AI for Social Good:
Decision aids for Countering Terrorism, Extinction, Homelessness

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

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 (Messaging) Resources

Public Health Awareness:
Influence Maximization as a Game against Nature


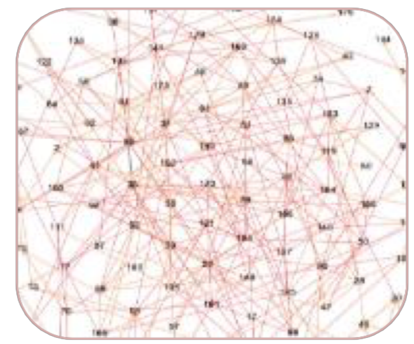

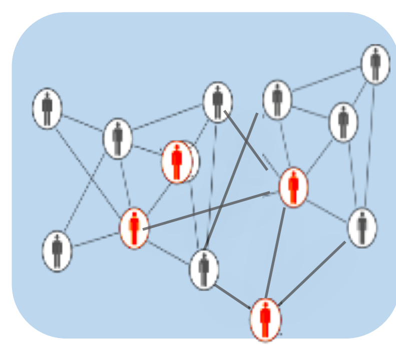


- Social networks to enhance intervention, e.g., HIV information
- Real-world pilot tests: Homeless youth shelters in Los Angeles

Overall Research Framework, Partnerships and Publications

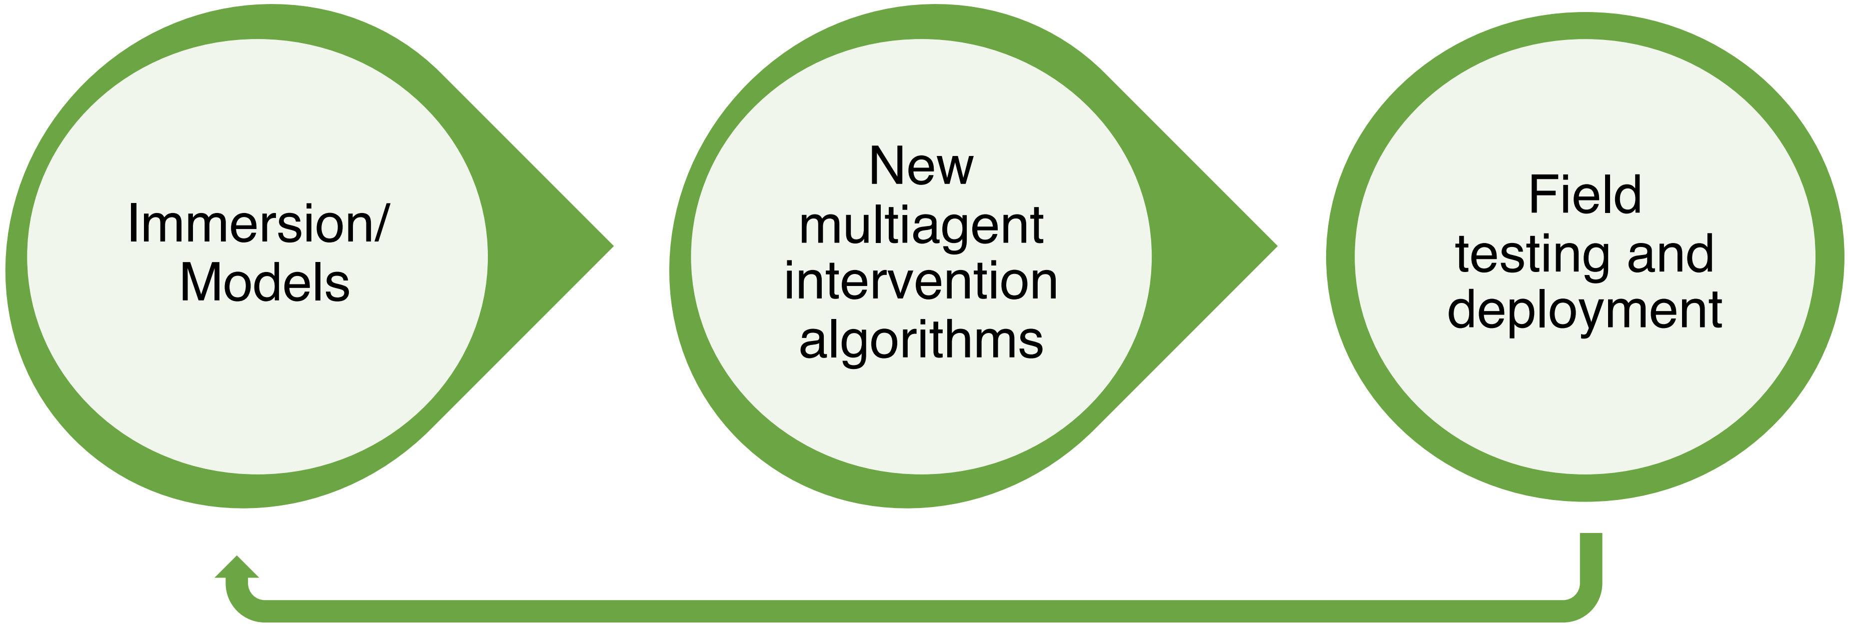

Outline

Public Safety and Security
Stackelberg Security Games

Conservation/Wildlife Protection
Green Security Games

Public Health
Influence maximization/Game against nature

- **AAMAS,AAAI,IJCAI evaluation + Real world evaluation**
- **PhD students and 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, key aspects for tractability

Set of targets, payoffs based on targets covered or not

Stackelberg Leader-Follower formulation

Adversary

Defender

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

Model: Stackelberg Security Games

Stackelberg: Defender commits to randomized strategy, adversary responds

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

Challenges faced: Massive scale games

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	….	….	….

+ Handling Uncertainty

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

Maximize defender expected utility

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

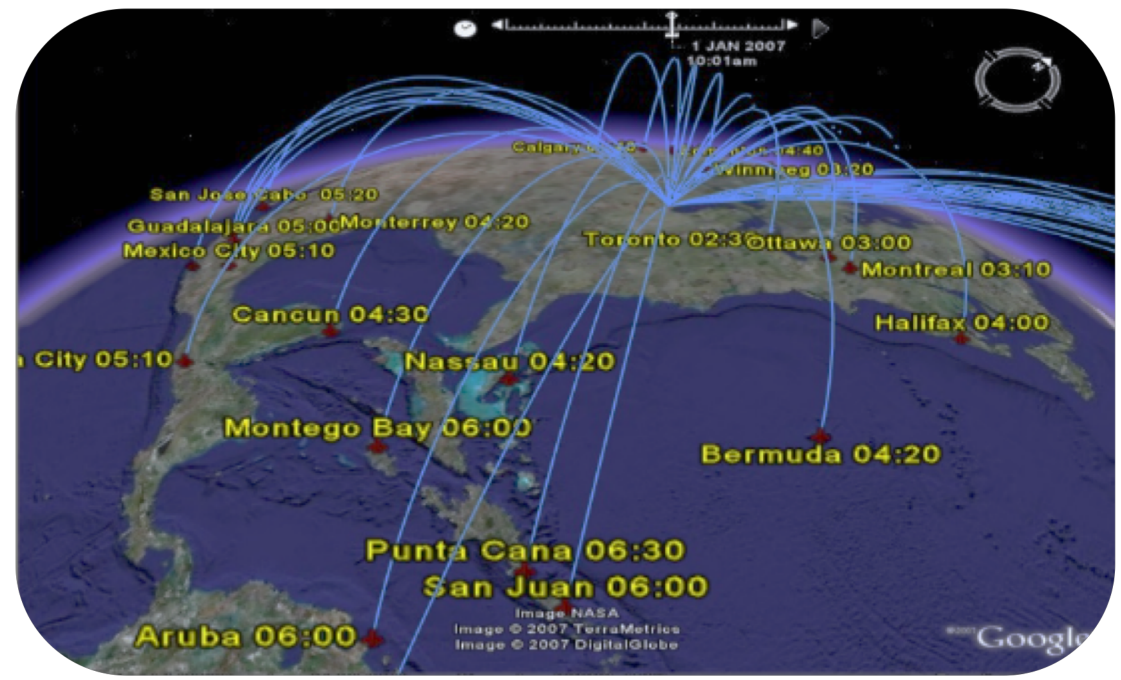

<table>
<tr><th></th><th>Strategy 1</th><th>Strategy 2</th><th>Strategy 3</th><th>Strategy 4</th></tr>
<tr><td>Strategy 1</td><td colspan="4" rowspan="4">IRIS 1000 flights/day
Actions: ~10^{41}</td></tr>
<tr><td>Strategy 2</td></tr>
<tr><td>Strategy 3</td></tr>
<tr><td>Strategy 4</td></tr>
</table>

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

Theorem: For T targets, solutions exist where support set size is T+1

Small support set size:
Most x_i variables zero

1000 flights, 20 air marshals:

10^{41} combinations

		Attack 1	Attack 2	Attack …	Attack 1000
$\cancel{X_{123} = 0.0}$	~~1 ,2, 3 ..~~	~~5,-10~~	~~4,-8~~	~~…~~	~~-20,9~~
$X_{124} = 0.239$	1, 2, 4 ..	5,-10	4,-8	…	-20,9
$\cancel{X_{135} = 0.0}$	~~1, 3, 5 ..~~	~~5,-10~~	~~-9,5~~	~~…~~	~~-20,9~~
$X_{378} = 0.123$	…				
	… ←	**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

Theory)
e strategy

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

GLOBAL OPTIMAL
1000 defender strategies
NOT 10^{41}

IRIS: Deployed FAMS [2009-]

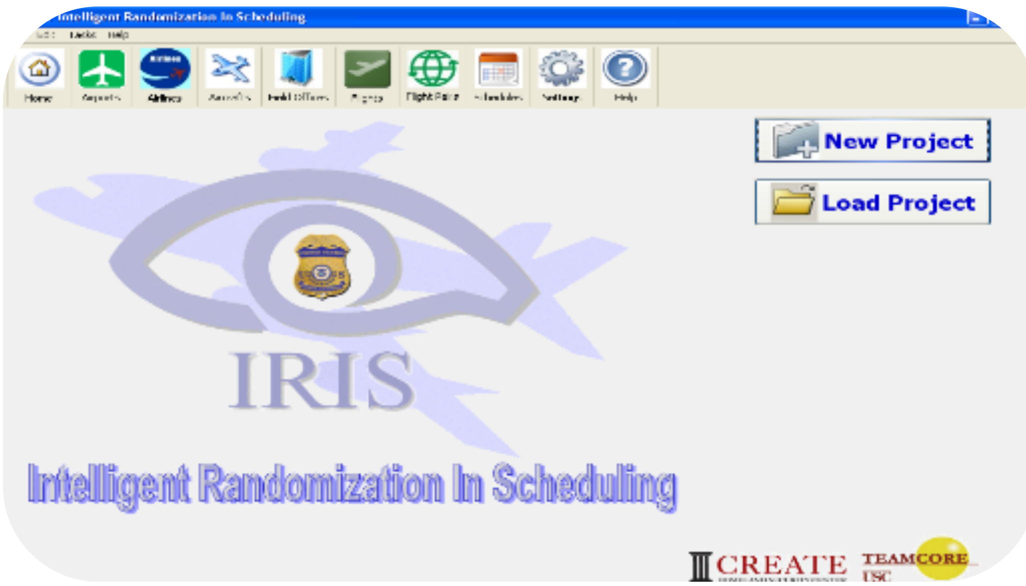

Significant change in FAMS operations

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)

26 Nov 2008, Mumbai
Police Checkpoints: Network Security Game

Jain

Road networks:

20,000 roads, 15 checkpoints

150 edges
2 Checkpoints
150-choose-2 strategies

Zero-Sum Network Security Game [2013]

Jain

Double oracle: New exact optimal algorithm for scale-up

	Path #1	Path #2	Path #3
Checkpoint strategy #1	5, -5	-1, 1	-2, 2
Checkpoint strategy #2	-5, 5	1, -1	-2, 2

Defender oracle

	Path #1	Path #2
Checkpoint strategy #1	5, -5	-1, 1
Checkpoint strategy #2	-5, 5	2, -1

Attacker oracle

	Path #1	Path #2	Path #3
Checkpoint strategy #1	5, -5	-1, 1	-2, 2
Checkpoint strategy #2	-5, 5	1, -1	-2, 2

Presentation at the Indian National Police Academy: Network Security Game [2016]

Road networks:

20,000 roads, 15 checkpoint:
Solved under 20 min

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
Transition Graph Representation

Fang

Jiang

Marginal strategy: New scale-up approach for Stackelberg Security Games

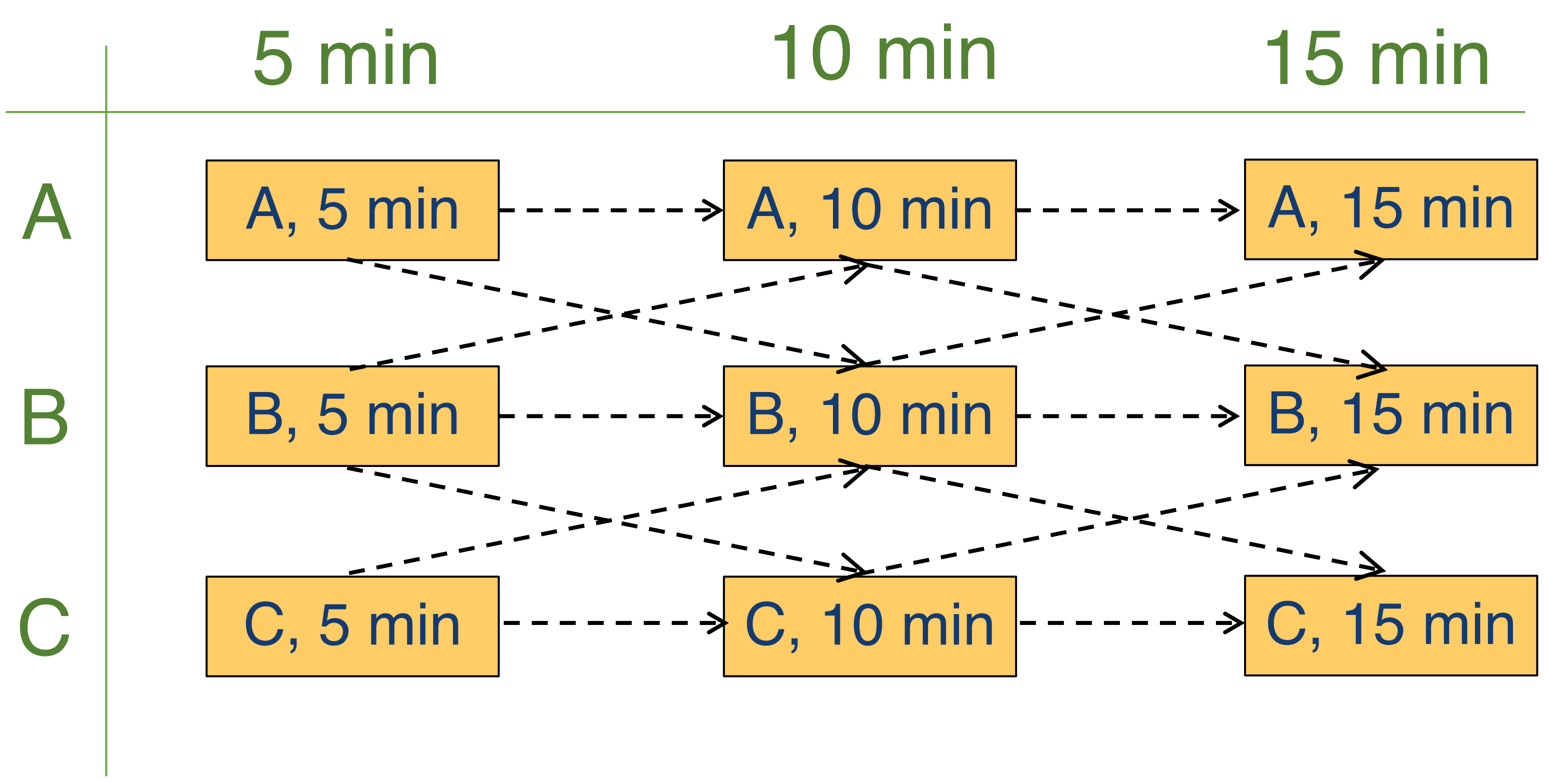

FERRIES: Mobile Resources & Moving Targets
Transition Graph Representation


Fang


Jiang

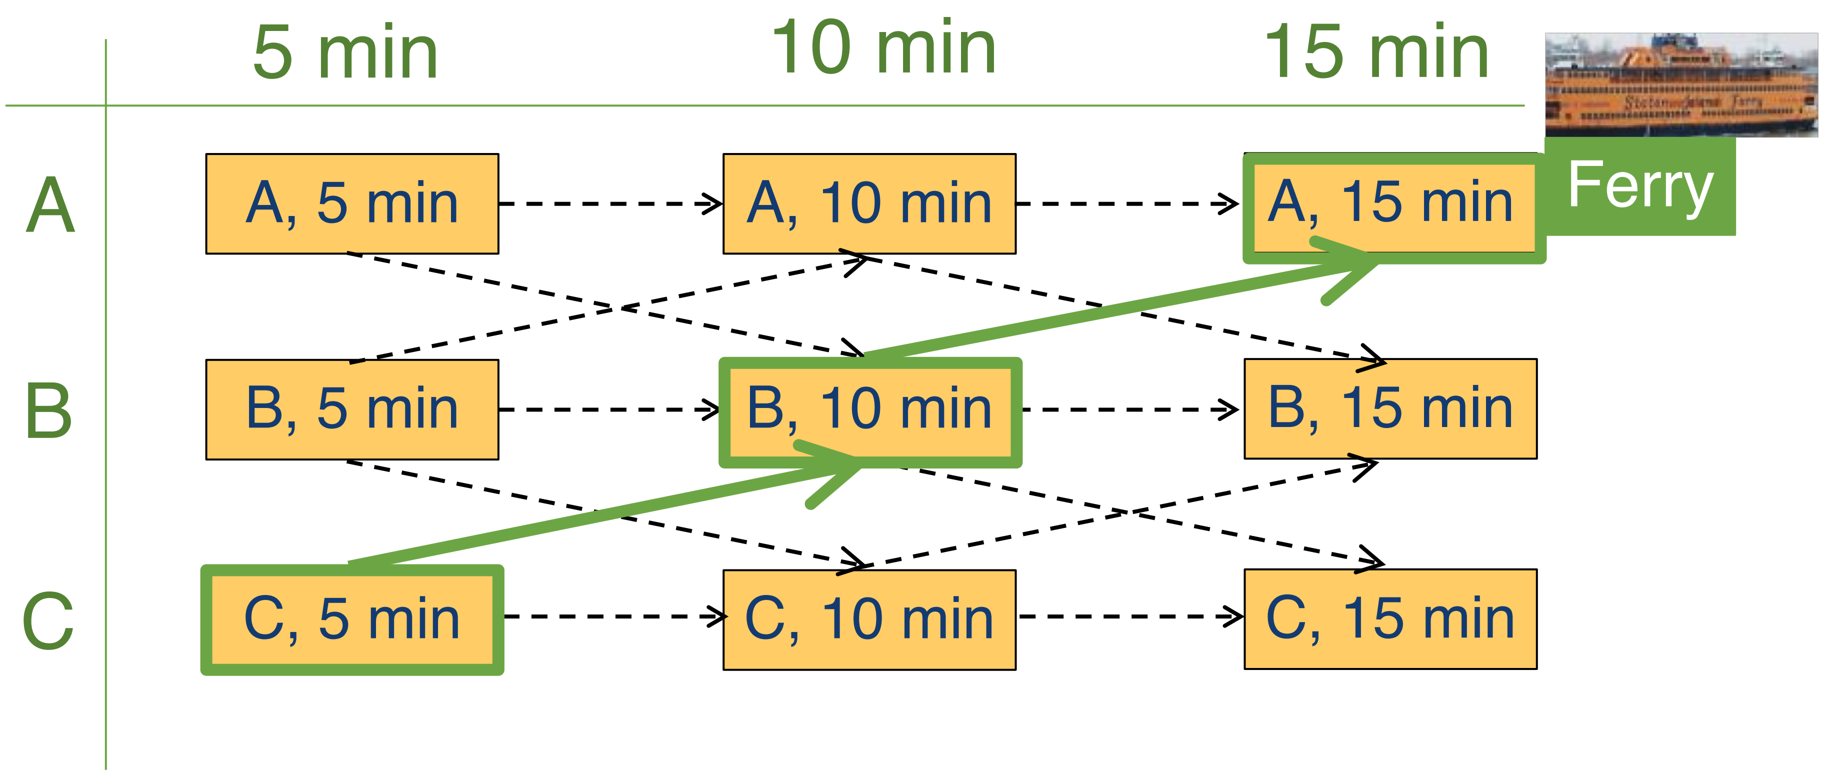

FERRIES: Mobile Resources & Moving Targets Transition Graph Representation

Fang

Jiang

Patrol protects nearby ferry locations

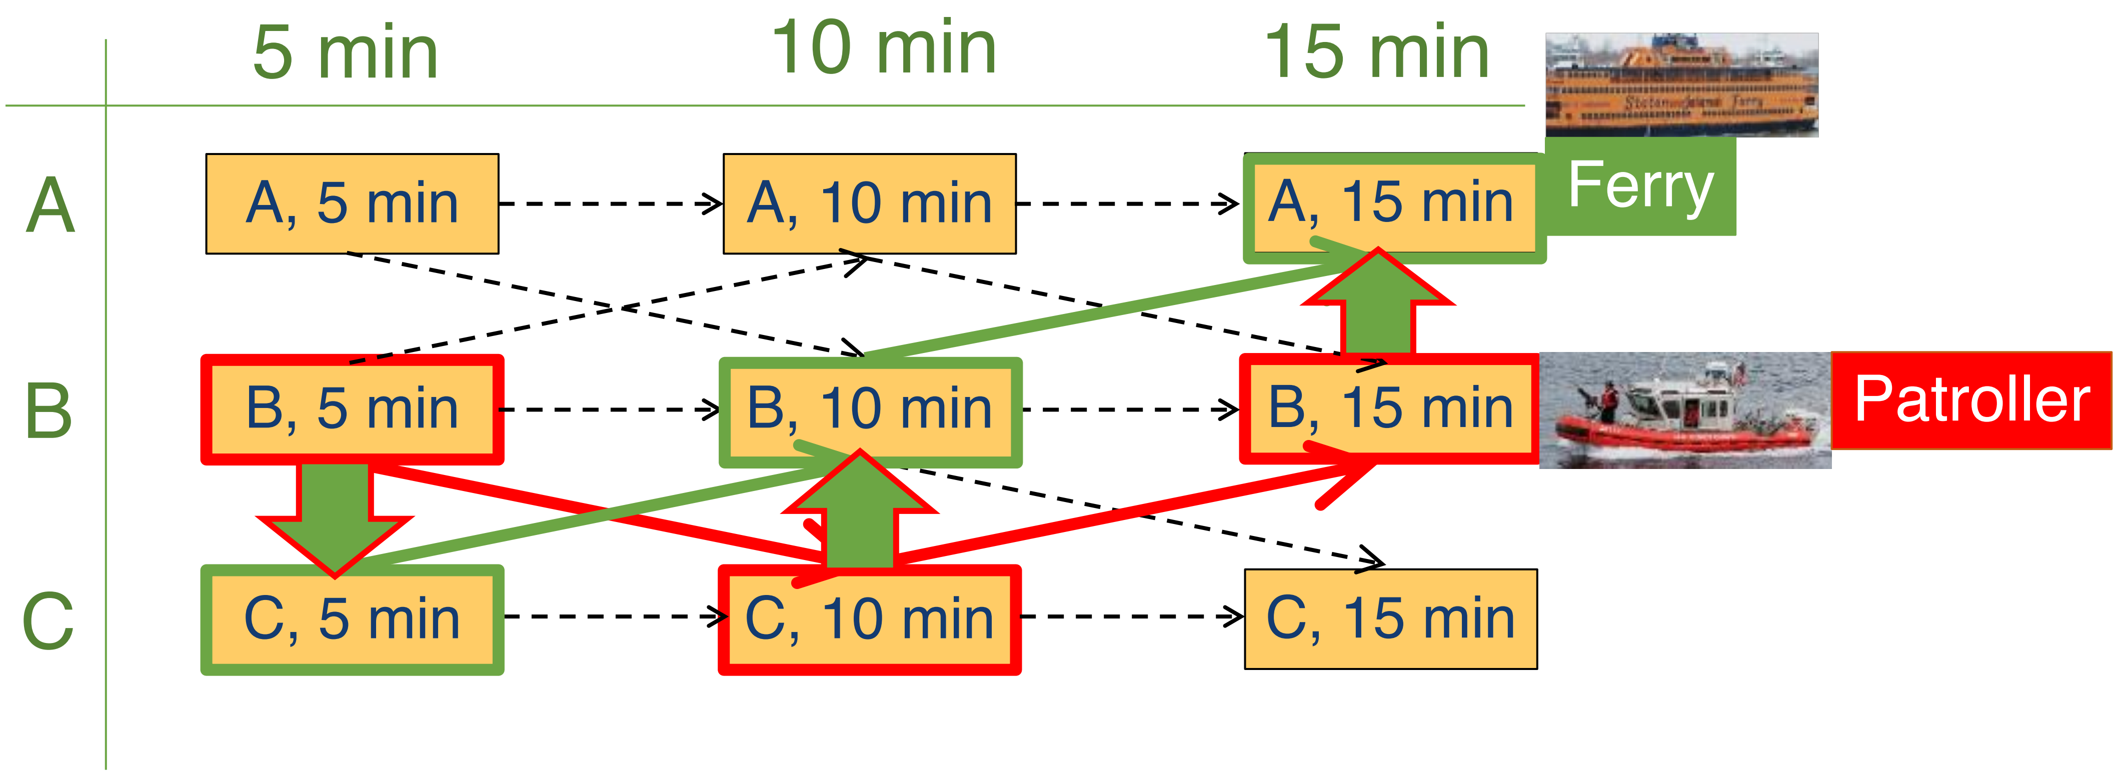

FERRIES: Mobile Resources & Moving Targets Transition Graph Representation


Fang


Jiang

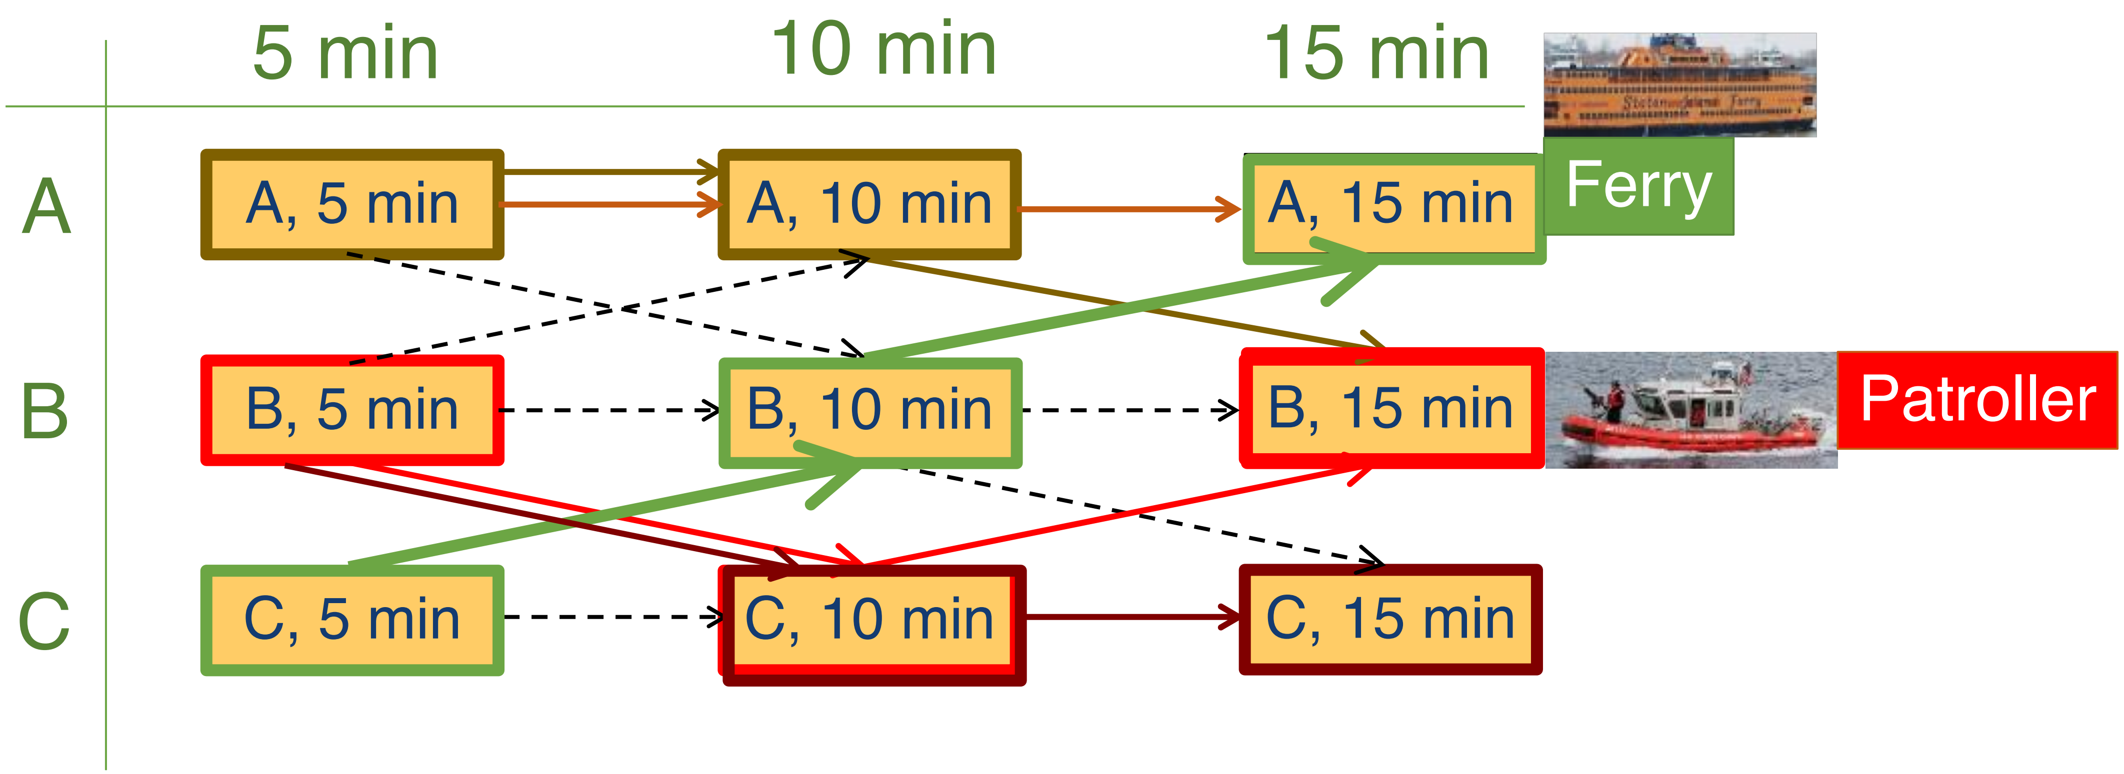

FERRIES: Mobile Resources & Moving Targets Transition Graph Representation

Fang

Jiang

ARMOR style LP: Determine probability for each route

Defender

N^T variables

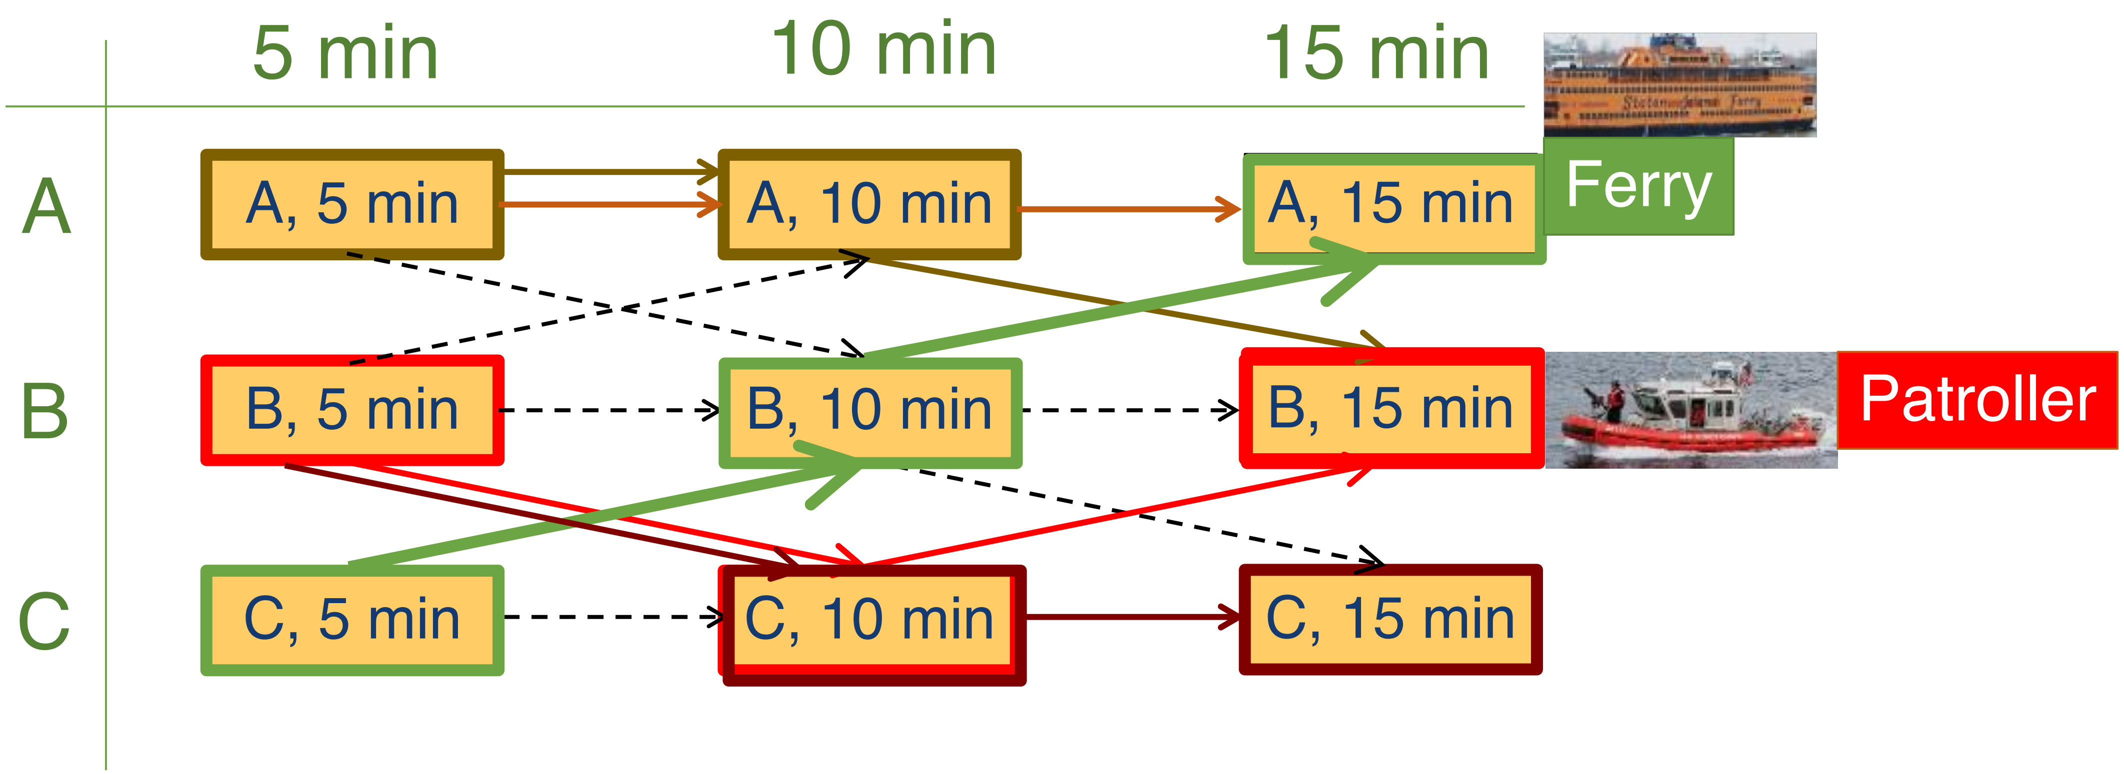

FERRIES: Scale-Up
Transition Graph Representation

Variables: NOT routes, but marginal probability over each segment

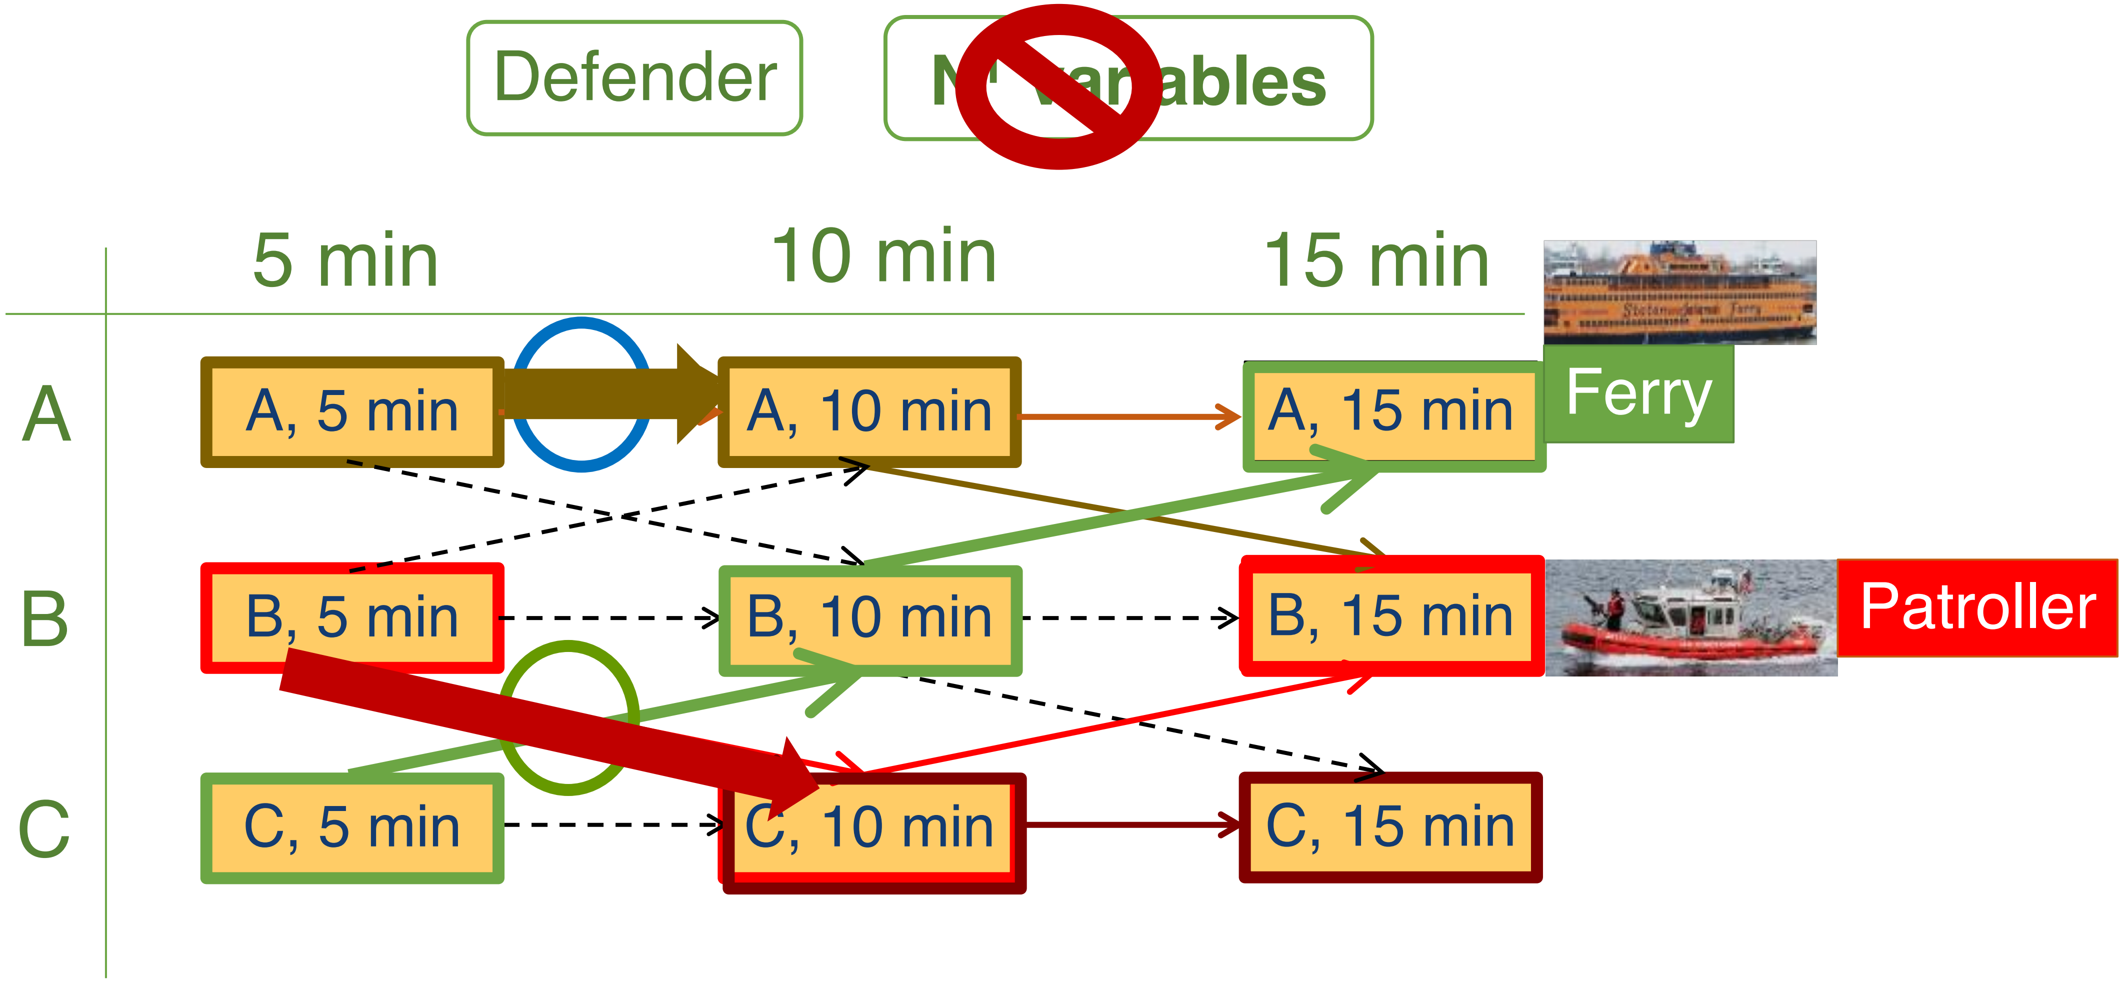

FERRIES: Scale-Up
Transition Graph Representation

Theorem: Marginal representation does not lose any solution quality

Defender

$N^2.T$ variables

~~**N^T variables**~~

Extract: *Pr([(B,5), (C, 10), (C,15)]) = 0.47*
Pr([(B,5), (C,10), (B,15)]) =0.23

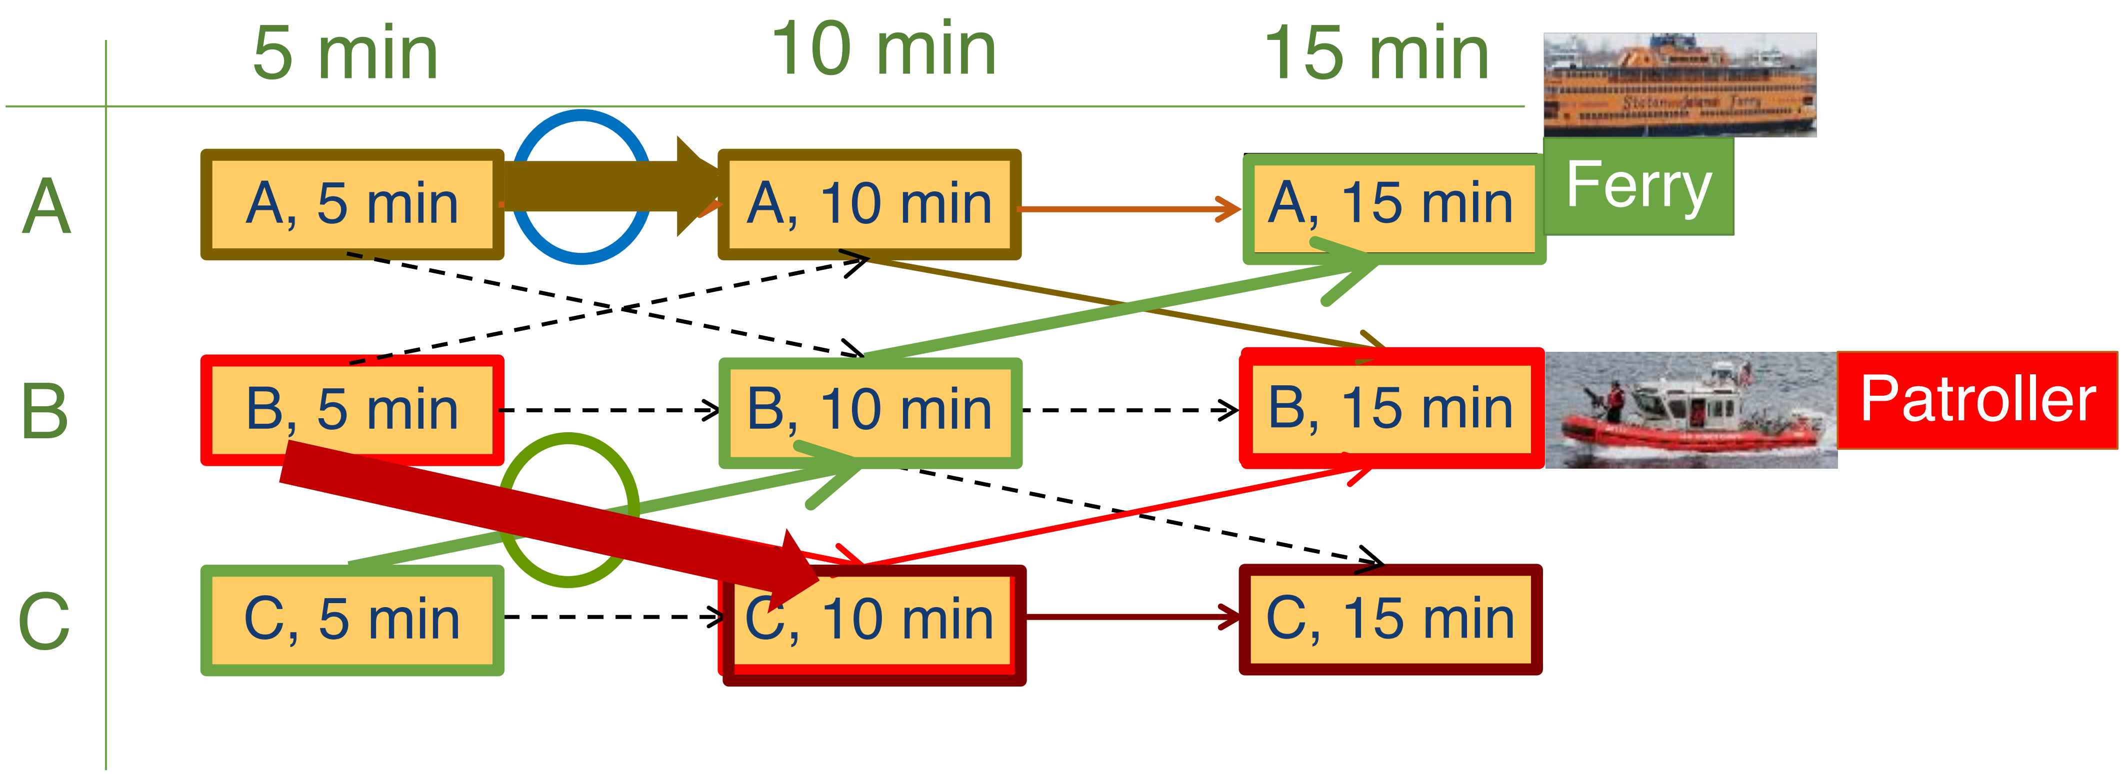

PROTECT: Port Protection Patrols [2013]
Congressional Subcommittee Hearing

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

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

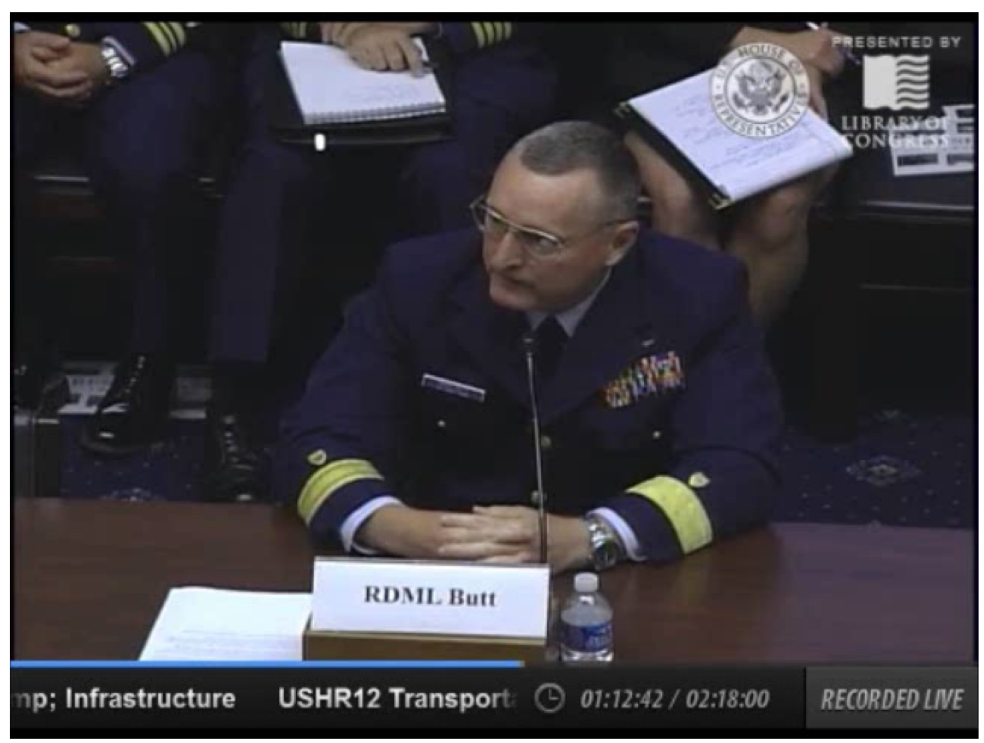

Train Patrols
Execution Uncertainty: MDPs

Jiang

Delle Fave

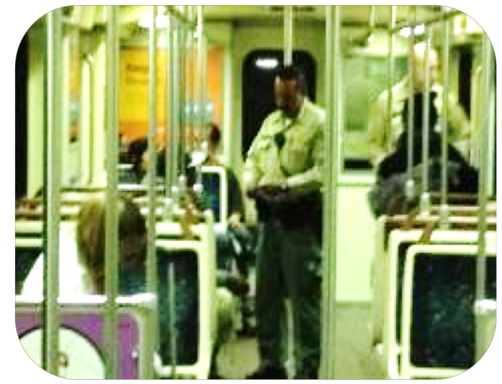

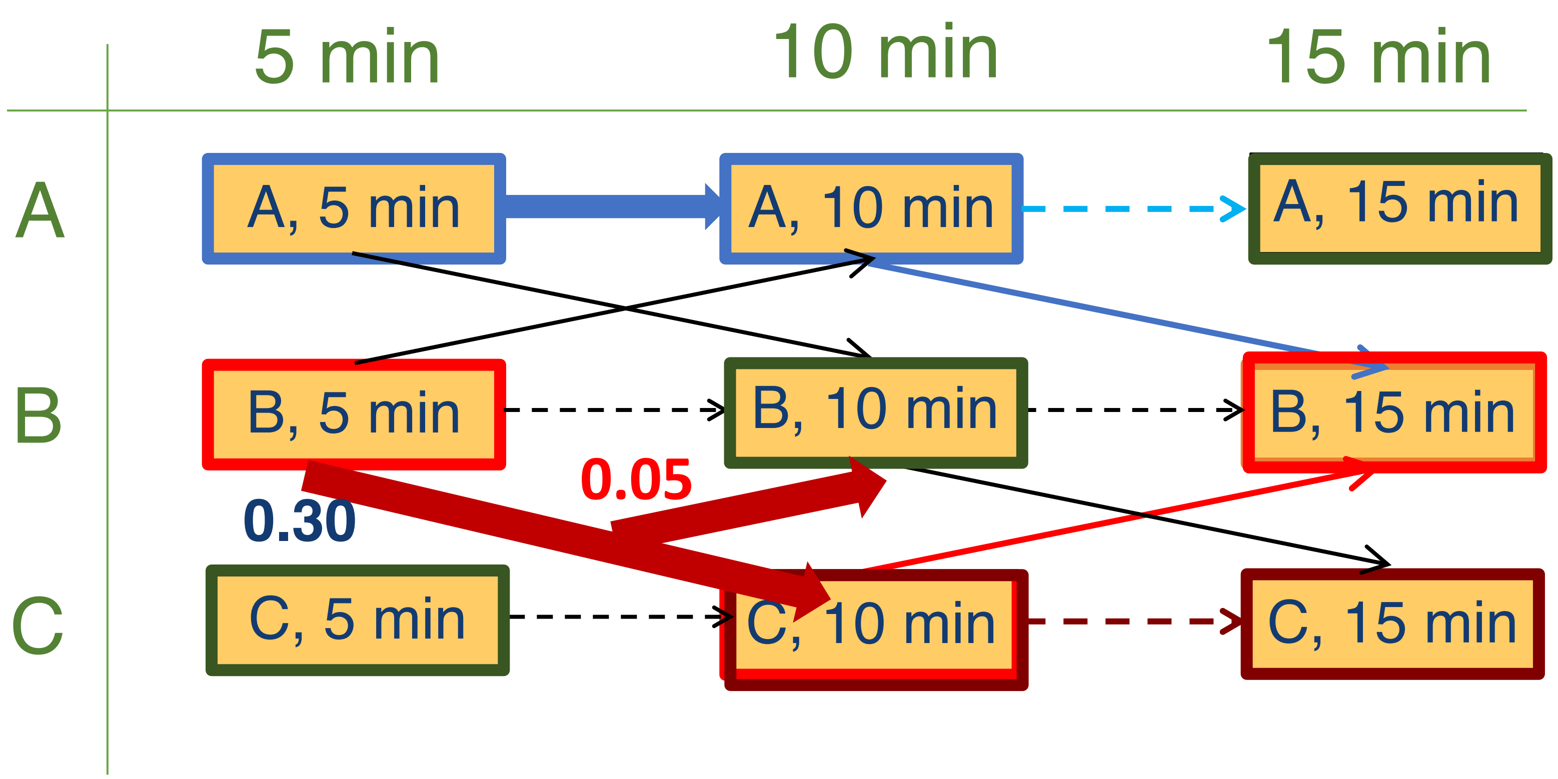


Nguyen


An

Handling Payoff Uncertainty: Optimal Defender Strategy Minimizing Max Regret

- **Payoff uncertainty**

Adversary

Defender

	Target #1	Target #2
Target #1	4, **[-4,-2]**	-1, **[0,2]**
Target #2	-5, **[4,6]**	2, **[-2,0]**

- **DefenderUtility(c): -2.3**
- **Optimal utility: 0.4**
- **Regret (c, payoff): 2.7**

Adversary

Defender

	Target #1	Target #2	c
Target #1	4, -3	-1, 1	**0.3**
Target #2	-5, 5	2, -2	**0.7**

Minimizing Maximum Regret: New Iterative Constraint Generation Algorithm


Nguyen An

Infinite #regret constraints

$$\min_{c,\, r} r$$
$$\text{where } r \geq \text{regret}(c, \text{payoff}), \forall \text{payoff} \in \text{Interval}$$

Master: Compute Lower Bound
Minimax Regret with sample set of attacker payoffs

payoff

c*

Slave: Compute Upper Bound
New attacker payoff causing max regret

Global Presence of Security using Game Theory [2015-2017]

Evaluating Deployed Security Systems Not Easy

How Well Optimized Use of Limited Security Resources?

Security Games superior
vs
Human Schedulers/"simple random"

- Lab evaluation
- *Scheduling competitions: Patrol quality unpredictability? Coverage?*
- Field evaluation: Tests against real adversaries
- *Economic cost-benefit analysis*
- *…*

Field Evaluation of Schedule Quality

Improved Patrol Unpredictability & Coverage for Less Effort

Patrols Before PROTECT: Boston

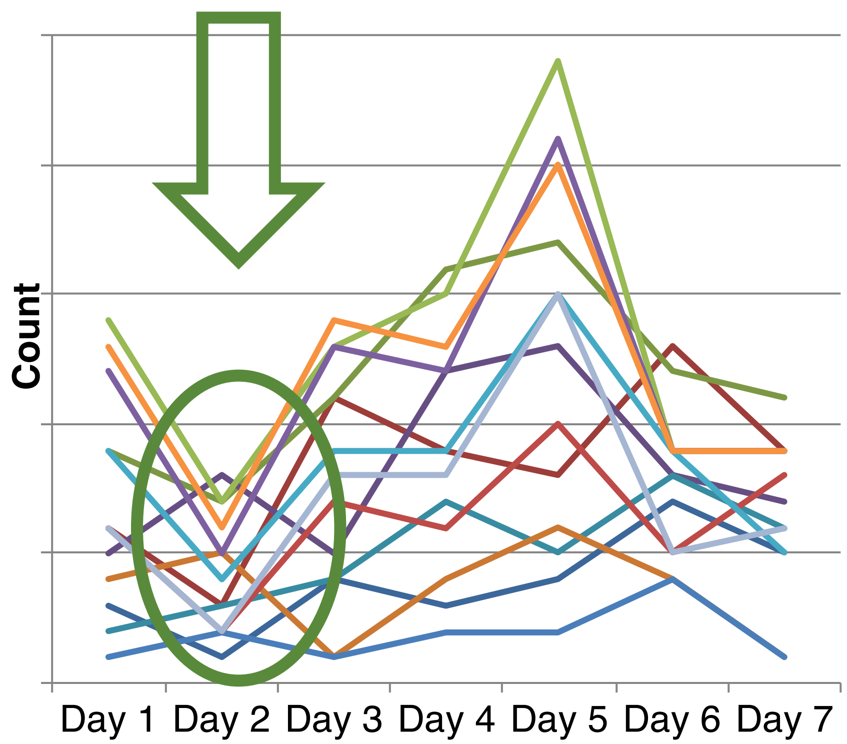

Patrols After PROTECT: Boston

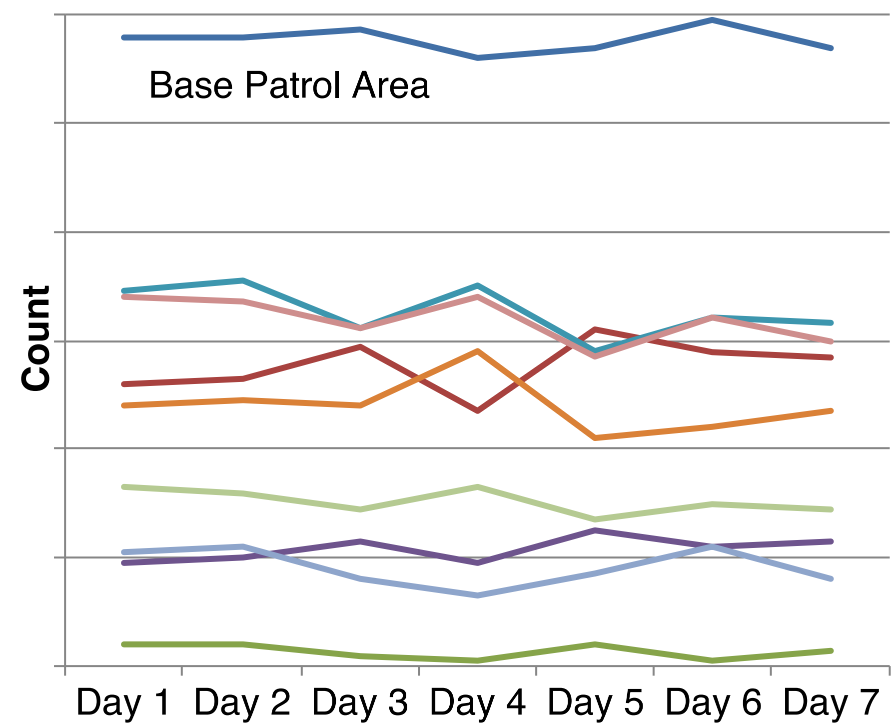

350% increase in defender expected utility

Field Evaluation of Schedule Quality

Improved Patrol Unpredictability & Coverage for Less Effort

FAMS: IRIS Outperformed expert human over six months

Report:GAO-09-903T

IRIS

Intelligent Randomization In Schedu

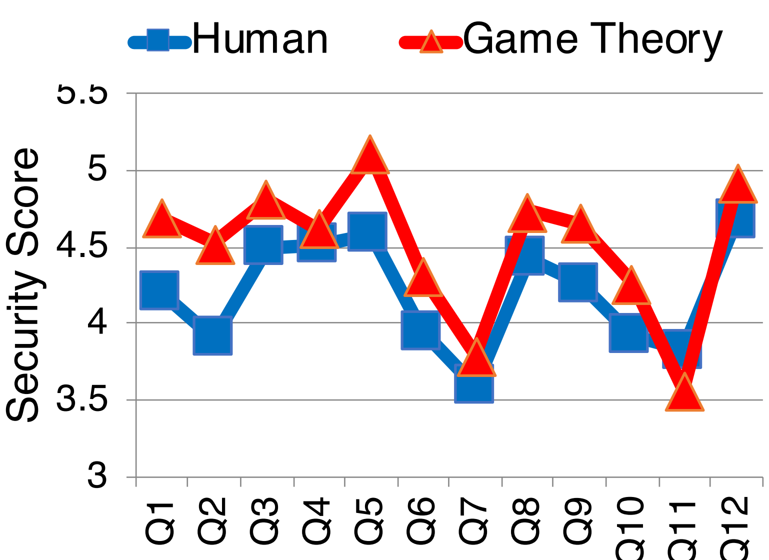

Trains: TRUSTS outperformed expert humans schedule 90 officers on LA trains

Field Tests Against Adversaries

Computational Game Theory in the Field

Controlled

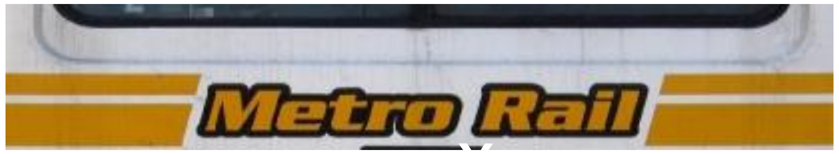

- 21 days of patrol, identical conditions
- Game theory vs Baseline+Expert

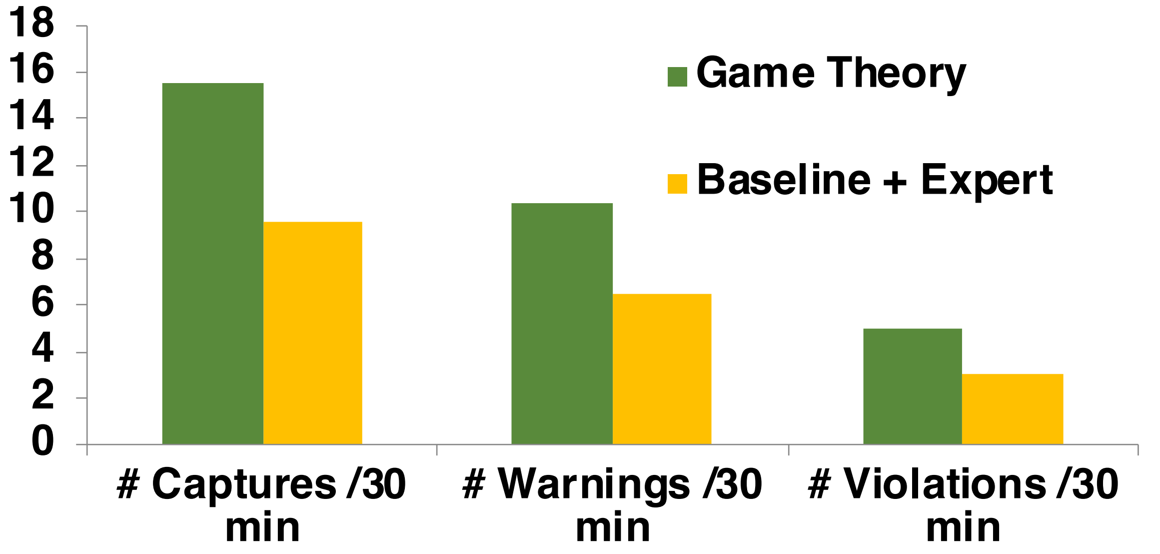

Not Controlled

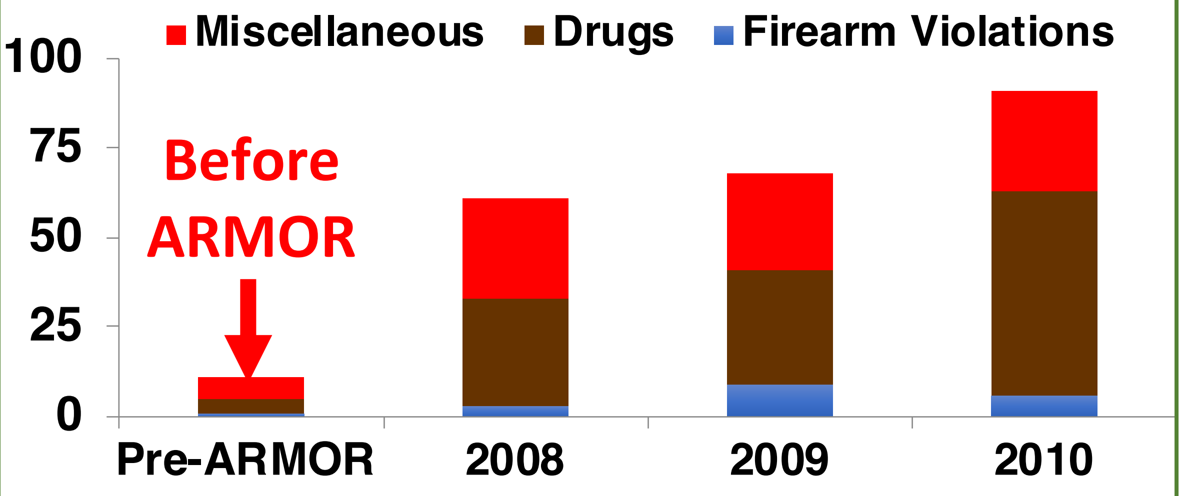

New Directions in Stackelberg Security Games


McCarthy


Schlenker


Sinha

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

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

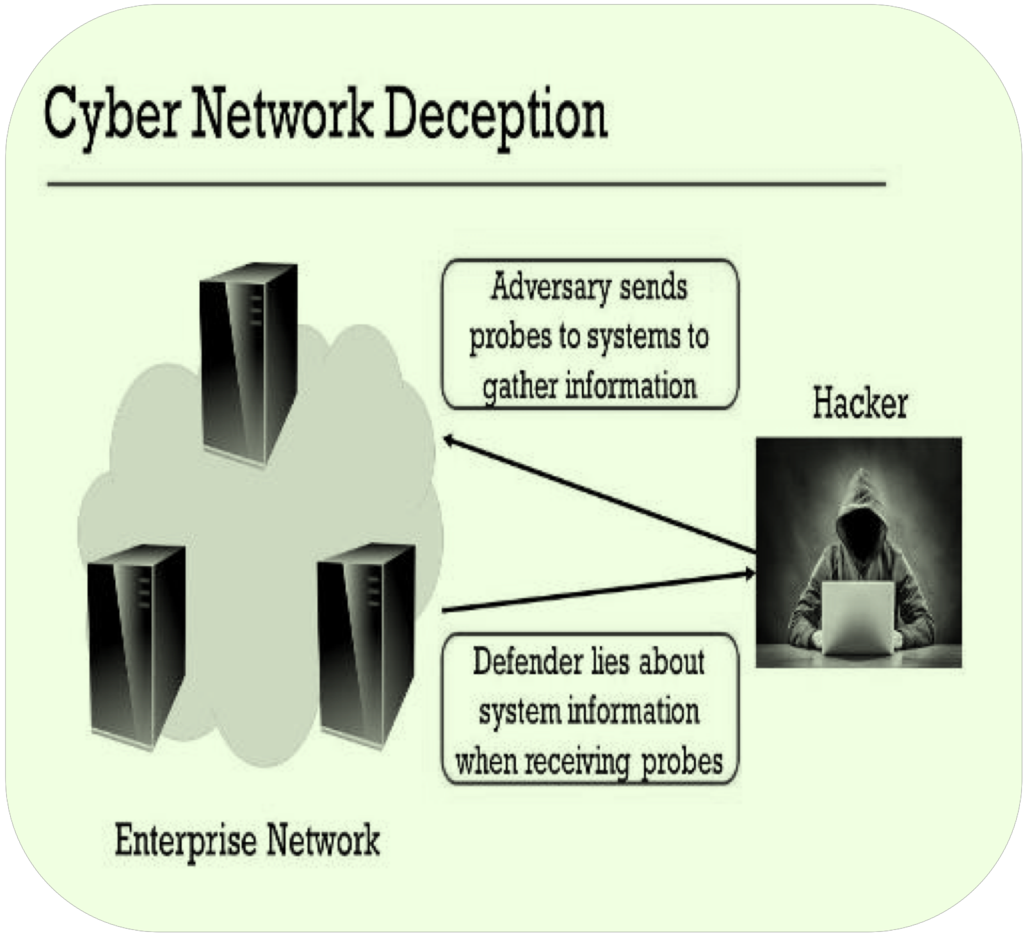

Outline

Public Safety and Security
Stackelberg Security Games

Conservation/Wildlife Protection:
Green Security Games

Dr Andy Plumptre
Conservation Biology

Public Health/Social Work:
Influence maximization/Game against nature

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

Green Security Games
Limited Ranger Resources to Protect Forests

Adversary not fully strategic; multiple "bounded rational" poachers

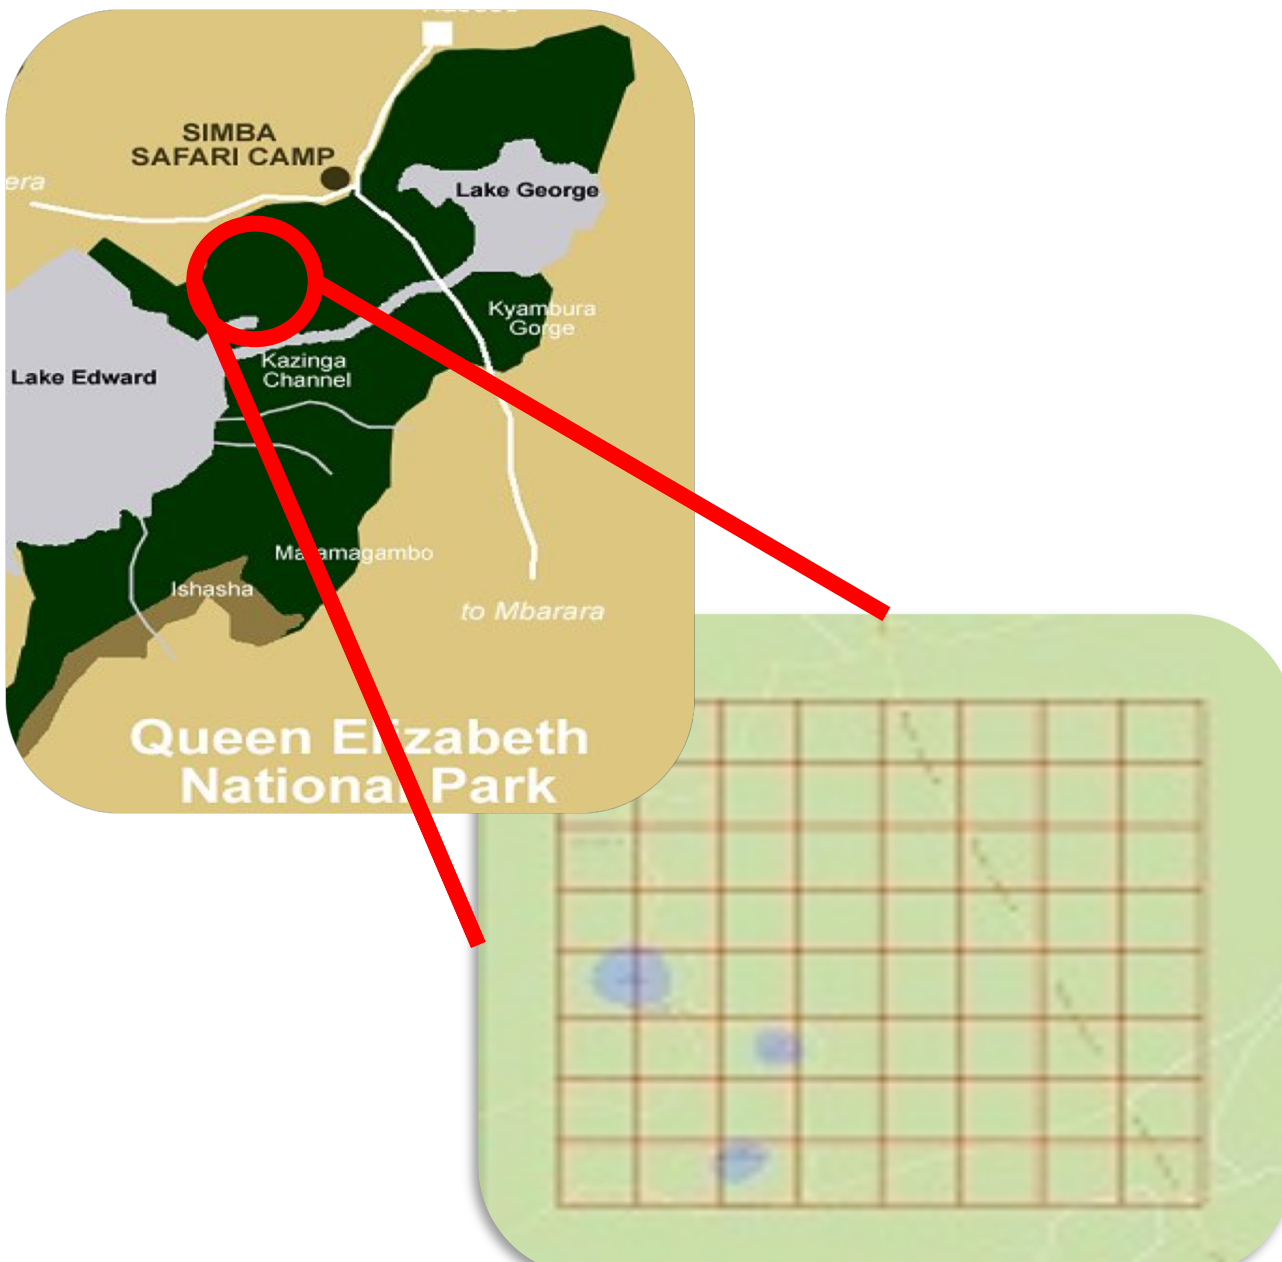

$$\max_{x,q} \sum_{i \in X} \sum_{j \in Q} R_{ij} x_i q_j$$

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
Game Theory + Machine Learning Poacher Behavior


Xu

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

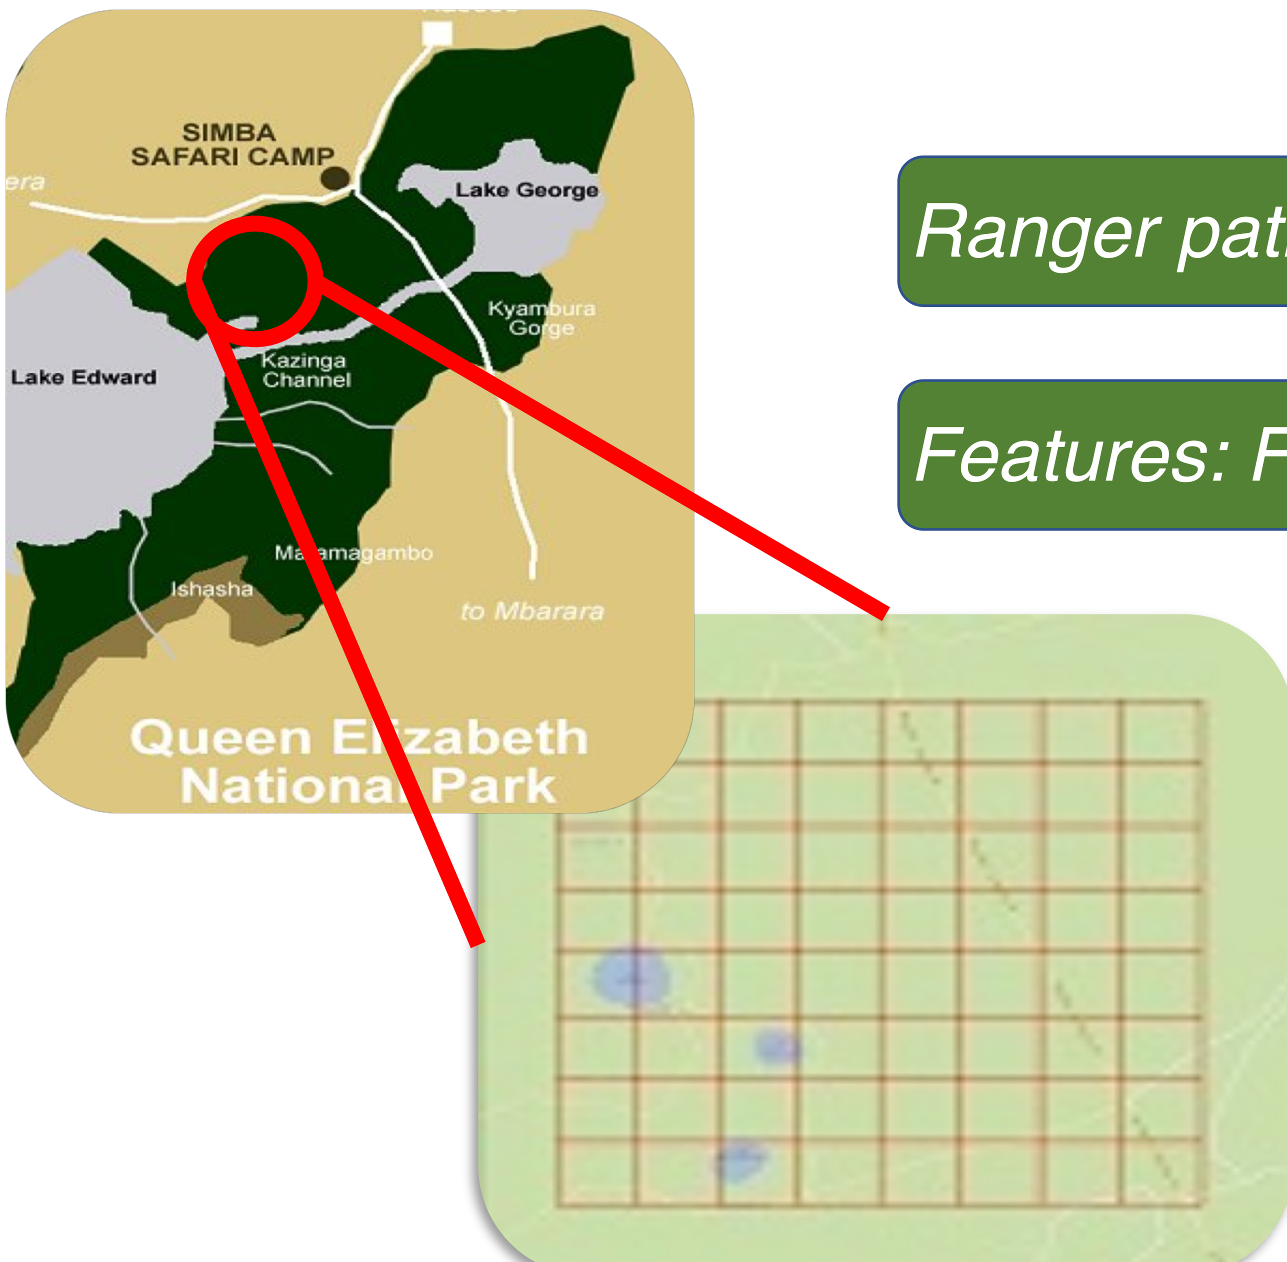

Ranger patrols: $X_{(i)}$

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

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

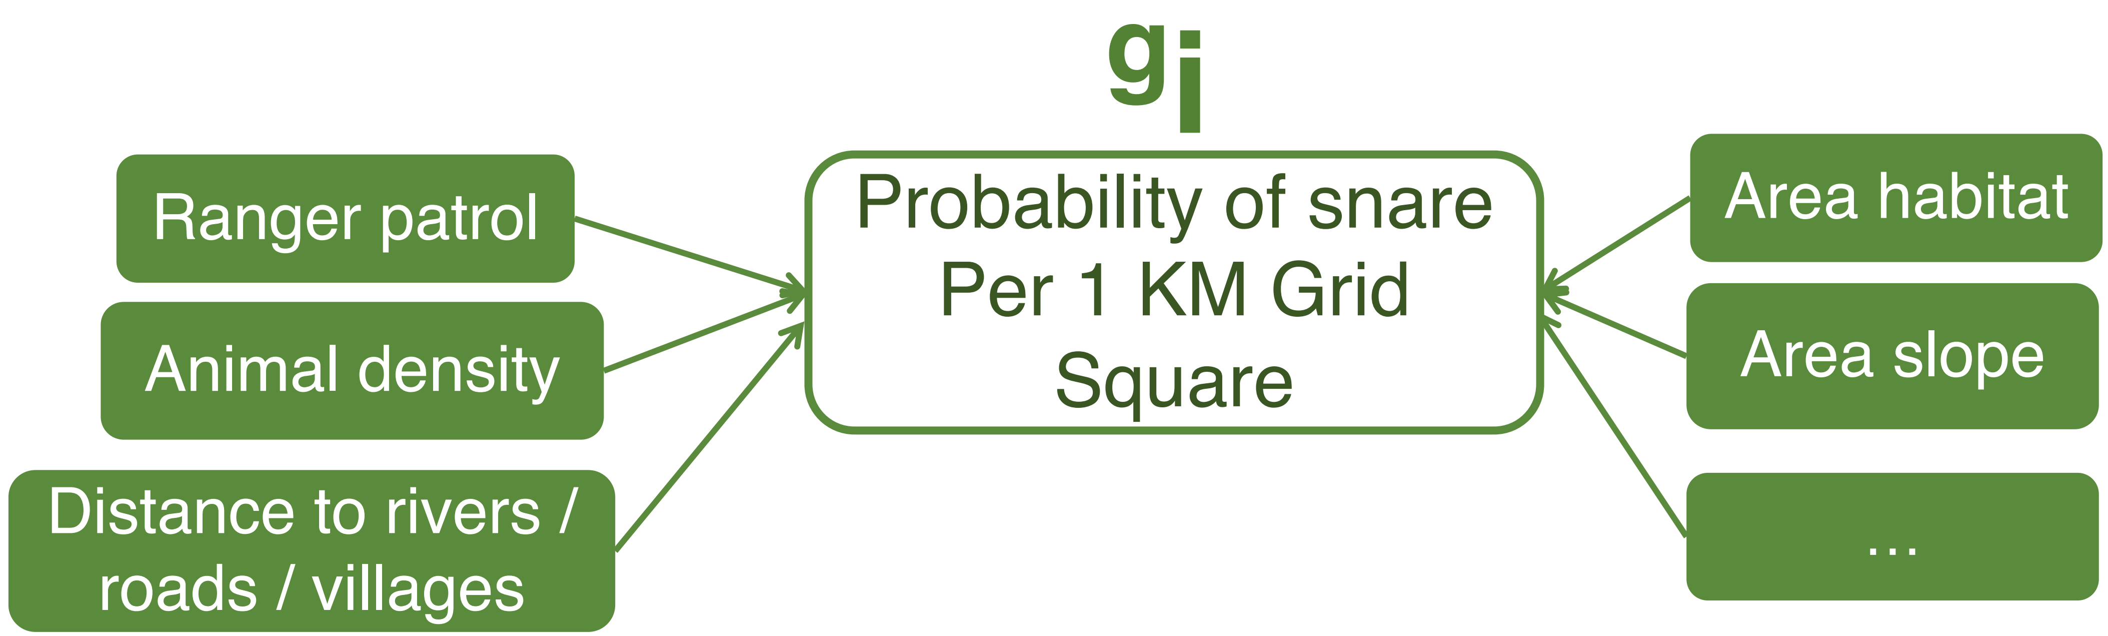

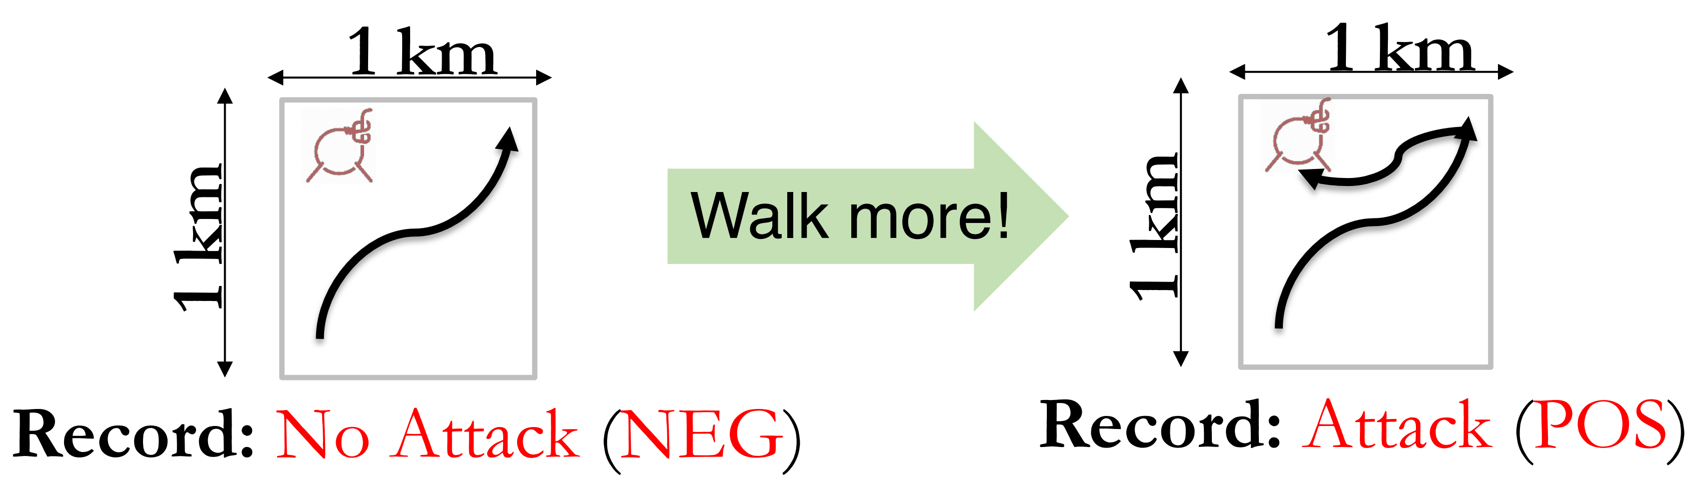

Adversary Modeling
Imperfect Crime Observation-aware Ensemble Model

Gholami

Training: Filtered Datasets

Predict: Ensemble of Classifiers

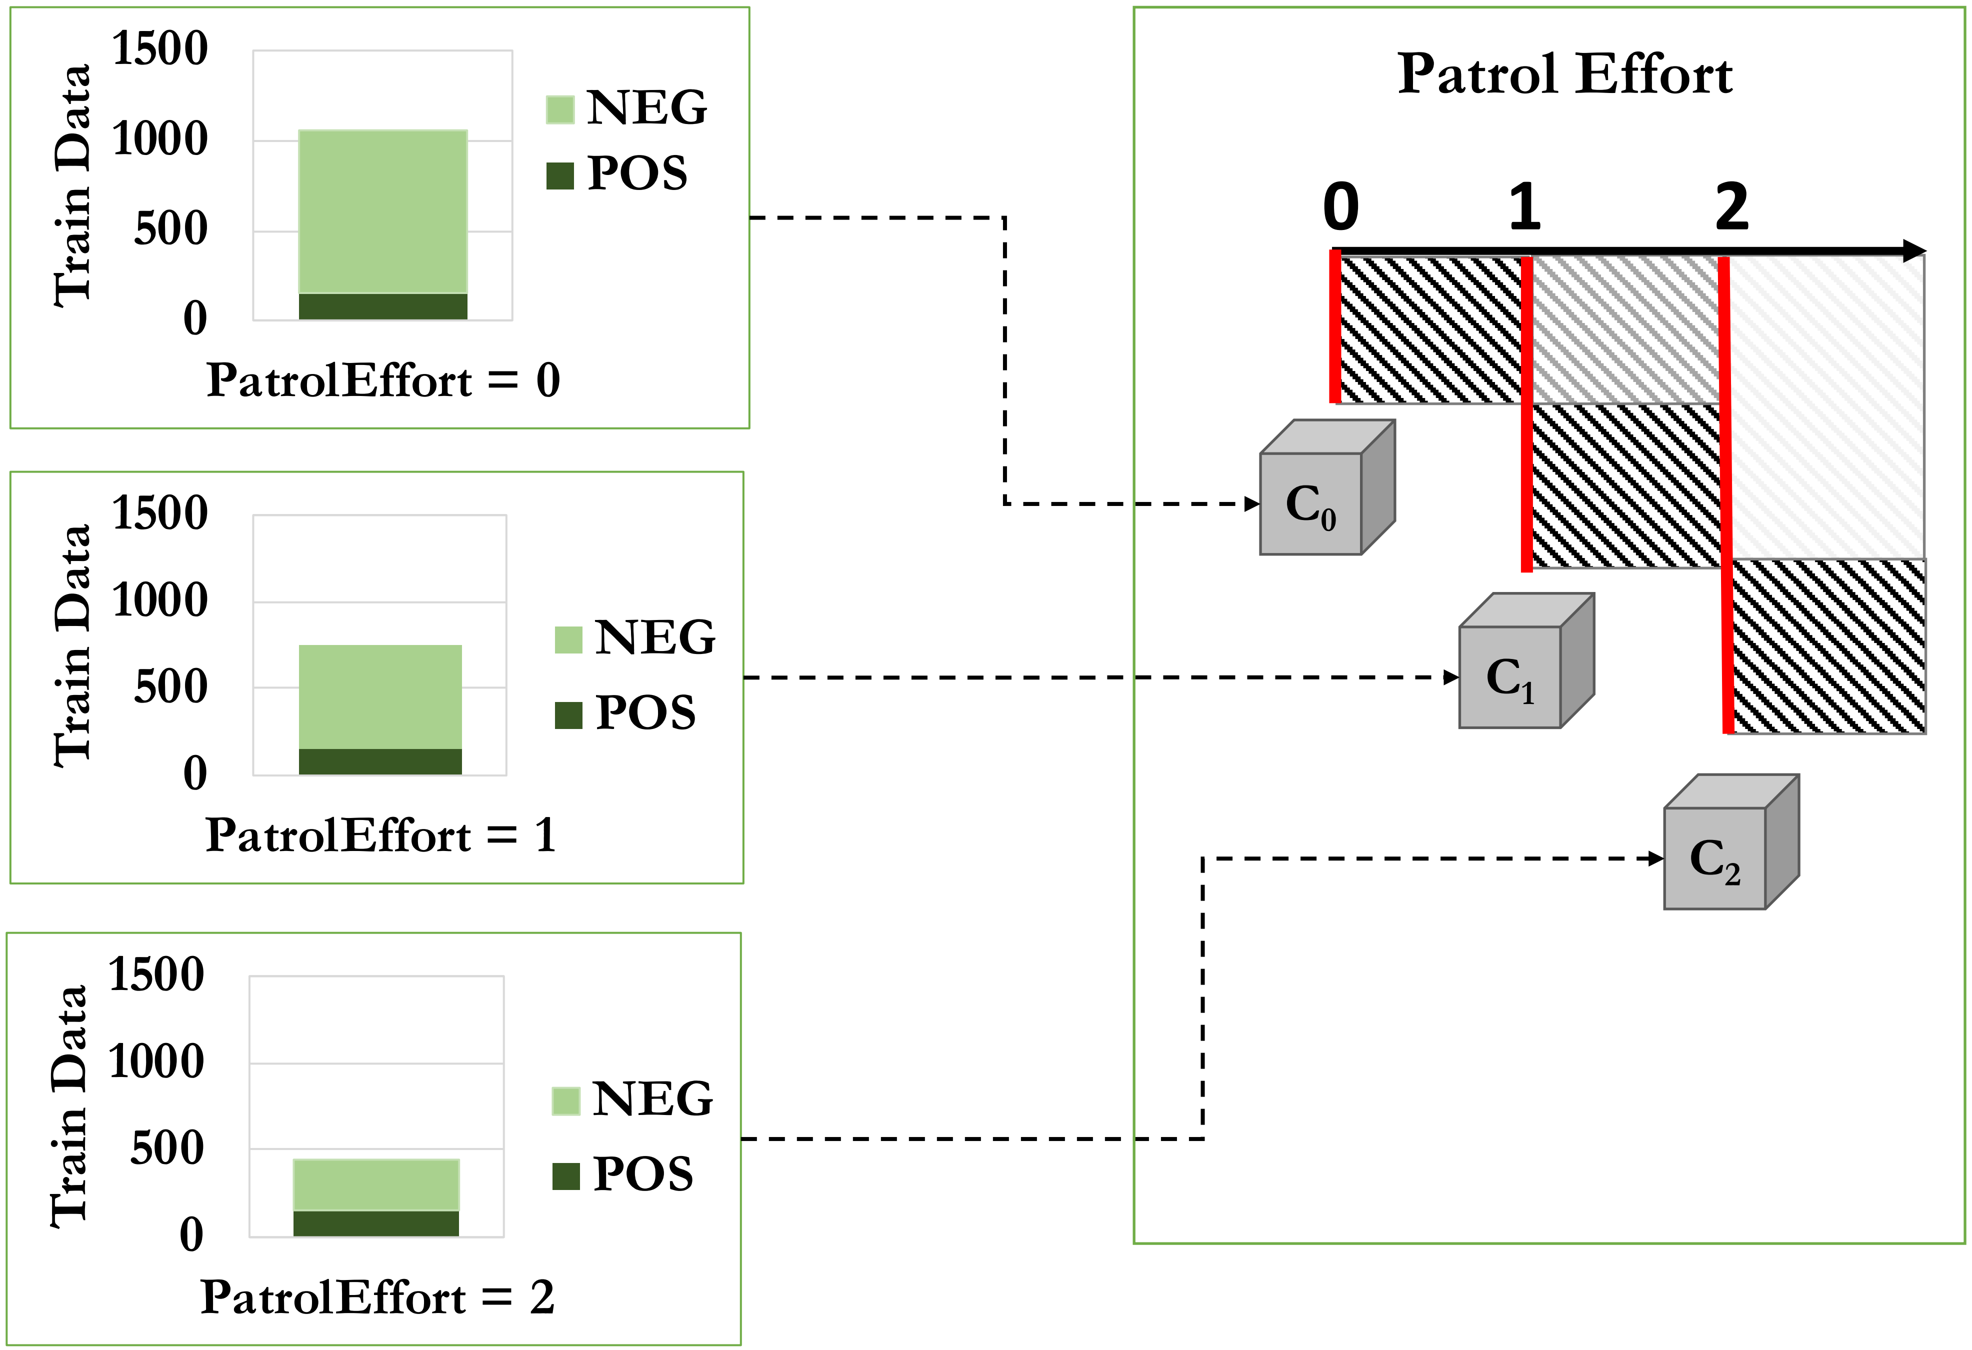

Poacher Attack Prediction in the Lab

Gholami

Poacher Behavior Prediction

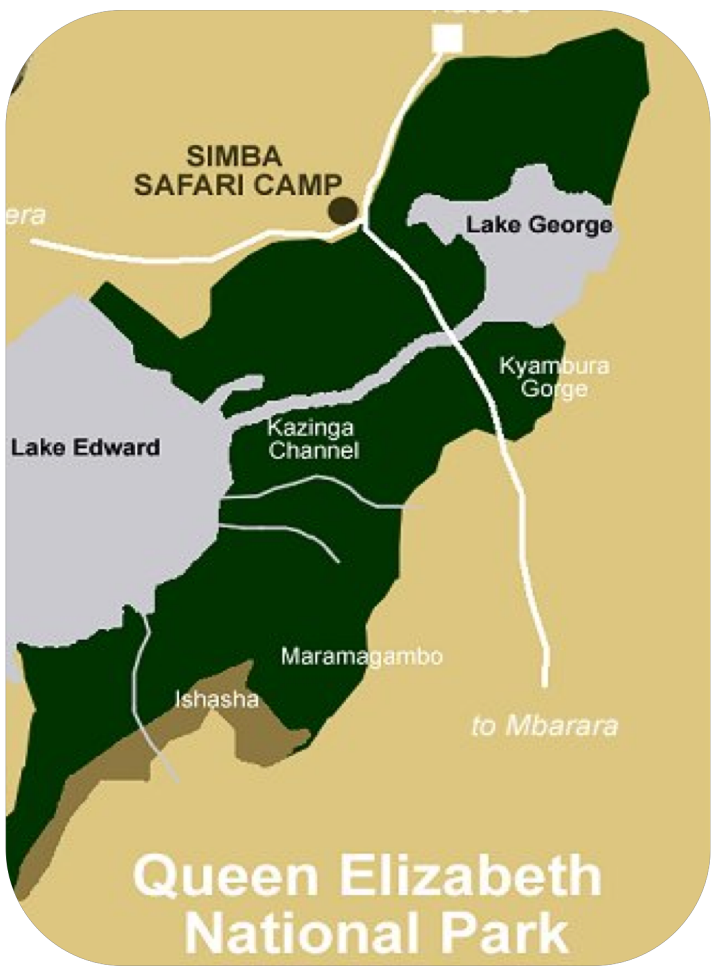

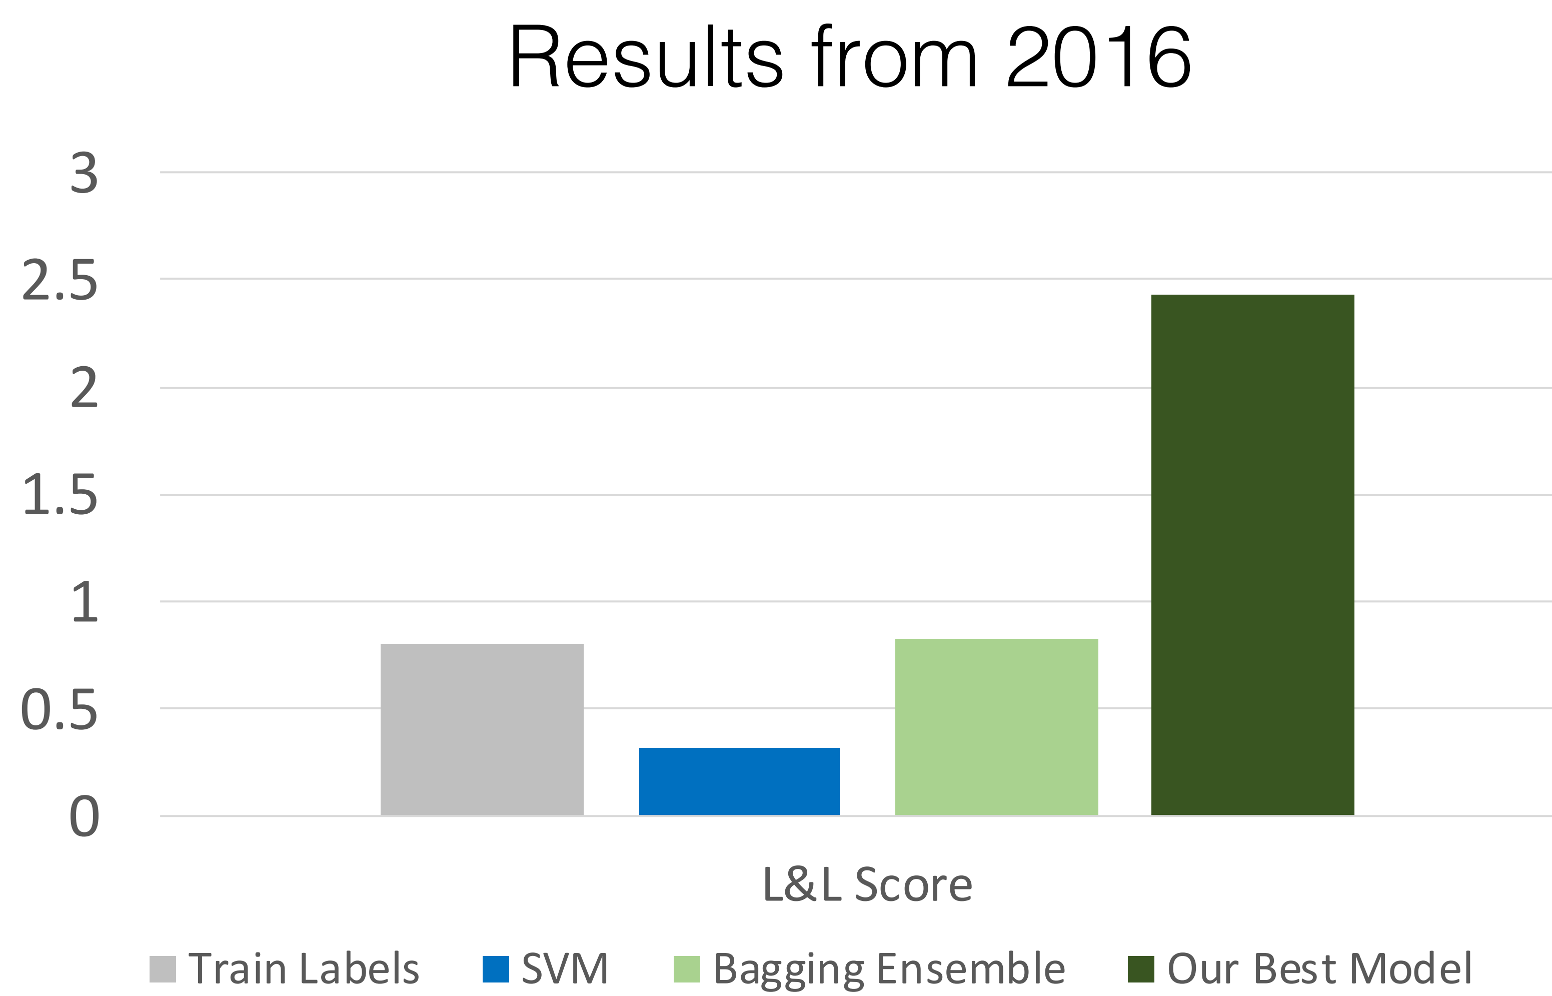

Real-world Deployment 2016: First Trial


Ford


Gholami

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

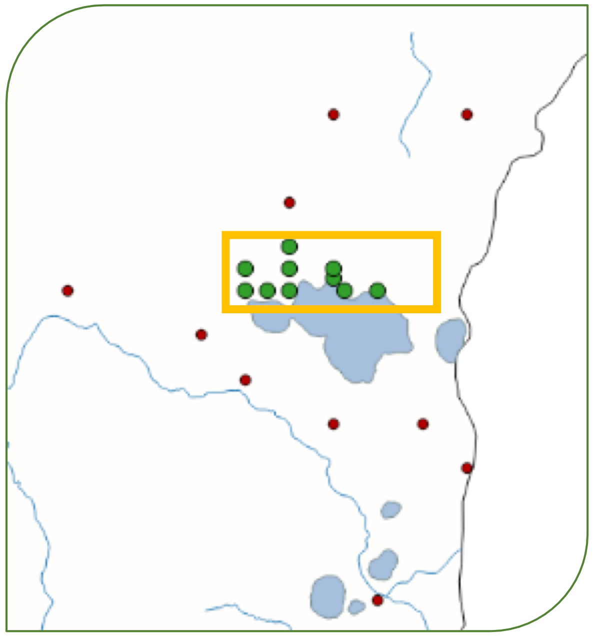

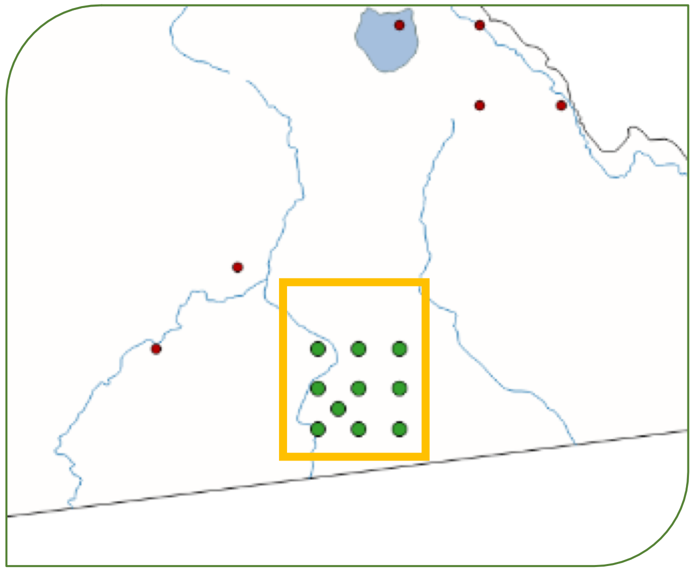

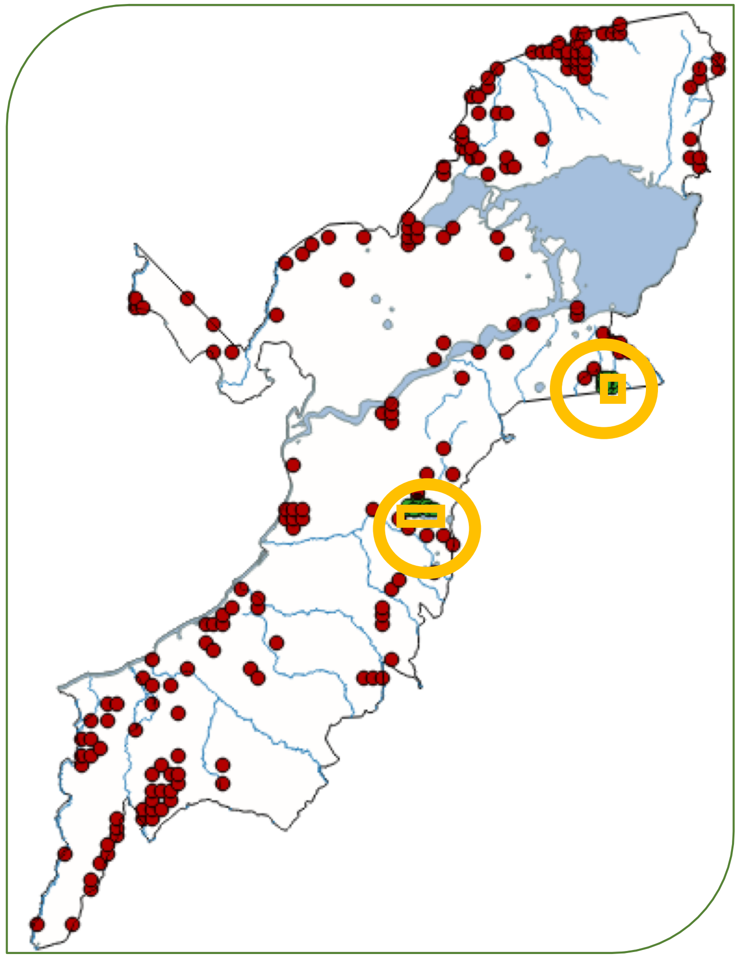

Real-world Deployment
Two Hot Spots Predicted

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

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

Gholami

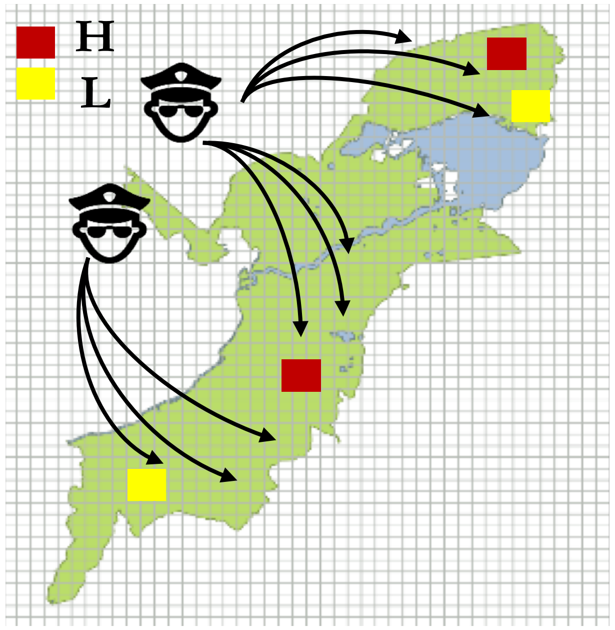

Queen Elizabeth National Park

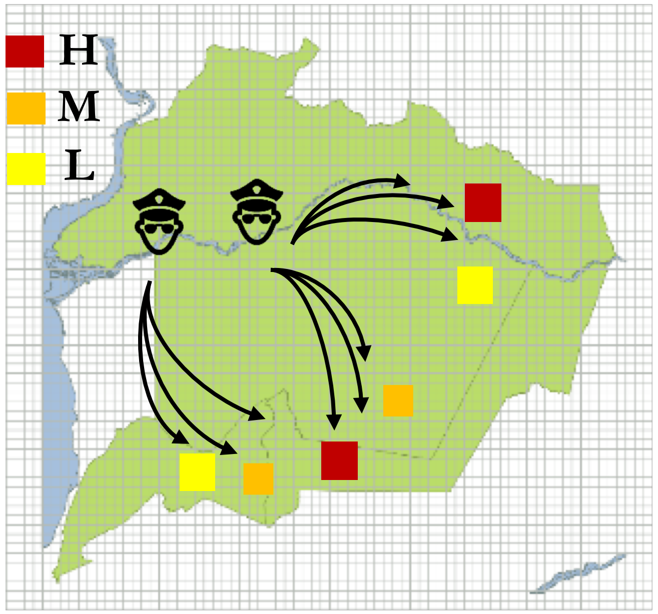

Murchison Falls National Park

Snares per patrolled sq. KM

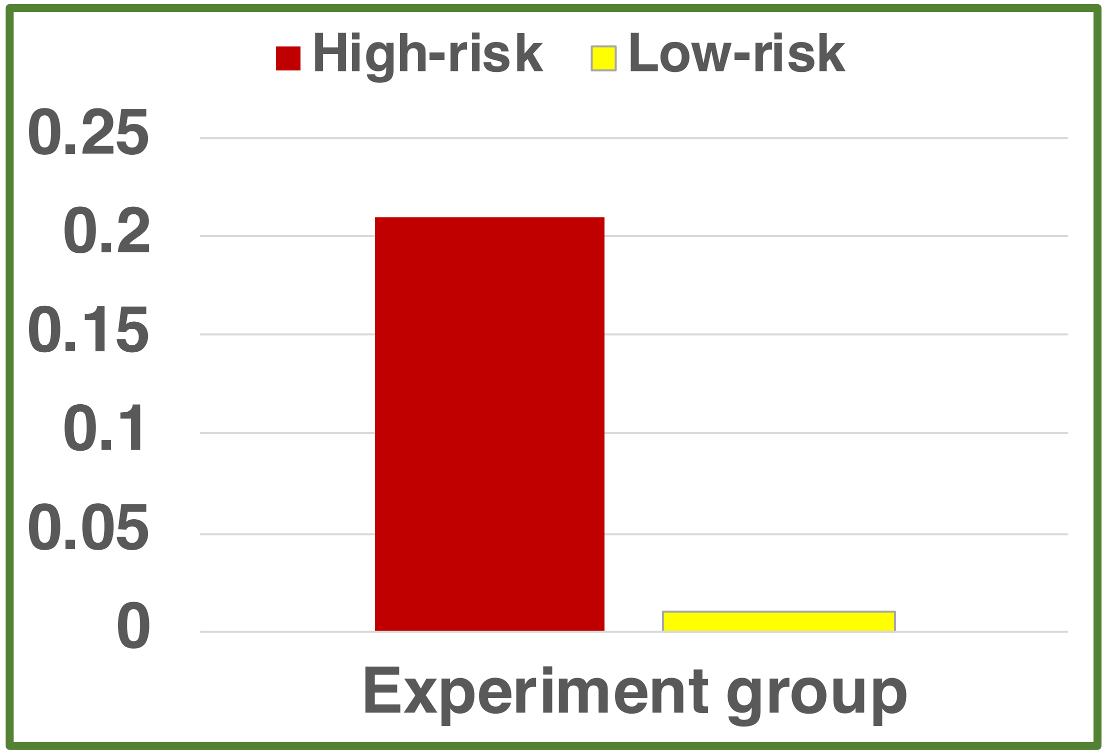

Snares per patrolled sq. KM

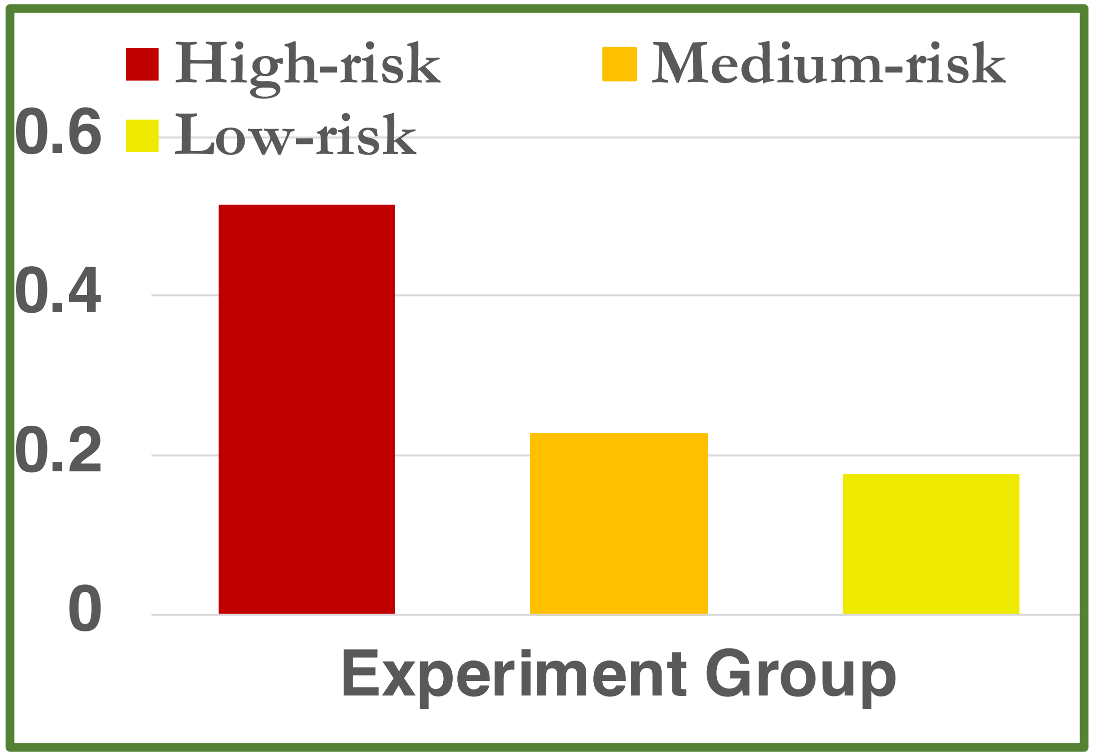

Green Security Games: Incorporating Real Time Information


Xu


Bondi

- Drones in Green Security Games (AAAI18, IAAI18, GameSec17…)

Air Shepherd
The Lindbergh Foundation

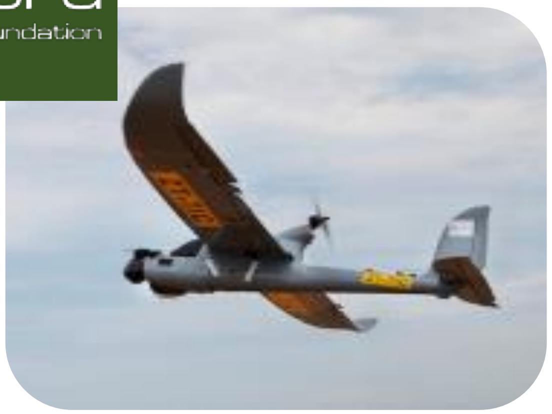

Ranger patrols: $X(i)$

Features: $F(i)$

Machine Learning

g_i

Probability of finding snare in cell i

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

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

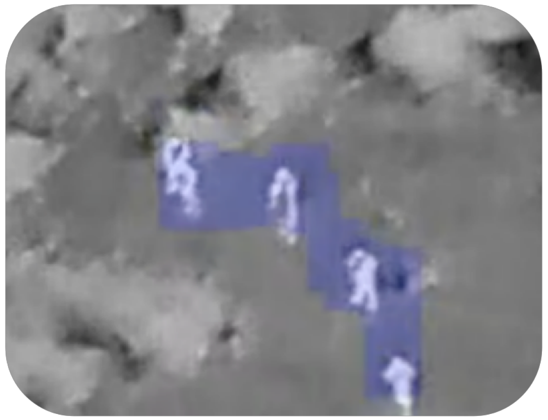

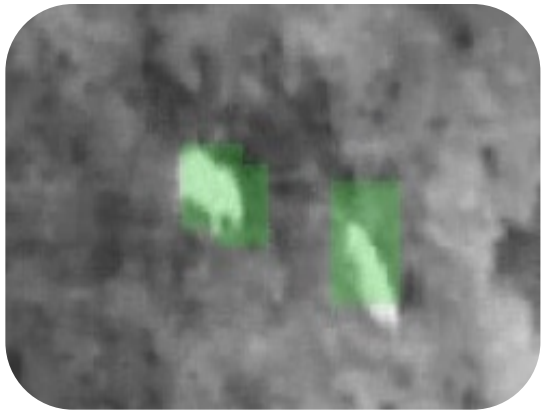

Green Security Games:
Around the Globe with SMART partnership

600
National Parks
Around the Globe

Wildlife, Forests, Fisheries…

Outline

Public Safety and Security
Stackelberg Security Games

Conservation/Wildlife Protection:
Green Security Games

Public Health:
Influence maximization/Game against nature

Prof Eric Rice
Social Work

Public Health
Optimizing Limited Intervention (Messaging) 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*

Influence Maximization Background

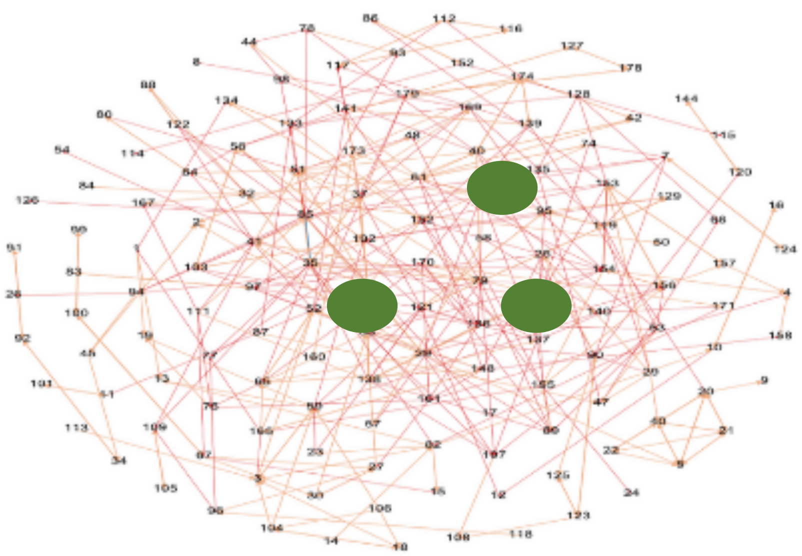

- Given:
  - Social network Graph G
  - Choose K “peer leader” nodes
- Objective:
  - Maximize expected number of influenced nodes
- *Assumption: Independent cascade model of information spread*

Independent Cascade Model and Real-world Physical Social Networks

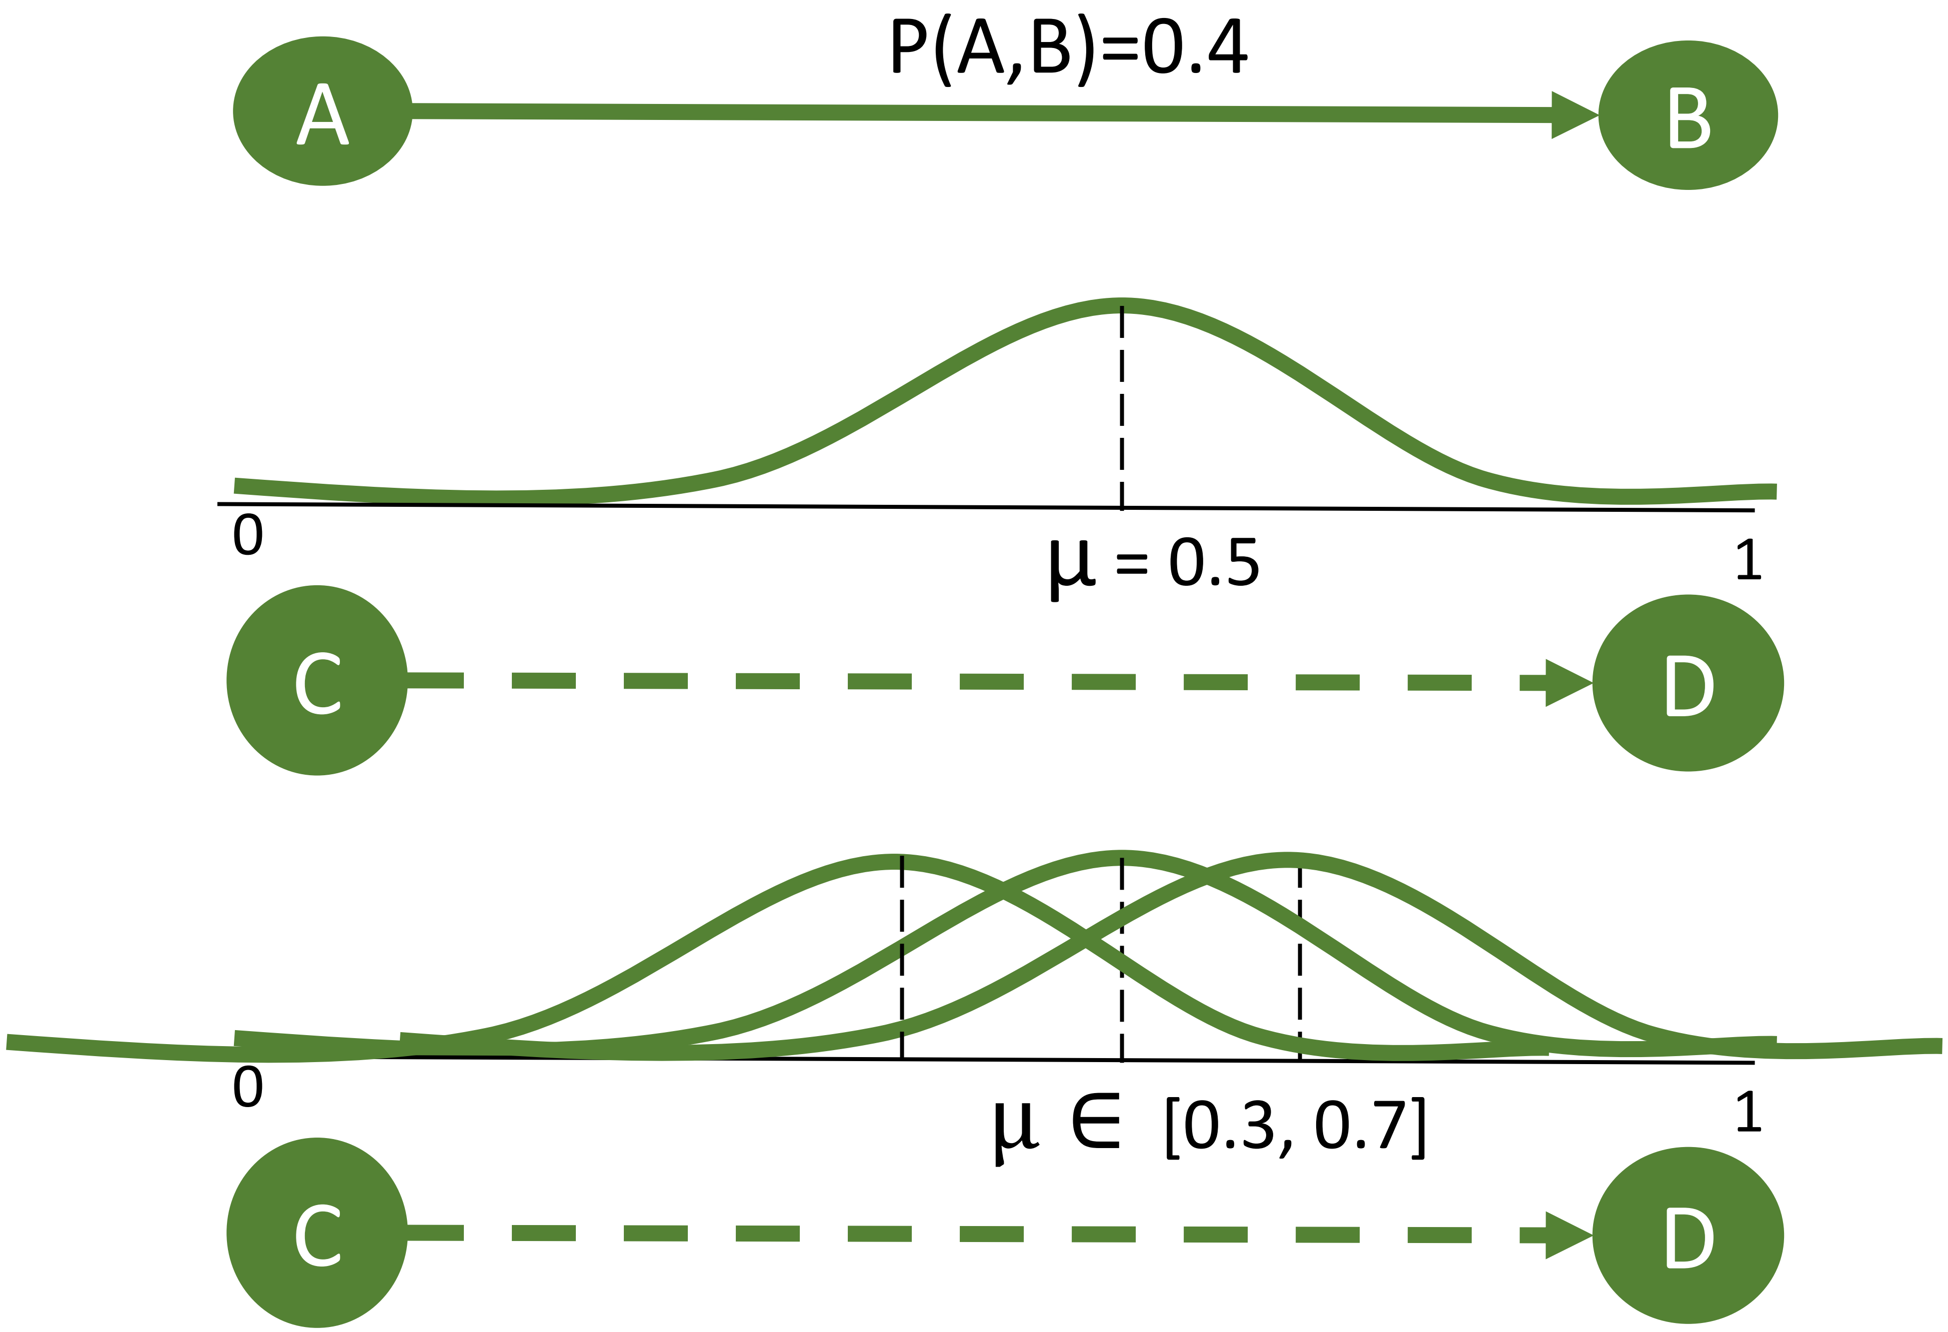

Robust, Dynamic Influence Maximization

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

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

Influencer'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

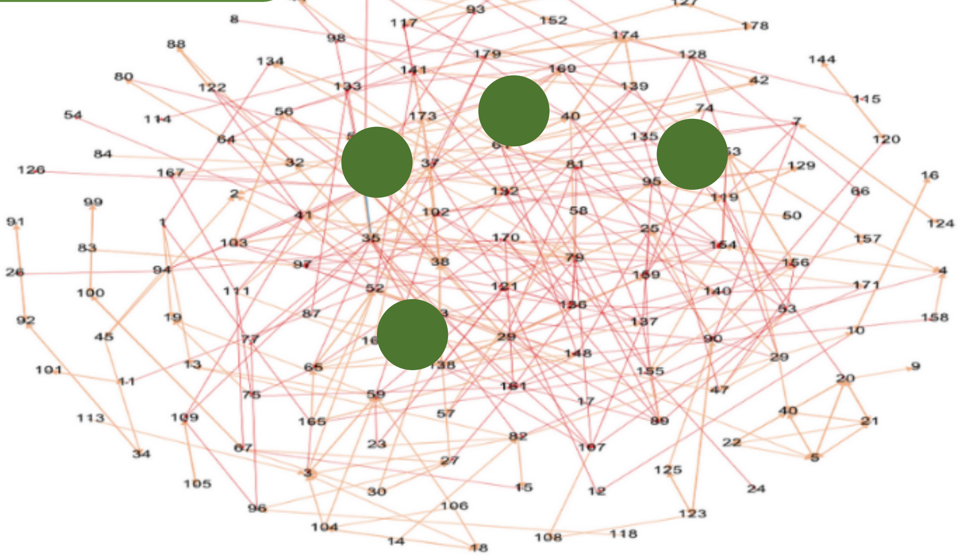

K = 4
2nd time step

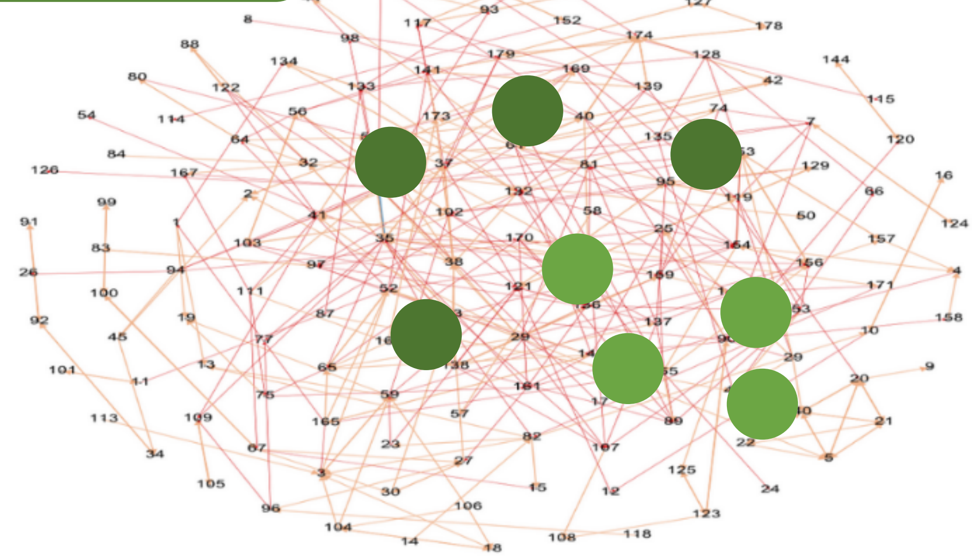

HEALER: POMDP Model for Multi-Step Policy [2015]
Robust, Dynamic Influence Maximization

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

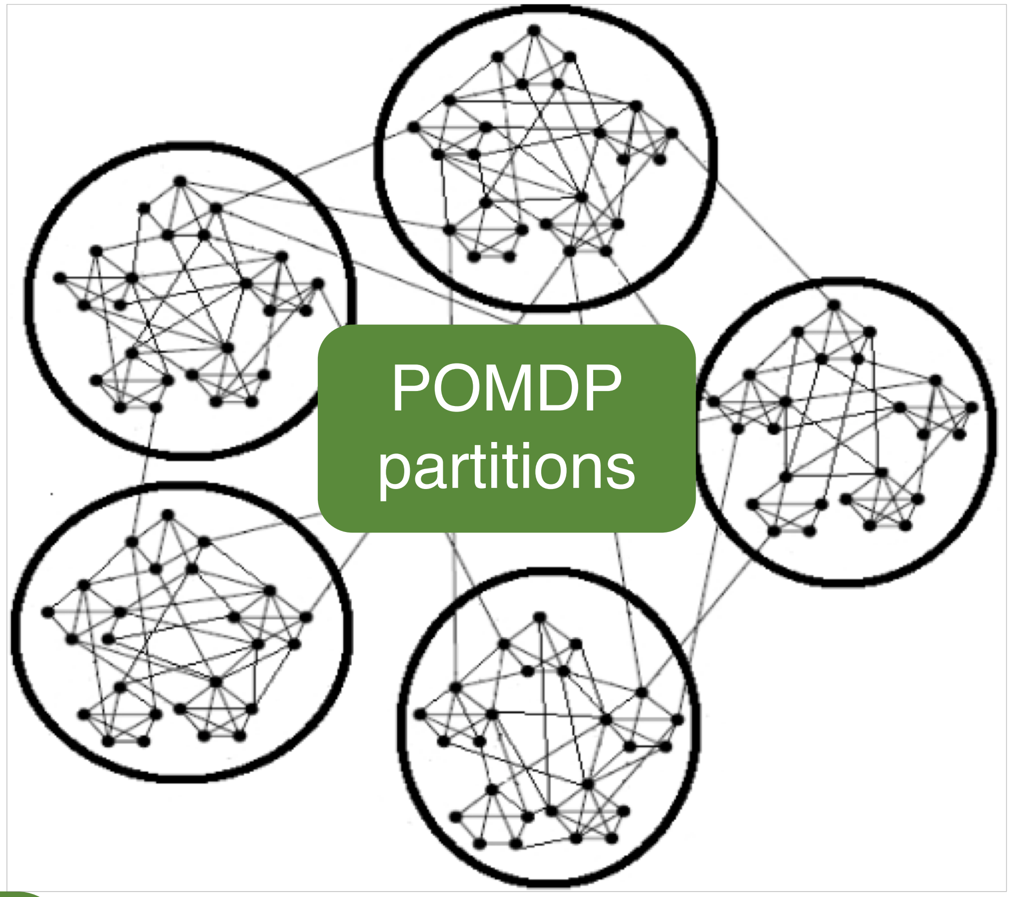

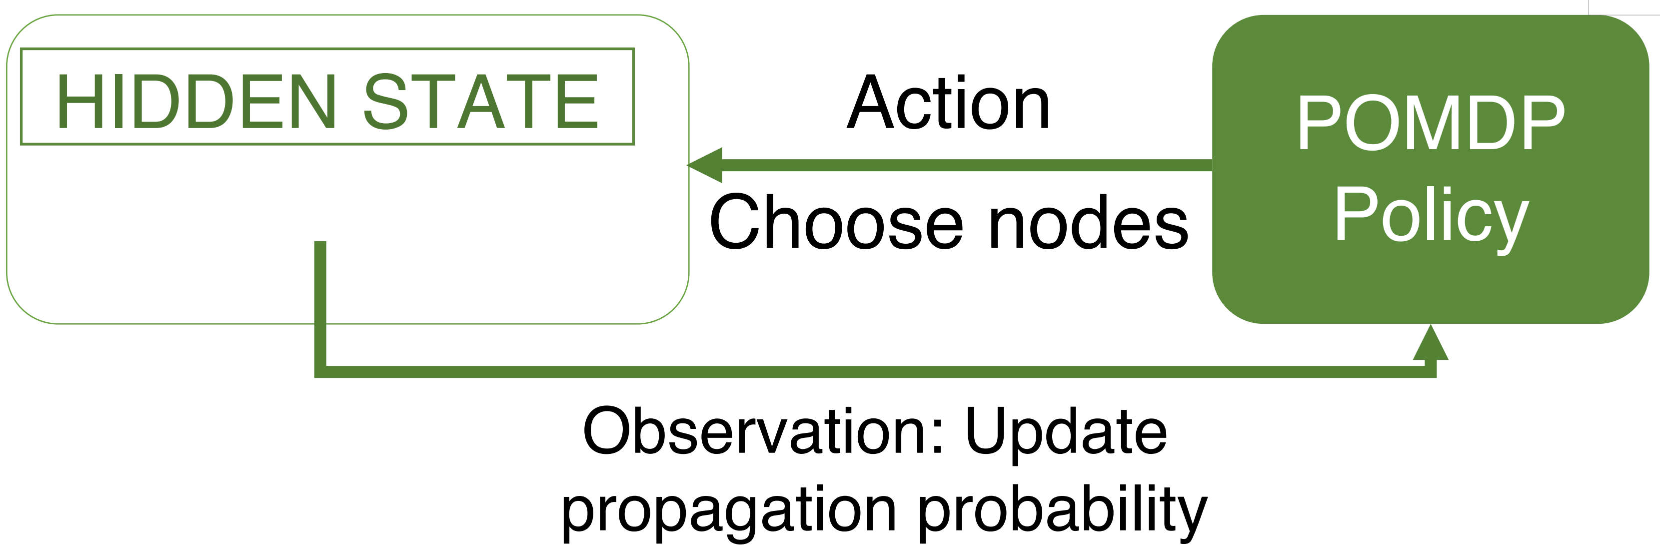

Pilot Tests with HEALER with 170 Homeless Youth [2017]

Yadav

Wilder

Recruited youths:

HEALER	HEALER++	DEGREE CENTRALITY
62	56	55

12 peer leaders

Results: Pilot Studies

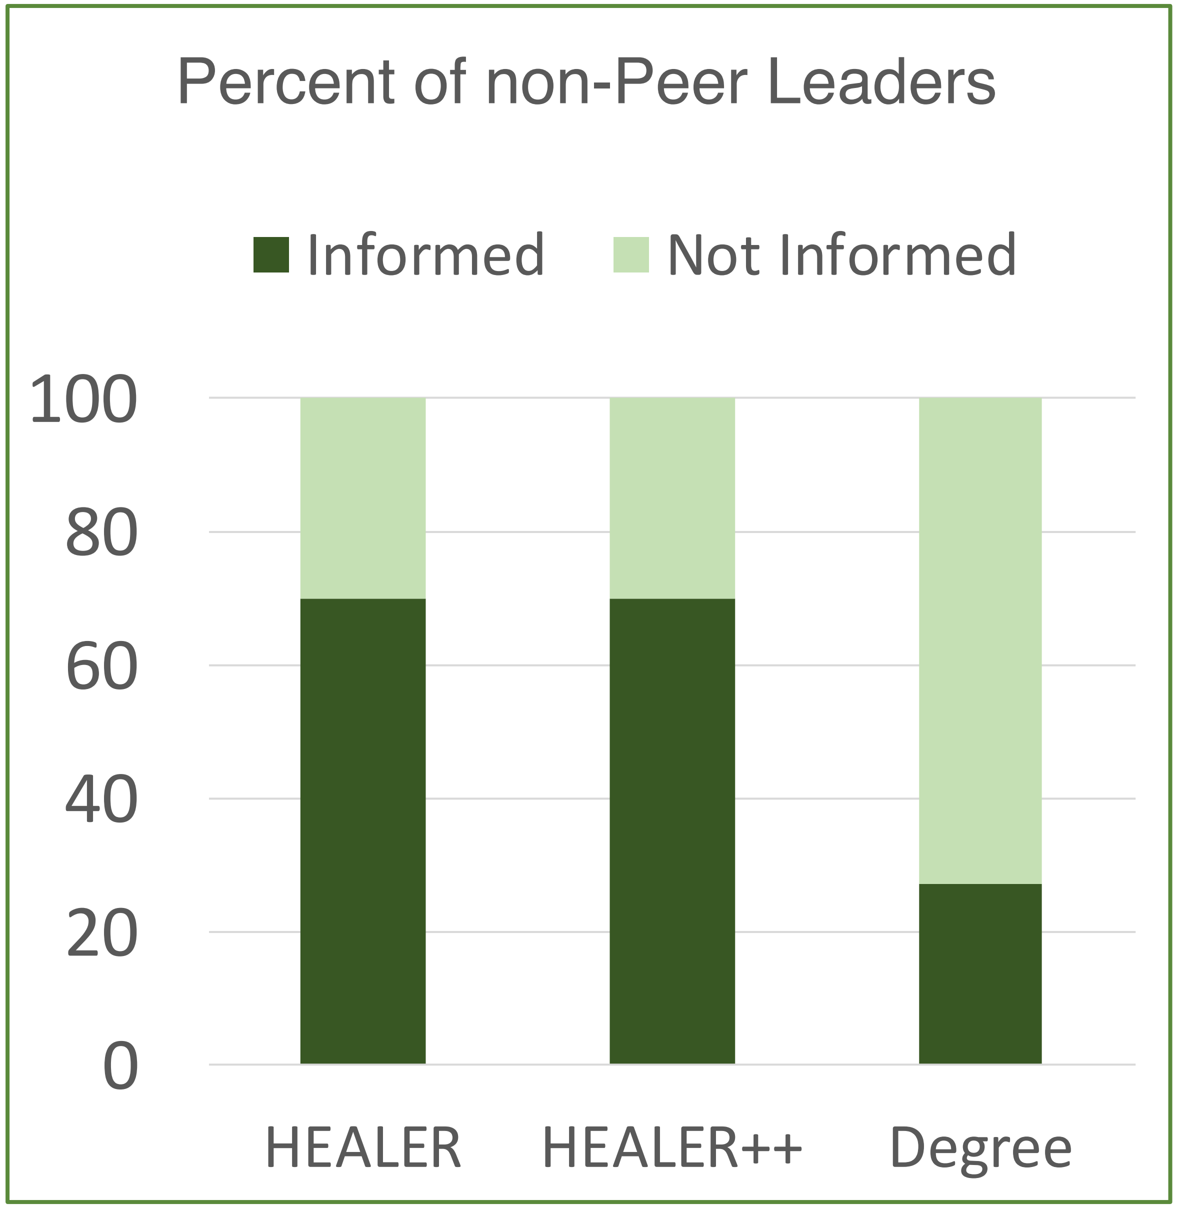

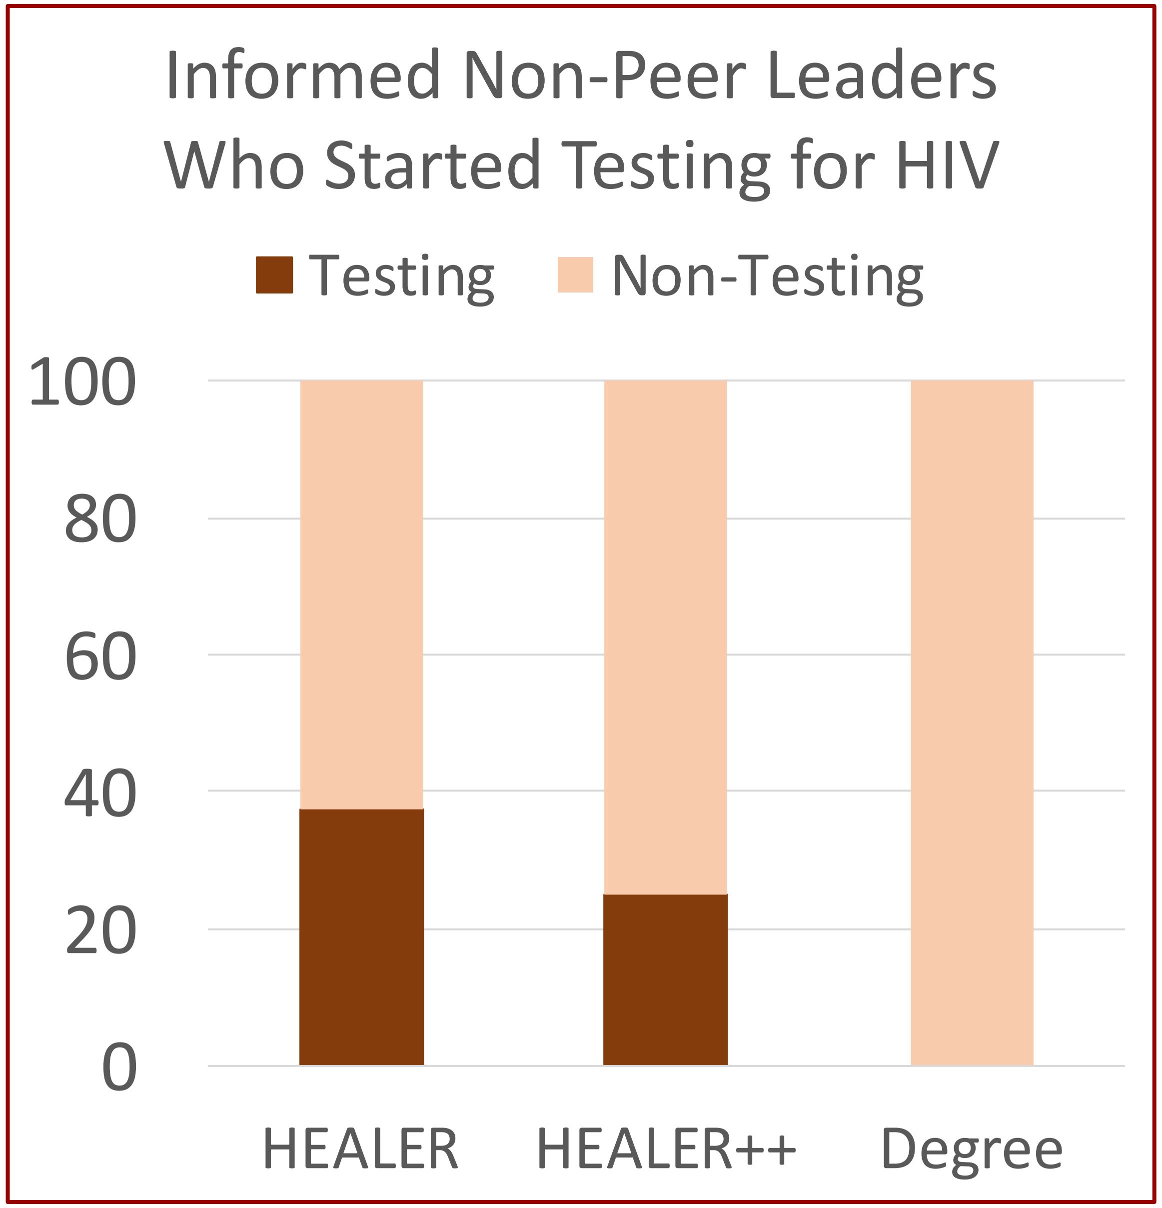

AI Assistant: HEALER

New Directions: Los Angeles From an Angeleno

900 youth study

(AAAI18, AAMAS18)

Mayor Garcetti @ USC

New Directions: Mumbai From a Mumbaikar

(AAAI18)

महाराष्ट्र शासन

Prime Minister Modi @ Mumbai
AI for Social Good

Key Lessons:
Directing Multiagent Systems Research towards Social Good

Multiagent systems research helps address complex social problems:

- 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
- *New models*: Stackelberg security games, green security games…
- *Key algorithms*: Incremental strategy generation, marginals, double oracle…

Immersion/Deployment helps identify crucial research challenges

Future: Multiagent Systems and AI Research for Social Good

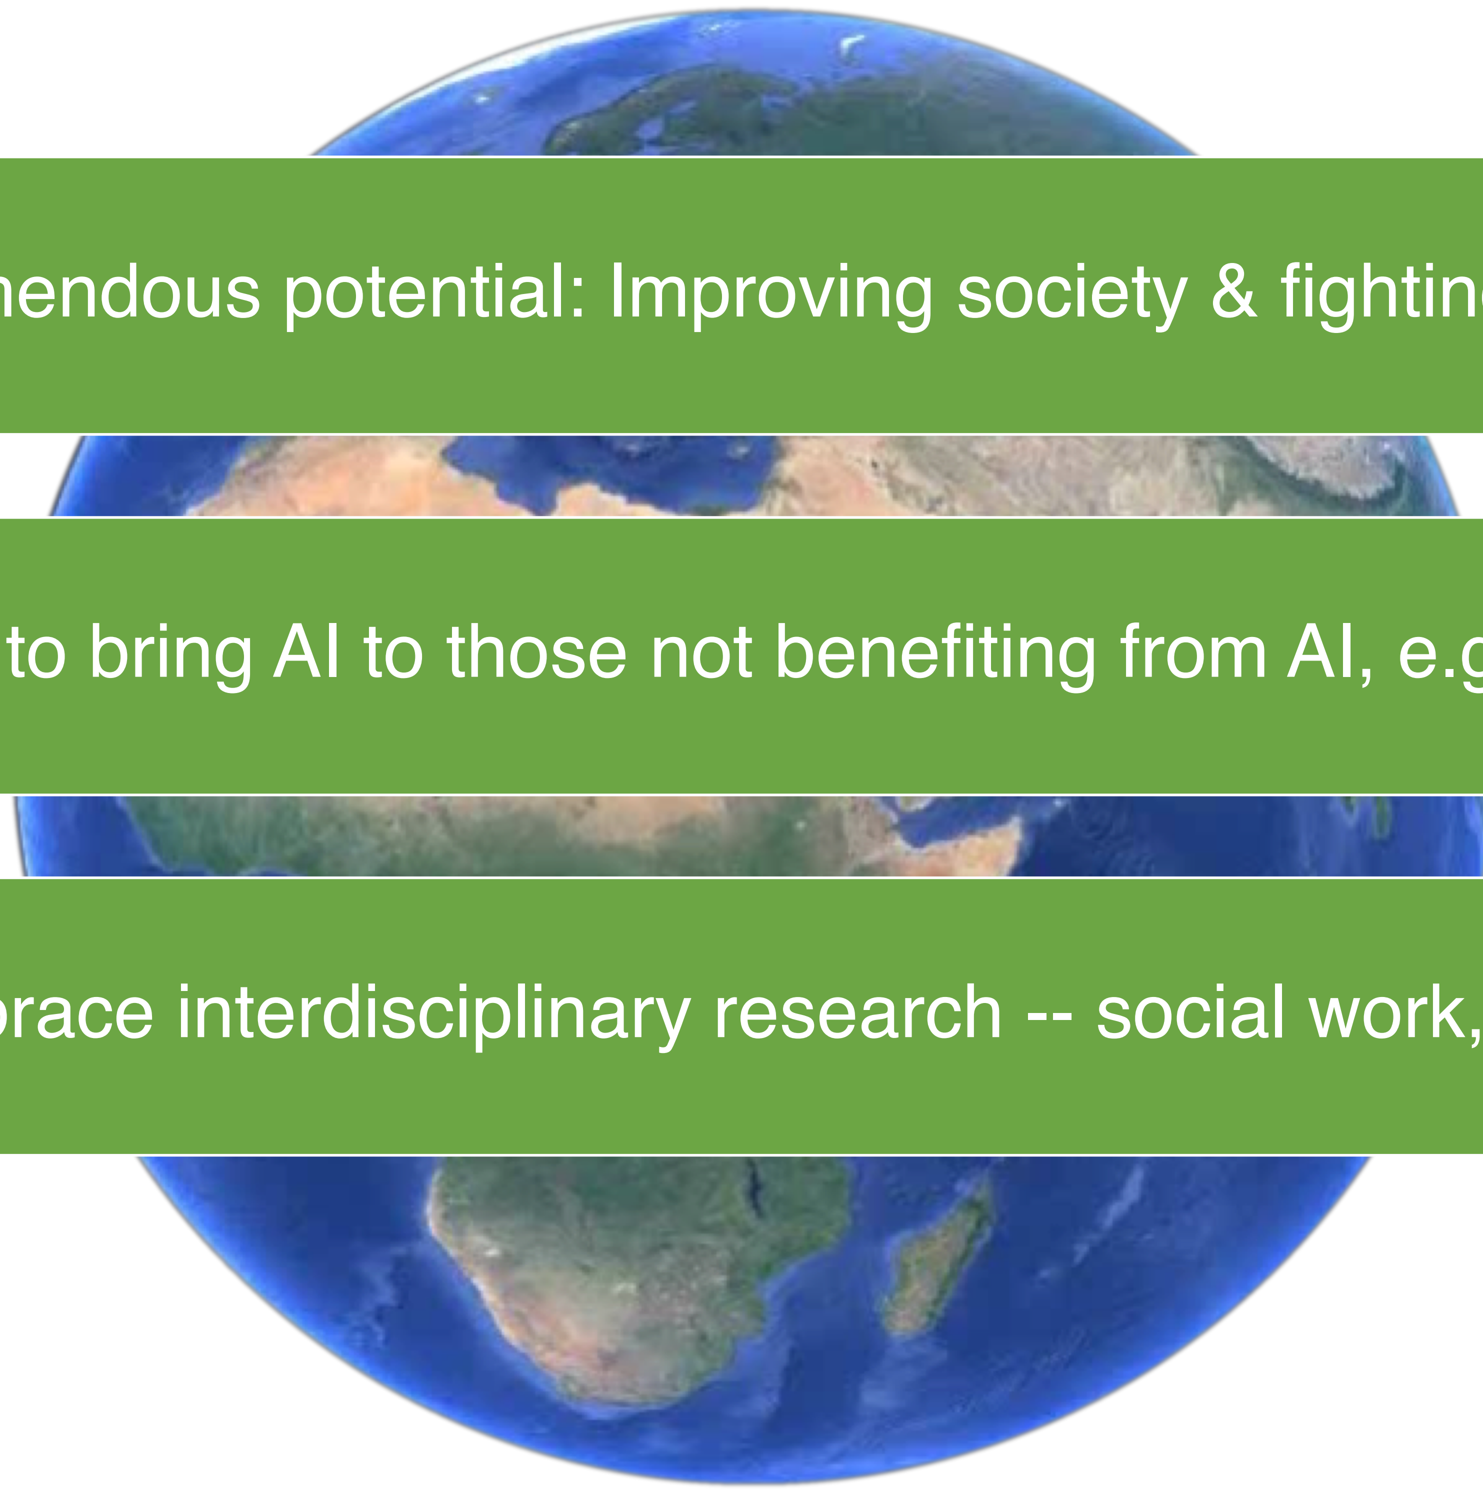

Tremendous potential: Improving society & fighting social injustice

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*
- *Actual problem for societal benefit?*
- *Model deficiencies for new research directions?*

Thank you

Mentor: Barbara Grosz

Collaborators:

Sarit Kraus

Vince Conitzer

Eugene Vorobeychik

Andy Plumptre

USC Collaborators:

Eric Rice

Bistra Dilkina

Phebe Vayanos

Fernando Ordonez

BACKUP

tambe@usc.edu