

Lecture 16:
CS 5306 / INFO 5306:
Crowdsourcing and
Human Computation

Netflix Prize Postscript: 2007

Robust De-anonymization of Large Sparse Datasets

Arvind Narayanan and Vitaly Shmatikov

The University of Texas at Austin

Abstract

We present a new class of statistical de-anonymization attacks against high-dimensional micro-data, such as individual preferences, recommendations, transaction records and so on. Our techniques are robust to perturbation in the data and tolerate some mistakes in the adversary's background knowledge.

We apply our de-anonymization methodology to the Netflix Prize dataset, which contains anonymous movie ratings of 500,000 subscribers of Netflix, the world's largest online movie rental service. We demonstrate that an adversary who knows only a little bit about an individual subscriber can easily identify this subscriber's record in the dataset. Using the Internet Movie Database as the source of background knowledge, we successfully identified the Netflix records of known users, uncovering their apparent political preferences and other potentially sensitive information.

and sparsity. Each record contains many attributes (*i.e.*, columns in a database schema), which can be viewed as dimensions. Sparsity means that for the average record, there are no “similar” records in the multi-dimensional space defined by the attributes. This sparsity is empirically well-established [7, 4, 19] and related to the “fat tail” phenomenon: individual transaction and preference records tend to include statistically rare attributes.

Our contributions. Our first contribution is a formal model for privacy breaches in anonymized micro-data (section 3). We present two definitions, one based on the probability of successful de-anonymization, the other on the amount of information recovered about the target. Unlike previous work [25], we do not assume *a priori* that the adversary's knowledge is limited to a fixed set of “quasi-identifier” attributes. Our model thus encompasses a much broader class of de-anonymization attacks than simple cross-database correlation.

Our second contribution is a very general class of de-anonymization algorithms, demonstrating the fundamental limits of privacy in public micro-data (section 4). Under very mild assumptions about the distribution from

1 Introduction

Netflix Prize Postscript: 2009



Press Releases

Current Press Releases

<< [\[Back to Press Releases\]](#)

Netflix Awards \$1 Million Netflix Prize and Announces Second \$1 Million Challenge

NEW YORK, Sept. 21 [/PRNewswire/](#) -- After almost three years and submissions by more than 40,000 teams from 186 countries, Netflix, Inc., the world's largest online movie rental service (NASDAQ: NFLX), today awarded the \$1 million Netflix Prize to a team of engineers, statisticians and researchers who achieved the competition's goal of a 10 percent improvement over the accuracy of the Netflix movie recommendation system when the competition was launched in Oct. 2006. Netflix members already are benefiting from improvements Netflix Prize contestants have contributed to the recommendations system.

Moments after bestowing the \$1 million prize, Netflix announced a second \$1 million challenge, asking the world's computer science and machine learning communities to keep the improvements coming.

The team "BellKor's Pragmatic Chaos," the merging of three teams that had previously competed against one another in the contest, received the \$1 million Netflix Prize in an award ceremony hosted here today by Netflix Co-Founder and CEO Reed Hastings and Chief Product Officer Neil Hunt.

"We had a bona fide race right to the very end," said Mr. Hastings. "Teams that had previously battled it out independently joined forces to surpass the 10 percent barrier. New submissions arrived fast and furious in the closing hours and the competition had more twists and turns than 'The Crying Game,' 'The Usual Suspects' and all the 'Bourne' movies wrapped into one."

The winning team is comprised of software and electrical engineers, statisticians and machine learning researchers from Austria, Canada, Israel and the United States. All seven team members - Bob Bell, Martin Chabbert, Michael Jahrer, Yehuda Koren, Martin Piotte, Andreas Toscher and Chris Volinsky - attended the awards ceremony. It was the first time all seven had met one another in person. How the \$1

MEDIA CENTER

[Press Kit](#)

[Press Releases](#)

[Management](#)

[Image Library](#)

[Upcoming Events](#)

[Investor Relations](#)

[Contact Public Relations](#)

[RSS](#)

[◀ Back to Homepage](#)

Netflix Prize Postscript: 2009

NETFLIX

Press Releases

Current Prize

Netflix Award:

NEW YORK, Sep 15, 2009 - Netflix today awarded the grand prize to the team that achieved the top recommendation benefiting from its new data set.

Moments after becoming the world's computer

The team "BellKor's Pragmatic Chaos" one another in the ceremony by Netflix Co-Fo

"We had a bona fide race right to the very end," said Jim Hastings. "Teams that had previously backed out independently joined forces to surpass the 10 percent barrier. New submissions arrived fast and furious in the closing hours and the competition had more twists and turns than 'The Crying Game,' 'The Usual Suspects' and all the 'Bourne' movies wrapped into one."

The winning team is comprised of software and electrical engineers, statisticians and machine learning researchers from Austria, Canada, Israel and the United States. All seven team members - Bob Bell, Martin Chabbert, Michael Jahrer, Yehuda Koren, Martin Piotte, Andreas Toscher and Chris Volinsky - attended the awards ceremony. It was the first time all seven had met one another in person. How the \$1

MEDIA CENTER

- Press Releases
- Internet
- Library
- Events
- Public Relations
- Homepage

Netflix Prize Postscript: 2009

1 Scott A. Kamber
2 David A. Stampley
3 KamberEdelson, LLC
4 11 Broadway, 22nd Floor.
5 New York, NY 10004
6 Telephone: (212) 920-3072
7 skamber@kamberedelson.com
8 dstampley@kamberedelson.com

9 Joseph H. Malley
10 Law Office of Joseph H. Malley
11 1045 North Zang Blvd
12 Dallas, TX 75208
13 Telephone: (214) 943-6100
14 malleylaw@gmail.com

15 David Parisi (SBN 162248)
16 Suzanne Havens Beckman (SBN 188814)
17 Parisi & Havens LLP
18 15233 Valleyheart Drive
19 Sherman Oaks, California 91403
20 Telephone: (818) 990-1299
21 dcparsi@parisihavens.com
22 shavens@parisihavens.com

23 Attorneys for Plaintiffs

24 **IN THE UNITED STATES DISTRICT COURT
25 FOR THE NORTHERN DISTRICT OF CALIFORNIA
26 SAN JOSE DIVISION**

27 JANE DOE, individually; NELLY VALDEZ-
28 MARQUEZ, ANTHONY SINOPOLI, PAUL
29 NAVARRO, individually and on behalf of a
30 class of similarly situated individuals
31 Plaintiffs,

32 v.
33 NETFLIX, INC., a Delaware Corporation,
34 and DOES 1 THROUGH 50, inclusive,

35 Defendants.

36 **JURY DEMAND**

37 **CLASS ACTION COMPLAINT FOR:**

- 38 1) Video Privacy Protection Act, 18 U.S.C. § 2710
- 39 2) Video Privacy Protection Act, 18 U.S.C. § 2710
- 40 3) California Consumers Legal Remedies Act, Civil Code § 1750
- 41 4) California Customer Records Act, Civil Code § 1798.80
- 42 5) California Unfair Competition Law, Business and Professions Code § 17200
- 43 6) California False Advertising Law, Business and Professions Code § 17500
- 44 7) Unjust Enrichment
- 45 8) Public Disclosure of Private Facts

FILED

2009 DEC 17 A 10:16

RICHARD M. STEWART
CLERK, U.S. DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN JOSE

ADR

E-filing

009 05903

JW

PVT

Netflix Prize Postscript: 2009

7 | I. NATURE OF THE ACTION

8 | 1. On October 2, 2006, Netflix perpetrated the largest voluntary privacy breach to
9 | date, disclosing sensitive and personal indentifying consumer information. The information was
10 | not compromised by malicious intruders. Rather, it was given away to the world freely, and
11 | with fanfare, as part of a contest intended to benefit its trusted custodian, Netflix.

12 | 2. This right to privacy does not appear to be significant to Netflix. This lawsuit is
13 | brought as a class action by and on behalf of similarly situated Netflix subscribers, qualified by
14 | the class definition and class period, whose privacy was violated by the actions of Netflix, Inc.,
15 | (“Netflix”) pursuant to their contest, “Netflix Prize.” Jane Doe, a lesbian, who does not want
16 | her sexuality nor interests in gay and lesbian themed films broadcast to the world, seeks ano-
17 | nymity in this action. Paul Navarro files this action to prevent Netflix from going through with
18 | its announced intentions to make additional disclosures of personal identifying information in-
19 | cluding, but not limited to, users’ video renting history and rating habits.

20 | 3. Netflix knowingly authorized, directed, ratified, approved, acquiesced, or par-
21 | ticipated in the disclosure to third parties of the sensitive information and/or personal identify-
22 | ing information derived from the activity of the Netflix subscribers’ online electronic commu-
23 | nications, when they accessed the Netflix website to rent and rate videos.

24 | 4. Netflix is an “Electronic Communication Service Provider” to its subscribers
25 | and knowingly disclosed to third parties the contents of Netflix’s subscribers’ communications,
26 | including but not limited to, subscribers’ rental and rating videos information, while in elec-

Netflix Prize Postscript: 2010



UNITED STATES OF AMERICA
FEDERAL TRADE COMMISSION
WASHINGTON, D.C. 20580

Maneesha Mithal
Associate Director
Division of Privacy & Identity Protection

Direct Dial: 202.326.2771
Fax : 202.326.3062
E-mail: mmithal@ftc.gov

March 12, 2010

BY E-MAIL & FEDERAL EXPRESS

Reed Freeman
Morrison & Foerster LLP
2000 Pennsylvania Ave., NW
Washington, DC 20006

Dear Mr. Freeman:

On October 13, 2009, staff from the FTC's Division of Privacy and Identity Protection contacted your client, Netflix, Inc. ("Netflix"), regarding the privacy implications of Netflix's planned release of customer movie viewing data in connection with the company's efforts to improve its movie recommendation algorithm. Specifically, staff expressed concern that, despite Netflix's efforts to "anonymize" the customer data prior to its release, it would be possible to re-identify specific customers and thereby associate them with their movie viewing histories and preferences.

Staff's concerns about Netflix's planned release stemmed from research published after the

Netflix Prize Postscript: 2010

NETFLIX

The Netflix Blog

FRIDAY, MARCH 12, 2010

Netflix Prize Update

This is Neil Hunt, Chief Product Officer for Netflix.



About five months ago we announced that Netflix would sponsor a sequel to the Netflix Prize. We've given a lot thought to how to sponsor a contest that discovers more about the predictability of Netflix members' movie watching behavior while always ensuring we protect Netflix members' privacy.

In the past few months, the Federal Trade Commission (FTC) asked us how a Netflix Prize sequel might affect Netflix members' privacy, and a lawsuit was filed by KamberLaw LLC pertaining to the sequel. With both the FTC and the plaintiffs' lawyers, we've had very productive discussions centered on our commitment to protecting our members' privacy.

We have reached an understanding with the FTC and have settled the lawsuit with plaintiffs. The resolution to both matters involves certain parameters for how

LINKS

- [Netflix Tech Blog](#)
- [Netflix America Latina](#)
- [Netflix Brasil](#)
- [Facebook Netflix Page](#)
- [Netflix Website](#)
- [--> RSS Feed Page](#)
- [--> DVDs Releasing This Week](#)
- [--> Logo and Media Materials](#)
- [--> Jobs at Netflix](#)

ABOUT THE NETFLIX BLOG

Hello and welcome to the official Netflix Blog! We the blog authors are various members of the Netflix team. We're also rabid movie fans. We hope this will be a great forum for us to talk about what we are doing, and for you to tell us what you think.

BLOG ARCHIVE

- ▶ [2011](#) (40)
- ▼ [2010](#) (27)
 - ▶ [December](#) (3)
 - ▶ [November](#) (4)
 - ▶ [October](#) (3)

2006 AOL Data Release

- August 4, 2006
- “anonymized” data for 20 million US-based searches for 658,000 users in March-May 2006 (1/3% of all queries, 1.5% of monthly users)
- Sample searches:
 - Names
 - SSNs
 - “how to kill oneself by natural gas”
 - “child porno”
 - “date rape” “is it normal to cook your friend breakfast after he rapes you”
“the morning after being raped” “sexual assault”

2006 AOL Data Release

- NY Times article: “A Face Is Exposed for AOL Searcher No. 4417749”
 - Searches included:
 - “numb fingers”, “60 single men”, “dog that urinates on everything”, “school supplies for Iraq children”, “safest place to live”, “the best season to visit Italy”, “termites”, “tea for good health”, “mature living”, “hand tremors”, “nicotine effects on the body”, “dry mouth”, “bipolar”
 - “landscapers in Lilburn, Ga”, “homes sold in shadow lake subdivision gwinnett county Georgia”, searches for multiple people with last name Arnold
 - Identified Thelma Arnold, age 62 (with her permission)

2006 AOL Data Release

- Data removed August 7, 2006
- CTO and two others lost their jobs
- AOL Spokesperson:

This was a screw up, and we're angry and upset about it. It was an innocent enough attempt to reach out to the academic community with new research tools, but it was obviously not appropriately vetted, and if it had been, it would have been stopped in an instant.

Although there was no personally-identifiable data linked to these accounts, we're absolutely not defending this. It was a mistake, and we apologize. We've launched an internal investigation into what happened, and we are taking steps to ensure that this type of thing never happens again.

- Copies still on the Internet
- Other discoveries...

2006 AOL Data Release

UNITED STATES OF AMERICA
FEDERAL TRADE COMMISSION

_____))
In the Matter of))
))
AOL LCC,) DOCKET NO.
a majority-owned subsidiary of))
TIME WARNER INC.))
_____)

REQUEST FOR INVESTIGATION AND
COMPLAINT FOR INJUNCTIVE RELIEF

The Electronic Frontier Foundation (“EFF”), having reason to believe that AOL LLC (“AOL”) has violated the Federal Trade Commission Act, and that investigation and injunctive relief is in the public interest, alleges that AOL committed unfair and deceptive trade practices by intentionally and publicly disclosing Internet search histories of more than half a million AOL users. Press reports, the analysis of commentators, and EFF’s own research show that these data include sensitive, personal information that can be linked to individuals.

In support of its complaint, EFF alleges as follows.

1. The Electronic Frontier Foundation is a 501(c)(3) nonprofit organization founded in 1990 to protect civil liberties in the digital age. Based in San Francisco, CA, EFF is a membership-supported organization that litigates and educates the public on issues such as free expression, freedom of the press, fair use of copyrighted works, anonymity, security, and privacy as they relate to computing and the Internet.
2. AOL is a Delaware corporation and a majority-owned subsidiary of Time Warner Inc. AOL maintains its principal place of business at 22000 AOL Way, Dulles, VA 20166. Time Warner Inc. is a Delaware corporation and maintains its principal place of business at One Time Warner Center, New York, NY 10019. AOL describes itself to the public as “a Web portal that provides a variety of custom content on top of linking you to relevant information available on the Internet.”¹ AOL also provides Internet connectivity options and specialized client software to its registered users.² AOL’s mission statement says that the company is dedicated “to the simple premise

User 17556639

17556639 how to kill your wife
17556639 how to kill your wife
17556639 wife killer
17556639 how to kill a wife
17556639 poop
17556639 dead people
17556639 pictures of dead people
17556639 killed people
17556639 dead pictures
17556639 dead pictures
17556639 dead pictures
17556639 murder photo

17556639 steak and cheese
17556639 photo of death
17556639 photo of death
17556639 death
17556639 dead people photos
17556639 photo of dead people
17556639 www.murderdpeople.com
17556639 decapitated photos
17556639 decapitated photos
17556639 car crashes3
17556639 car crashes3
17556639 car crash photo

User 672368

- "you're pregnant he doesn't want the baby"
- ...
- "foods to eat when pregnant"
- ...
- "abortion clinics charlotte nc"
- ...
- "can christians be forgiven for abortion"

Others

- "i hate my job"
- "why am i so ugly"
- "i hurt when i think too much i love roadtrips i hate my weight i fear being alone for the rest of my life"
- User 927

Types of Crowdsourcing

- Overt
 - Collecting (Amazon Reviews)
 - Labor Markets (Amazon Mechanical Turk)
 - Collaborative Decisions (Prediction Markets)
 - Collaborative Creation (Wikipedia)
 - Smartest in the Crowd (Contests)
 - Games with a Purpose
- Covert / Crowd Mining
 - Web page linkage, search logs, social media, collaborative filtering

Types of Crowdsourcing

- Overt
 - Collecting (Amazon Reviews)
 - Labor Markets (Amazon Mechanical Turk)
 - Collaborative Decisions (Prediction Markets)
 - Collaborative Creation (Wikipedia)
 - Smartest in the Crowd (Contests)
 - Games with a Purpose
- Covert / Crowd Mining
 - Web page linkage, search logs, social media, collaborative filtering

Types of Crowdsourcing

- Overt
 - Collecting (Amazon Reviews)
 - Labor Markets (Amazon Mechanical Turk)
 - Collaborative Decisions (Prediction Markets)
 - Collaborative Creation (Wikipedia)
 - Smartest in the Crowd (Contests)
 - Games with a Purpose
- Covert / Crowd Mining
 - Web page linkage, search logs, social media, collaborative filtering
- Dark side of crowdsourcing and human intelligence
- Collective intelligence in animals

Types of Crowdsourcing

- Overt
 - Collecting (Amazon Reviews)
 - Labor Markets (Amazon Mechanical Turk)
 - Collaborative Decisions (Prediction Markets)
 - Collaborative Creation (Wikipedia)
 - Smartest in the Crowd (Contests)
 - Games with a Purpose
- Covert / Crowd Mining
 - Web page linkage, search logs, social media, collaborative filtering
- Dark side of crowdsourcing and human intelligence
- Collective intelligence in animals

Games with a Purpose (GWAPs)


- Create a game that people play because they like the game
- As a side-effect of the game, perhaps without the players' awareness, the players are helping solve problems for the game designers

GWAPs: ESPGame

- Defunct
- Videos on YouTube

GWAPs: FoldIt

- 2010
 - "Predicting protein structures with a multiplayer online game"
Seth Cooper, Firas Khatib, Adrien Treuille, Janos Barbero, Jeehyung Lee, Michael Beenen, Andrew Leaver-Fay, David Baker and Foldit players, *Nature* 44p . 756-760 (05 August 2010).
- 2011
 - Structure of Mason-Pfizer monkey virus (M-PMV) retroviral protease (AIDS-causing monkey virus)
 - Unsolved for 15 years
 - Solved in 10 days



GWAPs: EteRNA

- Solve puzzles related to RNA folding
- 100,000 registered players
- 2014:
 - "RNA design rules from a massive open laboratory"
Lee, Jeehyung; Kladwang, Wipapat; Lee, Minjae; Cantu, Daniel; Azizyan, Martin; Kim, Hanjoo; Limpaecher, Alex; Yoon, Sungroh; Treuille, Adrien; Das, Rhiju; EteRNA participants. *PNAS*: Jan 17, 2014
- Feb 2016:
 - Beat supercomputers on RNA design challenge

GWAP: EteRNA


Science AAAS

Home News Journals Topics Careers

Latest News ScienceInsider ScienceShots Sifter From the Magazine About News Quizzes

SHARE

- f 272
- t
- g+
- 134
- @



Artist's diagram of an RNA molecule.

Vossman/Wikimedia Commons

For RNA paper based on a computer game, authorship creates an identity crisis

By **John Bohannon** | Feb. 17, 2016, 5:30 PM

A journal published a paper today that reveals a set of folding constraints in the design of RNA

GWAPs: Principles

- ESPGame:
 - Independence: Two unrelated mutually unknown players
 - Agreement: Same tag
 - Shared information: All they get is shared image
 - Computes labels on images
 - “Output agreement” design pattern

GWAPs: Principles

- TagATune
 - Two players are each given a song, maybe same, maybe not
 - Exchange tags for their songs
 - Decide if it's the same song

 - Computes tags for songs

 - “Input agreement” design pattern

GWAPs: Principles

- Peekabom:
 - One player given image and a secret word, must slowly unveil pieces of the image as clues for the word
 - Second player gets image unveiled bit by bit, must guess word as quickly as possible
 - Computes salient aspects of image for a word
 - “Inversion problem” design pattern