Privacy

1

Exposing users: techniques

- You Might Also Like: Privacy Risks of Collaborative Filtering, Calandrino, J.A, Kilzer, A., Narayanan, A., Felten, E.W., and Shmatikov, V., *IEEE Sym. on Security and Privacy (SP)*, 2011, pp. 231 - 246.
- Various item-to-item collaborative filtering methods
- Practical algorithms

2

Set up

- attacker and target user
- attacker to infer unobservable transaction by target user
  - e.g. item purchased or rating given item
- attacker uses "auxilary information" about some transactions of target user
- attacker only observes
  - does not enter ratings/ make transactions
  - no fake users

3

Sources of auxiliary information

- provided by target system
  - e.g. public ratings by user
- "third-party sites"
  - partner with target site
  - e.g. embed playlist on blog
- other sites
  - user places related content
  - e.g. Facebook user profile

4

"Generic Inference Attacks"

- Auxiliary information
  - target system provides lists of related items
  - target system provides item-to-item covariance matrix used by collaborative filtering
- Auxiliary information & Active attack
  - target system uses k-nearest neighbor recommender

5

Using related items

- system gives list of related items for each item based on user selection
- auxiliary items: attacker knows certain items associated with target user
- attacker
  - monitors related-items lists of auxiliary items
  - scores changes in lists:
    - new items or items moving up on lists
  - if score for an item above threshold, infer item added to target user's record

6

Using covariance matrix

- item-item covariance matrix M available
  - Hunch questions to users
- user record containing items interacted with
- auxiliary information: attacker knows subset A of items associated with target user u
  - new item in record for u => covariances beween new item and (some) items in A goes up
  - subset unique to target user?

7

Using covariance matrix, cont.

- attacker
  - monitors changes in covariance submatrix
    - columns for A
    - rows AU {candidate new items}
  - scores changes in submatrix
  - if score for an item above threshold, infer item added to target user's record
- Lots of details concerning update delays in paper

8

Active attack: for kNN recommender systems

- Example target system
  - similarity measure on users
  - find k most similar users to user u
  - rank items purchased by one or more of k most similar users
    - ranking by number times purchased
  - recommend items to u in rank order

9

kNN recommender systems, cont.

- auxiliary information: subset of m items target user U has purchased
  - claim m of about O(log (# users)) suffices
- attacker
  - creates k sybil users
  - puts m auxiliary items on sibils' histories
    - "high probability" kNN of each sybil is other k-1 sybils and U
  - infer that any items recommended by system to any of sybils and not one of m a aux items is item U has purchased

10

Evaluation

- use
  - yield: number inferences per user per observation period
  - accuracy: percentage of inference that are correct
- need "ground truth"
- Several studies in paper

11

used on Amazon

- no ground truth
- API provides "Customers who bought x also bought y" and sales rank of items
- chose customers: top reviewers but not among top 1000 reviewers
- auxiliary info: entire set items previously reviewed by chosen customers
  - avg ~120 per customer
  - misses items purchased w/out reviewing

12

Inference for Amazon

- collected data for 6 mo
- only considered customers who reviewed in 6mo. before or during data collection
- each item, each user: retrieved top 10 most related items
- infer: customer purchased t if t appears or rises in related-items list associated with at least K auxiliary items for the customer
  - K parameter
- evaluate with case studies

13

How address privacy issues in search, recommendations, and other information services?

14