The Engineer's Guide to Data Privacy

SecAppDev 2025

4 June 2025 — Leuven

Vera Rimmer

DistriNet KU LEUVEN

About me

Research Expert in Security Analytics KU Leuven, Belgium

Area: Applied AI, Network and Systems Security, PETs

PhD: “*Applied Deep Learning in Security and Privacy*”

Industry: 4 years in Secure Software Engineering

Outline

- Privacy definitions:
  - societal, legal, and business perspectives
- Engineer's view on privacy
  - privacy vs. security
  - privacy threat modelling
- Privacy-enhancing technologies:
  - general landscape,
  - privacy management tools,
  - private data publishing,
  - measuring leakage

What is Privacy?

"Protect the privacy of the individual from invasion either by the too enterprising press, the photographer, or the possessor of any other modern device for recording or reproducing scenes or sounds."

"The Right to be Left Alone", Warren & Brandeis (1890)

Samuel D. Warren II

"The claim of individuals to determine for themselves when, how and to what extent information about them is communicated to others."

"Privacy and Freedom", Westin (1970)

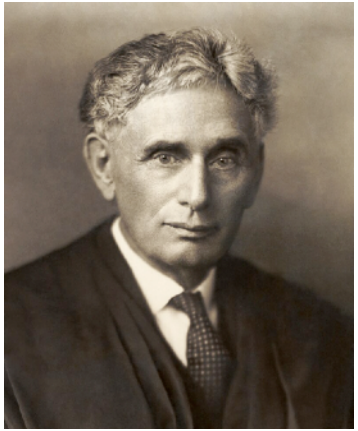

Louis Brandeis

"Protection of the individual against unlimited collection, storage, use and disclosure of their personal data."
"Capacity of the individual to determine in principle the disclosure and use of their personal data."

German constitutional ruling (1983)

Modern Legal Perspective

Art. 8 of the European Convention on Human Rights:

“Respect for private and family life, home and correspondence”

GDPR:

Transparency, consent, purpose, proportionality, accountability of data collection and use.

APRA (2024):

GDPR privacy definitions, but… mandates only opt-out (not opt-in),
data processing is permissible without consent if it aligns with the purposes in the law,
not generally applicable but solely to select businesses.

Business Perspective

Modern societal shifts create a data-driven world

- Data is easy/cheap to collect, store, replicate, transmit, disseminate:
  => IM, emails, files, payments, default location tracking, social graphs (contacts, social networks, meta-data), search histories, biometrics, behaviors…

- Growing computing power makes processing, profiling and inferences cheap too.

- Decision-making is becoming increasingly automated. New products emerge.

Business Perspective

Key flows

Intrinsically assymetrical relationship

users

Personal data

Trust

Functional value

Perceived benefit

organization

Additional risks: data breach? harm by third parties? data misuse? de-identification? surveillance?

Additional value: reselling to third parties, ads, profiling, influence, additional inferences

Business Perspective

Privacy implications can hit hard

Amazon's Massive GDPR Fine Shows the Law's Power–and Limits

It's the first significant GDPR ruling against Big Tech. But secrecy around the decision exposes the regulation's flaws.

WE WERE PROMISED huge fines, and GDPR has finally delivered. Last week Amazon's financial records revealed that officials in Luxembourg are fining the retailer €746 million ($883 million) for breaching the European regulation.

Google Agrees to $392 Million Privacy Settlement With 40 States

Under the agreement, which state attorneys general said was the largest U.S. internet privacy settlement, Google must also make its location-tracking practices clearer to users.

Share full article

Meta Fined $1.3 Billion for Violating E.U. Data Privacy Rules

The Facebook owner said it would appeal an order to stop sending data about European Union users to the United States.

Share full article 234

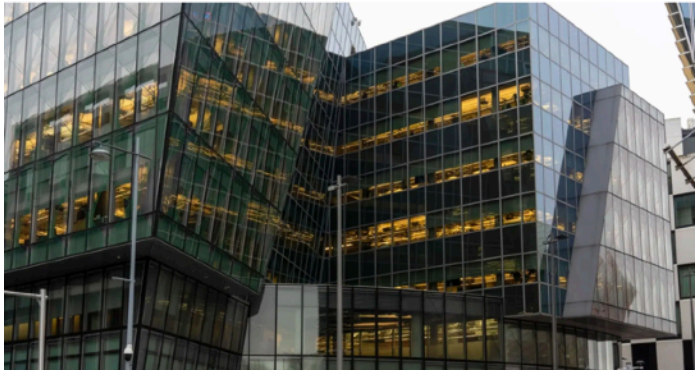

Business Perspective

Modern requirements for data privacy

before:

Data is an **asset.**
Obtain as much as we can store, use it for as long as we want.

now:

Data can be a **liability**.
The line between the **use** and **abuse** of user data is hard to navigate.

How can value be extracted from data while minimizing privacy risks?

Mitigation of Privacy Risks

Wide landscape of privacy measures

PETs

- Governance: regulations that organizations must follow to protect sensitive client information and stay transparent about its use.
- Informed consent mechanisms and user privacy controls and nudges.
- Data minimization.
- Privacy threat modelling.
- Privacy by design.
- **Privacy-enhancing technologies (PETs).**

The best way to do privacy right is proactively, in a multi-disciplinary approach, with various teams in an organization

Who Defines Privacy Goals and Measures?

Immediate social relevance,
but a limited “local” view on privacy

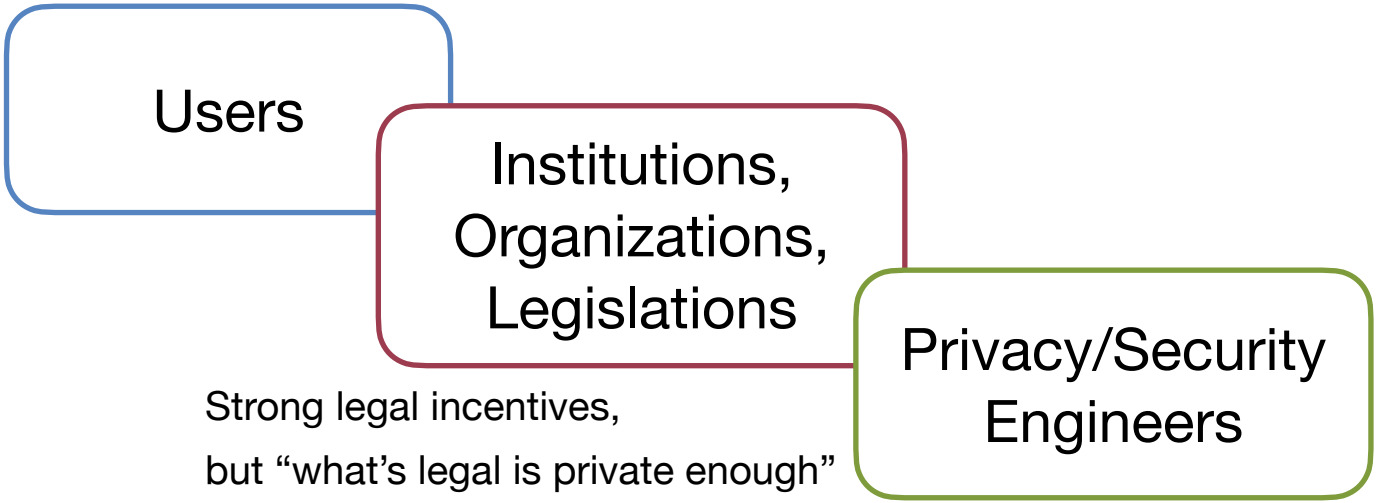

Strong legal incentives,
but “what’s legal is private enough”

Strong security mindset,
but relies on assumptions and strictly
formal/technical privacy definitions

Engineering/Technological Perspective

Compliance teams and/or privacy engineering teams set privacy guidelines and give recommendations… BUT:

- Software engineers build the software that uses the data.
- Software engineers know better what factually happens to the data than a privacy engineer is going to know.
- Software engineers need to correctly implement data management and privacy-enhancing techniques, correctly automate and scale up privacy evaluations, etc.

Engineering/Technological Perspective

- Include privacy as a foundational part of the organization's software development life cycle, and integrate it into code, minimize manual handling.
- Have to cover data throughout its entire life cycle, against various threats.
- Just like with security, privacy is not something that a single team can solve within a company; no single solution/technology exists.

> *"It's our thesis that privacy will be an integral part of the next wave in the technology revolution and that innovators who are emphasizing privacy as an integral part of the product life cycle are on the right track."*

The authors of The Privacy Engineer's Manifesto

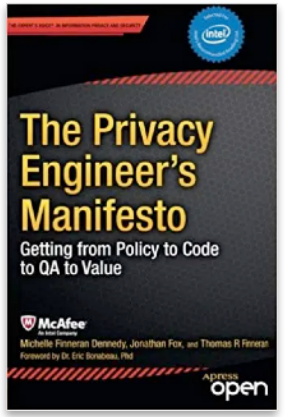

Engineering/Technological Perspective

How do Security (Technologies) and Privacy (Technologies) relate to each other?

Contradict each other?

Enable each other?

Security: protecting digital ASSETS.

Security technologies: Technologies and processes designed to protect general IT assets (i.e., systems, networks, programs, devices and data) from cybercrime (attacks). Protects corporate and national interests.

Carries financial, political and social value.

Privacy: protecting INDIVIDUALS.

Privacy technologies: Technologies and processes designed to protect individual personal and sensitive information, behaviors and preferences.

Fundamental right for privacy.

Engineering/Technological Perspective

How do **Security** (Technologies) and **Privacy** (Technologies) relate to each other?

Security and Privacy may or may not be in trade-off, depending on incentives among the stakeholders

Privacy is Security where the affected entity is an individual and the asset is their private data

Privacy-Enhancing Technologies (PETs)

PETs

- Objective:

  obtain privacy by addressing/mitigating privacy concerns (for individuals and organizations) through technological solutions.

- A broad class of tech. solutions and tools that make privacy enforceable and vary in goals and methods:

Threat models

Privacy goals

Assumptions

https://linddun.org/threat-types/

Privacy Technologies

Examples of Threats (LINDDUN)

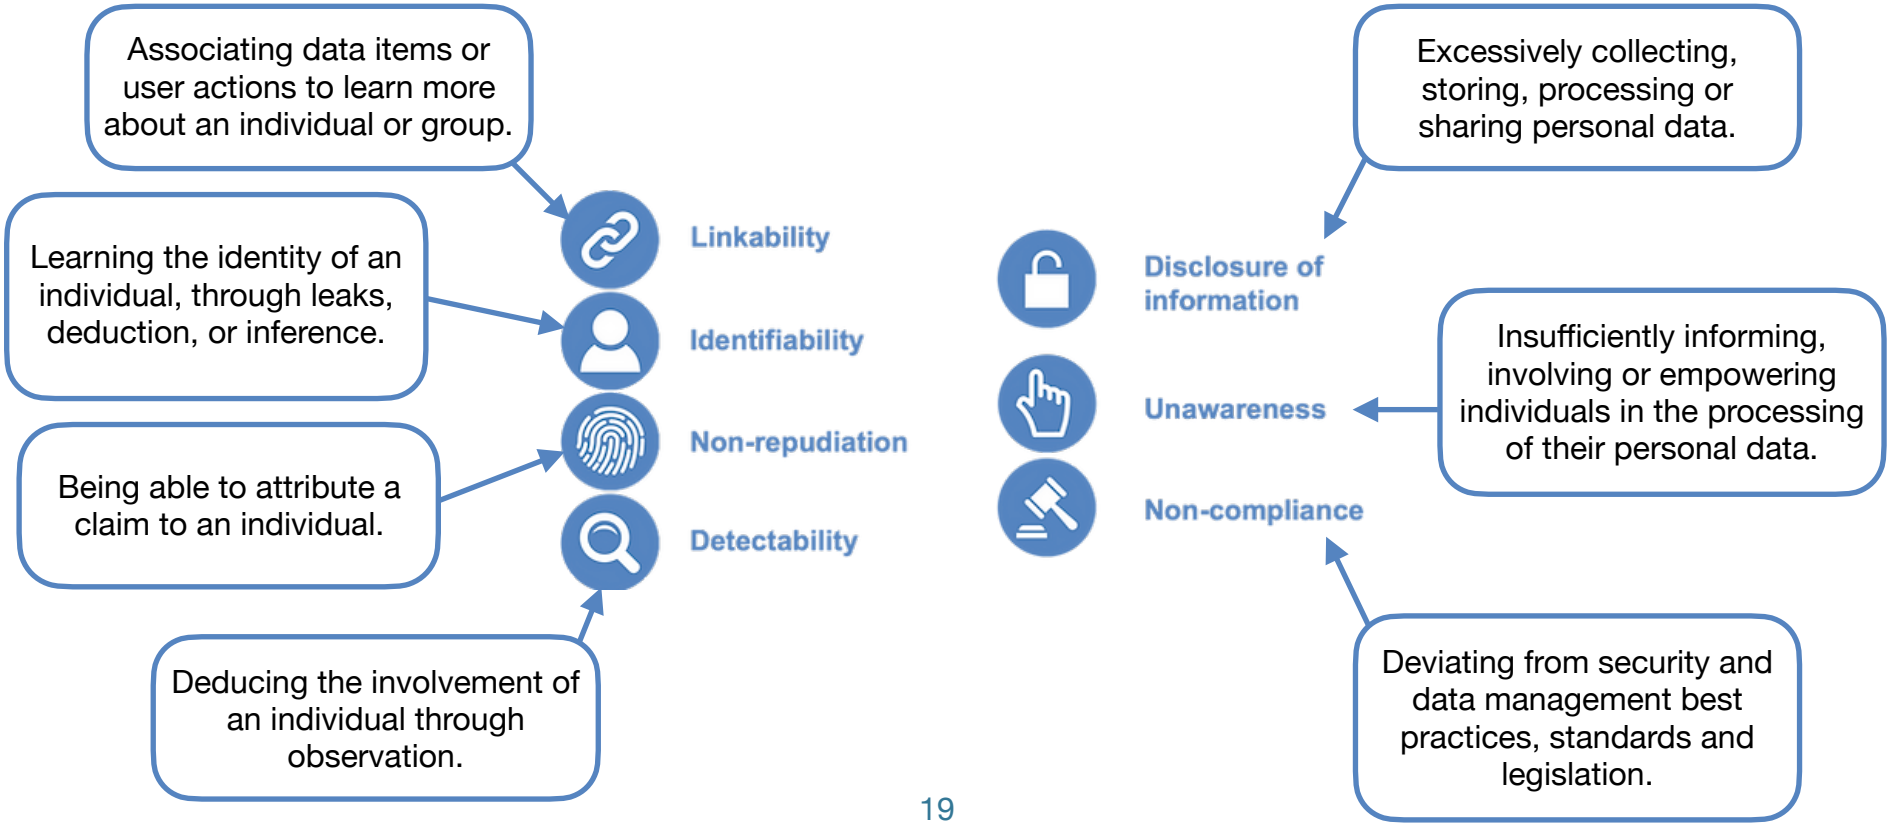

Privacy Technologies

Examples of Goals

- Provide individuals with (some) agency and control over their personal data
- Ensure legal compliance in information processing activities
- Make the data private against untrusted parties
- Resist AI-based inference (e.g., re-identification) and profiling
- Make the data private while at rest / in use / in transit
- …

Privacy Technologies

Examples of Assumptions

this talk

hard PETs:

- Protect privacy through different services and applications without trusting any third parties: no single point of trust, distributed trust.
- Freedom to conceal information and communicate privately, anti-surveillance.

VPN

onion routing

distributed architectures

mixnets

ZKP

soft PETs:

- Allows a trusted data controller to process data while having full control of how data is being used.
- Relies on regulatory compliance, auditing, certification, consent, and control.

SSL

data anonymization

data de-identification

data erasure

differential privacy

Privacy Technologies for Software/Data Engineering

this talk

Traditional protection techniques

Encryption, secure authentication and authorization, secure logging, purpose-based access control…

Privacy management tools

Data mapping, data deletion, data subject requests handling, consent tracking…

Leakage estimation with AI

Inference attacks, model stealing, data reconstruction, etc.

Private computation

homomorphic encryption, secure multi-party computation, zero-knowledge proofs, hardware tools

Private data publishing and analysis

de-identification, (pseudo-) anonymization, differential privacy, synthetic data

Private AI/ML

Differentially private learning, federated learning, etc.

I. Privacy Management Tools

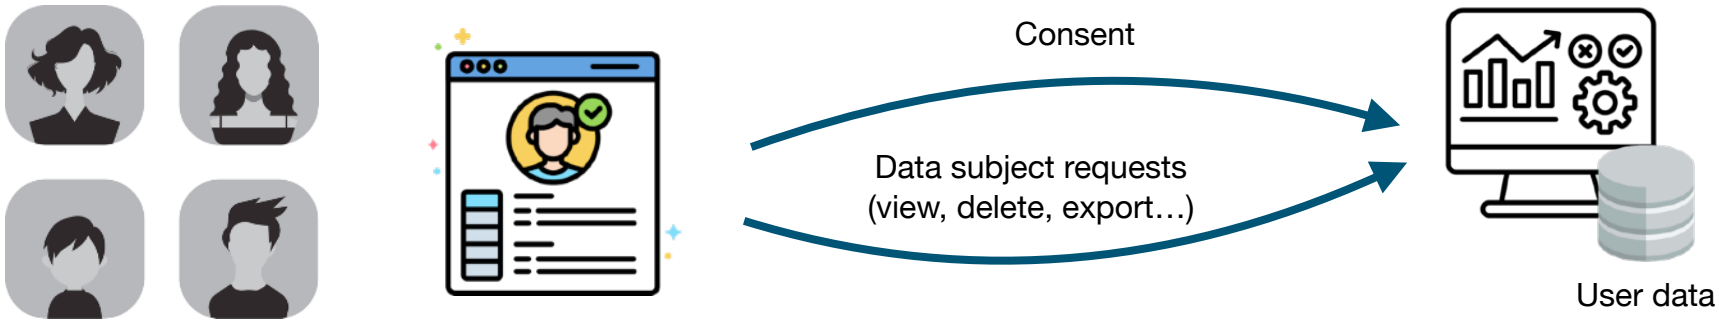

- Consent management platforms
  - Do the users consent? How can it be traced back and proven?
- Record of processing activities
  - What systems consume what type of data and why?
  - Network traffic and application code analysis.
- Data subject requests handling
  - Right to erasure, to restriction, to objection.
  - Right to access, to portability.

I. Privacy Management Tools

Tools and references

Klaro! A Simple Consent Manager

Klaro [klɛro] is a simple consent management platform (CMP) and privacy tool that helps you to be transparent about the third-party applications on your website. It is designed to be extremely simple, intuitive and easy to use while allowing you to be compliant with all relevant regulations (notably GDPR and ePrivacy).

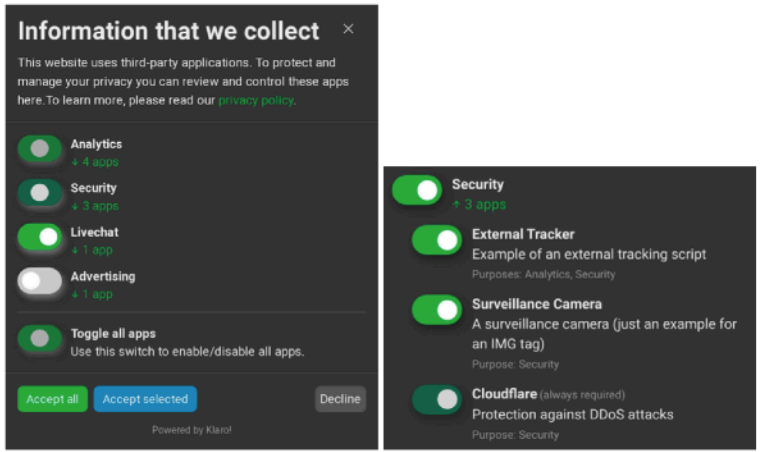

Klaro supports multiple modes of asking for consent and can display third-party apps individually or grouped by purpose.

Meet Fides: Privacy as Code

release v2.62.0 | Docker Build & Push no status | Publish Docs passing | Publish Fides passing | license Apache 2 | code style black | mypy checked | Follow @ethyca | codecov 87%

fides

Overview

Fides (pronounced */fee-dhez/*, from Latin: Fidēs) is an open-source privacy engineering platform for managing the fulfillment of data privacy requests in your runtime environment, and the enforcement of privacy regulations in your code.

- End-to-end data subject request automation
- Privacy-as-code
- Compliance-minded data mapping

I. Privacy Management Tools

Tools and references

What is Privado?

slack privado | docs gitbook

Privado is an open-source static code analysis tool to discover data flows in the code. It detects more than 110 personal data elements being processed and further maps the data flow from the point of collection to "sinks" such as external third parties, databases, logs, and internal APIs.

```
Exported output json to '/Users/source/my-app' folder

-------------------------
SUMMARY
-------------------------
DATA ELEMENTS       |   5
THIRD PARTIES FLOWS |   3
STORAGES            |   2
DATA LEAKAGES       |   4
ISSUES              |   2
-------------------------

> Successfully synchronized results with Privado Cloud
> Continue to view results on:
"https://community.privado.ai/repositories/overview/my-app"
```

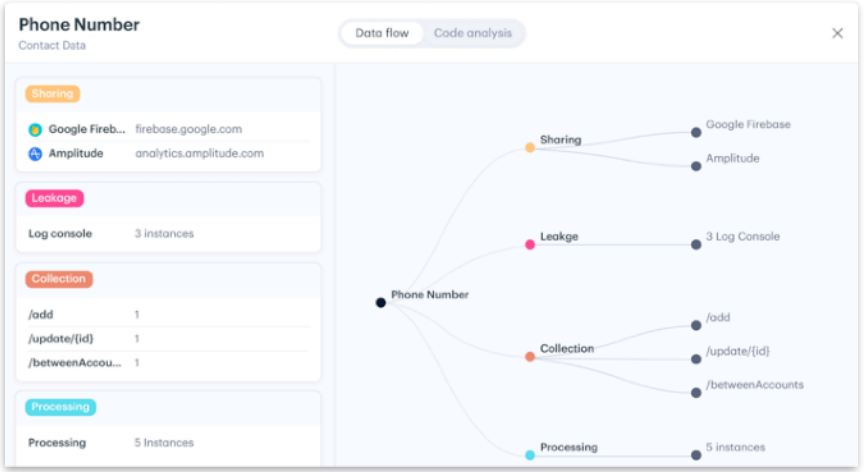

- Generate and maintain Data maps, Record of Processing Activity Reports, and the data-flow diagrams
- Identify and remove data leaks
- Find and fix unaccounted third-party sharing of data
- Do continuous monitoring for privacy and data issues
- …

II. Private Data Publishing and Analysis

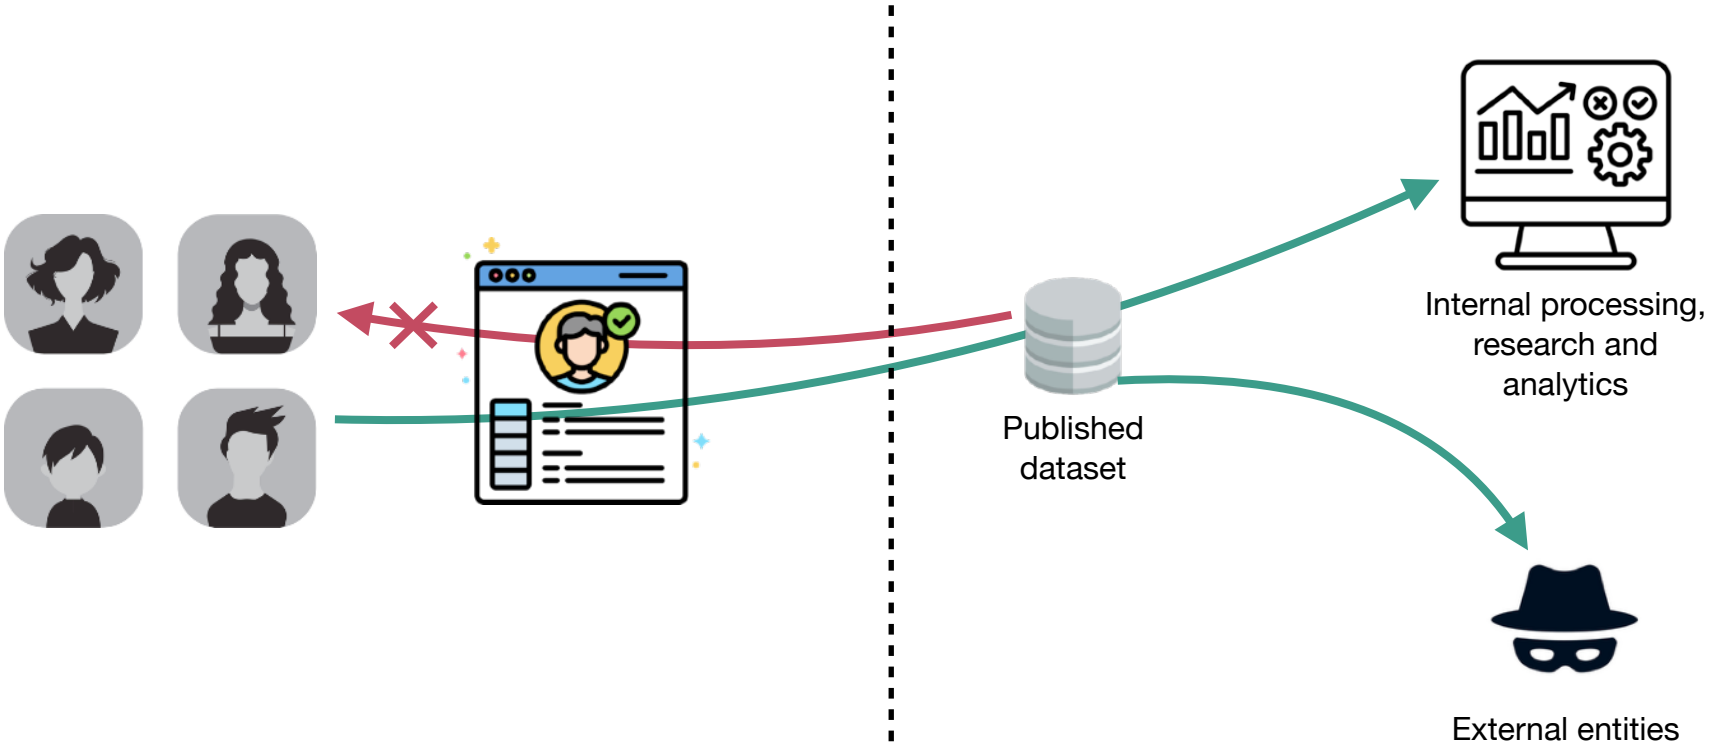

Anonymity

Anonymity is the state of being not identifiable within a set of subjects, the anonymity set

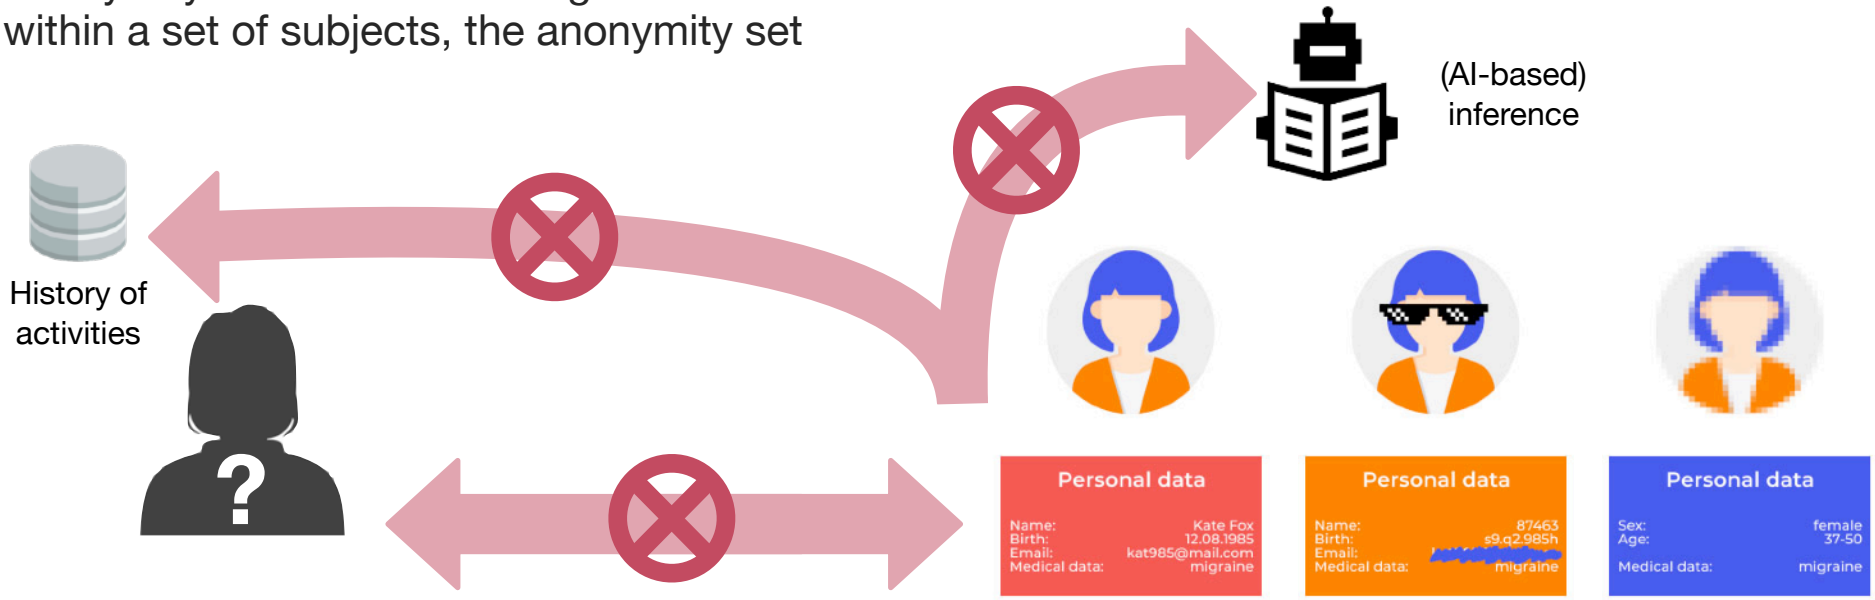

- Physically existing individual
- Device owner
- Account owner
- Subject related to a DB record
- Author of text online
- …

- Name
- National ID
- Biometric
- IP address
- Key to a DB
- …

Pseudoanonymization

=> reversible (w secret)

De-Identification

=> “irreversible”

Private Data Publishing and Analysis

Data anonymization / de-identification

Let's remove all unique identifiers… what can go wrong?

L. Sweeney, Simple Demographics Often Identify People Uniquely. Carnegie Mellon University, Data Privacy Working Paper 3. Pittsburgh 2000.

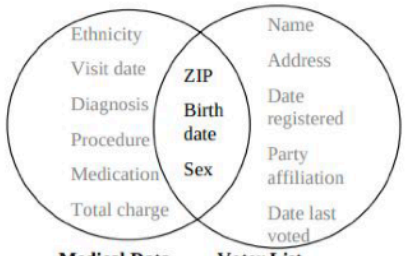

Figure 1 Linking to re-identify data

Access to aux information

Naive approach:

- Recognize and remove identifiers

Challenge:

- Linkage attacks: Cannot assume we know all the relevant information that is or might ever become available to the adversary!

=> attackers can join anonymized data records with other public or purchased data.
=> combination of (ZIP, Birth date, Sex) was unique for 87% of the corpus! (quasi-identifiers)

Private Data Publishing and Analysis

Data anonymization / de-identification

Other cases:

- The Netflix Prize dataset: joining with the IMDB dataset led to full re-identification (99% of records based on 8 movie preferences, out of which 2 can be wrong)
- Location data: Mobility traces used to identify 95% of individuals based on ...4... points
  And even coarser granularity does not improve anonymity by much.

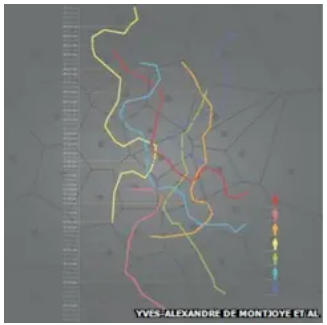

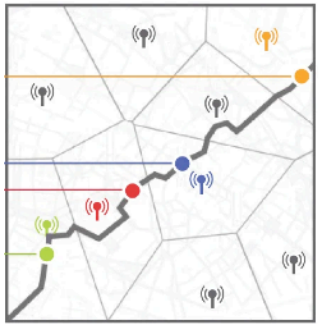

Huge universe of combinations
Sparsity/uniqueness
=> Re-identifiable even with noise!

Private Data Publishing and Analysis

k-Anonymity

Every distinct combination of quasi-identifiers designates at least k records.

Techniques:

- Generalization: roll up a field to a more common value
- Suppression: replace sensitive fields with a default value

Can we anonymize data without a significant loss in its utility?

Name	Beverage	Province	Age
Leslie	Espresso	BC	22
Lee	Latte	BC	23
Pat	Pour Over	AB	29
Lyn	London Fog	BC	37
Sam	Matcha Latte	ON	39

Name	Beverage	Country	Age
*	Coffee	Canada	20-30
*	Coffee	Canada	20-30
*	Coffee	Canada	20-30
*	Tea	Canada	30-40
*	Tea	Canada	30-40

Sweeney, Latanya. "k-anonymity: A model for protecting privacy."

Private Data Publishing and Analysis

k-Anonymity

Challenges:

- Have to define in advance what a potential adversary knows.
- Linkage attacks are still possible: ***all*** values are potentially identifying, depending on their prevalence in the population and on auxiliary data that the attacker may have.

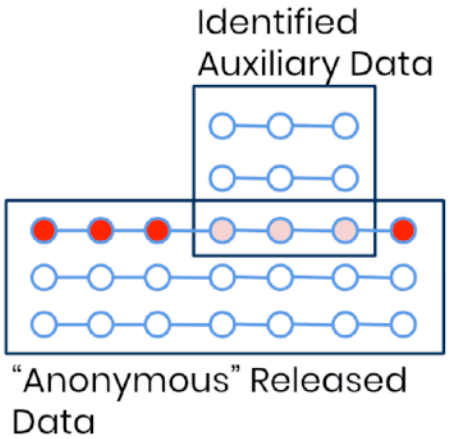

Can we have a technique with strong/provable protection?

Private Data Publishing and Analysis

Differential privacy

- Idea: if one individual is removed or added to the original dataset, it does not change what we can conclude from the final dataset.
- DP: adding random tunable “noise” to the data in a way that obscures the identity and data of the individuals but keeps the database useful overall as a source of statistical information.
- Great for large-scale aggregations (count, sum, etc.)
- **Strong mathematically provable privacy guarantees** within a certain budget

Private Data Publishing and Analysis

Differential privacy

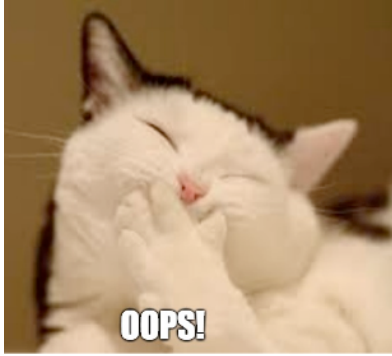

Challenges:

- Hard to do right, especially on non-trivial data structures.
- Might not work for some datasets: strips the data of its value.
- Querying the dataset multiple times diminishes its privacy.

SoK: Differentially Private Publication of Trajectory Data

Àlex Miranda-Pascual
Universitat Politècnica de Catalunya
Karlsruhe Institute of Technology
Karlsruhe, Germany
alex.miranda.pascual@upc.edu

Patricia Guerra-Balboa
Karlsruhe Institute of Technology
Karlsruhe, Germany
patricia.balboa@kit.edu

Javier Parra-Arnau
Universitat Politècnica de Catalunya
Karlsruhe Institute of Technology
Barcelona, Spain
javier.parra-arnau@kit.edu

Jordi Forné
Universitat Politècnica de Catalunya
Barcelona, Spain
jordi.forne@upc.edu

Thorsten Strufe
Karlsruhe Institute of Technology
Karlsruhe, Germany
thorsten.strufe@kit.edu

Private Data Publishing and Analysis

Synthetic data generation

- Idea: create a synthetic version of private/sensitive data for processing and analysis (and augmentation, imputation, etc.)
- Uses generators based on modern machine learning (GenAI).
- Can be combined with differential privacy.

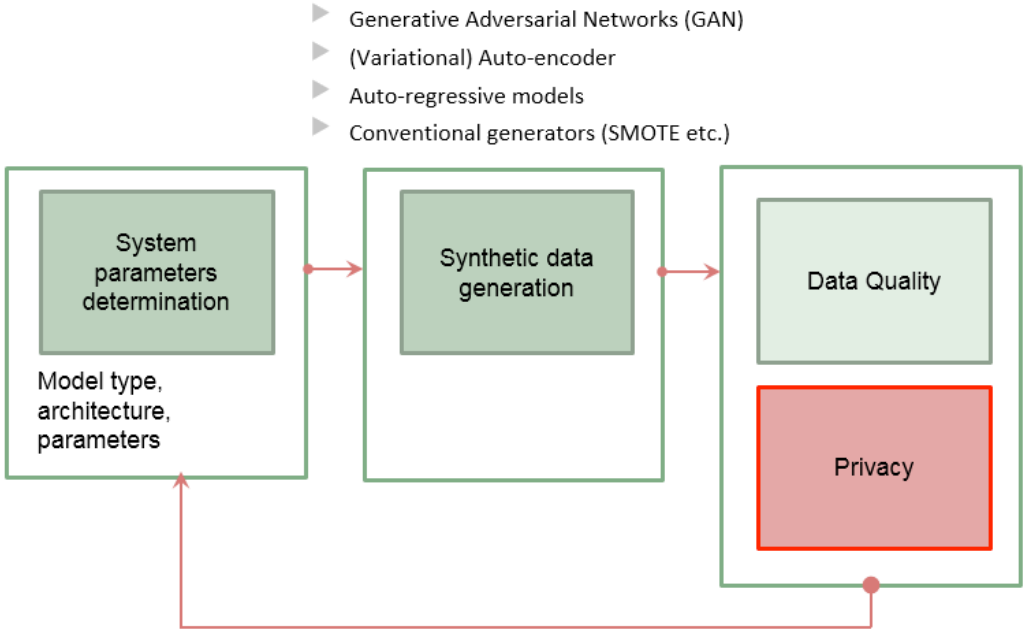

Private Data Publishing and Analysis

Synthetic data generation

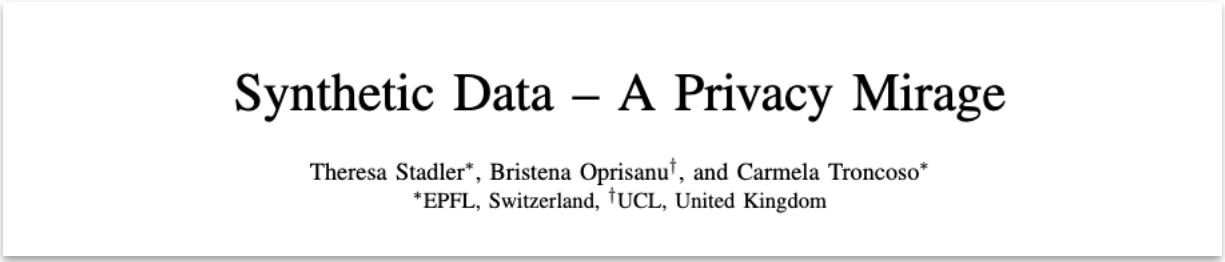

Synthetic Data – A Privacy Mirage

Theresa Stadler*, Bristena Oprisanu†, and Carmela Troncoso*
*EPFL, Switzerland, †UCL, United Kingdom

- High-dimensional synthetic datasets preserve a lot of information about the raw data.
- Privacy attacks work even on generators trained under differential privacy.
- Protection levels are not uniform across the data points.

No data anonymization/de-identification technique is a silver bullet!

Private Data Publishing and Analysis

Privacy-utility trade-off

Strong privacy primitives,

drastic impact on utility

Empirical privacy protection,

no mathematical guarantees,

optimal utility

Private Data Publishing and Analysis

Tools and references (1)

- k-anonymity
- k-map
- l-diversity
- delta-presence
- t-closeness

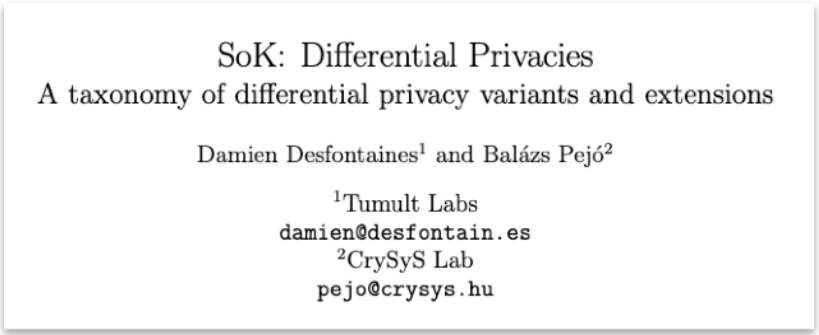

SoK: Differential Privacies

A taxonomy of differential privacy variants and extensions

Damien Desfontaines[1] and Balázs Pejó[2]

[1]Tumult Labs
damien@desfontain.es
[2]CrySyS Lab
pejo@crysys.hu

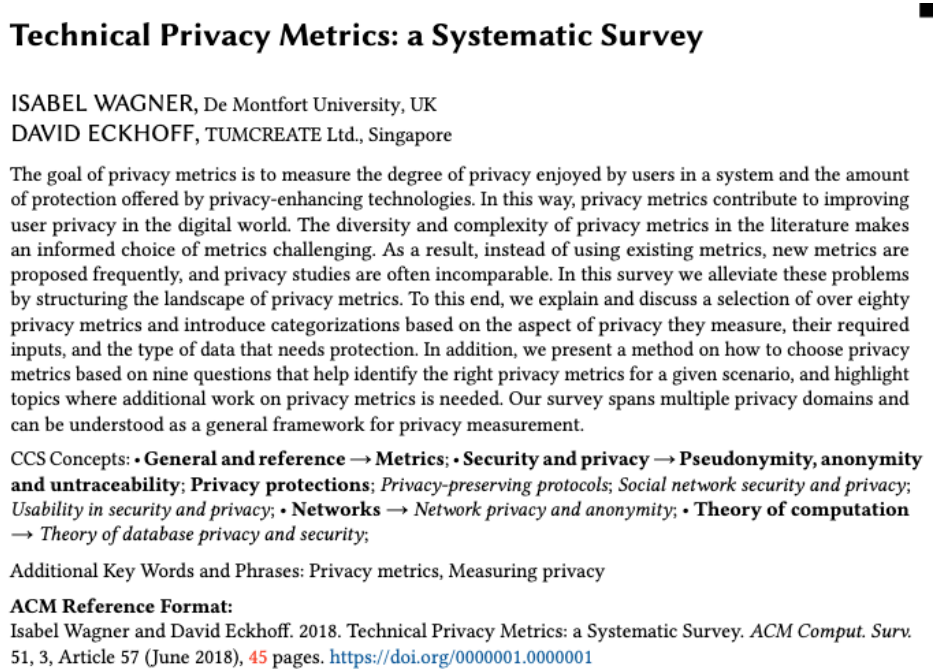

Technical Privacy Metrics: a Systematic Survey

ISABEL WAGNER, De Montfort University, UK
DAVID ECKHOFF, TUMCREATE Ltd., Singapore

The goal of privacy metrics is to measure the degree of privacy enjoyed by users in a system and the amount of protection offered by privacy-enhancing technologies. In this way, privacy metrics contribute to improving user privacy in the digital world. The diversity and complexity of privacy metrics in the literature makes an informed choice of metrics challenging. As a result, instead of using existing metrics, new metrics are proposed frequently, and privacy studies are often incomparable. In this survey we alleviate these problems by structuring the landscape of privacy metrics. To this end, we explain and discuss a selection of over eighty privacy metrics and introduce categorizations based on the aspect of privacy they measure, their required inputs, and the type of data that needs protection. In addition, we present a method on how to choose privacy metrics based on nine questions that help identify the right privacy metrics for a given scenario, and highlight topics where additional work on privacy metrics is needed. Our survey spans multiple privacy domains and can be understood as a general framework for privacy measurement.

CCS Concepts: • **General and reference** → **Metrics**; • **Security and privacy** → **Pseudonymity, anonymity and untraceability**; **Privacy protections**; *Privacy-preserving protocols*; *Social network security and privacy*; *Usability in security and privacy*; • **Networks** → *Network privacy and anonymity*; • **Theory of computation** → *Theory of database privacy and security*;

Additional Key Words and Phrases: Privacy metrics, Measuring privacy

ACM Reference Format:
Isabel Wagner and David Eckhoff. 2018. Technical Privacy Metrics: a Systematic Survey. *ACM Comput. Surv.* 51, 3, Article 57 (June 2018), 45 pages. https://doi.org/0000001.0000001

https://arxiv.org/pdf/1512.00327

Private Data Publishing and Analysis

Tools and references (2)

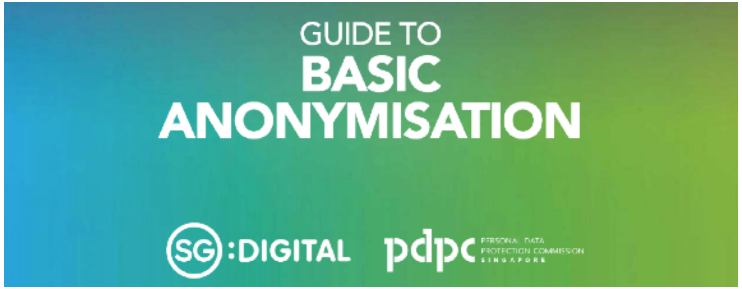

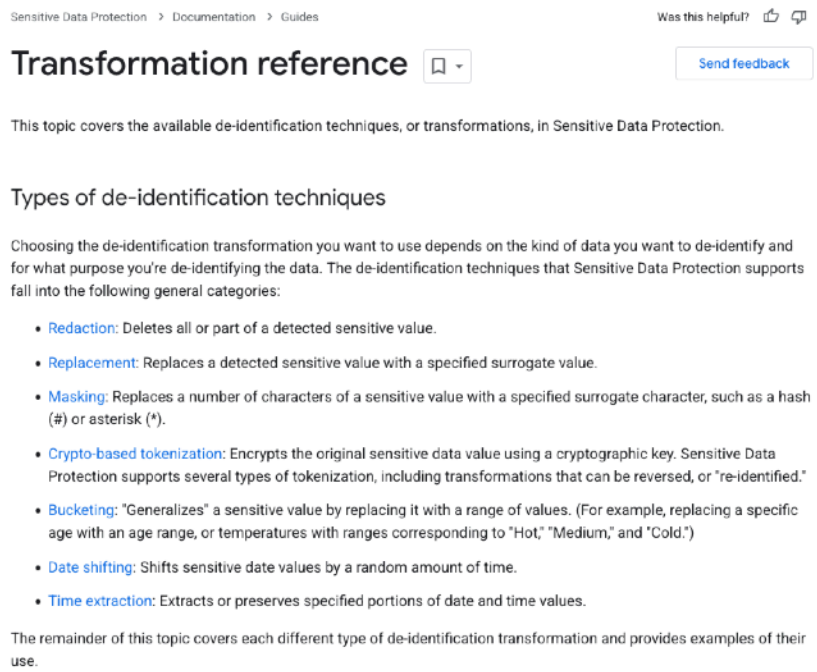

Sensitive Data Protection > Documentation > Guides

Was this helpful?

Transformation reference

Send feedback

This topic covers the available de-identification techniques, or transformations, in Sensitive Data Protection.

Types of de-identification techniques

Choosing the de-identification transformation you want to use depends on the kind of data you want to de-identify and for what purpose you're de-identifying the data. The de-identification techniques that Sensitive Data Protection supports fall into the following general categories:

- Redaction: Deletes all or part of a detected sensitive value.
- Replacement: Replaces a detected sensitive value with a specified surrogate value.
- Masking: Replaces a number of characters of a sensitive value with a specified surrogate character, such as a hash (#) or asterisk (*).
- Crypto-based tokenization: Encrypts the original sensitive data value using a cryptographic key. Sensitive Data Protection supports several types of tokenization, including transformations that can be reversed, or "re-identified."
- Bucketing: "Generalizes" a sensitive value by replacing it with a range of values. (For example, replacing a specific age with an age range, or temperatures with ranges corresponding to "Hot," "Medium," and "Cold.")
- Date shifting: Shifts sensitive date values by a random amount of time.
- Time extraction: Extracts or preserves specified portions of date and time values.

The remainder of this topic covers each different type of de-identification transformation and provides examples of their use.

Sensitive Data Protection (Google)

Private Data Publishing and Analysis

Tools and references (3)

differential-privacy Public

build https://github.com/badges/shields/issues/8671 pypi v1.1.5.rc4 license Apache-2.0

PyDP

programming-dp Public

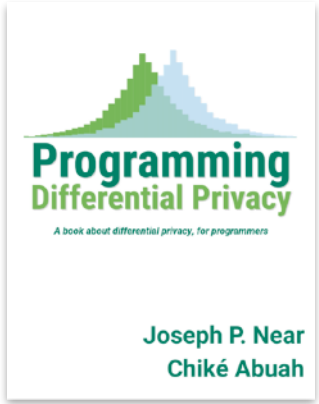

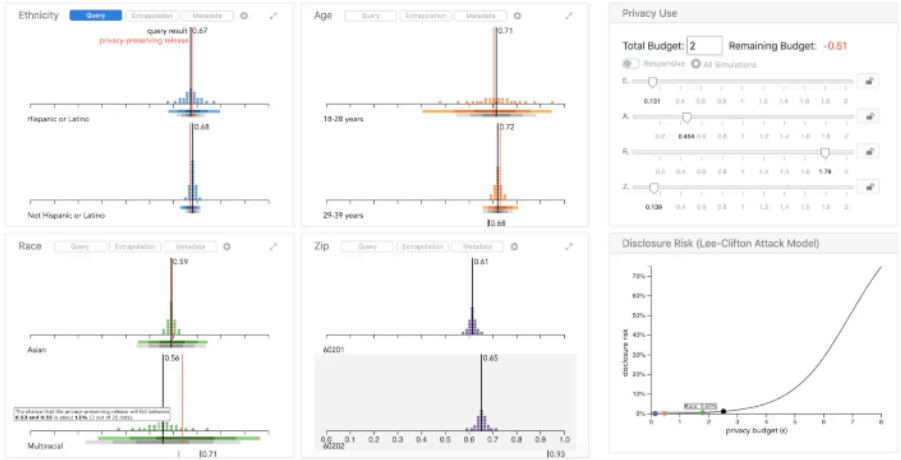

sciendo
Proceedings on Privacy Enhancing Technologies ; 2022 (2):601–618

Priyanka Nanayakkara*, Johes Bater, Xi He, Jessica Hullman, and Jennie Rogers

Visualizing Privacy-Utility Trade-Offs in Differentially Private Data Releases

III. Measuring Privacy Leakage

Applied best protection techniques… job done?

- Basic protections (e.g., pseudoanonymization, de-identification) are necessary but insufficient, they need to be properly vetted, typically through **audits** and **offensive techniques**.
- Leakage occurs via **correlation, overfitting, memorization, indirect inference.**
- Regulators (e.g., GDPR) increasingly recognize **inference as a privacy risk.**

Leakage types

- **Re-identification**: Linking records back to individuals
- **Profiling**: Building behavioral or demographic models
- **Memorization**: Unintended retention of data in models
- **Inference**: Deriving sensitive attributes (e.g., health, location)

Measuring Privacy Leakage

Re-identification

- Estimate the risk of record linkage using aux data and background knowledge
- Run inference attacks on noisy aggregates (e.g., with DP audit scripts)
- For unstructured data (text, images): de-obfuscate with AI
- Profiling-based re-identification / linkage:
  - Infer identity from past behavior and attributes (mobility traces, web activity, purchase logs, etc.)
  - Infer stable attributes from profiled behavior, then use those to link datasets.

De-identification is not bulletproof. Requires context-aware, threat-model-specific tests/attacks!

Measuring Privacy Leakage

Memorization: Training data reconstruction

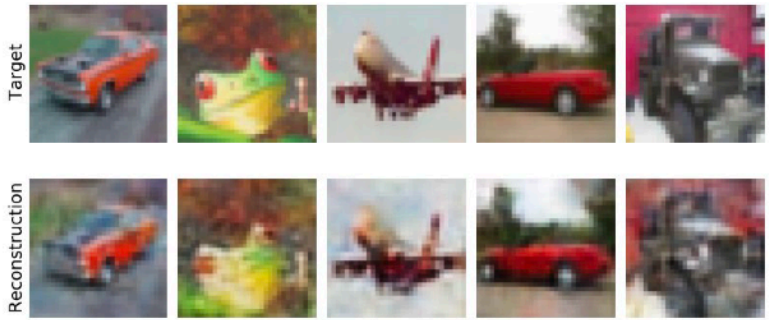

Fig. 1: Examples of training data points reconstructed from a 55K parameter CNN classifier trained on CIFAR-10.

ML models optimized for performance only memorize data!

Reconstructing training data with informed adversaries (S&P, 2022)

Measuring Privacy Leakage

Memorization: Membership inference

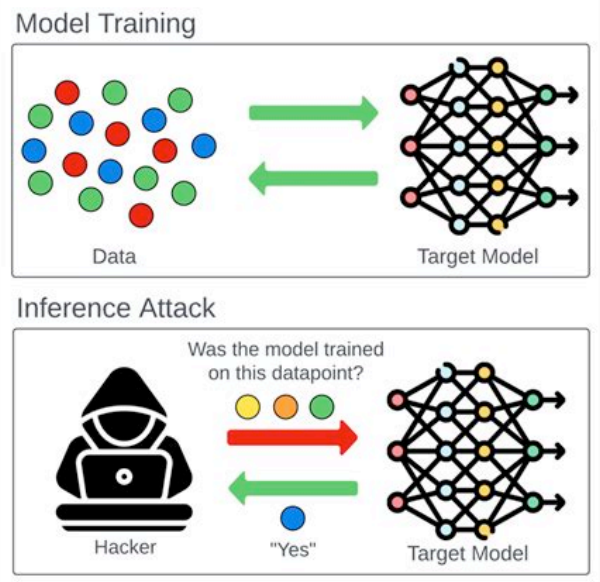

Membership information as an estimate of memorization

Measuring Privacy Leakage

Inference: Prediction of sensitive data

- Sensitive information can often be inferred from non-sensitive data.
- E.g.: from interaction data, algorithms can infer who a person's significant other is, their wealth, demographics, the propensity to overspend, personality traits, and other attributes.
- Arguably, behavioral data are as sensitive as the content of the communication: "metadata are data" (E.g., mobile phone metadata have been at the core of the Snowden revelations and their collection was later deemed illegal)

> *"Metadata absolutely tells you everything about somebody's life. If you have enough metadata, you don't really need content."*

NSA General Counsel Stewart Baker

Measuring Privacy Leakage

- Empirical testing of ML models under various threats:
  - Train shadow models to simulate attacks

(membership/attribute inference, reconstruction, model inversion, profiling-based, etc.)

  - Common in research papers and ML red-teaming.
  - Can be based on artificially inserted “canaries”.
  - Needs to be context-aware and threat-model-specific!

Measuring Privacy Leakage

Tools and references

Awesome Privacy Engineering

awesome

> A curated list of resources related to privacy engineering

Privacy Meter

python 3.12 | slack @privacy meter | License MIT | cite citation | contributors 14

Forks 111 | Stars 651 | Open in Colab

What is Privacy Meter?

Privacy Meter is an open-source library to audit data p
statistical and machine learning algorithms (classificat
vision, and natural language processing). The tool ena
assessment based on the state-of-the-art membershi

TensorFlow Privacy

Tensorflow Privacy (TF Privacy) is an open source library developed by teams in Google Research. The library includes implementations of commonly used TensorFlow Optimizers for training ML models with DP. The goal is to enable ML practitioners using standard Tensorflow APIs to train privacy-preserving models by changing only a few lines of code.

Takeaways

- **Privacy is multifaceted**; its understanding depends on stakeholders, social norms and expectations, and general context.
- **Diverse landscape** of PETs, different levels of dependency on technologies, different approaches to trust, assumptions, threat models, usability, etc.
- ***The word "Private" by itself does not tell us anything.*** Always ask what it means, who defined it, how it is achieved and how it is measured!
- Anonymity is difficult and fragile. Any "anonymized" dataset is probably **re-identifiable**. But we can learn and apply suitable privacy metrics!
- No need to "reinvent the wheel" given **existing tools**, but use with care, develop a profound understanding, and involve all relevant teams.

@DistriNet (KU Leuven, Belgium)

https://distrinet.cs.kuleuven.be/
https://blue41.cs.kuleuven.be

PhD Summer Schools

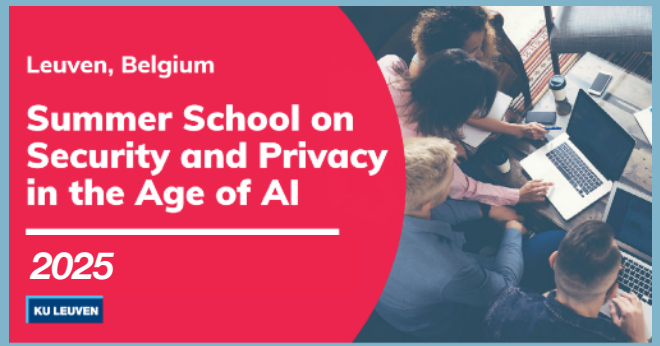

Training & Outreach

Advanced Master's