Bringing transparency to personalized services through statistical inference

Athanasios Andreou (EURECOM), Oana Goga (MPI-SWS), Patrick Loiseau (EURECOM), Krishna Gummadi (MPI-SWS)

Parties prenantes

EURECOM
Sophia Antipolis

Auteurs

Athanasios Andreou
Oana Goga
Patrick Loiseau
Krishna Gummadi

Partenaires

MAX PLANCK INSTITUTE FOR SOFTWARE SYSTEMS

Personalized services are omnipresent.

Motivation / Goals

- Personalized services are currently omnipresent (e.g., Facebook, Google, targeted advertising).
- Major privacy concerns due to lack of transparency and complexity of the systems.
- **Goals.** Develop methods/tools that:
  - Increase transparency:
  - Enable users to know what can be inferred about them.
  - Increase user control:
  - Enable users to understand and mitigate the information disclosure risks.

Identity vs. Attribute Disclosure Risks for Users with Multiple Social Profiles (2016)

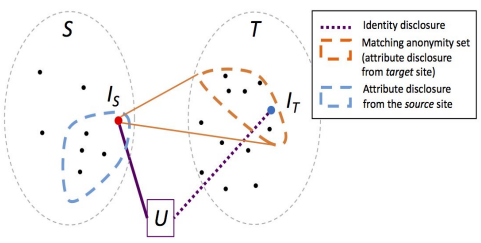

Identity and attribute disclosure in social computing systems.

- Users share a lot of data in several social computing systems.
- Potential for unwanted information disclosure.
- Attackers might employ two strategies in order to infer attributes across social networks for a user:
  - **Identity Disclosure.** Find her matching profile in a different social network and infer something based on it.
  - **Attribute Disclosure.** Infer the value of an attribute, by leveraging information from several users of a different social network.
- **Question.** What is the link between the two types of disclosure?

Contribution

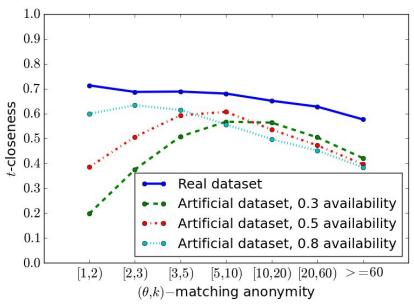

Effect of availability in t-closeness wrt (θ,k)-matching anonymity.

- Proposed a framework to quantify the risks based on k-anonymity/l-diversity/t-closeness.
  - k-anonymity cannot be applied directly since no two identities are identical.
- **(θ,k)-matching anonymity.** k users that are indistinguishable, based on the output of a matching classifier with threshold θ.
- Empirical evaluation on a real world dataset from Facebook and Twitter.
- **Results:** Lower identity disclosure does not always result in lower attribute disclosure. **There is a tradeoff**.
- **Tradeoff depends on:**
  - **Availability** of the targeted attribute.
  - **Uniformity** of the attribute's global distribution.
  - **Correlation** between the attribute and the features used for matching

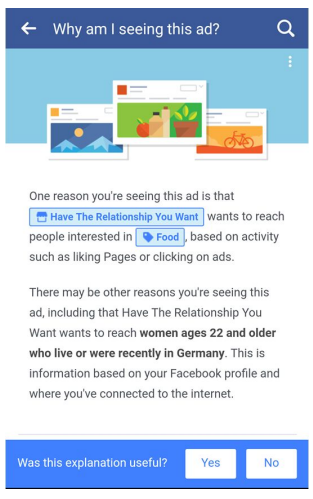

Example of explanations from the "Why Am I Seeing This Ad?" functionality from Facebook.

Future Work

- **Goal:** Reverse engineer why particular ads are targeted to particular individuals
- **Strategy:** Use case on Facebook and exploitation of the "Why Am I Seeing This Ad" feature and the explanations it provides.
- **Questions:**
  - are these explanations complete?
  - are they correct?
  - are they coherent with user interests?
- **Going beyond.** Creation of a framework for crowdsourcing transparency

Contact : andreou@eurecom.fr

Februar 2016 Future & Ruptures Paris