Privacy for Data Scientists

Utilizing Privacy-Aware Methods for Data Science

KIProtect

EuroPython 2018

Data Privacy: For Whom?

KOURTNEY KARDASHIAN
SCOTT DISICK

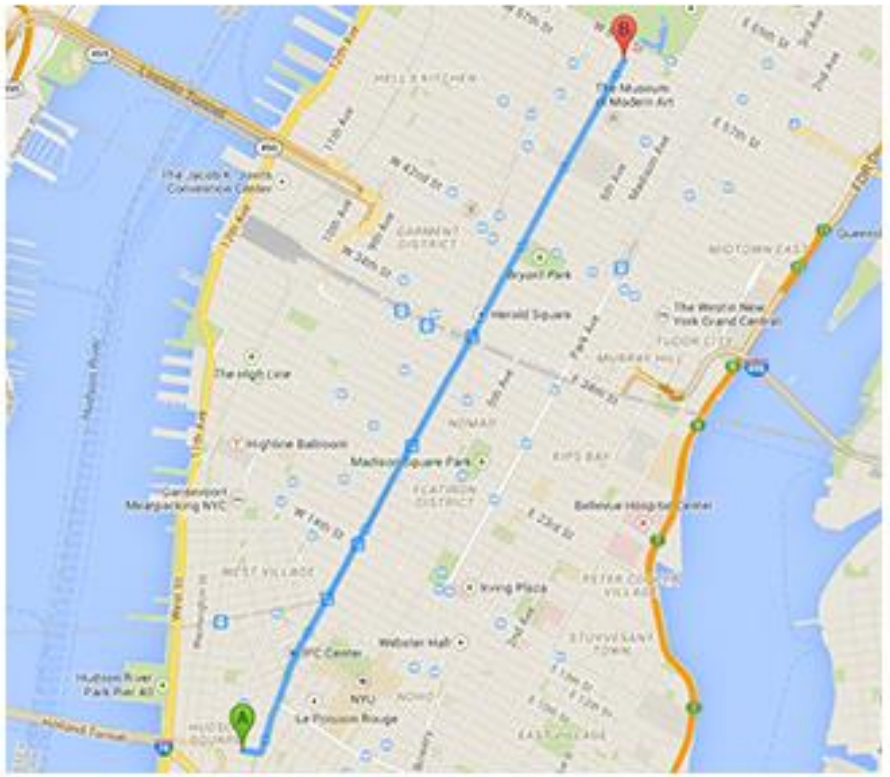

NOVEMBER 4, 2013 • 12:11 PM - 12:36 PM
246 SPRING ST. TO 1412 6TH AVE
$16.50 FARE • $3.40 TIP • ©SPLASH

Poorly “Anonymized” Data: NYC Taxi Rides

(released under FOIA and used a md5 hash)

Graphic: http://gawker.com/the-public-nyc-taxicab-database-that-accidentally-track-1646724546

Data Privacy: From Whom?

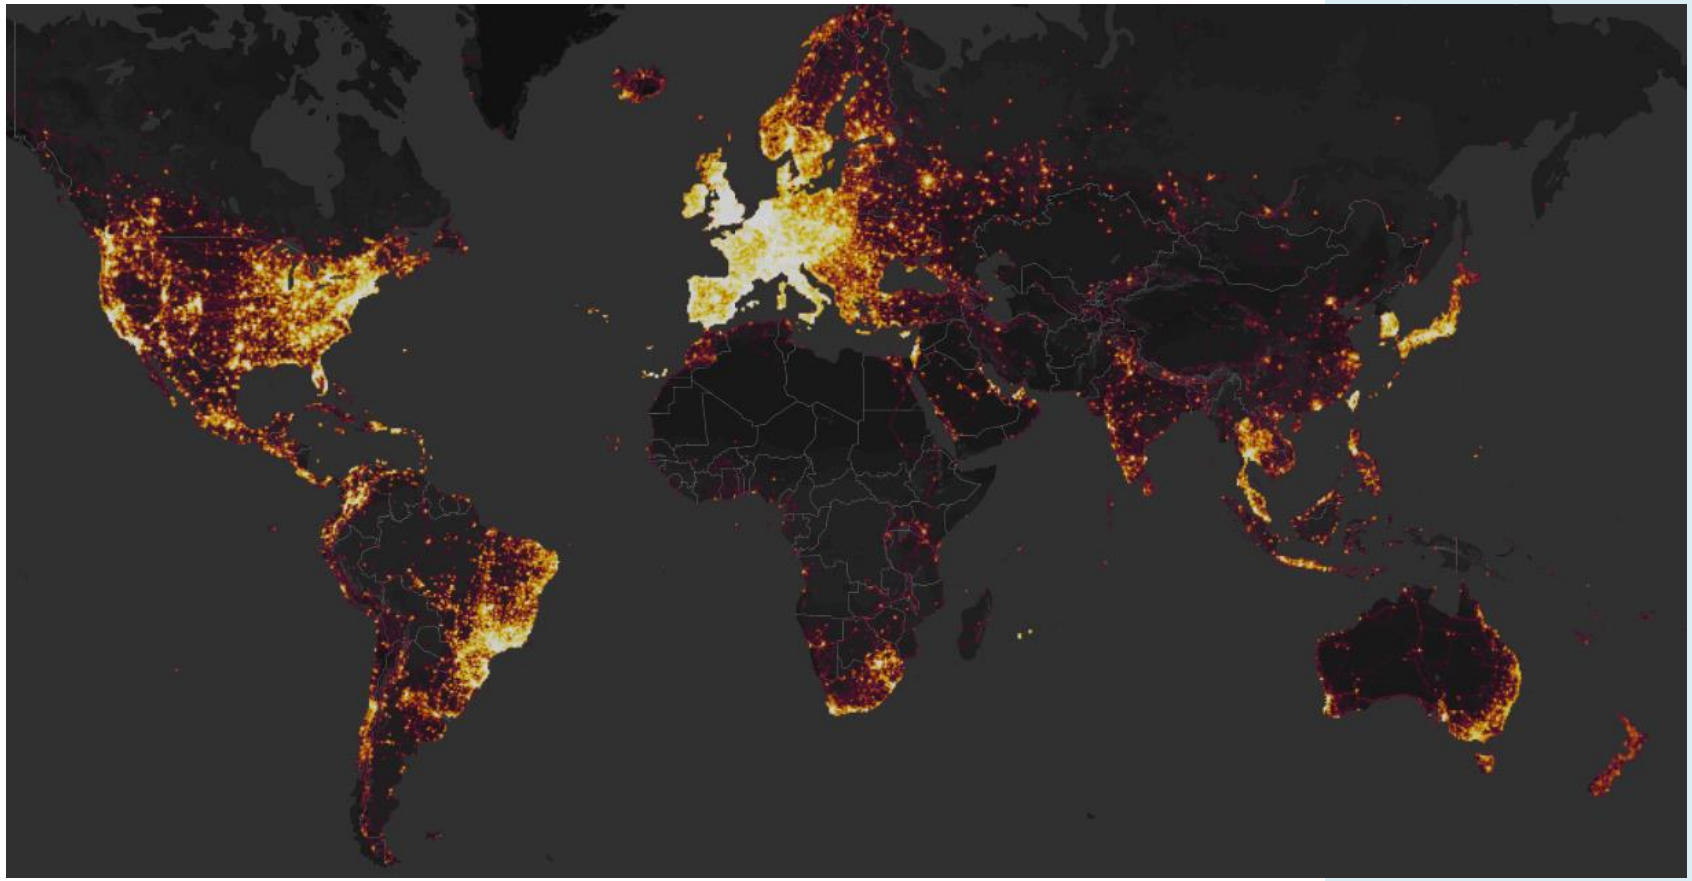

Aggregated
Anonymized Data
Leaks Secrets

(secret military bases revealed from aggregated sport data)

Strava Heatmap: https://www.strava.com/heatmap

About the Instructors

Dr. Andreas Dewes

Dr. rer. nat., Dipl. Phys., Dipl. Kfm.

Katharine Jarmul

KIPROTECT

What even **is** Privacy?

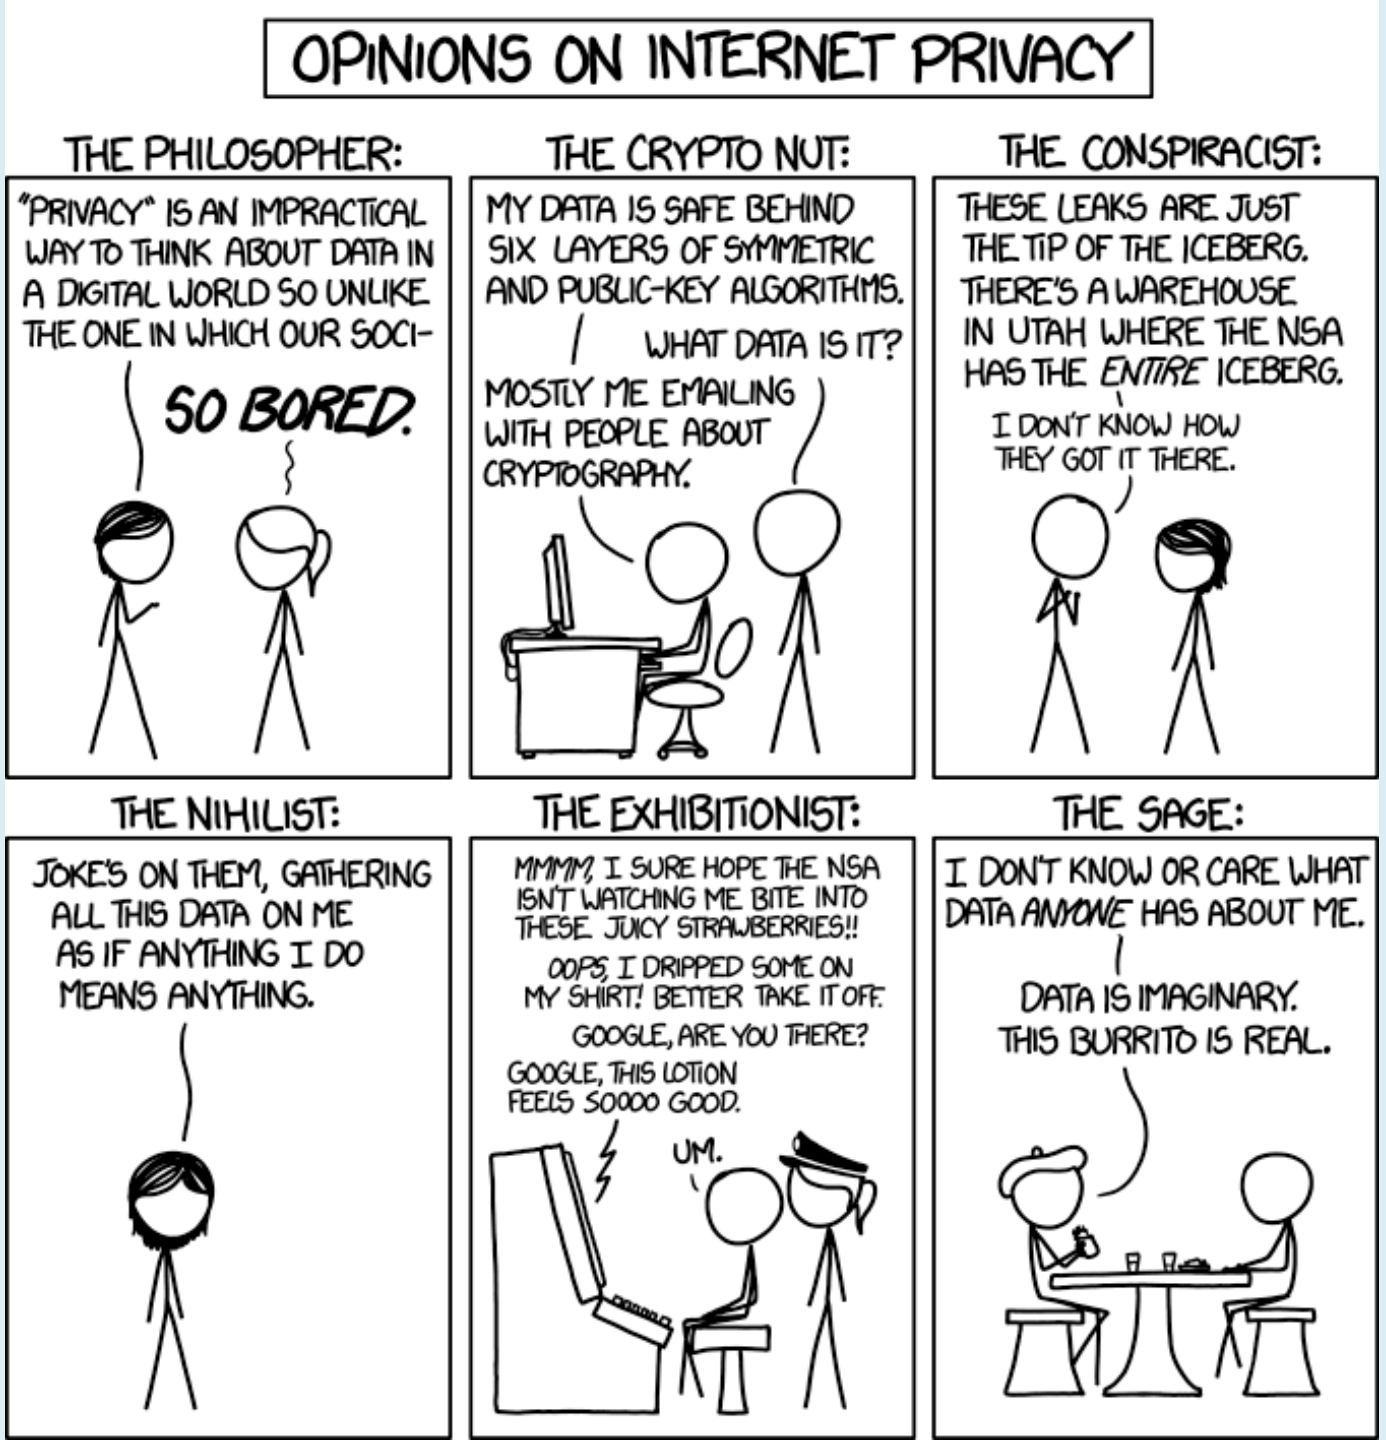

XKCD: https://xkcd.com/1269/

Privacy Definitions (which we will cover today)

- Pseudonymization: Personal information not directly disclosed, still vulnerable to statistical / informed attacks.

- K-Anonymity: Some anonymity guarantees (identifiable with outside information).

- Differential Privacy: “Gold Standard” of anonymity, may still leak group information.

Pseudonymization

- Originating from Greek (pseudonymos) "having a false name, under a false name“

- Under GDPR, it allows for processing of data in conditions not expressly defined in the initial collection. It is also recommended for Privacy by Design.

- Not safe for public release or release to unsecured partners

k-Anonymity

- Attack Model: Identify individuals via attribute combinations (e.g. age, gender, zip code and weight)

- If enough attributes (quasi identifiers) are available, a unique identifier can be constructed even in large data sets.

- Protect individuals by ensuring that for each possible identifier combination there are at least k individuals in the data set.

- Naive method has several drawbacks that can be partially fixed.

Differential Privacy

- Differential Privacy (DP) is not an anonymization method but a way to quantify information leakage in probabilistic querying or data transformation algorithms.

- It is more general than k-anonymity and has a much stricter risk model (as it does not distinguish between sensitive and non-sensitive attributes)

- For simple data types (e.g. binary data) we can easily implement differentially private querying / release mechanism.

Privacy-Preserving ML:

PATE

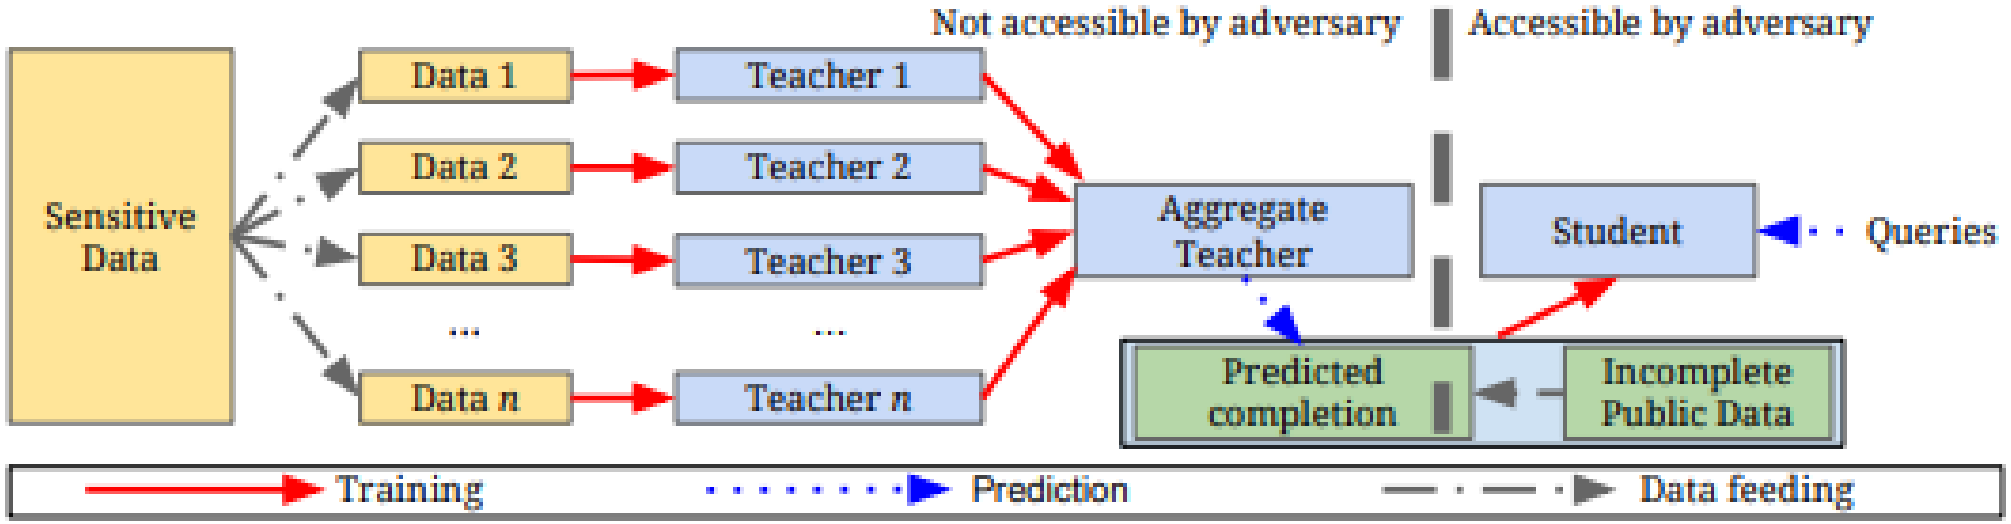

Figure 1: Overview of the approach: (1) an ensemble of teachers is trained on disjoint subsets of the sensitive data, (2) a student model is trained on public data labeled using the ensemble.

Papernot et al (2016):

https://arxiv.org/abs/1610.05755

Privacy-Preserving ML:

CryptoNets

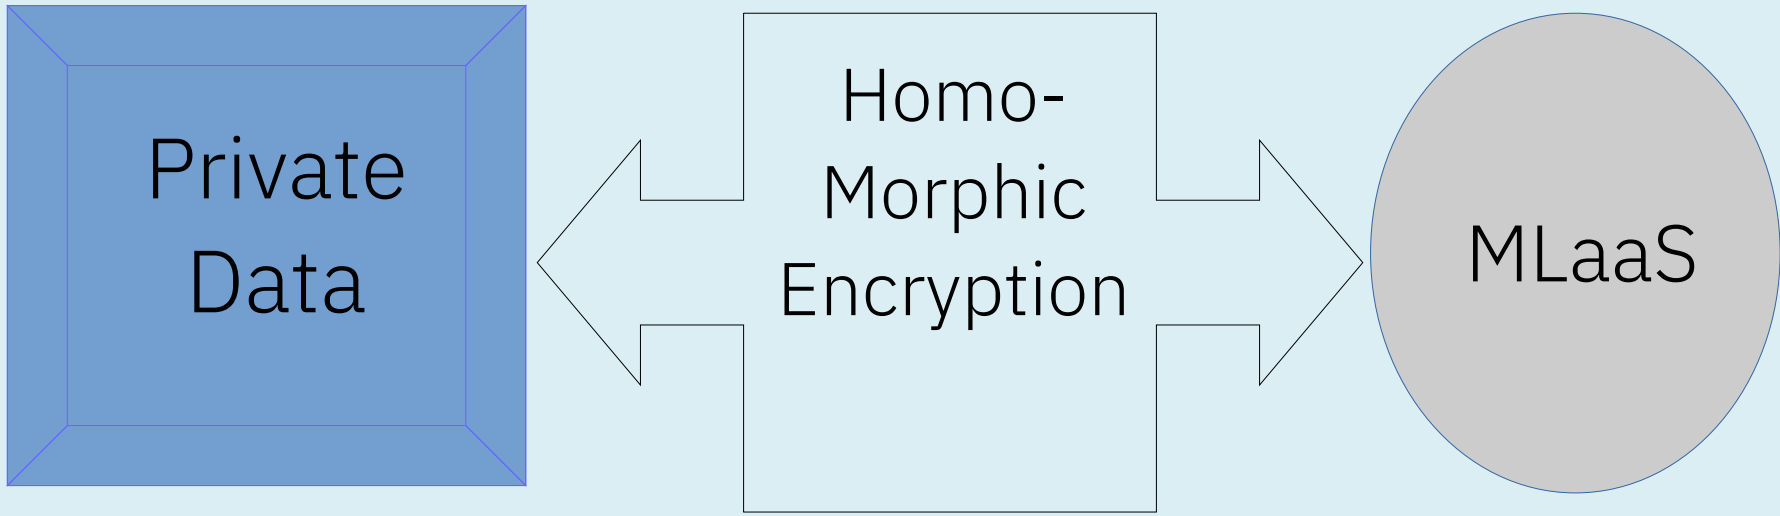

Dowlin et al (2016):
http://proceedings.mlr.press/v48/gilad-bachrach16.pdf

Privacy-Preserving ML

Individual-based fairness is enforced through a „Lipschitz property“: The statistical distance of classifications for two individuals should be bounded by their distance in the input space. Interesting takeway: **Fairness (usually) implies privacy.**

The input space distance is calculated based on a fair metric (that we need to define).

Note: **Implied „affirmative action“ not always legal!**

C. Dwork et. al.: Fairness Through Awareness

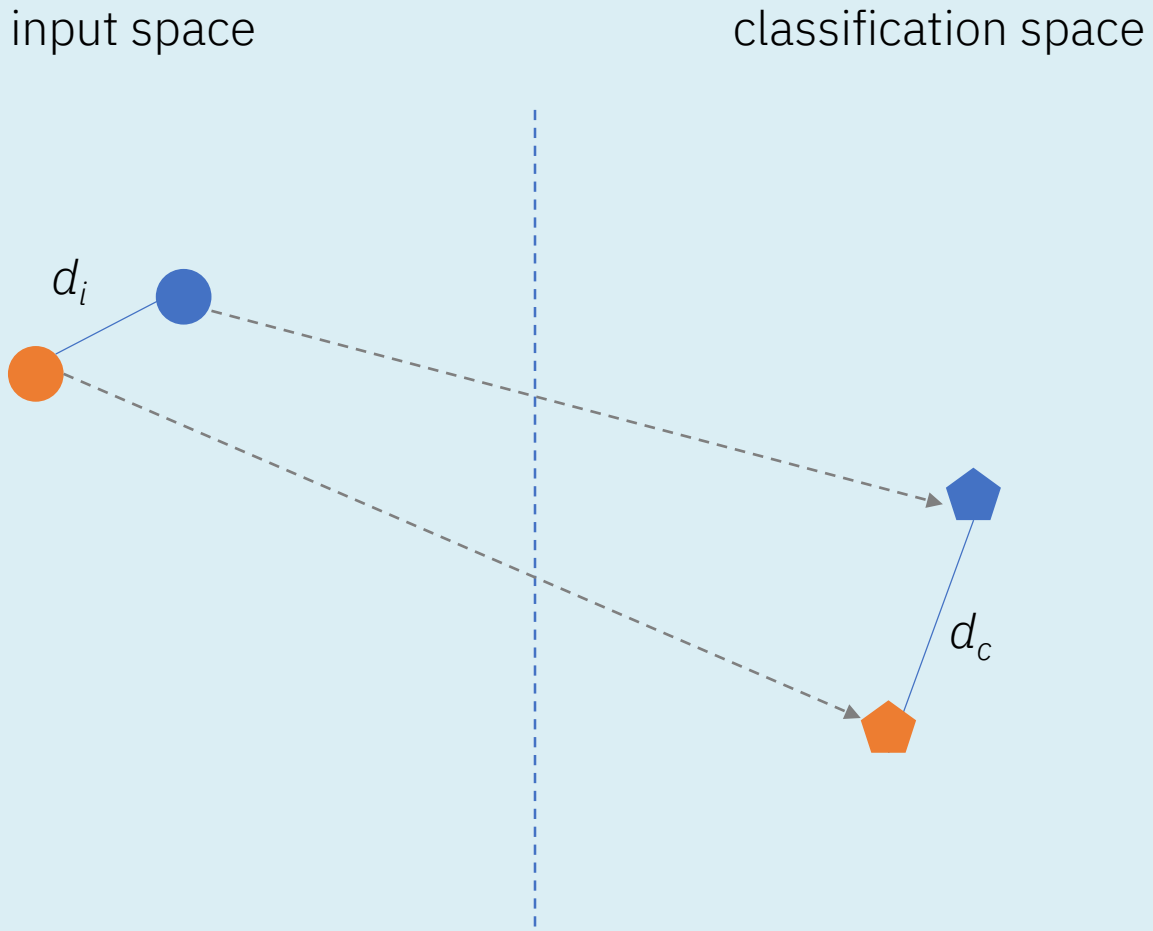

https://arxiv.org/pdf/1104.3913.pdf

If all else fails...

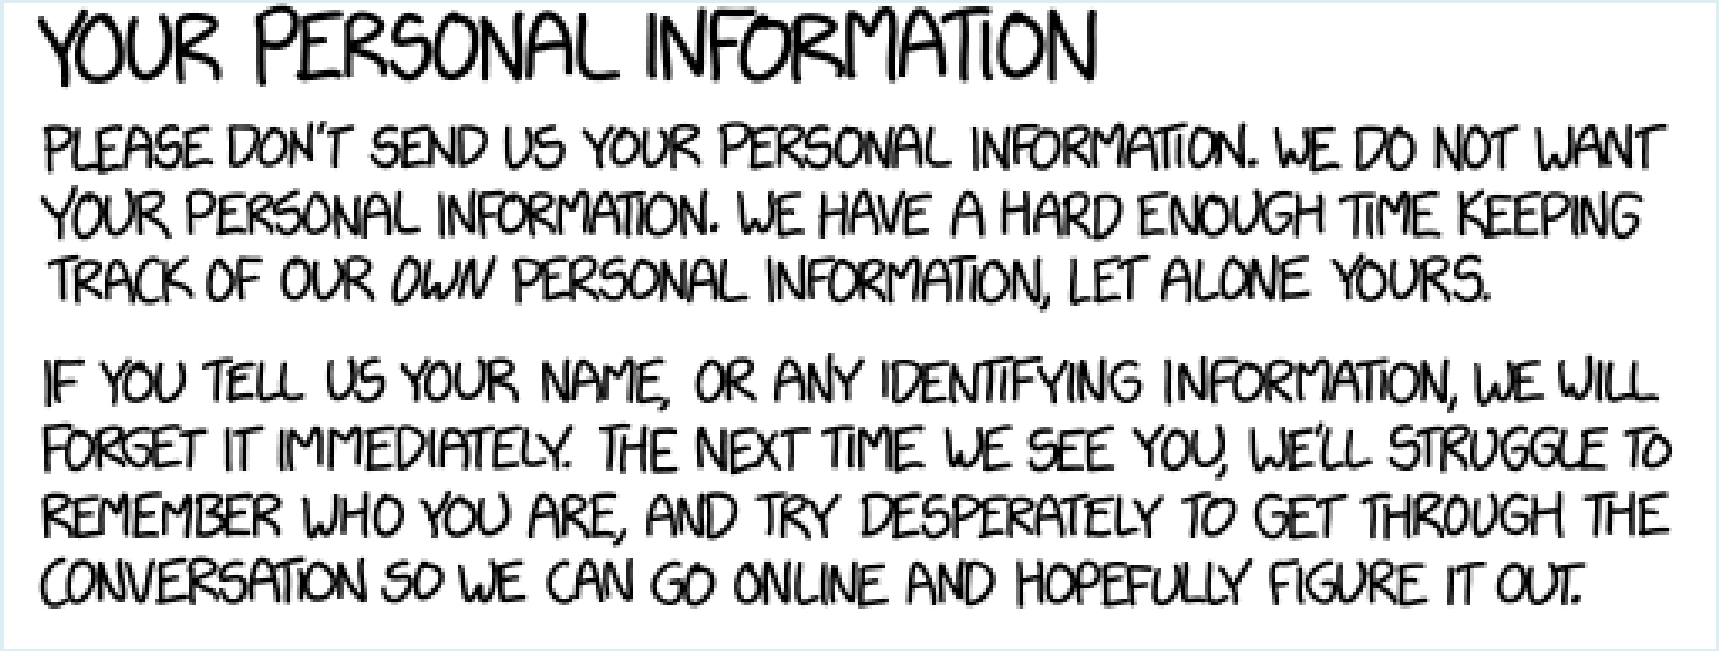

XKCD: https://xkcd.com/1998/

Thank you for attending!

Questions? We'd Love to hear them!

Or reach out anytime:

info@kiprotect.com
@KIProtect (Twitter)
https://github.com/kiprotect

Andreas Dewes
andreas@kiprotect.com
@japh44 (Twitter)

Katharine Jarmul
katharine@kiprotect.com
@kjam (Twitter)

7scientists GmbH
KIProtect
Bismarckstr. 10-12
10625 Berlin