Privacy and Recommendation Exploring the Tradeoffs

Davide Frey
joint work with Antoine Rault
as well as Anne-Marie Kermarrec, Rachid Guerraoui, François Taiani, Jingjing Wang

ASAP Team
Inria Rennes

Pervasive Recommendation

- Popular tool to
  - address information overload
  - target advertising
  - Identify interesting content

A Brief Recommender Taxonomy

- Content-based filtering
  - Exploit text-analysis, image processing
  - Limited privacy issues
- Collaborative filtering
  - Trace similarities between user interests
  - Memory based
    - User Based: Group users based on interesting items
    - Item Based: group items based on interested users
  - Model based
    - Bayesian, latent semantic models, SVD…
  - All need to combine the preferences of many users

Inherent Privacy Tradeoff

- Recommending requires information about users
- Users wish to protect their information from
  - Big Brothers (the recommender)
  - External Actors
  - Users

Architectures For Recommendation

centralized

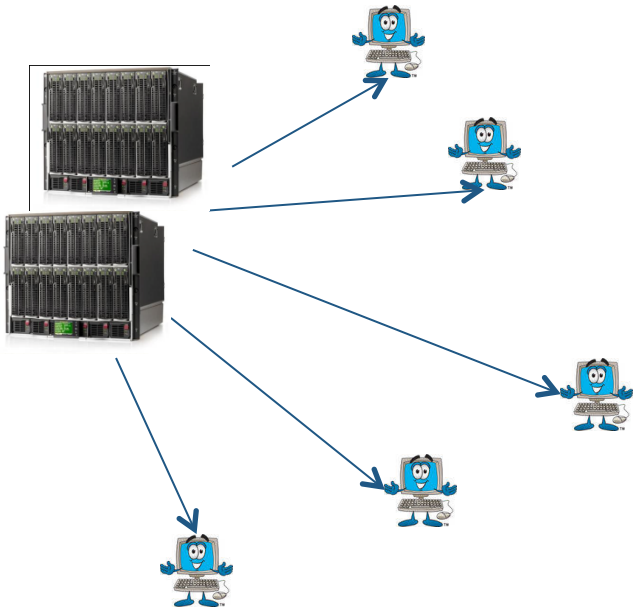

decentralized

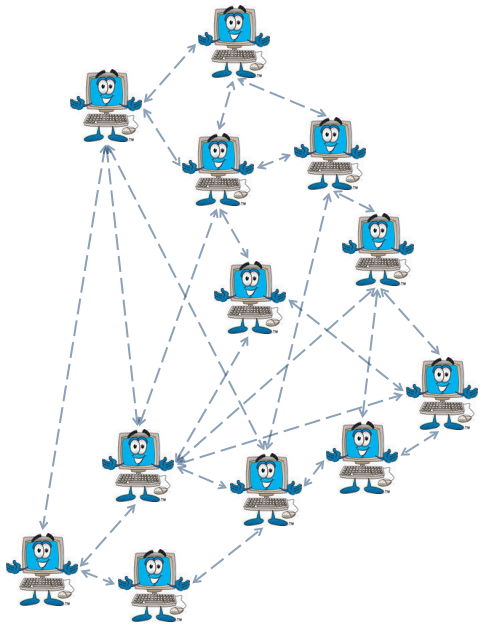

removes big brother, but …

This Talk

- A technique to improve privacy in decentralized collaborative filtering [to appear DSN 15]

- Analysis of an attack in centralized or decentralized collaborative filtering [to appear EuroSec 15]

Peer-to-Peer Collaborative Filtering

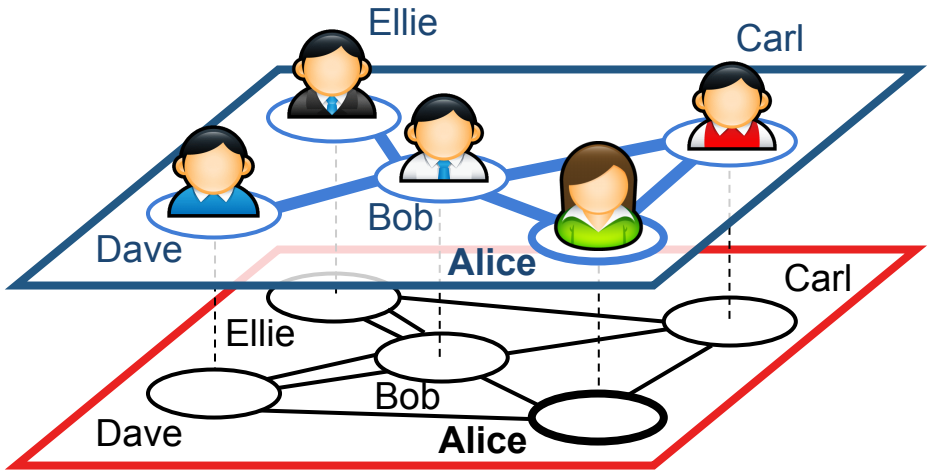

clustering layer
gossip-based
topology clustering

RPS layer providing
random sampling

Random link *Social link* *node*

- Remove Big Brother

Peer-to-Peer Collaborative Filtering

Build Knn graph through epidemic protocols

- RPS builds a random topology
  - Continuously provides new information
- Clustering identifies nearest neighbors
  - Similarity metric: e.g. cosine
- Recommendation based on neighbors' ratings

Key Privacy Leak: Similarity Computation

Computing similarities requires

knowledge of each other's profiles

Replace big brother by many little brothers

Attacker Model

- Goal: Discover a target user's interests
- Restricted active adversary
  - Passive information gathering
  - Some active steps:
    - Tap unencrypted communications
    - Try to bias multi-party computations
    - Unlimited similarity computations
- No collusion, no Sybil attack

Hide and Share

Main Insight: Landmark-based similarity

- Indirectly compare user profiles by exploiting their similarities with randomly generated profiles (landmarks)

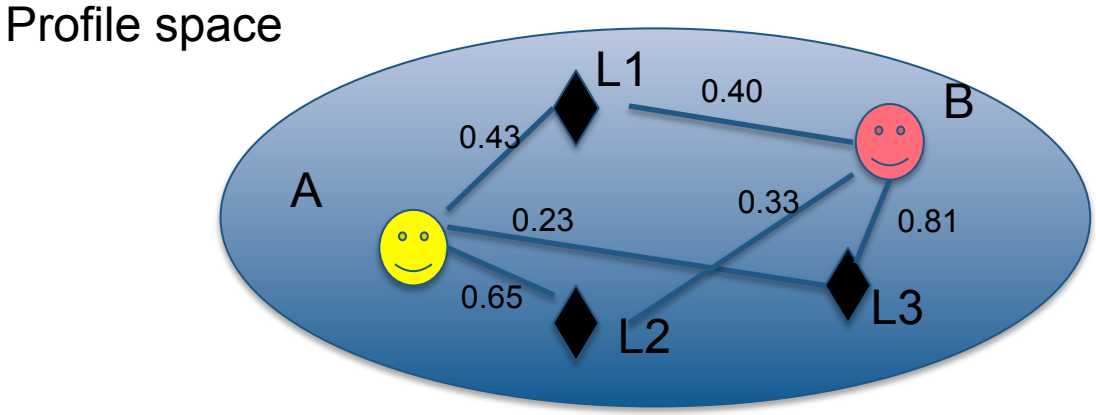

Coordinate system analogy

Hide and Share Requirements

- Computation Confidentiality
- Landmark-profile independence
- Fair Landmark generation
- Time-independent information release

Computation confidentiality

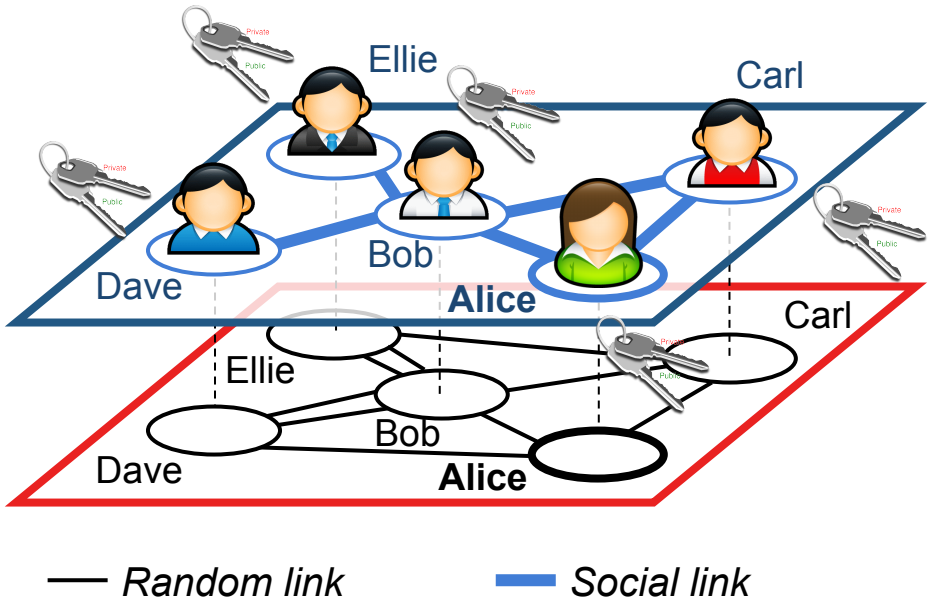

Attach Public Key to gossip messages

Generate secret key to exchange data for similarity computation

Landmark-profile Independence

- Need to generate random landmarks
- Need a way to describe the profile space!

- Represent profiles as binary vectors
  - Profile is a set of items
  - Compact profile in the form of bloom filters
    - Only count “liked” items (rating>threshold)

Fair Landmark Generation

- Need common seed
  - Bit-commitment – blum's protocol

P1 and P2 flip a coin
P1 sends f(conc(result, nonce))
P2 reveals result to P1
P1 reveals result to P1
If same result -> bit = 1

Time-independent information release

- Generate landmarks using common seed
- Store seed for future use
  - Will recompute the same landmarks the next time it meets peer.
- Overhead -> one seed per peer

Protocol Summary

A and B's first meeting

Set up secure communication channel

Protocol Summary

A and B's first meeting

Set up secure communication channel

Agree on common seed

Protocol Summary

A and B's first meeting

Set up secure communication channel

Agree on common seed

Derive L random profiles (landmarks) using the seed

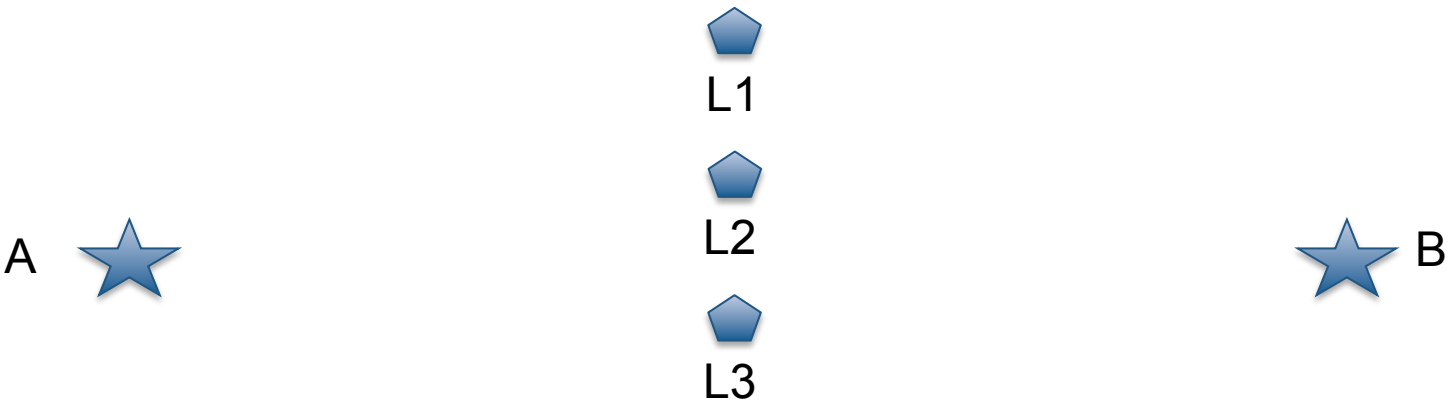

Protocol Summary

A and B's first meeting

Set up secure communication channel

Agree on common seed

Derive L random profiles (landmarks) using the seed

Compute similarity with the landmarks

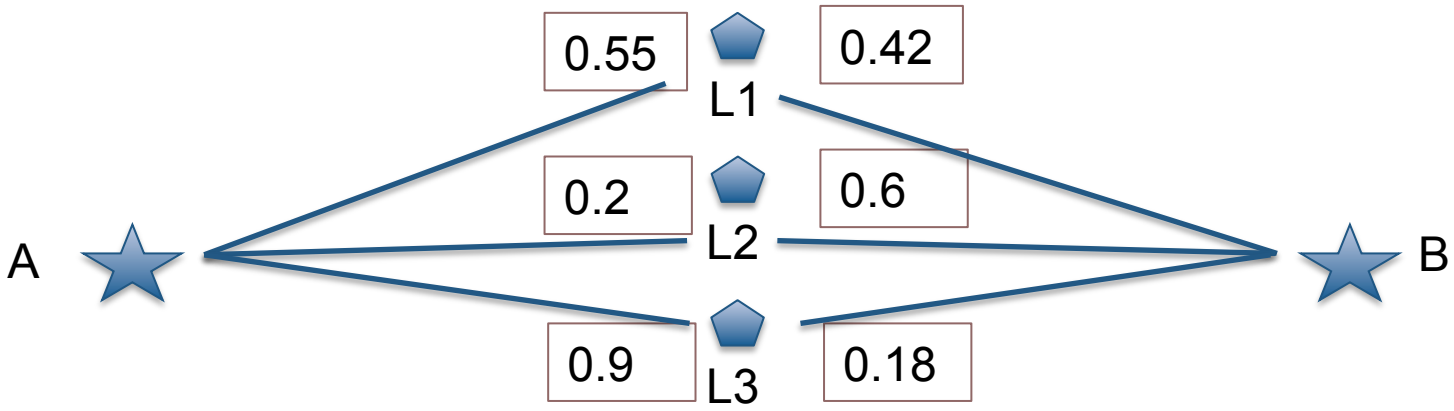

Protocol Summary

A and B's first meeting

Set up secure communication channel

Agree on common seed

Derive L random profiles (landmarks) using the seed

Compute similarity with the landmarks

Cosine similarity of coordinate vectors

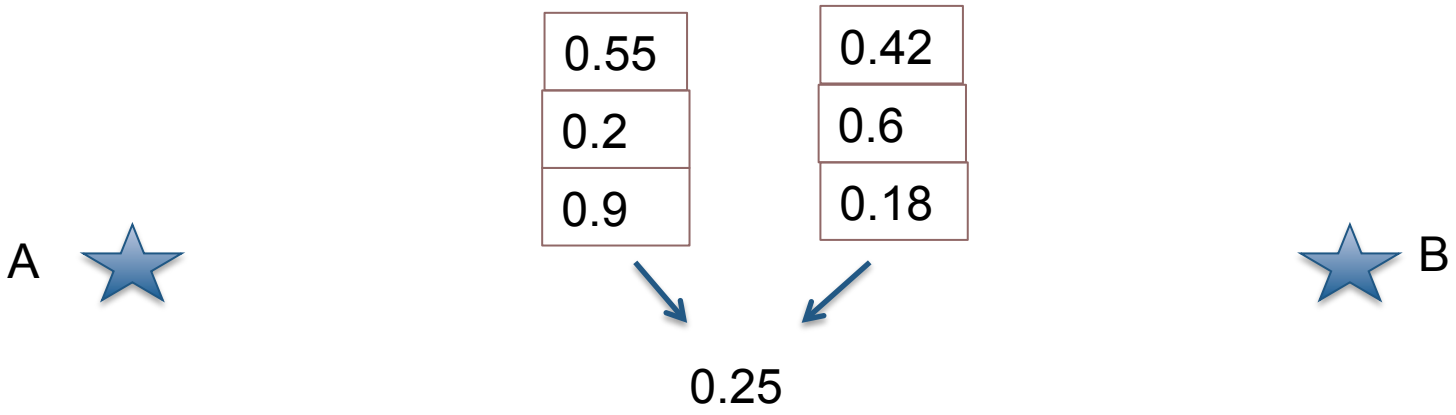

Protocol Summary

A and B meet again

Derive L random profiles (landmarks) using the seed

Compute similarity with the landmarks

Cosine similarity of coordinate vectors

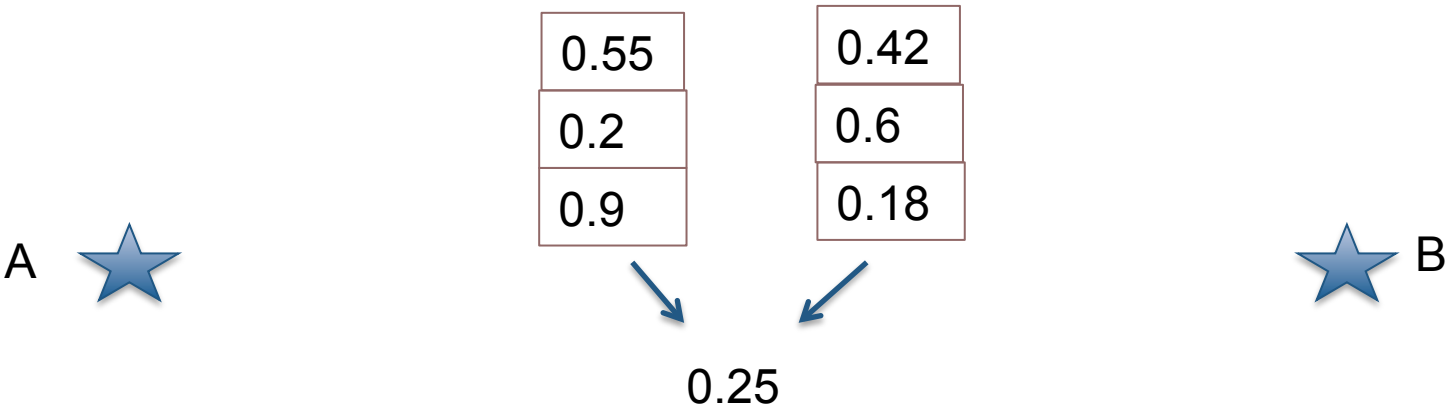

Evaluation

- **MovieLens**: movies recommendation datasets
- **Jester**: jokes recommendation dataset

	nb users	nb items	rating range
ML-100k[1]	943	1,682	1:5 (integers)
ML-1M[1]	6,040	3,900	1:5 (integers)
Jester[2]	24,983	100	-10:10 (continuous)

[1] MovieLens: `http://grouplens.org/datasets/movielens/`
[2] Jester: `http://eigentaste.berkeley.edu/dataset/`

Evaluation

1- Split dataset randomly

2- Use training set to fill profiles

3- Generate recommendations and check against training set

Metrics

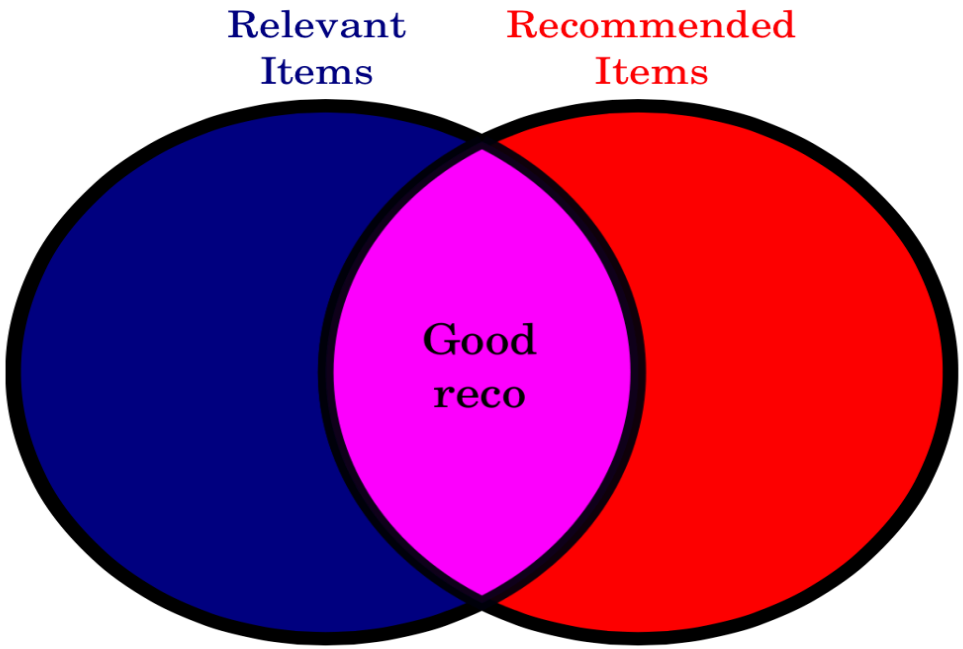

Recall = Good / Relevant

Precision = Good / Recommended

Recommendation Quality

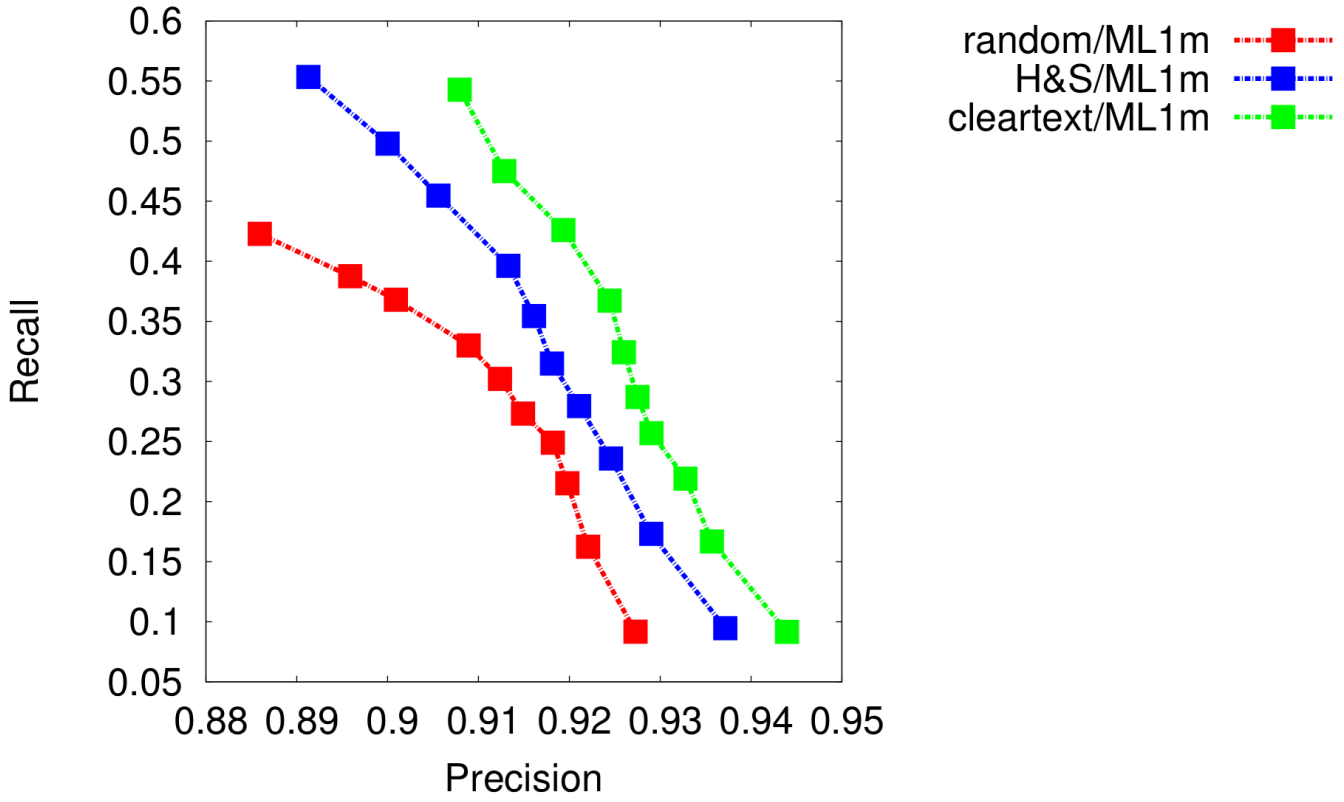

Neighborhood Quality

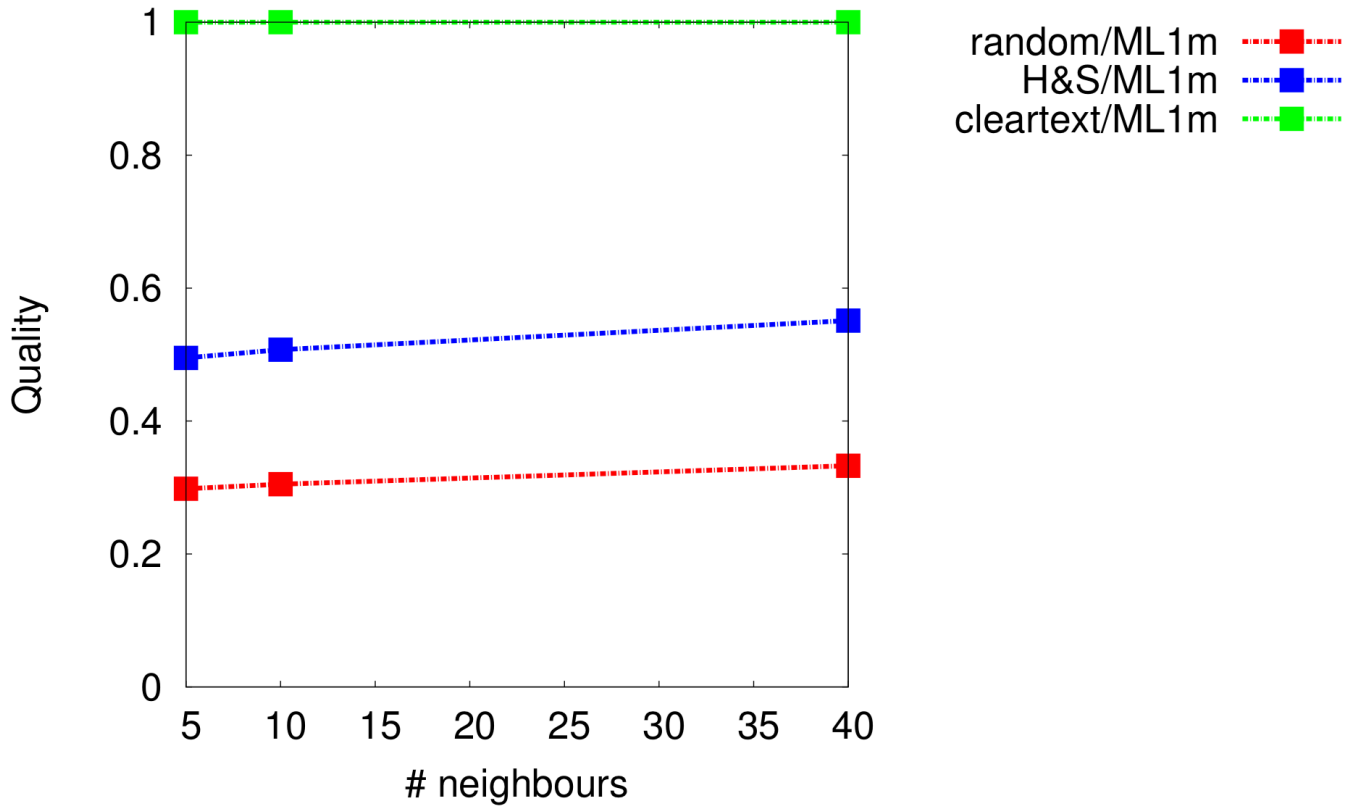

Privacy: Profile Reconstruction

Profile Reconstruction Attack

- Infer target profile from landmark similarities
- Guess
- items that form the target compact profile
  - Assumption: The attacker knows all the item signatures
- Attack:
  - Consider closest landmark profile as target profile
  - Guess all items that march target profile

Privacy

- How to measure privacy?
  - Simulation: set score
    - G = guessed profile
    - P = peer profile

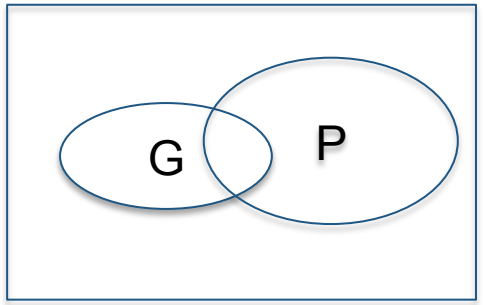

$$\text{SETSCORE}(G, P) = \frac{|G \Delta P| - |G \cap P|}{|G \cup P|}$$

    - Range [-1, 1]
      - -1 = exact and complete guess
      - 1 = completely wrong guess

Setup

- Baseline: Randomized profiles
  - Apply random perturbation to compact profiles
  - Varying percentage of randomized bits (5% to 100%)

- Hide and Share configuration
  - Vary landmarks between 2 to 100

Bandwidth Consumption

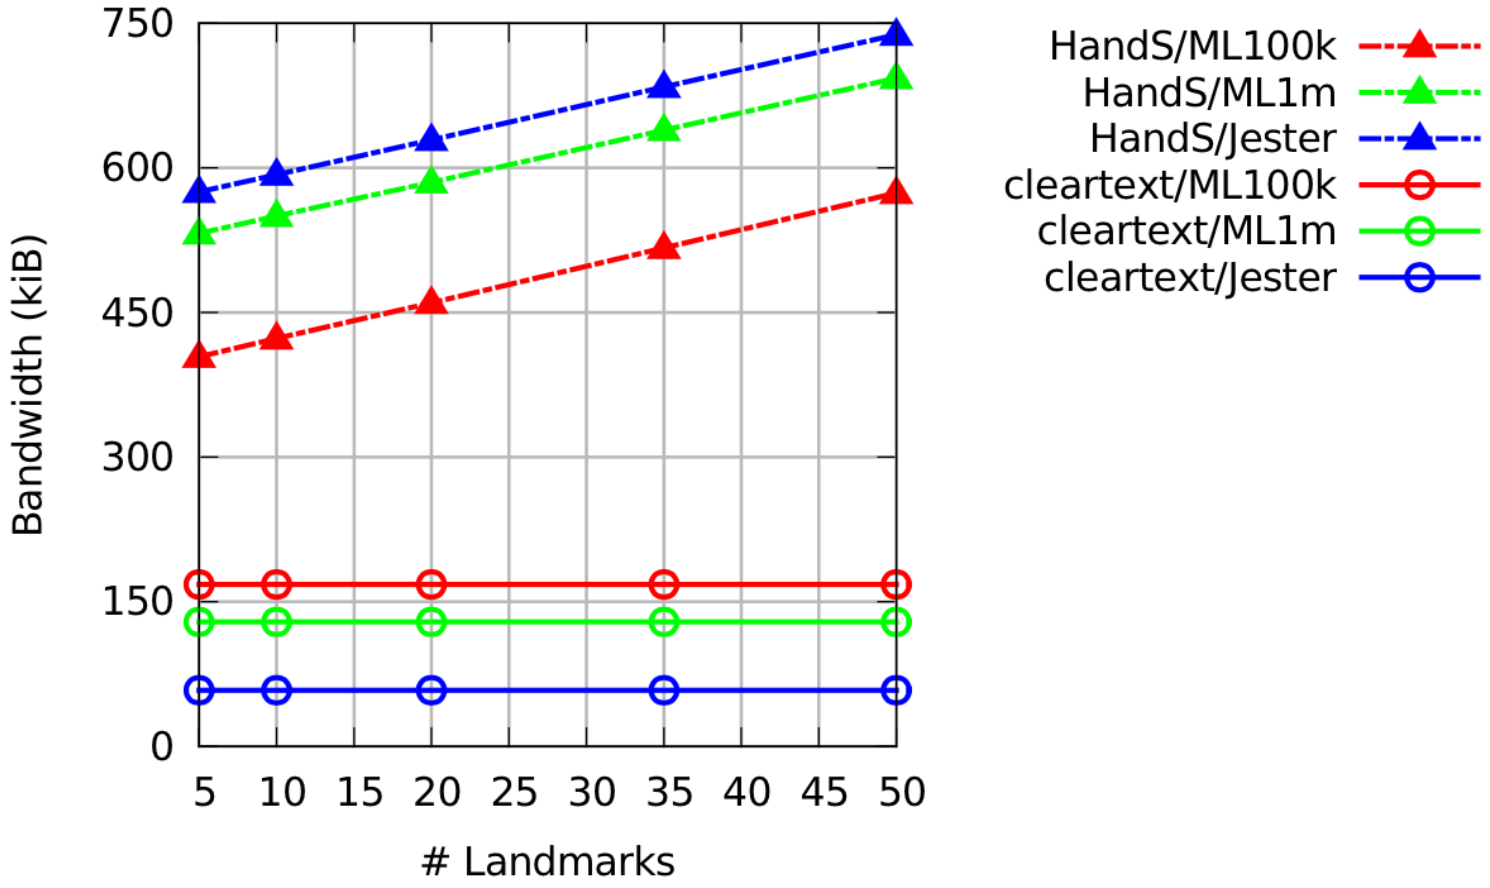

Results

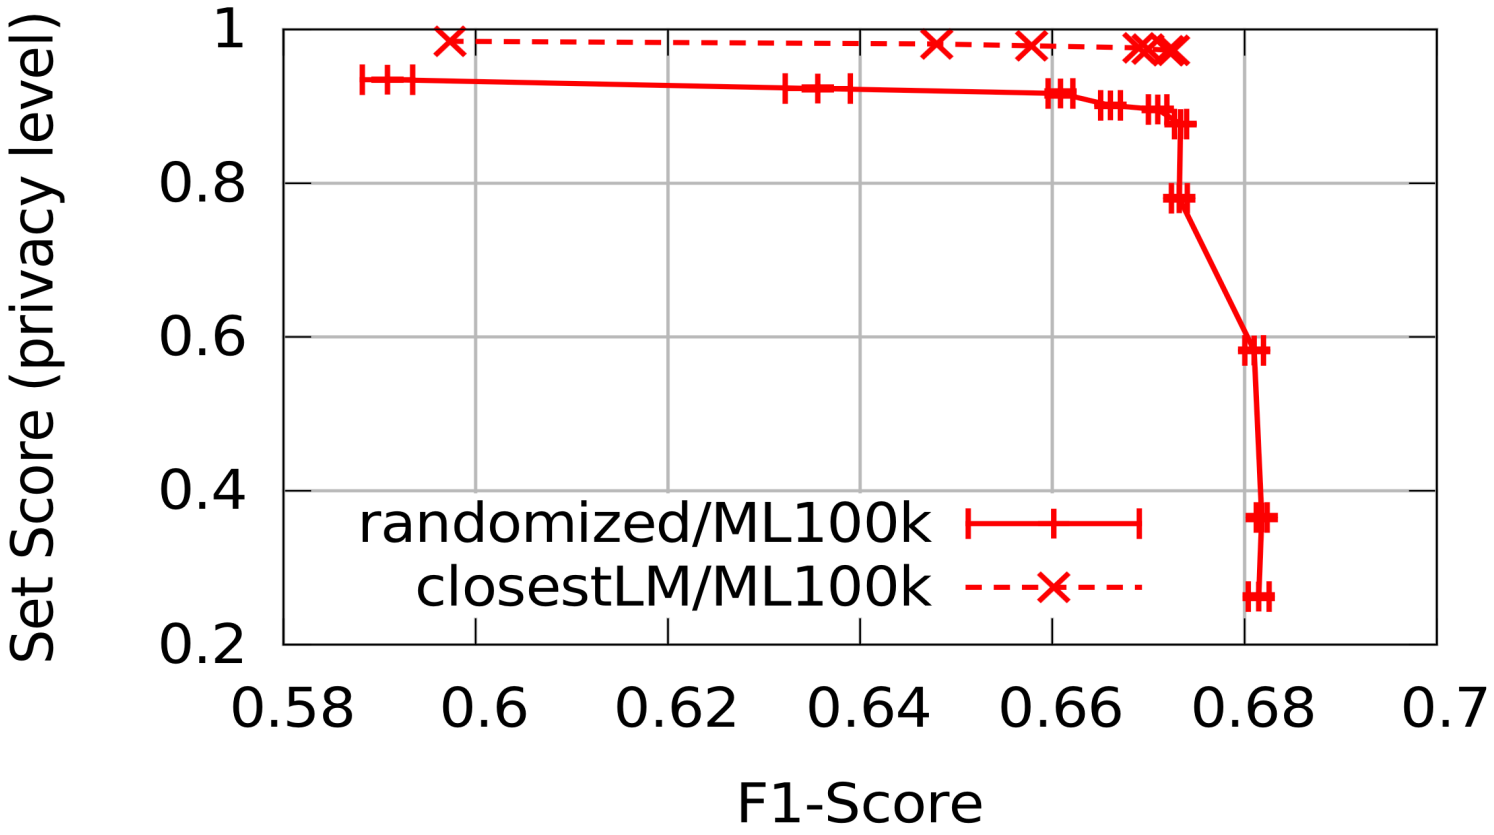
Set Score (privacy level)
1
0.8
0.6
0.4
0.2
0.58
0.6
0.62
0.64
0.66
0.68
0.7
F1-Score
randomized/ML100k
closestLM/ML100k

Storage Space

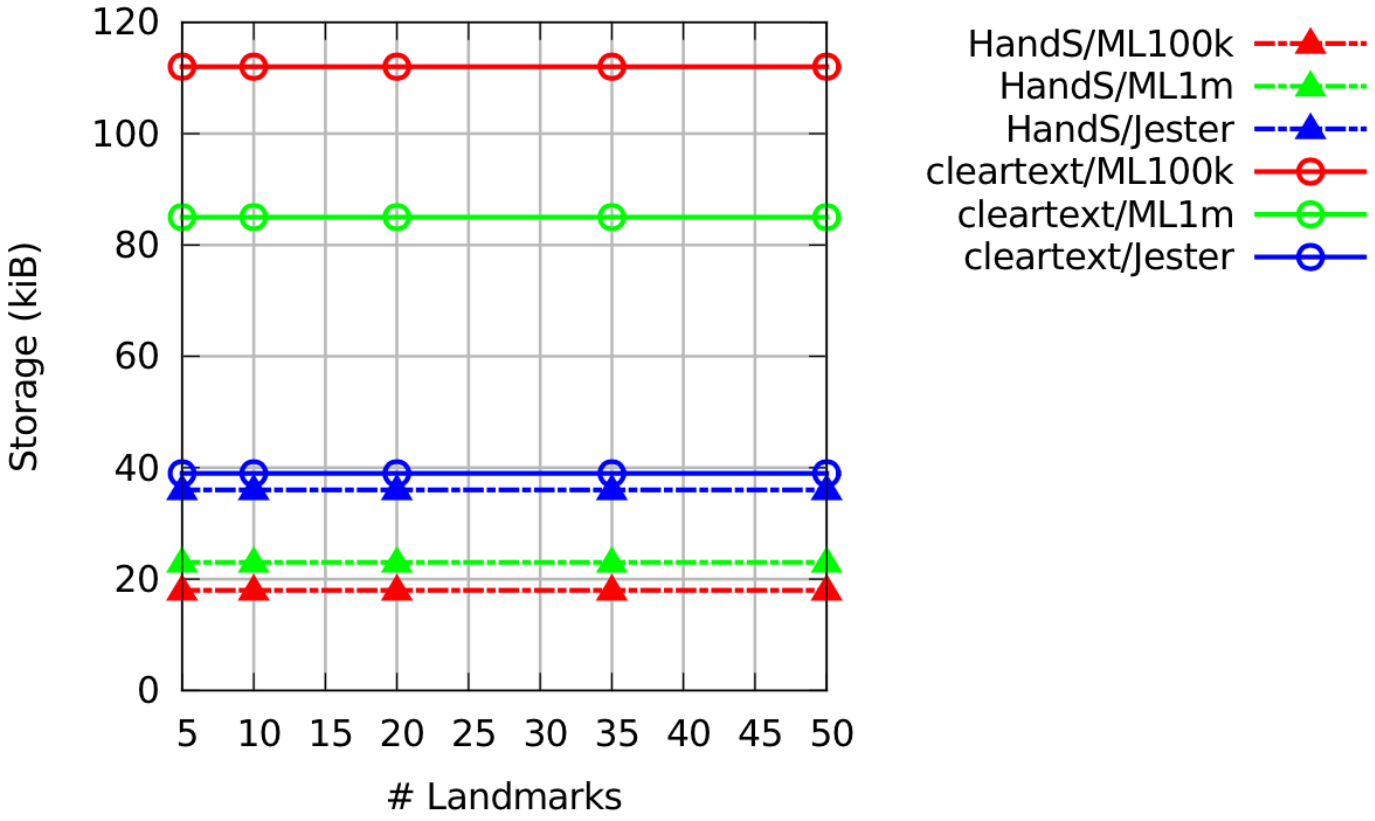

Leakage Analysis

- Leaked Information
  - Let M be landmark matrix
  - Let D(M) be the number of non-zero rows in M
  - - L information loss

$$\mathcal{L} \leq n - \frac{1}{2^n} \sum_{D(M)} \binom{n}{D(M)} (\rho_1)^{D(M)} (1-\rho_1)^{n-D(M)} T(D(M)).$$

$$\rho_1 = p(\vec{r} \neq \vec{0}) = 1 - (1-\eta)^m$$

η probability of element in M being 1

ρ_1 probability of having a non-zero row in M

Conclusion

- Recommendation
  - Useful
  - But at odds with privacy
- Data discoverable also through indirect means
- Research directions
  - Distinguishing information: to keep private, not to keep private (not necessarily personal data)
  - Let users choose what they share with whom

merci

LIEU
LOCALISATION

www.inria.fr

SetScore

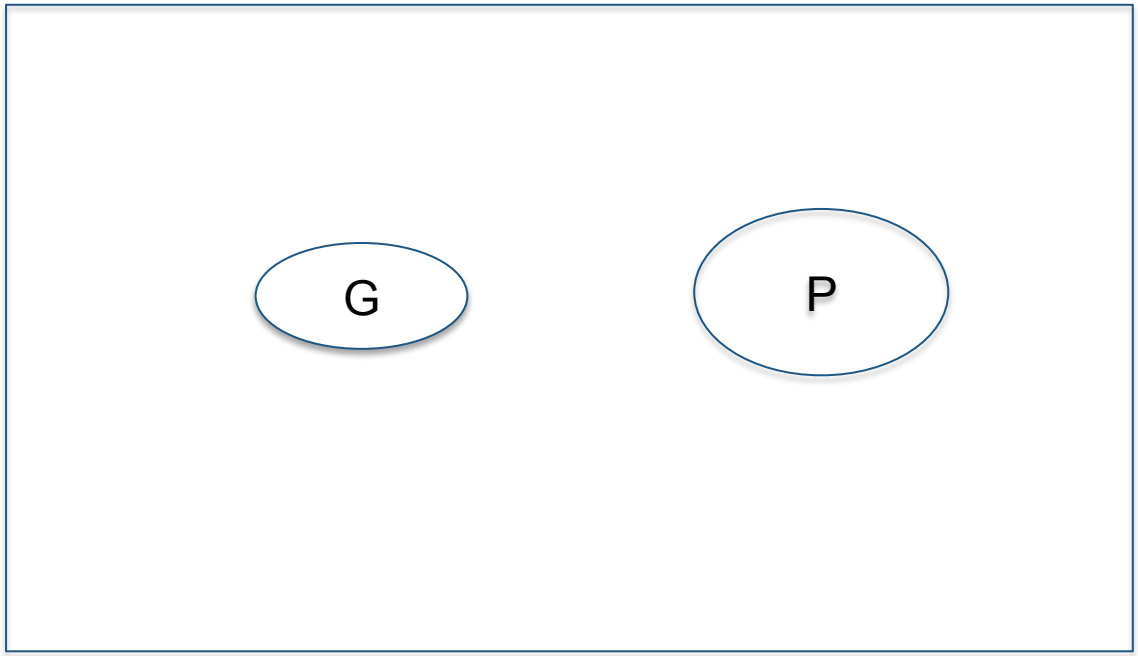

$$\textsc{setScore}(G, P) = \frac{|G \Delta P| - |G \cap P|}{|G \cup P|}$$