AI for Social Good:
Learning & Planning in End-to-End, Data to Deployment Pipeline

MILIND TAMBE

Founding Co-director, Center for Artificial Intelligence in Society (CAIS)

University of Southern California

tambe@usc.edu

Co-Founder, Avata Intelligence

AI and Multiagent Systems Research for Social Good

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

Computational Game Theory

End-to-End Data-to-Deployment

Multiagent Systems
Optimizing Limited Intervention (Security) Resources

Public Safety and Security
Stackelberg Security Games






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

Multiagent Systems
Optimizing Limited Intervention (Ranger) Resources

Conservation/Wildlife Protection:
Green Security Games







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

Multiagent Systems
Optimizing Limited Intervention (Social Worker) Resources

Public Health:
Games against Nature


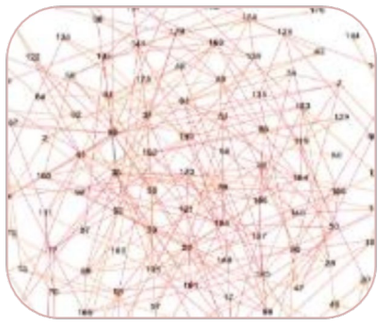

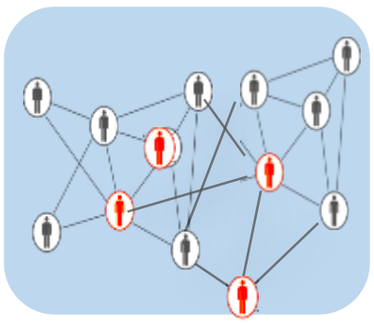


- Social networks to enhance intervention, e.g., HIV information
- Real-world pilot tests: Homeless youth shelters in Los Angeles

Solving Problems: Overall Research Framework
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

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

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

Outline: Overview of Past 10 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

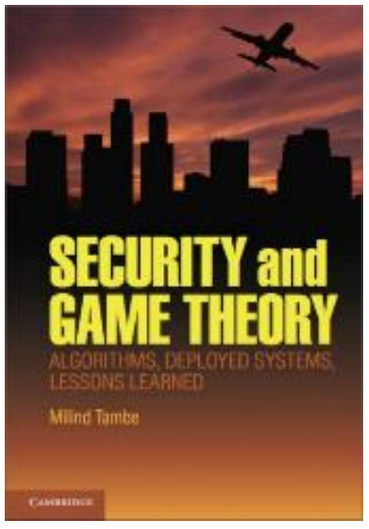

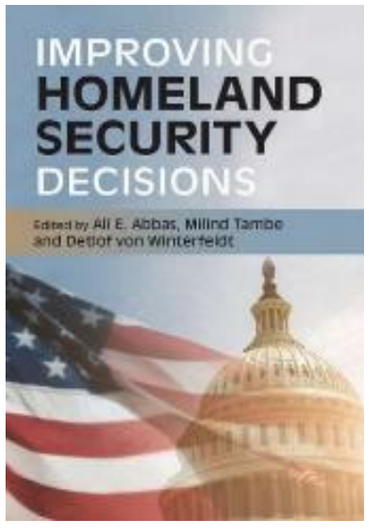

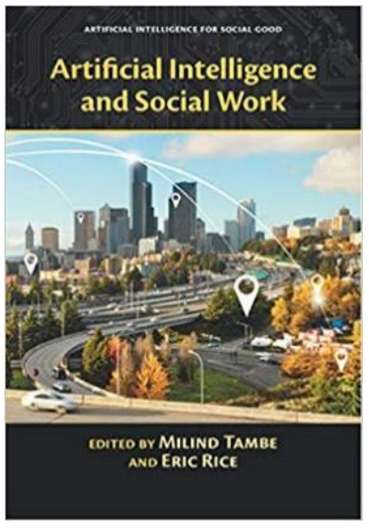

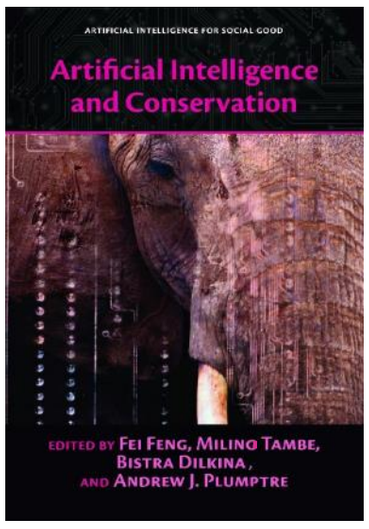

11 July 2006: Mumbai

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]


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….

Federal Air Marshals Service [2009]

Visiting Freedom Center: Home of Federal Air Marshals Service

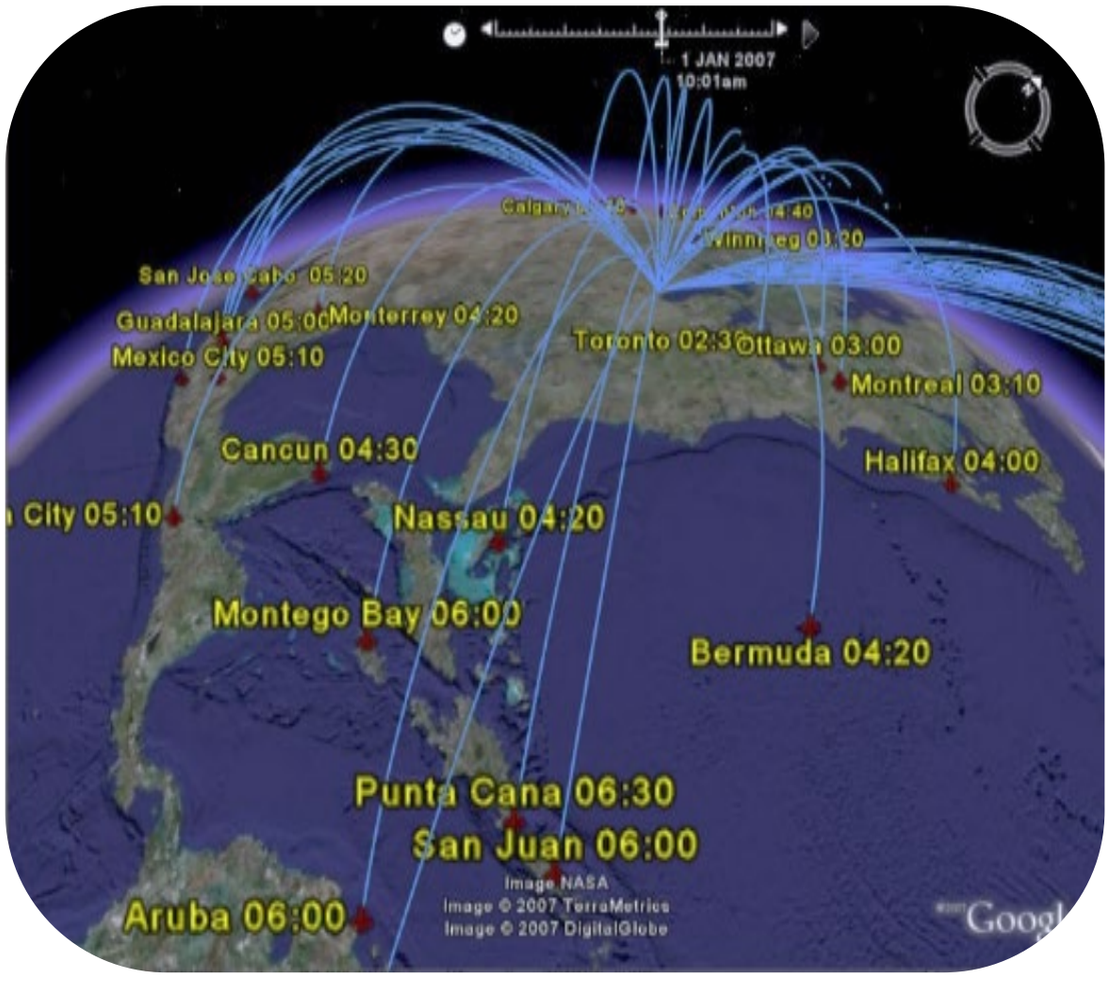

Scale Up Difficulty [2009]


Kiekintveld


Jain

x_i Defender mixed strategy

1000 flights, 20 air marshals:

10^{41} combinations

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

$$s.t. \sum_i x_i = 1, \sum_{j \in Q} q_j = 1$$

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

	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			

Scale Up [2009]
Exploiting Small Support Size


Kiekintveld


Jain

Theorem: For T targets, optimal solution of support set size T+1 always exists

Small support set size:
Most x_i variables zero

1000 flights, 20 air marshals:

10^{41} combinations

$X_{123} = 0.0$

$X_{124} = 0.239$

$X_{135} = 0.0$

$X_{378} = 0.123$

	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			

New Exact Algorithm for Scale up


Kiekintveld


Jain

Incremental strategy generation: First for Stackelberg Security Games

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

	Attack 1
1,2,4	5,-10
3,7,8	-8,10
…	…

Theory)
e strategy

GLOBAL OPTIMAL
1000 defender strategies
NOT 10^{41}

IRIS: Deployed FAMS [2009-]

Significant change in FAMS operations

Transportation Security Administration

Office of Law Enforcement/Federal Air Marshal Service

Milind Tambe

In recognition and appreciation of your outstanding achievement in developing the Intelligent Randomization In Scheduling (IRIS) program to advance the mission of the Office of Law Enforcement/Federal Air Marshal Service.

This 2nd day of September, 2011

James B. Curren
Studies, Research and Analysis
Office of Flight Operations

September 2011: Certificate of Appreciation (Federal Air Marshals)

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


Shieh


An

Boston

Los Angeles

New York

PROTECT: Ferry Protection Deployed [2013]

FERRIES: Mobile Resources & Moving Targets

Spatio-Temporal Security Games: Transition Graphs

Fang

Jiang

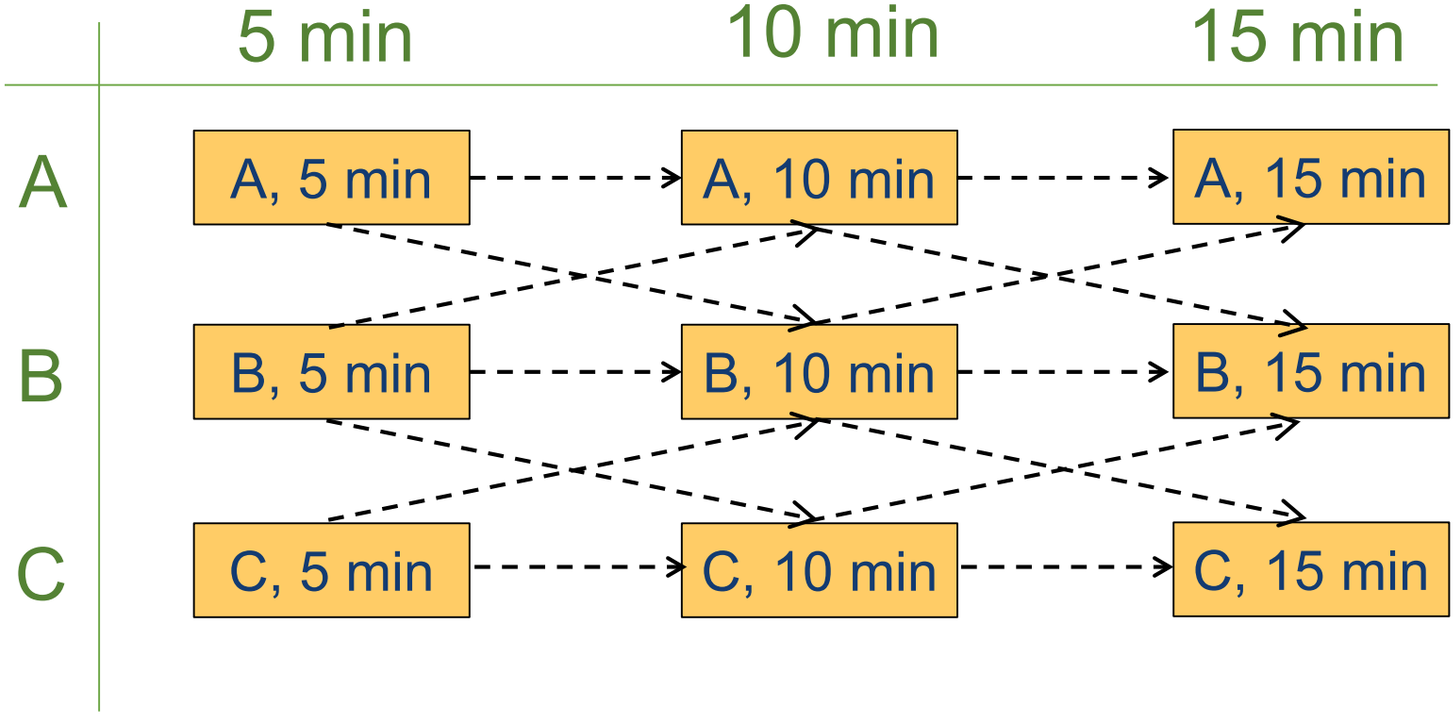

FERRIES: Mobile Resources & Moving Targets

Spatio-Temporal Security Games: Transition Graphs

Fang

Jiang

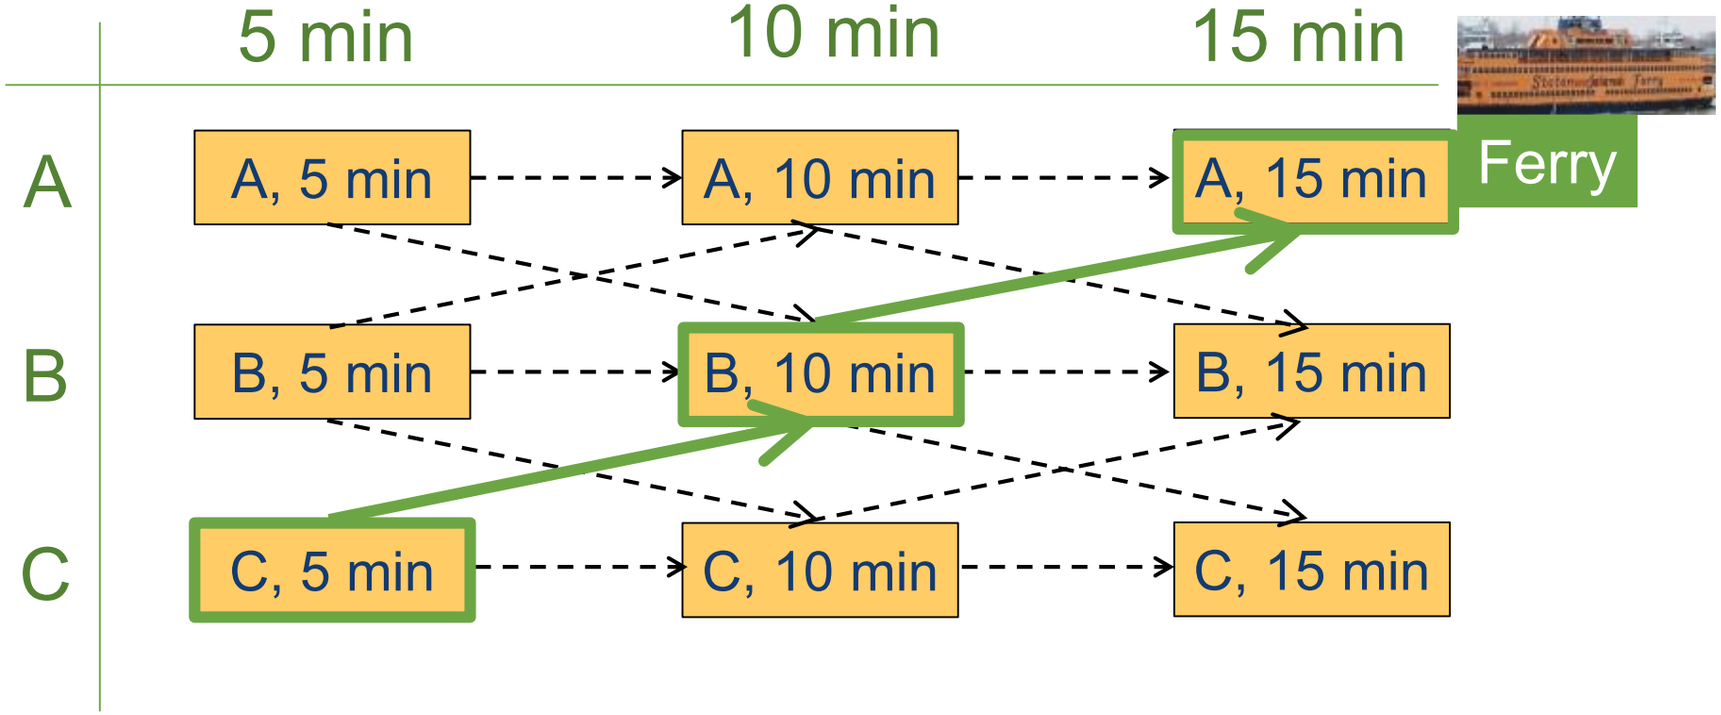

FERRIES: Mobile Resources & Moving Targets

Spatio-Temporal Security Games: Transition Graphs


Fang


Jiang

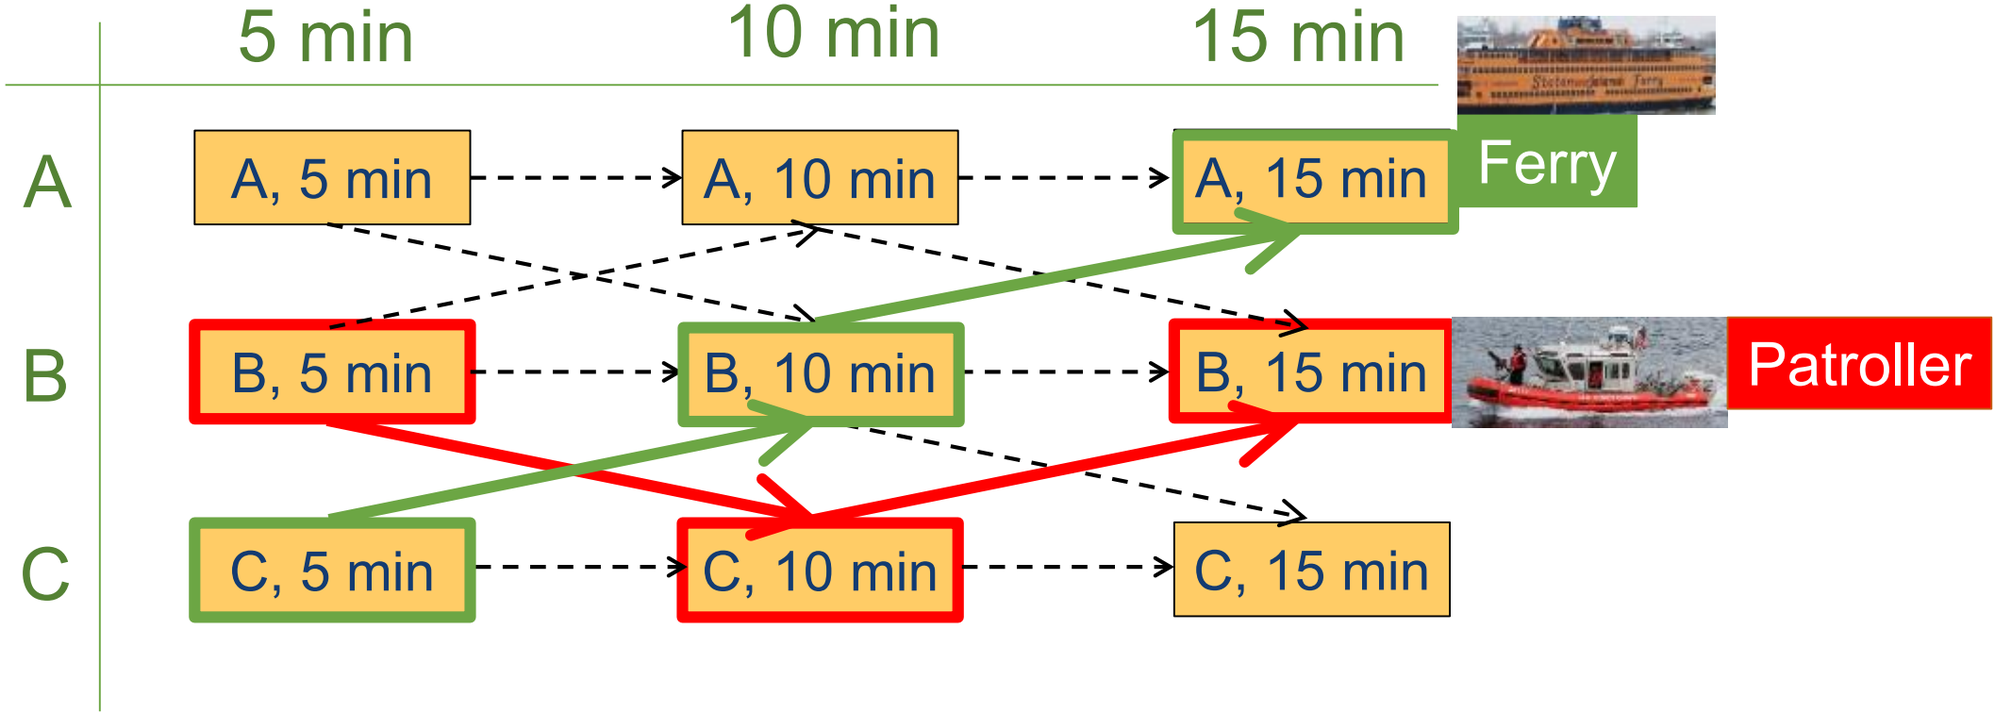

FERRIES: Scale up Difficulties


Fang


Jiang

*Theorem: **Marginals** enable scale-up with no solution quality loss*

Exponential N^T routes

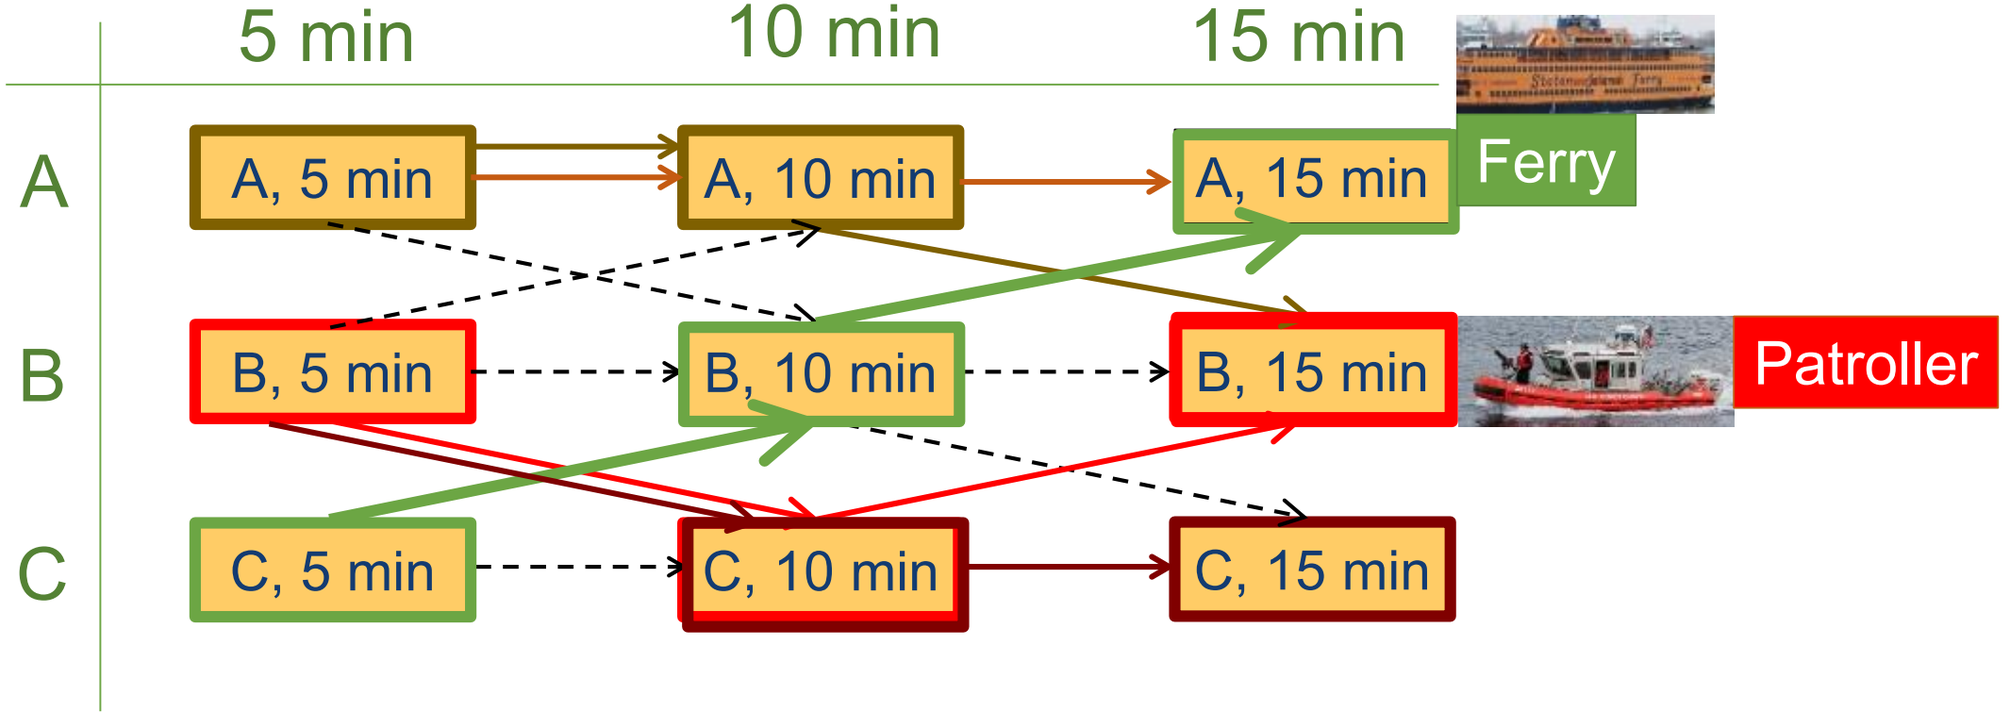

PROTECT: Port Protection Patrols [2013] Congressional Subcommittee Hearing

June 2013: Meritorious Team Commendation from Commandant (US Coast Guard)

July 2011: Operational Excellence Award (US Coast Guard, Boston)

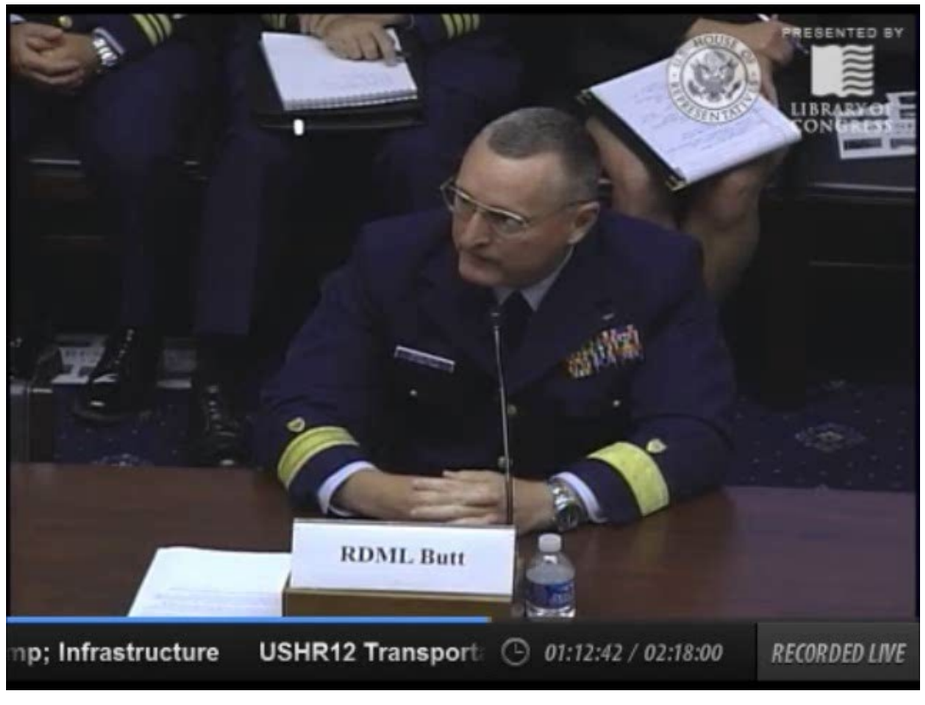

Solving Problems: Overall Research Framework
End-to-End, Data to Deployment Pipeline

Immersion

Data Collection

Predictive model

Learning/ Expert input

Prescriptive algorithm

Game theory Intervention

Field tests & deployment

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

Patrols Before PROTECT: Boston

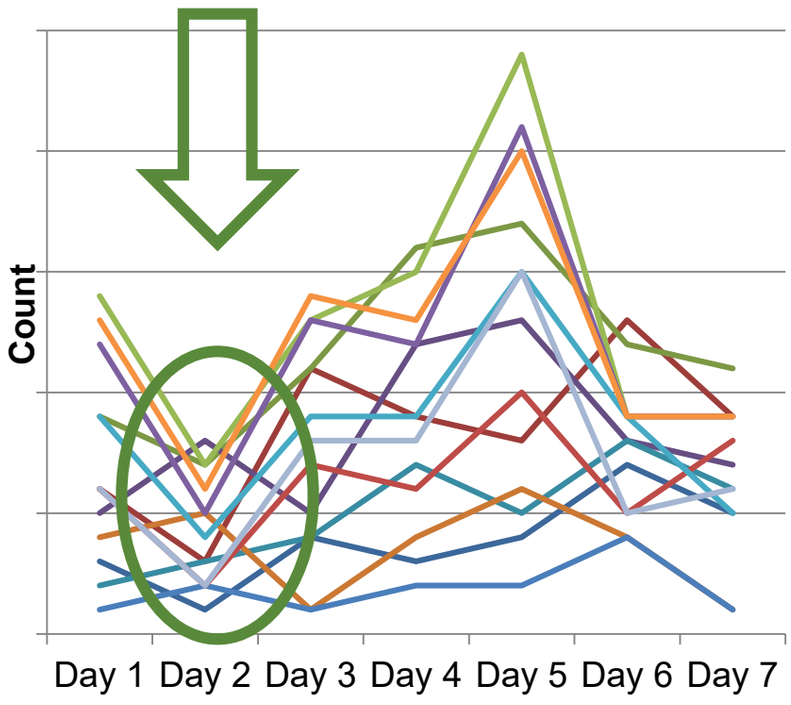

Patrols After PROTECT: Boston

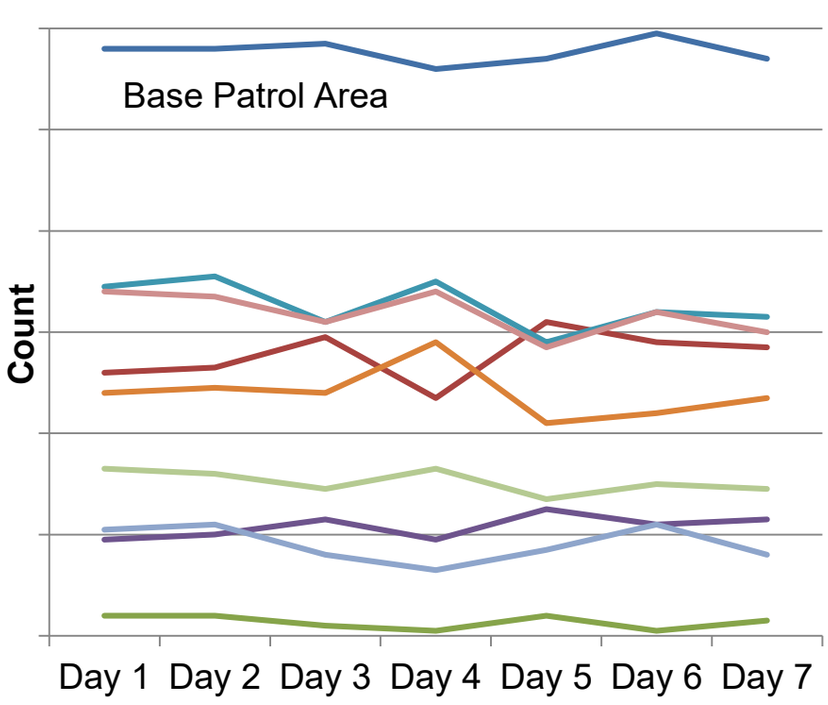

350% increase in defender expected utility

Field Evaluation of Schedule Quality

Improved Patrol Unpredictability & Coverage for Less Effort

FAMS: IRIS Outperformed expert human over six months

Report:GAO-09-903T

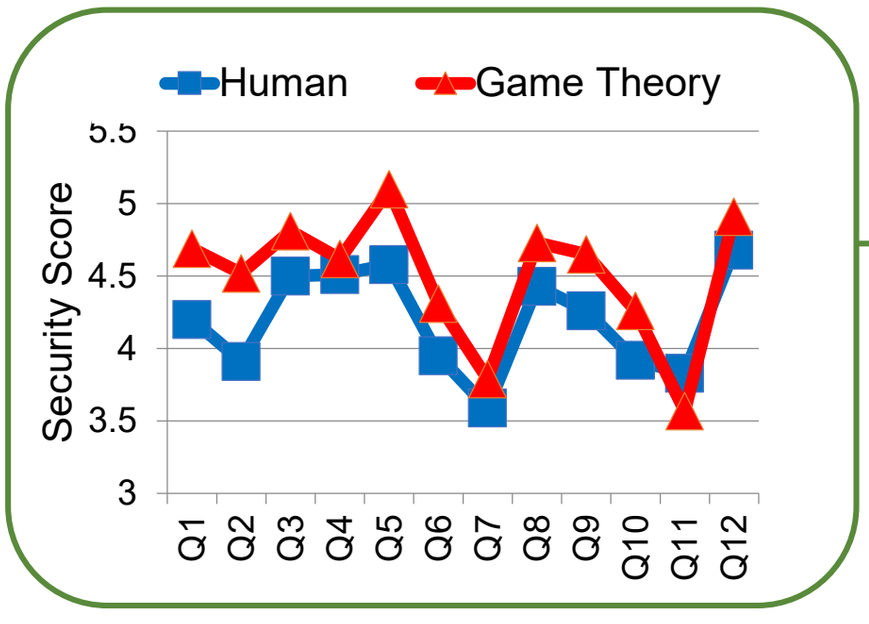

Train patrols: Game theory outperformed expert humans schedule 90 officers

Field Tests Against Adversaries

Computational Game Theory in the Field

Controlled

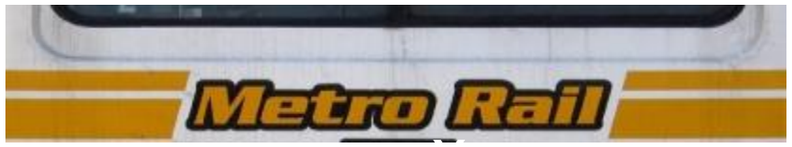

- 21 days of patrol, identical conditions
- Game theory vs Baseline+Expert

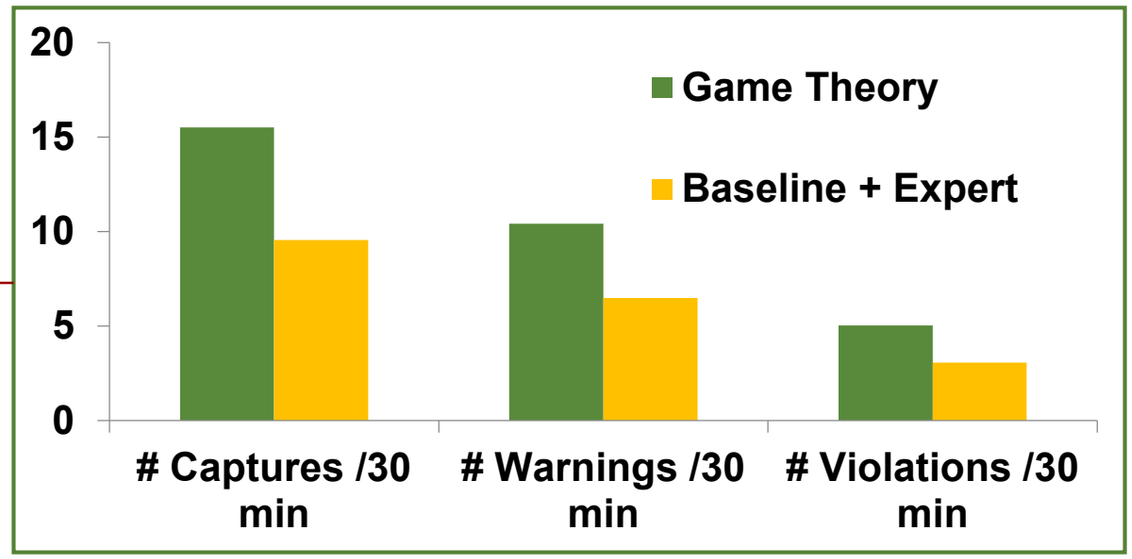

Not Controlled

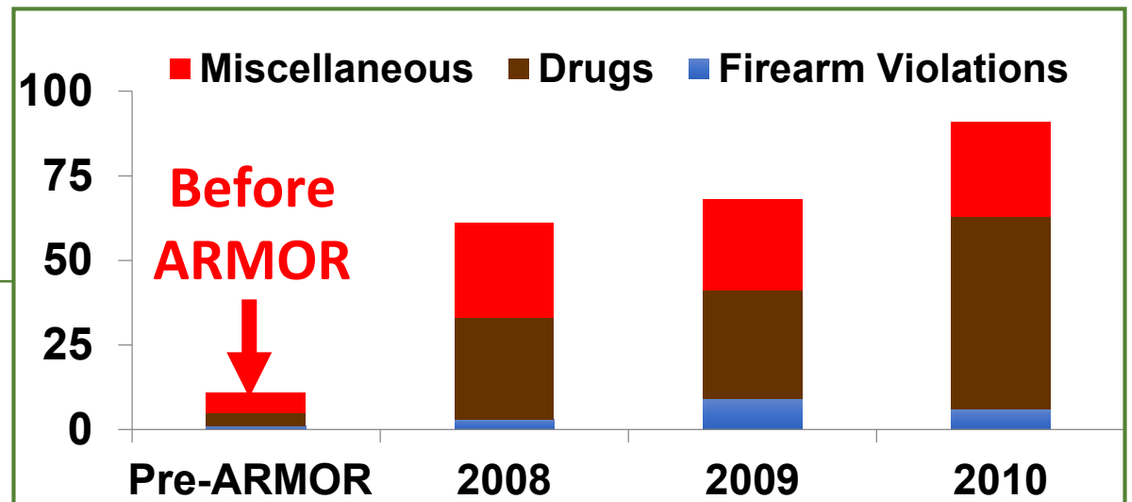

New Directions in Stackelberg Security Games [2018]


McCarthy


Schlenker


Sinha

- Threat Screening Games
  (AAAI16, IJCAI17, IJCAI18…)

- Cyber Security Games
  (IJCAI17, AAMAS18, CogSci18…)

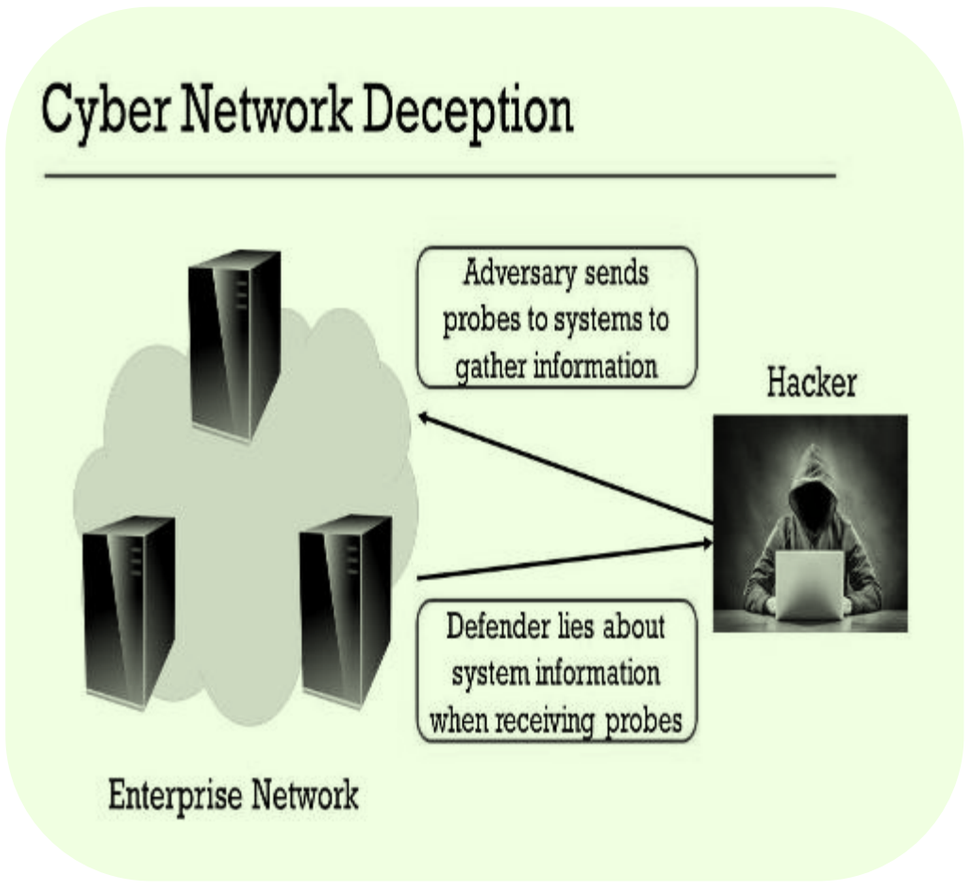

Outline

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

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

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

Xu

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

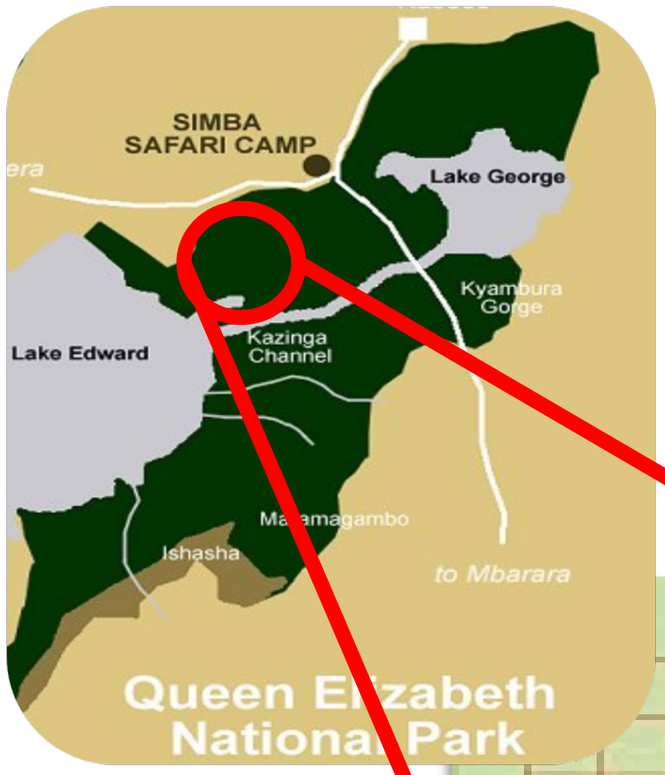

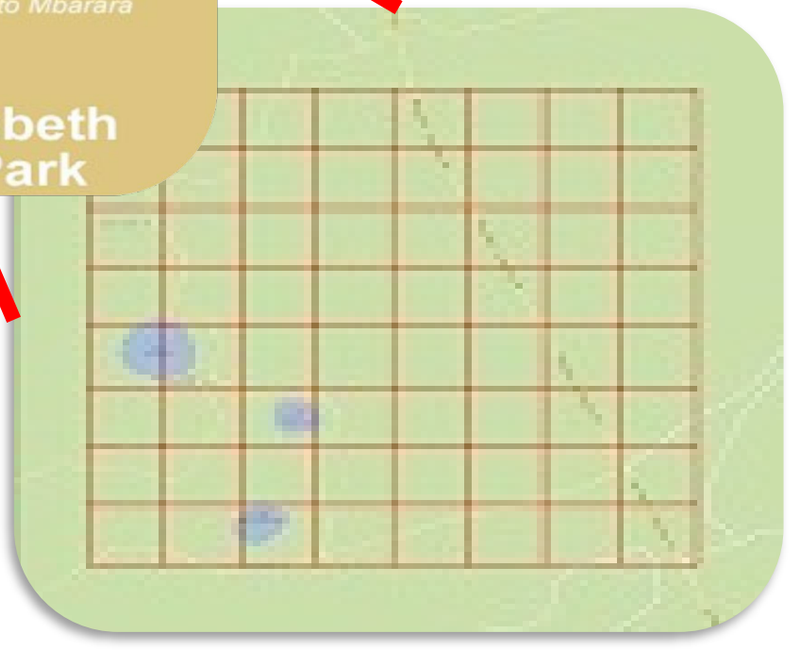

Ranger patrols: $X_{(i)}$

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

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

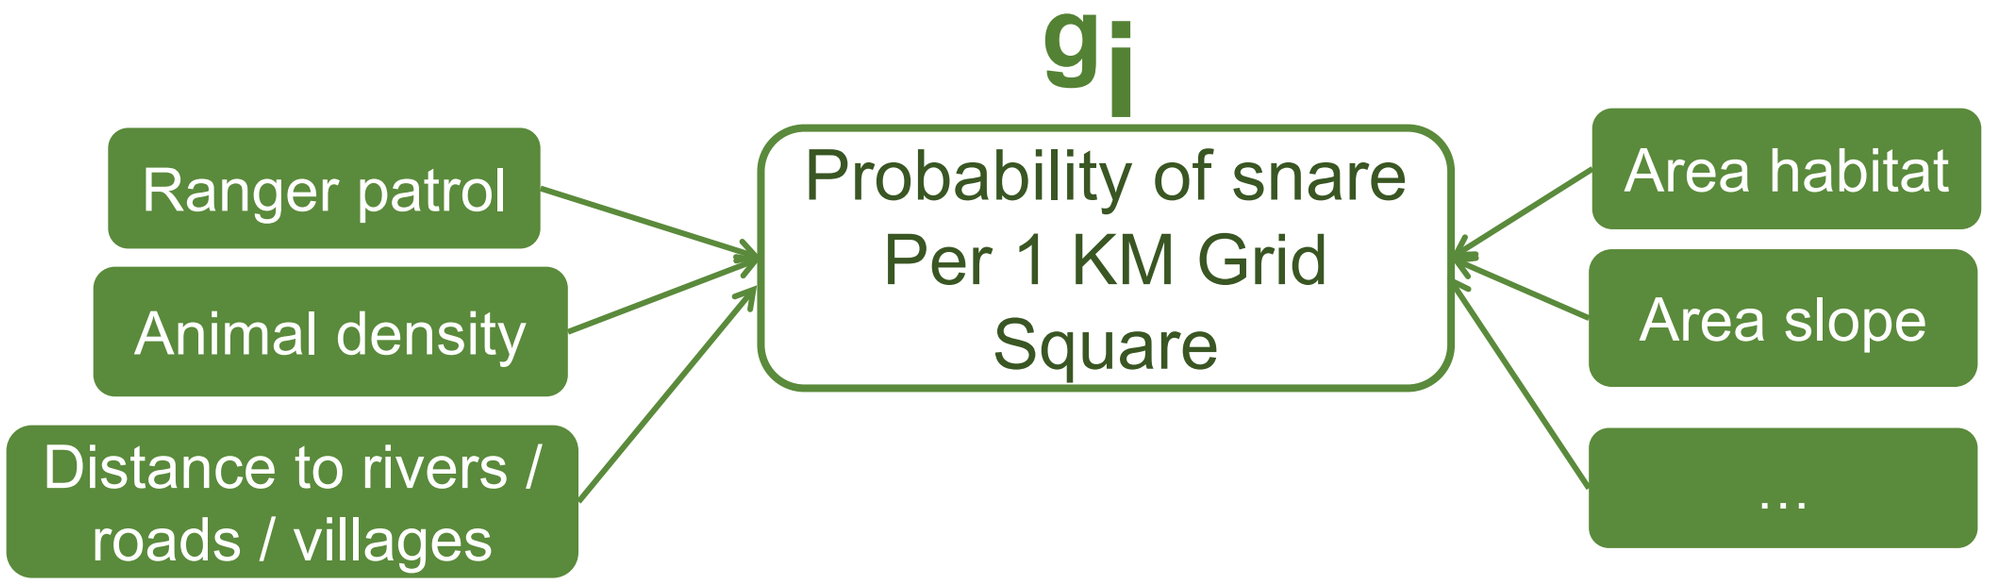

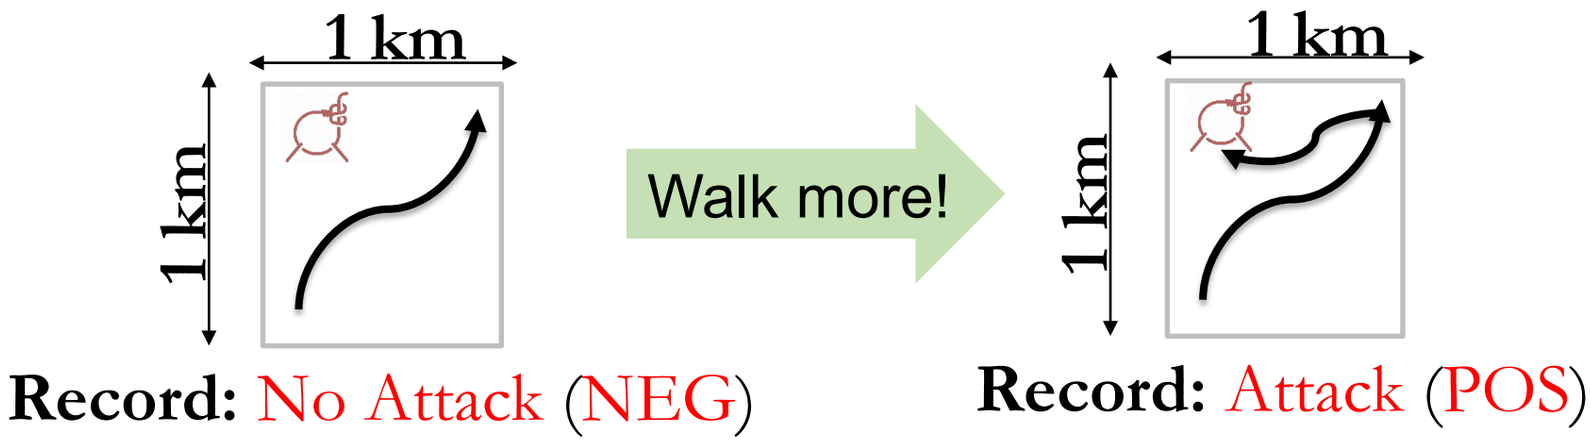

Adversary Modeling [2016]
Imperfect Crime Observation-aware Ensemble Model

Gholami

Training: Filtered Datasets

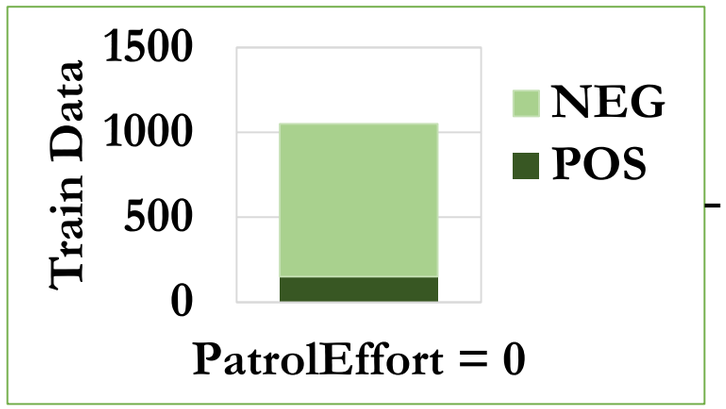

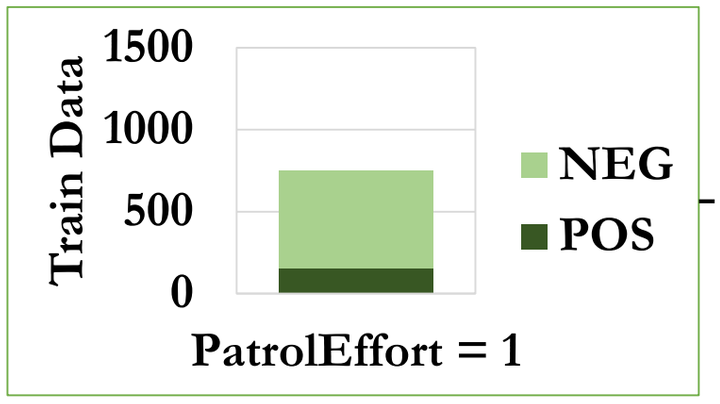

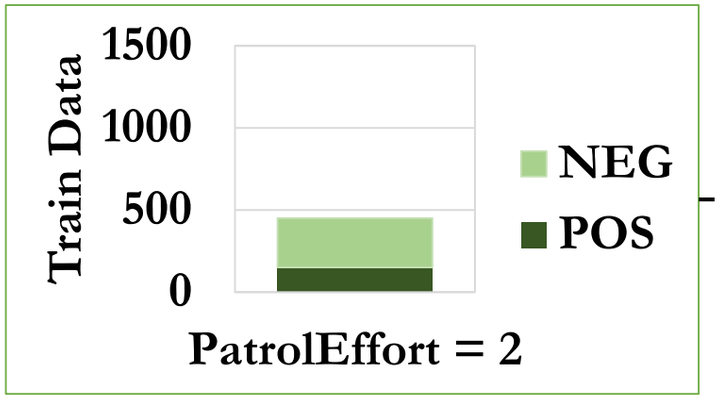

Predict: Ensemble of Classifiers

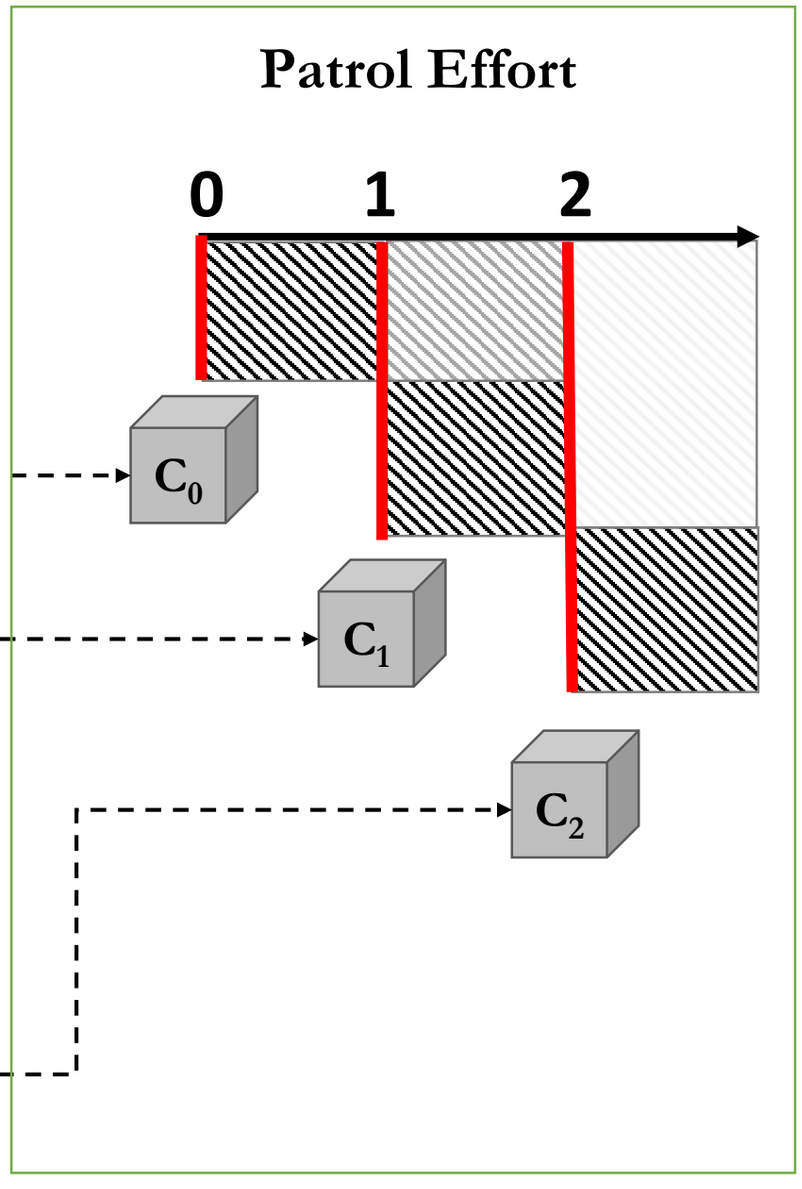

PAWS: Protection Assistant for Wildlife Security
Poacher Attack Prediction in the Lab


Gholami

Poacher Behavior Prediction

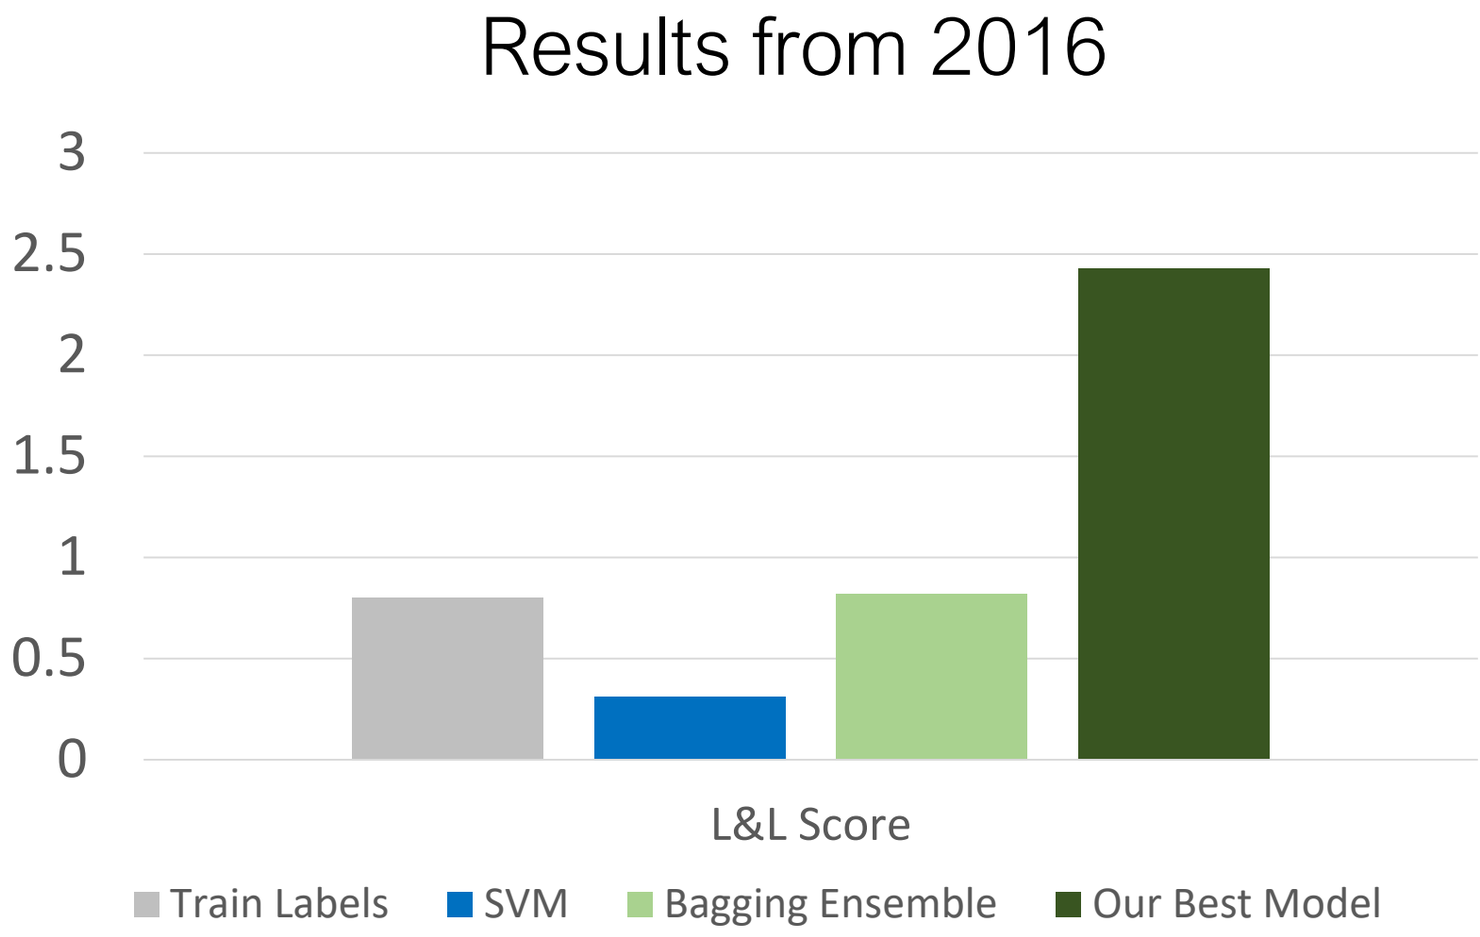

PAWS:
Real-world Deployment 2016: First Trial


Ford


Gholami

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

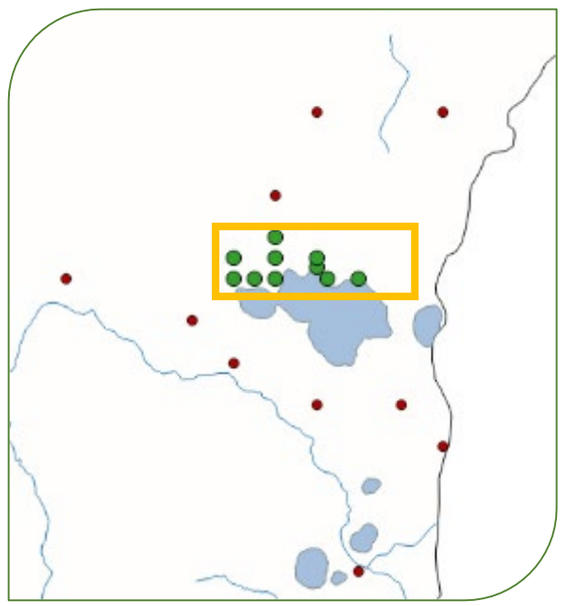

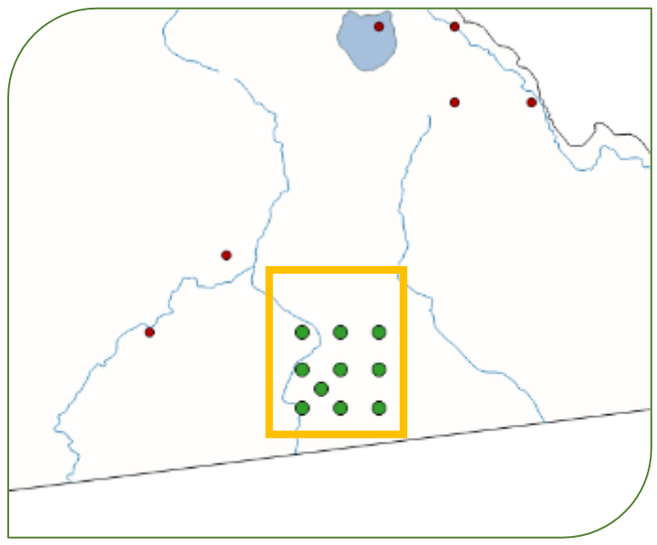

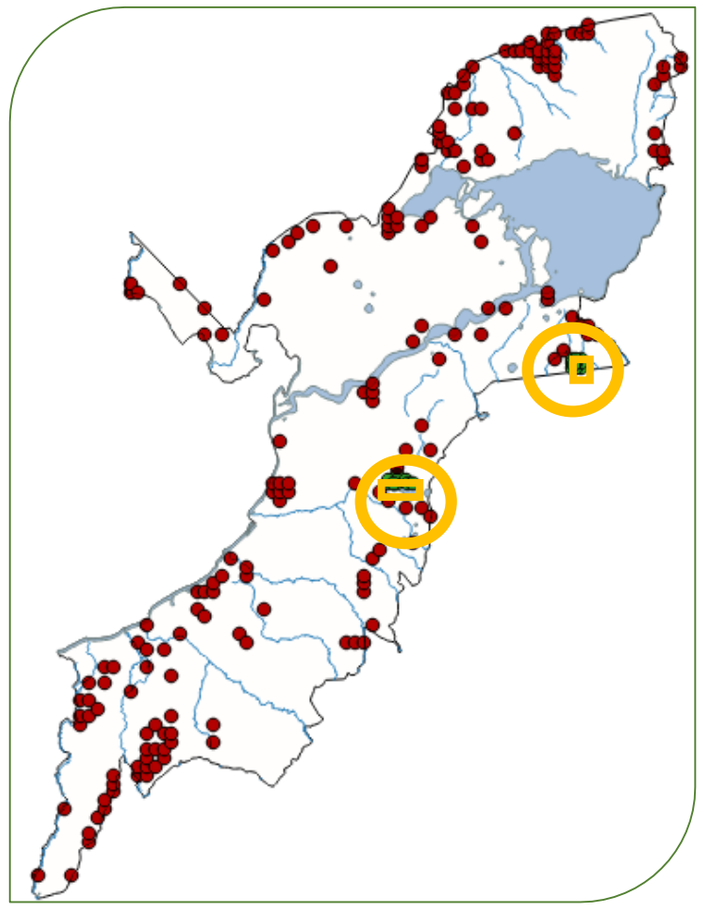

PAWS Real-world Deployment
Two Hot Spots Predicted


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

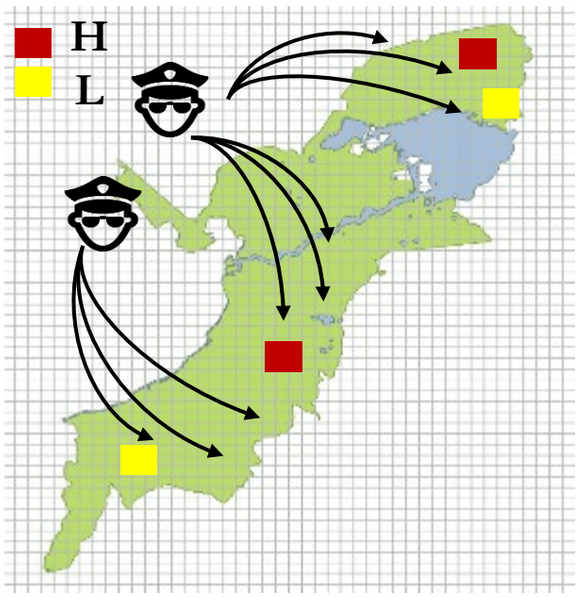

Queen Elizabeth National Park

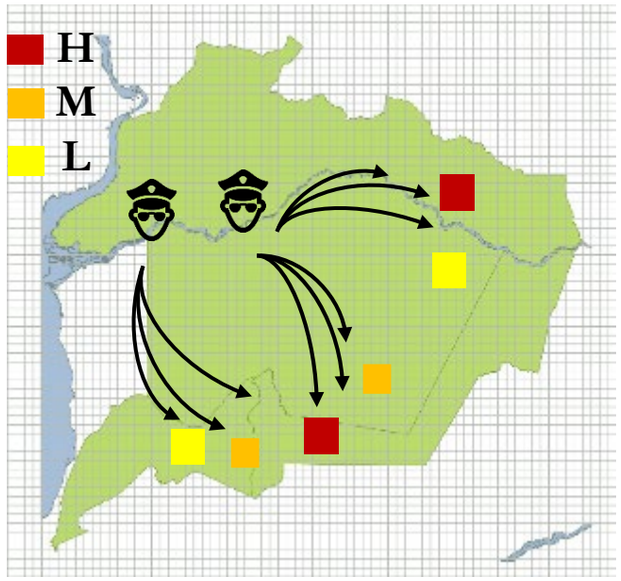

Murchison Falls National Park

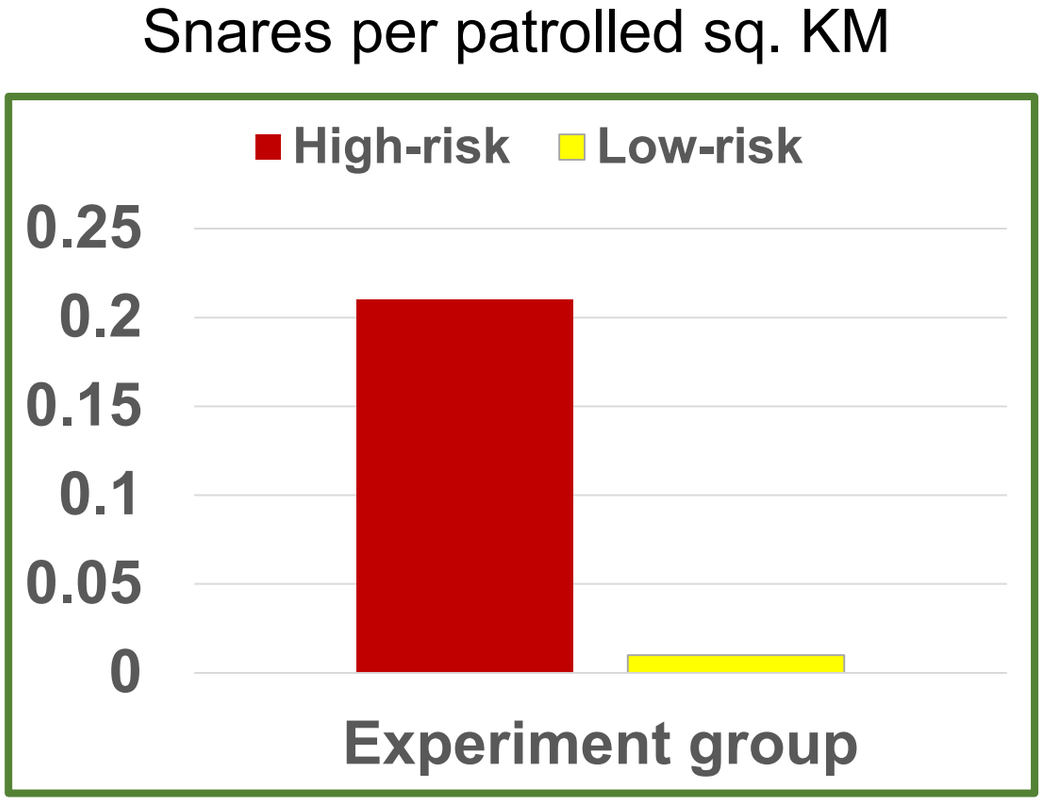

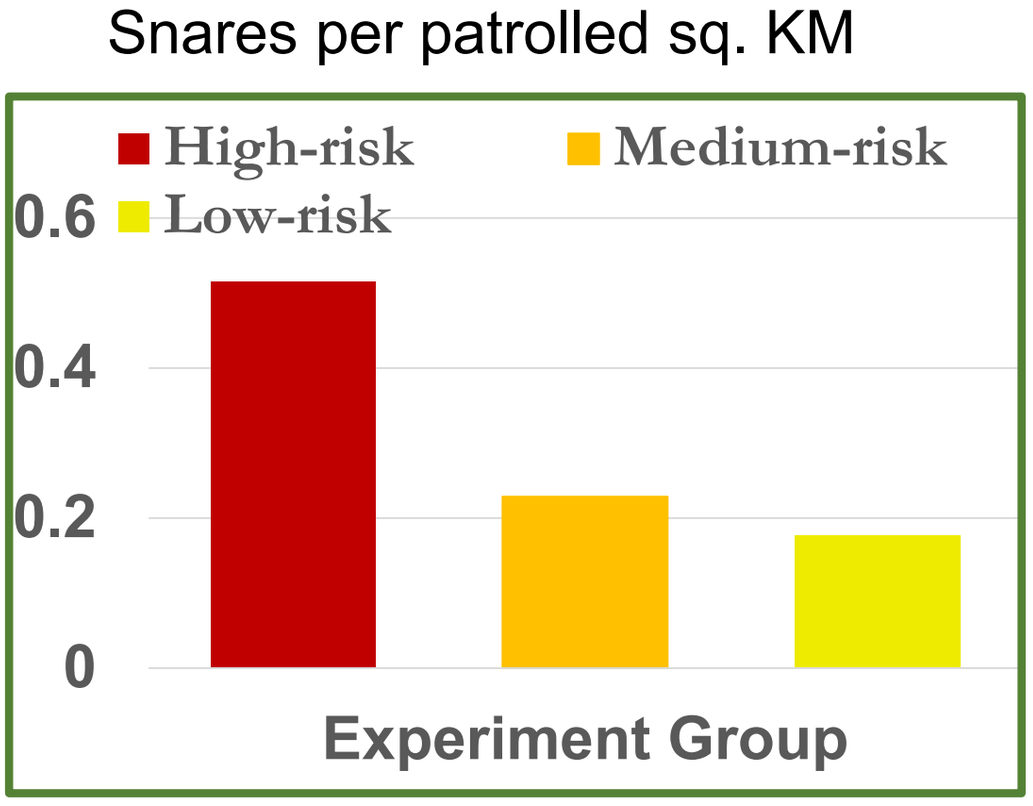

PAWS Real-world Deployment
Cambodia: Srepok Wildlife Sanctuary [2018-2019]


Xu

Date: 3/18/2019

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


Xu

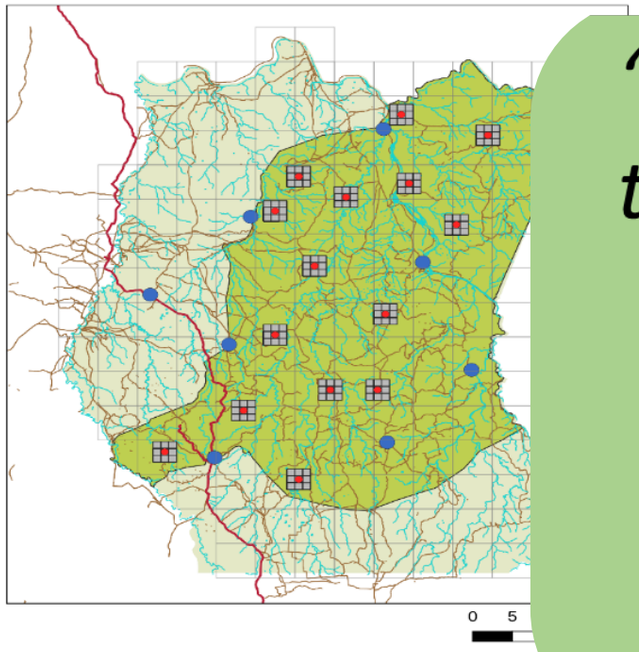

"@Milind: I am Super excited with the results. Let's get this going on other countries too this year."

Rohit Singh, WWF (2019)

- *521 snares/month our tests*

vs

- *101 snares/month 2018*

Snares per patrolled sq. KM

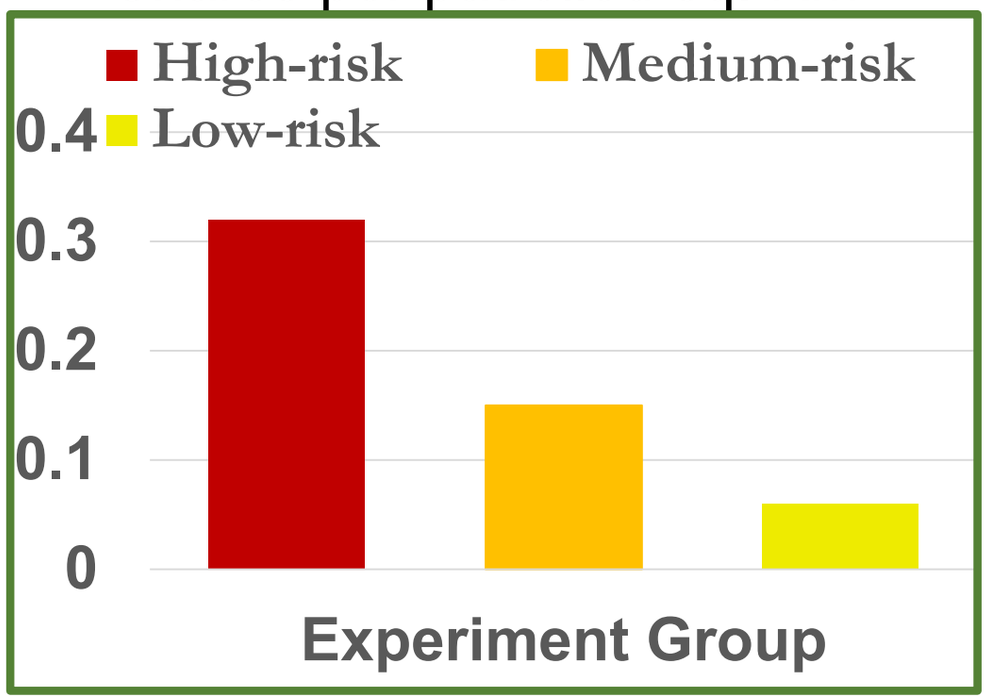

Green Security Games:
Integrating Real-Time Information in the Pipeline

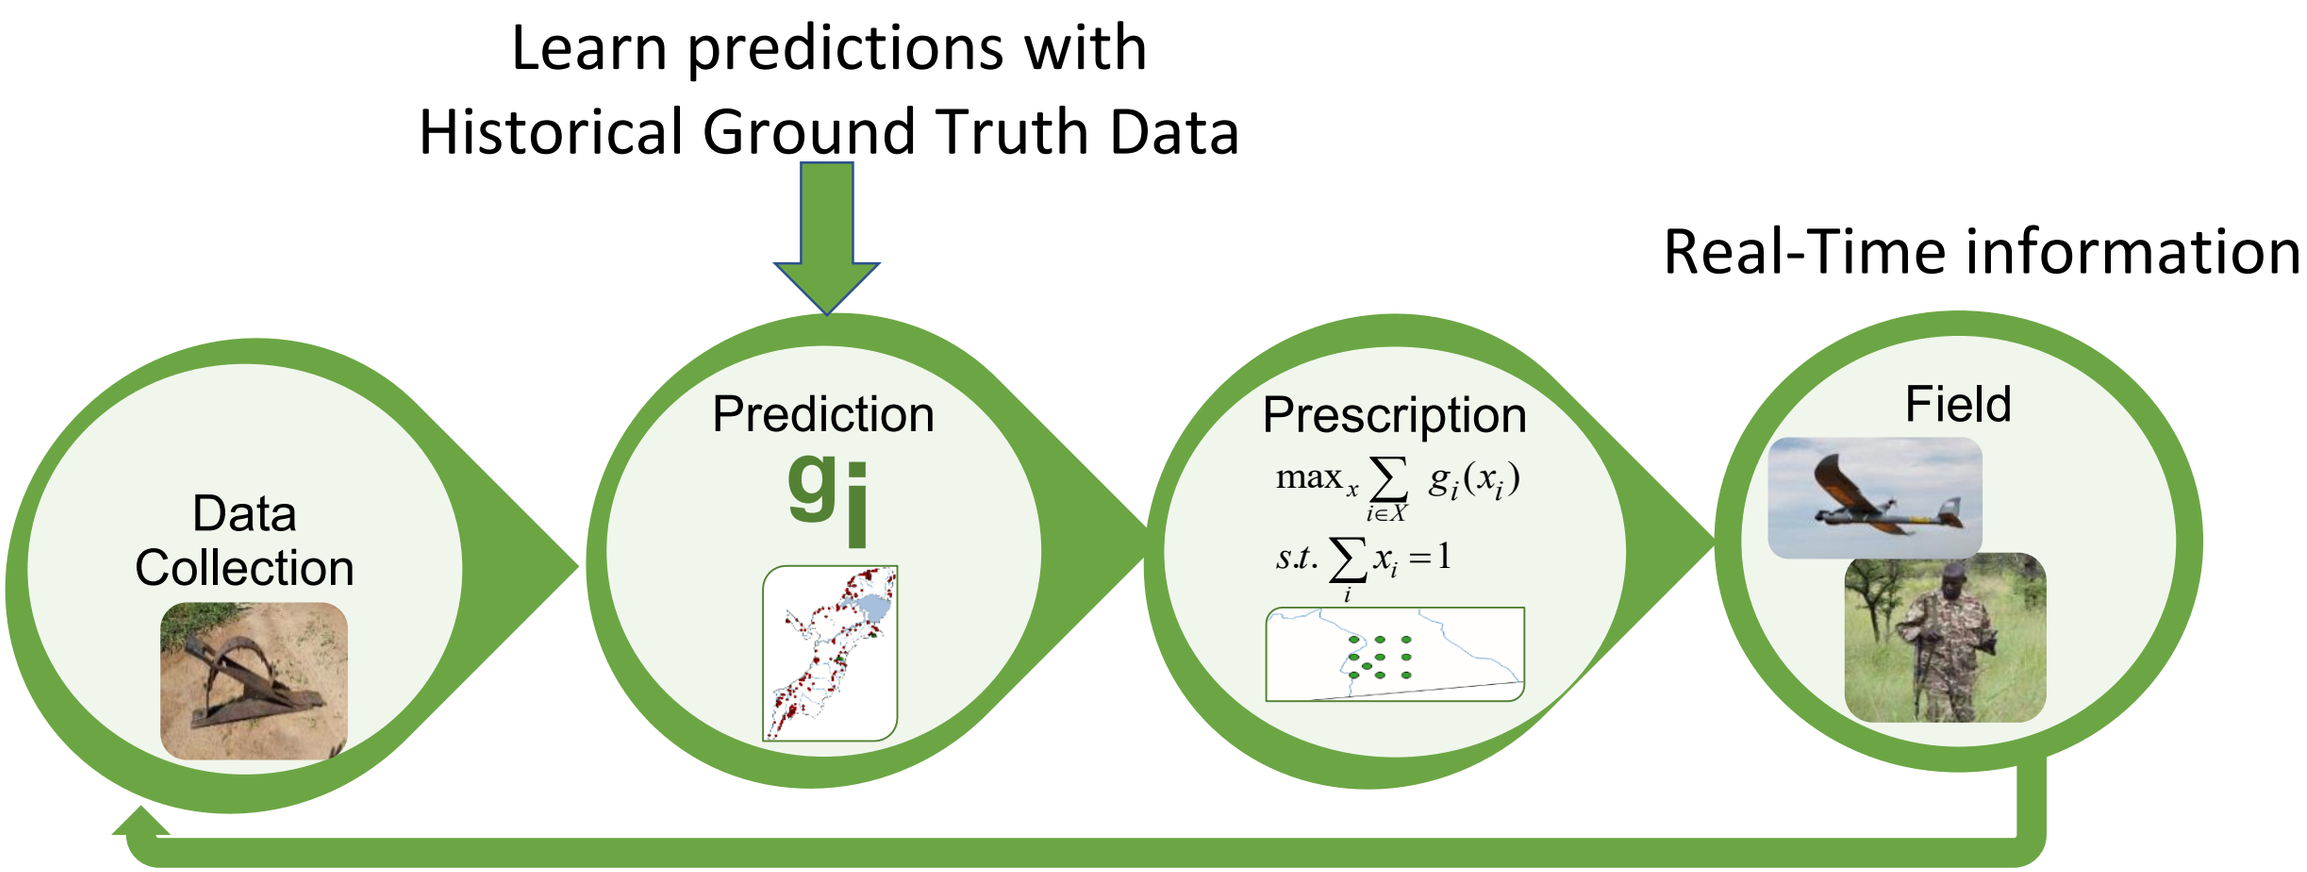

Green Security Games: Integrating Real-Time "SPOT" Information [2018]

Bondi

Air Shepherd
The Lindbergh Foundation

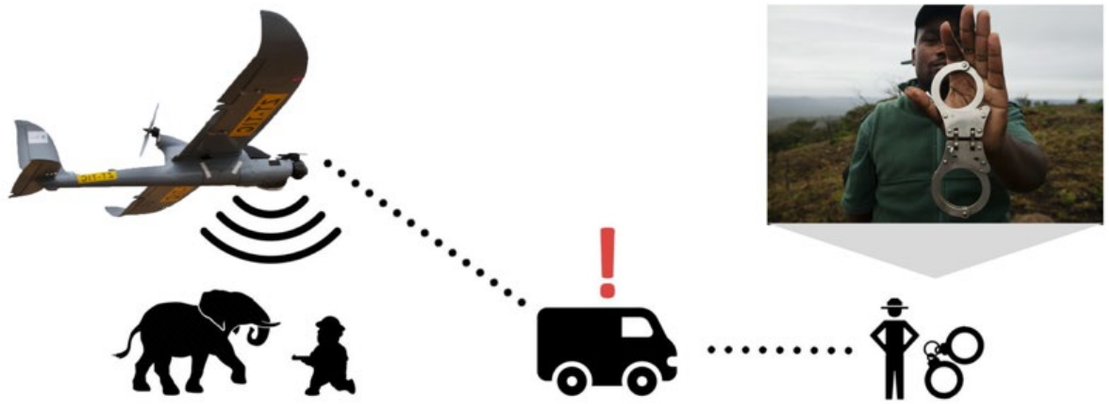

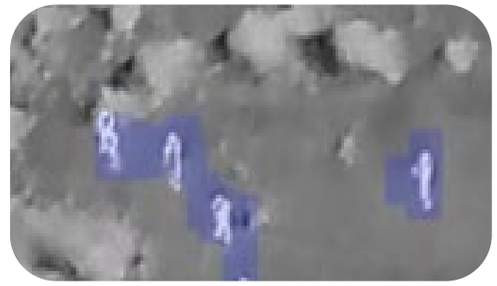

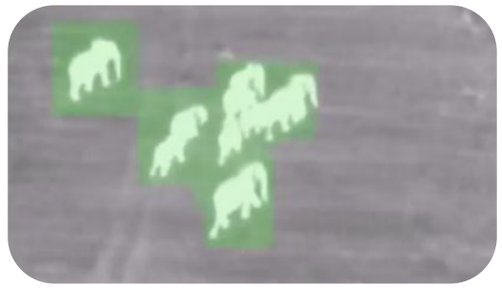

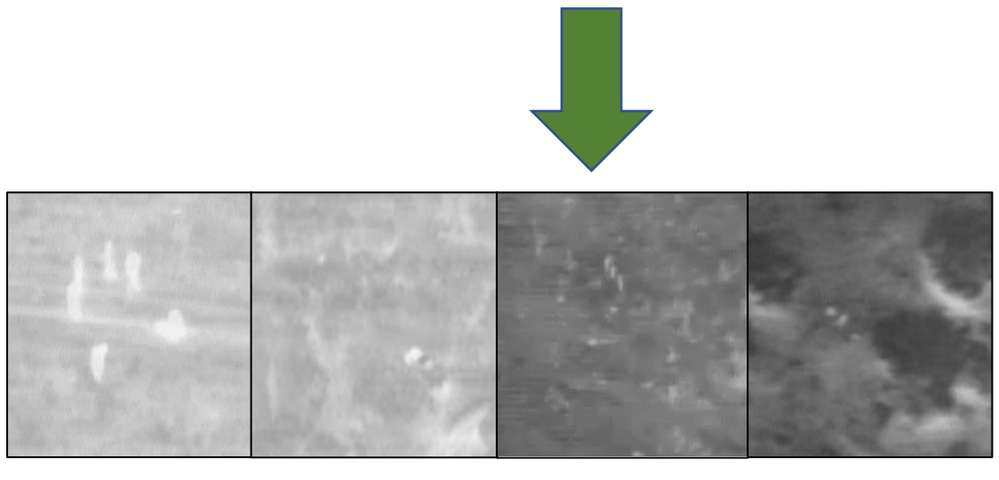

Goal: automatically find poachers

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

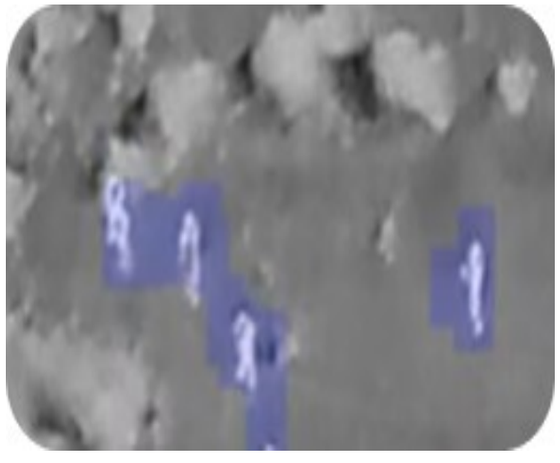

Prob(ranger) = 0.3

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

Prob(ranger) = 0.3

Air Shepherd
The Lindbergh Foundation

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

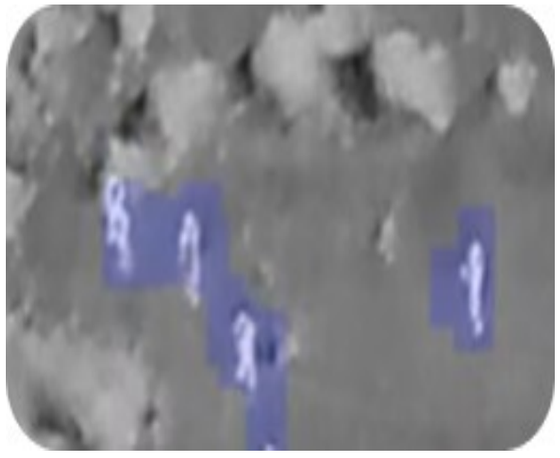

Prob(ranger) = 0.3

Air Shepherd
The Lindbergh Foundation

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

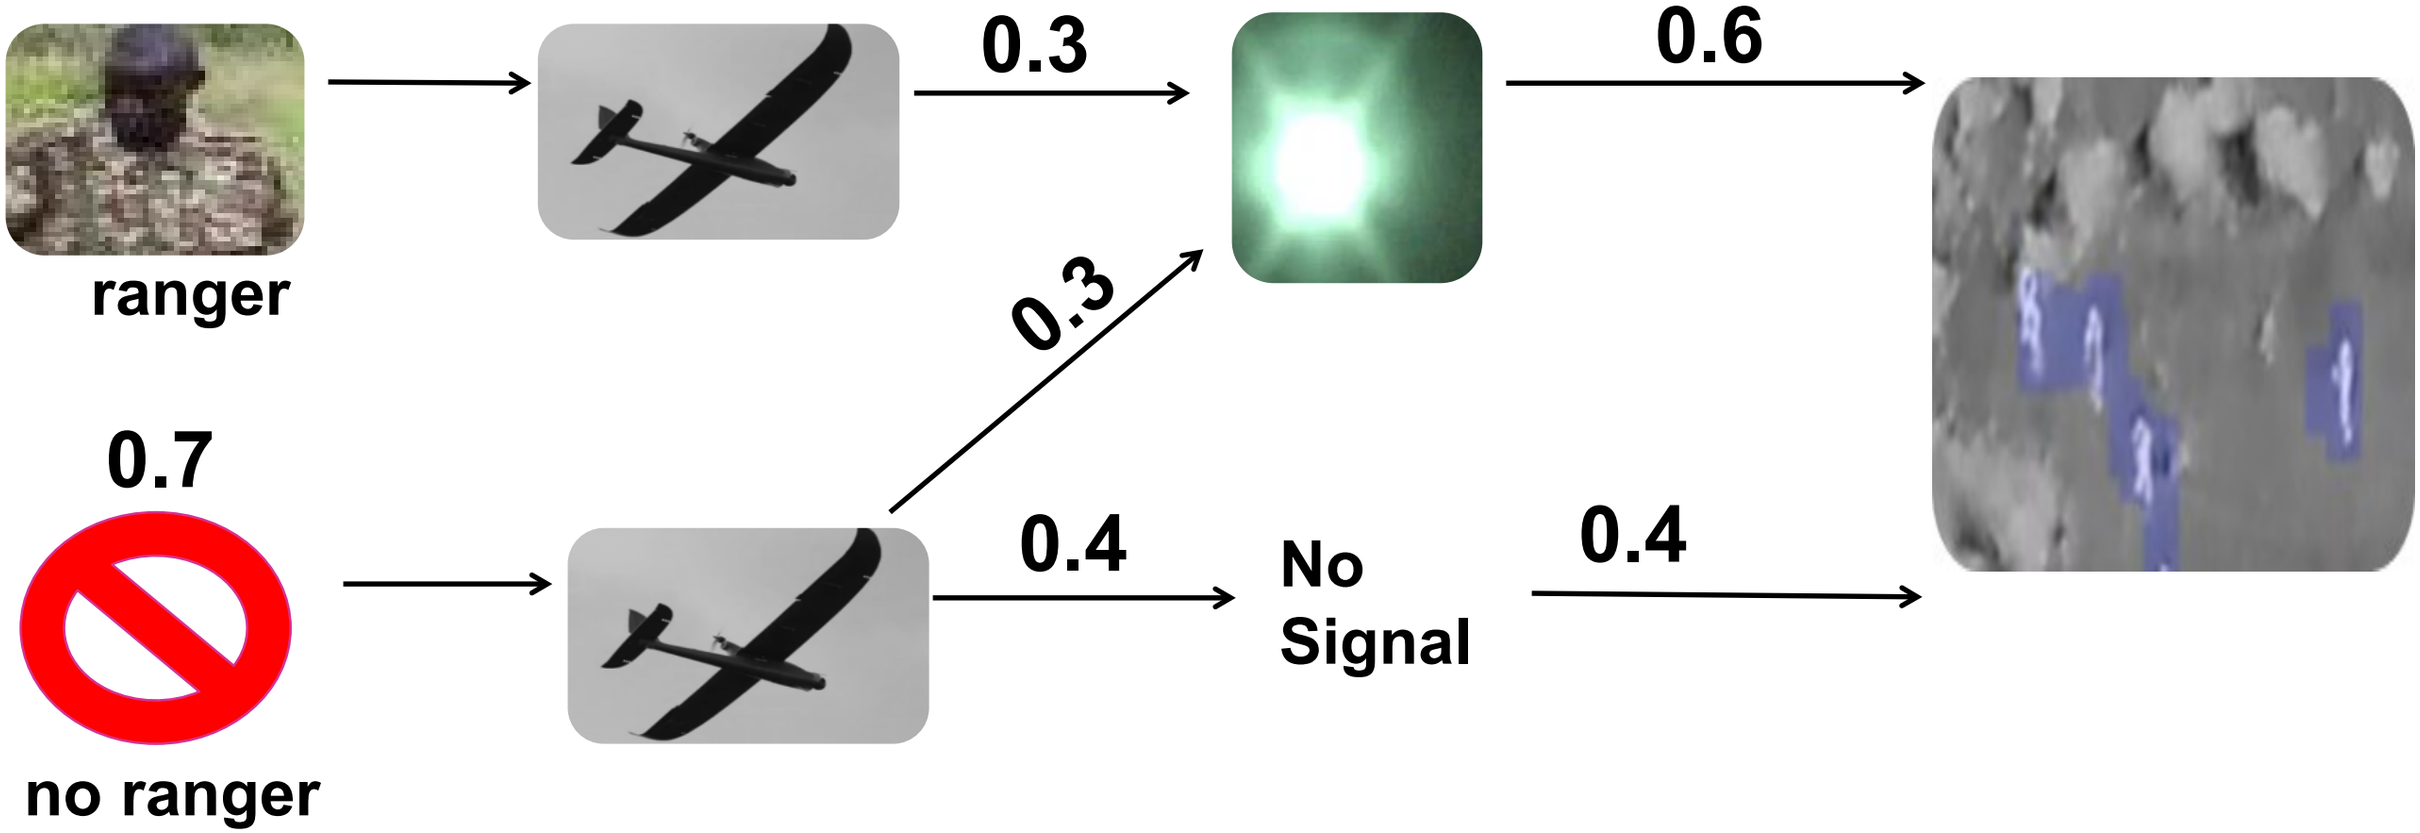

Strategic Signaling: Informational Advantage Defender Knows Pure & Mixed Strategy

Xu

Theorem: Signaling reduces complexity of equilibrium computation

- Poacher best interest to "believe signal" even if know 50% time defender is lying

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

Protect Wildlife
600
National Parks
Around the Globe

Also: Protect Forests, Fisheries…

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

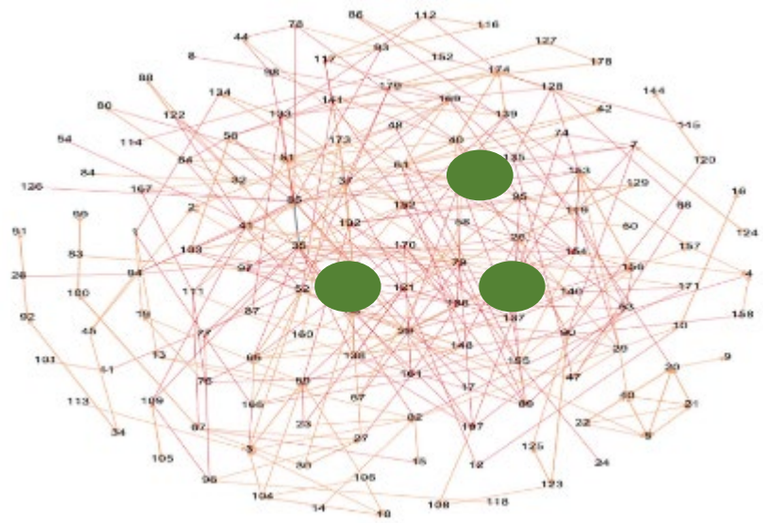

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
- Objective:
  - Maximize expected number of influenced nodes
- *Assumption: Independent cascade model of information spread*

Independent Cascade Model and Real-world Physical Social Networks

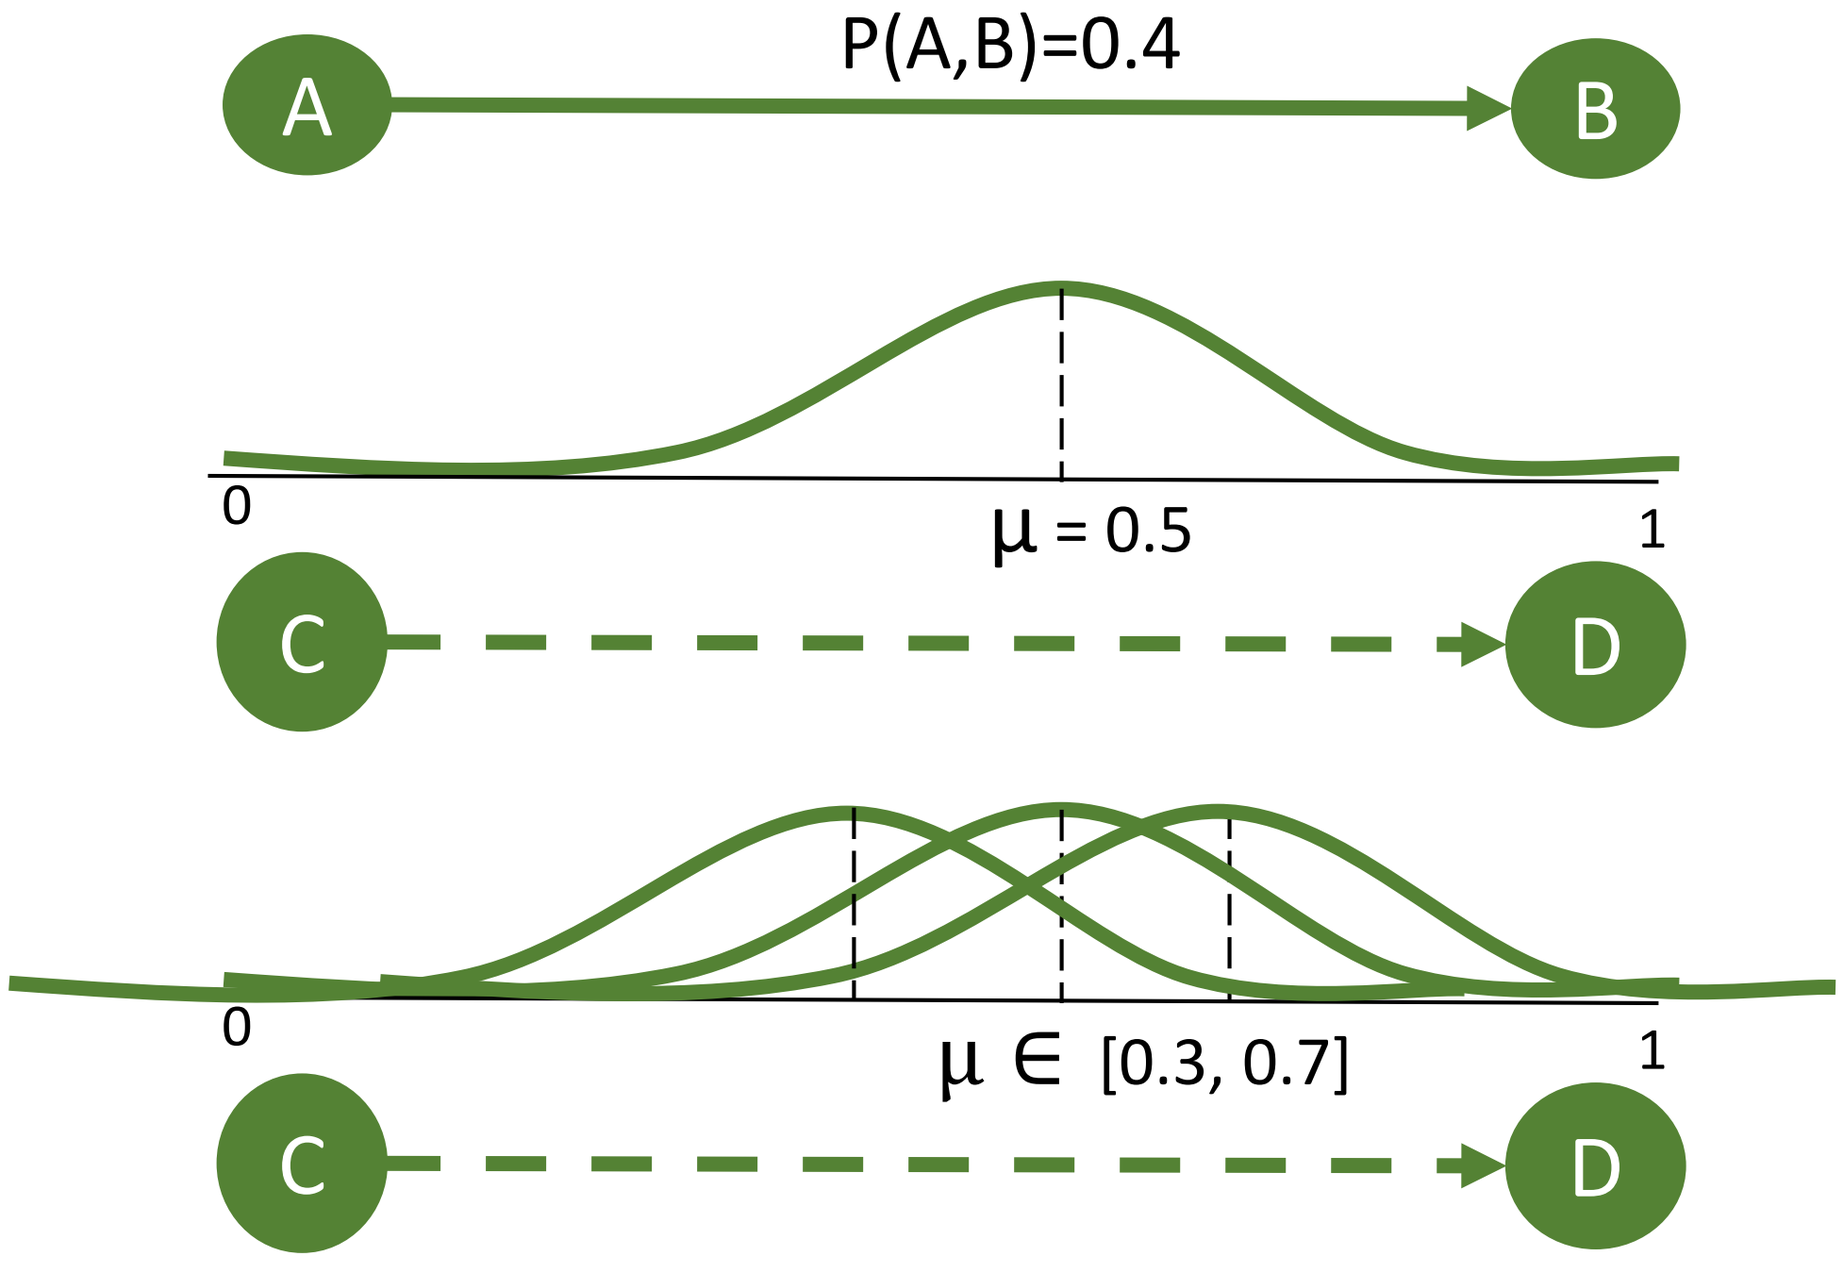

Robust, Dynamic Influence Maximization


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

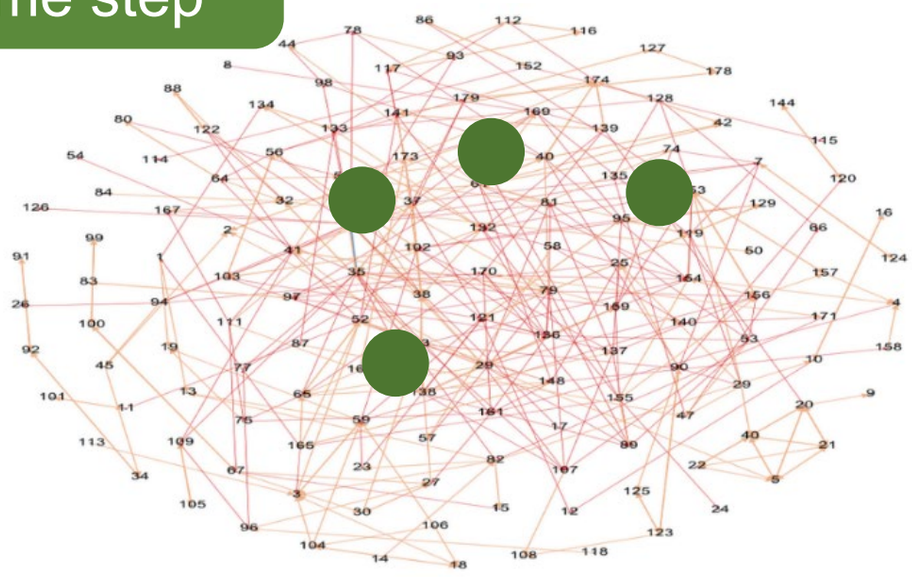

K = 4
2nd time step

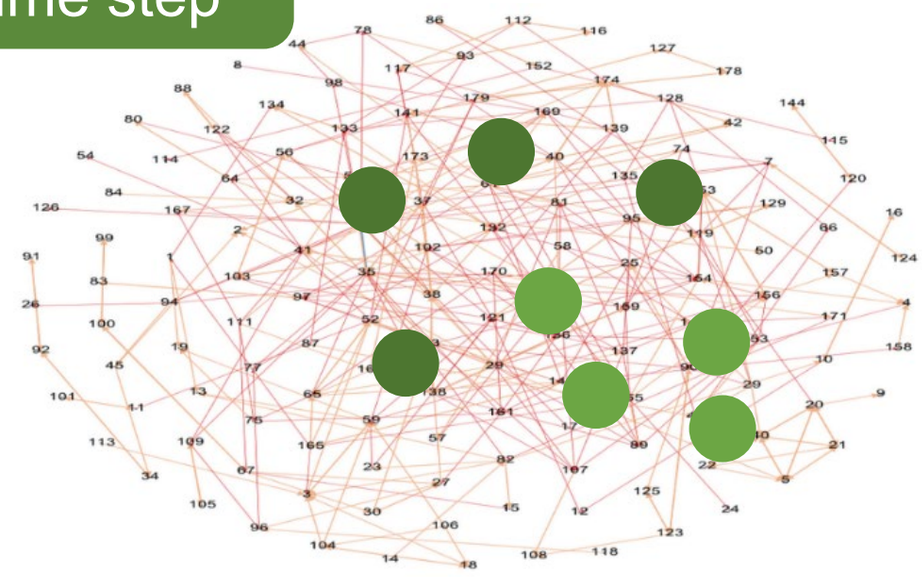

HEALER: POMDP Model for Multi-Step Policy Robust, Dynamic Influence Maximization

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

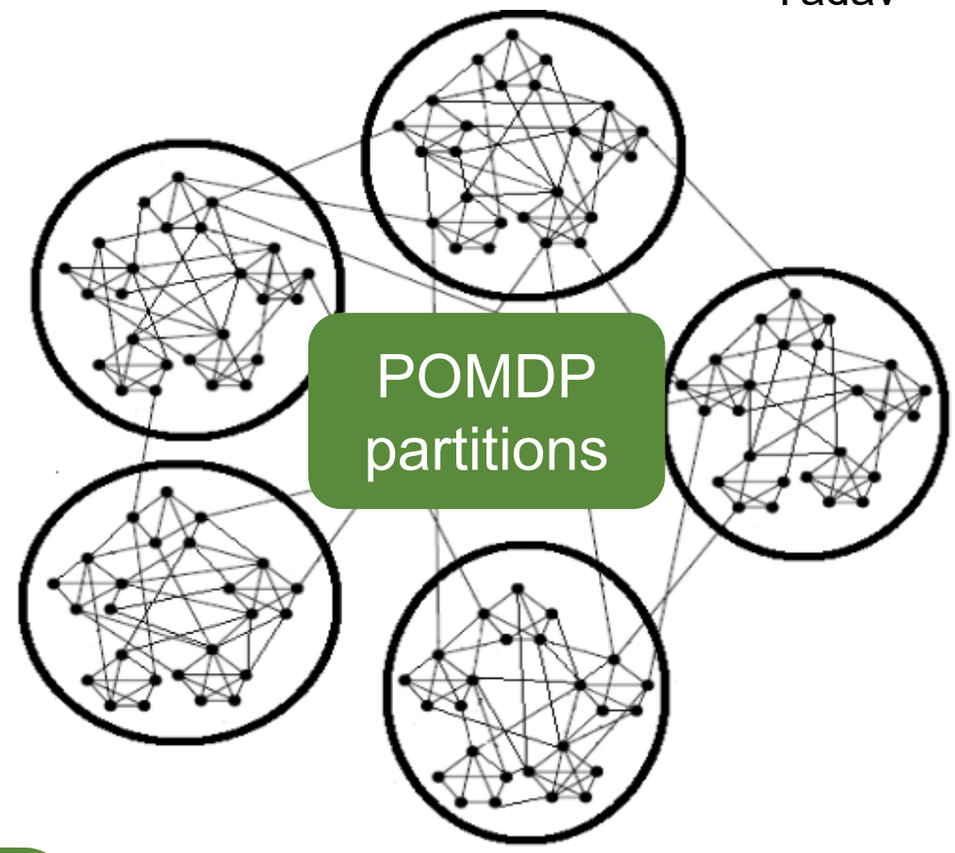

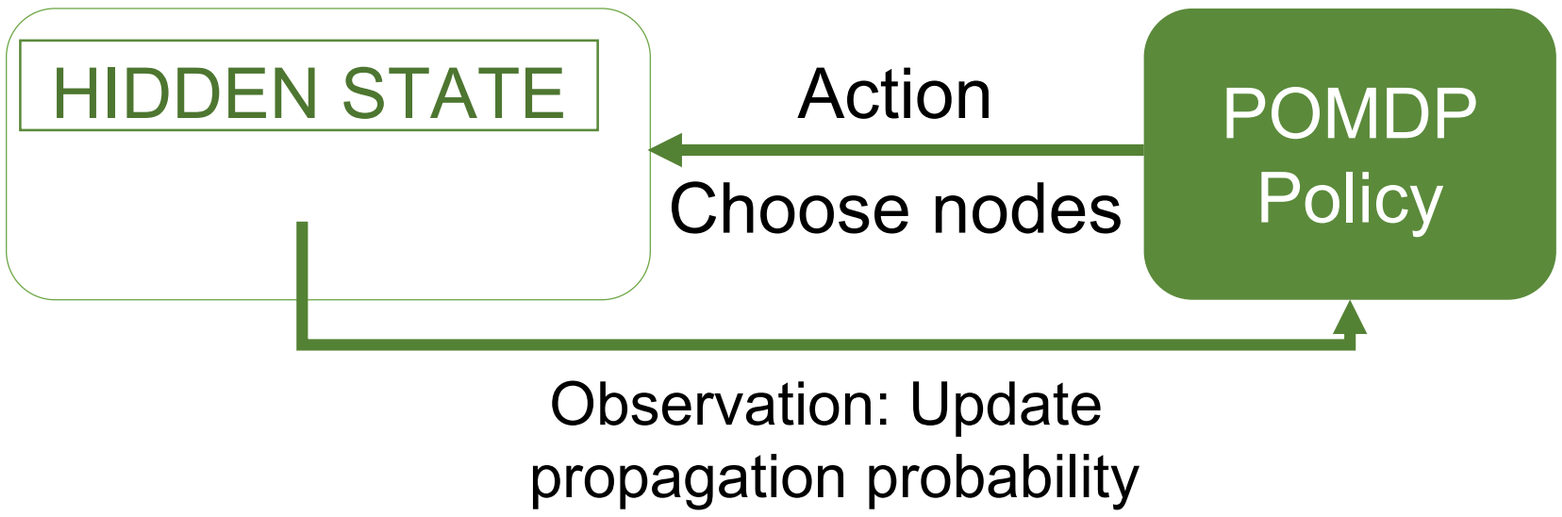

Pilot Tests with HEALER with 170 Homeless Youth [2017]


Yadav


Wilder

Recruited youths:

HEALER	HEALER++	DEGREE CENTRALITY
62	56	55

12 peer leaders

Results: Pilot Studies [2017]


Yadav


Wilder

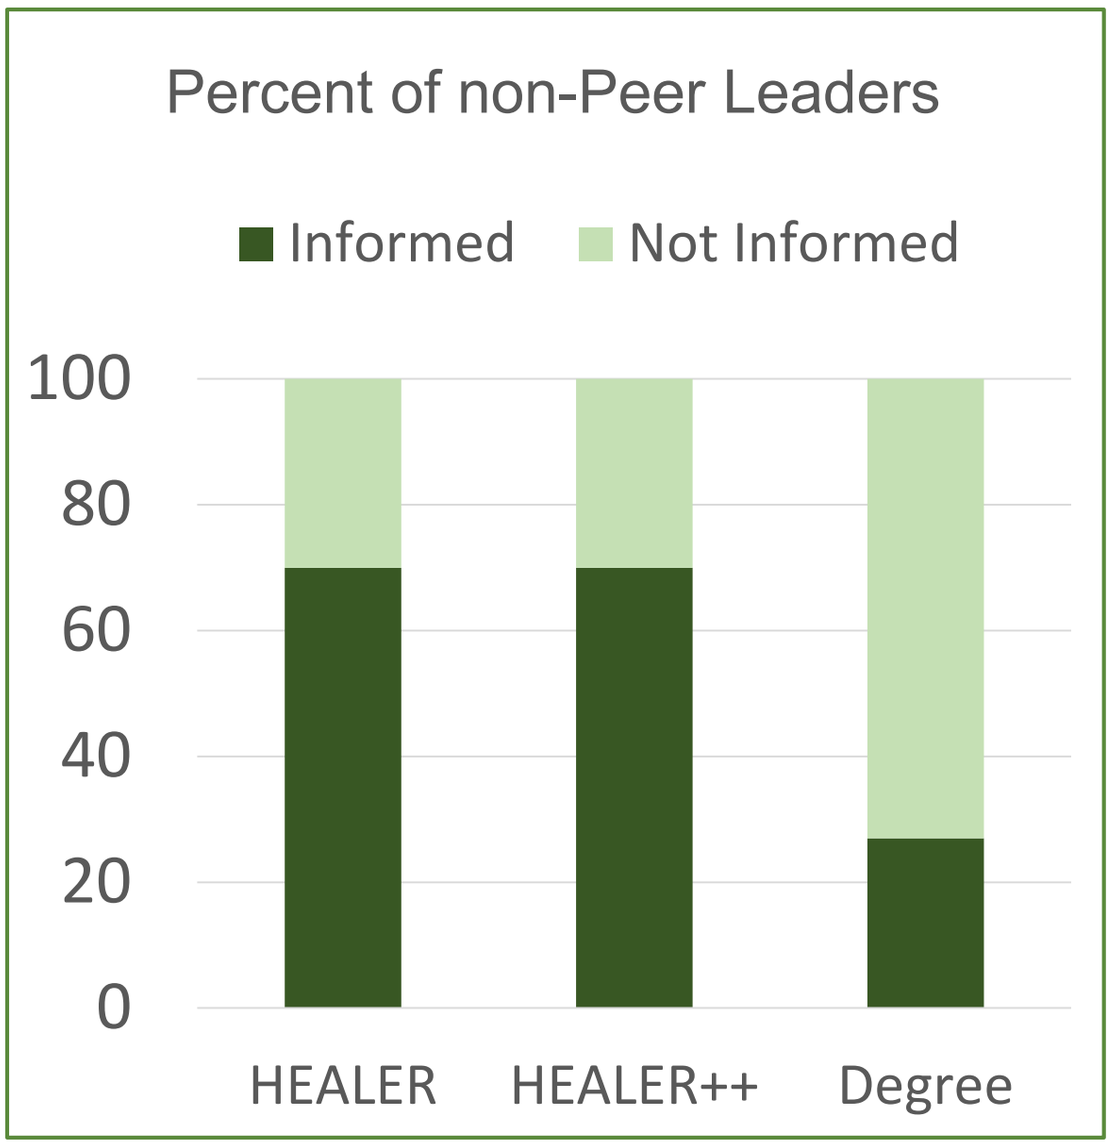

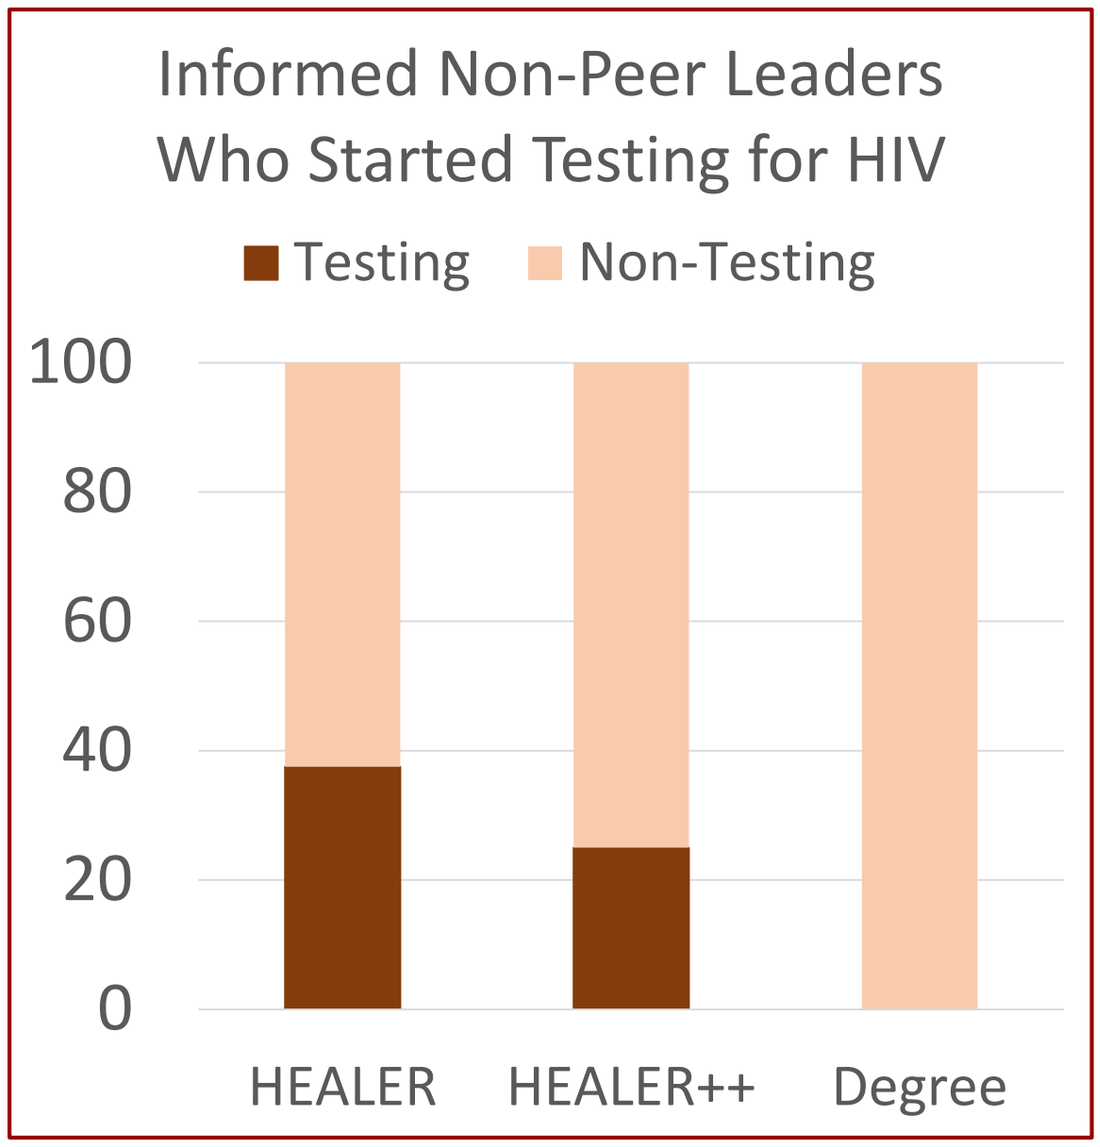

More details: Journal of Society of Social Work & Research (Nov 2018)

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


Wilder

Data collection costly

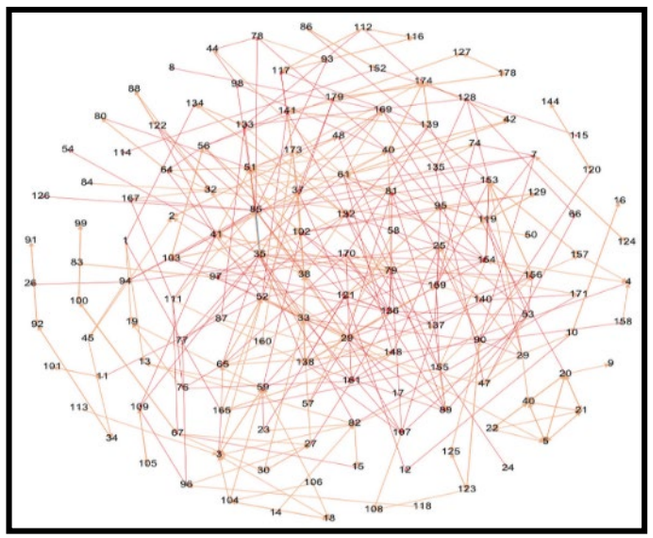

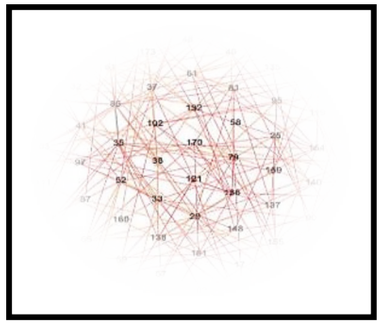

Sample 18%

Sampling from largest communities

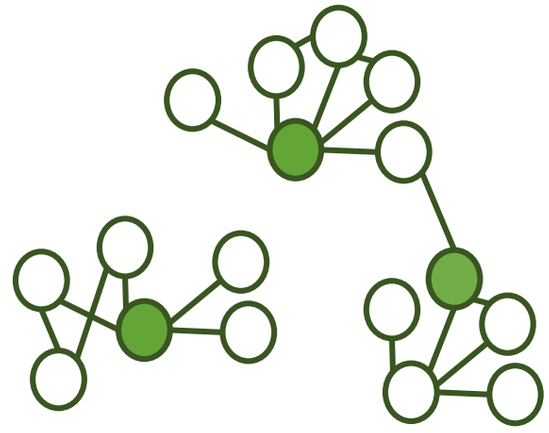

New experiment With 60 homeless youth

12 peer leaders

Results: Pilot Studies with New Sampling Algorithm [2018]


Wilder

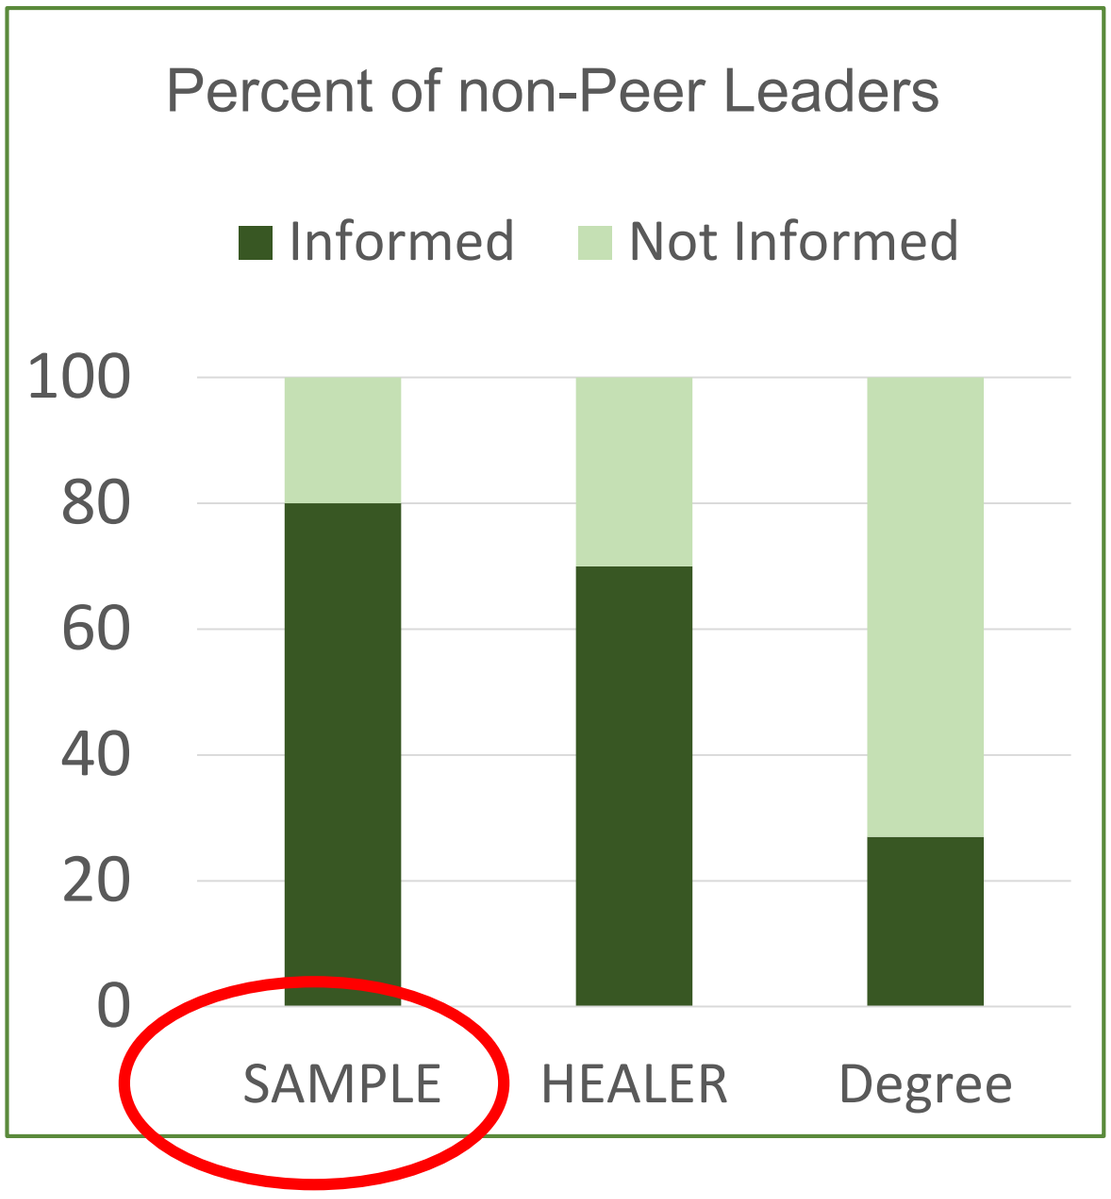

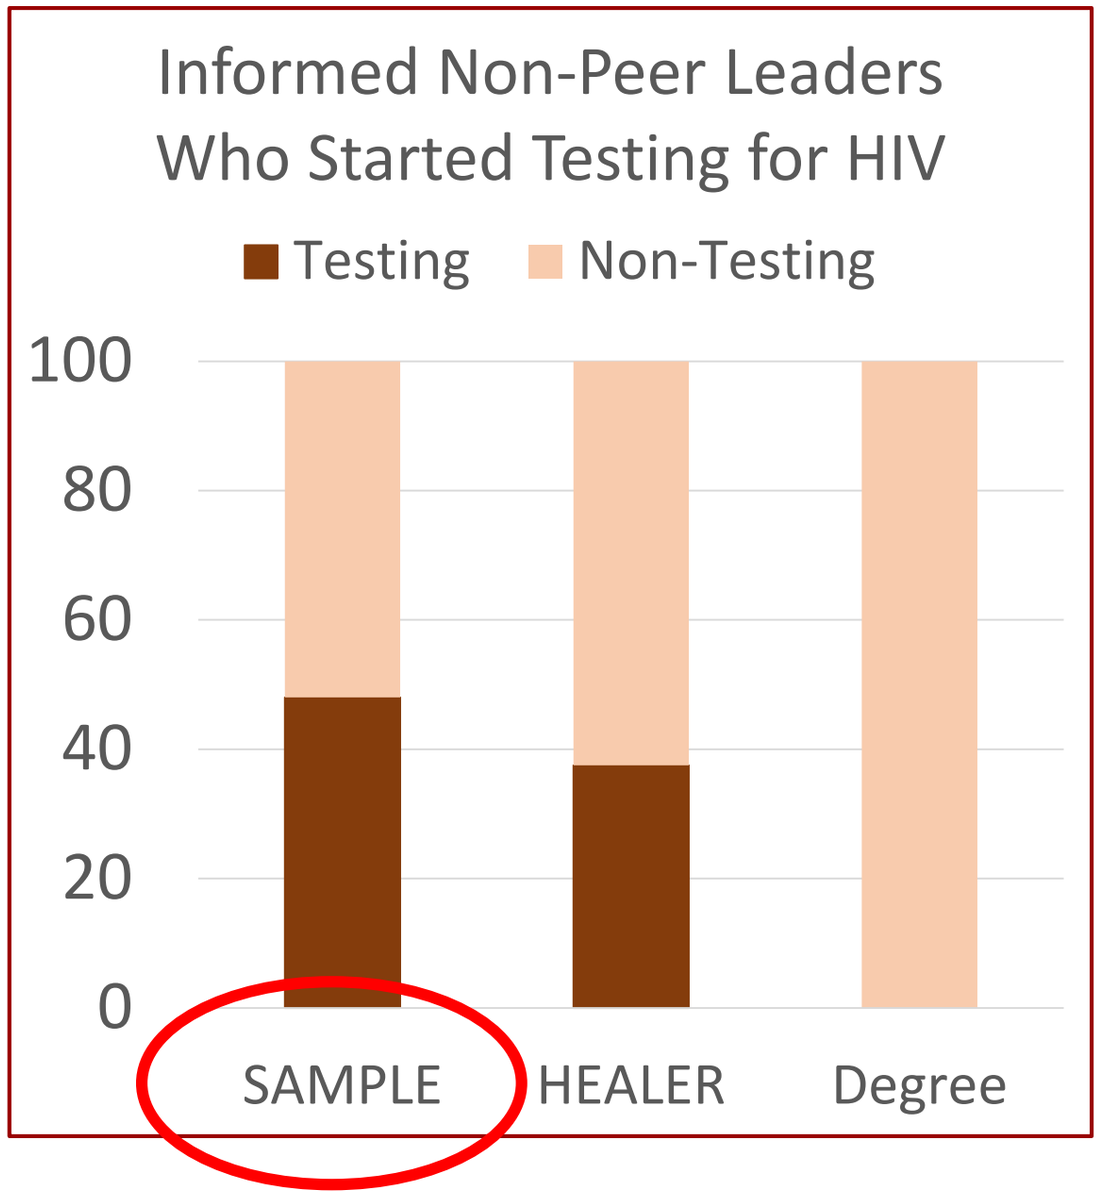

AI Assistant: HEALER

Continuing Research on HIV prevention [2019]

- Completing 900 youth study at three homeless shelters

my friend's place

S.P.Y

safe place for youth

LOS ANGELES LGBT CENTER

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*

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

Killian

- *Working jointly with Everwell Health Solutions & Microsoft Research India*
- *Everwell collaborates on software: Serves millions of TB patients in India, other countries*

everwell

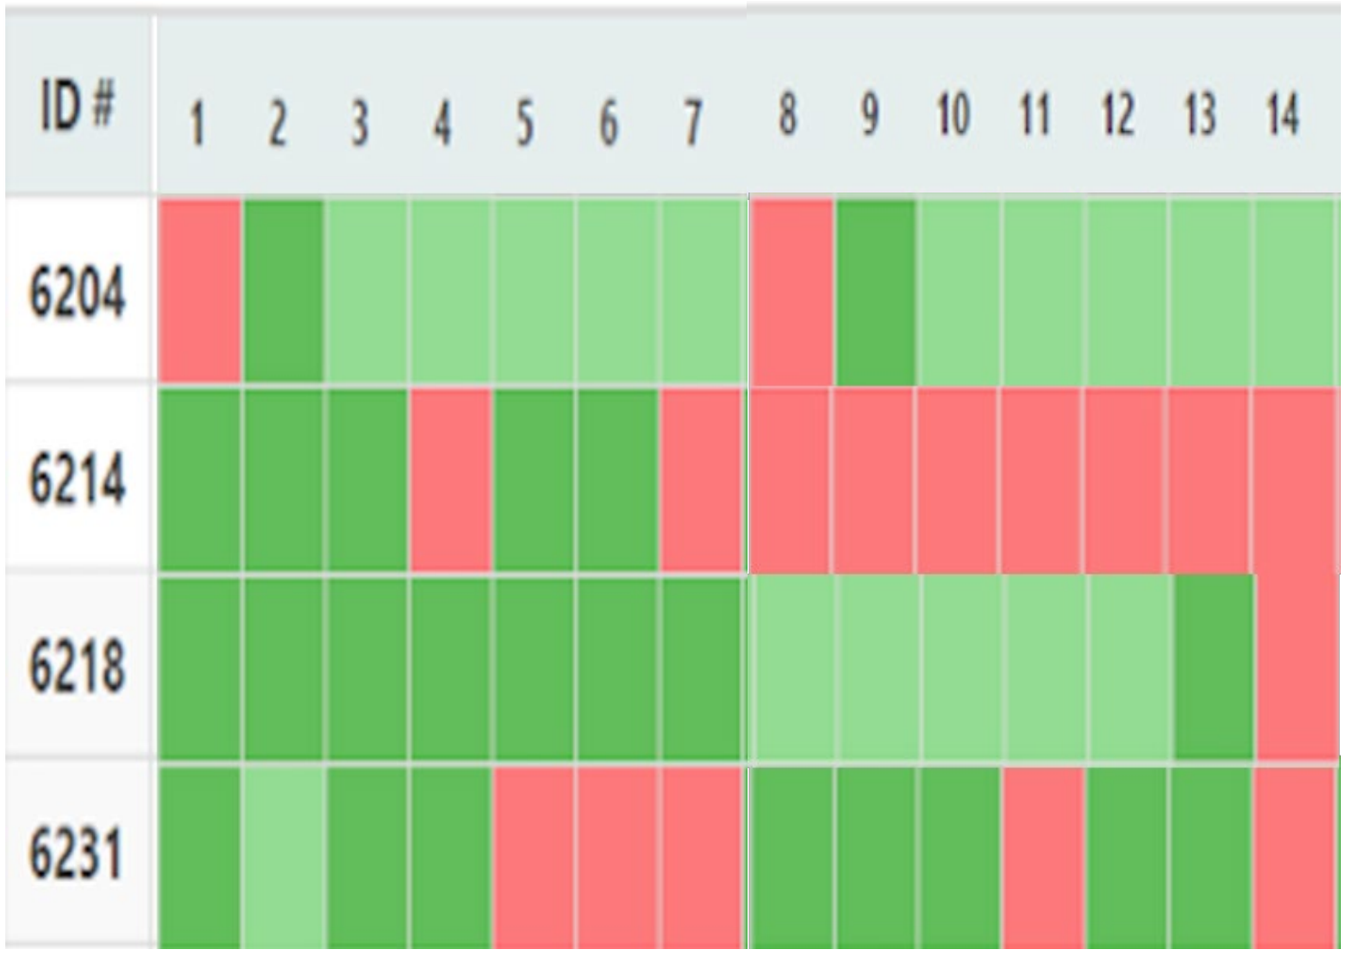

TB Treatment Adherence but Limited Resources: Intervening Selectively before patients miss doses

Killian

Data Collect

Phone logs

Predict high risk patients

RF or LSTM

Prescription

Constraint Top K

Field

- *15K patients, 1.5M calls*

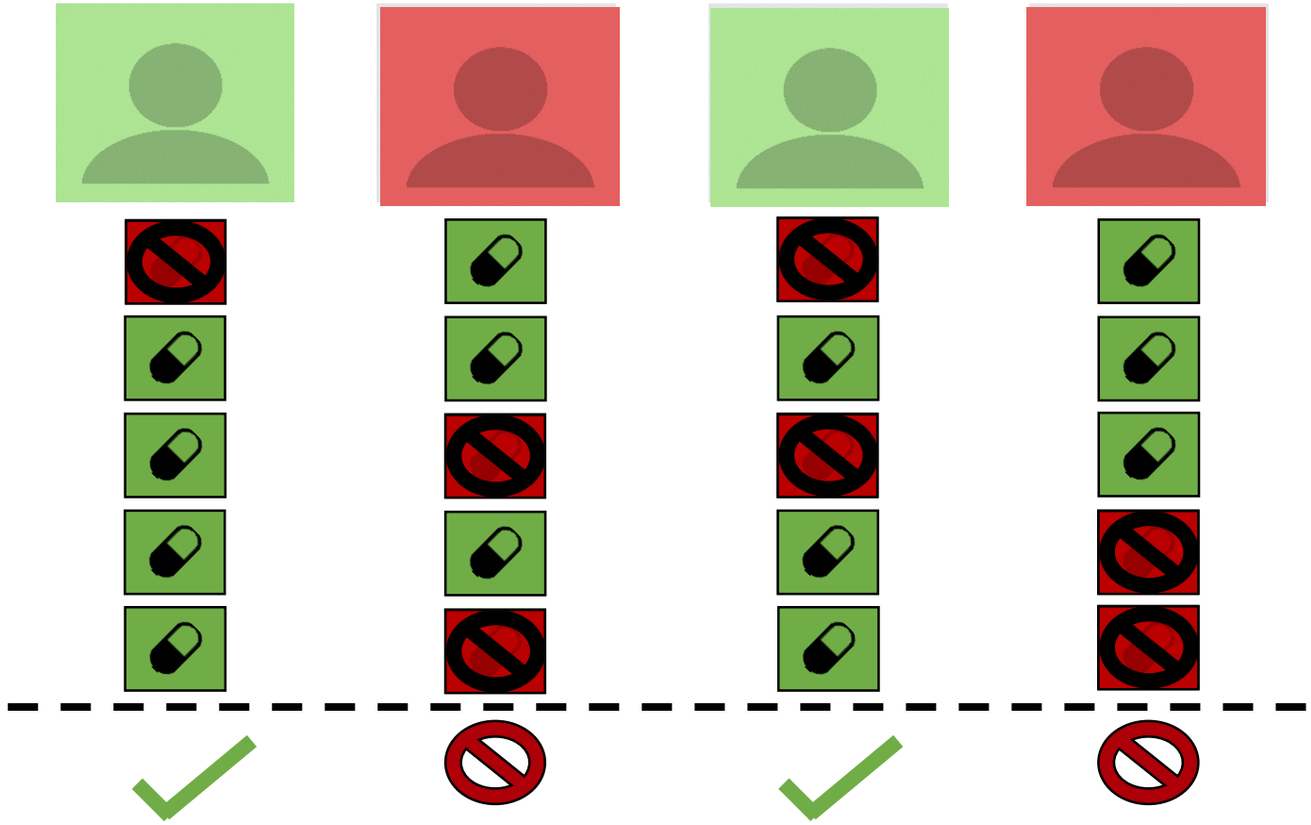

Increasing TB Treatment Adherence: Intervening before patients miss doses [2019]

Killian

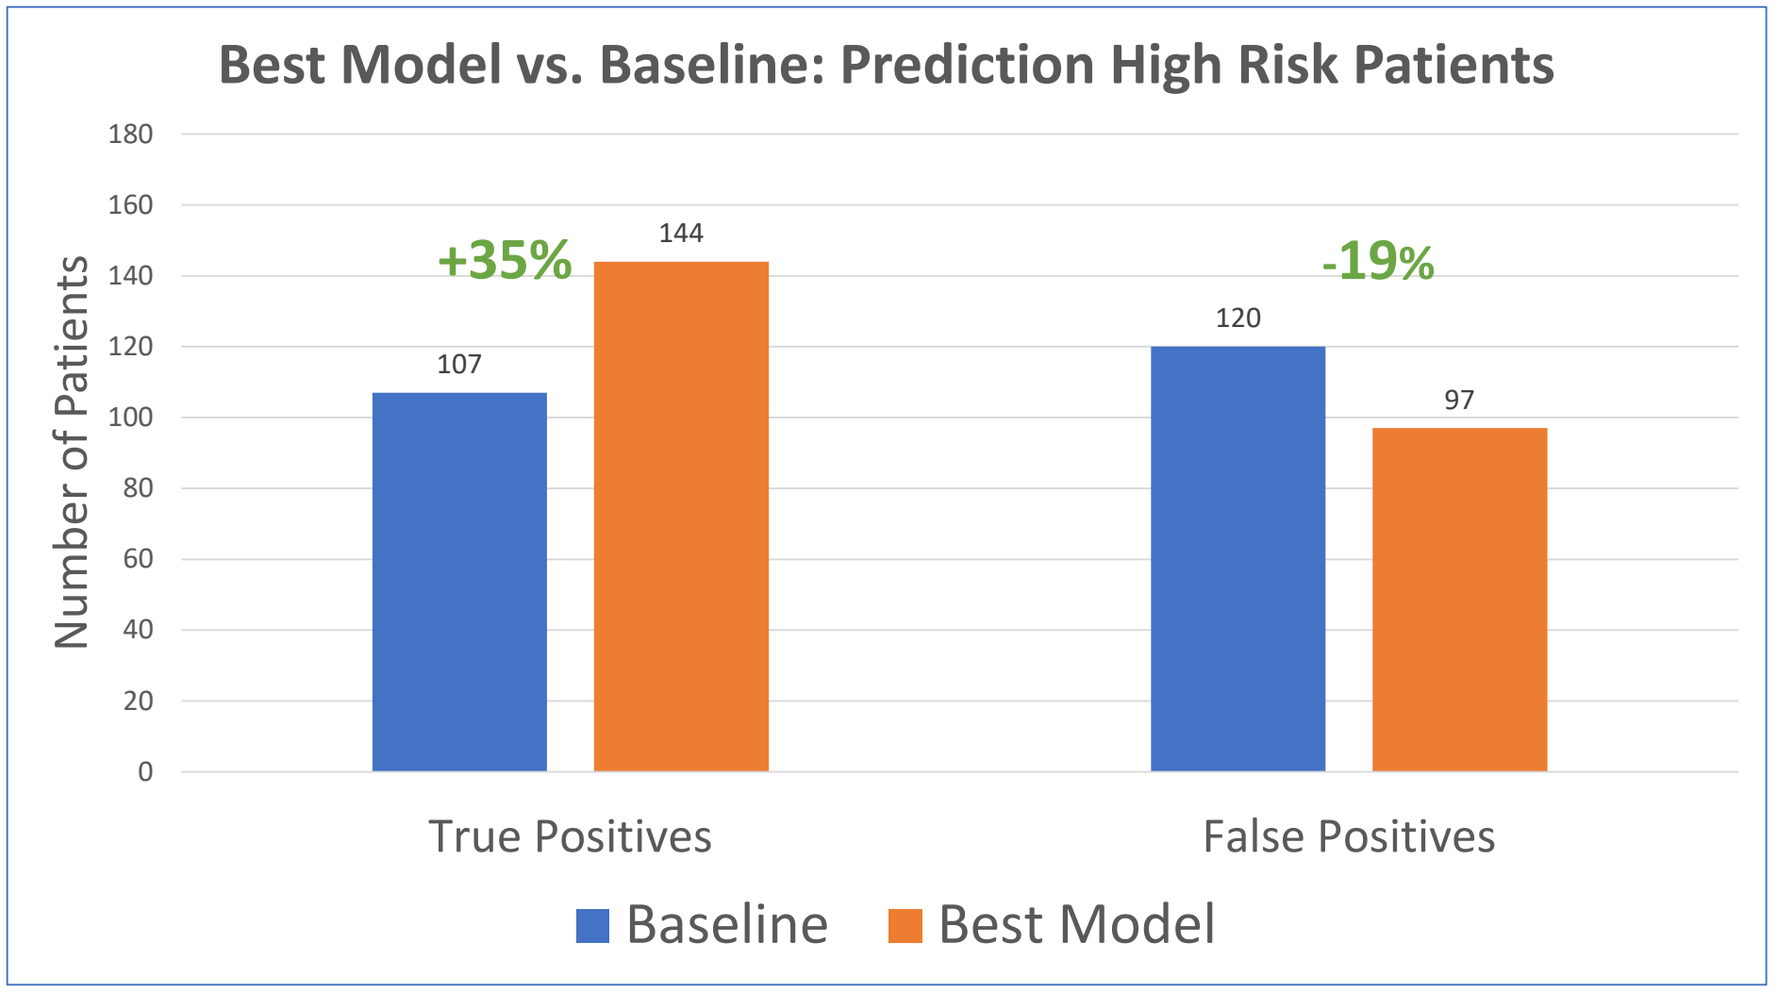

Data from

State of Maharashtra

India

Improving TB interventions
Stage by Stage Methods in Data to Deployment Pipeline


Wilder

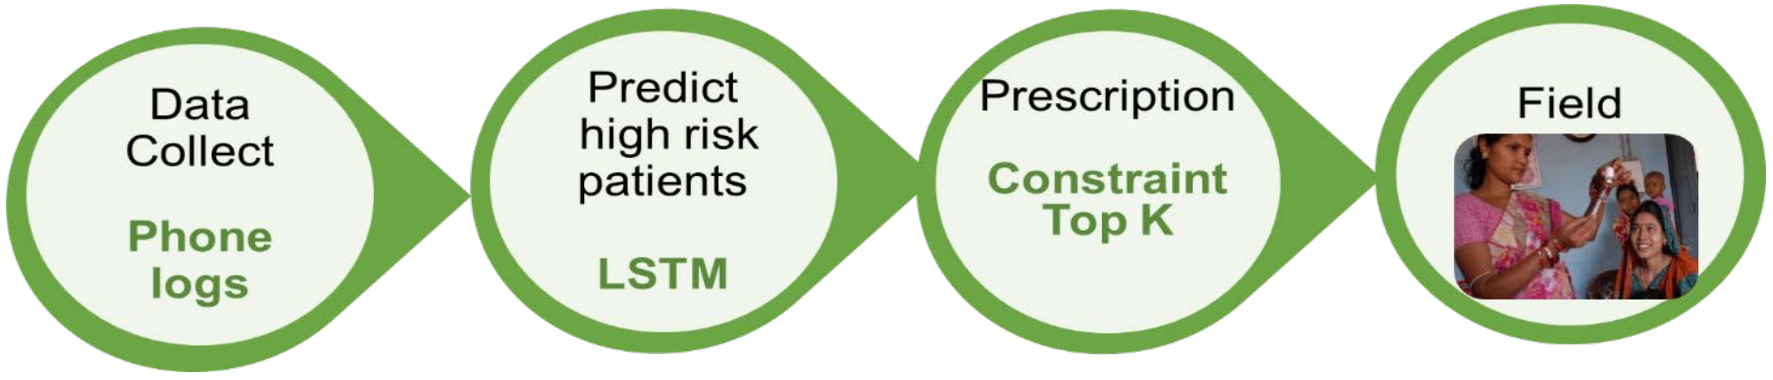

- **Maximizing accuracy ≠ Maximizing decision quality**

Predict: Machine learning

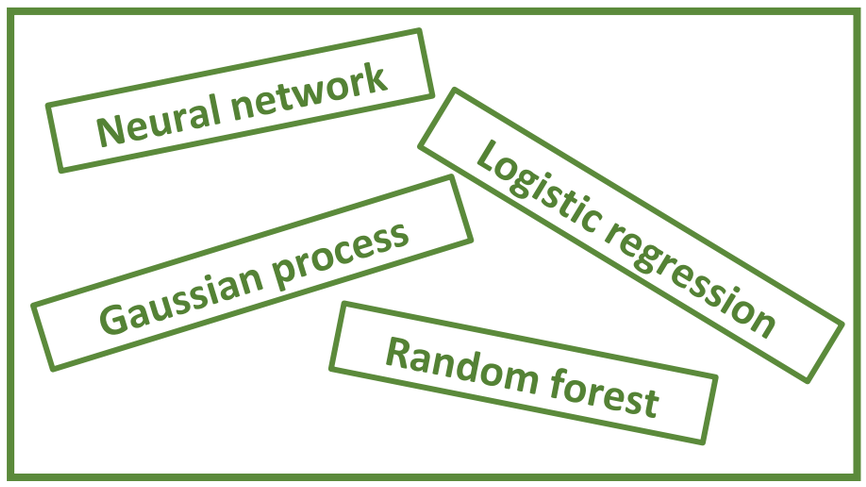

Goal: maximize accuracy

Prescription: Optimization

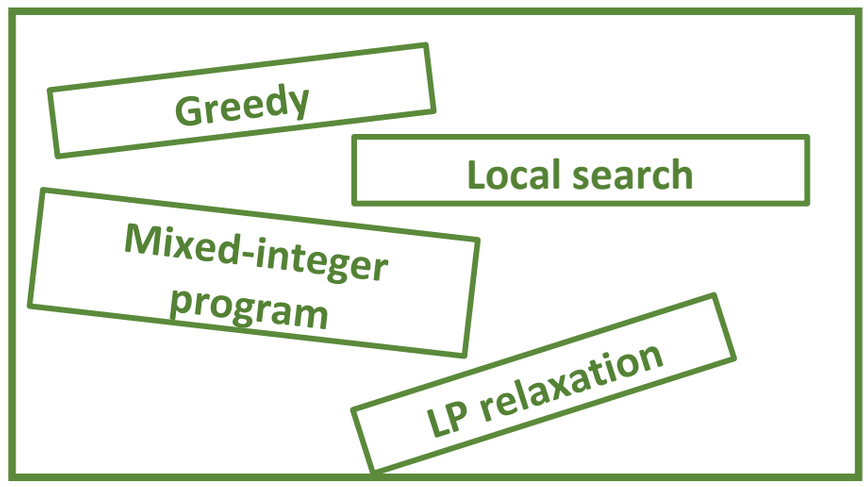

Goal: maximize decision quality

Improving TB interventions
Decision-Focused Method in Data to Deployment Pipeline


Wilder

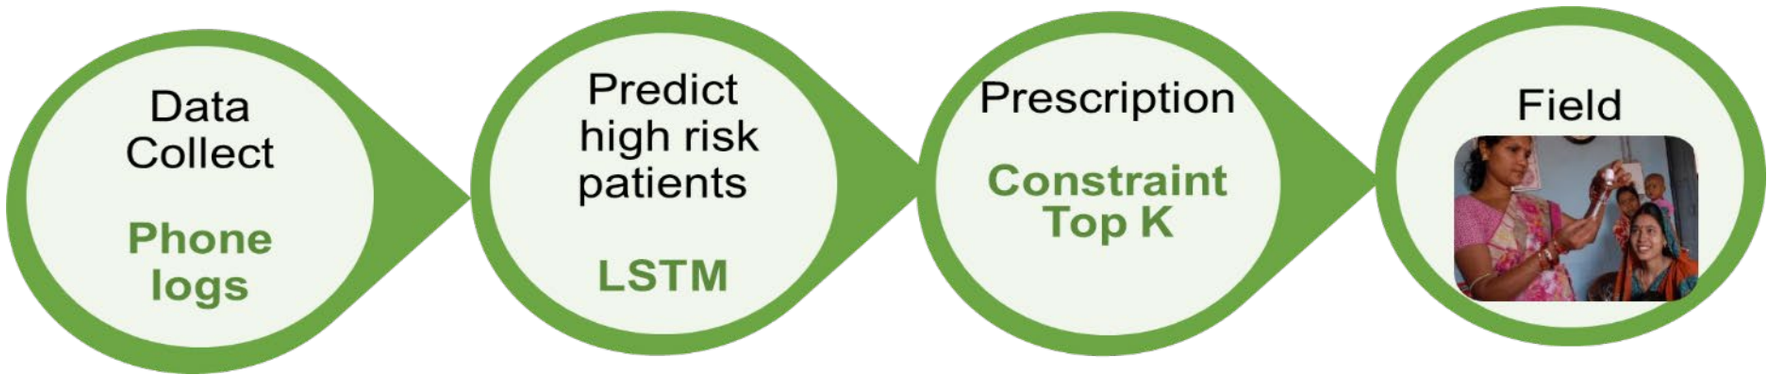

Automatically shape model loss: Optimization problem in training loop

Predict: Machine learning

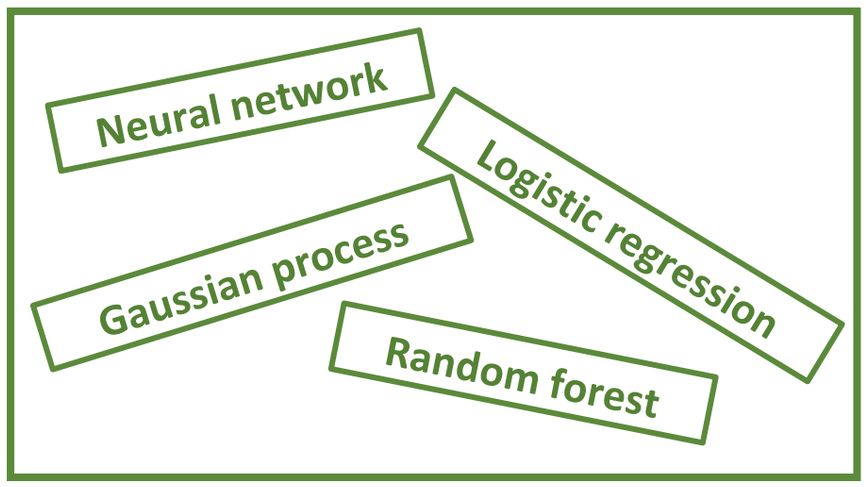

Goal: maximize accuracy

Prescribe: Optimization

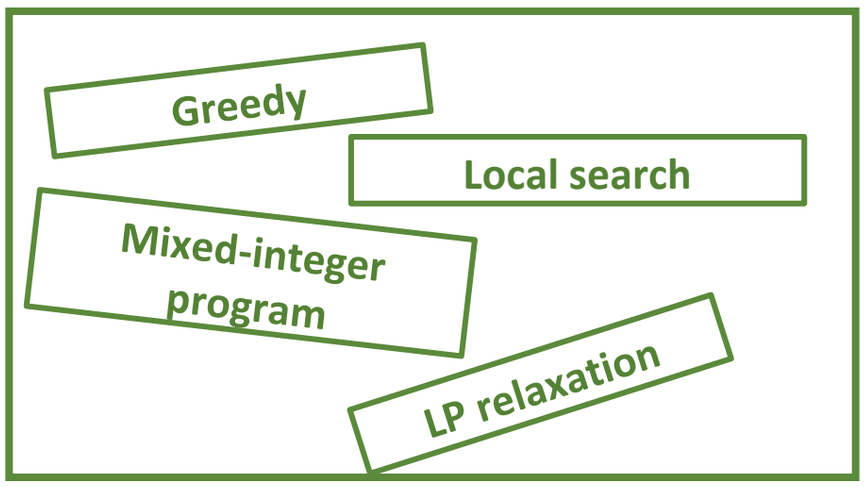

Goal: maximize decision quality

Improving TB interventions
Decision-Focused vs Stage by Stage Methods


Wilder

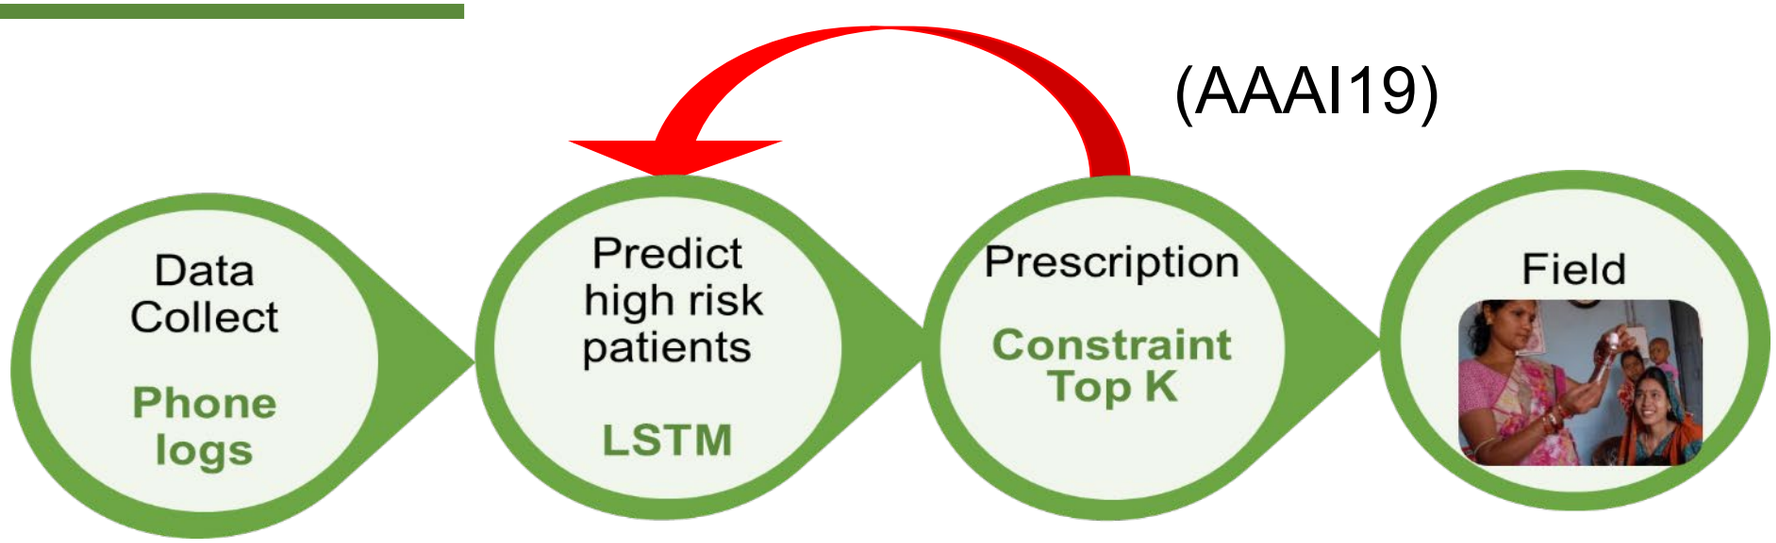

Decision focused learning improves TB interventions

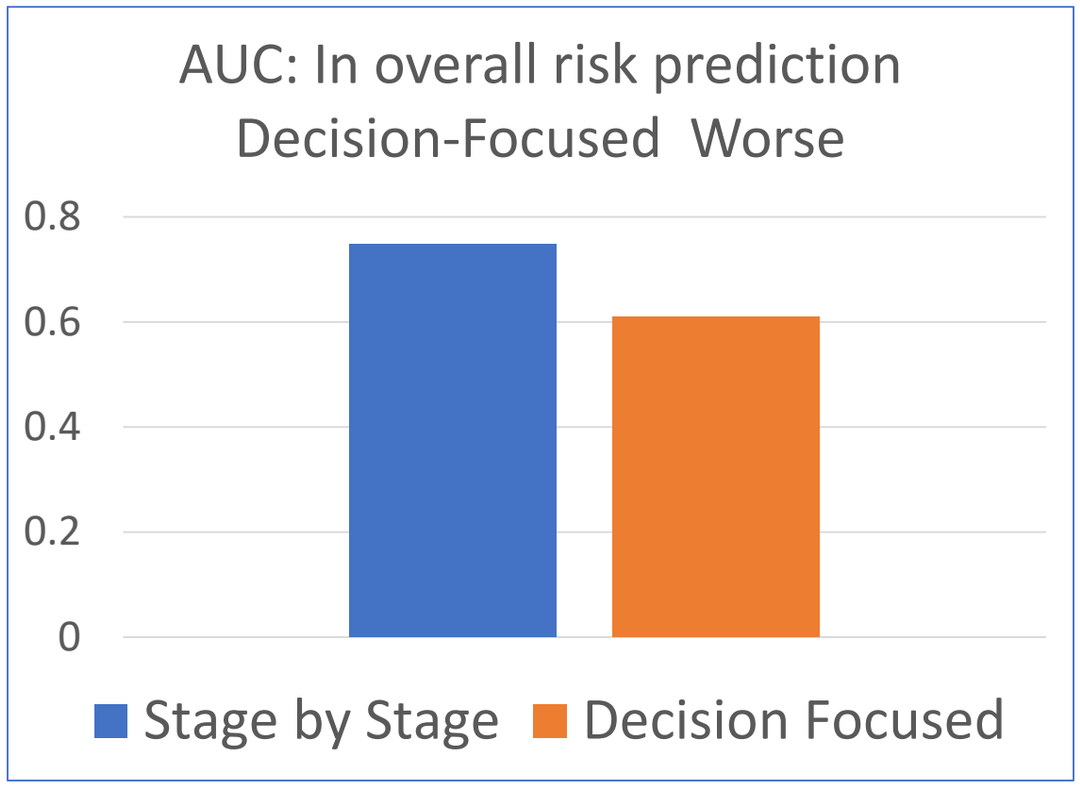

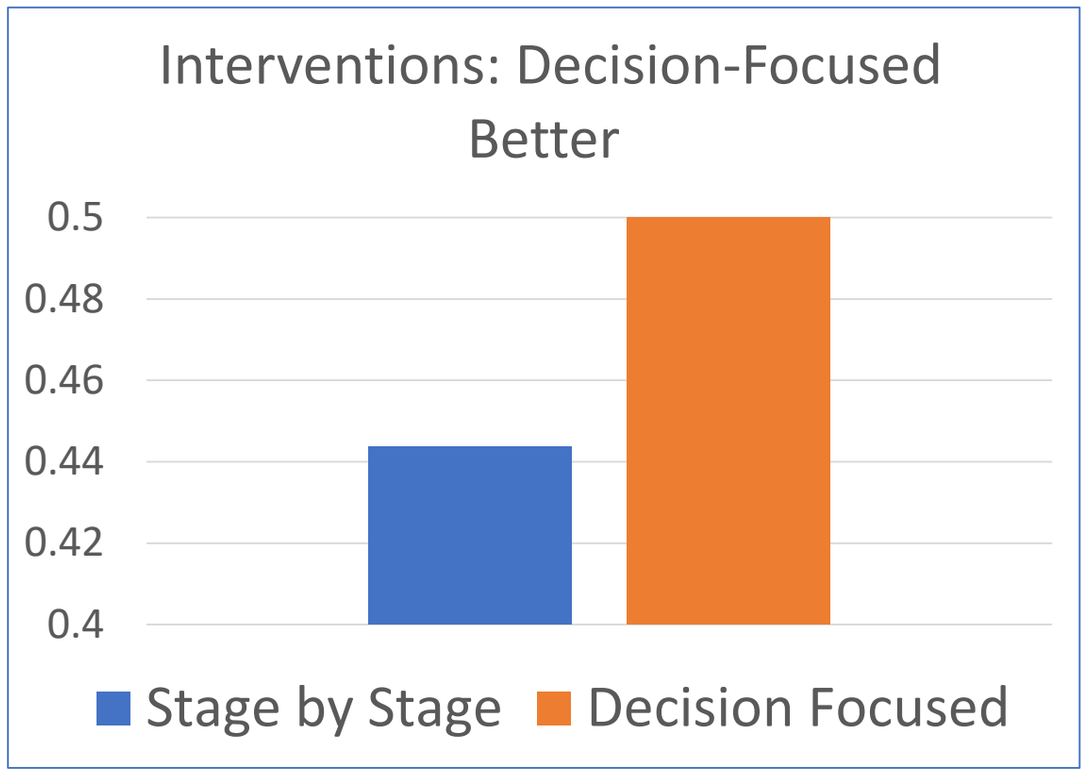

Integrating with Everwell's Platform

Killian

everwell

This work has a lot of potential to save lives.

Bill Thies

Co-founder, Everwell Health Solutions

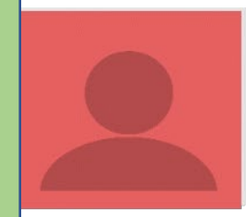

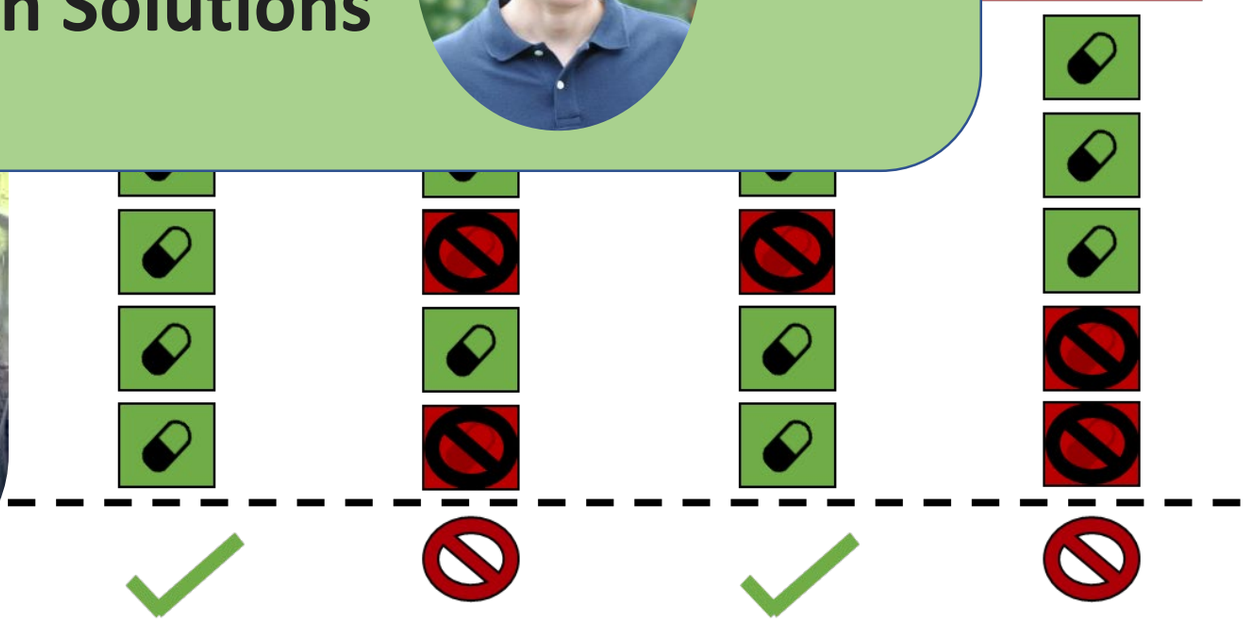

Childhood Obesity Prevention via Network Optimization


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*

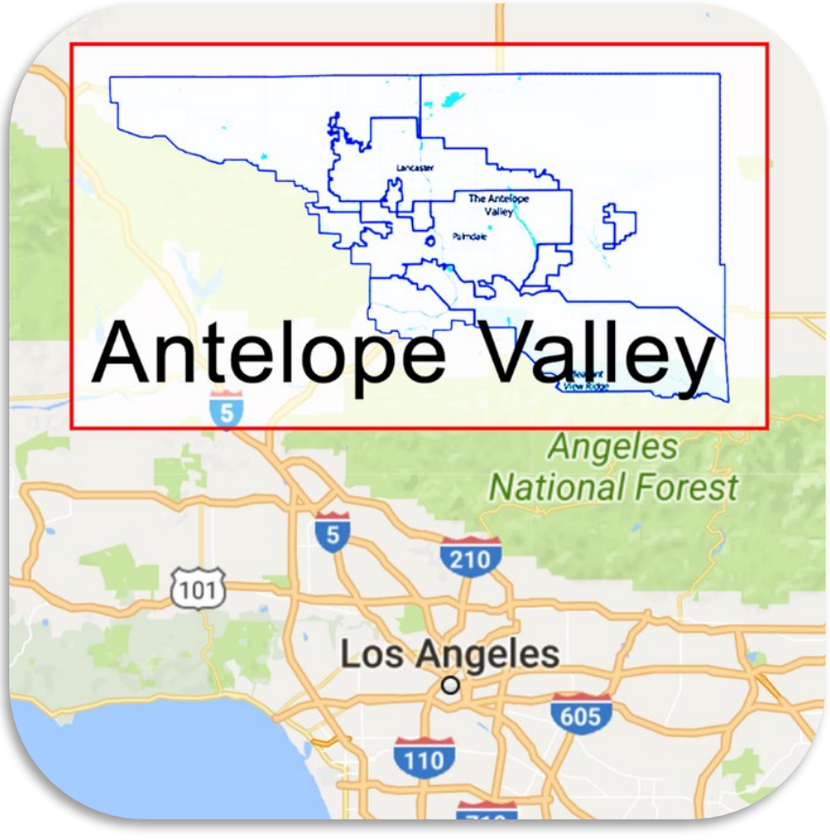

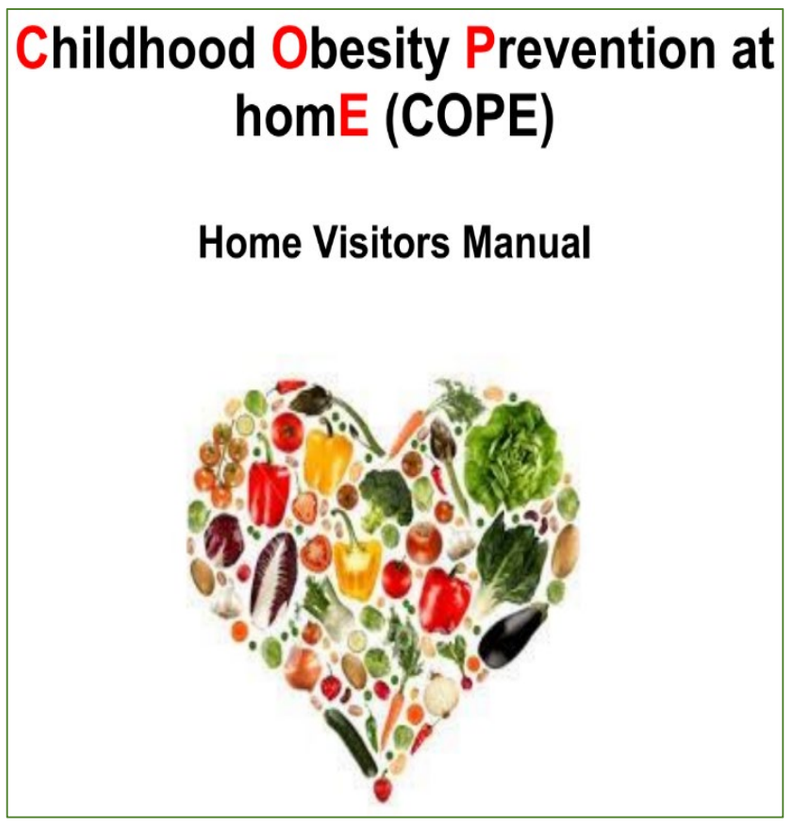

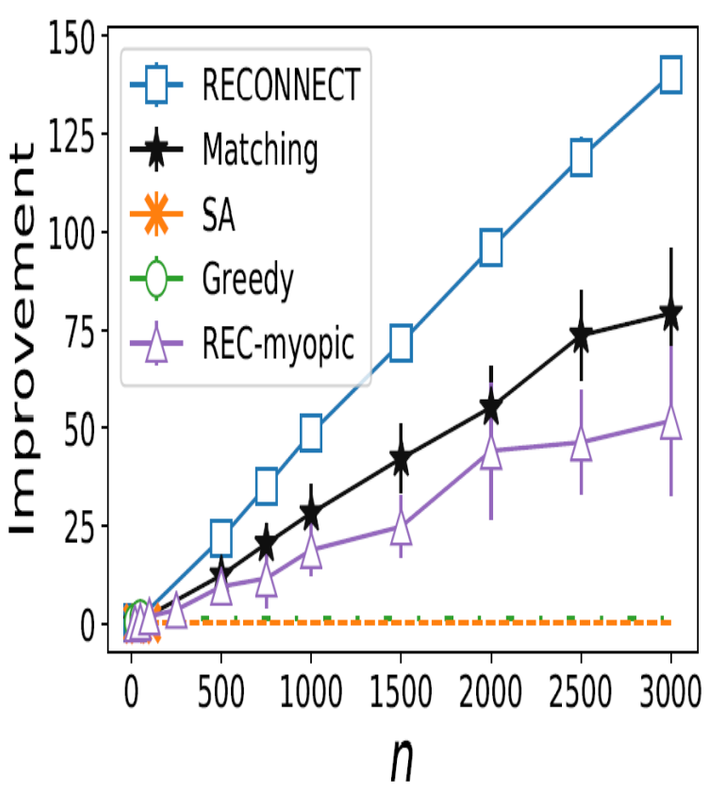

Suicide Prevention in Marginalized Populations: Choose Gatekeepers in social networks

Rahmattalabi

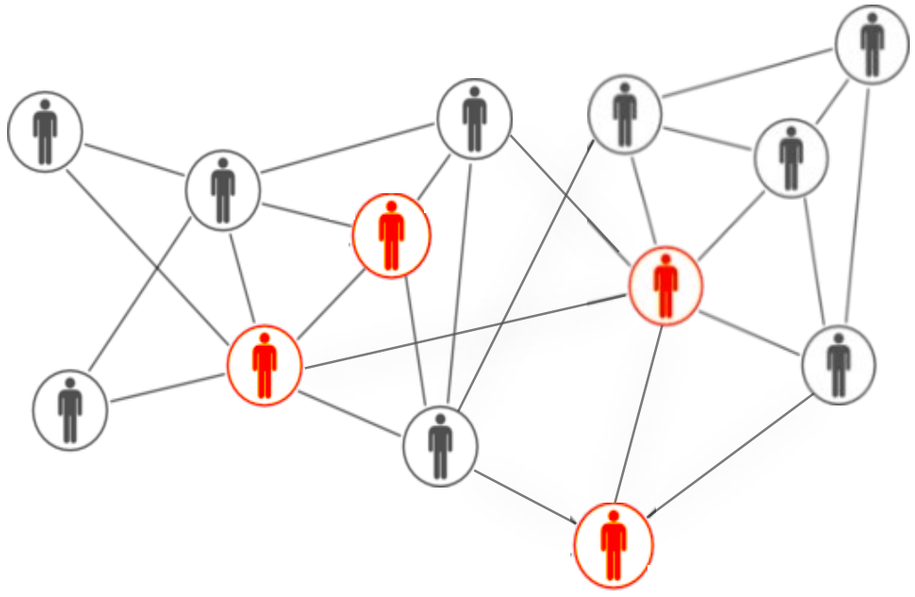

- Worst case parameters: a zero-sum game against nature

Algorithm
Chooses K gatekeepers

vs

Nature
Chooses some gatekeepers to not participate

New Directions: Los Angeles From an Angeleno [2019]

(AAMAS18)

Mayor Garcetti @ USC

New Directions: Mumbai

From a Mumbaikar [2019]

(AAAI18)

Government of Maharashtra
महाराष्ट्र शासन

Chief Minister Maharashtra
@ Mumbai
AI for Social Good

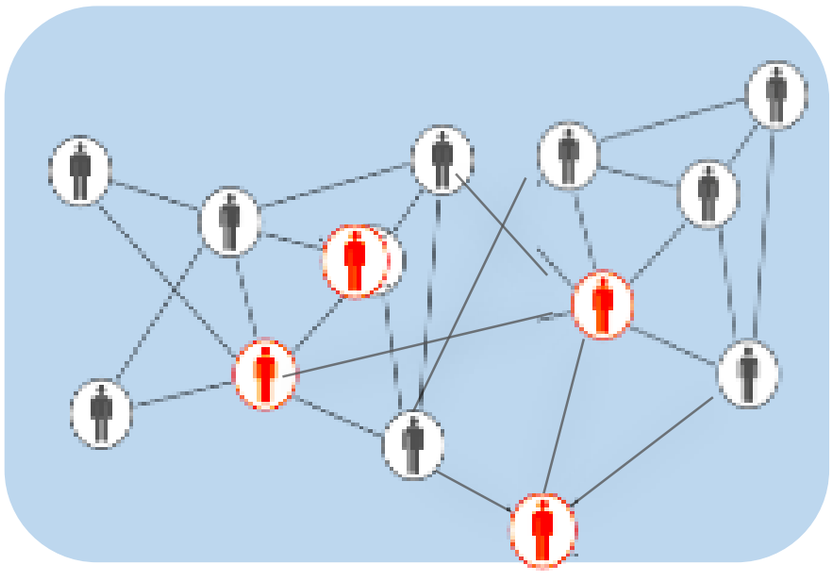

Key Lessons

Directing Multiagent Systems Research towards Social Good:

- Public safety & security, conservation, public health

Shared multiagent research challenges, solutions across problem areas:

- *Challenge*: Optimize limited intervention resources in interacting with others
- *Solution*: Computational game theory models/algorithms

Research contributions that arise from the domain:

- *Models*: Stackelberg Security Games/Green Security Games
- *Algorithms*: Incremental strategy generation, marginals, double oracle

Future: Multiagent Systems and AI Research for Social Good

Tremendous potential: Improving society & fighting social injustice

Vital to bring AI to those not benefiting from AI, e.g., global south

Embrace interdisciplinary research -- social work, conservation

Future Multiagent Systems and AI for Social Good in the FIELD

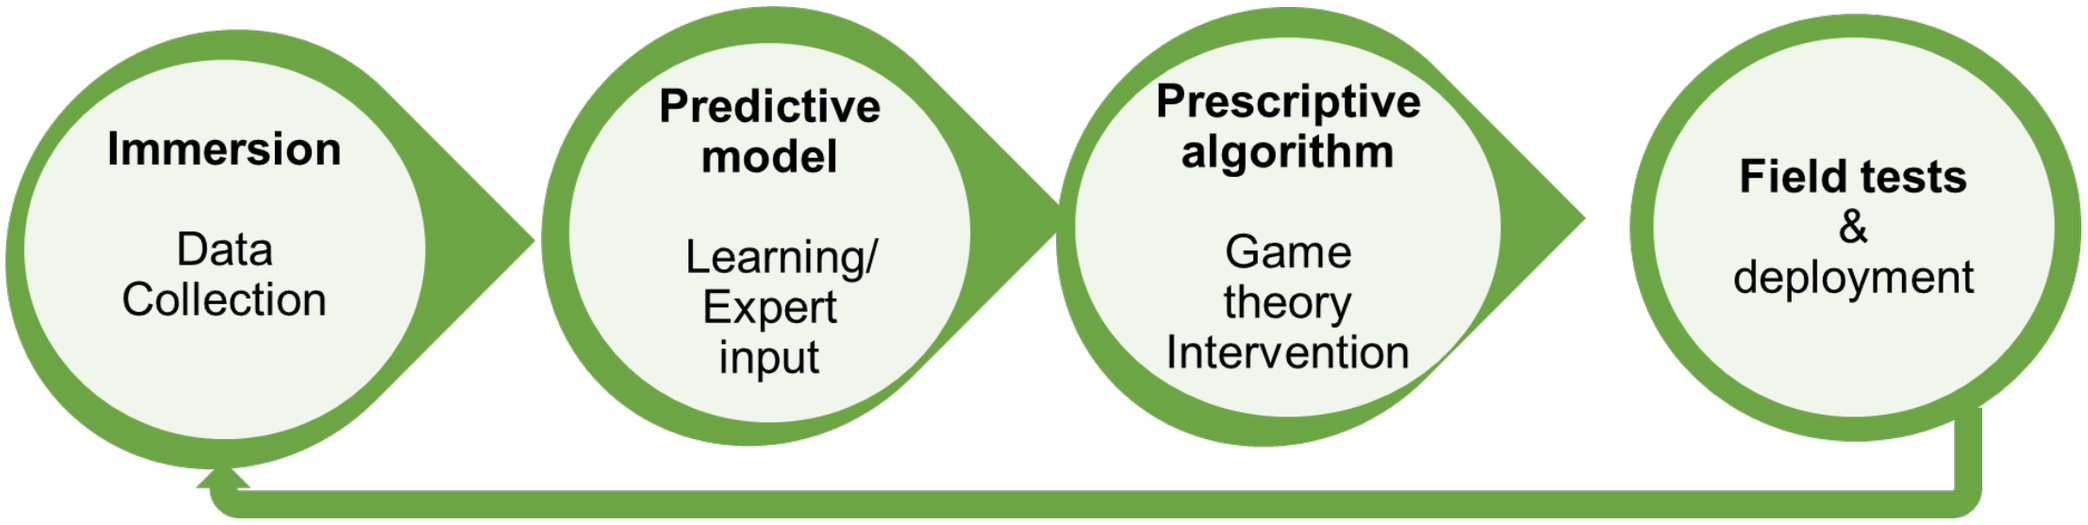

When working on AI for Societal Benefits:

Important step out of lab & into the field

- ⇨ *Societal impact*
- ⇨ *Model deficiencies for new research*

Thank you for Inspiring Us

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

Transportation Security Administration

Wildlife Conservation Society

WWF

SMART

Air Shepherd
The Lindbergh Foundation

LOS ANGELES LGBT CENTER

S.P.Y
safe place for youth

my friend's place

everwell

THANK YOU

@MilindTambe_AI

CAIS.USC.EDU