

Privacy of profile-based ad targeting

Alexander Smal
and
Ilya Mironov

User-profile targeting

- Goal: increase impact of your ads by targeting a group potentially interested in your product.
- Examples:
 - Social Network
Profile = user's personal information + friends
 - Search Engine
Profile = search queries + webpages visited by user

Facebook ad targeting

Location

Country: [?]

Everywhere
 By City [?]
 By State/Province [?]
 By Zip Code [?]

Include cities within miles.

Demographics

Age: [?] -

Require exact age match [?]

Sex: [?] All Men Women

Advanced Demographics

Interested In: [?] All Men Women

Relationship: [?] All Single Engaged In a relationship Married

Languages: [?]

Education & Work

Education: [?] All College Grad

In College
 In High School

Workplaces: [?]

Interests

- Business/Technology
- Family Status
- Interests
- Mobile
- Movie/Film
- Music
- Retail/Shopping
- Sports
- University
- Business/Technology
- Family Status
- Interests
- Mobile
- Movie/Film
- Music
- Retail/Shopping
- Activities 2
- Birthday
- Business/Technology
- Family Status
- Interests
- Mobile
- Movie/Film

- Home & Garden
- News
- Pets (All)
- Pets (Cats)
- Pets (Dogs)
- Politics (US Active)
- Politics (US Conservative)
- Politics (US Liberal)
- Pop Culture
- Mobile (All)
- Android
- iPhone
- RIM/Blackberry
- Windows Phone
- Baby Boomers
- Engaged (<6 months)
- Newlywed (<1 year)
- Parents (All)
- Parents (child: 0-3yrs)
- Parents (child: 4-12yrs)
- Parents (child: 13-15yrs)
- Parents (child: 16-19yrs)

Characters



Advertising company

My system
is private!

Privacy researcher





Advertising company

My system
is private!

How can I
target
privately?

Privacy researcher

Unless your
targeting is not
private, it is not!



How to protect information?

- Basic idea: add some noise
 - Explicitly
 - Implicit in the data
 - noiseless privacy [BBGLT11]
 - natural privacy [BD11]
- Two types of explicit noise
 - Output perturbation
 - Dynamically add noise to answers
 - Input perturbation
 - Modify the database



Advertising company

I like input
perturbation
better...

Privacy researcher



Input perturbation

- Pro:
 - Pan-private (not storing initial data)
 - Do it once
 - Simpler architecture



Advertising company

I like input
perturbation
better...



Privacy researcher

Signal is
sparse and
non-random

Adding noise

- Two main difficulties in adding noise:

- Sparse profiles

1	0	0	0	1	0	0	0	1	0
0	0	0	0	0	0	0	0	1	1
0	0	0	1	0	0	0	0	0	1
0	0	0	1	0	0	1	0	0	0
1	1	1	1	1	0	0	1	0	1
0	0	0	0	1	0	0	1	0	0
0	0	1	0	0	0	0	1	1	0
0	0	0	0	0	0	0	1	0	1
1	0	1	0	0	0	0	0	0	0
0	1	0	1	0	1	0	0	0	0

differential privacy

- Dependent bits

1	0	0	1	1	1	0	1	0	0
0	1	0	0	0	0	0	0	1	1
1	0	0	0	0	0	0	0	1	0
1	0	0	1	1	1	1	0	0	0
0	1	1	0	0	0	0	0	1	1
1	0	0	0	1	0	0	1	0	0
0	1	1	1	1	1	0	1	1	1
0	1	0	0	0	0	0	1	0	0
1	0	1	1	1	1	0	0	0	0
0	1	0	0	0	0	0	1	0	1

“Smart noise”

deniability



Advertising company

I like input perturbation better...

Let's shoot for deniability, and add "smart noise"!

Privacy researcher



Signal is sparse and non-random

“Smart noise”

- Consider two extreme cases
 - All bits are independent
independent noise
 - All bits are correlated with correlation coefficient 1
correlated noise
- “Smart noise” hypothesis:
“If we know the exact model we can add right noise”

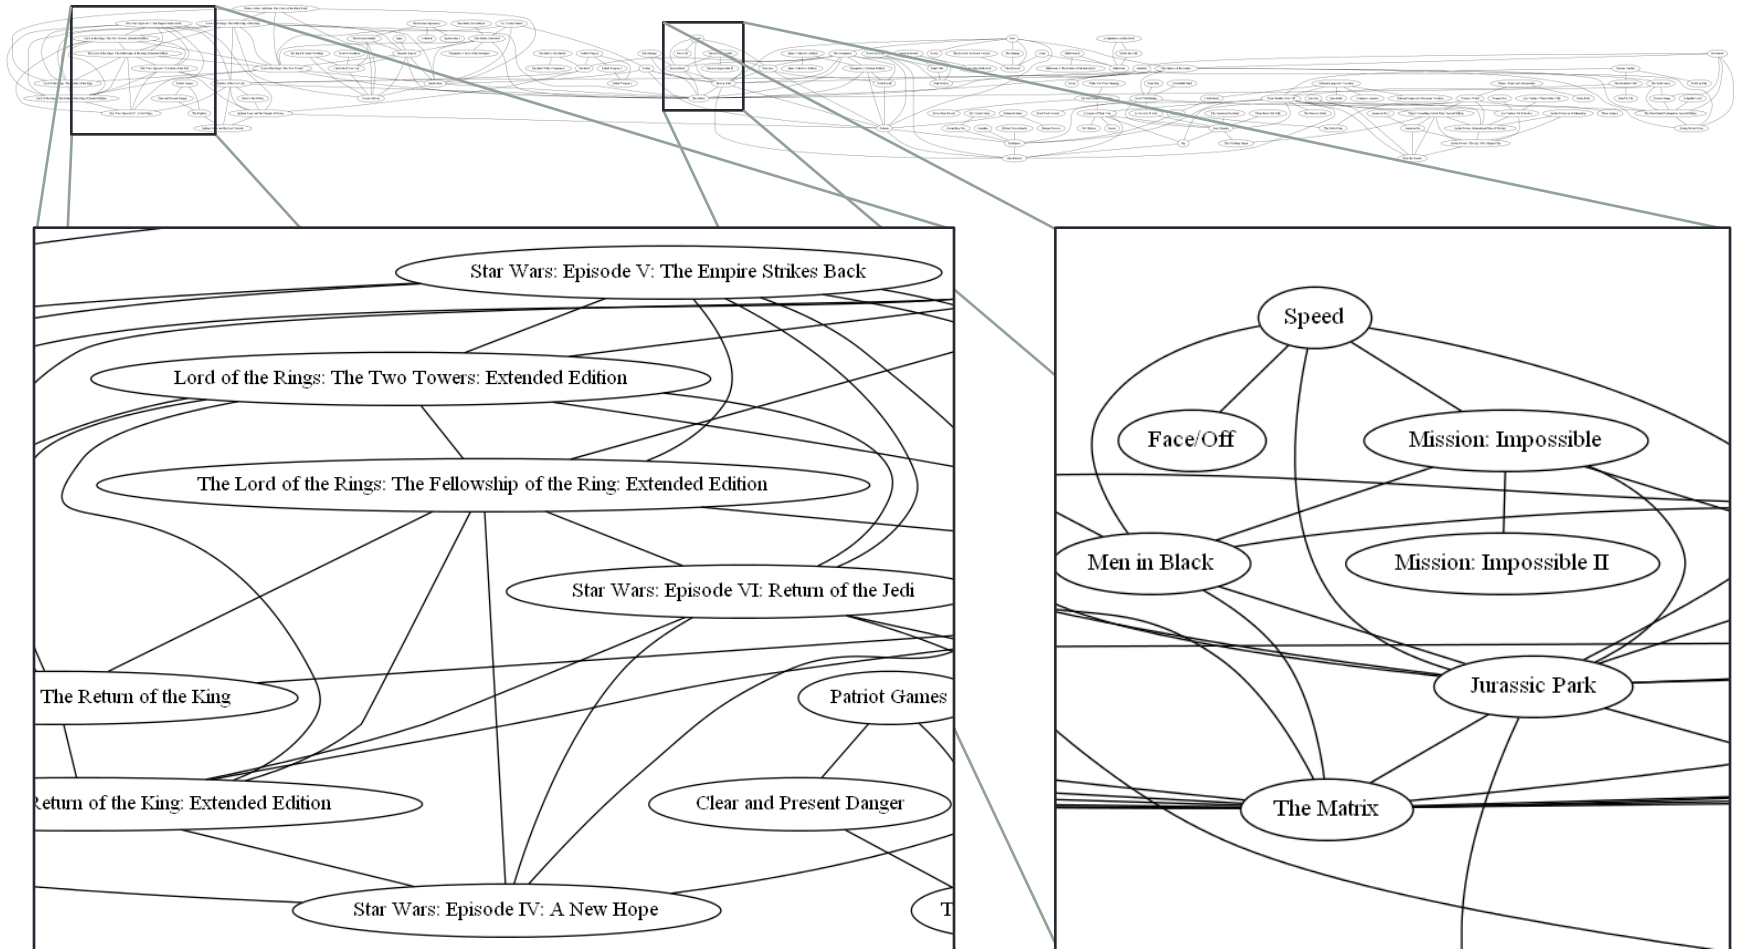


Aha!

Dependent bits in real data

- Netflix prize competition data
 - ~480k users, ~18k movies, ~100m ratings
- Estimate movie-to-movie correlation
 - Fact that a user rated a movie
- Visualize graph of correlations
 - Edge – correlation with correlation coefficient > 0.5

Netflix movie correlations





Advertising company


Let's shoot for deniability, and add "smart noise"!



Privacy researcher

Let's construct models where "smart noise" fails


How can “smart noise” fail?



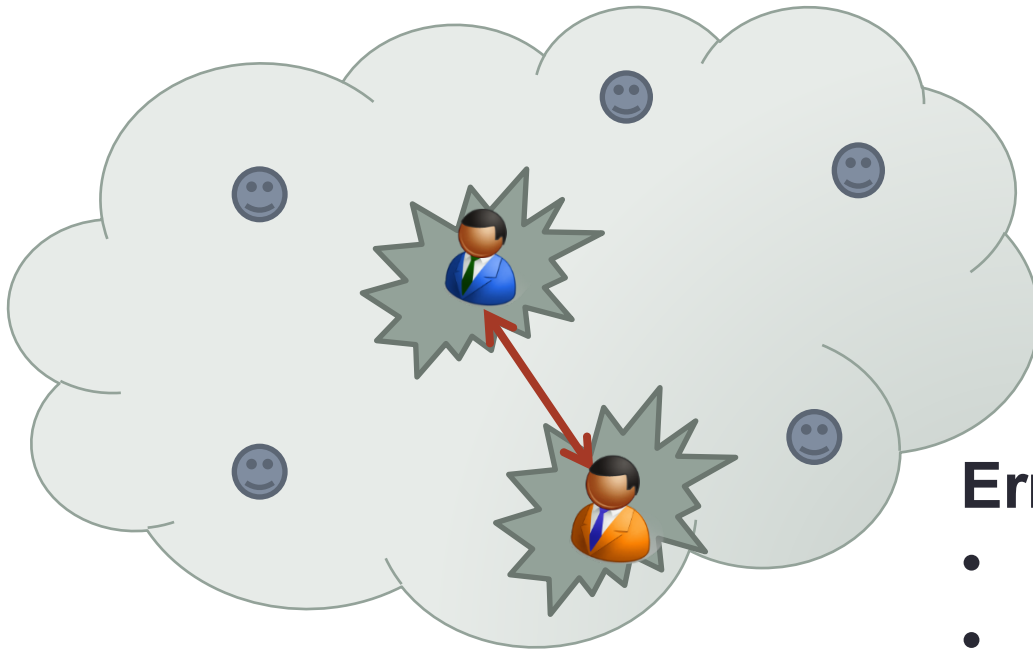
large = relative distance $\Omega(1)$

Models of user profiles

- M hidden independent bits
- N public bits




- Public bits are some functions of hidden bits
- Are users well separated?



Error-correcting codes

- Constant relative distance
- Unique decoding
- Explicit, efficient



Advertising company

But this model is unrealistic!



Privacy researcher

See — unless the noise is $>25\%$, no privacy

Let me see what I can do with monotone functions...

Monotone functions

- Monotone function: for all i and for all values of $x_j, j \neq i$
$$f(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n) \geq f(x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n)$$
- Monotonicity is a natural property
$$[wants\ Kindle] \leftrightarrow [likes\ reading] + [likes\ gadgets] \times [uses\ Amazon]$$
- Monotone functions are bad for constructing error-correcting codes

Approximate error-correcting codes

- α -approximate error-correcting code with distance δ :

function $f: \{0,1\}^n \rightarrow \{0,1\}^m$

$\forall x, x'$, such that $\|x - x'\|_1 \geq \alpha n$:

$$\|f(x) - f(x')\|_1 \geq \delta m.$$

- If less than δ fraction of $f(x)$ is corrupted then we can reconstruct x within α fraction of bits.
- We need $o(1)$ -approximate error-correcting code with constant distance.

blatant non-privacy


Noise sensitivity

- *Noise sensitivity* of function f :


$$NS_{\epsilon}(f) = \mathbf{Pr}_x[f(x) \neq f(y)],$$

where x is chosen uniformly at random, y is formed by flipping each bit of x with probability ϵ .

- If $NS_{\epsilon}(f)$ is big \Rightarrow



- If $NS_{\epsilon}(f)$ is small \Rightarrow



Monotone functions

- There exist highly sensitive monotone functions [MO'03].
- **Theorem:** there exists monotone $o(1)$ -approximate error-correcting code with constant distance on average.
- **Idea of proof:** Let f_1, f_2, \dots, f_m be random independent monotone boolean functions, such that $\text{NS}_\epsilon(f_i) \geq c$ and f_i depends only on $o(n)$ bits of x .
- Let $F(x) = \langle f_1(x), \dots, f_m(x) \rangle$.
- With high probability for random x there is no x' such that $\|x - x'\|_1 \geq \epsilon n$ and $\|F(x) - F(x')\|_1 \leq \frac{cn}{2}$.
- For Talagrand $o(1/\sqrt{n})$ -approximate error-correcting code with constant distance on average.



Advertising company

Hmmm. Does smart noise ever work?

Privacy researcher



If the model is monotone, blatant non-privacy is still possible

Linear threshold model

- Function $f: \{-1,1\}^n \rightarrow \{0,1\}$ is a *linear threshold function*, if there exist real numbers α_i 's such that

$$f(x) = \text{sgn}(\alpha_0 + \alpha_1 x_1 + \cdots + \alpha_n x_n).$$

- **Theorem** [Peres'04]: Let f be a linear threshold function, then $\text{NS}_\delta(f) \leq 2\sqrt{\delta}$.



No $o(1)$ -approximate error-correcting code with $O(1)$ distance

Conclusion

- Two separate issues with input perturbation:
 - Sparseness
 - Dependencies
- “Smart noise” **fallacy** :

Even for a publicly known, relatively simple model, constant corruption of profiles may lead to blatant non-privacy.
- Connection between noise sensitivity of boolean functions and privacy
- Open questions:
 - Linear threshold privacy-preserving mechanism?
 - Existence of interactive privacy-preserving solutions?

Arbitrary
Monotone
~~Linear threshold~~

Thank for your attention!

Special thanks for Cynthia Dwork, Moises Goldszmidt, Parikshit Gopalan, Frank McSherry, Moni Naor, Kunal Talwar, and Sergey Yekhanin.

- Events $[f_i(x) \neq f_i(y)]$ and $[f_j(x) \neq f_j(y)]$ are independent for random $x, y = N_\epsilon(x)$, i and j .
- Chernoff bounds: $\Pr_{x,y=N_\epsilon(x)} \left[\sum_{i=1}^m |f_i(x) - f_i(y)| < \frac{mc}{2} \right] < e^{-\frac{mc}{8}}$.