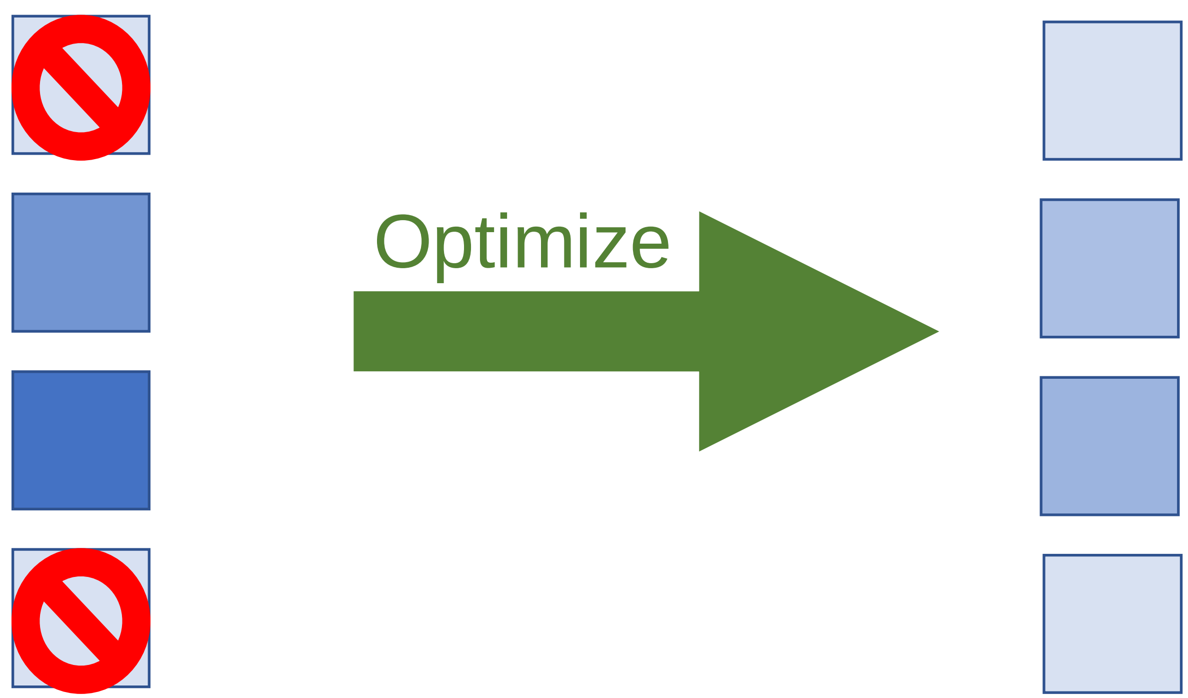

Maximize defender's expected utility $\sum (1 - p_i(\hat{q}))\, q_{\text{empirical}}$

Previous Two-Stage Method: Gradient Descent

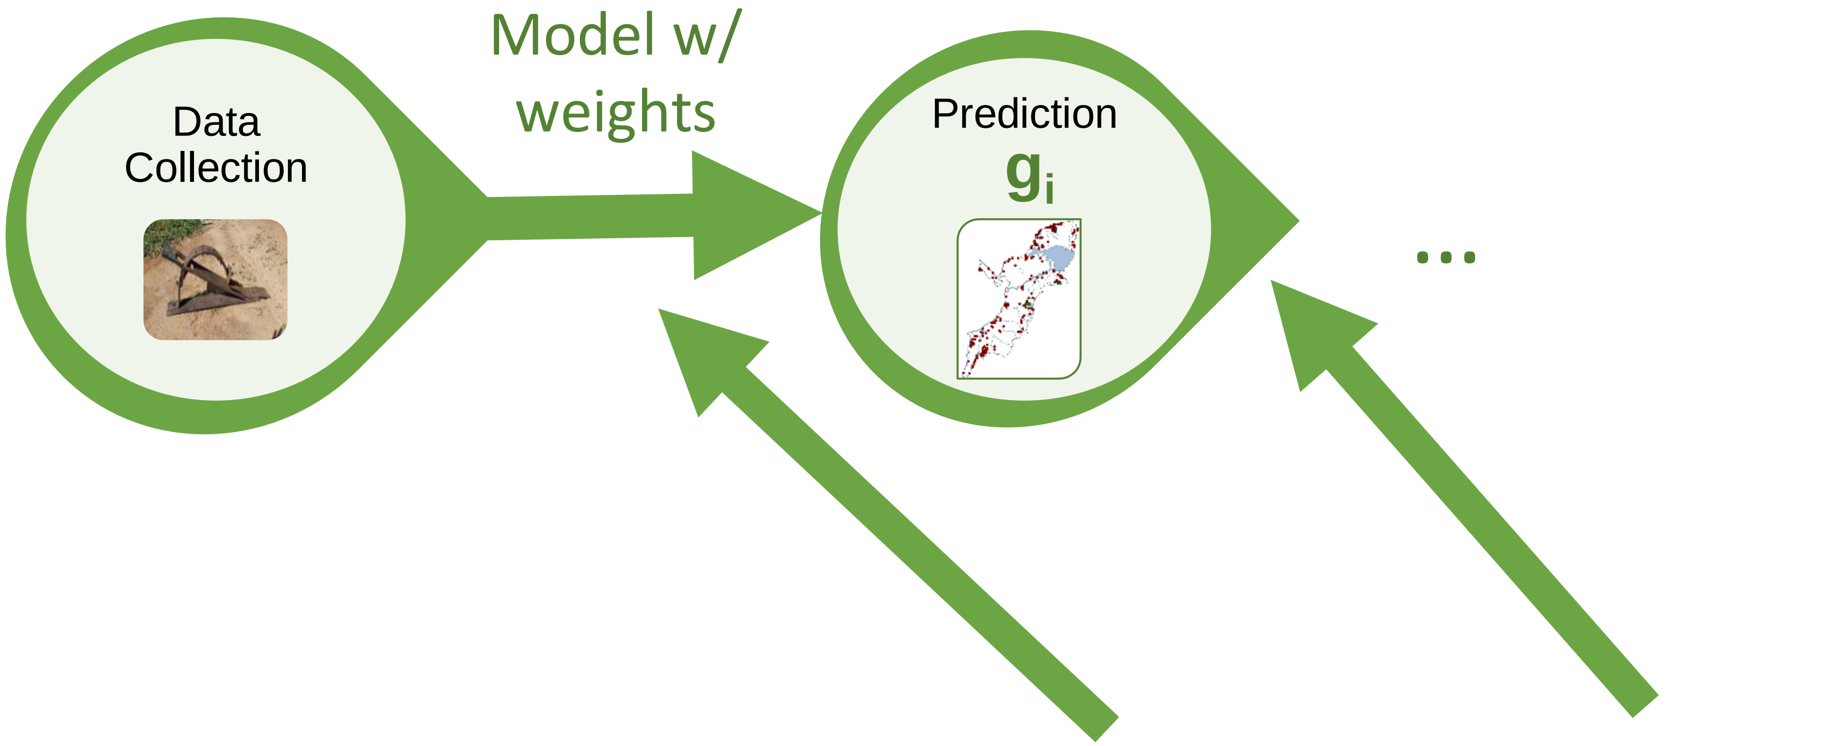

- Max accuracy gradient descent: $$\frac{\partial \text{accuracy}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{accuracy}}{\partial \text{prediction}}$$

Game-Focused Learning: End-to-End Method


Perrault Wilder Mate

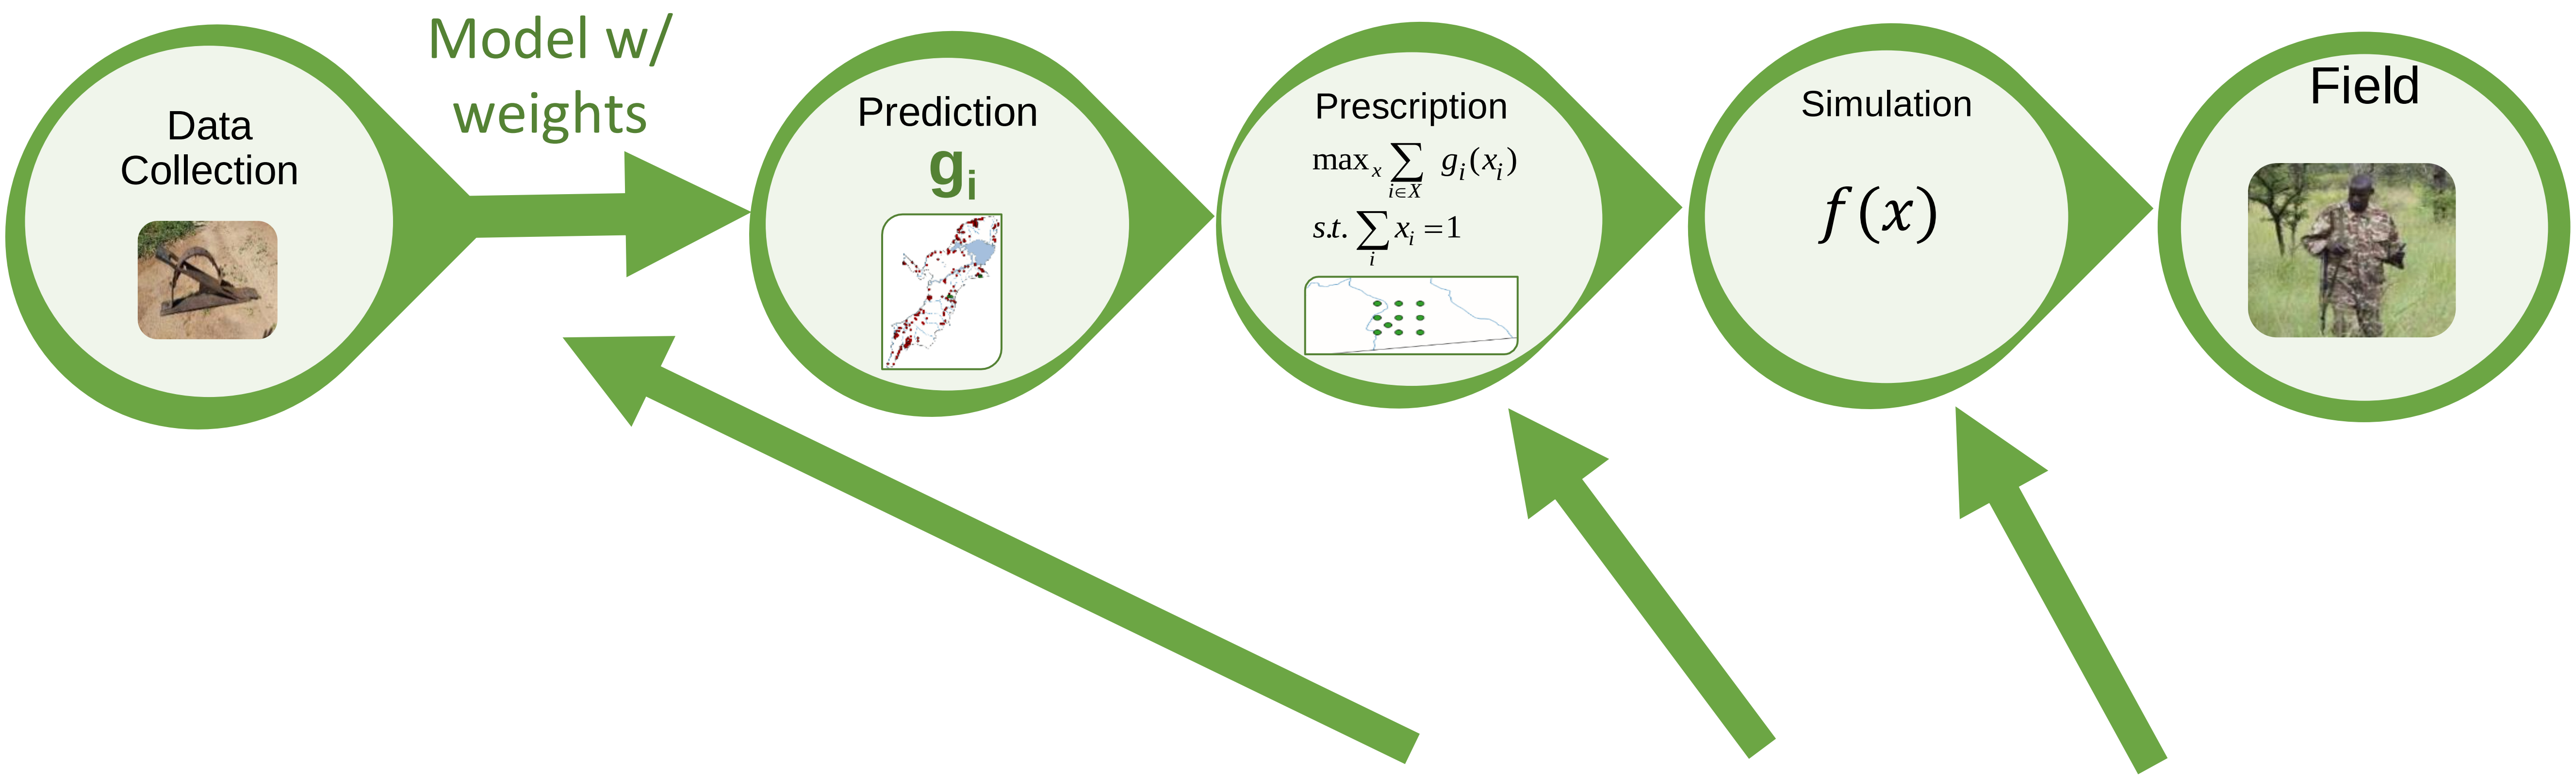

- Game-focused gradient descent: $$\frac{\partial \text{obj(decision)}}{\partial \text{weights}} = \frac{\partial \text{prediction}}{\partial \text{weights}} \frac{\partial \text{decision}}{\partial \text{prediction}} \frac{\partial \text{obj(decision)}}{\partial \text{decision}}$$

Game-Focused Learning:
End-to-End Method

- Focusing learning on important targets increases defender utility

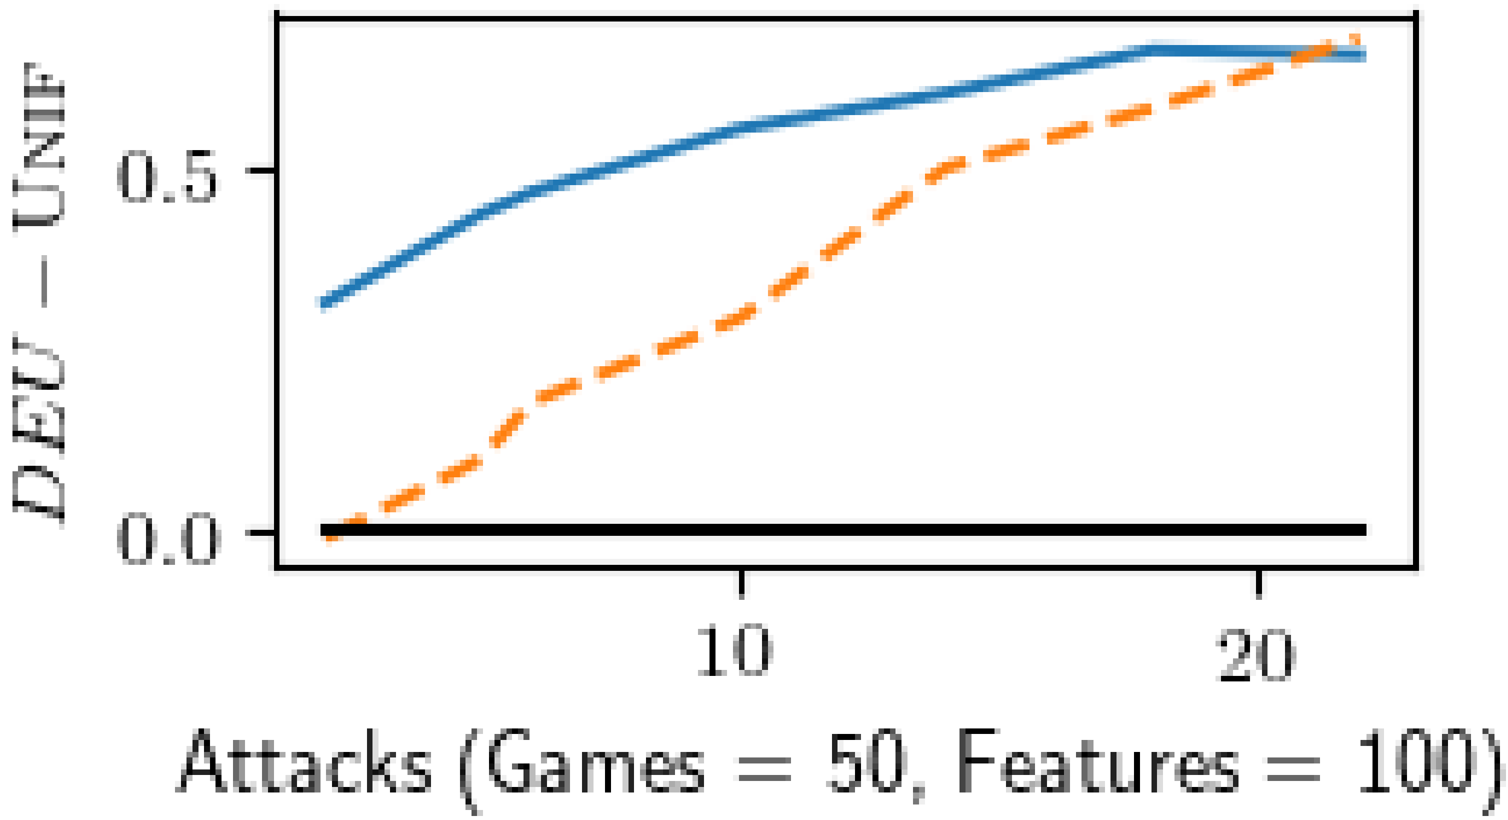

Outline

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

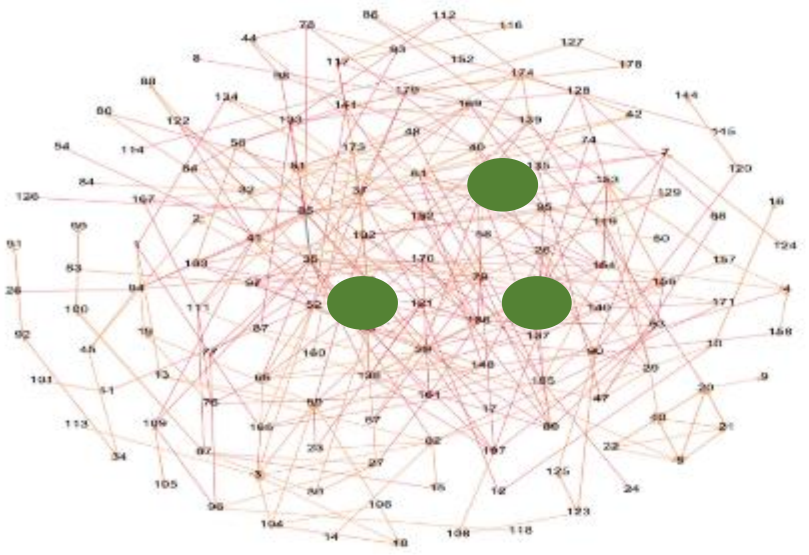

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
- Objective:
  - Maximize expected number of influenced nodes
- *Assumption: Independent cascade model of information spread*

Independent Cascade Model and Real-world Physical Social Networks

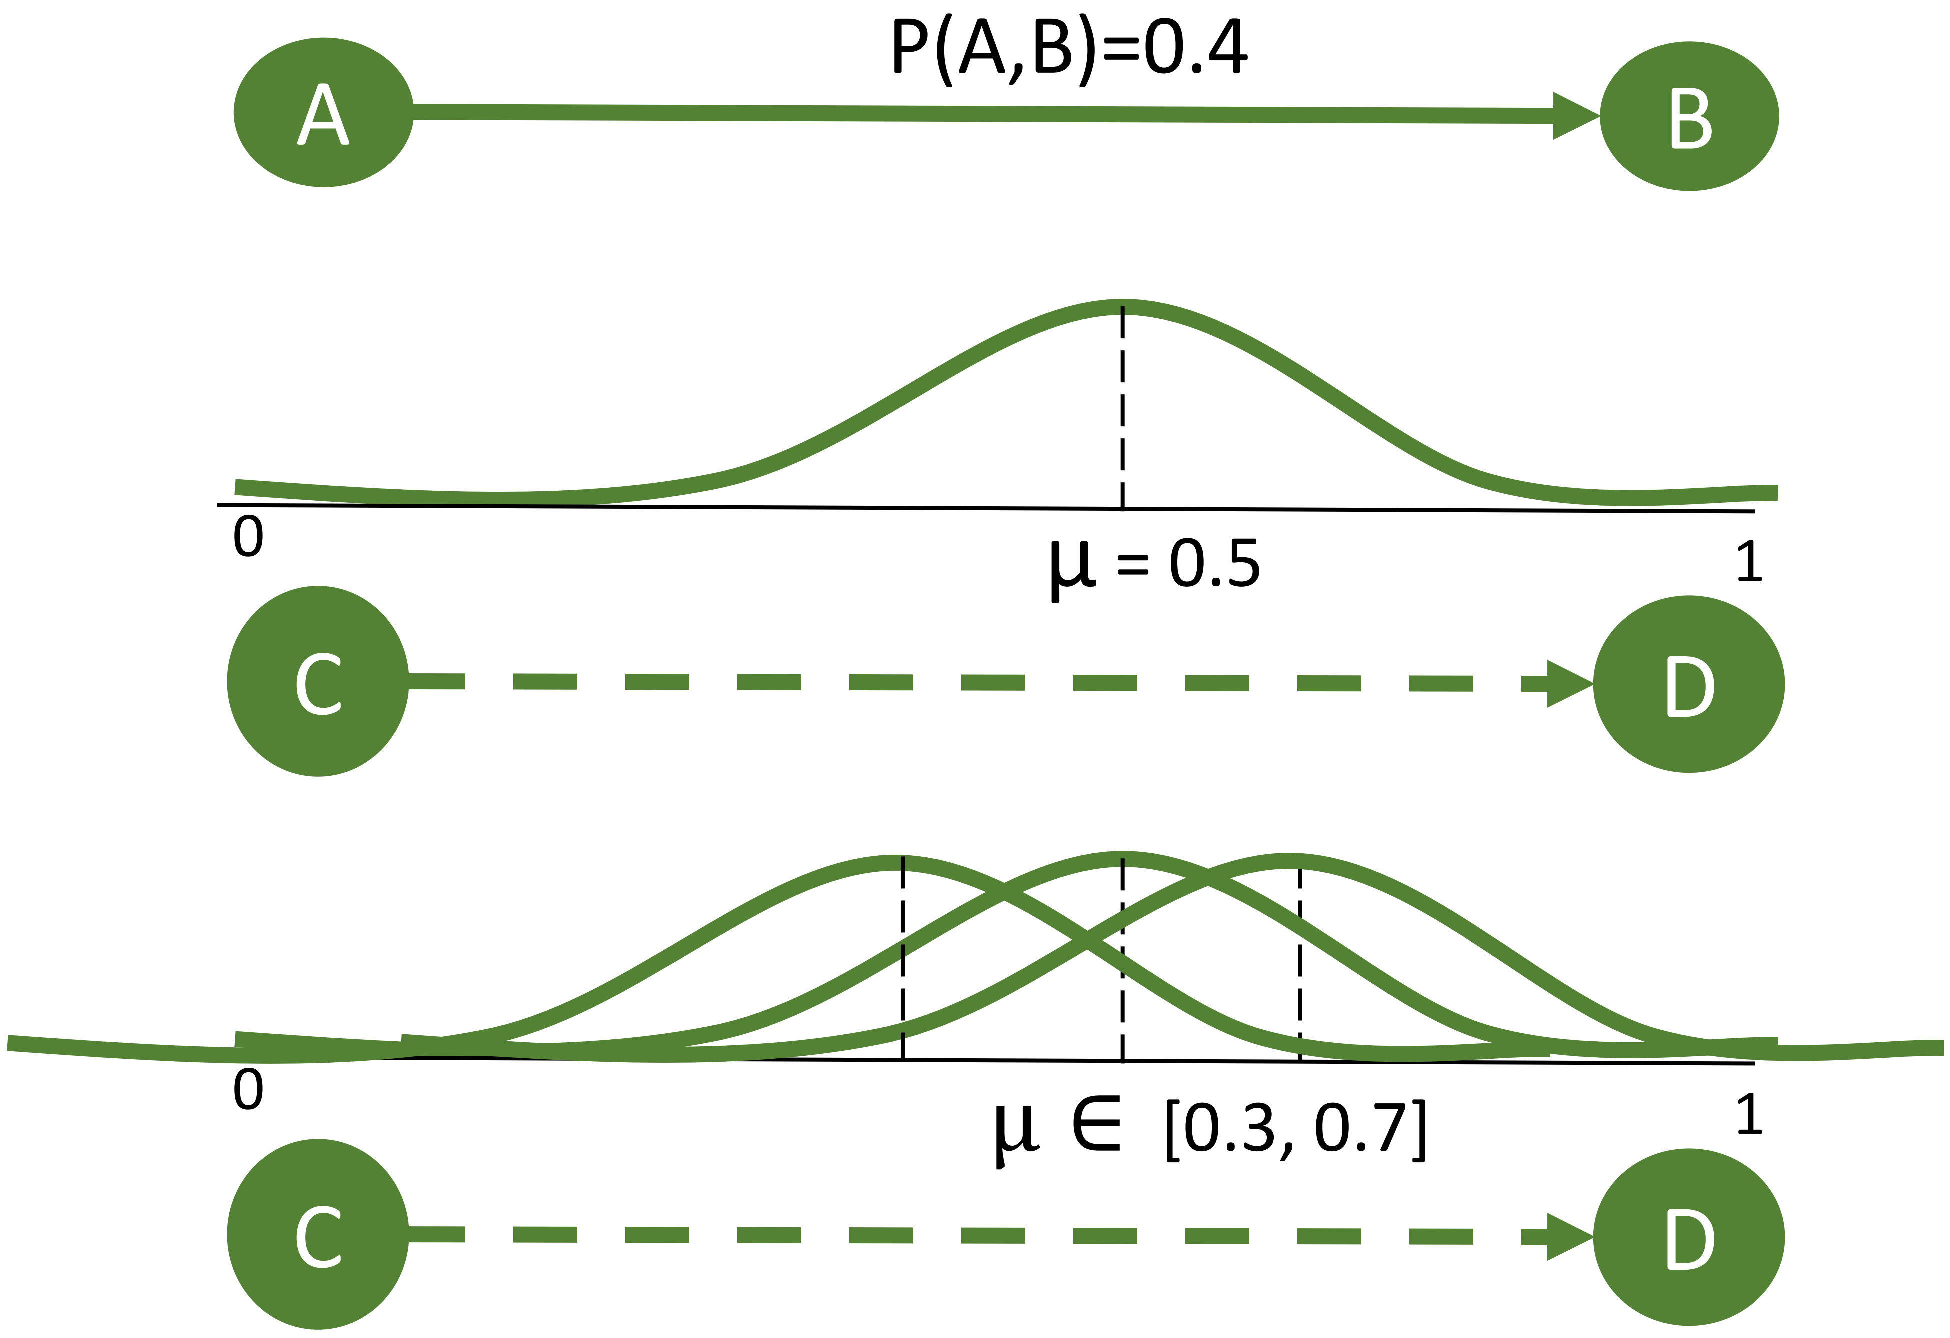

Robust, Dynamic Influence Maximization

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


Yadav


Wilder

Challenge: Multi-step Policy

	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

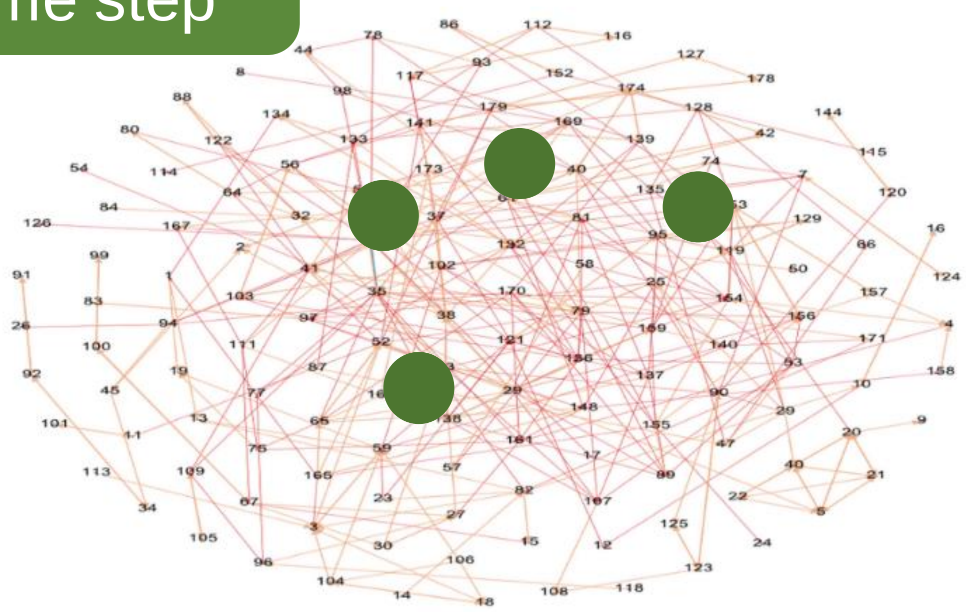

K = 4
2nd time step

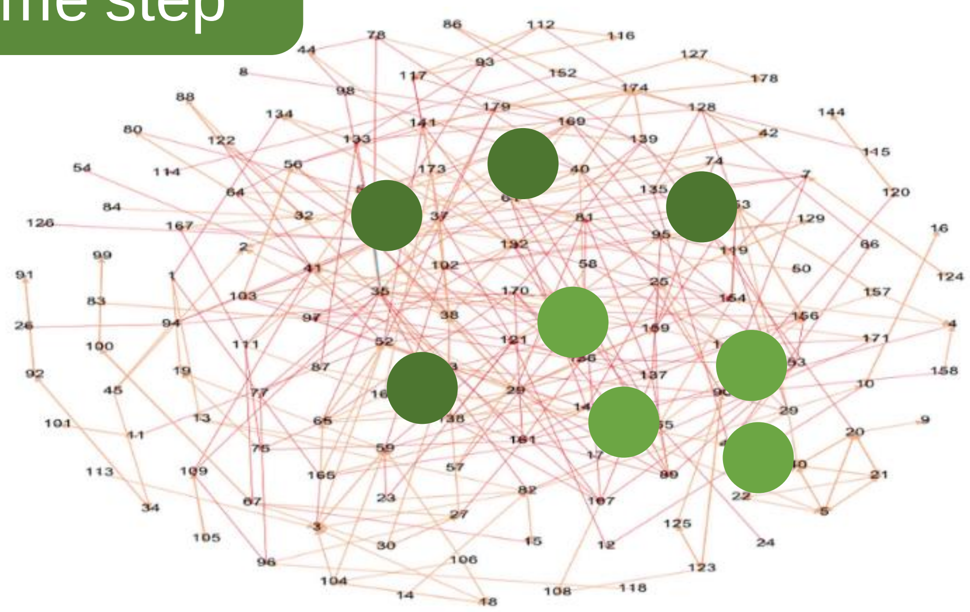

HEALER: POMDP Model for Multi-Step Policy Robust, Dynamic Influence Maximization


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

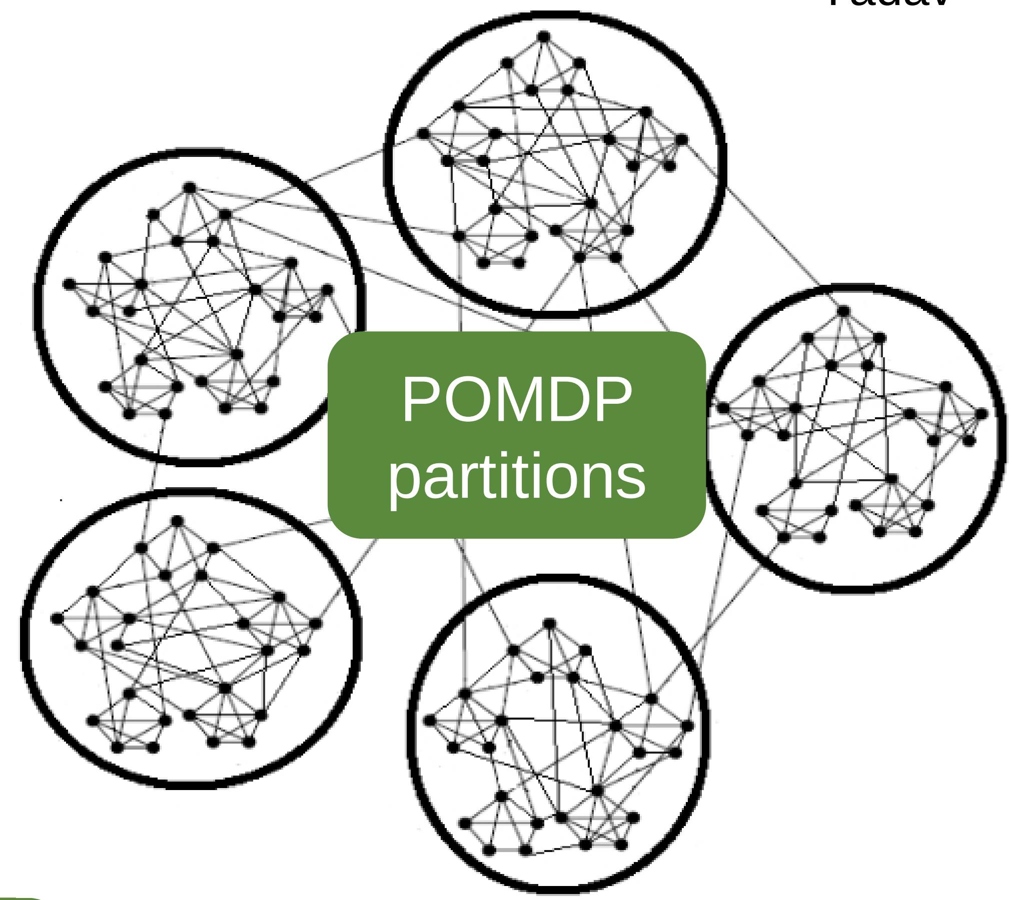

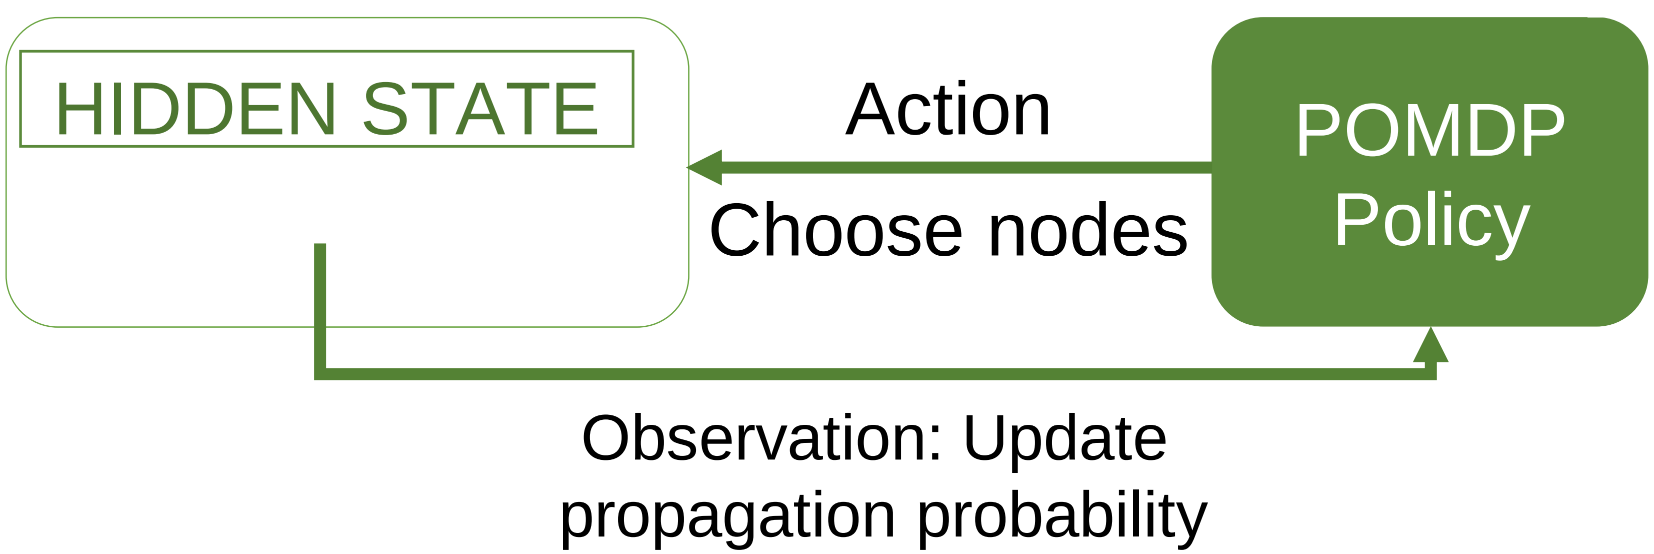

Pilot Tests with HEALER with 170 Homeless Youth [2017]


Yadav


Wilder

Recruited youths:

HEALER	HEALER++	DEGREE CENTRALITY
62	56	55

12 peer leaders

Results: Pilot Studies

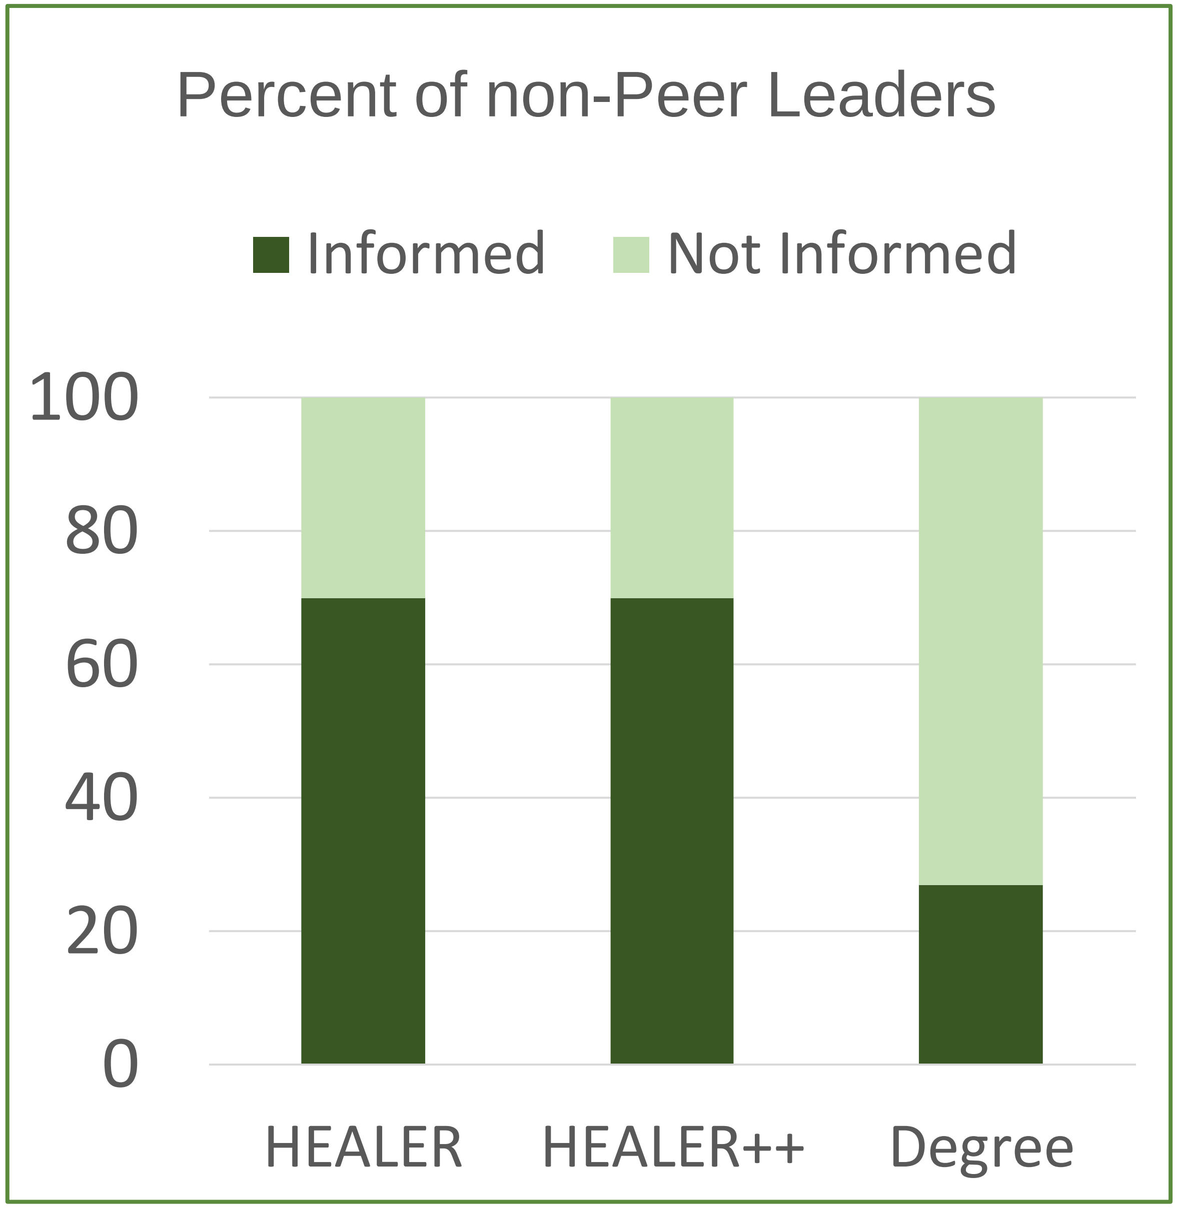

Data to Deployment Pipeline:
Network Sampling to avoid Data Collection Bottleneck

Wilder

Data collection costly

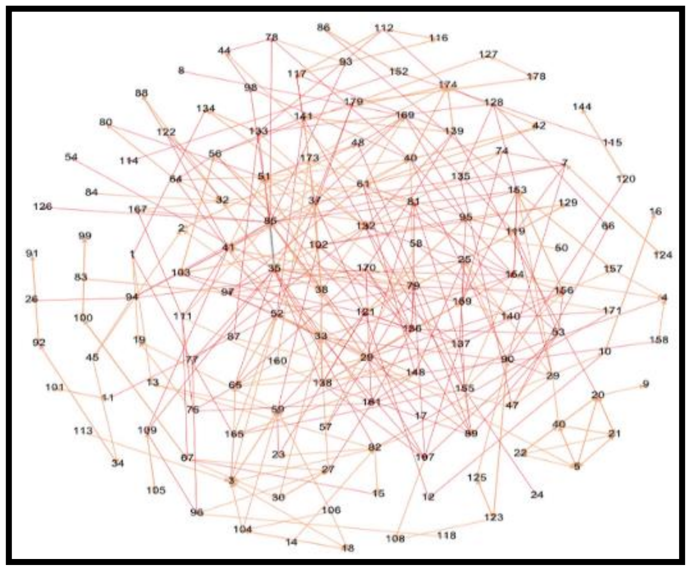

Sample 18%

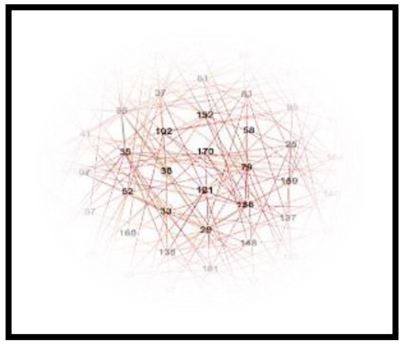

Sampling from largest communities

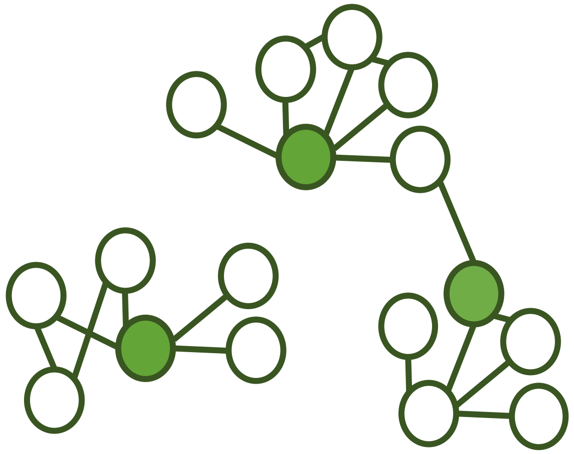

New experiment With 60 homeless youth

12 peer leaders

Results: Pilot Studies with New Sampling Algorithm [2018]


Wilder

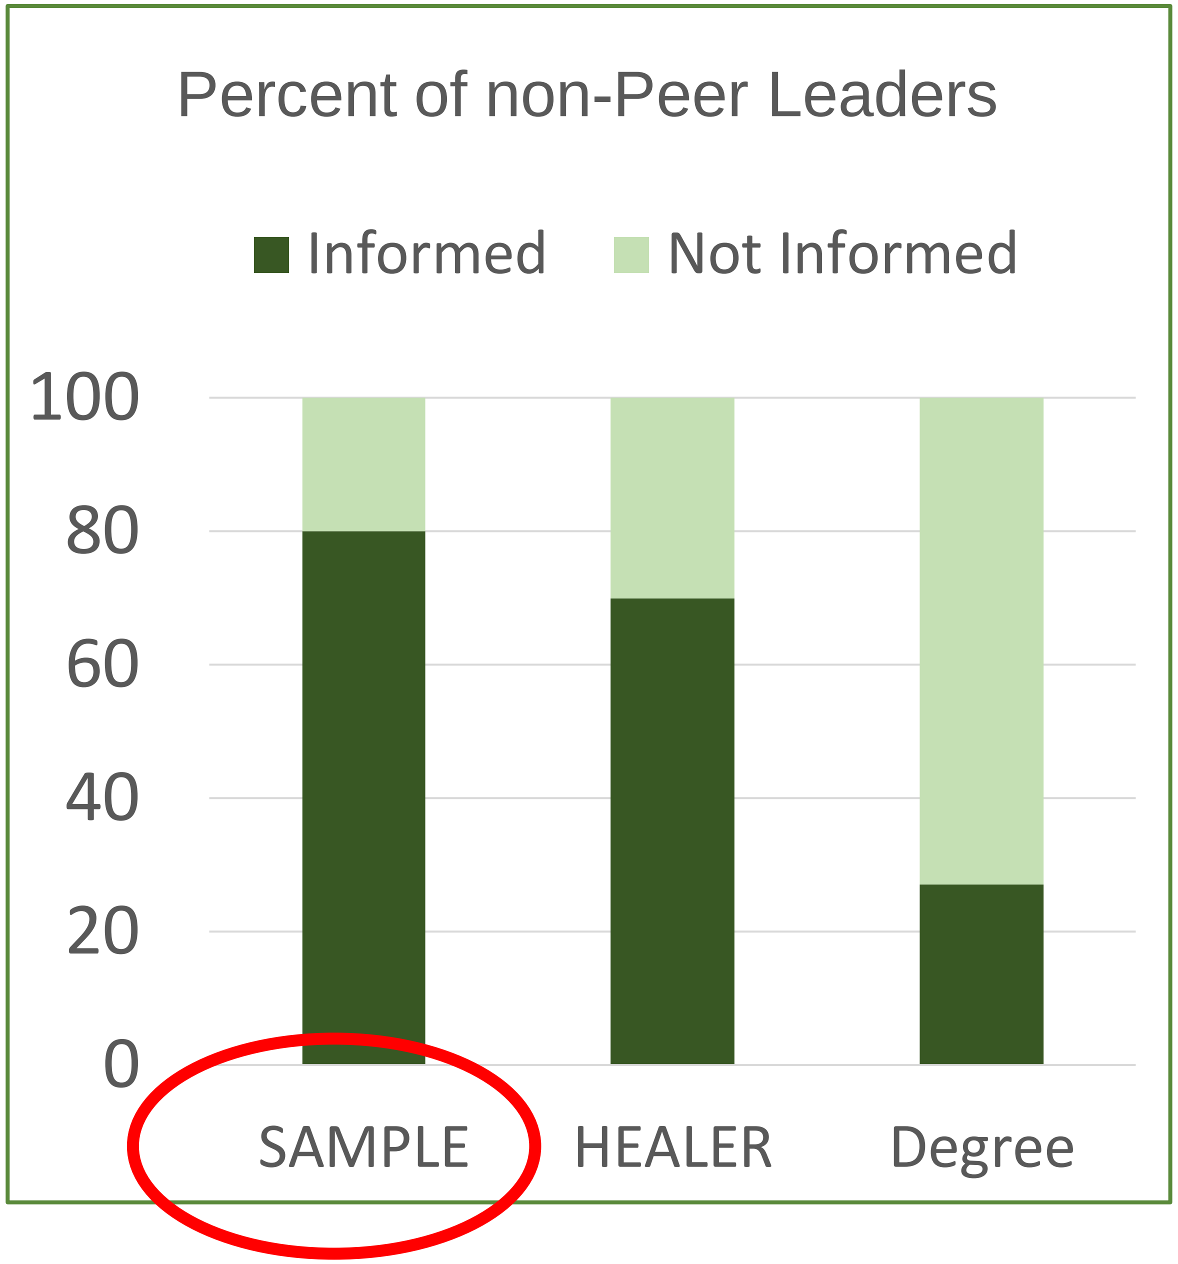

Results of 900 Youth Study [RECENT]

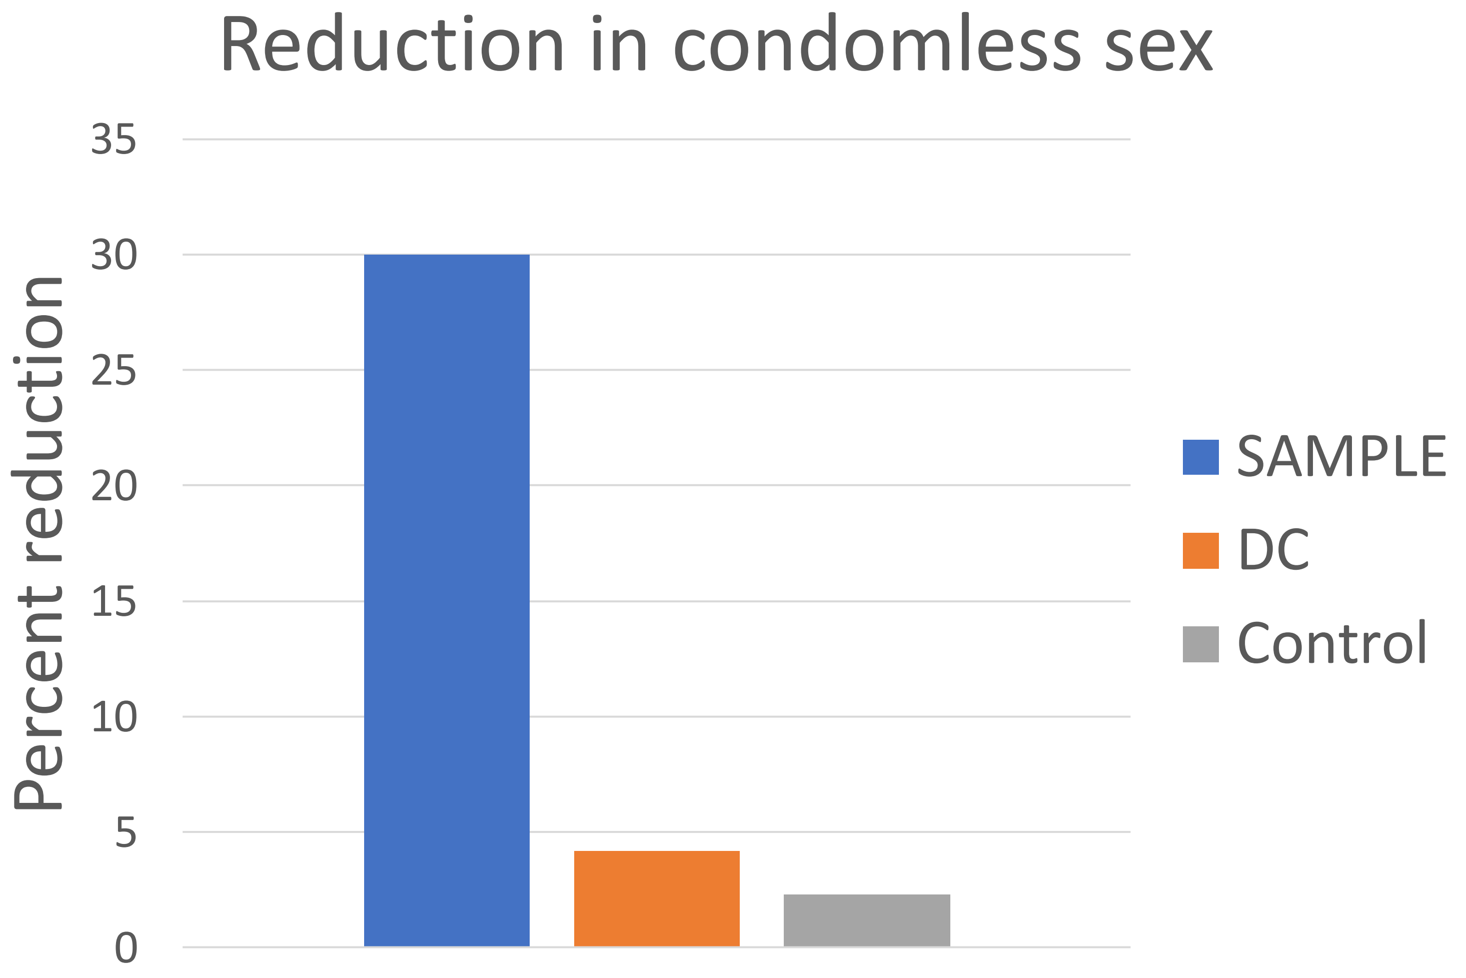

AI Assistant: HEALER

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

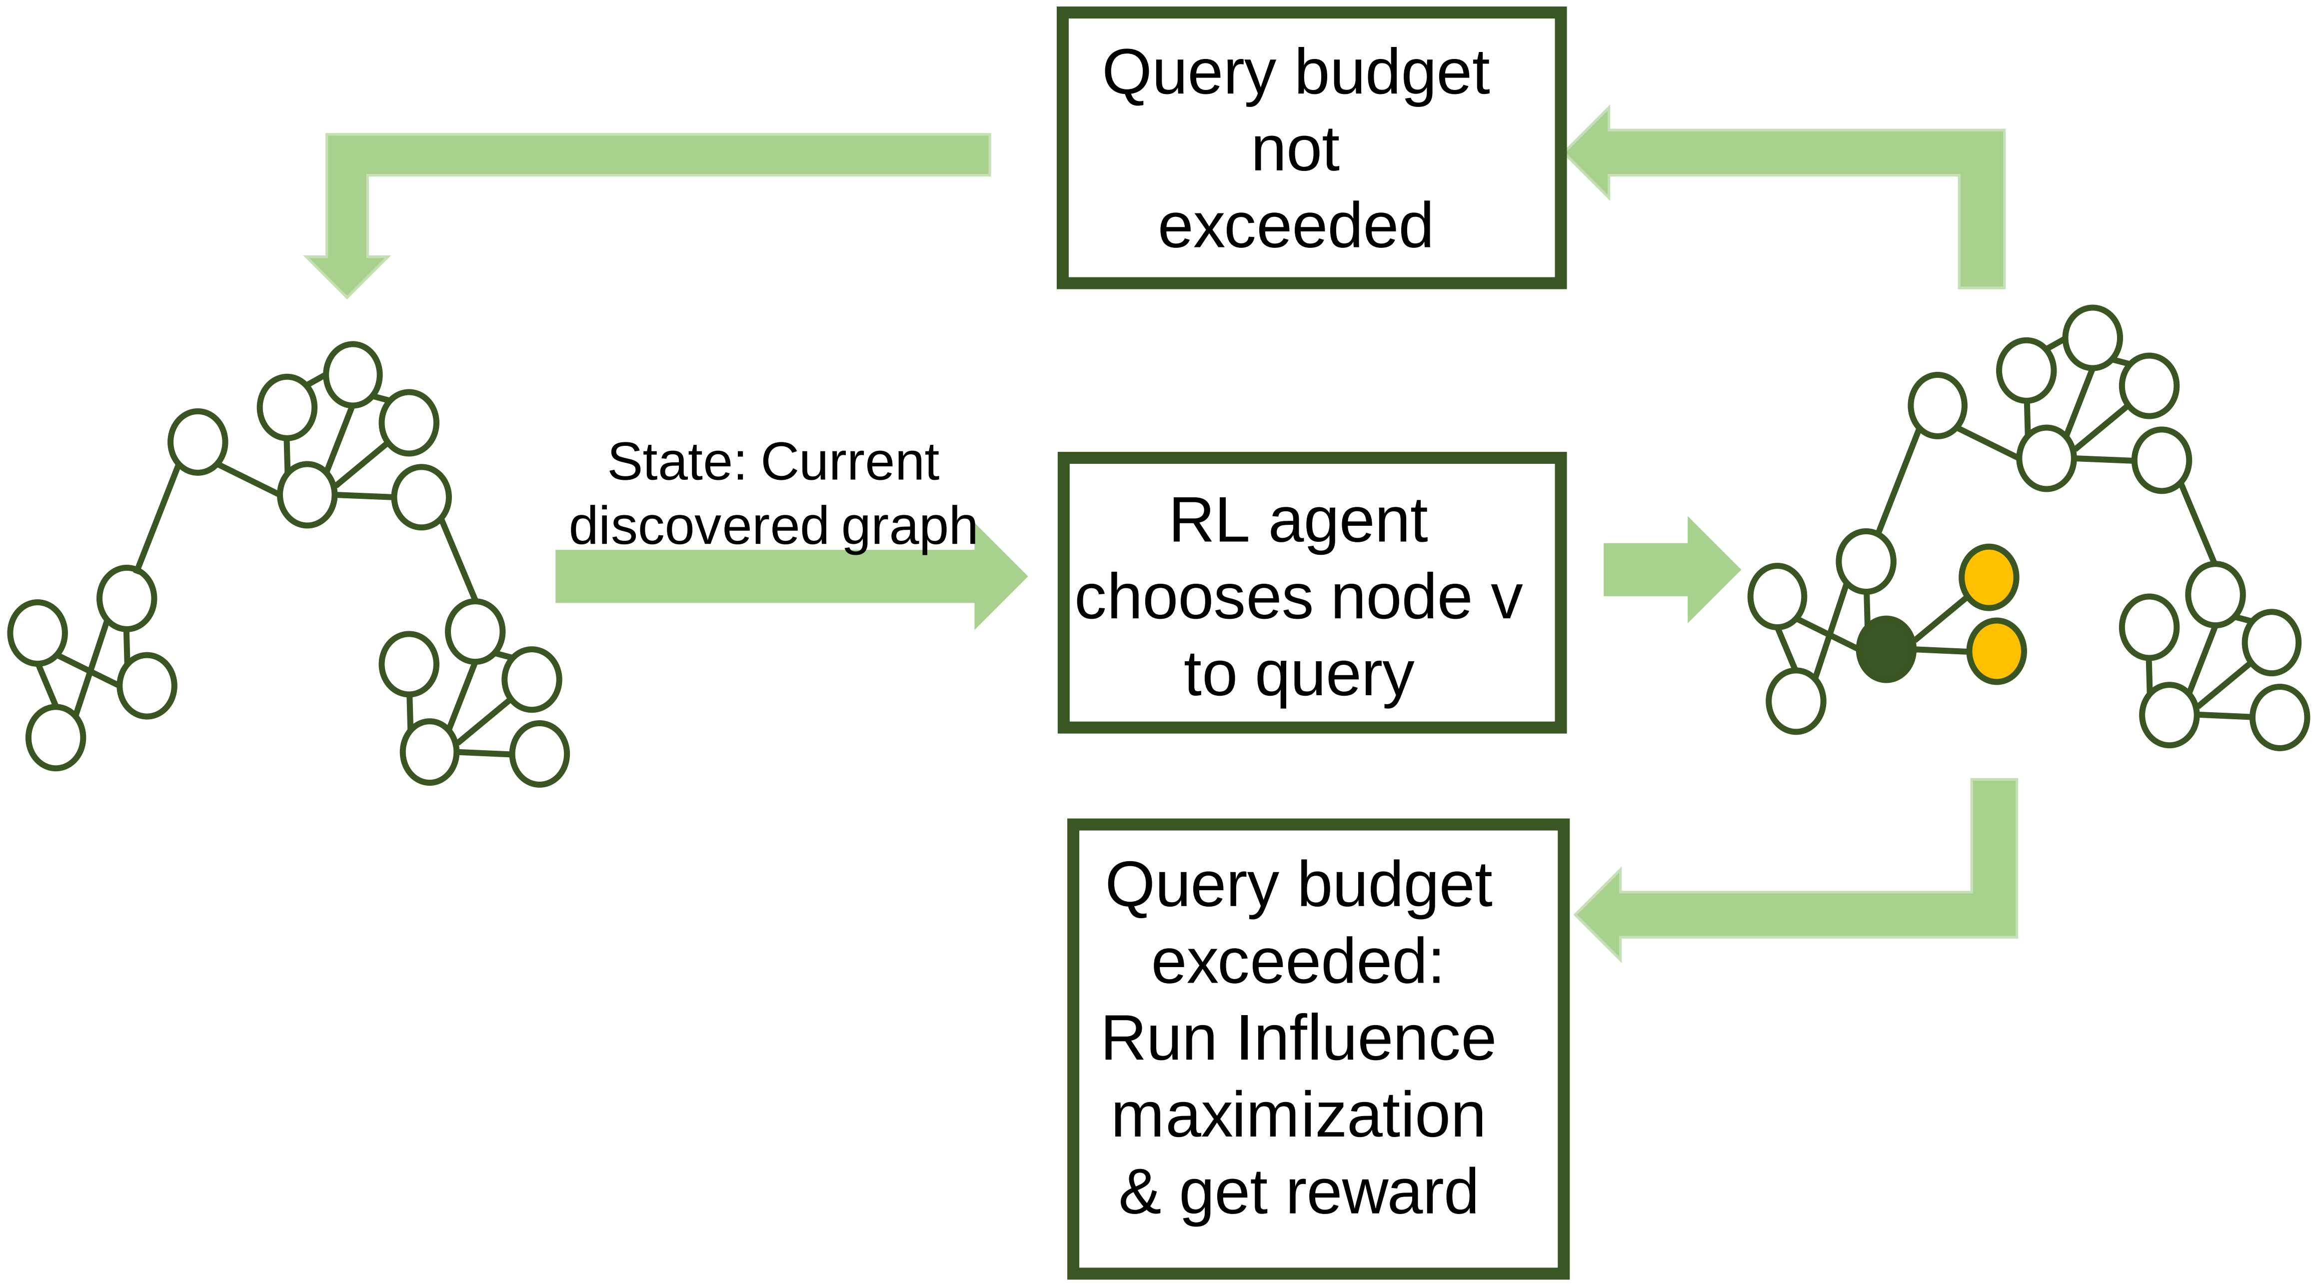

Data to Deployment Pipeline: Using an RL agent

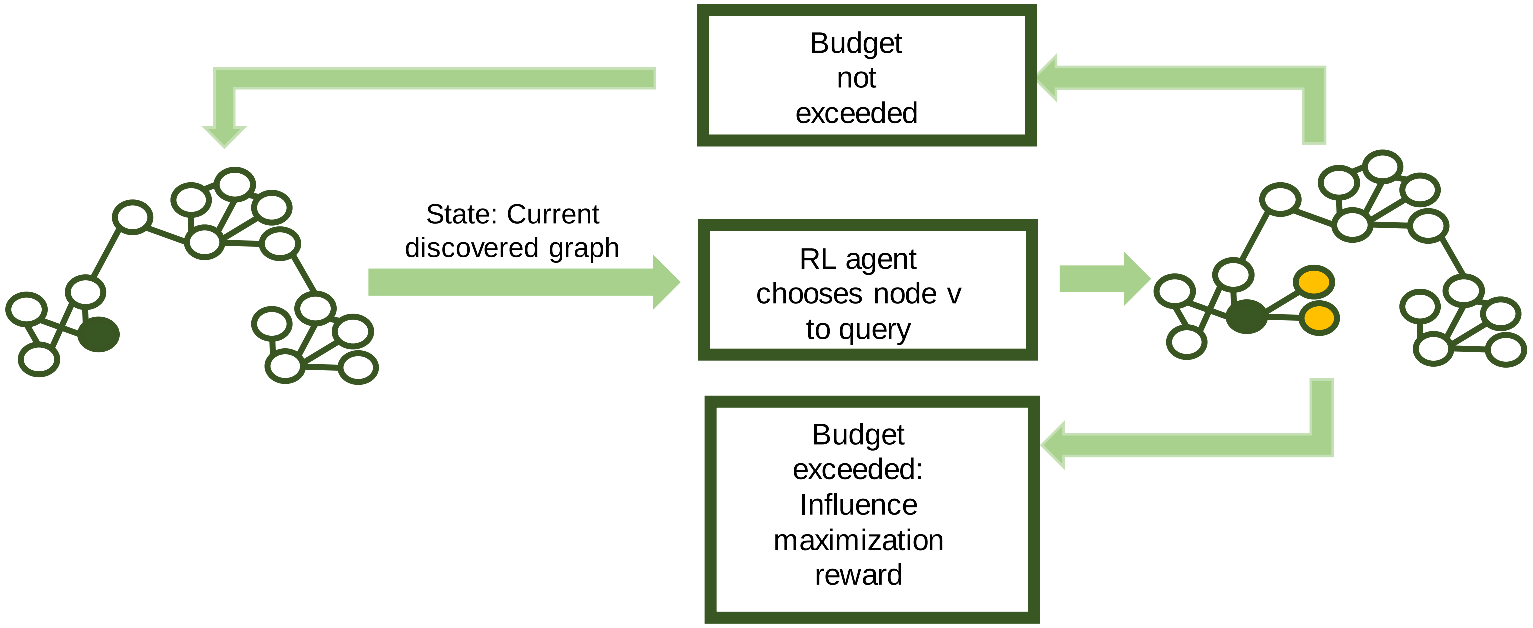

Network Family	Improve %
Rural	23.76
Animal	26.6
Retweet	19.7
Homeless	7.91

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

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

- *Non-adherence to TB Treatment*
- *Active case finding using social networks*

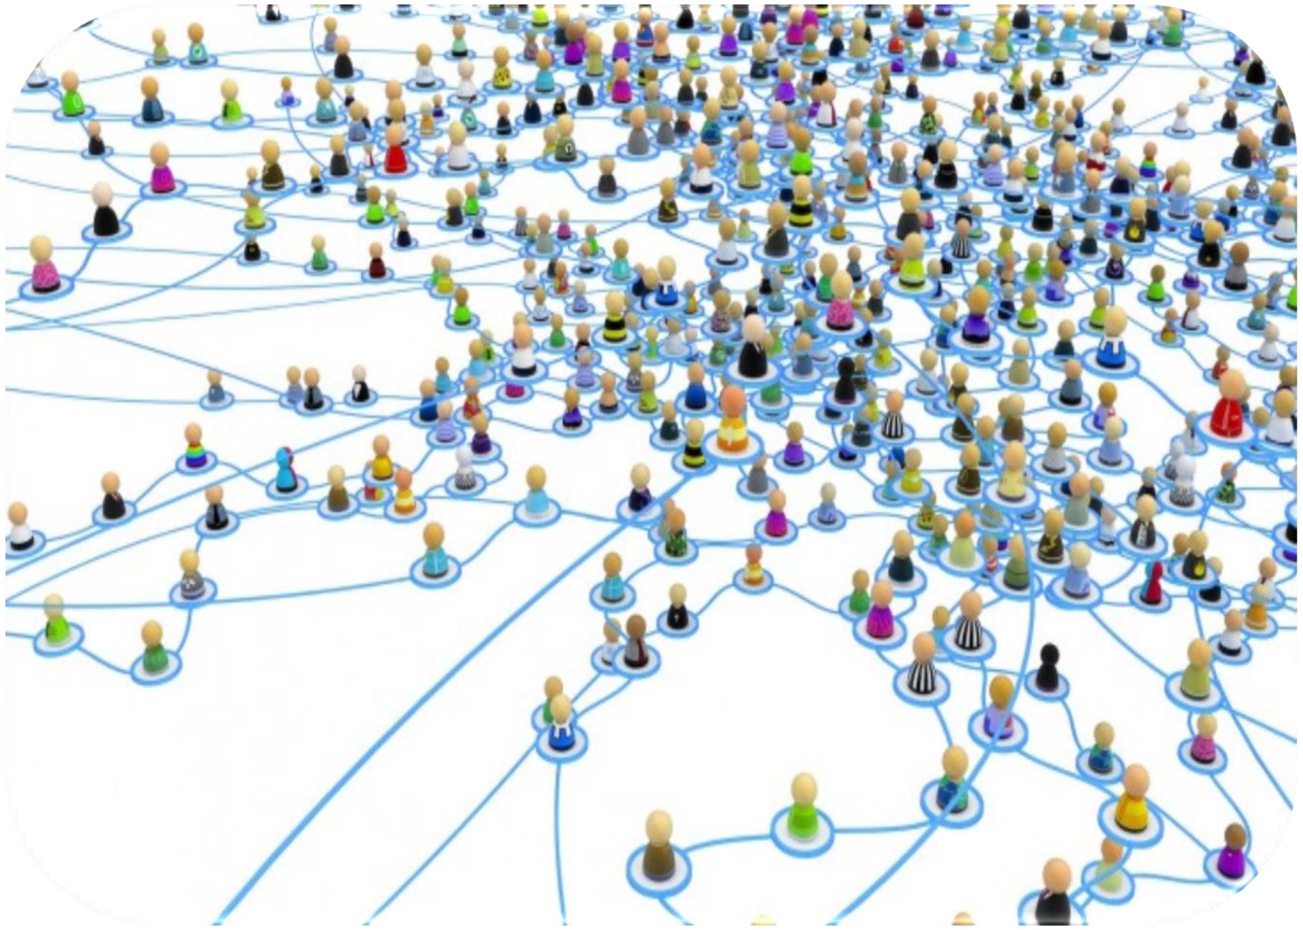

Non-Adherence to TB treatment
Preventing Tuberculosis in India [2019]

- *Digital adherence tracking: Patients call phone #s on pill packs; many countries*
- *Health workers track patients on a dash board*
- *Predict adherence risk from phone call patterns? Intervene before patients miss dose*

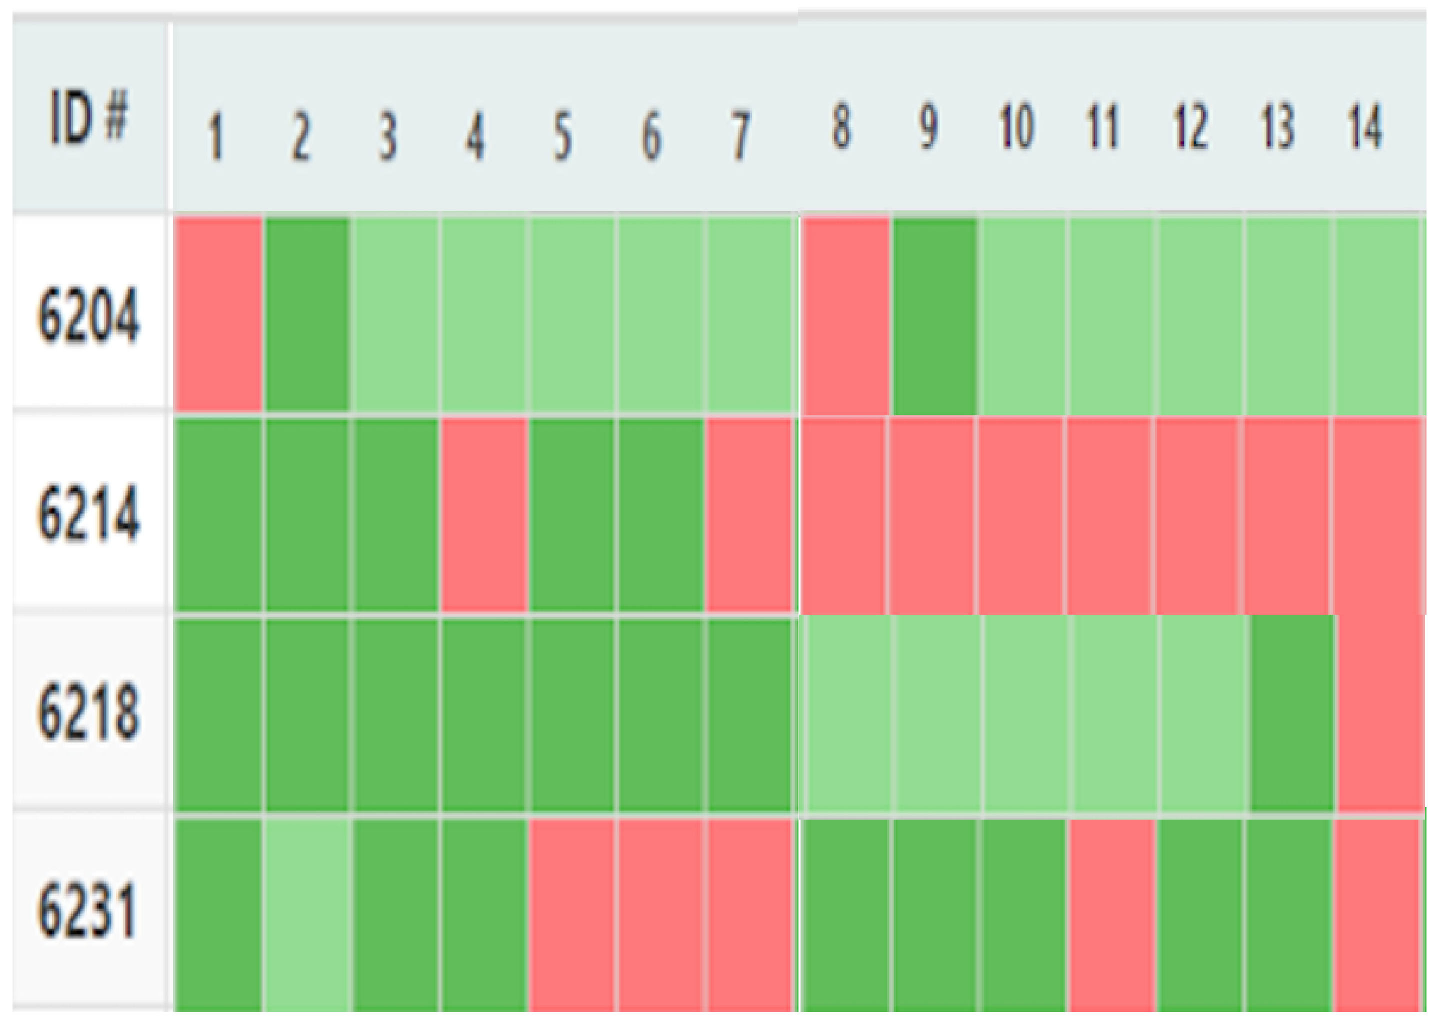

TB Treatment Adherence but Limited Resources: Intervening Selectively before patients miss doses

Killian

Data Collect

Phone logs

Predict high risk patients

RF or LSTM

Prescription

Constraint Top K

Field

- *15K patients, 1.5M calls*

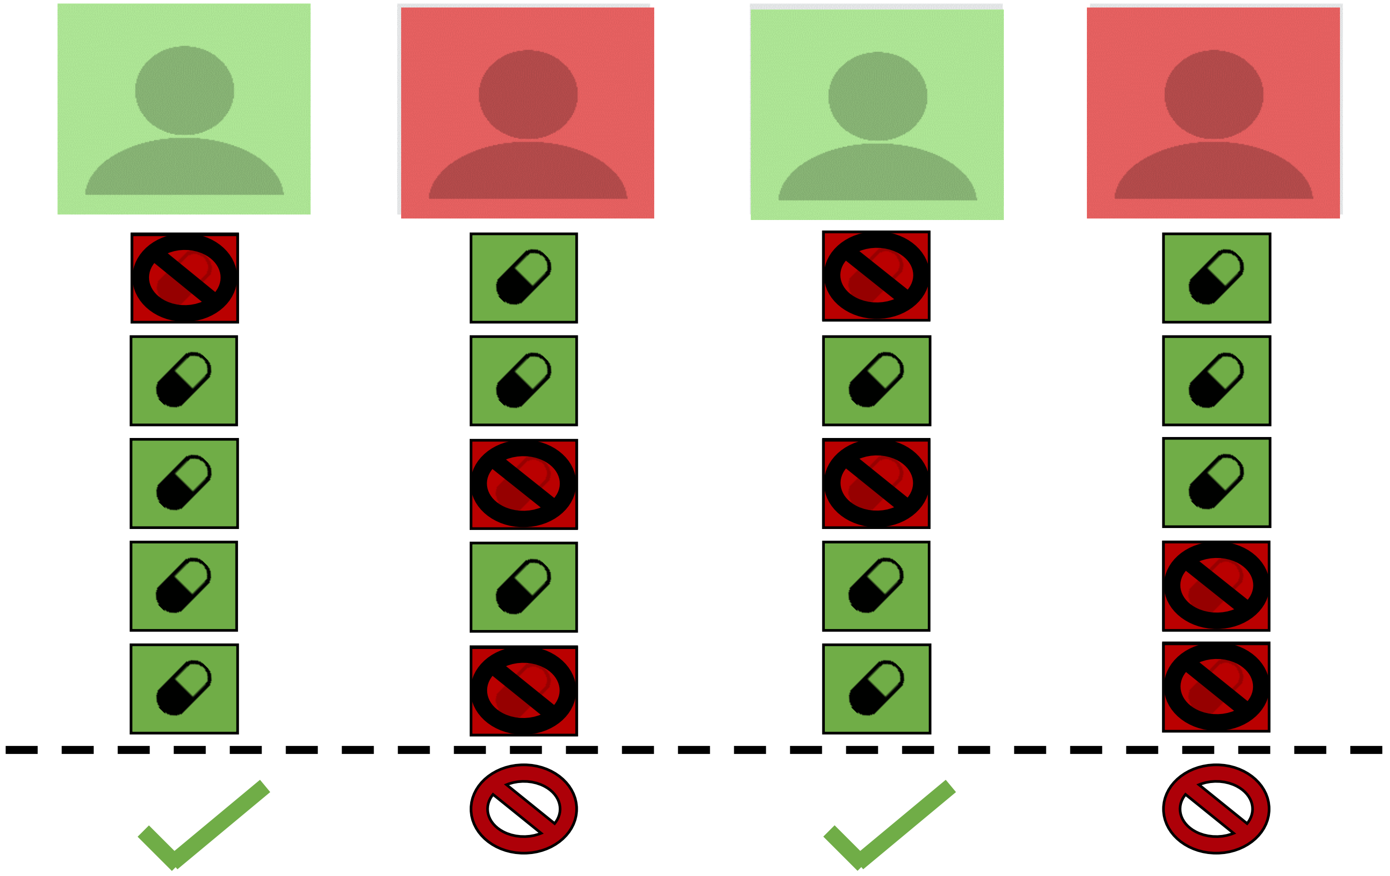

Increasing TB Treatment Adherence: Intervening before patients miss doses [2019]

Killian

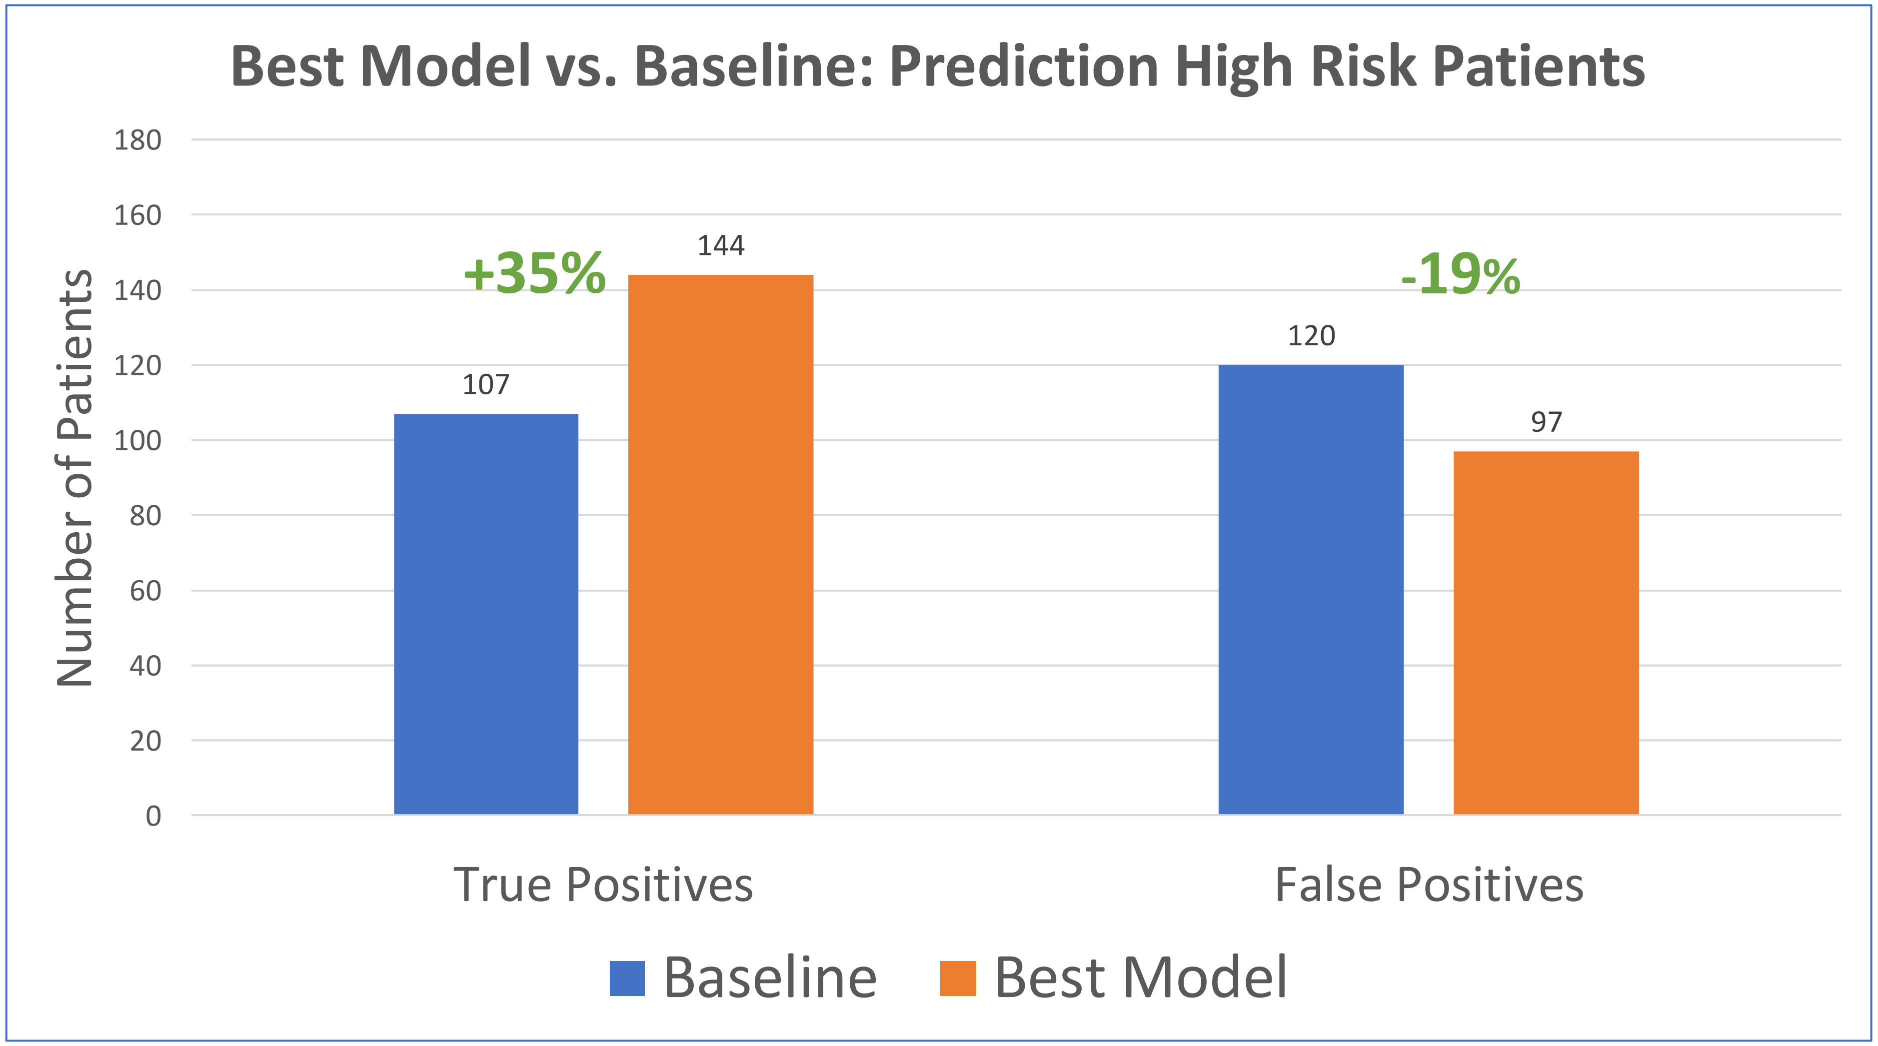

Data from

State of Maharashtra

India

Improving TB interventions
Decision-Focused vs Stage by Stage Methods

Wilder

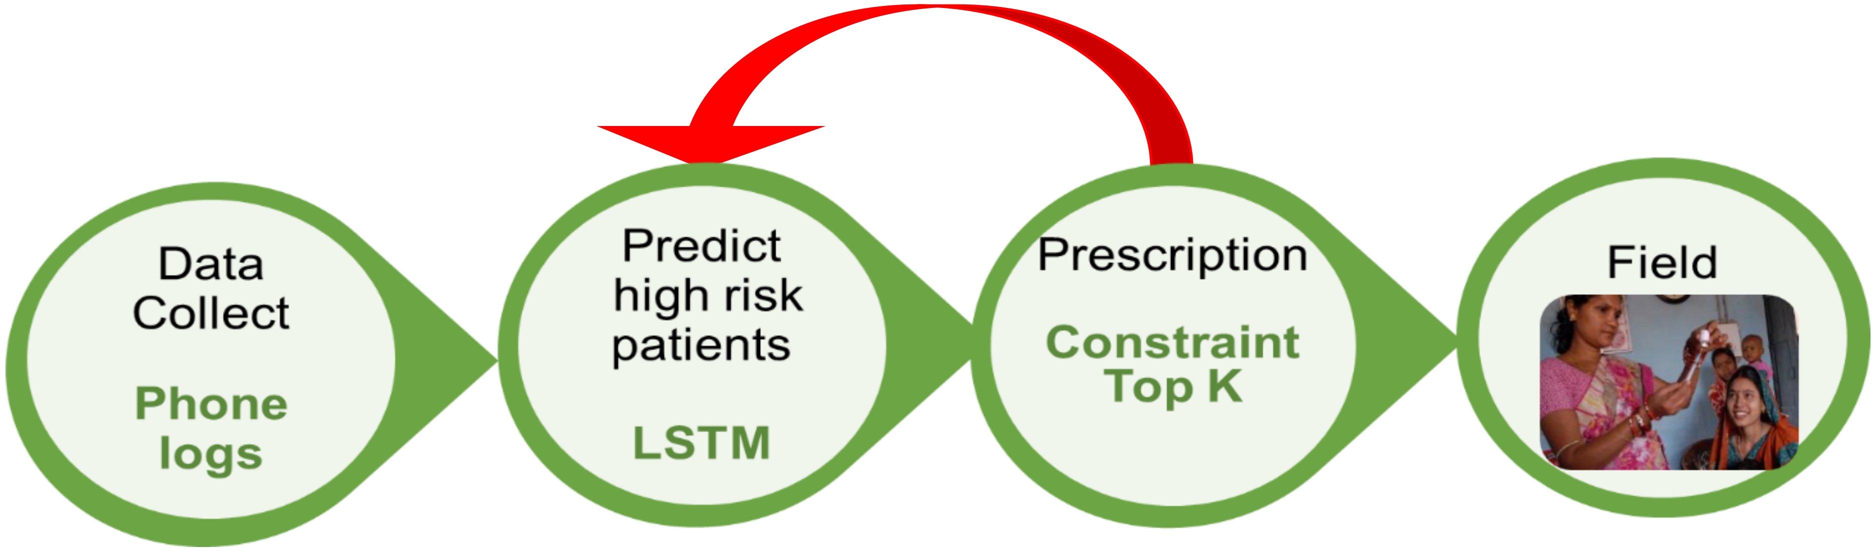

Decision focused learning improves TB interventions

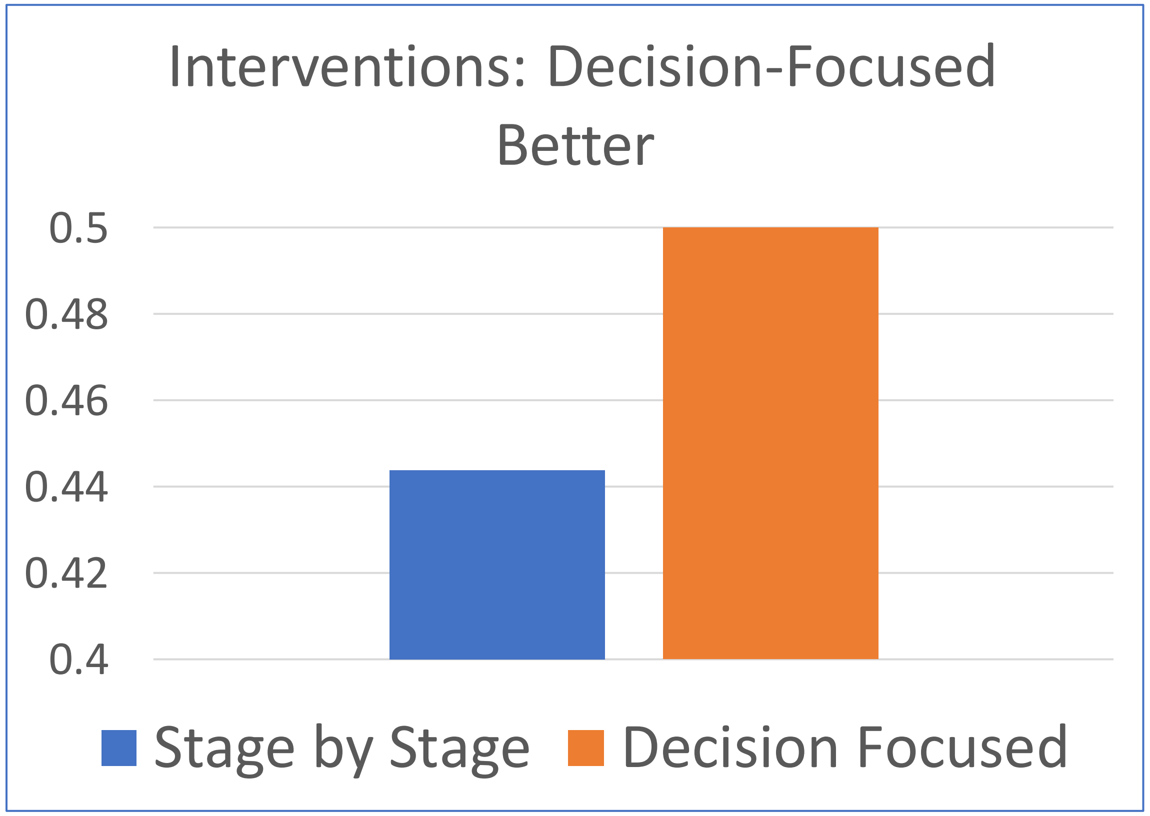

Integrating with Everwell's Platform


Killian

This work has a lot of potential to save lives.

Bill Thies

Co-founder, Everwell Health Solutions

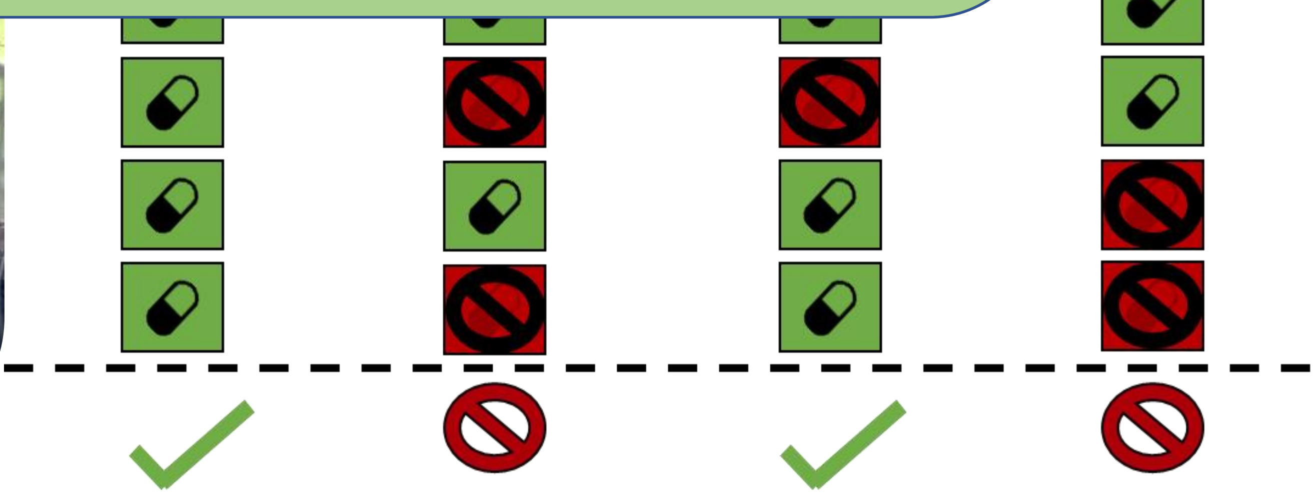

Active Case Finding via Network Optimization

Ou

Perrault

- *Active case finding: Find key nodes in contact network to cure*
- *Active screening: How to allocate limited resourced?*
- *Work with Wadhwani AI*

Wadhwani AI
AI for Social Good

Active Case Finding in India

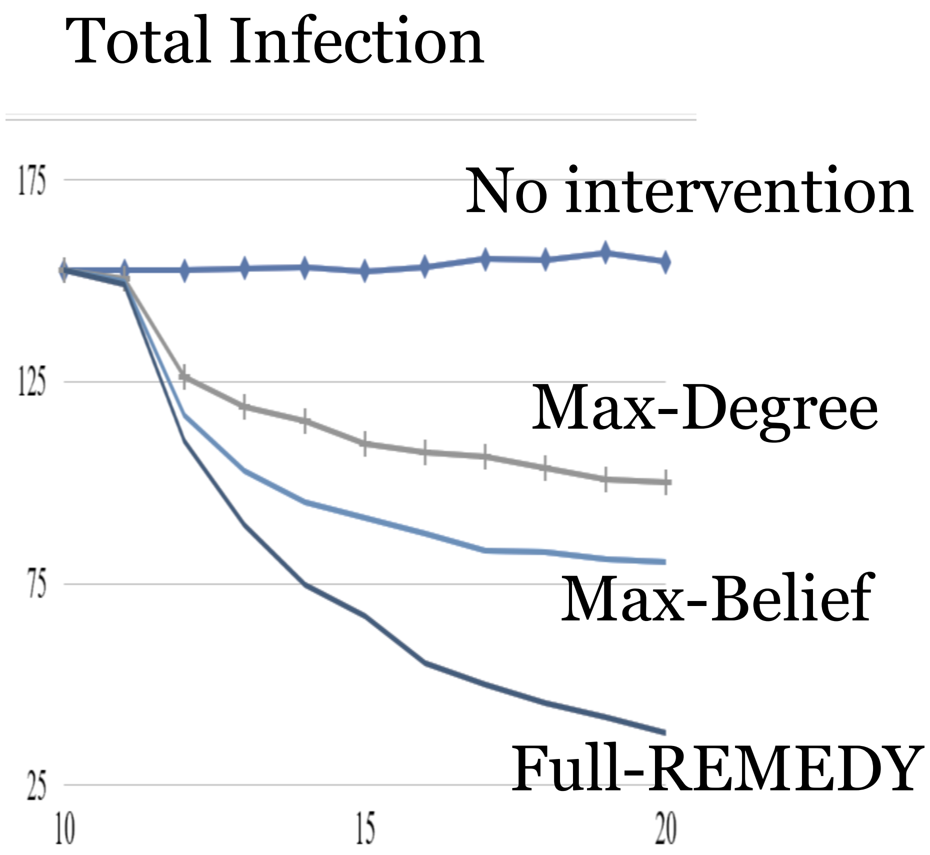

Suicide Prevention in Marginalized Populations: Choose Gatekeepers in social networks

Rahmattalabi

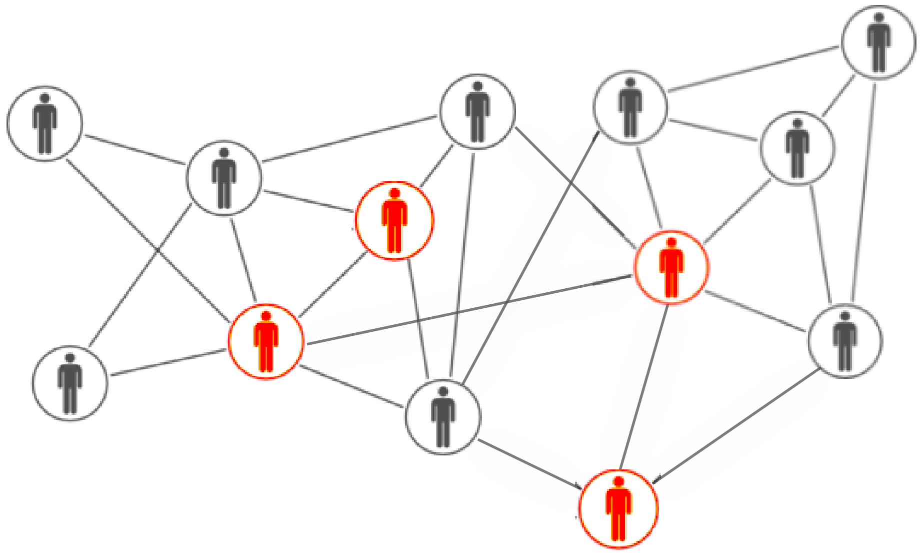

- Worst case parameters: a zero-sum game against nature
- Fairness of coverage

Algorithm
Chooses K gatekeepers

vs

Nature
Chooses some gatekeepers to not participate

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: AI Research for Social Good


It is possible to simultaneously advance AI research & do 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:
Reduces Errors on Important Targets

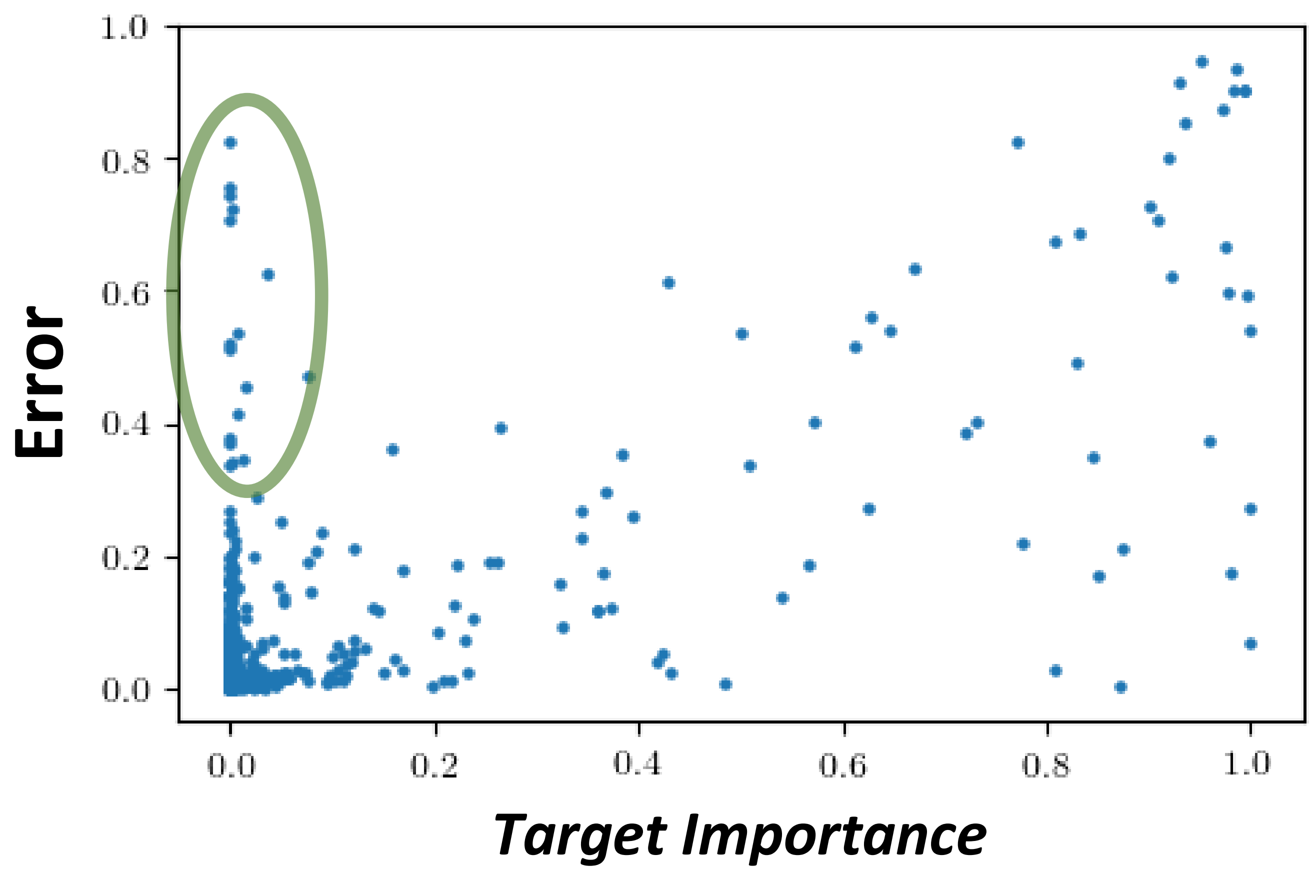

Game-focused

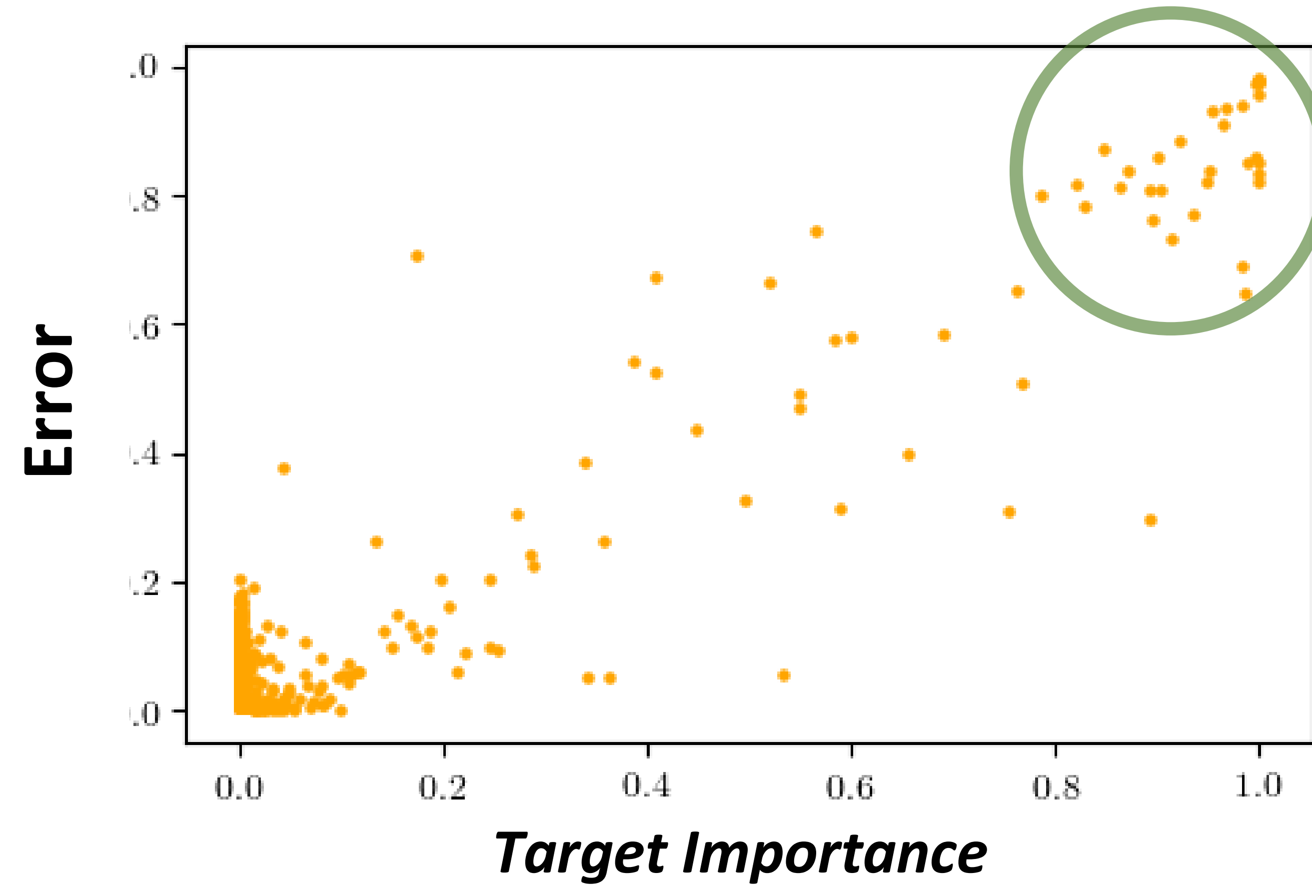

Two-stage

Field Evaluation of Schedule Quality

Improved Patrol Unpredictability & Coverage for Less Effort

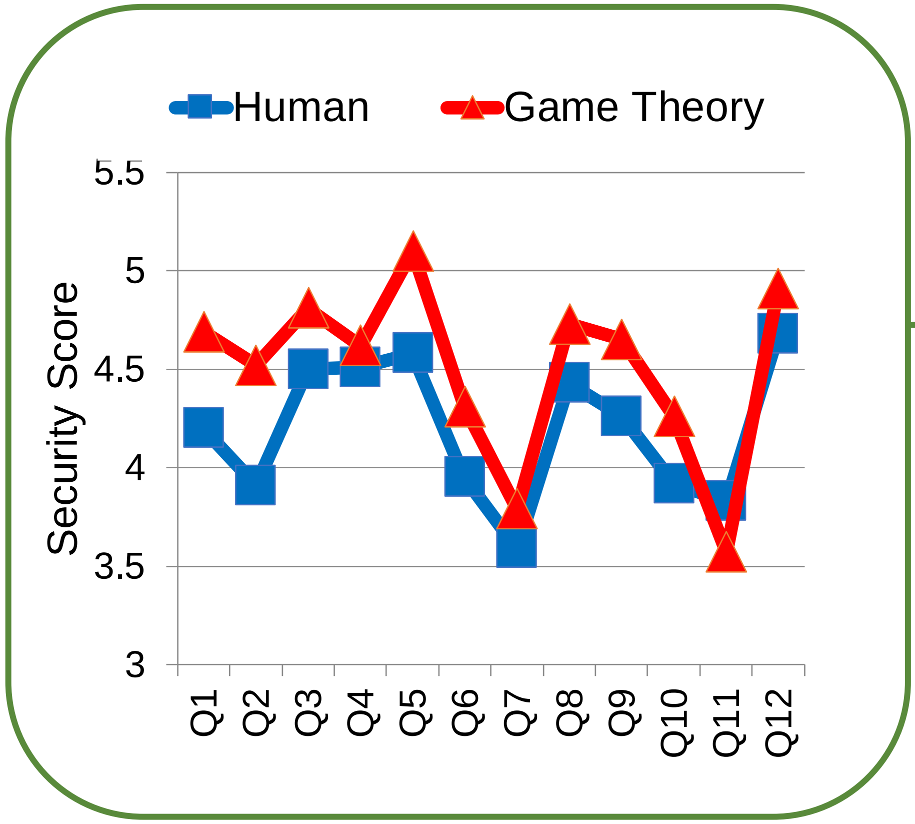

Train patrols: Game theory outperformed expert humans schedule 90 officers

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*


Childhood Obesity Prevention at homE (COPE)

Home Visitors Manual

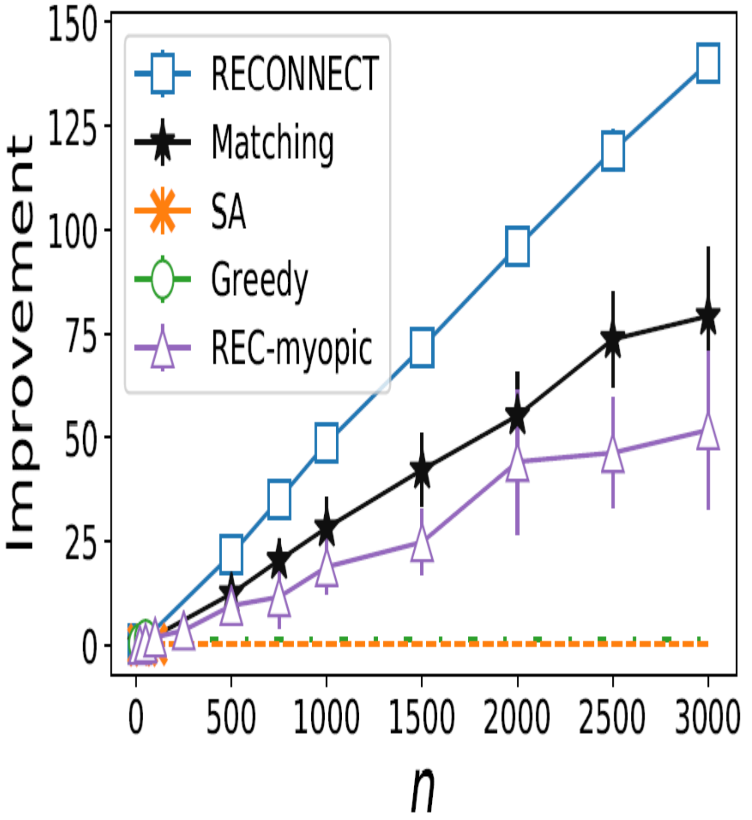

Solving Problems: Overall Research Framework
End-to-End, Data to Deployment Pipeline

Immersion

Data Collection

Predictive model

Learning/ Expert input

Prescriptive algorithm

Game theory Intervention

Field tests
&
deployment