Optimal Patrol Planning for Green Security Games with Black-Box Attackers

Haifeng Xu[1], Ben Ford[1], Fei Fang[2], Bistra Dilkina[1], Andrew Plumptre[3], Milind Tambe[1], Margaret Driciru[4], Fred Wanyama[4], Aggrey Rwetsiba[4], Mustapha Nsubaga[4], Joshua Mabonga[4]

[1]University of Southern California (USC)
[2]Carnegie Mellon University (CMU)
[3]Wildlife Conservation Society
[4]Uganda Wildlife Authority

The Classical Stackelberg Security Game Paradigm

- Stackelberg Game
  - Defender (leader): use limited resources to protect critical targets
  - Attacker (follower): long-term surveillance, well-planned (thus perfectly rational)

Flights

Ferries

Airports

Road Networks

A Rapidly Growing Trend: Green Security Domains

Endangered Wildlife

Fisheries

Environmental Resources

Challenges for Patrol Planning in Green Security Games

- Attacker's bounded rationality $\Rightarrow$ intricate attacker (behavior) models
  - E.g., graphical model [Nguyen et al.'16], ensemble of decision trees [Kar et al.'17], Markov random field [Gholami et al.'17] . . .

Challenge 1:

How to optimize patrolling against these complicated attacker models?

Do we have to design a different algorithm for each attacker model?

Challenges for Patrol Planning in Green Security Games

- Attackers may have partial real-time surveillance
  - Can observe rangers' current move and infer where they go next

" Those (poachers) would simply observe the rangers and base their offending patterns on the schedules of the rangers "

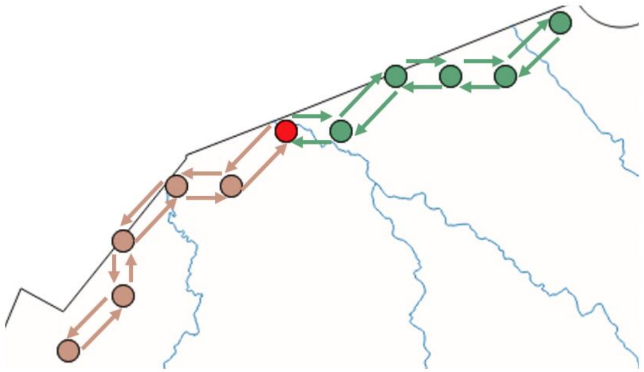

Challenge 2:

How to deal with attacker's (partial) real-time surveillance?

Our Contributions:

- A new patrol planning framework OPERA (**O**ptimal patrol **P**lanning with **E**nhanced **RA**ndomness)
  - Work for *any attacker model* (under mild assumptions)
  - Mitigate negative effects of attacker's real-time surveillance with enhanced randomness
- Test performances on real-world data from Uganda

Outline

- **Motivation and Game Model**
- **Optimal Patrol Planning Against Black-Box Attackers**
- **Experimental Evaluation**

Outline

- Motivation and Game Model
- Optimal Patrol Planning Against Black-Box Attackers
- Experimental Evaluation

Motivation Domain: Wildlife Protection in Uganda

Forest Area: QEPA

- Covers 2520 sq. km
- Divided into grids of 1km×1km

Poachers: set trapping tools (e.g., snare)

Rangers: conduct patrols

Our Goal: maximize catches of snares

Collaborators: Wildlife Conservation Society, Uganda Wildlife Authority,

Motivation Domain: Wildlife Protection in Uganda

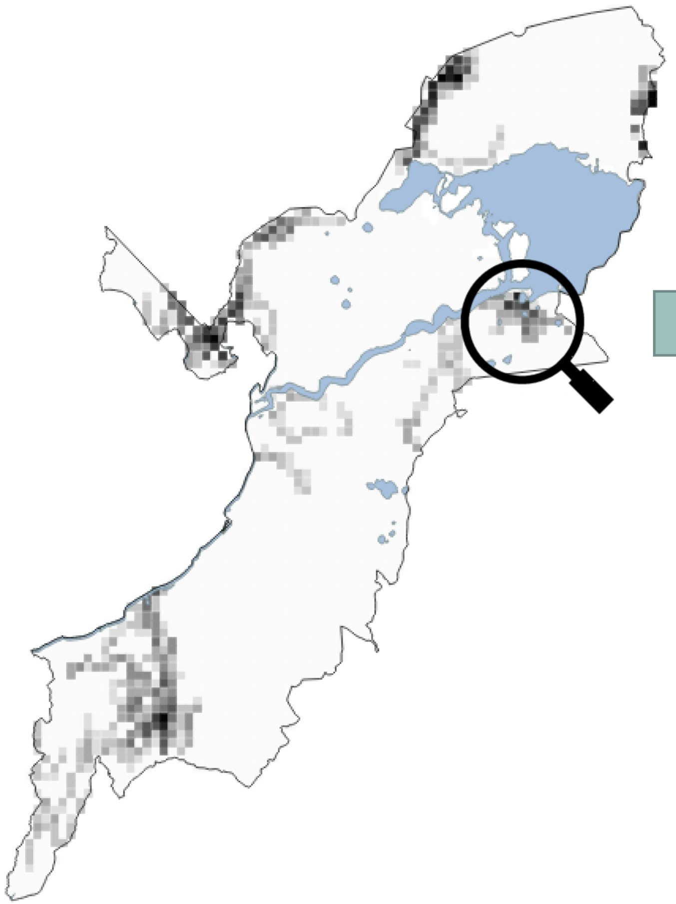

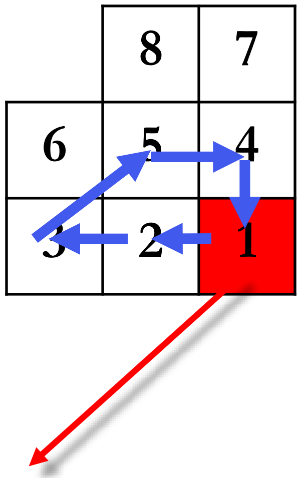

Patrol post
(one patroller)

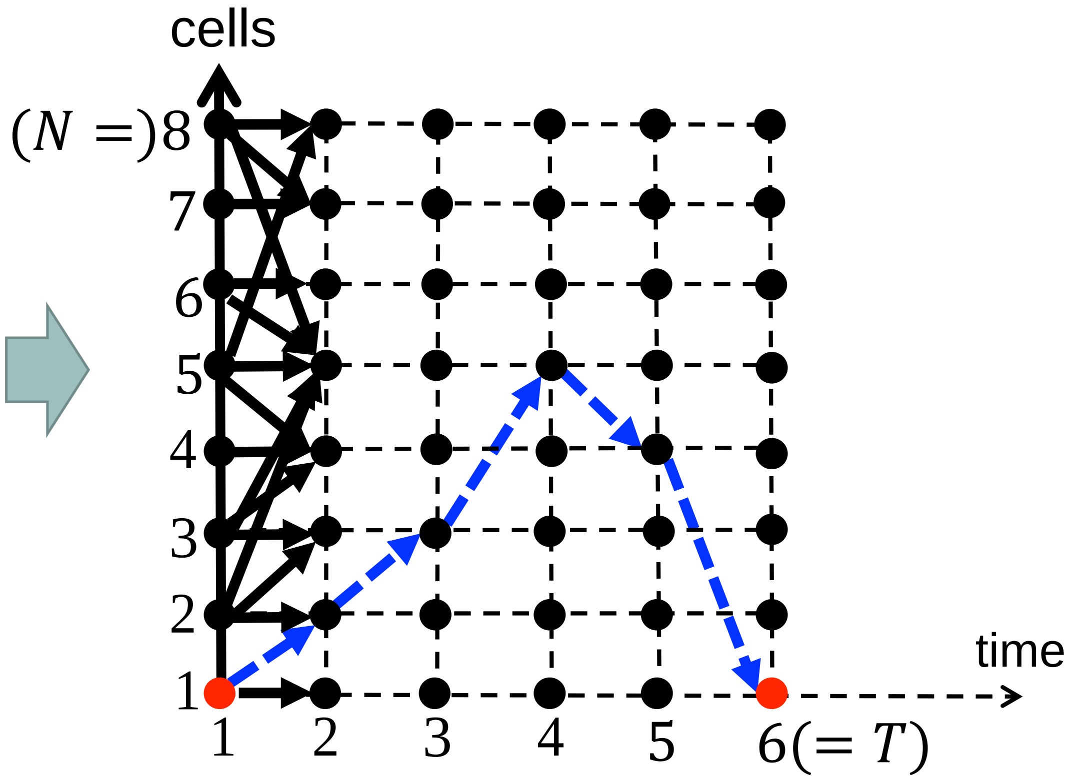

Time-unrolled Graph

Defender Strategy

Observe: a pure strategy = a path from v_{11} to v_{1T}

Claim: a mixed strategy $\Leftrightarrow$ one-unit fractional flow from v_{11} to v_{1T}

Def: patrol effort at cell i = the aggregated flow through cell i

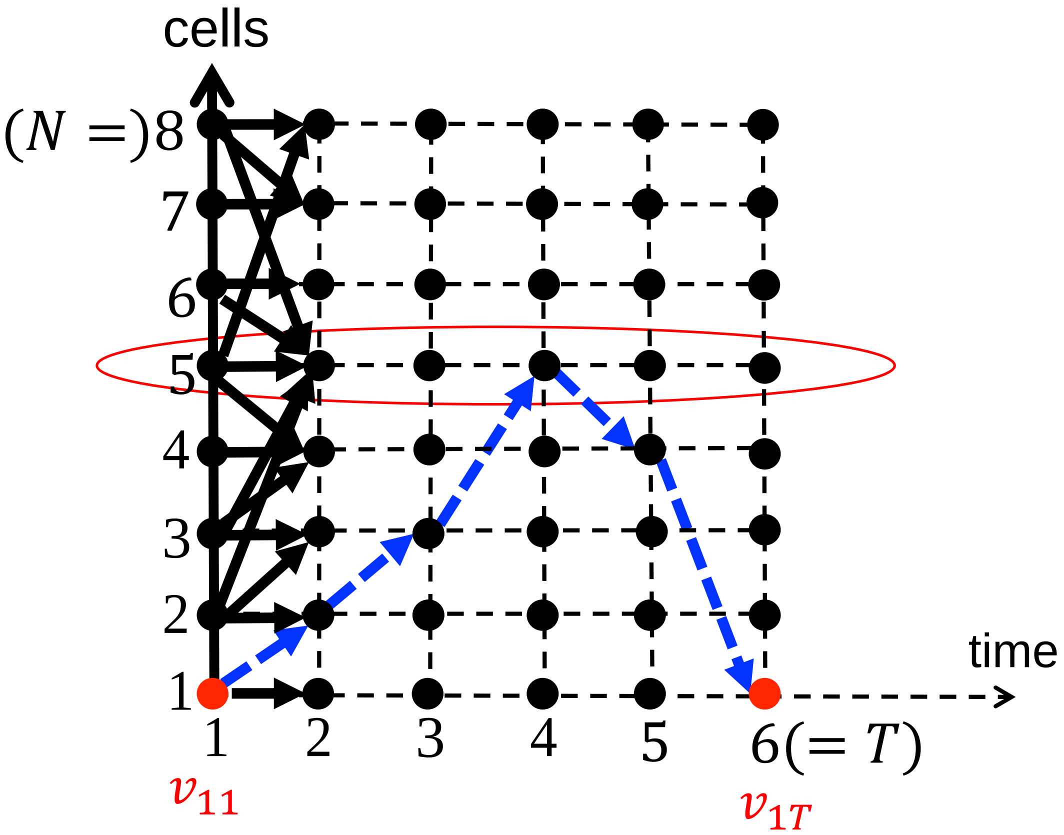

Time-unrolled Graph

Outline

- Motivation and Game Model
- Optimal Patrol Planning Against Black-Box Attackers
- Experimental Evaluation

The Single-Step Planning Task

Timeline:

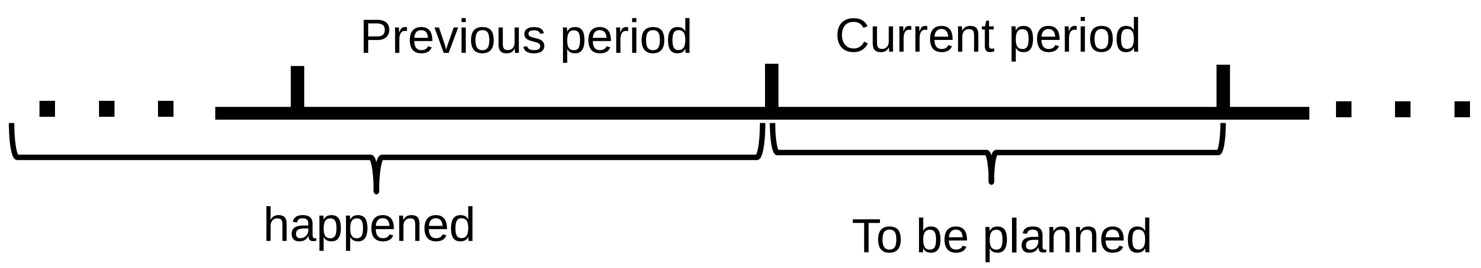

Goal: **maximize** catches of snares **against** any given attacker model

- Attacker model: (current patrolling effort + other features) → predicted snare presence

But…Many Complicated Attacker Models

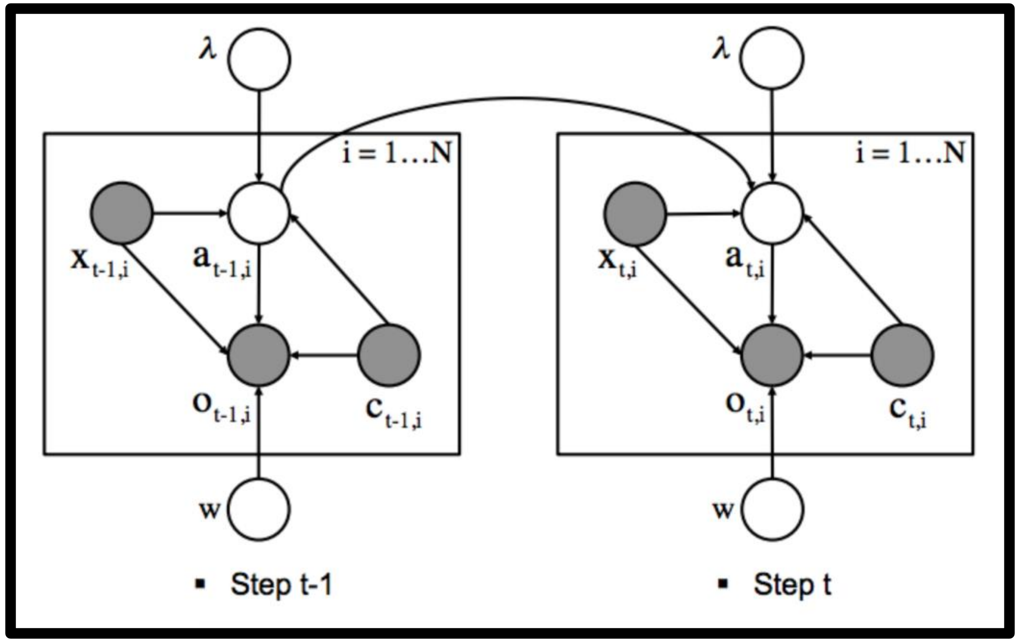

Graphical Model [Nguyen et al.’16]

But…Many Complicated Attacker Models

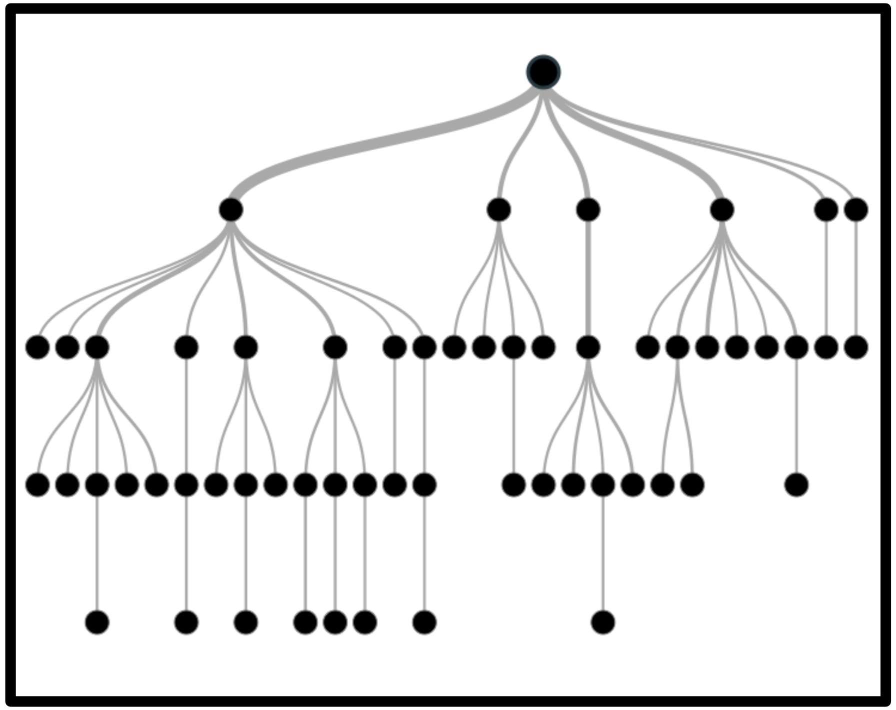

Decision Trees [Kar et al.' 17]

But…Many Complicated Attacker Models

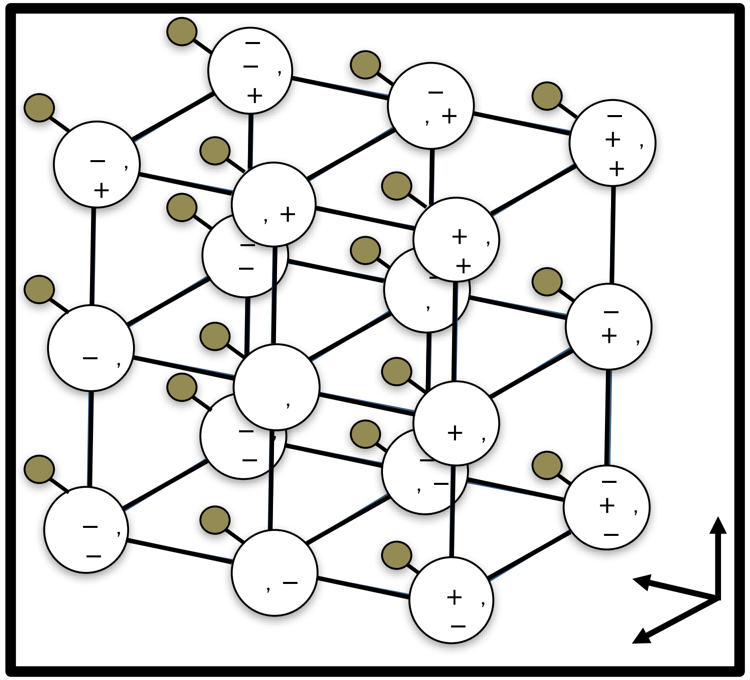

Markov Random Field [Gholami et al.'17]

But…Many Complicated Attacker Models

More are coming…

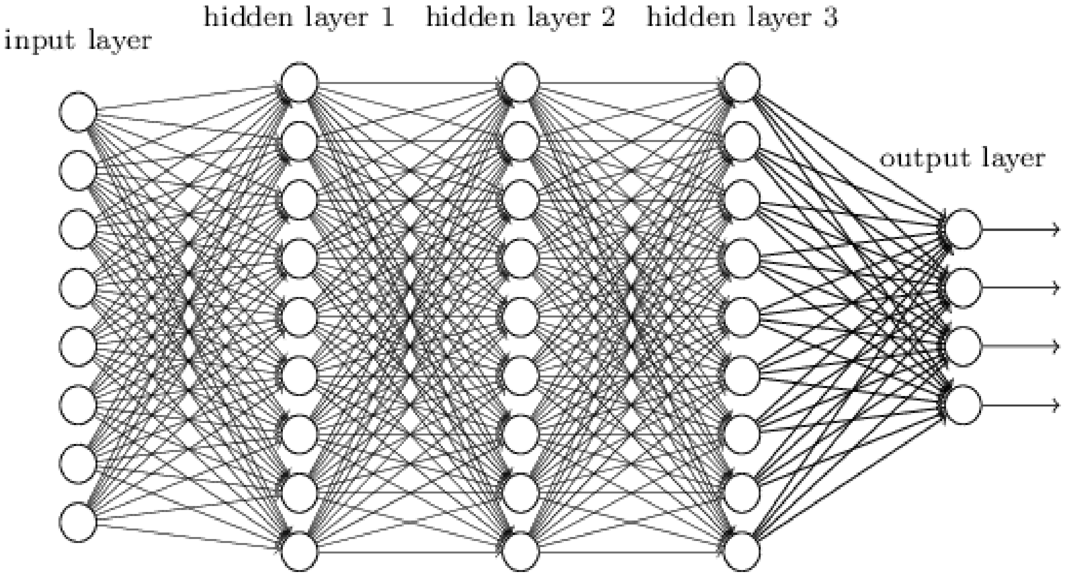

Deep Neural Networks ???

How to optimize over these complicated attacker models?

Our Idea: Treat It as a Black-Box Function

For each cell i:

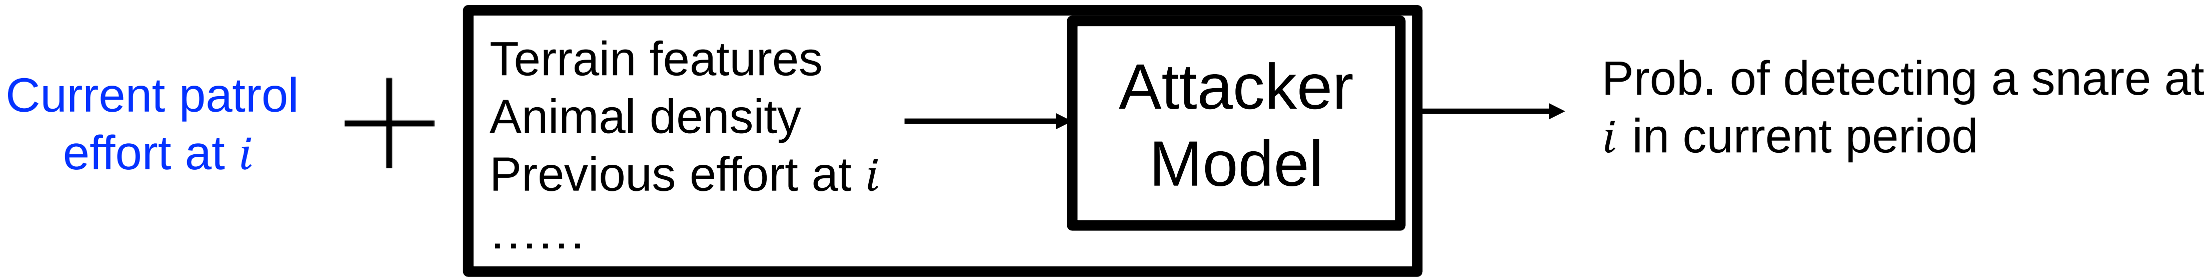

Our Idea: Treat It as a Black-Box Function

For each cell i:

Current patrol effort at i → $\boxed{g_i}$ → Prob. of detecting a snare at i in current period

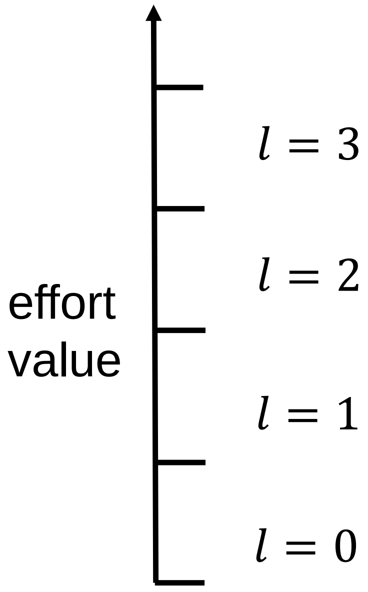

Assumption: g_i depends discretely on the current patrol effort

- Patrol levels in $\{0, 1, 2, \ldots, m\}$
  - **Thresholds** to classify patrol efforts into levels
- $g_i(0), g_i(1), \ldots, g_i(m)$ are the predicted probabilities for each level
- A good approximation when g_i is Lipchitz continuous in effort and m sufficiently large

The Optimization Task

Design patrol levels $l_1, \ldots, l_m$ (induced by patrol efforts) to

$$\text{maximize } \sum_{i=1}^{N} g_i(l_i)$$

- Main Challenge: black-box representation results in combinatorial decision making problem under constraints

NP-Hardness

Theorem: Computing optimal mixed strategy is NP-hard.

Idea: reduction from Knapsack Problem

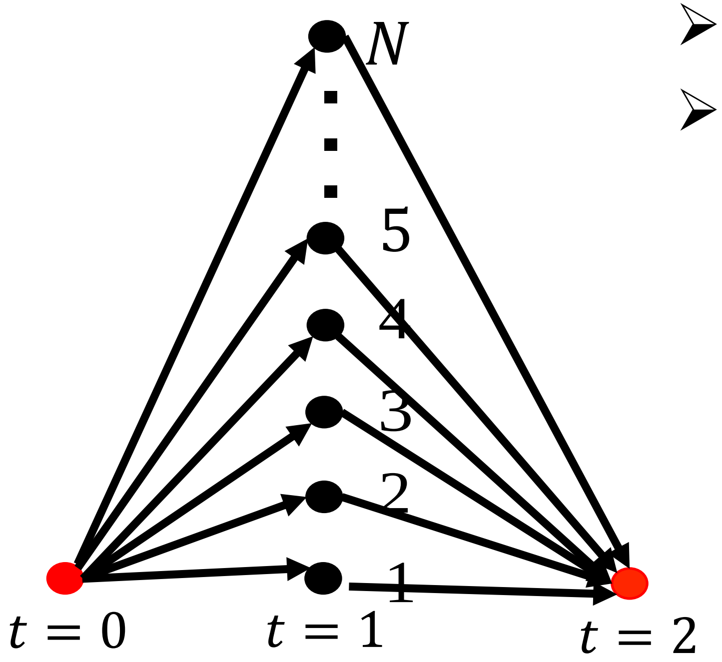

- m patrol levels with thresholds: $\alpha_0 < \alpha_1, \dots, < \alpha_m$
- $g_i(i) = p_i$ and $g_i(j) = 0, \ \forall j \ \neq i$

Goal: with 1 unit patrol budget, decide for each i to patrol with α_i (reward p_i) or patrol with 0 (reward 0)

$$\Longleftrightarrow$$

Packing m items (weight α_i, value p_i) to a 1 unit bag

Our Solution

A compact *mixed integer linear program* formulation for the optimization problem

$$
\begin{array}{lll}
\text{maximize} & \sum_{i=1}^{N} \left(g_i(0) + \sum_{j=1}^{m} z_i^j \cdot [g_i(j) - g_i(j-1)] \right) & \\
\text{subject to} & x_i \geq \sum_{j=1}^{m} z_i^j \cdot [\alpha_j - \alpha_{j-1}], & \text{for } i = 1, ..., N. \\
& x_i \leq \alpha_1 + \sum_{j=1}^{m} z_i^j \cdot [\alpha_{j+1} - \alpha_j], & \text{for } i = 1, ..., N. \\
& z_i^1 \geq z_i^2 ... \geq z_i^m, & \text{for } i = 1, ..., N. \\
& z_i^j \in \{0, 1\}, & \text{for } i = 1, ..., N, j = 1, ..., m. \\
& x_i = \sum_{t=1}^{T} \left[\sum_{e \in \sigma^+(v_{t,i})} f(e) \right], & \text{for } i = 1, ..., N. \\
& \sum_{e \in \sigma^+(v_{t,i})} f(e) = \sum_{e \in \sigma^-(v_{t,i})} f(e), & \text{for } i = 1, ..., N; t = 2, ..., T-1. \\
& \sum_{e \in \sigma^+(v_{T,1})} f(e) = \sum_{e \in \sigma^-(v_{1,1})} f(e) = 1 & \\
& 0 \leq x_i \leq 1, \qquad 0 \leq f(e) \leq 1, & \text{for } i = 1, ..., N; e \in E.
\end{array}
$$

Our Solution

A compact *mixed integer linear program* formulation for the optimization problem

- Involve a particular technique to linearize the problem
- Scalable to problems with, e.g., 100 targets and 5 patrol levels

However

- Output a mixed strategy randomizing over only a few paths
- Unavoidable – efficient solvers are *designed to* find small-support solutions
- Vulnerable to attacker's (partial) real-time surveillance

Add Extra Randomness by Entropy Maximization

- Many mixed strategies implement the same patrolling effort

- We compute the one that maximizes (Shannon) entropy
  - Usually support on a much larger set of paths
  - Difficult to learn

- There is an efficient algorithm to compute max-entropy distribution here
  - Convex analysis, combinatorial optimization, duality theory

Extension: Multi-Step Planning

Timeline:

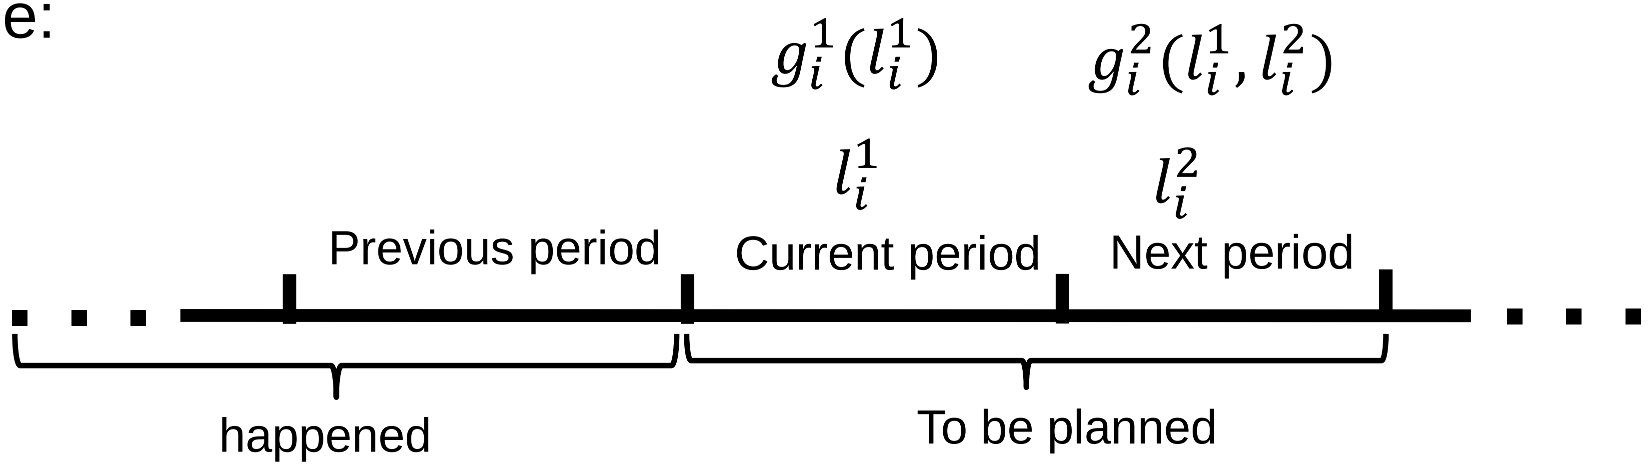

Goal: maximize aggregated total catch

$$\text{maximize} \ \sum_{i=1}^{N} g_i^2(l_i^2, l_i^1) + \sum_{i=1}^{N} g_i^1(l_i^1)$$

Outline

- Motivation and Game Model
- Optimal Patrol Planning Against Black-Box Attackers
- **Experimental Evaluation**

Real-World Data Set from QEPA

Rangers record captures of snares

- From 2003 – 2017
- 39 patrol posts
- We test on post 11, 19, 24 (the mostly attacked)

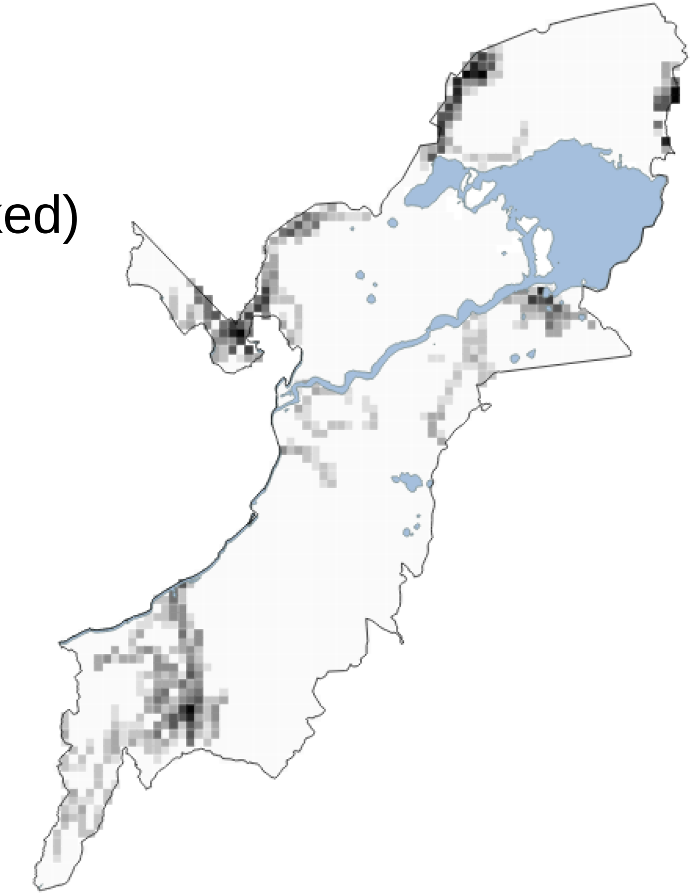

Collaborators: Wildlife Conservation Society, Uganda Wildlife Authority,

Experiment 1: Compare with Baseline Algorithms

OPERA: bagging ensemble model [Gholami et al.'17]

(two levels: low and $high$)

OPP: **O**ptimal **P**atrol **P**lanning

Another Two Baselines

- GREED: greedily pick the next reachable cell to patrol
- RAND: randomly pick the next reachable cell to patrol

Experiment 1: Compare with Baseline Algorithms

	#Detection	#Cover	#Routes	Entropy
OPERA	15/19	20/47	61	4.0
OPP	15/19	20/47	10	2.0
GREED	5/19	4/47	84	4.4
RAND	4/19	6/47	89	4.5

Comparisons of Different Criteria for Patrol Post 11

- #Detection: a/b → out of b predicted attacks, the algorithm detects a attacks
- #Cover: a/b → out of b cells, a of them are covered with $high$

Experiment 1: Compare with Baseline Algorithms

	#Detection	#Cover	#Routes	Entropy
OPERA	6/6	24/72	22	2.6
OPP	6/6	24/72	6	1.3
GREED	2/6	2/72	1	0
RAND	2/6	6/72	90	4.5

Comparisons of Different Criteria for Patrol Post 19

Experiment 2: Compare with Past (Real) Patrolling

Criteria	Post 11		Post 19		Post 24	
	OPERA	Past	OPERA	Past	OPERA	Past
#Detections	15/19	4/19	6/6	5/6	4/4	3/4
#Cover	20/47	6/47	24/72	11/72	20/59	14/59

Take-Away Message

- An efficient patrol planning tool that
  - Optimize against very general class of attacker models
  - Mitigate attacker real-time surveillance by adding extra randomness

Special Thanks to Wildlife Conservation Society, Uganda Wildlife Authority

Thank You