SAND2012-8685P

Security Games with Limited Surveillance

Bo An[1], David Kempe[1], Christopher Kiekintveld[2], Eric Shieh[1],
Satinder Singh[3], Milind Tambe[1], Yevgeniy Vorobeychik[4]

[1]University of Southern California
[2]University of Texas, El Paso
[3]University of Michigan
[4]Sandia National Laboratories

* Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Motivation: Game Theory for Security

- Limited security resources
- Adversary monitors defenses, exploits patterns

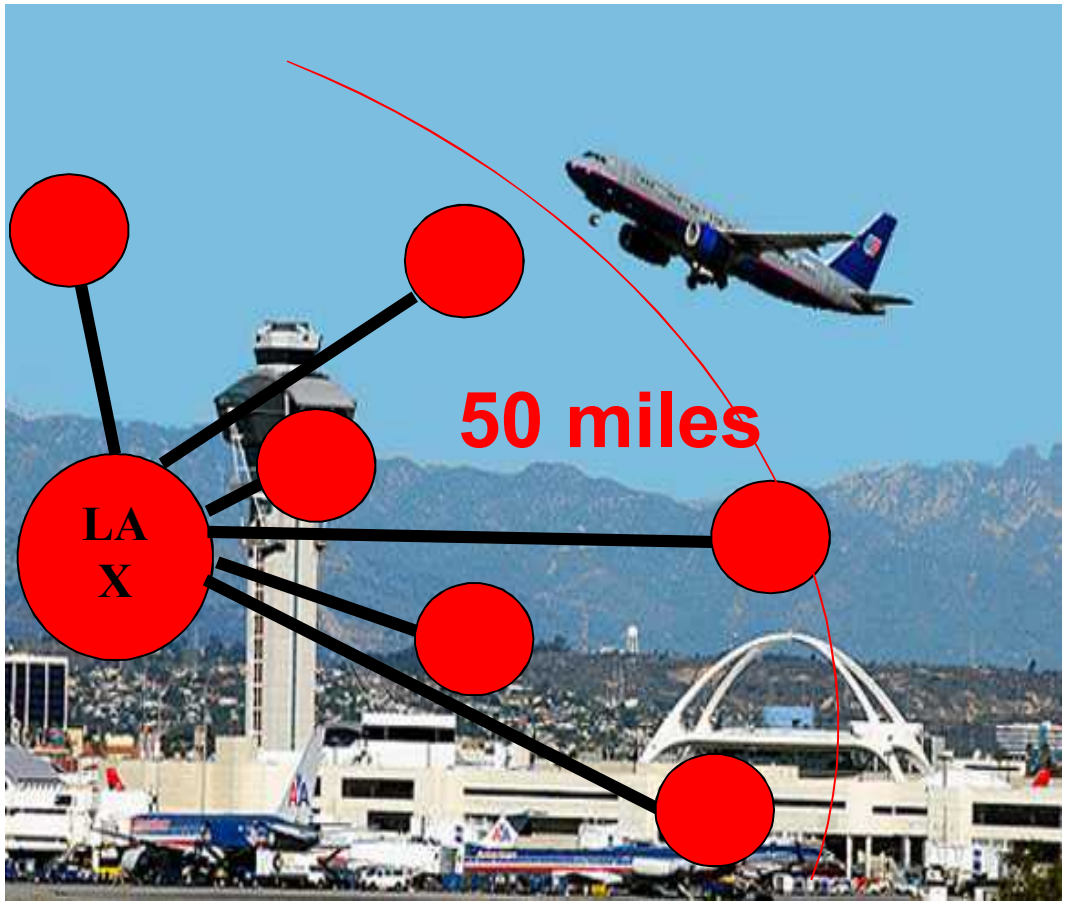

Game Theory: Bayesian Stackelberg Games

- Security allocation: (i) Target weights; (ii) Opponent reaction
- *Stackelberg*: Security forces commit first
- *Bayesian*: Uncertain adversary types
- *Optimal security allocation:* Weighted random
- **Strong Stackelberg Equilibrium (Bayesian)**
  - **NP-hard**

Adversary

	Target #1	Target #2
Target #1	5, -3	-1, 1
Target #2	-5, 5	2, -1

Sandia National Laboratories

Problem Statement

How does the attacker
conduct surveillance and build belief ?

- **Previous work:**
  - ***Perfect surveillance***
  - ***Imperfect surveillance***
    - **COBRA [Pita et al. 2010], RECON [Yin et al. 2011]**
    - **Rely on hand-tuned parameters, cannot reason about surveillance cost**

Our Contributions

- Security game with limited observation model
- Analysis
- Algorithms
  - Compute optimal defender strategy
  - Estimate attacker observation length
- Experiments

Security Games with Limited Surveillance

will make τ observations

decide mixed strategy X

observe τ pure strategies O, update belief, attack the best target

Sandia National Laboratories

Illustrating Example

- **The scenario:**
  - **3 targets t_1, t_2, t_3: 3 attacker pure strategies**
  - **1 resource: 3 defender pure strategies A_1, A_2, A_3**
- **Defender strategy $X=<x_1, x_2, x_3>$**
  - **stationary**
- **Attacker has a prior belief**
  - **e.g., uniform distribution**
- **Attacker makes τ (e.g., 5) observations**
  - **e.g., $O= A_1, A_2, A_1, A_2, A_3$**
  - **compact representation: $O=<o_1, o_2, o_3>=<2, 2, 1>$**
- **Best response based on posterior belief**

Defender's Optimization Problem

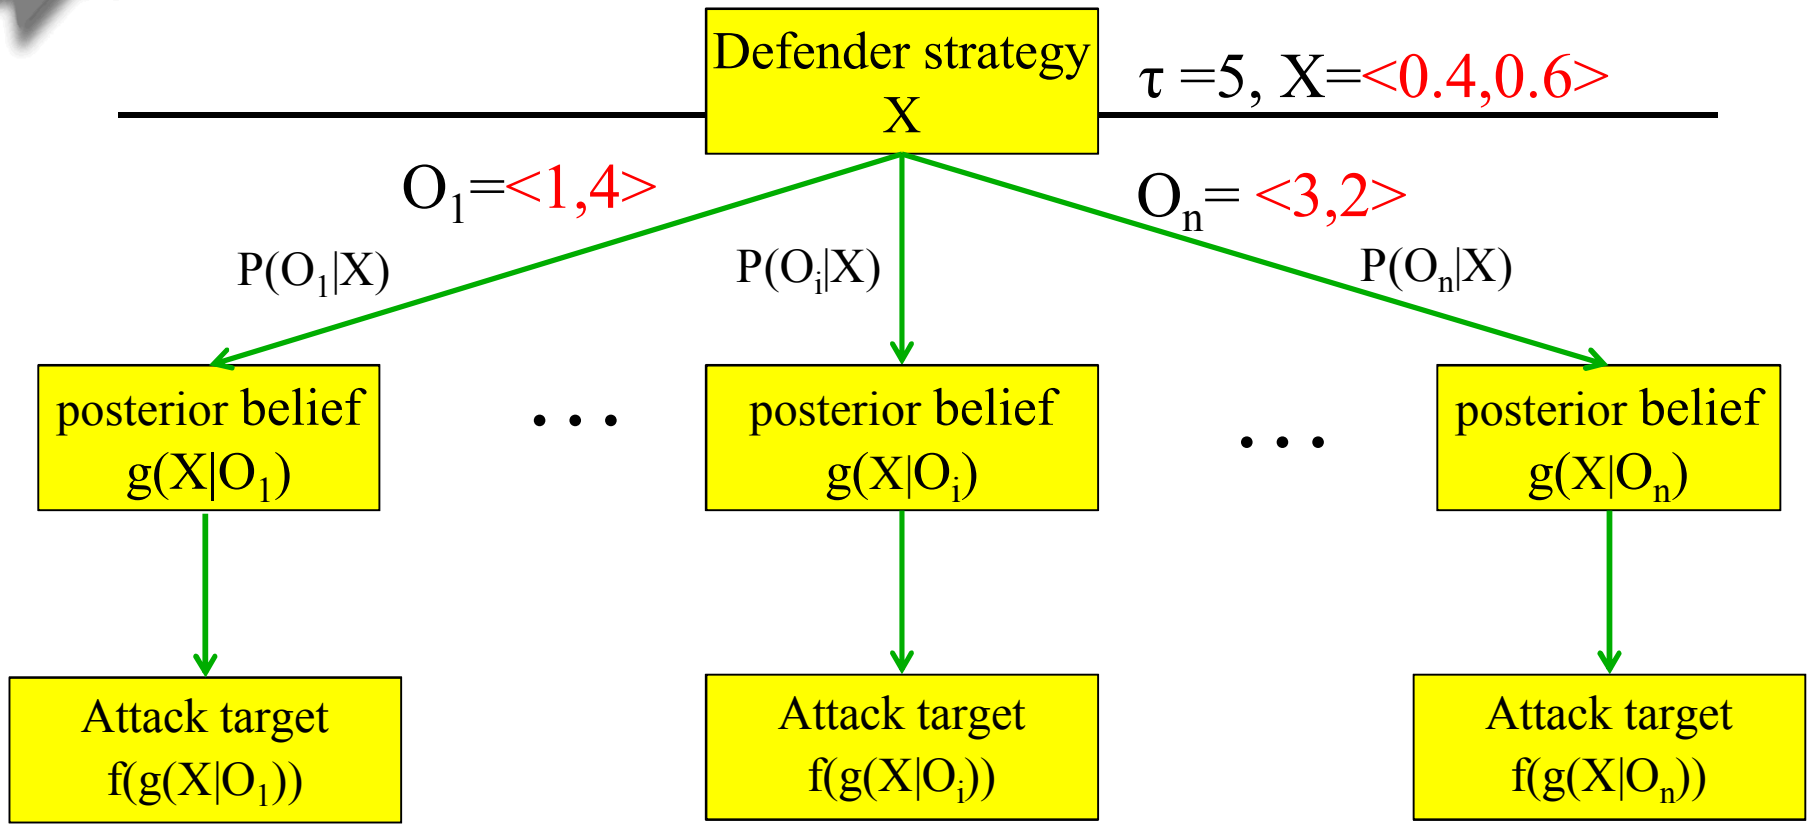

$$\max_X \sum_O P(O|X)U(X,f(g(X|O)))$$

Sandia National Laboratories

Optimal Defender Strategy

$$\max \sum_{\mathbf{o} \in \mathcal{O}_\tau} \frac{\tau!}{\prod_{A \in \mathcal{A}} o_A!} \prod_{A \in \mathcal{A}} (x_A)^{o_A} d^{\mathbf{o}}$$

$$\textbf{s.t.} \quad x_A \in [0,1] \qquad \forall A \in \mathcal{A}$$

$$\sum_{A \in \mathcal{A}} x_A = 1$$

$$c_i = \sum_{A \in \mathcal{A}} x_A A_i \qquad \forall i \in T$$

$$d^{\mathbf{o}} = c_{\psi(\mathbf{o})} \left(R^d_{\psi(\mathbf{o})} - P^d_{\psi(\mathbf{o})} \right) + P^d_{\psi(\mathbf{o})} \qquad \forall \mathbf{o} \in \mathcal{O}_\tau$$

Sandia National Laboratories

Analysis

Defender utility is always monotone non-increasing with larger τ ✗

More targets protected with larger τ ✗

Safe targets should always not be protected ✓

Unsafe targets should always be protected ✗

Estimate Attacker Observation Length

$$\operatorname*{argmax}_{\tau}\left(\sum_{\mathbf{o}\in\mathcal{O}_\tau}\frac{\tau!}{\prod_{A\in\mathcal{A}}o_A!}\prod_{A\in\mathcal{A}}x_A^*(\tau)^{o_A}k^{\mathbf{o}}-\lambda\cdot\tau\right)$$

$\mathbf{x}^*(\tau)$: defender's optimal strategy

$k^{\mathbf{o}}$: attacker's utility when he observes o and the defender plays $\mathbf{x}^*(\tau)$ strategy

Sandia National Laboratories

Experimental Results

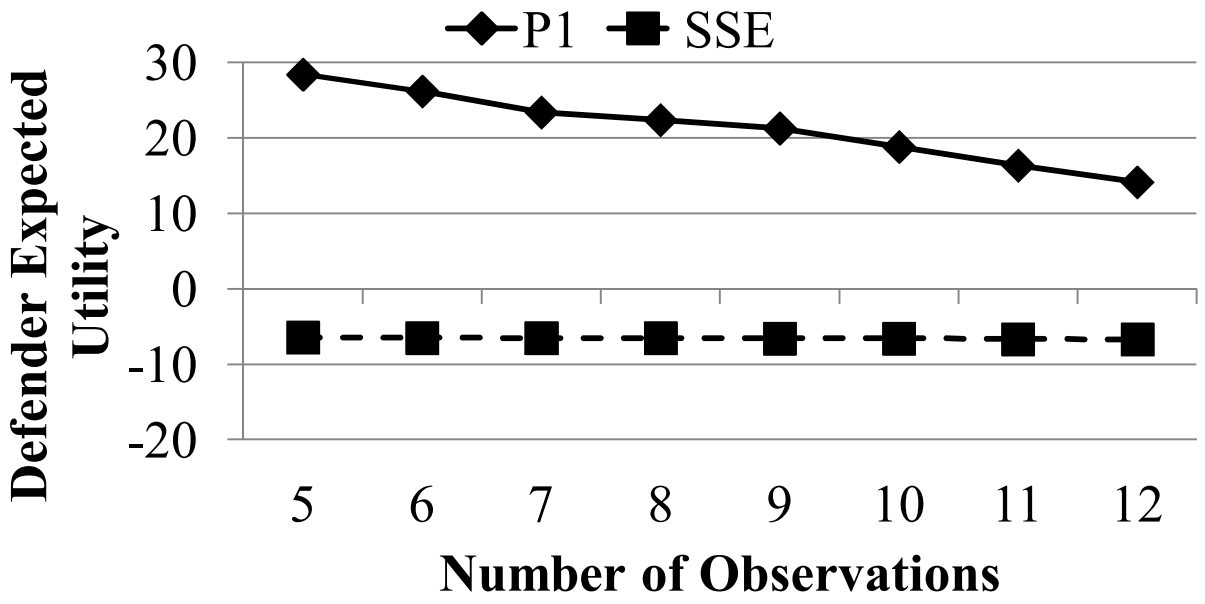

Sandia National Laboratories

Conclusions

- **Contributions**
  - Security game with limited observation model
  - Analysis
  - Algorithm for computing optimal defender strategy
  - Algorithm for estimating observation length

- **Future work**
  - Stopping problem
  - Human-subject experiments
  - Scalability