Combining game theory and statistical learning for security, privacy and networked systems

Patrick Loiseau

HDR defense

EURECOM, December 2016

Digital world opportunities and threats

- The Internet brings many opportunities to build useful services
  - Social medias, forums, daily apps (running, shopping, taxi), surveys, medical services, etc.

- Based on personal data
  - Exploited through machine learning

- But also important threats and issues
  - Security
  - Privacy
  - Reliability / performance

My approach: methodology combining game theory and statistical learning

- Security, privacy and performance are strongly impacted by strategic behavior of humans
  - Need to take into account incentives

➢ Game theory: mathematical tool to model users interactions and incentives

- Statistical learning: at the core of services based on personal data (privacy and security)

➢ Combination of game theory and statistical learning to design better digital systems

Contributions areas and types

- Development of models/methods/theoretical results combining game theory and statistical learning for…

  1. Security
  2. Privacy
  3. Networked systems

Theory ←→ Applications

- Other works not covered in this HDR:
  - Large deviations [Stoc. Proc. Appl. '11]
  - Heart-rate analysis [Physica A '12]
  - Resource provisioning [IEICE Trans on Com '12], Internet cooperation [IJCS '16]

Roadmap

- Game theory and statistical learning for
  - Security
  - Privacy
  - Networked systems
- Perspectives

Roadmap

- Game theory and statistical learning for
  - Security
  - Privacy
  - Networked systems

- Perspectives

Overview of security games

Defender

chooses *def*

to maximize

$Utility_{Defender}(def, att)$

Attacker

chooses *att*

to maximize

$Utility_{Attacker}(def, att)$

– Two players game modeling interaction attacker/defender

– Strategies (att, def) and utilities depend on the particular scenario at stake

  - Learning algorithm for defense, defense resource allocation

➔ The game solution helps building better defenses

Summary of my contributions in security

- A game-theoretic study of adversarial classification

Key papers: [CDC ’12, GameSec ’12, ArXiv ’16]
Key collaborations: UC Santa Cruz
1 student unofficially advised

- A new solution of the Blotto game (resource allocation)

Key papers: [Netgcoop ’14]
Key collaborations: UC Berkeley

- Regret minimization in repeated games with discounted losses

Key papers: [StonyBrooks ’16/ArXiv ’16]
Key collaborations: UC Berkeley
1 intern

Summary of my contributions in security

- **A game-theoretic study of adversarial classification**

  ***Key papers*: [CDC '12, GameSec '12, ArXiv '16]**
  ***Key collaborations*: UC Santa Cruz**
  1 student unofficially advised

- A new solution of the Blotto game (resource allocation)

  Key papers: [Netgcoop '14]
  Key collaborations: UC Berkeley

- Regret minimization in repeated games with discounted losses

  Key papers: [StonyBrooks '16/ArXiv '16]
  Key collaborations: UC Berkeley
  1 intern

Attack detection through classification

- Need to differentiate attacks from normal behavior
  - Spam detection, malware detection, fraud detection, etc.

- Standard tools from supervised machine learning
  - Logistic regression, SVM, Naive Bayes, etc.

Cats

vs

Dogs

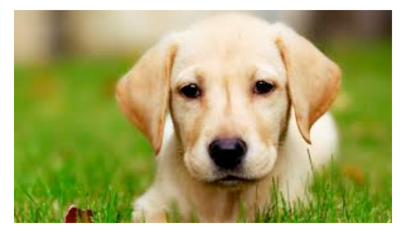

Cat or dog?

- In security: [dog=normal, cat=attack]
  - Looking for best features, implementing/testing in real life

Key limitation of supervised learning in security

- Standard learning algorithms based on "iid assumption"

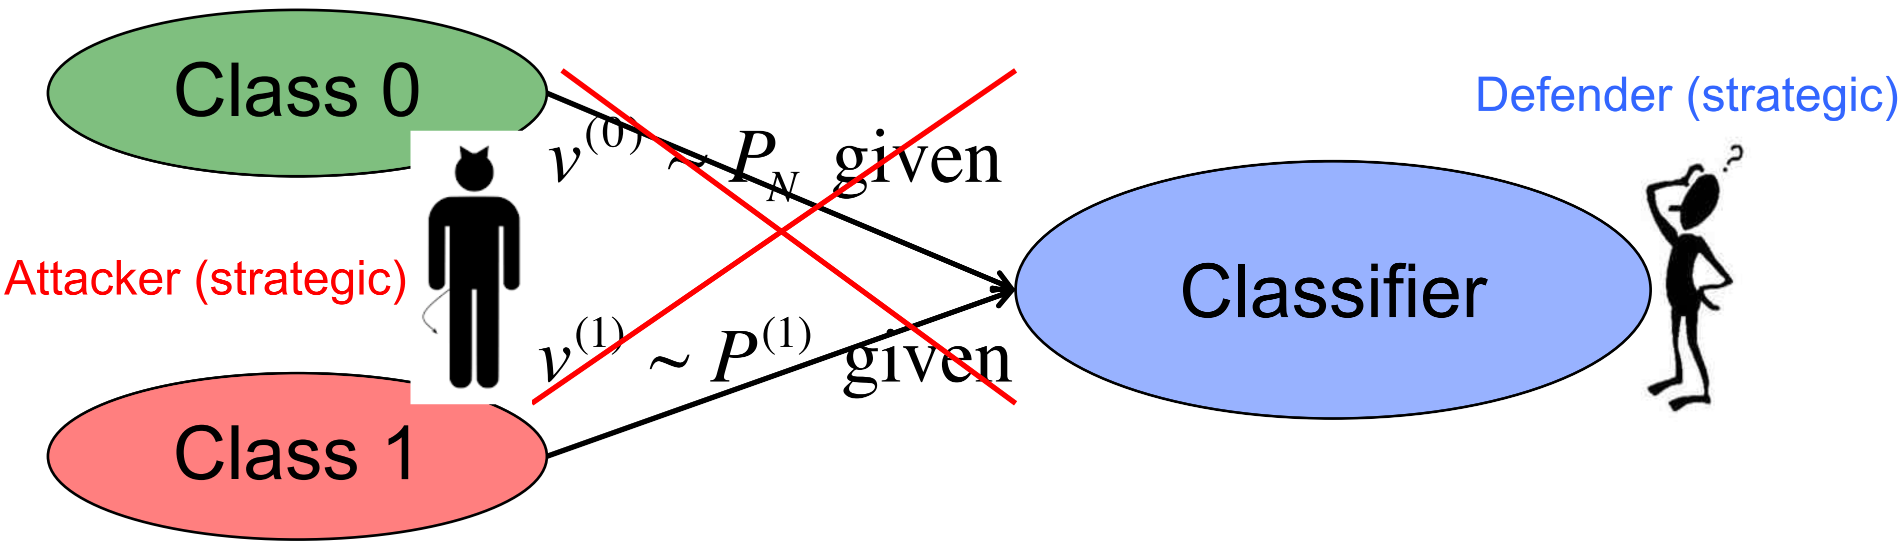

- Security: data generated by an adversary
  ➔ iid assumption fails, standard algorithms work poorly

➔ How to learn in these situations?
What can game-theory bring to this question?

EURECOM Sophia Antipolis

Literature & contribution

- Large literature on "adversarial learning"
  [Dalvi et al. '04], [Lowd, Meek '05], [Globerson, Roweis '06], [Huang, Biggio, Nelson, Laskov, Barreno, Joseph, Rubinstein, Tygar et al. '08-'15], [Wang, Zhao et al. '14], [Zhou, Kantarcioglu et al. '12-'14], [Vorobeychik, Li '14-'15], …
  - Simple, worst-case solutions
  - Proposes randomization as defense but without justification
- Large literature on game-theory for security
  [Alpcan Basar, CUP 2011], [Alpcan, Basar, CDC '04, Int Symp Dyn Games '06], [Zhu et al., ACC '10], [Liu et al, Valuetools '06], [Chen, Leneutre, IEEE TIFS '09], [Tambe et al. '09-'15], …
  - Simple payoff, no learning
- Our work:
  - Flexible game-theoretic model of classification
  - Game solution ➔ insights on "how to learn"

Model: players and actions

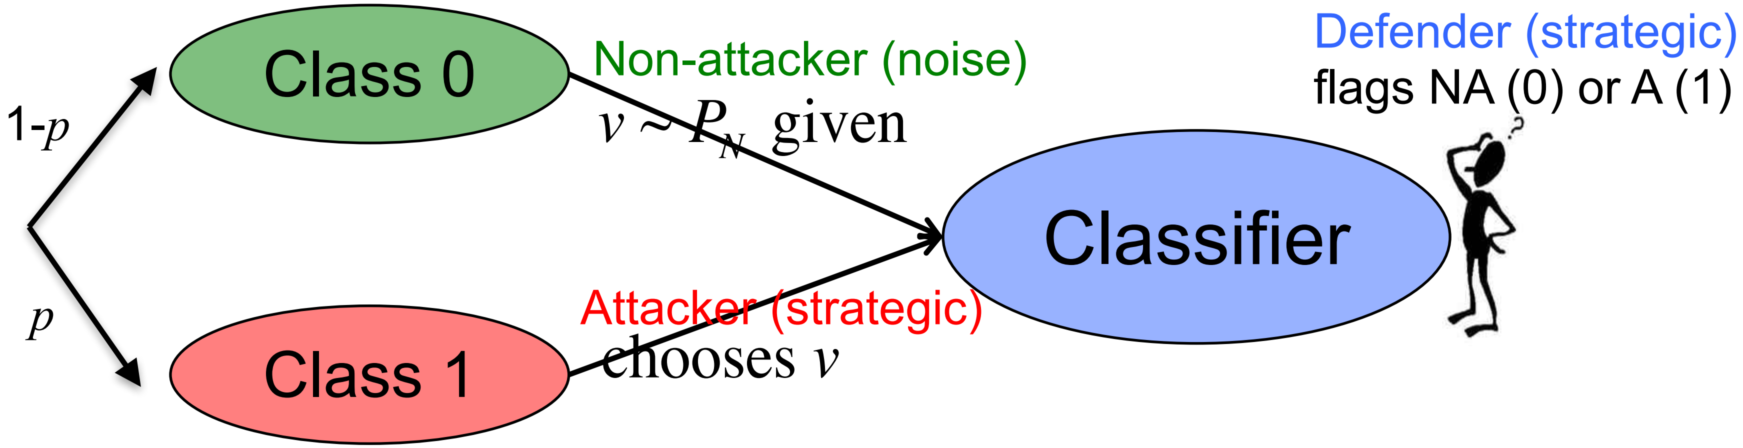

- Attacker chooses $v \in V$ → Set of feature vectors
- Defender chooses $c \in C$ → Set of classifiers $\{0,1\}^{|V|}$
  - Classifier $c : V \rightarrow \{0,1\}$
- Two-players game $G = \langle V, C, \overbrace{P_N, p, c_d, c_{fa}}^{\text{Payoff-relevant Parameters}} \rangle$

Model: payoffs

- Attacker's payoff:

cost if detected

$$U^A(v,c) = R(v) - c_d 1_{c(v)=1}$$

reward from attack

- Defender's payoff:

cost of false alarm

$$U^D(v,c) = p\left(-R(v) + c_d 1_{c(v)=1}\right) + (1-p)c_{fa}\left(\sum_{v' \in V} P_N(v') 1_{c(v')=1}\right)$$

Rescaling

$$U^D(v,c) = -U^A(c,v) + \frac{(1-p)}{p} c_{fa}\left(\sum_{v' \in V} P_N(v') 1_{c(v')=1}\right)$$

EURECOM Sophia Antipolis

Nash equilibrium in the classification game

- Mixed strategies:
  - Attacker: probability distribution α on V
  - Defender: probability distribution β on C

- Utilities extended: $U^A(\alpha,\beta) = \sum_{v\in V}\sum_{c\in C}\alpha_v U^A(v,c)\beta_c$

- Nash equilibrium: (α^*,β^*) s.t. each player is at best-response:

$$\alpha^* \in \underset{\alpha}{\operatorname{argmax}}\, U^A(\alpha,\beta^*)$$

$$\beta^* \in \underset{\beta}{\operatorname{argmax}}\, U^D(\alpha^*,\beta)$$

Best-response equivalence to a zero-sum game

$$U^A(v,c) = R(v) - c_d 1_{c(v)=1} - \frac{(1-p)}{p} c_{fa} \left(\sum_{v' \in V} P_N(v') 1_{c(v')=1} \right)$$

$$U^D(v,c) = -R(v) + c_d 1_{c(v)=1} + \frac{(1-p)}{p} c_{fa} \left(\sum_{v' \in V} P_N(v') 1_{c(v')=1} \right)$$

- The non-zero-sum part depends only on $c \in C$
- Best-response equivalent to zero-sum game
- Solution can be computed by LP, BUT
  - The size of the defender's action set is large
  - Gives no information on the game and solution structure

Main result 1: defender combines features based on attacker's reward

- Define C^T: set of threshold classifiers on $R(v)$

$$C^T = \left\{ c \in C : c(v) = 1_{R(v)\geq t}\ \forall v, \text{ for some } t \in \Re \right\}$$

Theorem:

For every NE of $G = \langle V, C, P_N, p, c_d, c_{fa} \rangle$, there exists a NE of $G^T = \langle V, C^T, P_N, p, c_d, c_{fa} \rangle$ with the same attacker's strategy and the same equilibrium payoffs

- Classifiers that compare $R(v)$ to a threshold are optimal for the defender
  - Different from know classifiers (logistic regression, etc.)
  - Reduces a lot the size of the defender's strategy set

Main result 1: proof's key steps

1. The utilities depend on β only through the probability of class 1 classification:

$$\pi_d^{\beta}(v) = \sum_{c \in C} \beta_c 1_{c(v)=1}$$

1. At NE, if $P_N(v) > 0$ for all v, then
   $\pi_d^{\beta}(v)$ increases with $R(v)$

2. Any $\pi_d^{\beta}(v)$ that increases with $R(v)$ can be achieved by a mix of threshold strategies in C^T

Main result 2: Nash equilibrium structure

Theorem:

At a NE of $G^T = \langle V, C^T, P_N, p, c_d, c_{fa} \rangle$, for some k:

- The attacker's strategy is $\left(0,\cdots,0,\alpha_k,\cdots,\alpha_{|V|}\right)$
- The defender's strategy is $\left(0,\cdots,0,\beta_k,\cdots,\beta_{|V|},\beta_{|V|+1}\right)$

where $$\beta_i = \frac{r_{i+1}-r_i}{c_d}, \text{ for } i \in \left\{k+1,\cdots,|V|\right\} \qquad \left(r_i = R(v_i) < r_{i+1} = R(v_{i+1})\right)$$

$$\alpha_i = \frac{1-p}{p}\frac{c_{fa}}{c_d}P_N(v_i), \text{ for } i \in \left\{k+1,\cdots,|V|-1\right\}$$

EURECOM Sophia Antipolis

Nash equilibrium illustration

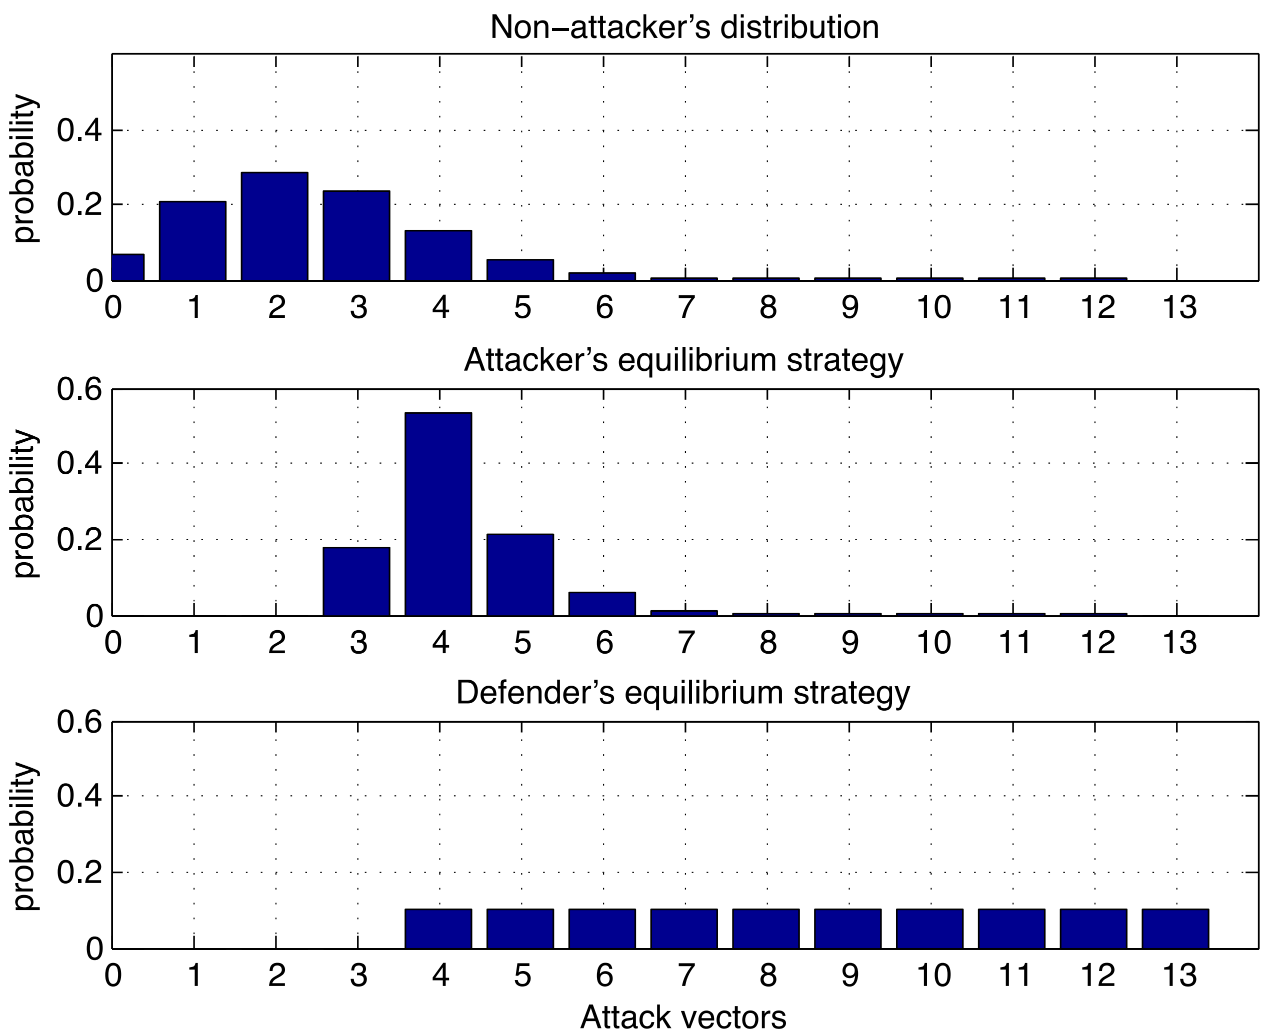

- Case

$$r_i = i \cdot c_a$$

NE computation

- Defender: try all vectors β of the form (for all k)

β:
Mix of defender threshold strategies

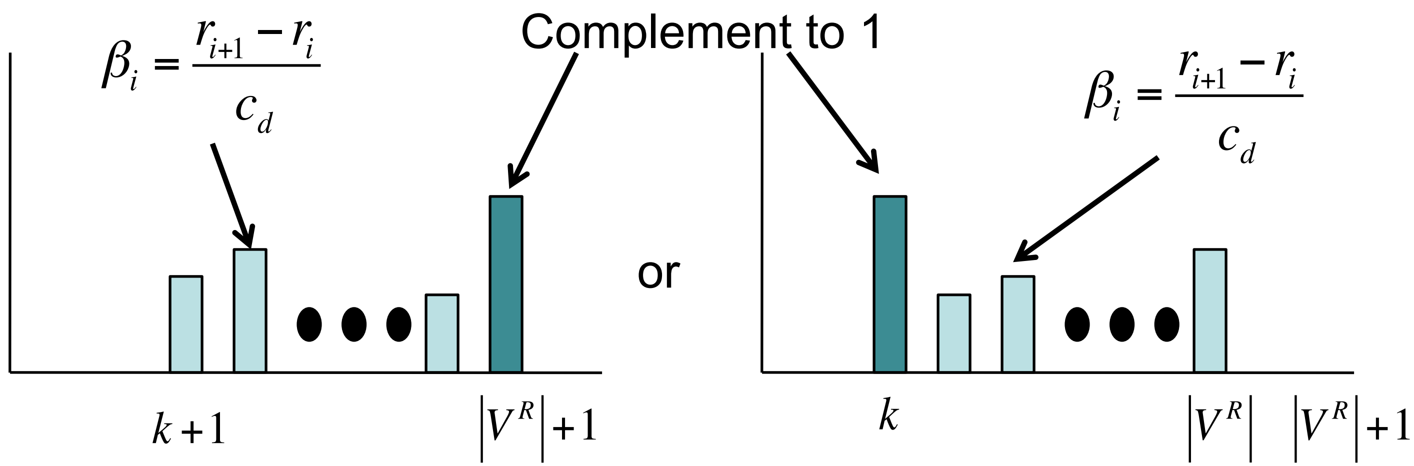

- Take the one maximizing payoff
  - Unique maximizing β → unique NE.
  - Multiple maximizing β → any convex combination is a NE
- Attacker: use the formula
  - Complete first and last depending on β

Main result 2: proof's key steps

1. Matrix formulation $U^A(\alpha,\beta) = -\alpha'\Lambda\beta$ and $U^D = \alpha'\Lambda\beta - \mu'\beta$

2. At NE, β is solution of LP:

$$\begin{aligned} &\text{maximize} && z - \mu'\beta \\ &\text{s.t.} && \Lambda\beta \geq z \cdot 1_{|V^R|},\ \beta \geq 0,\ 1_{|V^R|+1} \cdot \beta = 1 \end{aligned}$$

- extreme points of $\Lambda x \geq 1_{|V^R|},\ x \geq 0 \quad \left(\beta = x/\|x\|\right)$

3. Look at polyhedron and eliminate points that are not extreme

$$\begin{aligned} c_d x_1 + (r_{|V^R|} - r_1 + \varepsilon)\|x\| &\geq 1 \\ &\vdots \\ c_d(x_1 + x_2 + \cdots + x_{|V^R|}) + \varepsilon\|x\| &\geq 1 \end{aligned}$$

Summary: binary classification from strategic data

- Simple game model of classification from strategic data

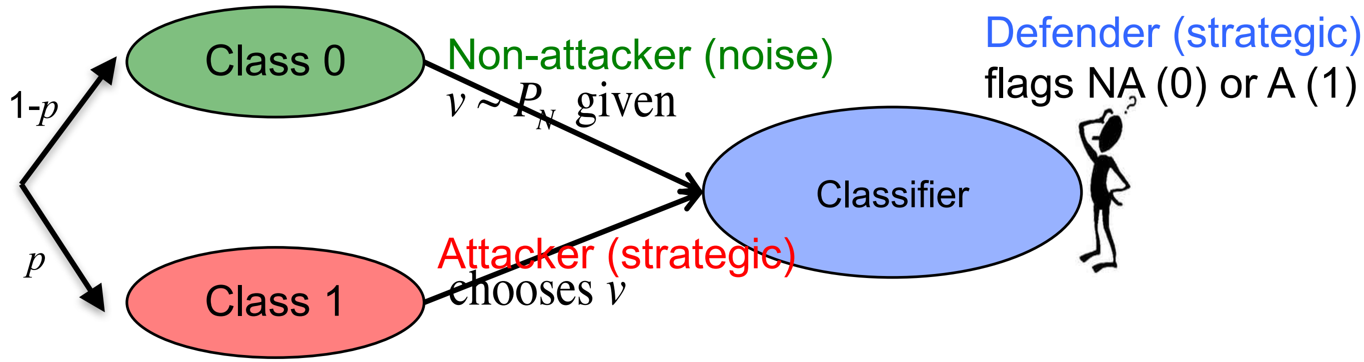

- Nash equilibrium brings insights on learning question:
  - Defender: combine features according to attacker's reward
    - Mix on thresholds prop. to marginal reward, up to highest threshold
  - Attacker: mimic non-attacker on defender's support
  - Answer questions: "is it worth investing in extra sensors?"
- Preliminary results for more complex scenarios

EURECOM
Sophia Antipolis

Roadmap

- Game theory and statistical learning for
  - Security
  - Privacy
  - Networked systems
- Perspectives

Overview of my research in privacy

- Users revealing data are worried about privacy losses

1. Mechanisms to learn better from personal data while allowing users to reveal less data
   - A new game-theoretic model treating information as a public good

Key papers: [WINE ’13, FC ’15, CSF ’15, SING ’15/ArXiv ’16]
Key collaborations: Technicolor, Northeastern, PennState
1 postdoc

2. Estimation of privacy risk from data already public
   - Matching user profiles across multiple sites

Key papers: [KDD ’15, ArXiv ’16]
Key collaborations: MPI-SWS
1 student

Overview of my research in privacy

- Users revealing data are worried about privacy losses

1. **Mechanisms to learn better from personal data while allowing users to reveal less data**
   - **A new game-theoretic model treating information as a public good**

***Key papers*: [WINE ’13, FC ’15, CSF ’15, SING ’15/ArXiv ’16]**
***Key collaborations*: Technicolor, Northeastern, PennState**
1 postdoc

2. Estimation of privacy risk from data already public
   - Matching user profiles across multiple sites

Key papers: [KDD ’15, ArXiv ’16]
Key collaborations: MPI-SWS
1 student

How to learn from personal data?

- Personal data is special:
  - Privacy concerns, revealed by privacy-conscious human beings
- ➢ Large literature on incentives through payments

- Users reveal data without being paid, because they have an interest in the learning result
  - ➢ Learning outcome (information) is a public good
- ➢ Personal data is strategic!
  - ➢ How much can we learn? At which privacy cost?
  - ➢ Can we increase learning accuracy without payment?
  - ➢ How to find optimal learning algorithm?

Model (1): linear model of user data

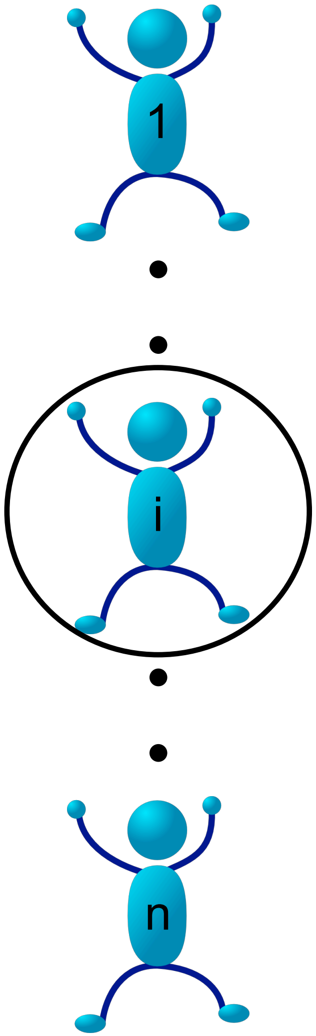

$$y_i = \beta^T x_i + \varepsilon_i$$

Private data of *i*
$\in R$

Model parameter
$\in R^d$
(unknown)

Public features of *i*
$\in R^d$

Inherent noise of *i*
mean 0, variance σ^2

User *i* adds noise

Added noise of *i*
mean 0, variance σ_i^2

$$\tilde{y}_i = \beta^T x_i + \underbrace{\varepsilon_i + z_i}$$

Data reported by *i*
$\in R$

Total noise of *i*
mean 0, variance $\sigma^2 + \sigma_i^2$

Model (2): analyst's learning

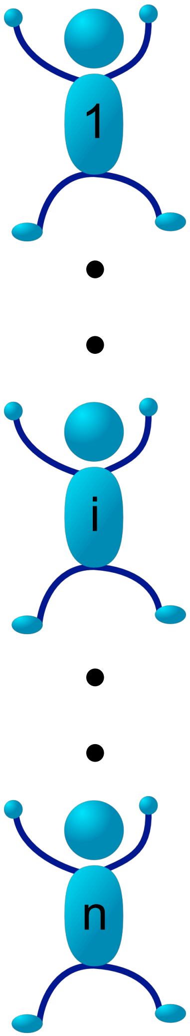

$(n \times 1)$ vector of reported data

$(n \times d)$ matrix of public features

$$\hat{\beta} = \left(X^T \Lambda X\right)^{-1} X^T \Lambda \tilde{y}$$

inverse variance of $\tilde{y}_1$

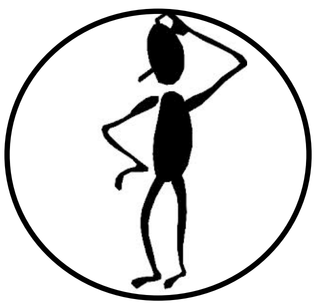

weights

$$\Lambda = \begin{pmatrix} \frac{1}{\sigma^2 + \sigma_1^2} & & 0 \\ & \ddots & \\ 0 & & \frac{1}{\sigma^2 + \sigma_n^2} \end{pmatrix}$$

- Generalized least-square estimator
  - Unbiased, covariance $V = \left(X^T \Lambda X\right)^{-1}$
  - Gauss-Markov/Aitken thm: smallest covariance amongst all linear unbiased estimators

Model (3): utilities/cost functions

- User *i* chooses precision of data revealed

$$\lambda_i = \frac{1}{\sigma^2 + \sigma_i^2} \in \left[0, 1/\sigma^2\right]$$

  – "contribution to result accuracy (public good)"

- Minimize cost

$$J_i(\lambda_i, \lambda_{-i}) = c_i(\lambda_i) + f(\lambda_i, \lambda_{-i})$$

Privacy cost
Increasing convex

Estimation cost

$$f(\lambda_i, \lambda_{-i}) = F(V(\lambda_i, \lambda_{-i}))$$

F, hence *f*, increasing convex

Examples: $F_1(\cdot) = trace(\cdot), \quad F_2(\cdot) = \|\cdot\|_F^2 = trace(\cdot^2)$

Nash equilibrium results for the linear model

- If $<d$ users contribute, infinite estimation cost
  ➔ trivial equilibria

- Main equilibrium result

Theorem:

There exists a unique non-trivial equilibrium

- Proof:
  - Potential game
  - Potential is convex

$$\Phi(\lambda_i, \lambda_{-i}) = \sum_i c_i(\lambda_i) + f(\lambda_i, \lambda_{-i})$$

Equilibrium efficiency

- Social cost: sum of cost of all users

$$C(\vec{\lambda}) = \sum_i c_i(\lambda_i) + nf(\vec{\lambda})$$

- Inefficiency of eq. measure by price of stability:

$$PoS = \frac{C(\vec{\lambda}^{NE})}{C(\vec{\lambda}^{SO})}$$

Social cost at the non-trivial Nash equilibrium

Minimal social cost

- Remarks:
  - Same as *PoA* if we remove the trivial equilibria
  - $PoS \geq 1$, "large *PoS:* inefficient", "small *PoS*: efficient"

Efficiency results for the linear model

- A first result:

Theorem:

The *PoS* increases at most linearly: $PoS \leq n.$

  - Obtained only from potential structure: by positivity of the estimation and privacy costs:

$$\frac{1}{n}C(\vec{\lambda}^{NE}) \leq \Phi(\vec{\lambda}^{NE}) \leq \Phi(\vec{\lambda}^{SO}) \leq C(\vec{\lambda}^{SO})$$

  - Works for any estimation cost, i.e., any scalarization F
  - But quite rough!

Efficiency results for the linear model (2)

- Monomial privacy costs: $c_i(\lambda_i) = c_i \cdot \lambda_i^k,\ c_i > 0, k \geq 1$

Theorem: (monomial costs)

If the estimation cost is $F_1(\cdot) = trace(\cdot)$, then $PoS \leq n^{1/(k+1)}$

If the estimation cost is $F_2(\cdot) = \|\cdot\|_F^2$, then $PoS \leq n^{2/(k+2)}$

  - Worst case (linear cost): $n^{1/2}$ for trace, $n^{2/3}$ for Frobenius

Theorem: (general costs)

With $F_1(\cdot) = trace(\cdot)$: if $nc_i'(\lambda) \leq c_i'(n^{1/2}\lambda)$, then $PoS \leq n^{1/2}$

With $F_2(\cdot) = \|\cdot\|_F^2$: if $nc_i'(\lambda) \leq c_i'(n^{1/3}\lambda)$, then $PoS \leq n^{2/3}$

EURECOM Sophia Antipolis

Population average case

- Case d=0: $y_i = \beta_0 + \varepsilon_i$ (β_0 is the population average)

Theorem (monotonicity):

When the number of agent increases, at equilibrium:
- each agent gives a smaller precision (λ_i decreases)
- the estimator's precision improves ($Var(\hat{\beta}_0)$ decreases)

– Note: If $c_i(\lambda) = \lambda^k$, then $Var(\hat{\beta}_0) \sim n^{-1+2/(k+1)}$ (slower than iid)

Theorem (improved learning accuracy):

For a well chosen η, the analyst can strictly improve the estimator's variance by restricting the users precision choice to {0}U[η, 1/σ²]

EURECOM Sophia Antipolis

Roadmap

- Game theory and statistical learning for
  - Security
  - Privacy
  - Networked systems
- Perspectives

Overview of my contributions in networked systems

- Causal analysis of network performance
  - A new bootstrap inference algorithms and application to TCP, DNS

Key papers: [AlgoTel ’14, ITC ’15, TIST ’16, Comnet ’16]
1 student

- Robust incentives for decongestion
  - Lottery-based scheme robust to utility estimation errors
  - Study of day-ahead pricing schemes in smart grids

Key papers: [Netgcoop ’12, Allerton ’12, ToN ’14, ACC ’16]
Key collaborations: UC Santa Cruz, UC Berkeley, Inria
1 student

- Approximation algorithms for cloud resource allocation

Key papers: [Allerton ’15, TPDS maj. rev., ArXiv ’16]
1 student

Roadmap

- Game theory and statistical learning for
  - Security
  - Privacy
  - Networked systems

- Perspectives: *humans vs machine learning*

Learning from strategic data

- How to learn from strategic data? (not iid)
  ➔ using solutions of game-theoretic models

- Learning from personal data of privacy-conscious users
  - Find algorithms that optimize learning accuracy at equilibrium
    - Incomplete information
    - Non-linear regression, recommendation
  - A statistical learning theory for strategic data
    - Risk bounds, sample complexity
- Learning from strategic data in security
  - Incomplete information
  - Dynamic models

Human-friendly learning algorithms

- Learning algorithms have a major impact on humans life…
  - Online services, hiring, justice, etc.
- …but we often can't understand how they work

➢ Bringing transparency
  - Collaborative transparency tool
  - Definition of explanation

➢ Bringing fairness
  - Designing algorithms under constraints of acceptability

Professional activities & visibility

- Teaching
  - 3 courses / year: game theory, network economics, statistical data analysis
  - Responsible networking track
- Students supervision
  - 5 PhD students (2 graduated)
  - 1 postdoc (graduated)
  - 5 interns (graduated)
- Funding (total ~800k)
  - Projects: IMT F&R, Labex UCN@Sophia, ANR Tremplin-ERC
  - Industry: Symantec faculty gift, Data Transparency lab, Cifre SAP, Cifre Nokia
- Sabbatical visits
  - UC Berkeley, summer 2012
  - MPI-SWS, summer 2014 and 2016-17
- Awards
  - Humbold Research Award 2016
- Editorial activities
  - Associate editor ACM TOIT
  - Lead guest editor of 2 special issues
- Steering committees
  - Chair NetEcon SC
  - Member SC Labex
- Conference organization
  - PC chair NetEcon ’12-’15
  - Registration chair SIGMETRICS ’13, ’16
  - Chair sophia-networking seminar
- PhD committees and grant panels
  - PhD reviewer and committee member
  - Grant panel expert FRS Belgium, ARC Singapore
- Keynotes and invited talks/lectures
  - Keynote AEP ’16
  - Invited lectures UCLA IPAM summer school, RESCOM summer school, SIGMETRICS tutorial
  - Invited talks In’Tech, MIT, Harvard, Northeastern, Berkeley, IHP, AlgoGT, UCLA, Caltech, etc.

Main achievements since PhD

Oana, July 2016

Luca, November 2016

THANK YOU! QUESTIONS?