Differential Privacy as a Fairness Intervention

Johannes Kaiser

Technical University of Munich

Munich University Hospital “Rechts der Isar”

Klinikum rechts der Isar
Technische Universität München

Klinikum rechts der Isar
Technische Universität München

Johannes Kaiser

- PhD Student since May 2023
- Affiliated with Technical University of Munich & Hospital Clinic
- Supervised by Daniel Rückert and George Kaissis
- Research Interests: Everything that sparks joy (interactions of ai an humans, and theory)

Klinikum rechts der Isar
Technische Universität München

Warren and Brandeis (1890):
Privacy is a right of individuals to be protected from the unsolicited distribution of information regarding their private life, particularly via publications.

Legally (Non-Formal)

13 U.S.C. §9.
Prohibits any publication whereby the data furnished by. . .[an] individual. . .can be identified

HIPAA Privacy Rule
Permits the disclosure of health information that has been de-identified (removal of information from a list of 18 identifiers)

Critique

Too strict – prohibits the sharing of any aggregate statistics[1]

Too loose – de-identification is known to be faulty[2]

[1] Kifer, Daniel, and Ashwin Machanavajjhala. "No free lunch in data privacy." Proceedings of the 2011 ACM SIGMOD International Conference on Management of data. 2011.
[2] Benitez, Kathleen, and Bradley Malin. "Evaluating re-identification risks with respect to the HIPAA privacy rule." Journal of the American Medical Informatics Association 17.2 (2010): 169-177.

Society separates into three categories with respect to their privacy via self-assessment (Westin Studies 1978 – 2004)

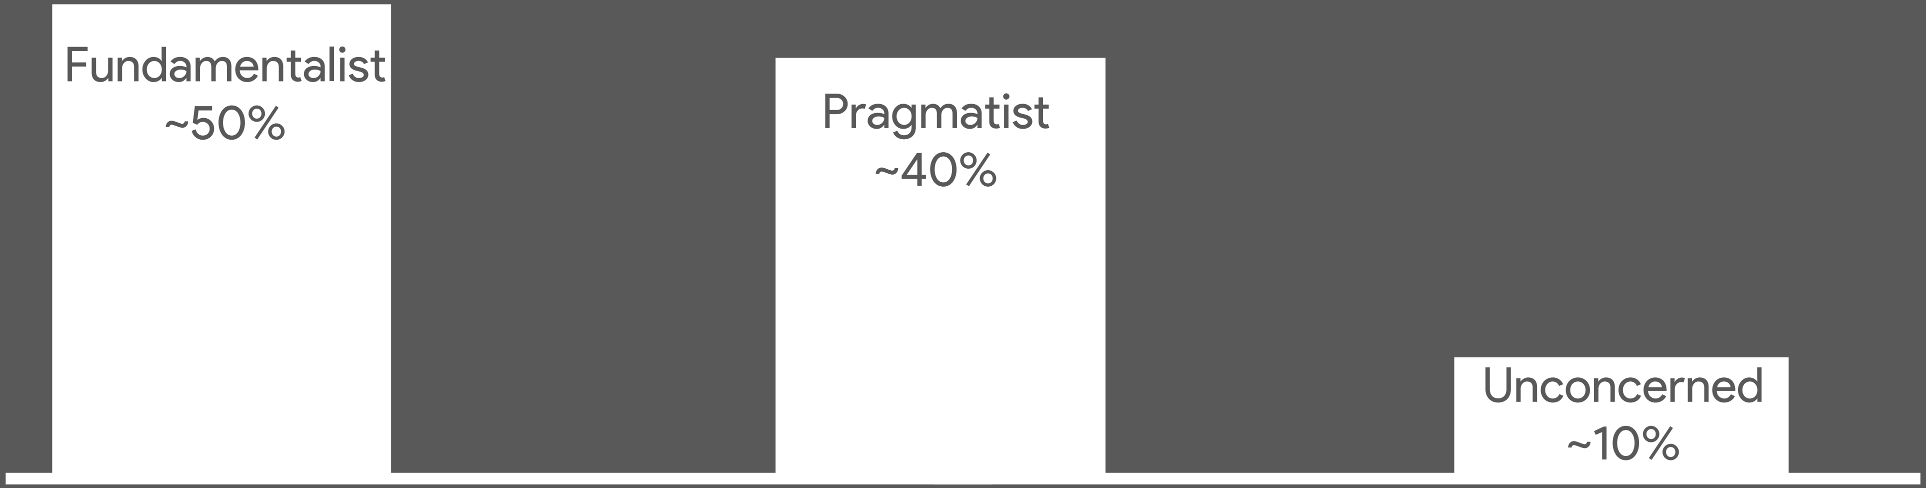

- Protective of their privacy
- Individuals should be proactive
- Support stronger laws

- Weight the pros and cons
- Evaluate protection and trust

- Expect benefits to outweigh risk

P. Kumaraguru and L. F. Cranor, 'Privacy Indexes: A Survey of Westin's Studies'
A. Woodruff, V. Pihur, S. Consolvo, L. Schmidt, L. Brandimarte, and A. Acquisti, 'Would a privacy fundamentalist sell their DNA for $1000... if nothing bad happened as a result? The Westin categories, behavioral intentions, and consequences'.

Klinikum rechts der Isar
Technische Universität München

Introduction
Privacy
ivacy vs. Fairness
DP as a Tool
Individual DP
Results
Ongoing Research
Discussion

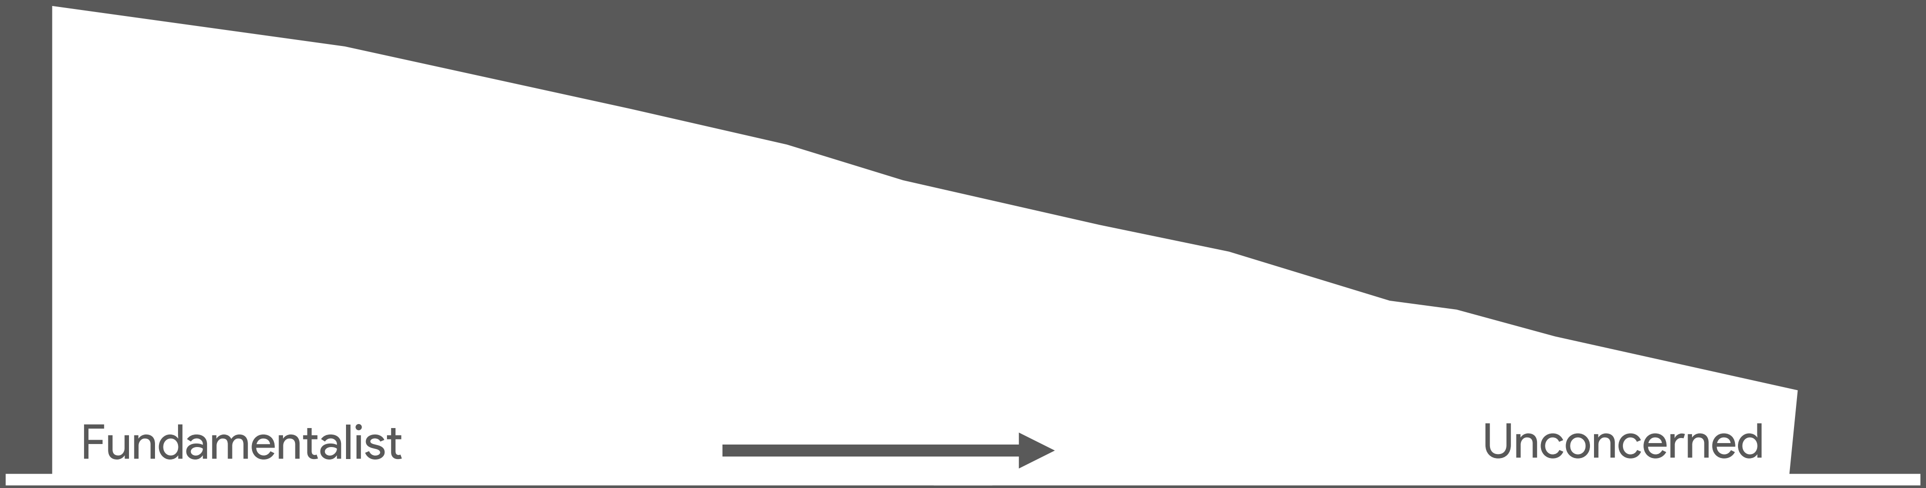

Privacy Segmentation is difficult

- Weigh the potential pros and cons
- Evaluate protection and trust
- Privacy perception is context-specific
- Influenced by cost-benefit, morals, responsibility to share
- Self-assessment fails (knowledge vs. motivation)

Woodruff, V. Pihur, S. Consolvo, L. Schmidt, L. Brandimarte, and A. Acquisti, 'Would a privacy fundamentalist sell their DNA for $1000... if nothing bad happened as a result? The Westin categories, behavioral intentions, and consequences'.
Cynthia E Schairer, Cynthia Cheung, Caryn Kseniya Rubanovich, Mildred Cho, Lorrie Faith Cranor, Cinnamon S Bloss, Disposition toward privacy and information disclosure in the context of emerging health technologies, Journal of the American Medical Informatics Association, Volume 26, Issue 7, July 2019,

Introduction
Privacy
rivacy vs. Fairness
DP as a Tool
Individual DP
Results
Ongoing Research
Discussion

Qualitatively from:
W. M. P. Steijn and A. Vedder, ‘Privacy under Construction: A Developmental Perspective on Privacy Perception’, Science, Technology, & Human Values, vol. 40, no. 4, pp. 615–637, Jul. 2015

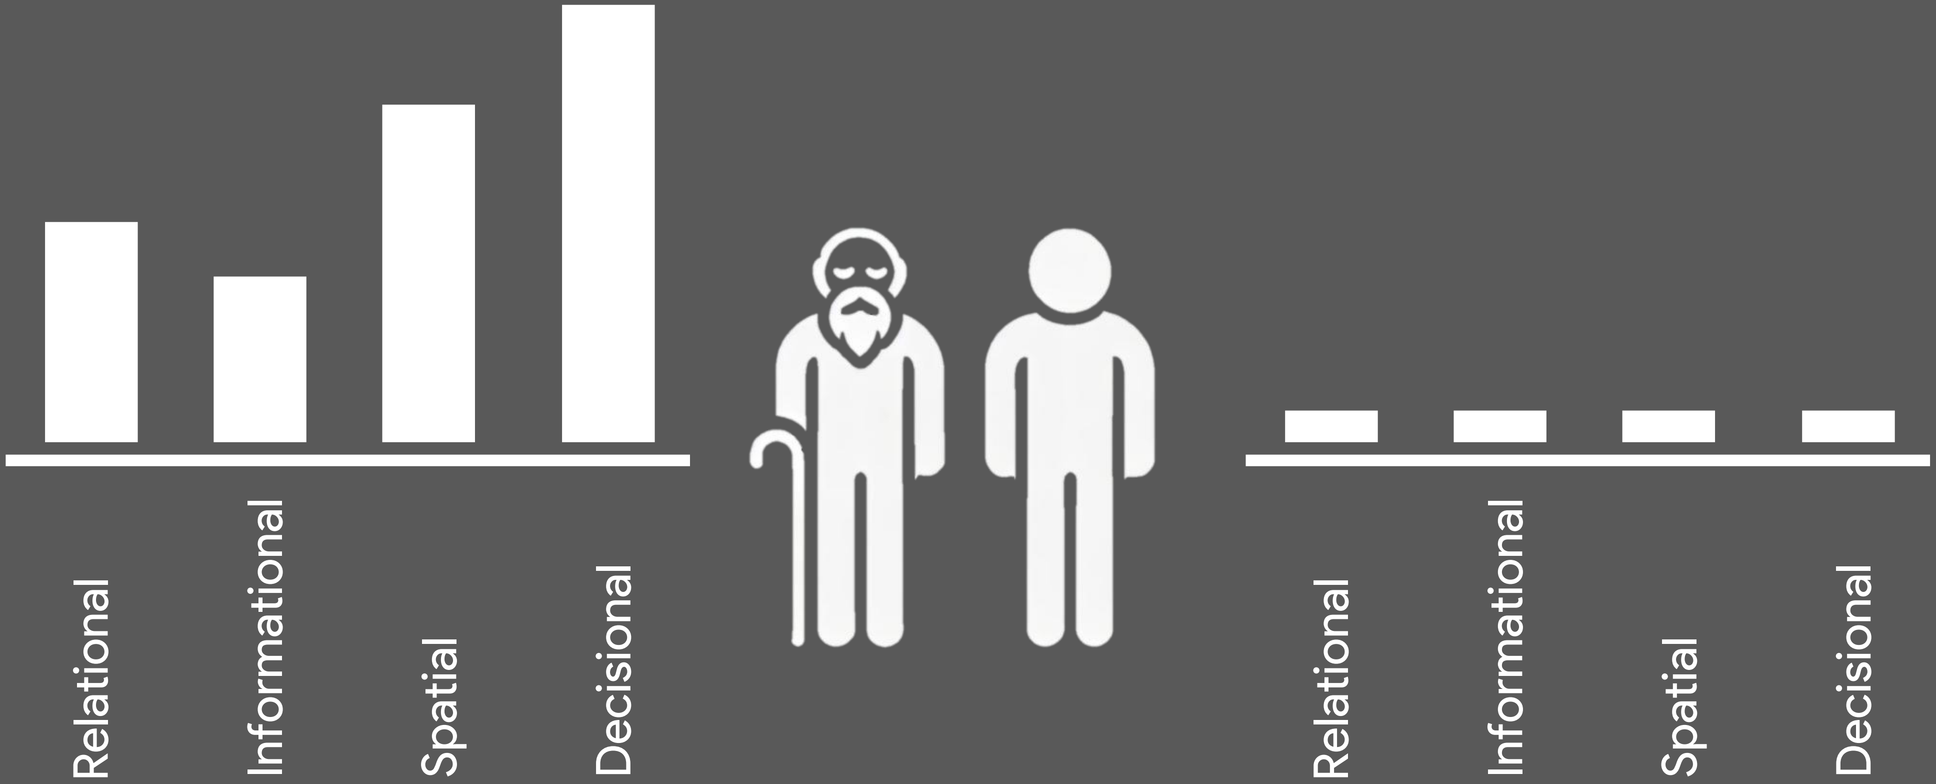

Qualitatively from:

W. M. P. Steijn and A. Vedder, ‘Privacy under Construction: A Developmental Perspective on Privacy Perception’, Science, Technology, & Human Values, vol. 40, no. 4, pp. 615–637, Jul. 2015

Perceived vulnerabilities transform due to:

- technological developments
- changes in socioeconomic conventions and traditions

....

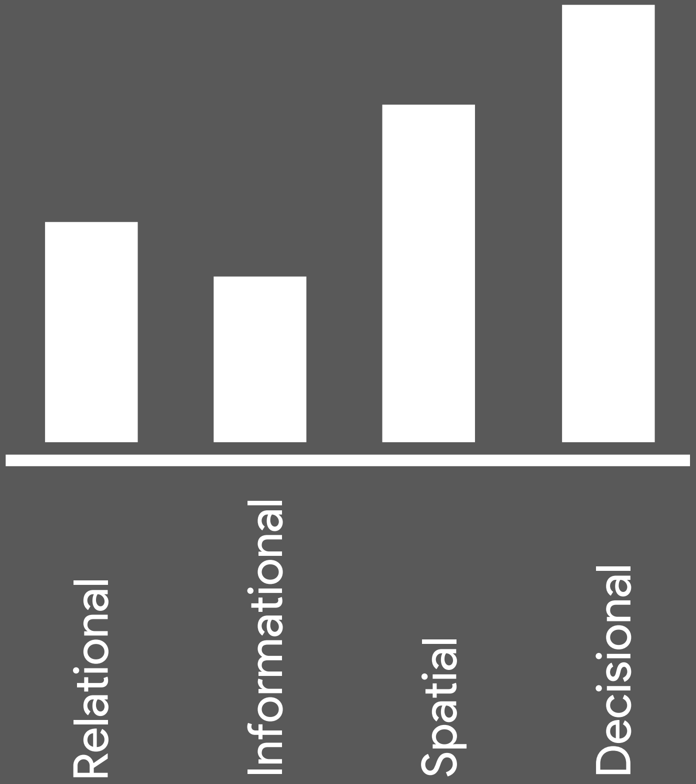

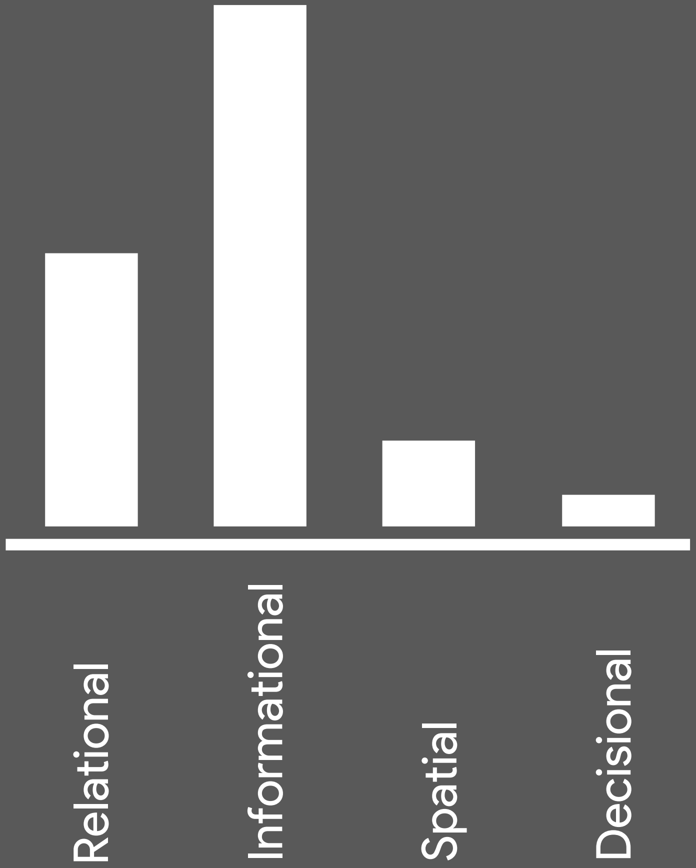

Privacy perception is a personal property
Strongly context dependent

Qualitatively from:
W. M. P. Steijn and A. Vedder, 'Privacy under Construction: A Developmental Perspective on Privacy Perception', Science, Technology, & Human Values, vol. 40, no. 4, pp. 615–637, Jul. 2015

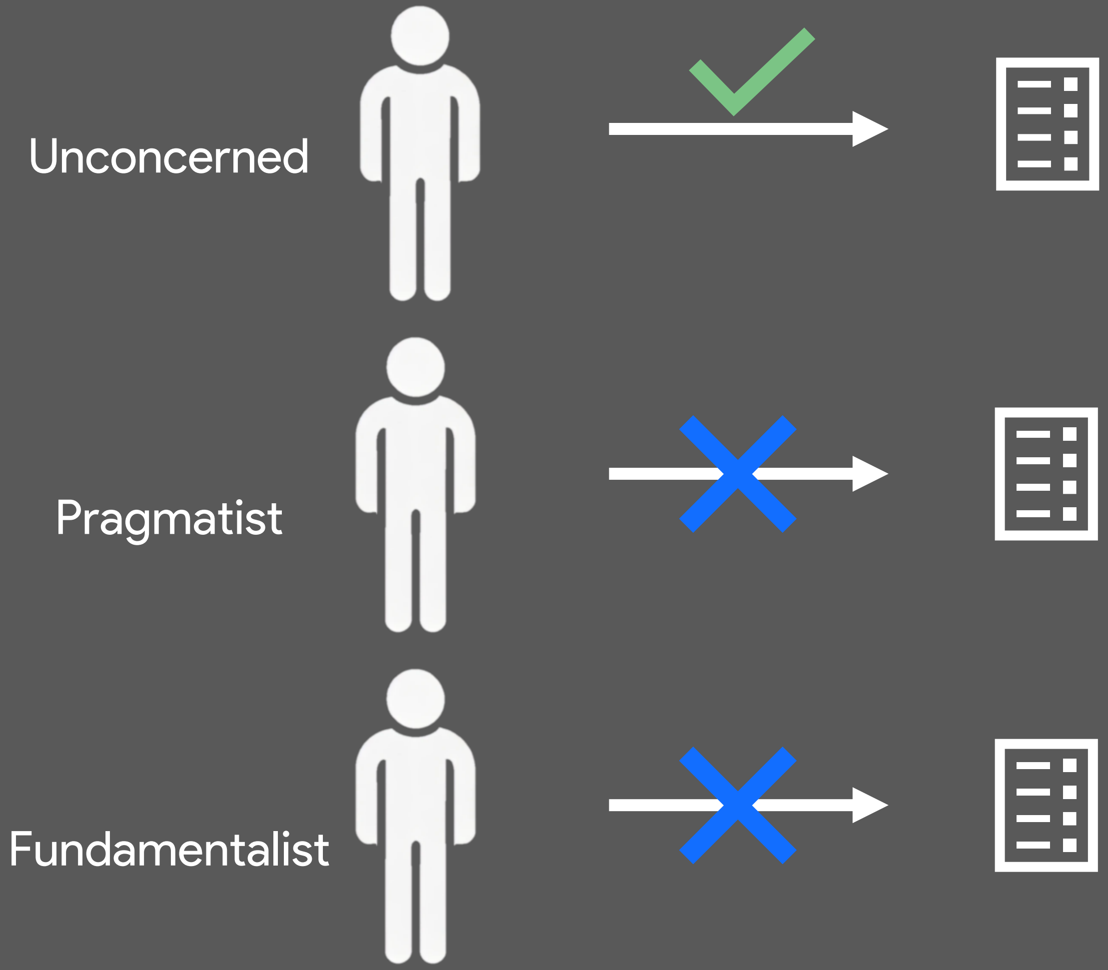

Taylor, Humphrey. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.
Ghosh, Arpita, and Aaron Roth. "Selling privacy at auction." Proceedings of the 12th ACM conference on Electronic commerce. 2011.

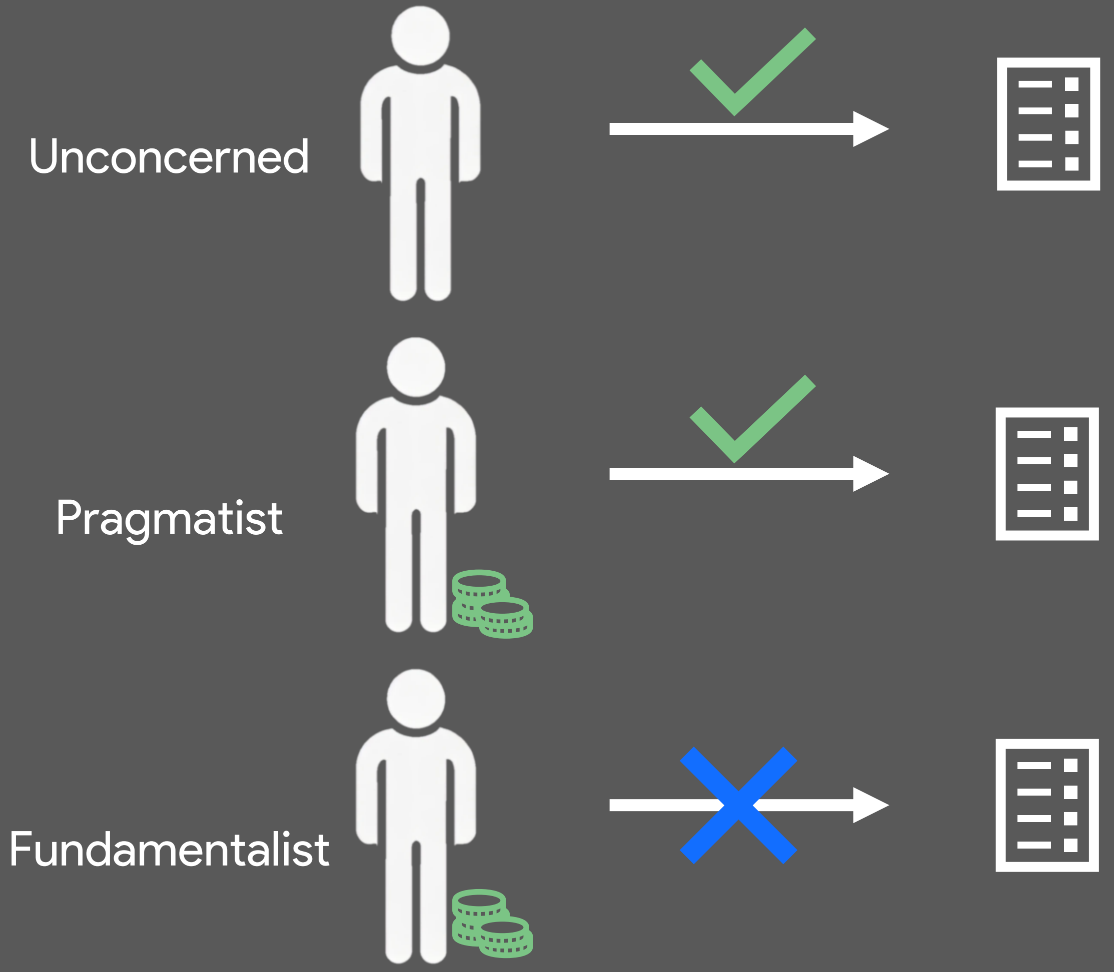

Taylor, Humphrey. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.
Ghosh, Arpita, and Aaron Roth. "Selling privacy at auction." Proceedings of the 12th ACM conference on Electronic commerce. 2011.

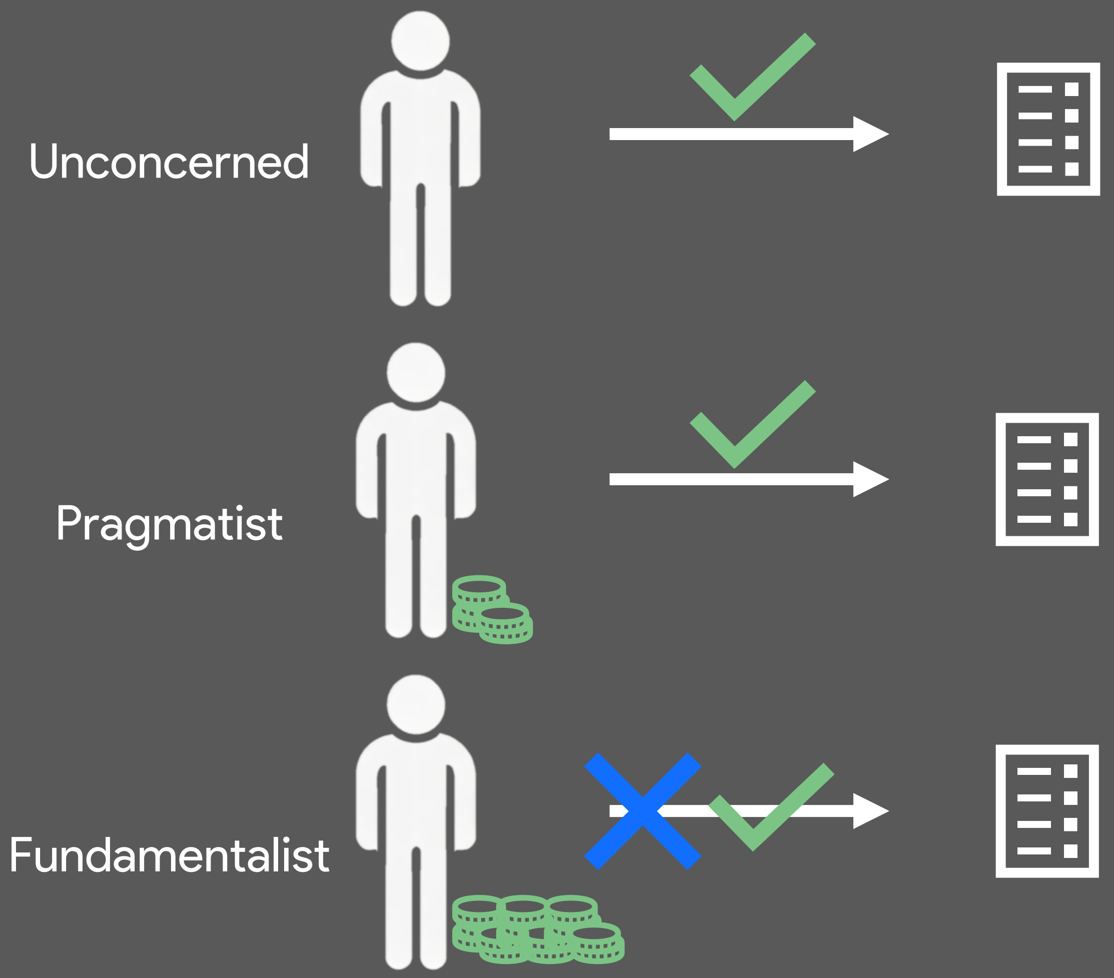

Taylor, Humphrey. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.
Ghosh, Arpita, and Aaron Roth. "Selling privacy at auction." Proceedings of the 12th ACM conference on Electronic commerce. 2011.

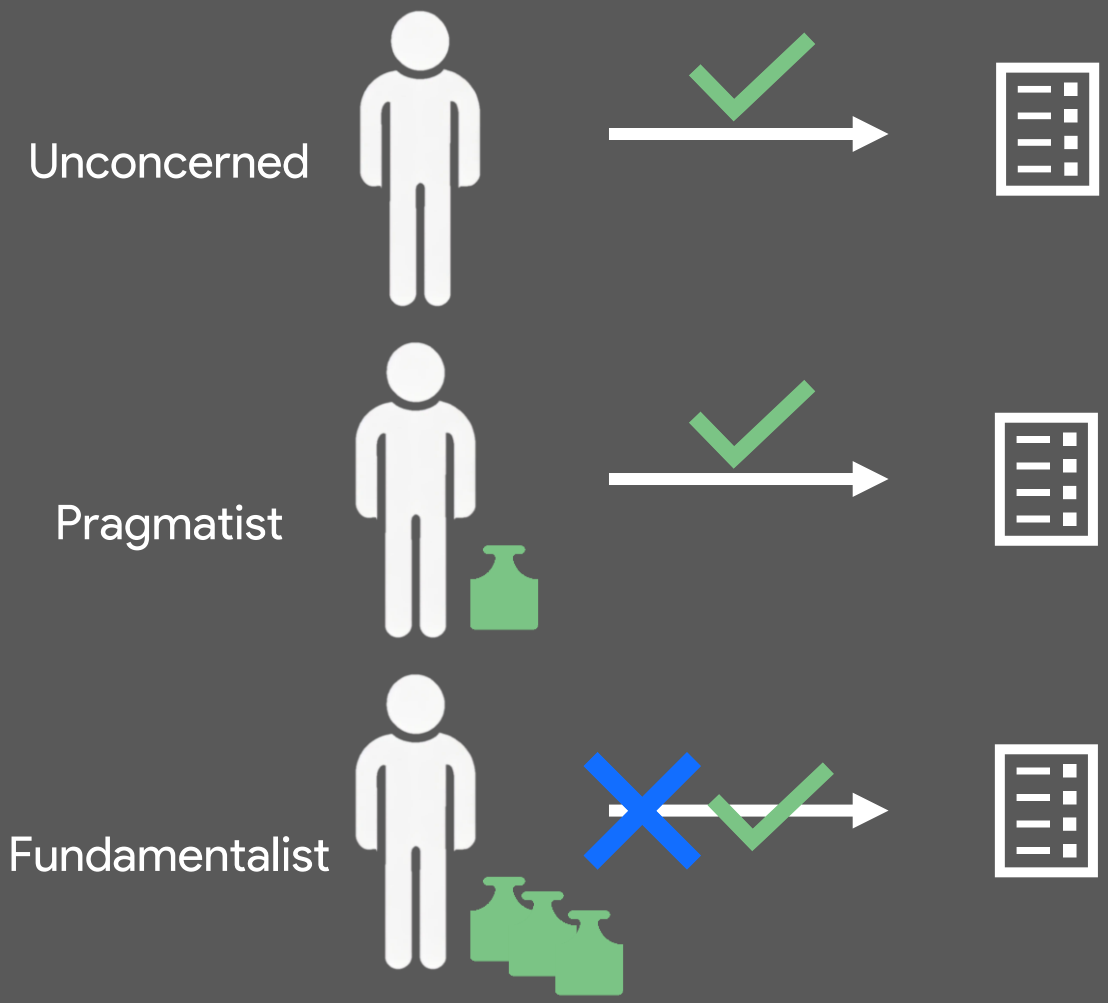

Taylor, Humphrey. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.
Ghosh, Arpita, and Aaron Roth. "Selling privacy at auction." Proceedings of the 12th ACM conference on Electronic commerce. 2011.

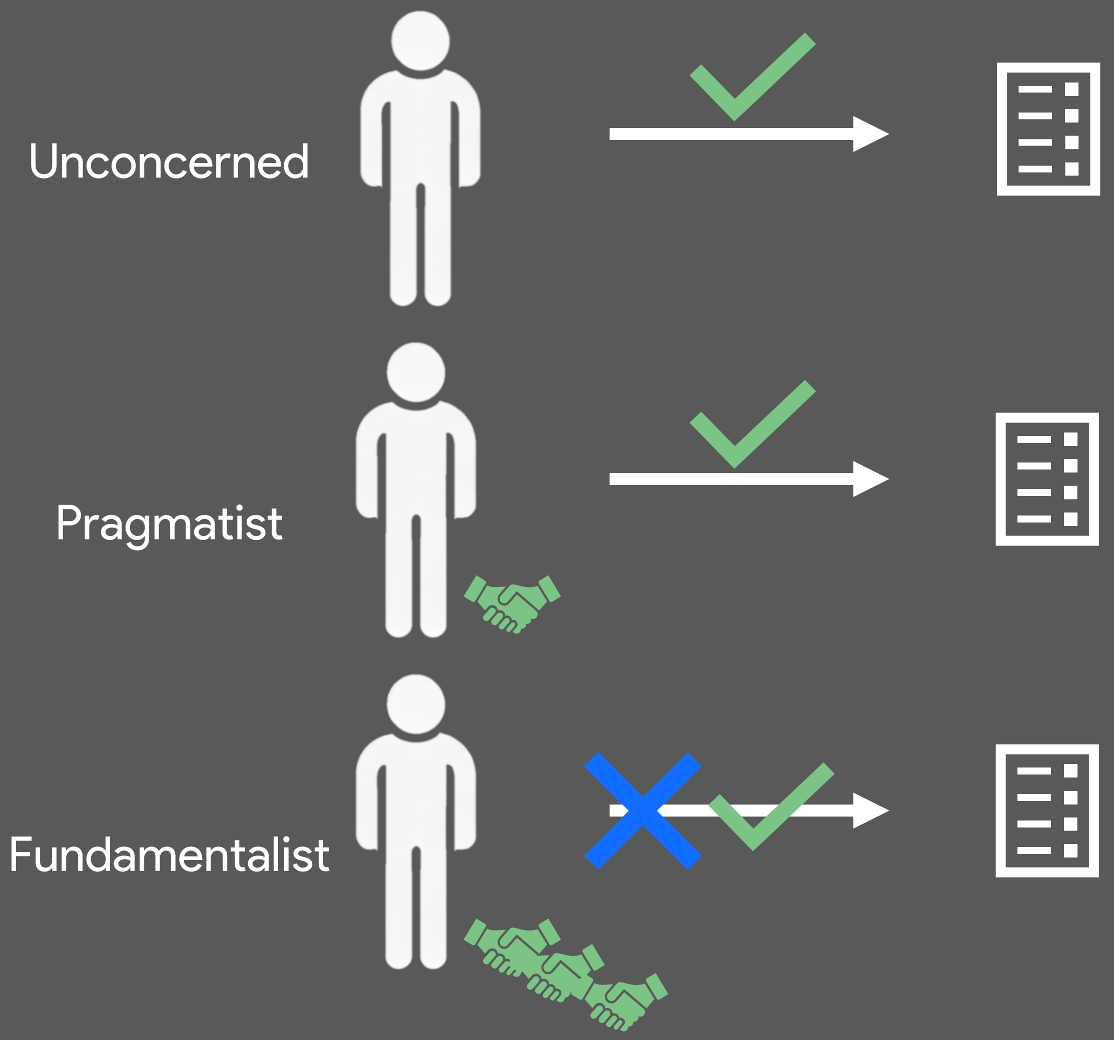

Taylor, Humphrey. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.
Ghosh, Arpita, and Aaron Roth. "Selling privacy at auction." Proceedings of the 12th ACM conference on Electronic commerce. 2011.

Unconcerned

Pragmatist

Fundamentalist

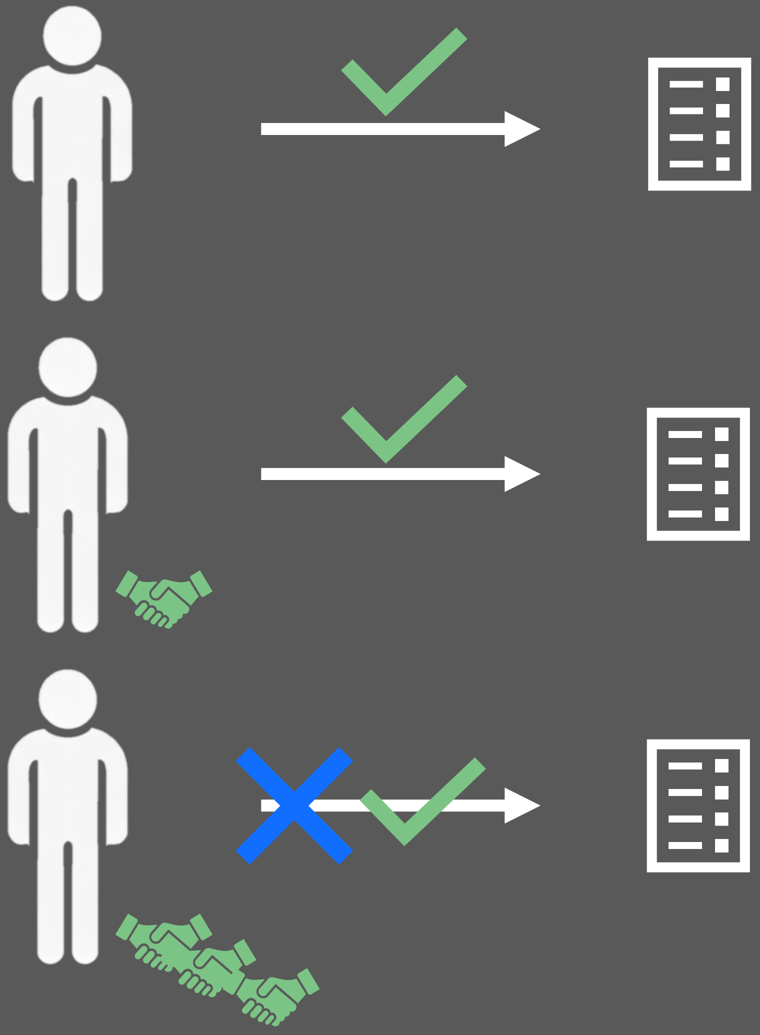

Limited research on “adequate” compensation [1]

Taylor, Humphrey. (2003). Most People Are "Privacy Pragmatists" Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits.
Ghosh, Arpita, and Aaron Roth. "Selling privacy at auction." Proceedings of the 12th ACM conference on Electronic commerce. 2011.
[1] Taylor, David G., Donna F. Davis, and Ravi Jillapalli. "Privacy concern and online personalization: The moderating effects of information control and compensation." Electronic commerce research 9 (2009): 203-223.

Klinikum rechts der Isar
Technische Universität München

- AI systems should be equitable across all demographic groups [1]
  - "Nobody should suffer worse or accuracy in ML solely due to them belonging to a specific group"
- Improving fairness to benefit one group should not hurt any other group [2, 3]
  - "Lowering predictive accuracy for one group because of the presence of another is also a fairness issue."
- Sensitive attributes are essential for many AI applications [4]

[1] Yang, Y., Zhang, H., Gichoya, J.W. *et al.* The limits of fair medical imaging AI in real-world generalization. *Nature Medicine* (2024)
[2] Ghassemi, M., Gusev, A. Limiting bias in AI models for improved and equitable cancer care. *Nature Reviews Cancer* (2024)
[3] Suriyakumar, Vinith M., Marzyeh Ghassemi, and Berk Ustun. "When personalization harms: Reconsidering the use of group attributes in prediction." arXiv preprint arXiv:2206.02058 (2022)
[4] Taylor S, Jaques N, Nosakhare E, Sano A, Picard R. Personalized Multitask Learning for Predicting Tomorrow's Mood, Stress, and Health. IEEE Trans Affect Comput. (2020)

Privacy

Privacy in machine learning refers to the protection of individuals' sensitive data during the training and deployment of models, ensuring that personal information is not exposed or inferred from the model's outputs.

Fairness

Fairness in machine learning refers to the design and deployment of models that ensure equitable treatment of all individuals or groups, avoiding biases and discrimination in predictions or outcomes.

Privacy

Identifier	Attr. 1	Attr. 2	Attr. 3	Sens. Attr.
0001	105	12	38	Black
0002	98	8	42	Blonde
0003	93	10	36	Red
0004	112	14	29	Blonde
0005	102	6	30	Brown
0006	96	11	31	Blonde

≈

Identifier	Attr. 1	Attr. 2	Attr. 3	Sens. Atr.
0001	105	12	38	Black
[redacted]				
0003	93	10	36	Red
0004	112	14	29	Blonde
0005	102	6	30	Brown
0006	96	11	31	Blonde

Fairness

Identifier	Attr. 1	Attr. 2	Attr. 3	Sens. Attr.
0001	105	12	38	Black
0002	98	8	42	Blonde
0003	93	10	36	Red
0004	112	14	29	Blonde
0005	102	6	30	Brown
0006	96	11	31	Blonde

≈

Identifier	Attr. 1	Attr. 2	Attr. 3	Sens. Attr.
0001	105	12	38	[redacted]
0002	98	8	42	
0003	93	10	36	
0004	112	14	29	
0005	102	6	30	
0006	96	11	31	

Privacy

Fairness

- Privacy concerns the training data
- Privacy concerns a limited group of people
- Privacy is a data usage property
- Privacy requirement is a personal property (it can be compensated for)

- Fairness concerns the output data
- Fairness concerns an unlimited group of people
- Fairness is a model property
- Fairness requirement is a societal property

Privacy

Differential Privacy

$$P(\theta(D_1) \in O) \leq e^{\epsilon} \cdot P(\theta(D_2) \in O) \; with \; D_1 \cong D_2, \forall O$$

$$L_{D_1 D_2}(O) = \ln\left(\frac{P(\theta(D_1) \in O)}{P(\theta(D_2) \in O)}\right) \; with \; D_1 \cong D_2, \forall O$$

Fairness

Disparate Impact

$$\phi_{DI} = \frac{P(\theta(x) = y \,|\, g(x) = a)}{P(\theta(x) = y \,|\, g(x) = b)} \; \forall a, b$$

Klinikum rechts der Isar
Technische Universität München

Decision Making with Differential Privacy under a Fairness Lens

Cuong Tran[1], **Ferdinando Fioretto**[1], **Pascal Van Hentenryck**[2] and **Zhiyan Yao**[3*]
[1]Syracuse University
[2]Georgia Institute of Technology
[3]Nanjing University of Science and Technology
{cutran, ffiorett}@syr.edu, pvh@isye.gatech.edu, zyao09@syr.edu

Abstract

Many agencies release datasets and statistics about groups of individuals that are used as input to a number of critical decision processes. To conform with privacy and confidentiality requirements, these agencies are often required to release privacy-preserving versions of the data. This paper studies the release of differentially private datasets and analyzes their impact on some critical resource allocation tasks under a fairness perspective. The paper shows that, when the decisions take as input differentially private data, the noise added to achieve privacy disproportionately impacts some groups over others. The paper analyzes the reasons for these disproportionate impacts and proposes guidelines

Although DP provides strong privacy guarantees, *it may induce biases and fairness issues in downstream decision processes*, as shown empirically in [Pujol *et al.*, 2020]. Since at least $675 billion are being allocated based on U.S. census data, the use of differential privacy without a proper understanding of these biases and fairness issues may adversely affect the health, well-being, and sense of belonging of many individuals. Indeed, the allotment of federal funds, apportionment of congressional seats, and distribution of vaccines should ideally be fair and unbiased. Similar issues arise in several other areas including election, energy, and food policies. The problem is further exacerbated by the recent recognition that *commonly adopted DP mechanisms for data release may introduce unexpected biases on their own, independently of a downstream decision process* [Zhu et al., 2021].

This paper builds on these observations and provides a step

DP introduces substantial bias in allotment problems due to the stronger perturbation of smaller values than larger values due to the noise addition

Introduction

Privacy

Privacy vs. Fairness

DP as a Tool

Individual DP

Results

Ongoing Research

Discussion

Decision Making with Differential Privacy under a Fairness Lens

Cuong Tran[1], Ferdinando Fioretto[1], Pascal Van Hentenryck[2] and Zhiyan Yao[3*]

[1]Syracuse University

Trade-Offs between Fairness and Privacy in Machine Learning

Sushant Agarwal
University of Waterloo, Canada
sushant.agarwal@uwaterloo.ca

Abstract

The concerns of fairness, and privacy, in machine learning based systems have received a lot of attention in the research community recently, but have primarily been studied in isolation. In this work, we look at cases where we want to satisfy both these properties simultaneously, and find that it may be necessary to make trade-offs between them. We

properties simultaneously, and analyse how they interact. We find that that these properties are at odds with each other, and it is necessary to make trade-offs between them. We show a theoretical result to demonstrate this, which talks about the clash between the requirements of differential privacy, accuracy, and fairness in learning algorithms. It is an impossibility theorem which states that even in a very simple binary classification setting, no learning algorithm that is ϵ-differentially private (for any $\epsilon < \infty$), and approximately fair (i.e., the

Theoretical proof, that pure DP and approximate fairness cannot achieve accuracy better than a constant classifier

Introduction
Privacy
Privacy vs. Fairnes
DP as a Tool
Individual DP
Results
ngoing Research
Discussion

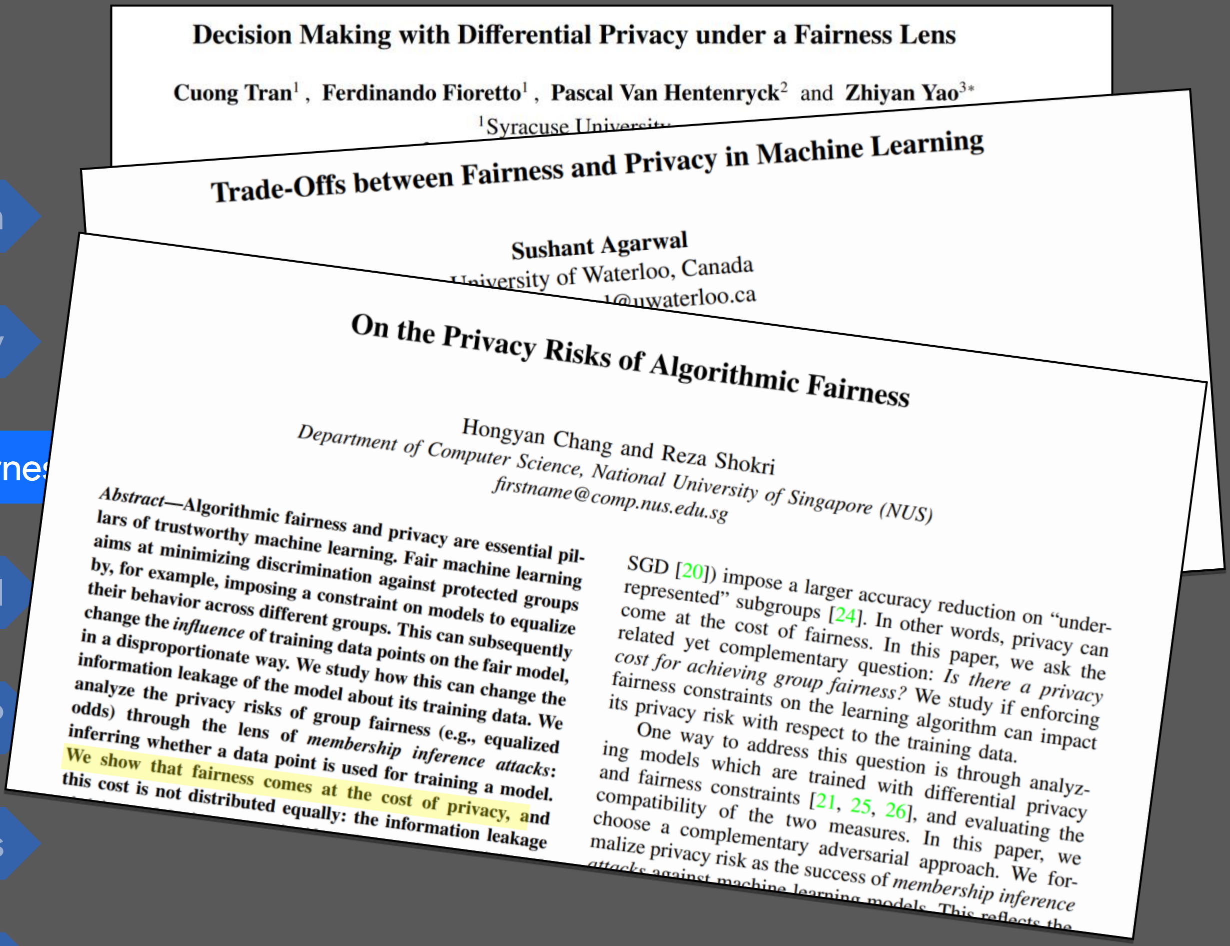

Decision Making with Differential Privacy under a Fairness Lens

Cuong Tran[1], Ferdinando Fioretto[1], Pascal Van Hentenryck[2] and Zhiyan Yao[3]*

Trade-Offs between Fairness and Privacy in Machine Learning

Sushant Agarwal
University of Waterloo, Canada

On the Privacy Risks of Algorithmic Fairness

Hongyan Chang and Reza Shokri
Department of Computer Science, National University of Singapore (NUS)
firstname@comp.nus.edu.sg

***Abstract*—Algorithmic fairness and privacy are essential pillars of trustworthy machine learning. Fair machine learning aims at minimizing discrimination against protected groups by, for example, imposing a constraint on models to equalize their behavior across different groups. This can subsequently change the *influence* of training data points on the fair model, in a disproportionate way. We study how this can change the information leakage of the model about its training data. We analyze the privacy risks of group fairness (e.g., equalized odds) through the lens of *membership inference attacks*: inferring whether a data point is used for training a model. We show that fairness comes at the cost of privacy, and this cost is not distributed equally: the information leakage**

SGD [20]) impose a larger accuracy reduction on "under-represented" subgroups [24]. In other words, privacy can come at the cost of fairness. In this paper, we ask the related yet complementary question: *Is there a privacy cost for achieving group fairness?* We study if enforcing fairness constraints on the learning algorithm can impact its privacy risk with respect to the training data.

One way to address this question is through analyzing models which are trained with differential privacy and fairness constraints [21, 25, 26], and evaluating the compatibility of the two measures. In this paper, we choose a complementary adversarial approach. We formalize privacy risk as the success of *membership inference*

Empirically show, that data of fairer models is more susceptible to MIA

Klinikum rechts der Isar
Technische Universität München

- Introduction
- Privacy
- Privacy vs. Fairnes
- DP as a Tool
- Individual DP
- Results
- Ongoing Research
- Discussion

Decision Making with Differential Privacy under a Fairness Lens

Cuong Tran[1], Ferdinando Fioretto[1], Pascal Van Hentenryck[2] and Zhiyan Yao[3*]
[1]Syracuse Universit

Trade-Offs between Fairness and Privacy in Machine Learning

Sushant Agarwal
University of Waterloo, Canada
@uwaterloo.ca

On the Privacy Risks of Algorithmic Fairness

Hongyan Chang and Reza Shokri
Department of Computer Science, National University of Singapore (NUS)
firstname@comp.nus.edu.sg

Abstract—Algorithmic fairness and privacy ar
lars of trustworthy machine le
aims at minimi

Differentially Private Fair Learning

Matthew Jagielski[1] Michael Kearns[2] Jieming Mao[2] Alina Oprea[1] Aaron Roth[2] Saeed Sharifi-Malvajerdi[2] Jonathan Ullman[1]

Abstract

Motivated by settings in which predictive models may be required to be non-discriminatory with respect to certain attributes (such as race), but even collecting the sensitive attribute may be forbidden or restricted, we initiate the study of fair learning under the constraint of differential privacy. Our first algorithm is a private implementation of the equalized odds post-processing approach of (Hardt et al., 2016). This algorithm is appealingly simple, but must be able to use protected group membership explicitly at test time, which can be viewed as a form of "disparate treatment". Our second algorithm is a differentially private version of the oracle-efficient in-processing approach of (Agarwal et al., 2018) which is more complex but need not have access to protected group membership at test time. We identify new tradeoffs between fairness, accuracy, and privacy that emerge only when requiring all three properties, and show that these tradeoffs can be milder if group membership may be used at test time. We conclude with a brief experimental evaluation.

regulations often restrict the use of "sensitive" or protected attributes in algorithmic decision-making. U.S. law prevents the use of race in the development or deployment of consumer lending or credit scoring models, and recent provisions in the E.U. General Data Protection Regulation (GDPR) restrict or prevent even the collection of racial data for consumers. These two developments — the demand for non-discriminatory algorithms and models on the one hand, and the restriction on the collection or use of protected attributes on the other — present technical conundrums, since the most straightforward methods for ensuring fairness generally require knowing or using the attribute being protected. It seems difficult to guarantee that a trained model is not discriminating against (say) a racial group if we cannot even identify members of that group in the data.

A recent line of work (Veale & Binns, 2017; Kilbertus et al., 2018) made these cogent observations, and proposed an interesting solution employing the cryptographic tool of *secure multiparty computation* (commonly abbreviated *MPC*). In this model, we imagine a commercial entity with access to consumer data that excludes race, but this entity would like to build a predictive model for, say, commer-

Empirically and theoretically show privacy-fairness-utility tradeoff for DP fairness postprocessing & DP oracle learner on tabular data

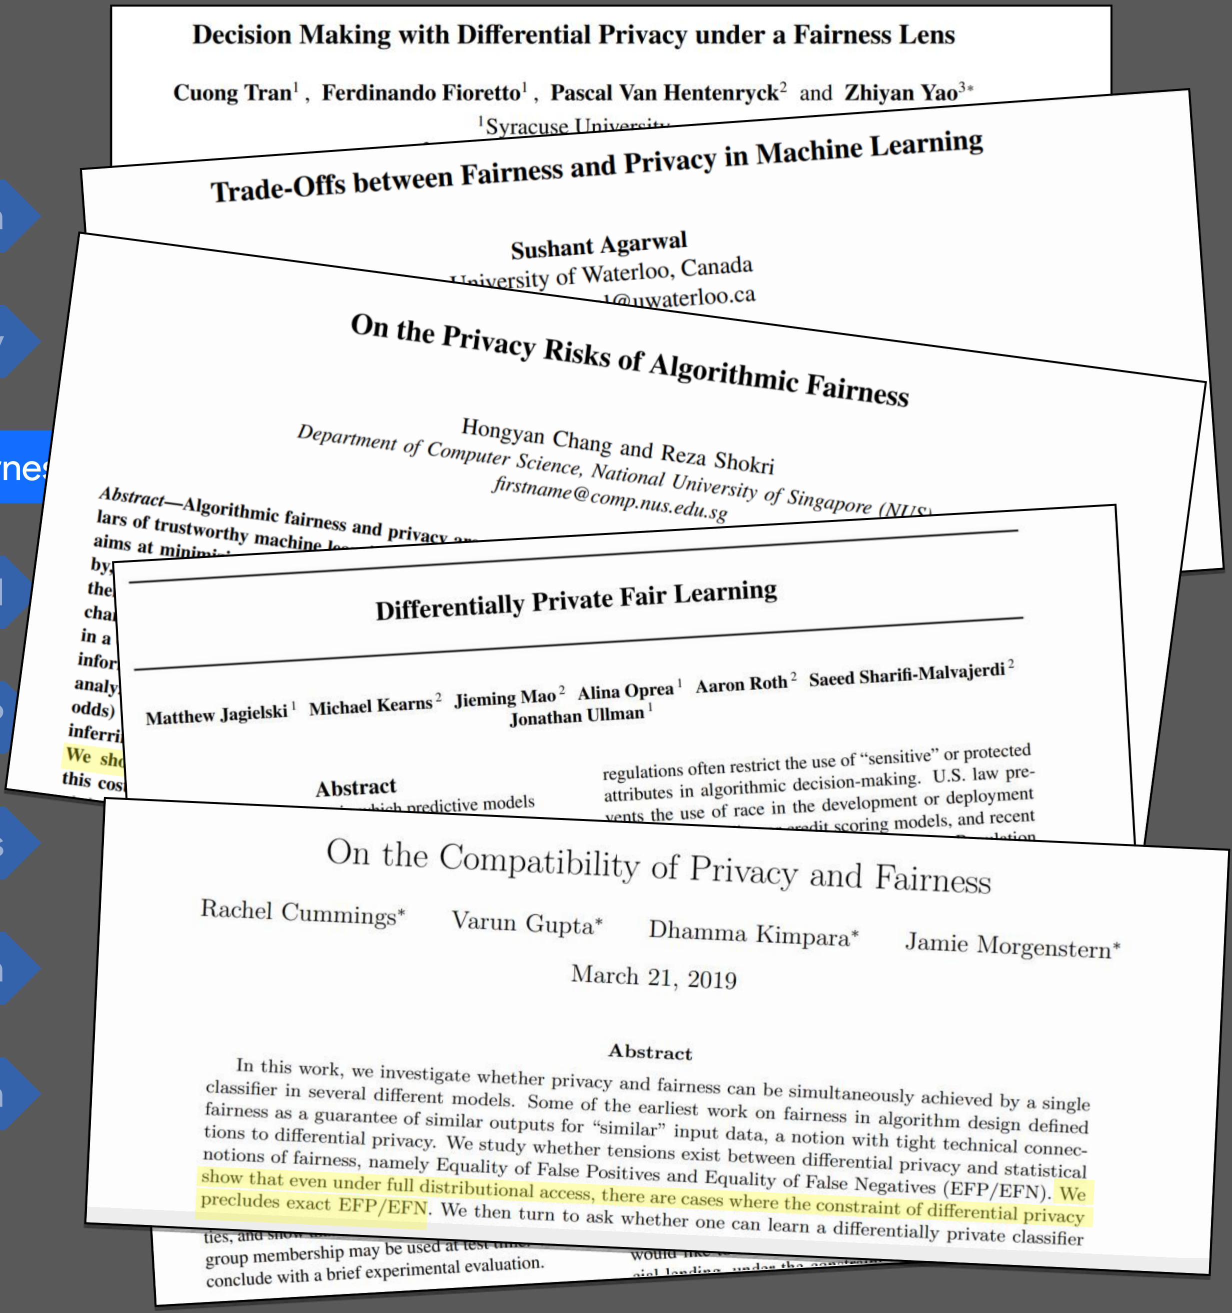

Theoretical proof, that pure DP and exact fairness cannot achieve accuracy better than a constant classifier

ε = 5
Accuracy = 95%

ε = 5
Accuracy = 93%

ε = 5
Accuracy = 80%

$\varepsilon = 5$
Accuracy = 95%

$\varepsilon = 5$
Accuracy = 93%

$\varepsilon = 5$
Accuracy = 80%

Due to:
- Higher variance in data

Klinikum rechts der Isar
Technische Universität München

ε = 5
Accuracy = 95%

ε = 5
Accuracy = 93%

ε = 5
Accuracy = 80%

Due to:
- Higher variance in data
- Poor data quality

$\varepsilon = 5$
Accuracy = 95%

$\varepsilon = 5$
Accuracy = 93%

$\varepsilon = 5$
Accuracy = 80%

Due to:
- Higher variance in data
- Poor data quality
- Lack of data

$\varepsilon = 5$
Accuracy = 95%

$\varepsilon = 5$
Accuracy = 93%

$\varepsilon = 5$
Accuracy = 80%

Due to:
- Higher variance in data
- Poor data quality
- Lack of data
- Unknown reasons

ε = 5
Accuracy = 95%

ε = 5
Accuracy = 93%

ε = 5
Accuracy = 80%

- Tune Privacy Budget w.r.t target accuracy

ε = 8
Accuracy = 92%

- Compensate for Additional Privacy Loss

Klinikum rechts der Isar
Technische Universität München

Introduction
Privacy
ivacy vs. Fairness
DP as a Tool
Individual DP
Results
Ongoing Research
Discussion

Individual Differentially Private SGD

Accounting

Individual Odometer

Individual Filters

Dropping data that exceeds the privacy budget

Assignment

Individual sensitivity (clipping bounds)

Individual sampling rates

Individual Differentially Private SGD

Accounting

Individual Odometer

Individual Filters

Dropping data that exceeds the privacy budget

Suffers from catastrophic forgetting
Usually drops most important data first

Assignment

Individual sensitivity (clipping bounds)

Individual sampling rates

Klinikum rechts der Isar
Technische Universität München

Individual Differentially Private SGD

Accounting

Individual Odometer

Individual Filters

Dropping data that exceeds the privacy budget

Suffers from catastrophic forgetting
Usually drops most important data first

Assignment

Individual sensitivity (clipping bounds)

Individual sampling rates

Sampling rates outperform clipping

Find q_p such that $\epsilon_p \leq \underbrace{I \cdot 2q_p^2 \frac{\alpha}{\sigma_{sample}}}$

Ensures, that the privacy budget is spent after I iterations

With q_p such that $\underbrace{\frac{1}{N}\sum_{p=1}^{P} |G_p| q_p} = q = \frac{B}{N}$ } Desired batch size

Sum over expected number of samples per group

Boenisch, Franziska, et al. "Have it your way: Individualized Privacy Assignment for DP-SGD." *Advances in Neural Information Processing Systems* 36 (2024).

Finding q_p

Require: Per-group target privacy budgets $\{\epsilon_1, \ldots, \epsilon_p\}$, target δ, Iterations I, number of total data points N, per-privacy group number of data points $\{|G_1|, \ldots, |G_p|\}$.

Init σ_{sample}: $\sigma_{sample} \leftarrow getNoise(\epsilon_1, \delta, q, I)$
Init $\{q_1, \ldots, q_p\}$ where for $p \in [P]$
$q_p \leftarrow getSampleRate(\epsilon_p, \delta, \sigma_{sample}, I)$

While $q \not\approx \frac{1}{N}\sum_{p=1}^{P} |G_p| q_p$:
  $\sigma_{sample} \leftarrow s_i \sigma_{sample}$ with $s_i < 1$
  $q_p \leftarrow getSampleRate(\epsilon_p, \delta, \sigma_{sample}, I)\ \forall\, p \in [P]$
Output: $\sigma_{sample}, \{q_1, \ldots, q_p\}$

Boenisch, Franziska, et al. "Have it your way: Individualized Privacy Assignment for DP-SGD." *Advances in Neural Information Processing Systems* 36 (2024).

FairFace

- Type: Facial Images
- Attributes: Gender, Age, Ethnicity labels (Not self-reported)
  Balanced w.r.t. Ethnicity

- Classification Target: Gender
- Sensitive Attribute: Ethnicity

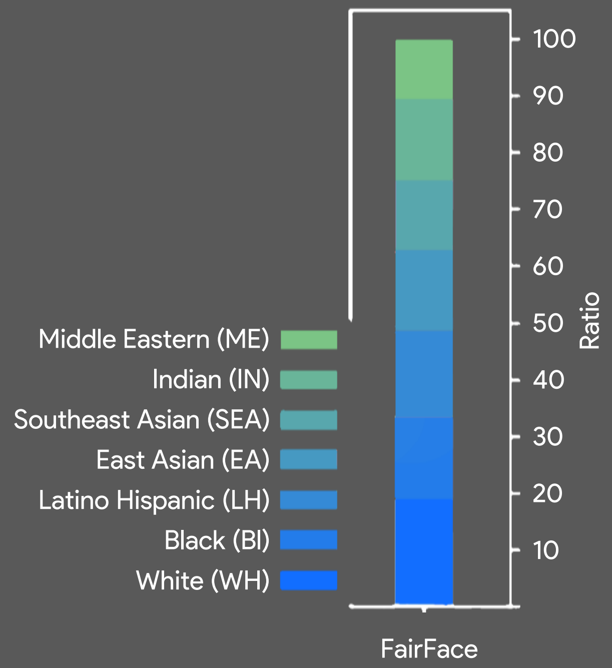

Karkkainen, K., & Joo, J. (2021). FairFace: Face Attribute Dataset for Balanced Race, Gender, and Age for Bias Measurement and Mitigation. In Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (pp. 1548-1558).

Klinikum rechts der Isar
Technische Universität München

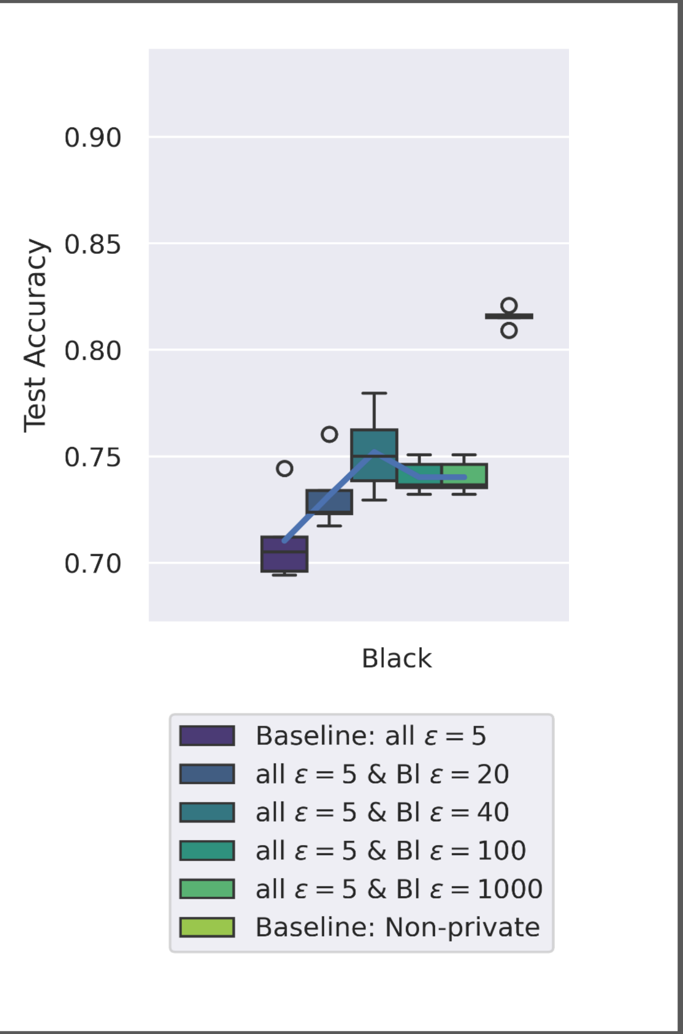

Introduction
Privacy
rivacy vs. Fairness
DP as a Tool
Individual DP
Results
Ongoing Research
Discussion

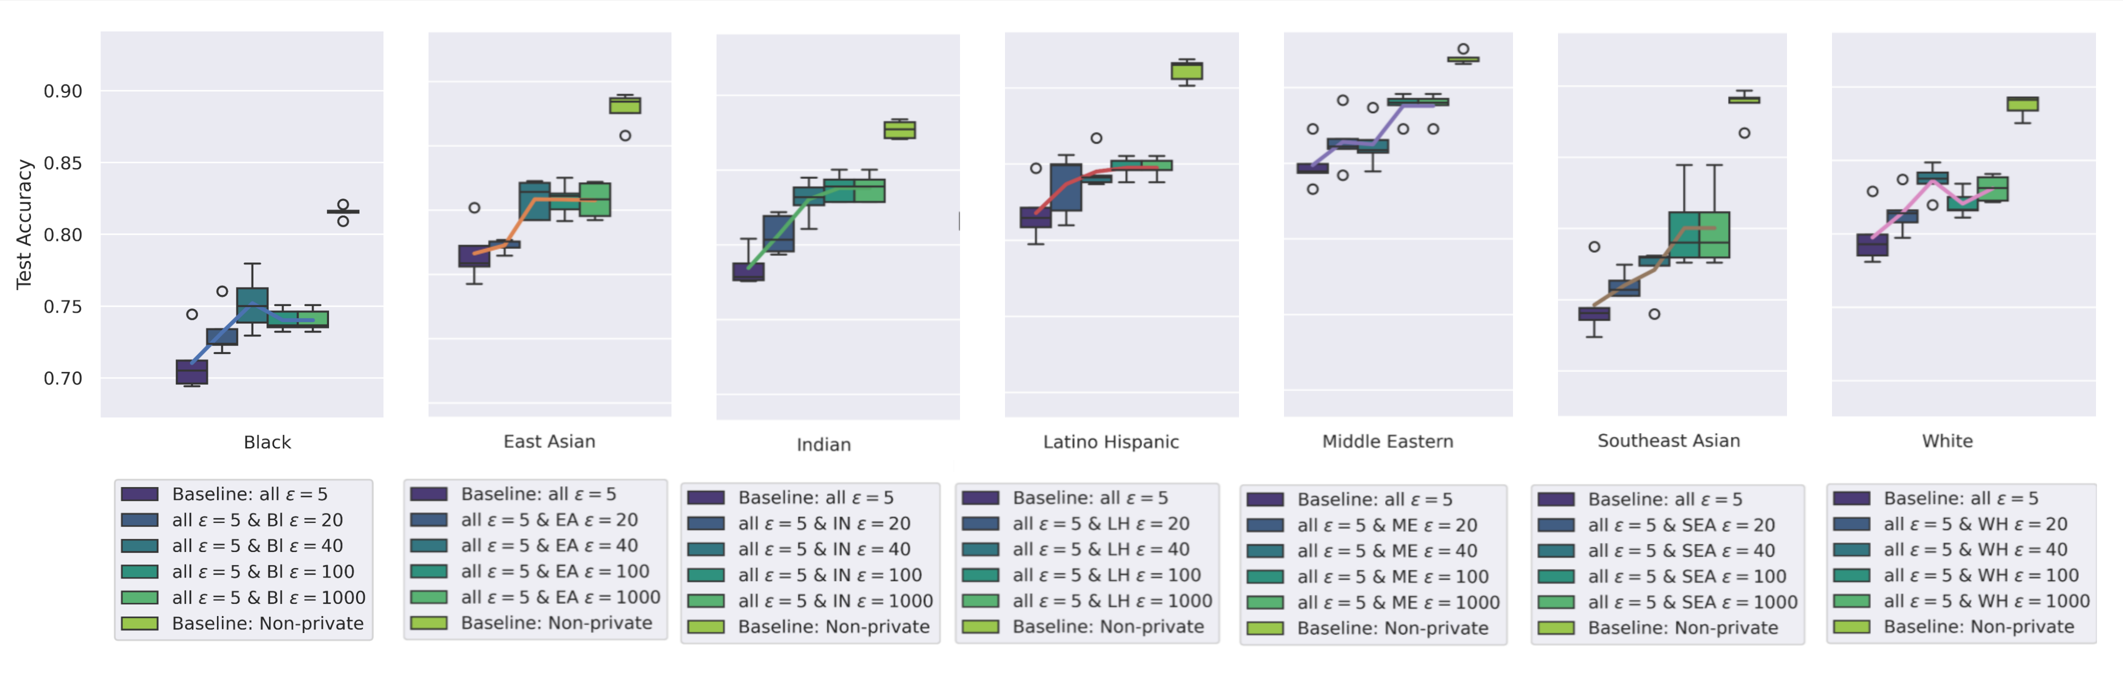

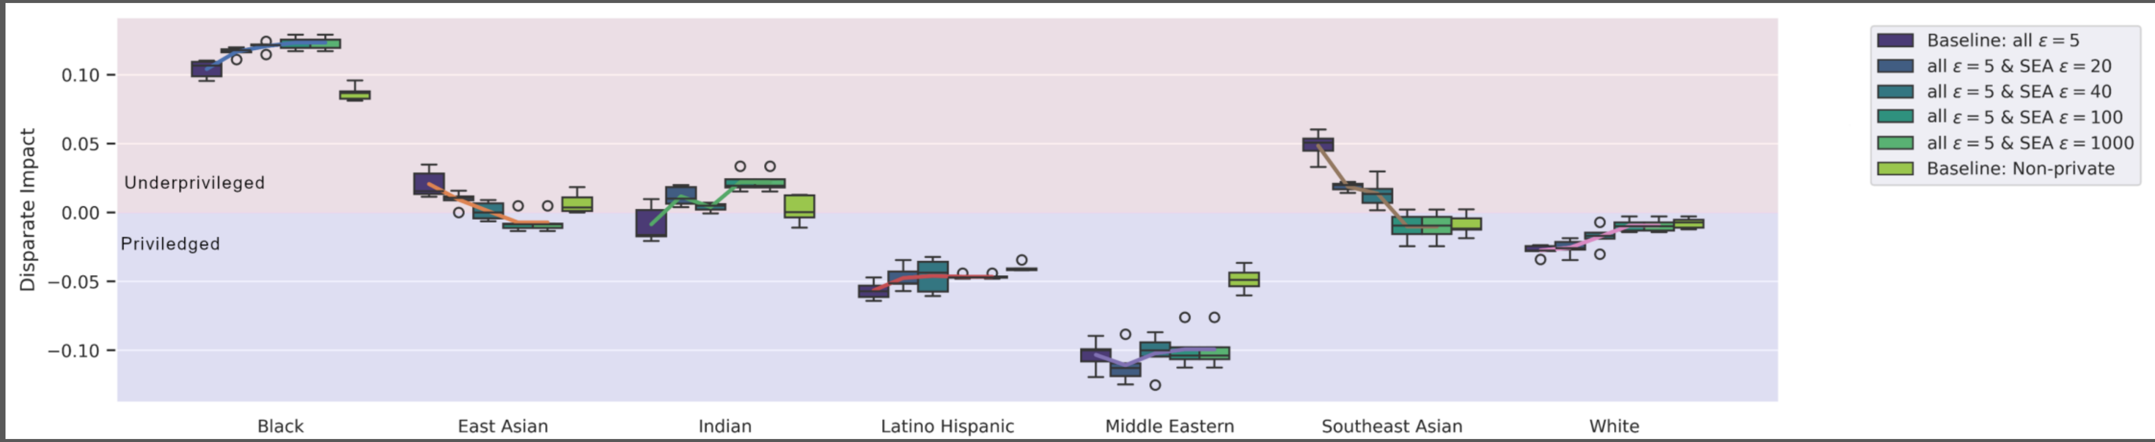
Disparate Impact
0.10
0.05
0.00
−0.05
−0.10
Underprivileged
Priviledged
Black
East Asian
Indian
Latino Hispanic
Middle Eastern
Southeast Asian
White
Baseline: all ε = 5
all ε = 5 & SEA ε = 20
all ε = 5 & SEA ε = 40
all ε = 5 & SEA ε = 100
all ε = 5 & SEA ε = 1000
Baseline: Non-private

Klinikum rechts der Isar
Technische Universität München

Introduction
Privacy
ivacy vs. Fairness
DP as a Tool
Individual DP
Results
ngoing Research
Discussion

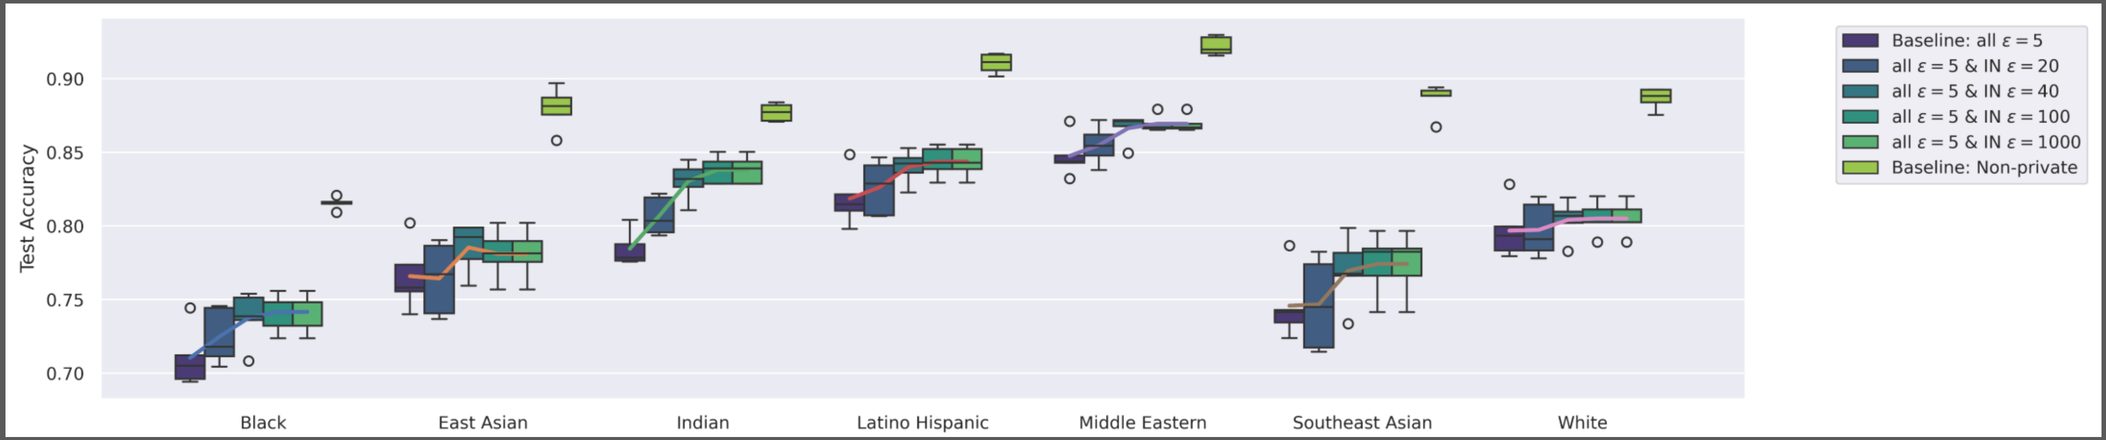

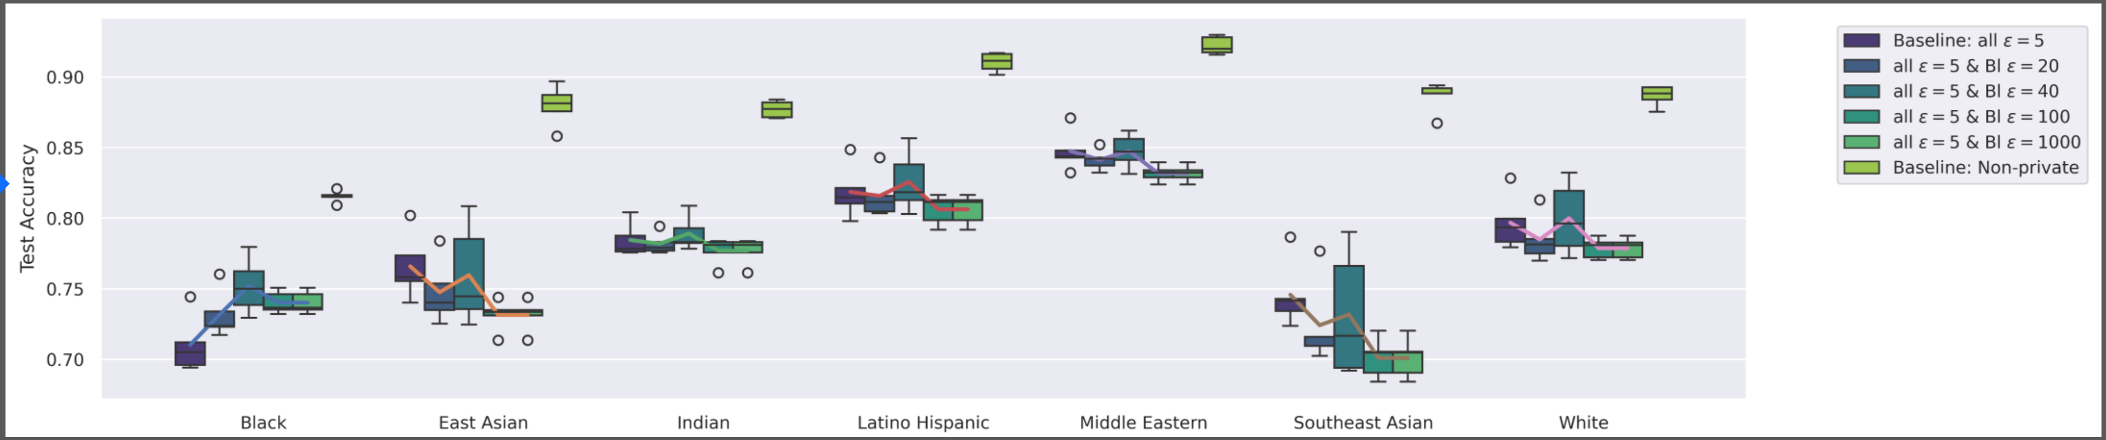

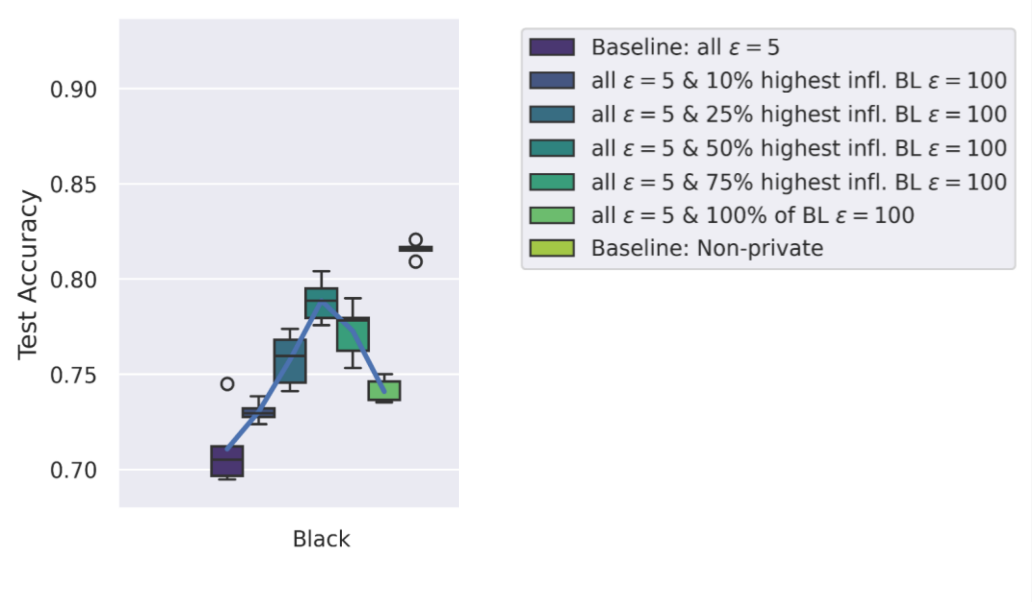
Test Accuracy
0.90
0.85
0.80
0.75
0.70
Black
Baseline: all $\varepsilon = 5$
all $\varepsilon = 5$ & 10% highest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 25% highest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 50% highest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 75% highest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 100% of BL $\varepsilon = 100$
Baseline: Non-private

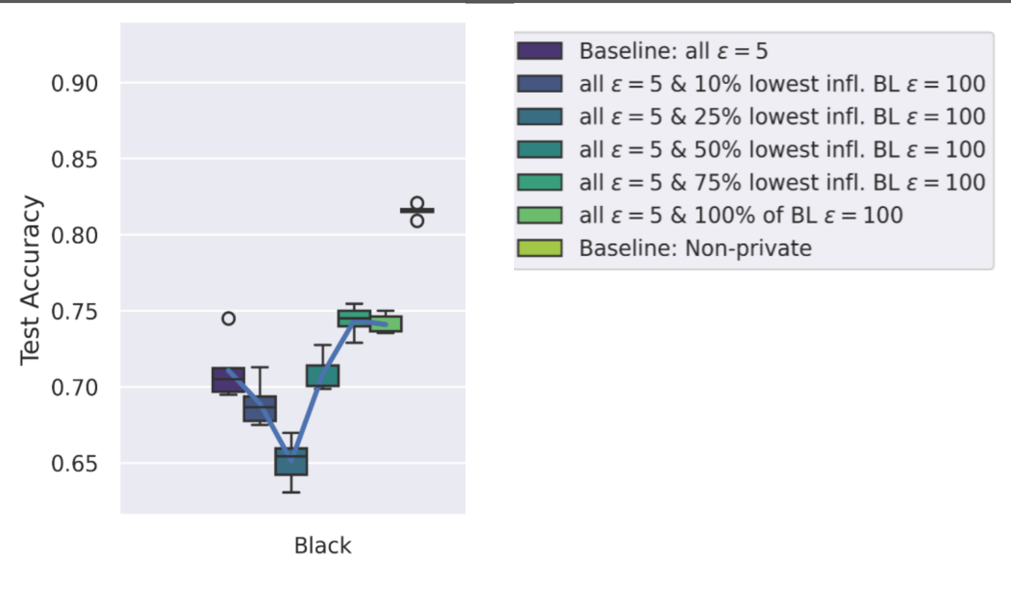
Test Accuracy
0.90
0.85
0.80
0.75
0.70
0.65
Black
Baseline: all $\varepsilon = 5$
all $\varepsilon = 5$ & 10% lowest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 25% lowest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 50% lowest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 75% lowest infl. BL $\varepsilon = 100$
all $\varepsilon = 5$ & 100% of BL $\varepsilon = 100$
Baseline: Non-private

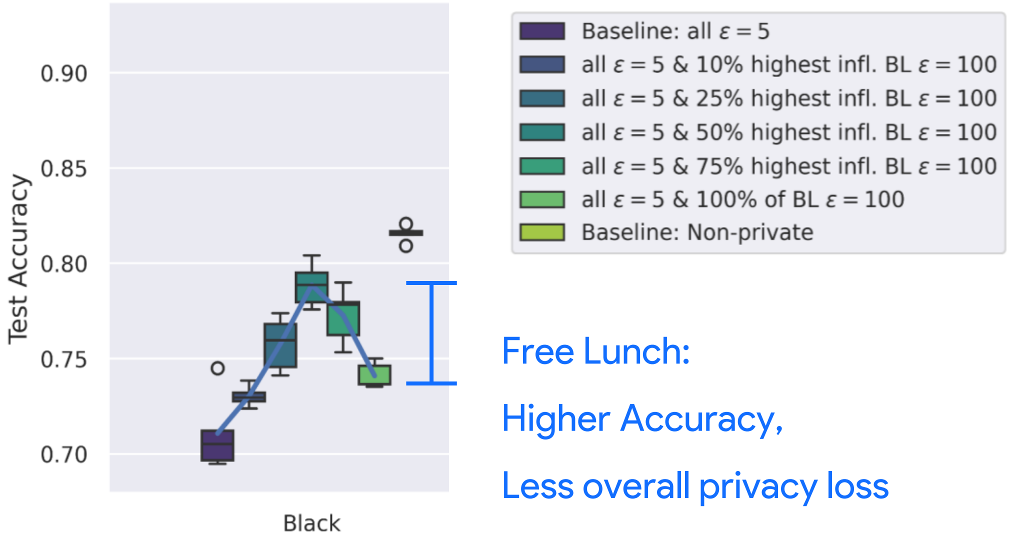

Free Lunch:

Higher Accuracy,

Less overall privacy loss

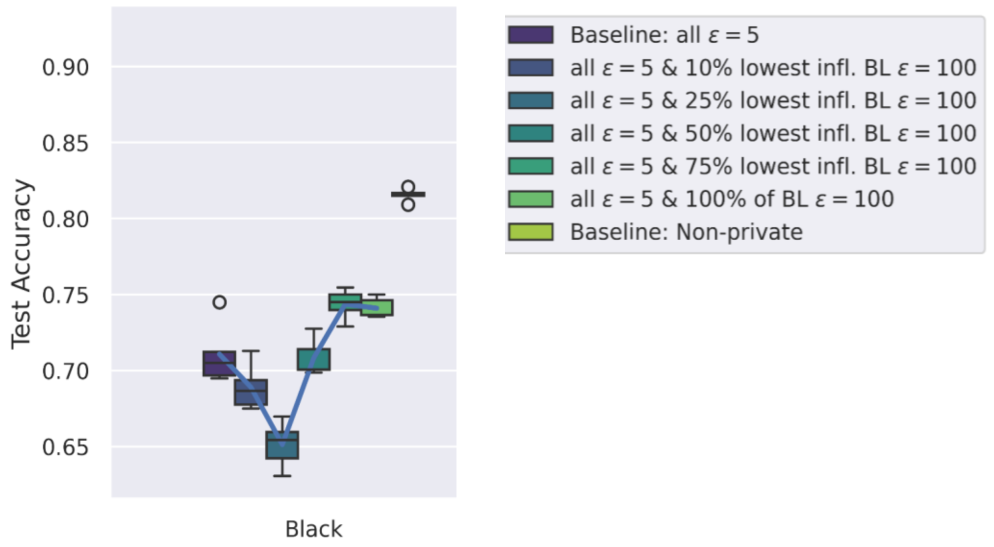

Investigate the effect of the intervention on groups of:

- Higher variance in data
- Poor data quality
- Lack of data

Problems:

- Finding a setting where these have a substantial effect
- Finding a setting that allows to compare the intervention on a dataset with the three different corruptions

---

Predicting the inter-group correlative behaviour

Problems:

- There is no good group correlation metric
- Averaging sample-wise cross-influence metrics yield “group-influences” magnitues smaller (i.e., close to zero)

Increasing group-specific privacy budget increases their theoretical upper bound on the risk
However, for many contributors, the true risk may be far smaller
Evaluate the change in risk using MIA

Problems:
None ☺

Klinikum rechts der Isar
Technische Universität München

Thank you!

Got further questions? Let's connect: johannes.kaiser@tum.de