38th Conference on Uncertainty in Artificial Intelligence
Eindhoven, Netherlands
August 1-5, 2022
uai2022

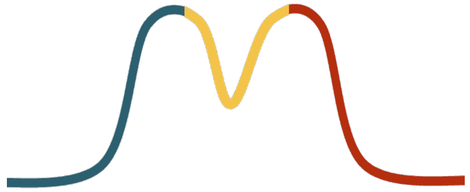

ETH zürich

ETH AI CENTER

How unfair is private learning ?

Amartya Sanyal, Yaxi Hu, Fanny Yang

Amartya

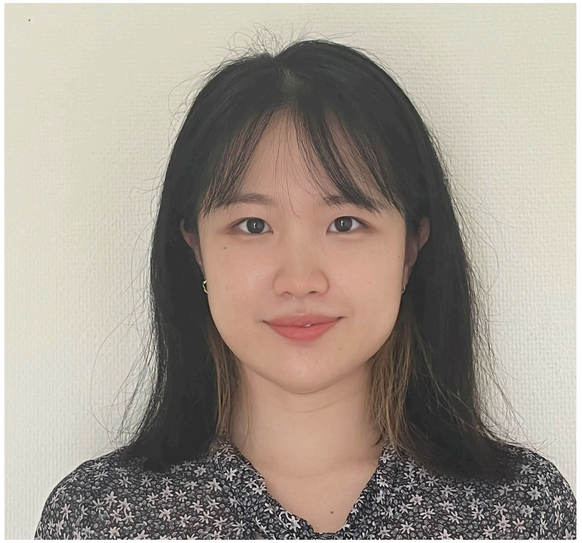

Yaxi

Fanny

Privacy and Fairness

Privacy and Fairness

Privacy and Fairness are both desirable properties in machine learning applications.

Privacy and Fairness

Privacy and Fairness are both desirable properties in machine learning applications.

Prior Work has mostly looked at the intersection:

Privacy and Fairness

Privacy and Fairness are both desirable properties in machine learning applications.

Prior Work has mostly looked at the intersection:

Privacy and Accuracy: Kasiviswanathan et al. 2008, Feldman and Xiao 2014, Alon et. al., 2022.

Privacy and Fairness

Privacy and Fairness are both desirable properties in machine learning applications.

Prior Work has mostly looked at the intersection:

Privacy and Accuracy: Kasiviswanathan et al. 2008, Feldman and Xiao 2014, Alon et. al., 2022.

Fairness and Accuracy: Sagawa et. al. 2019, Du et al. 2021, Goel et. al. 2021.

Privacy and Fairness

Privacy and Fairness are both desirable properties in machine learning applications.

Prior Work has mostly looked at the intersection:

Privacy and Accuracy: Kasiviswanathan et al. 2008, Feldman and Xiao 2014, Alon et. al., 2022.

Fairness and Accuracy: Sagawa et. al. 2019, Du et al. 2021, Goel et. al. 2021.

THIS WORK: The interaction of Privacy and Fairness of nearly **accurate algorithms**.

Differential Privacy

Differential Privacy

Differential Privacy

Differential Privacy

Differential Privacy

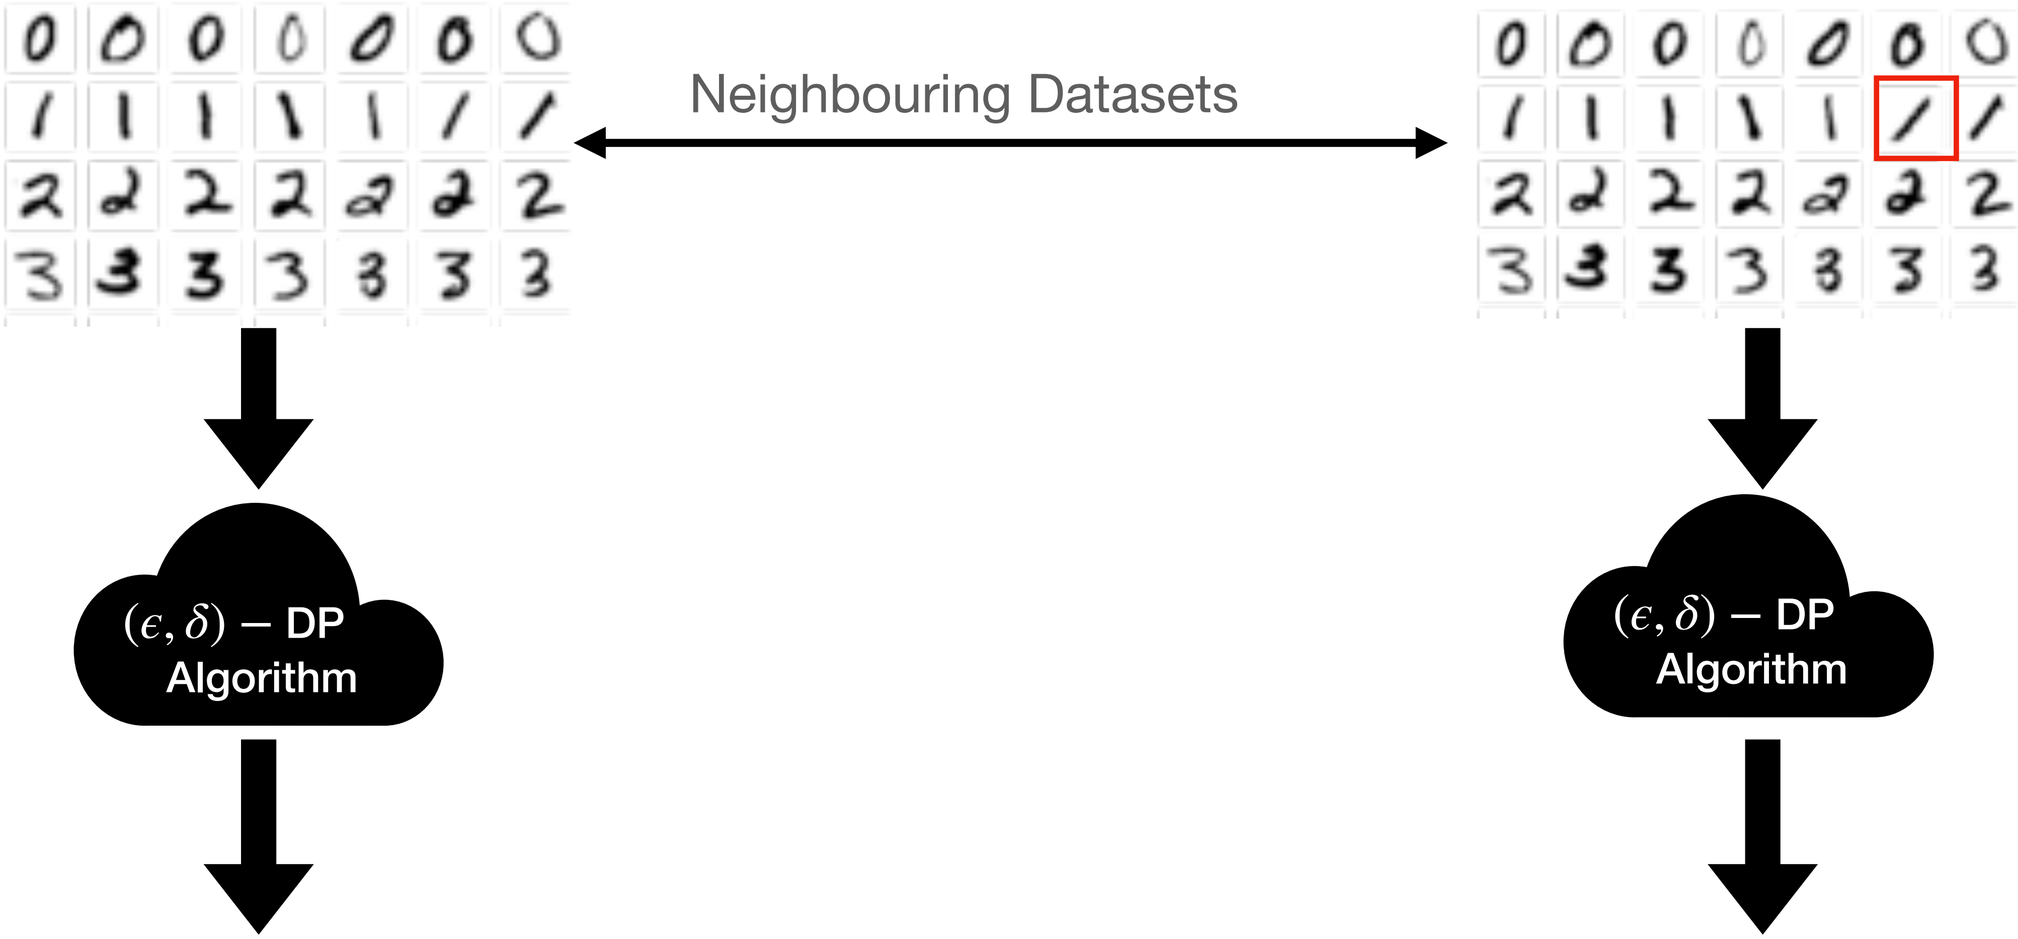

Differential Privacy

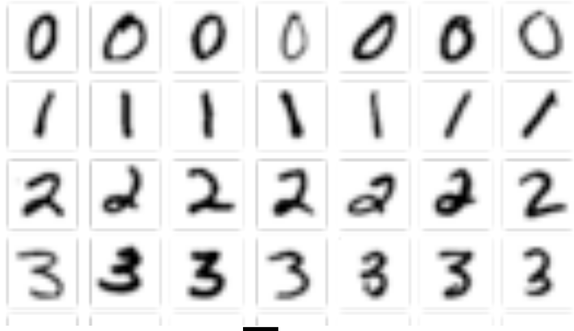

Neighbouring Datasets

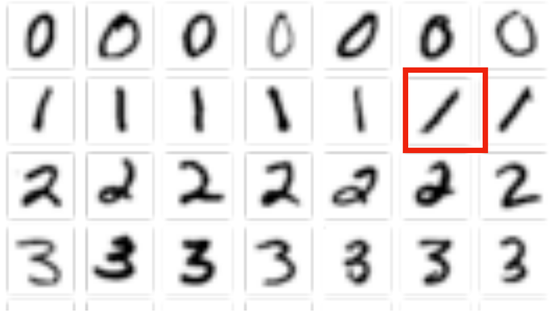

$(\epsilon, \delta) - \text{DP}$ Algorithm

$(\epsilon, \delta) - \text{DP}$ Algorithm

Model 1

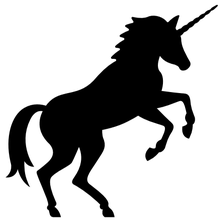

Differential Privacy

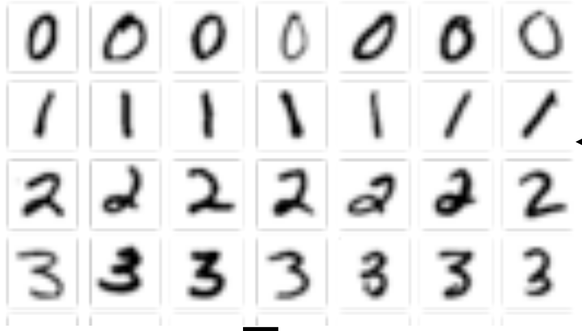

Neighbouring Datasets

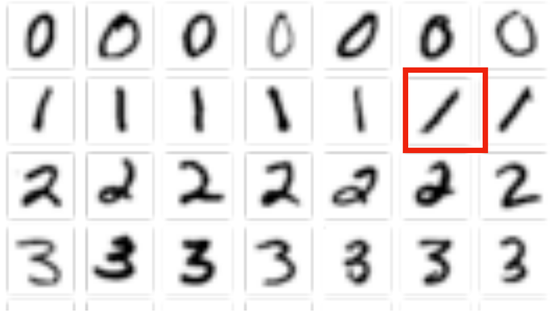

$(\epsilon, \delta) - \text{DP}$ Algorithm

$(\epsilon, \delta) - \text{DP}$ Algorithm

Model 1

Model 2

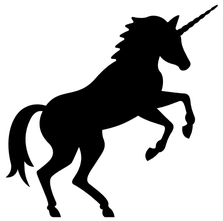

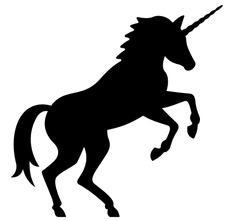

Differential Privacy

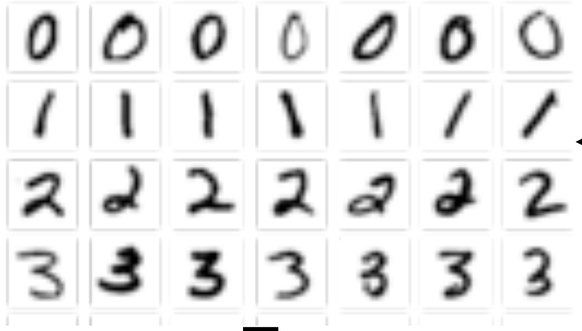

Neighbouring Datasets

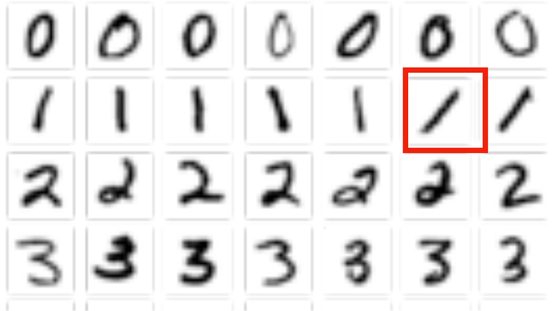

$(\epsilon, \delta) - \text{DP}$ Algorithm

$(\epsilon, \delta) - \text{DP}$ Algorithm

Model 1

Model 2

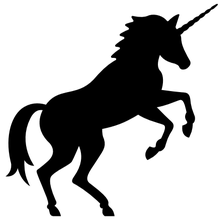

Basically same

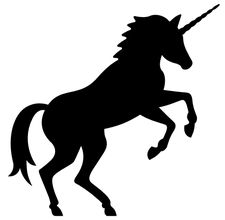

Differential Privacy

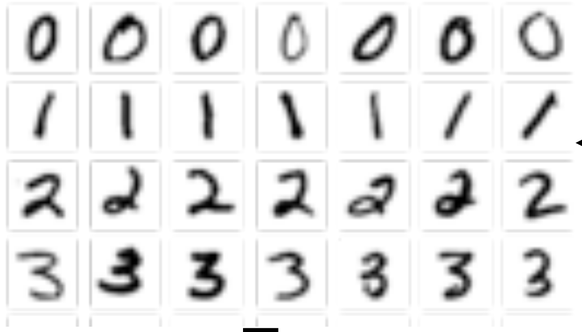

Neighbouring Datasets

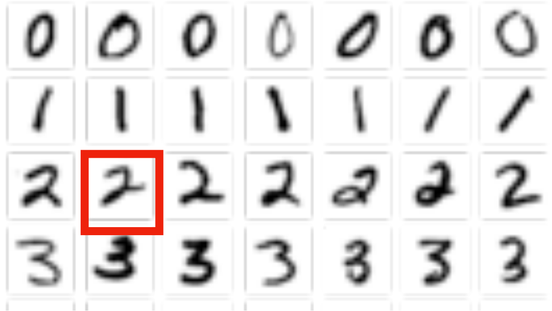

$(\epsilon, \delta) - \text{DP}$ Algorithm

$(\epsilon, \delta) - \text{DP}$ Algorithm

Model 1

Model 2

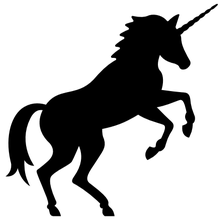

Basically same

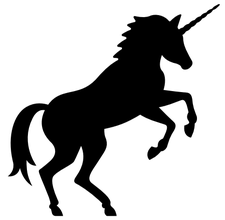

Differential Privacy

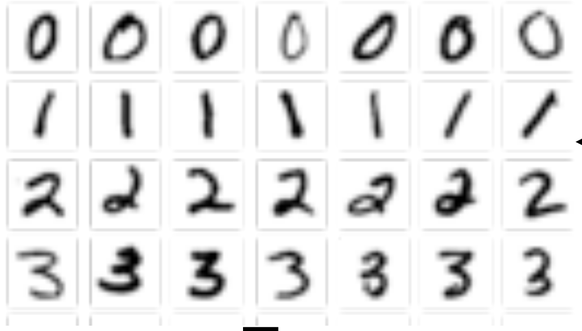

Neighbouring Datasets

$(\epsilon, \delta) - \text{DP}$ Algorithm

$(\epsilon, \delta) - \text{DP}$ Algorithm

Model 1

Model 2

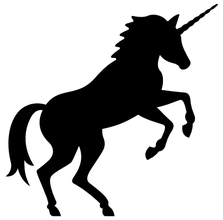

Basically same

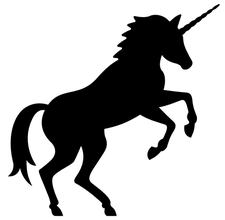

(Un) Fairness (Accuracy Discrepancy)

(Un) Fairness (Accuracy Discrepancy)

(Un) Fairness (Accuracy Discrepancy)

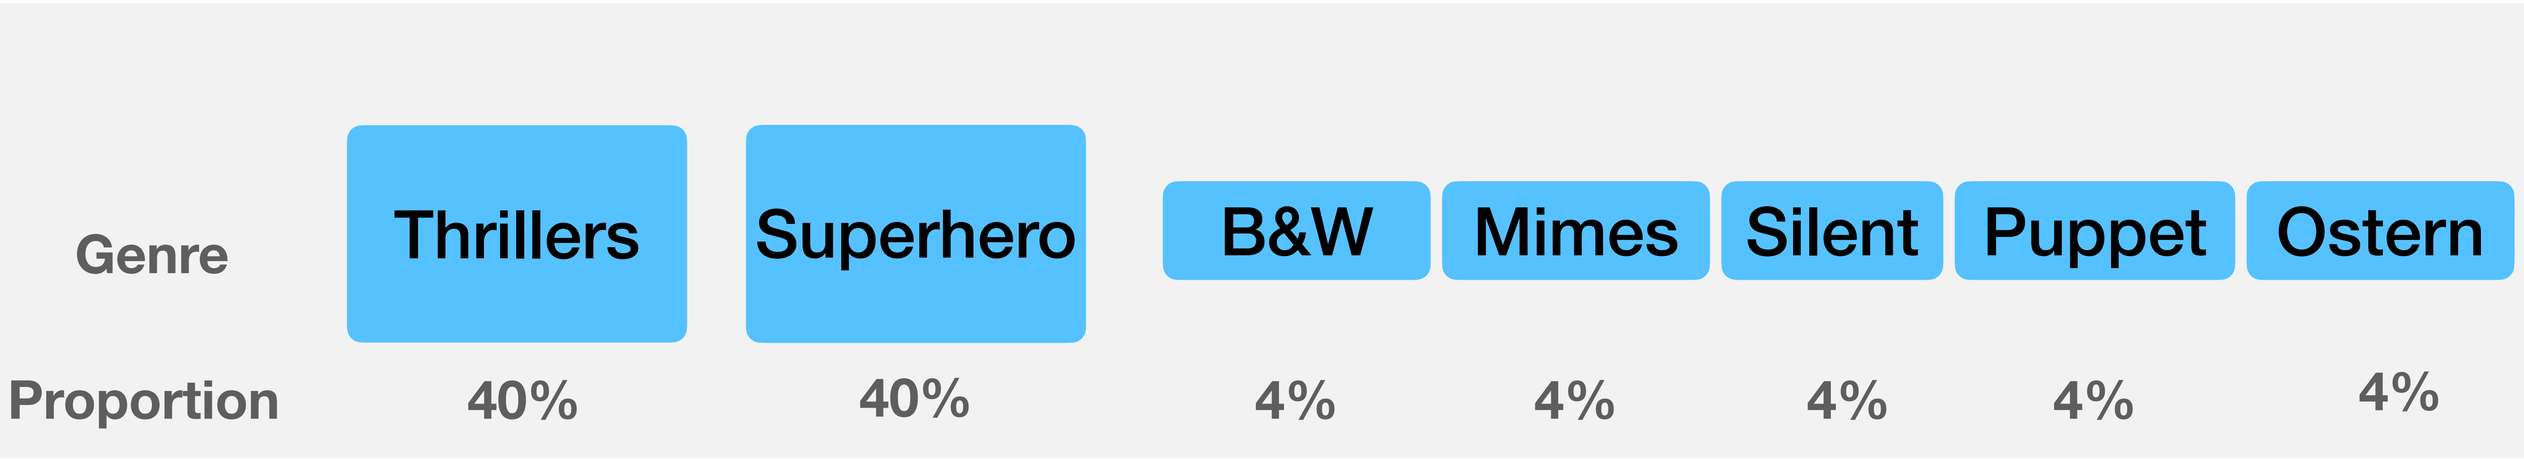

(Un) Fairness (Accuracy Discrepancy)

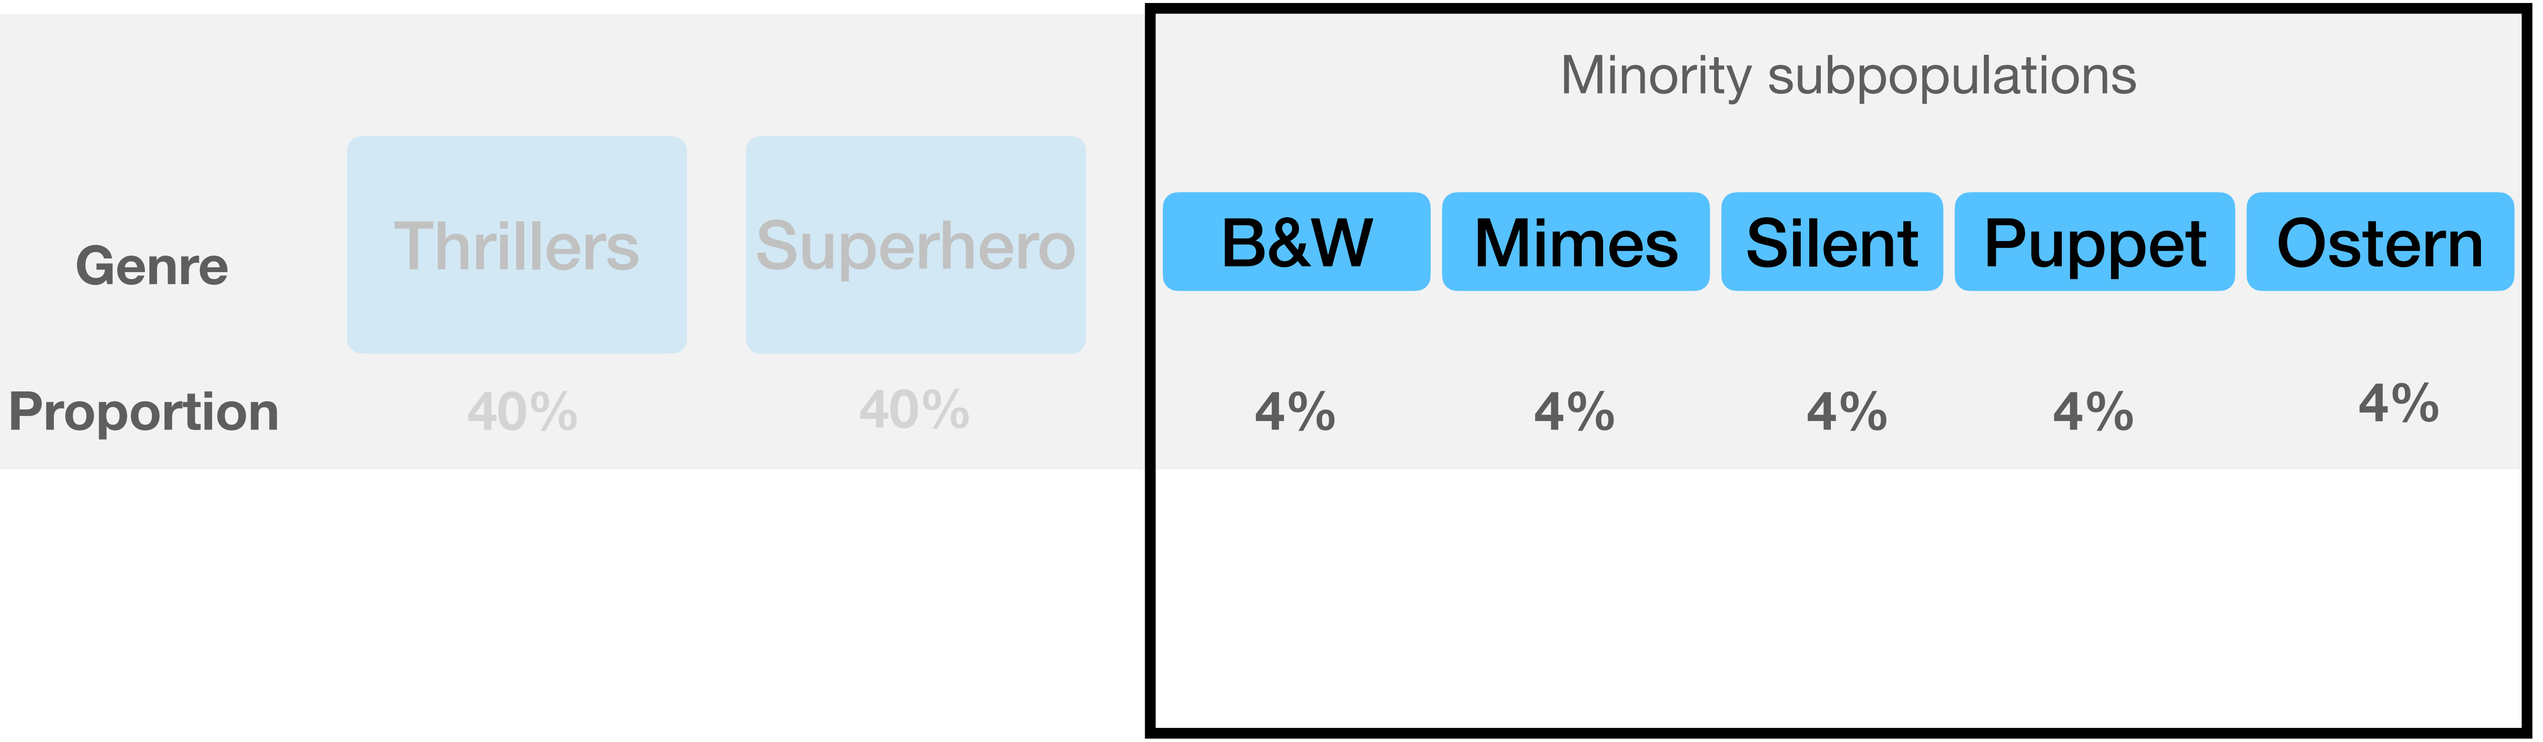

(Un) Fairness (Accuracy Discrepancy)

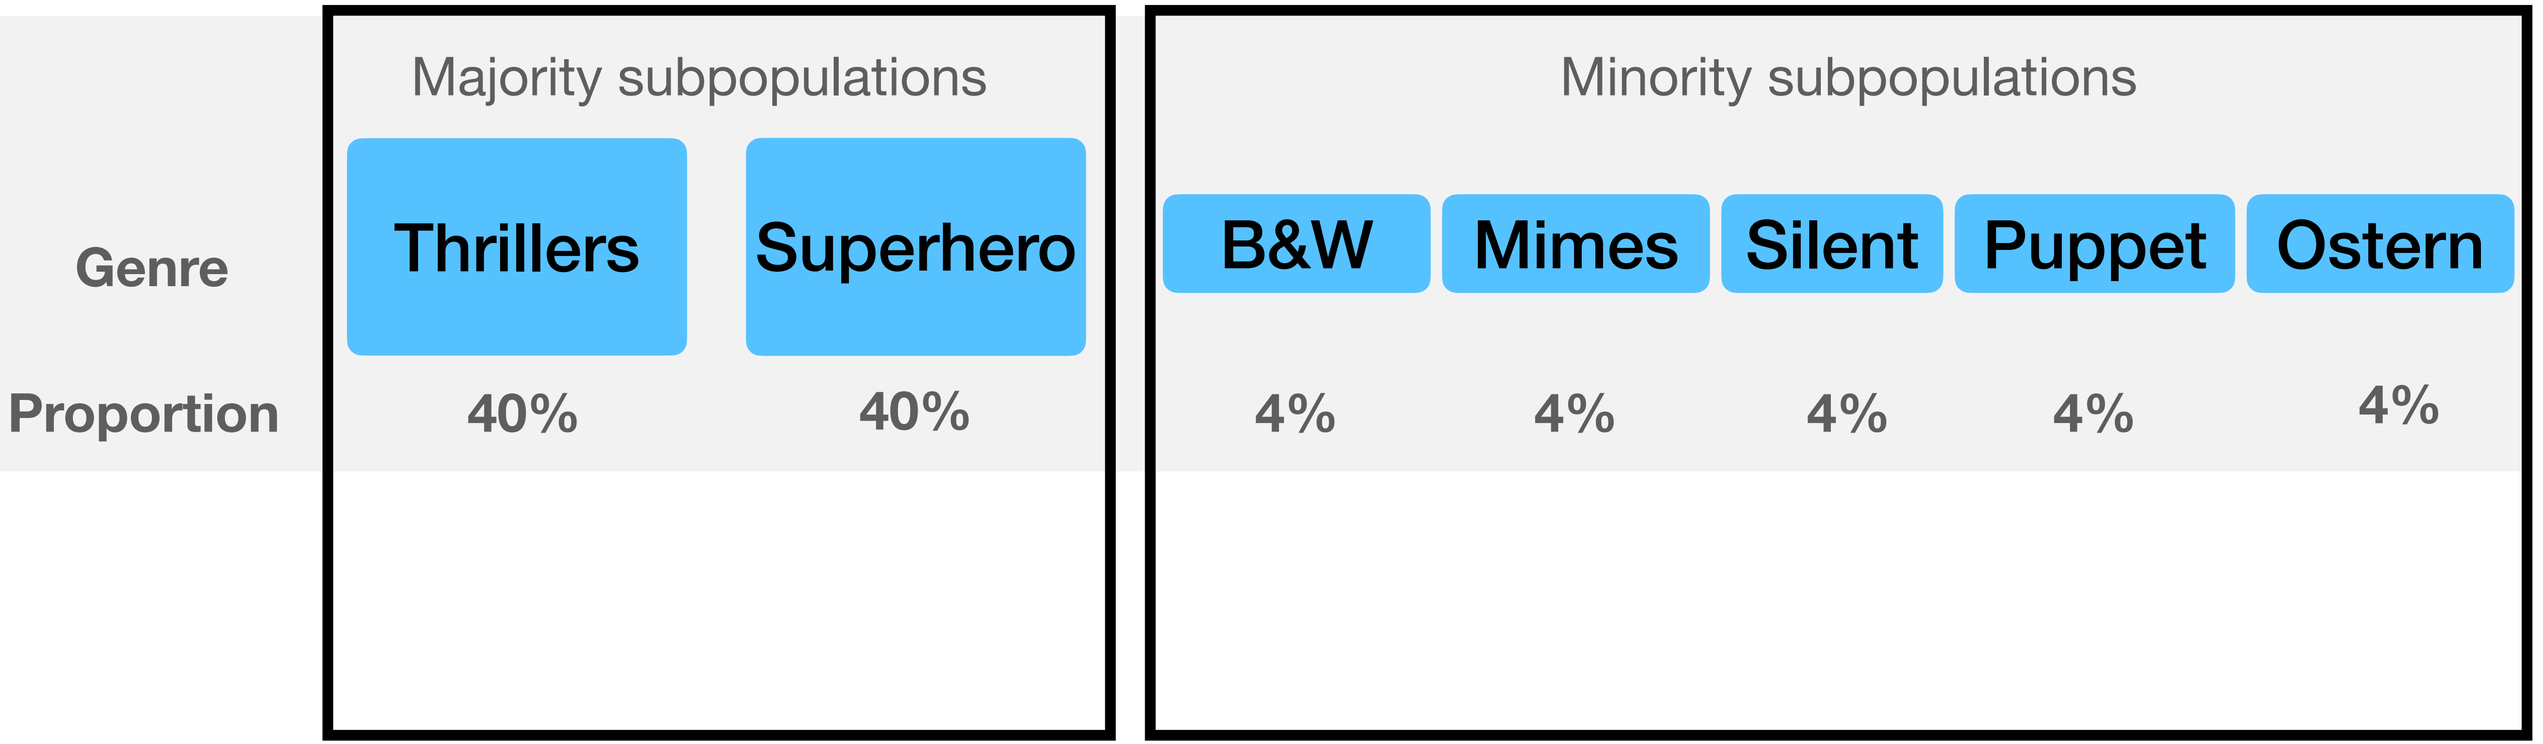

(Un) Fairness (Accuracy Discrepancy)

ML Problem: Is the movie safe to watch for kids ?

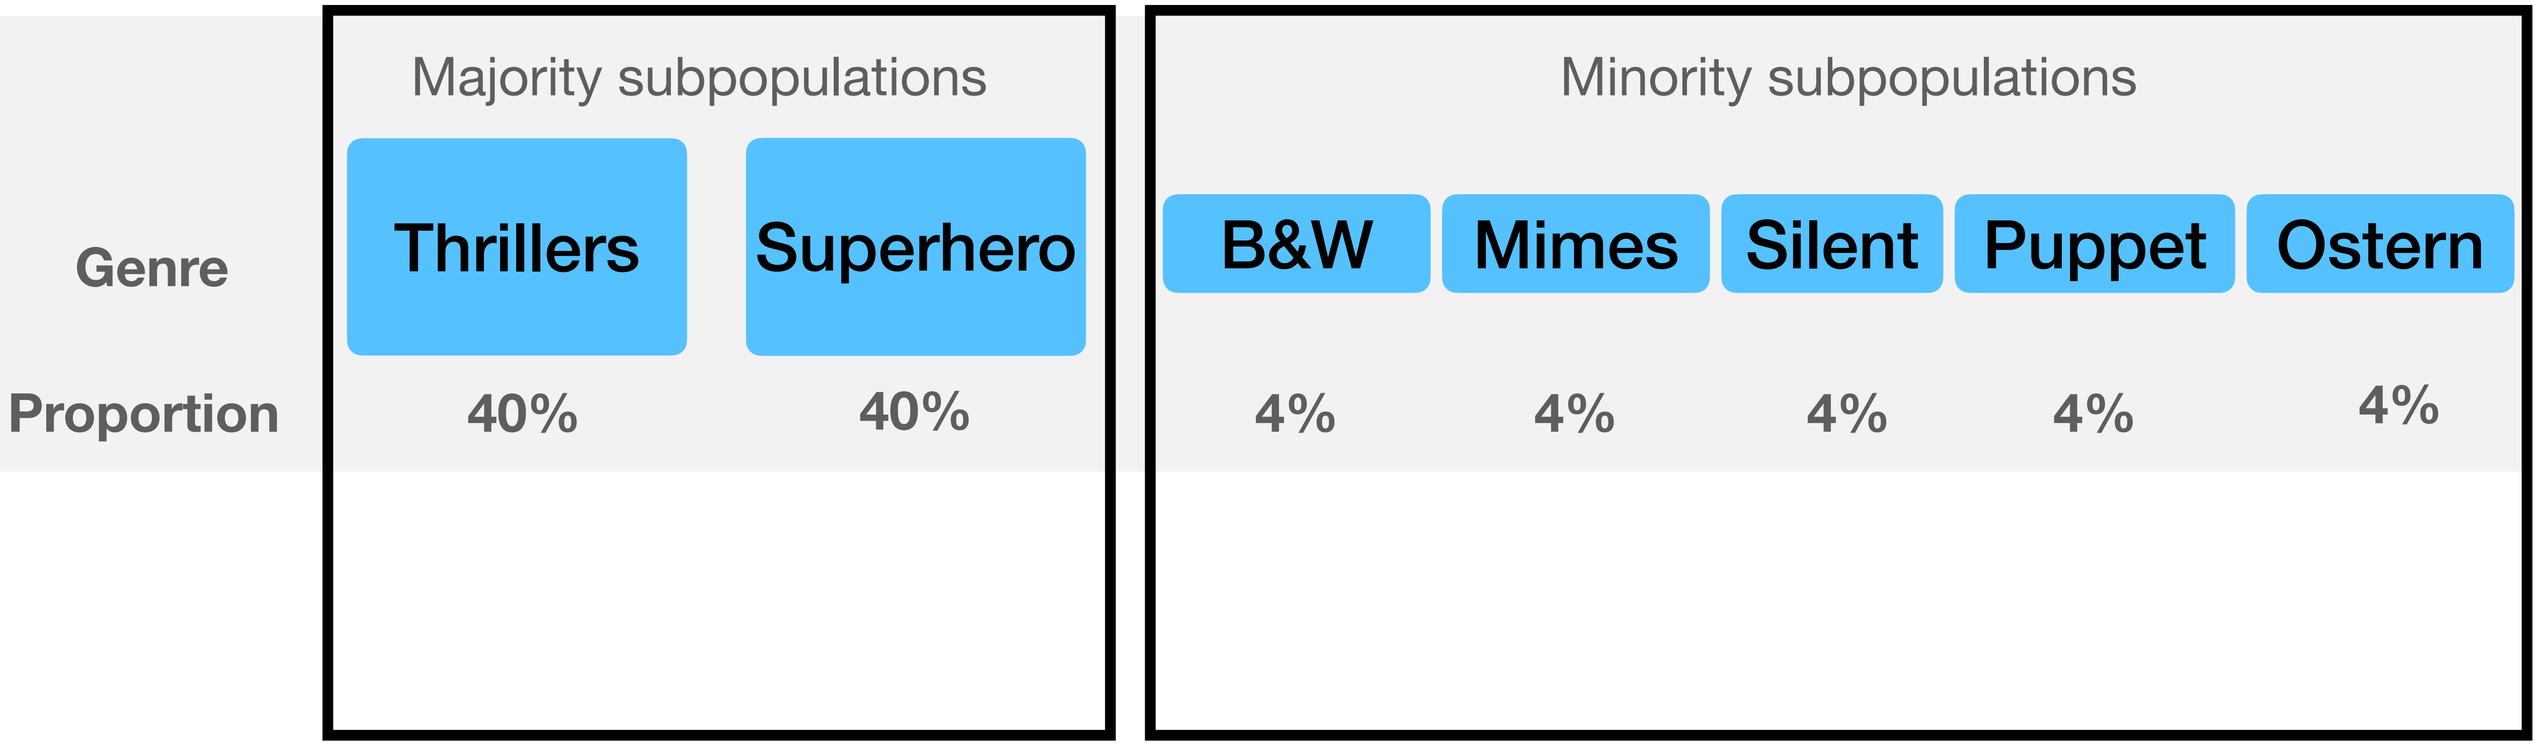

(Un) Fairness (Accuracy Discrepancy)

ML Problem: Is the movie safe to watch for kids ?

	Majority subpopulations		Minority subpopulations				
Genre	Thrillers	Superhero	B&W	Mimes	Silent	Puppet	Ostern
Proportion	40%	40%	4%	4%	4%	4%	4%
Error	5%	5%	65%	75%	80%	80%	50%

(Un) Fairness (Accuracy Discrepancy)

ML Problem: Is the movie safe to watch for kids ?

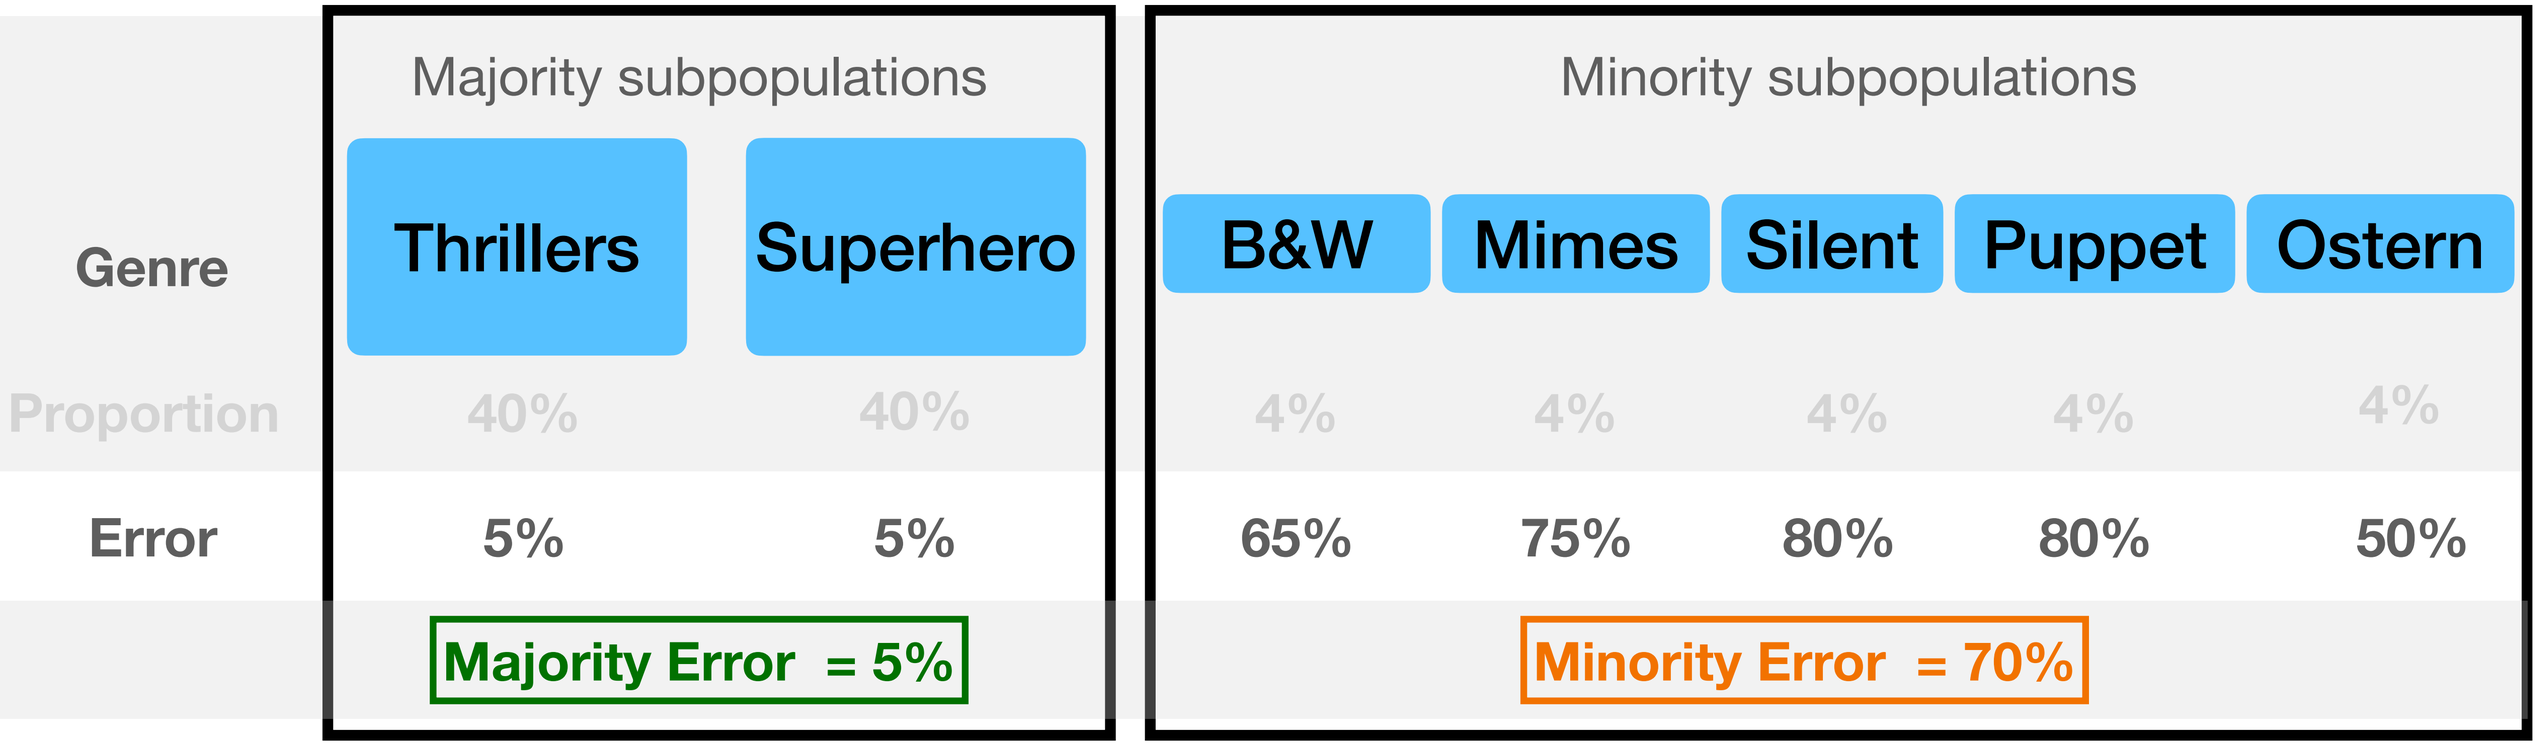

(Un) Fairness (Accuracy Discrepancy)

ML Problem: Is the movie safe to watch for kids ?

	Majority subpopulations		Minority subpopulations				
Genre	Thrillers	Superhero	B&W	Mimes	Silent	Puppet	Ostern
Proportion	40%	40%	4%	4%	4%	4%	4%
Error	5%	5%	65%	75%	80%	80%	50%
	Majority Error = 5%		Minority Error = 70%				

Total Error = 18%

(Un) Fairness (Accuracy Discrepancy)

ML Problem: Is the movie safe to watch for kids ?

	Majority subpopulations		Minority subpopulations				
Genre	Thrillers	Superhero	B&W	Mimes	Silent	Puppet	Ostern
Proportion	40%	40%	4%	4%	4%	4%	4%
Error	5%	5%	65%	75%	80%	80%	50%
	Majority Error = 5%		Minority Error = 70%				

Total Error = 18%

Accuracy Discrepancy = Minority Error - Total Error

(Un) Fairness (Accuracy Discrepancy)

ML Problem: Is the movie safe to watch for kids ?

	Majority subpopulations		Minority subpopulations				
Genre	Thrillers	Superhero	B&W	Mimes	Silent	Puppet	Ostern
Proportion	40%	40%	4%	4%	4%	4%	4%
Error	5%	5%	65%	75%	80%	80%	50%
	Majority Error = 5%		Minority Error = 70%				

Total Error = 18%

Accuracy Discrepancy = 70 - 18 = 52%

Example dataset

CelebA

Example dataset

CelebA

Example dataset

CelebA

40 binary attributes with each image

Eyeglasses

Bangs

Pointy Noise

Example dataset

CelebA

40 binary attributes -> **2^{40} subpopulations.**

40 binary attributes with each image

Eyeglasses

Bangs

Pointy Noise

Example dataset

CelebA

40 binary attributes -> 2^{40} **subpopulations.**

- **Subpopulation 1**: Eyeglasses, bangs, …, pointy nose.

40 binary attributes with each image

Eyeglasses

Bangs

Pointy Noise

Example dataset

CelebA

40 binary attributes -> 2^{40} **subpopulations.**

40 binary attributes with each image

- **Subpopulation 1**: Eyeglasses, bangs, …, pointy nose.
- **Subpopulation 2:** No eyeglasses, bangs,……,pointy noise.

Eyeglasses

Bangs

Pointy Noise

Example dataset

CelebA

40 binary attributes with each image

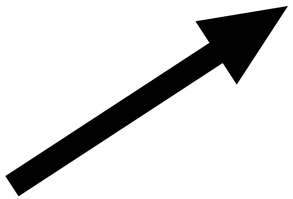

40 binary attributes -> **2^{40} subpopulations.**

- **Subpopulation 1**: Eyeglasses, bangs, …, pointy nose.
- **Subpopulation 2:** No eyeglasses, bangs,……,pointy noise.
- …
- …
- **Subpopulation 2^{40}:** No eyeglasses, no bangs,…, no pointy nose.

Example dataset

CelebA

40 binary attributes with each image

40 binary attributes -> 2^{40} **subpopulations.**

- **Subpopulation 1**: Eyeglasses, bangs, …, pointy nose.
- **Subpopulation 2:** No eyeglasses, bangs,……,pointy noise.
- …
- …
- **Subpopulation 2^{40}:** No eyeglasses, no bangs,…, no pointy nose.

Example dataset

CelebA

40 binary attributes with each image

Eyeglasses

Bangs

Pointy Noise

40 binary attributes -> **2^{40} subpopulations.**

- **Subpopulation 1**: Eyeglasses, bangs, …, pointy nose.
- **Subpopulation 2:** No eyeglasses, bangs,……,pointy noise.
- …
- …
- **Subpopulation 2^{40}:** No eyeglasses, no bangs,…, no pointy nose.

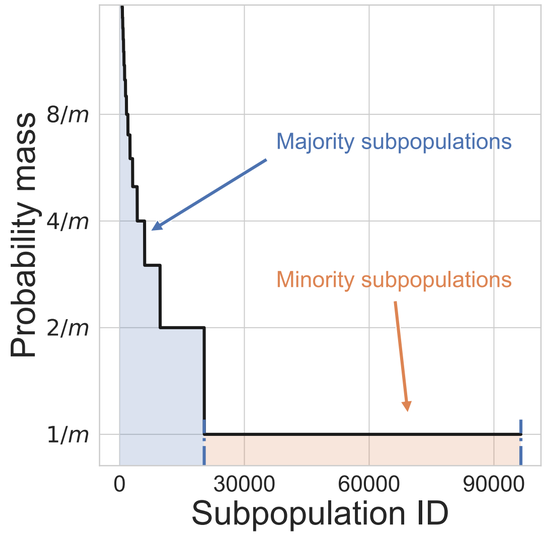

Example dataset

CelebA

40 binary attributes with each image

Eyeglasses

Bangs

Pointy Noise

40 binary attributes -> **2^{40} subpopulations.**

- **Subpopulation 1**: Eyeglasses, bangs, …, pointy nose.
- **Subpopulation 2:** No eyeglasses, bangs,……,pointy noise.
- …
- …
- **Subpopulation 2^{40}:** No eyeglasses, no bangs,…, no pointy nose.

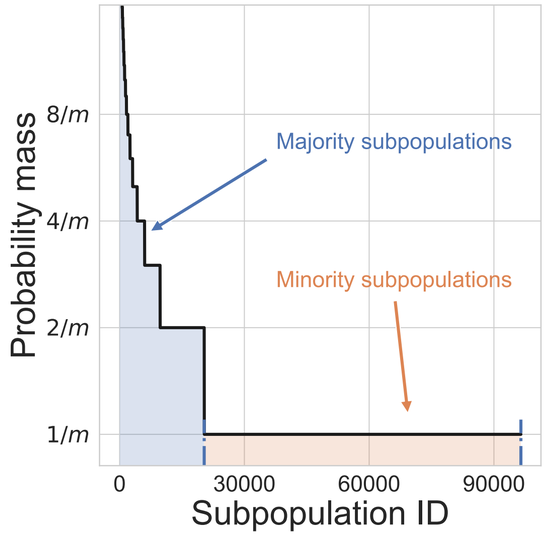

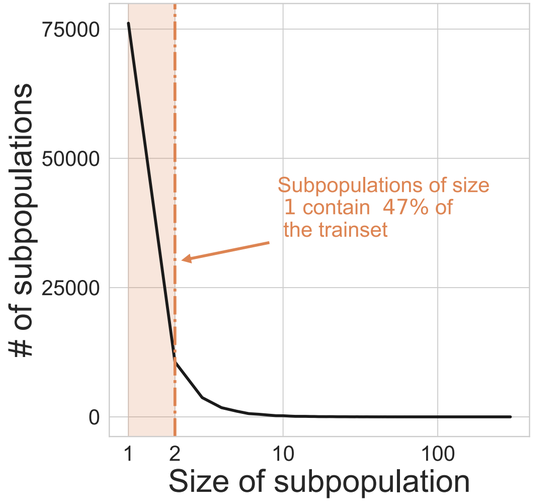

Privacy vs Fairness

CelebA

Privacy vs Fairness

CelebA

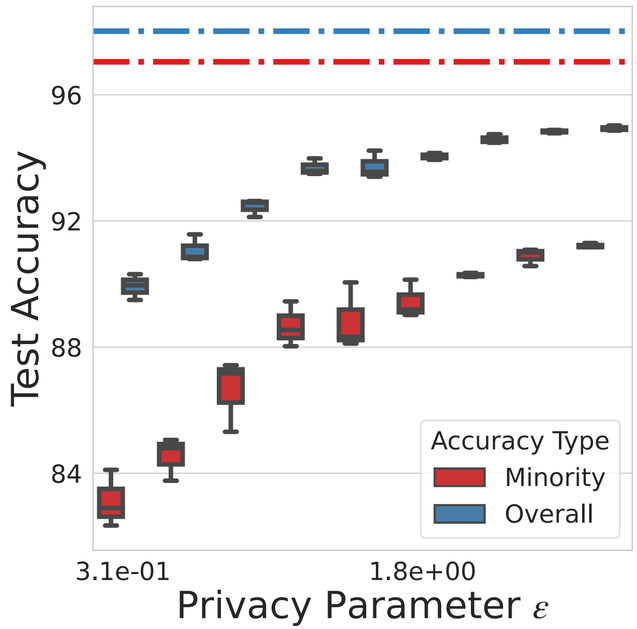

Privacy vs Fairness

CelebA

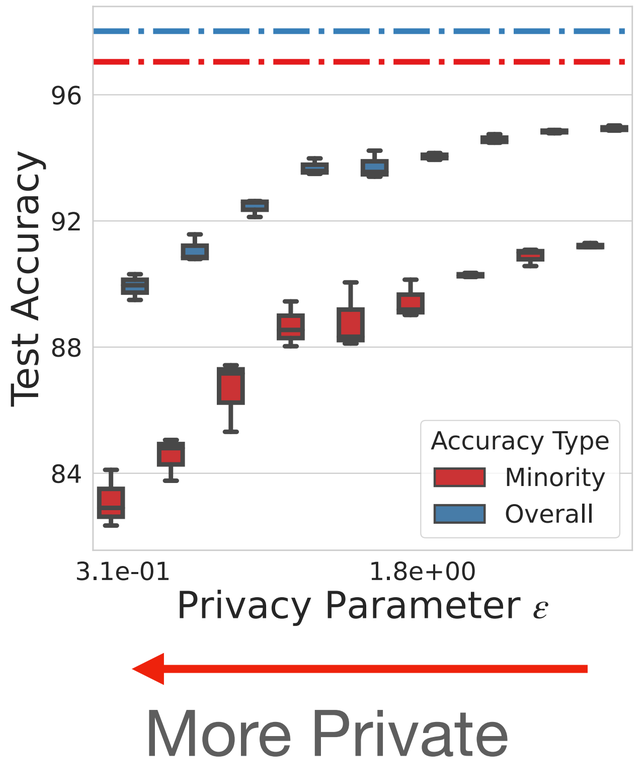

Privacy vs Fairness

CelebA

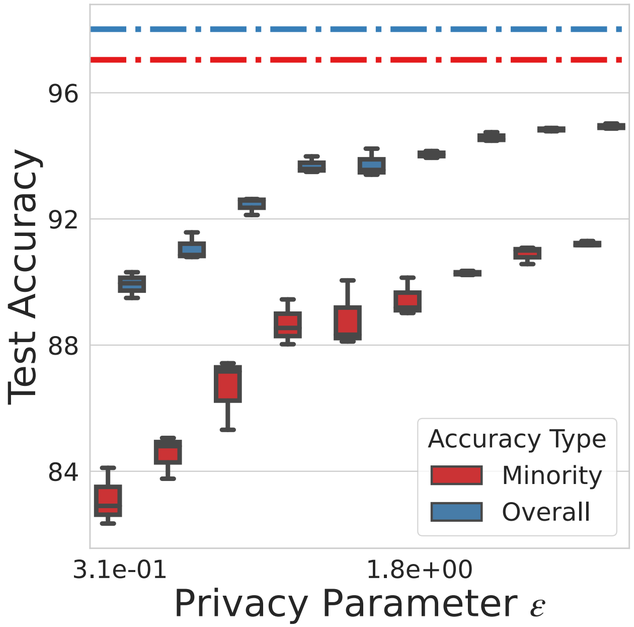

More Private

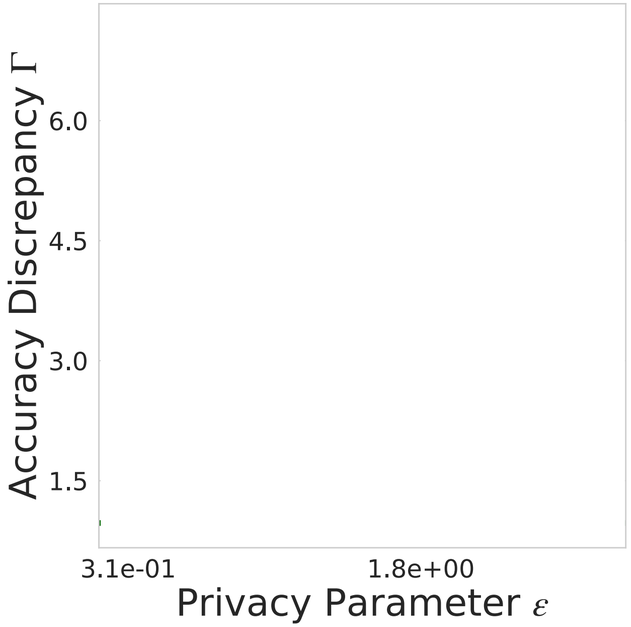

Privacy vs Fairness

CelebA

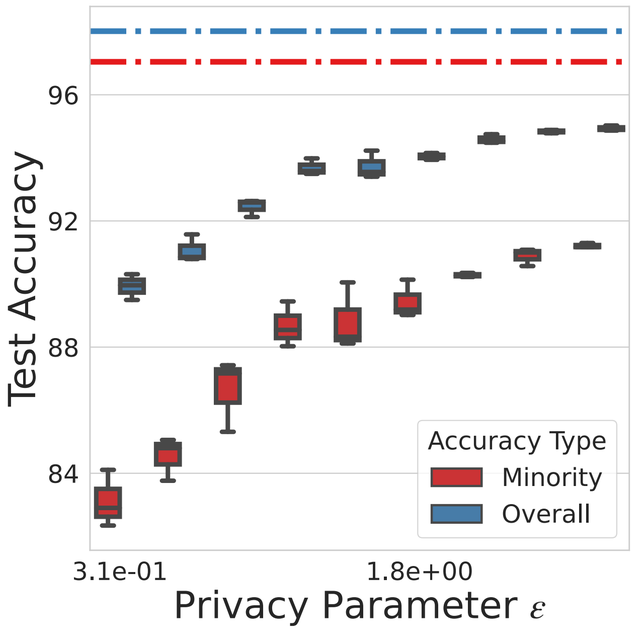

More Private

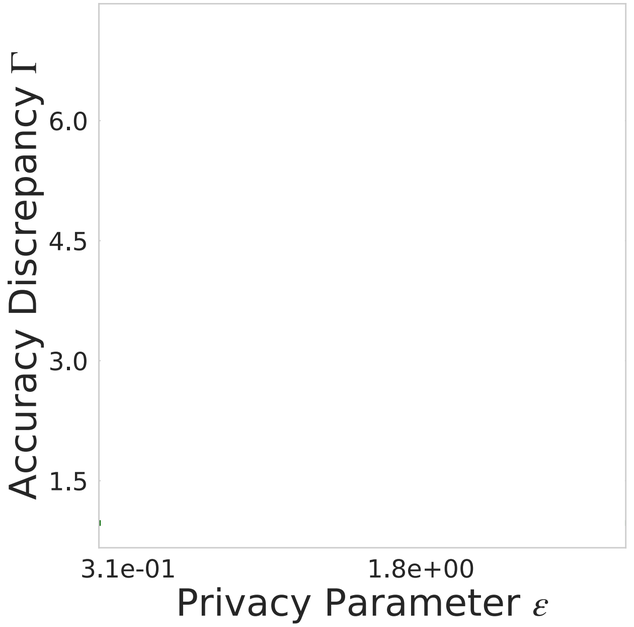

More Private

Privacy vs Fairness

CelebA

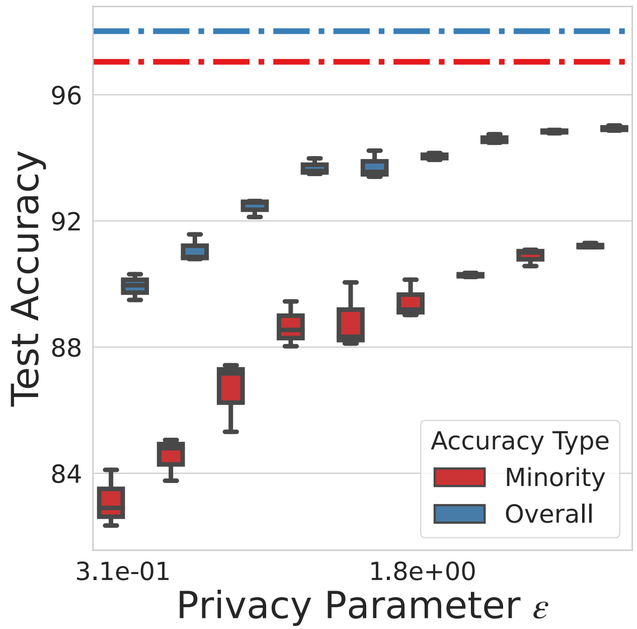

More Private

Less Fair

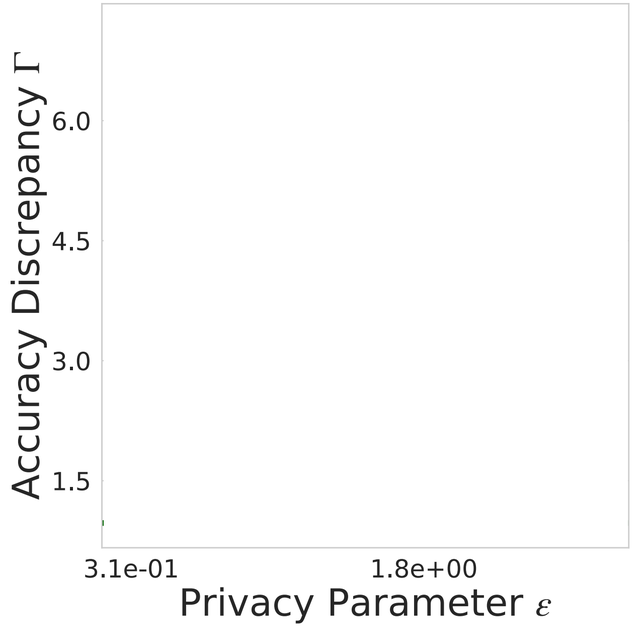

More Private

Privacy vs Fairness

CelebA

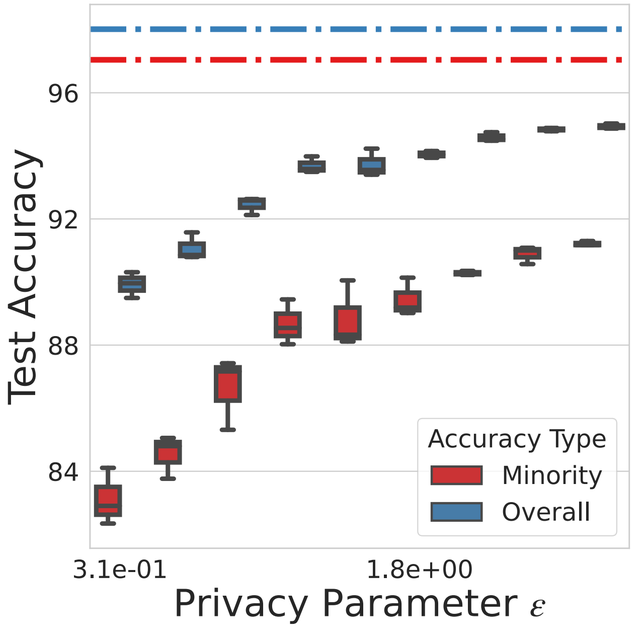

More Private

Less Fair

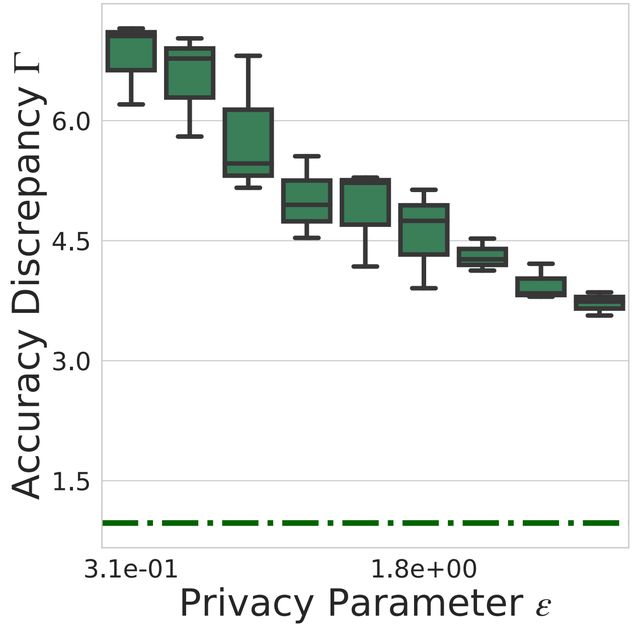

More Private

Is this trend systematic?

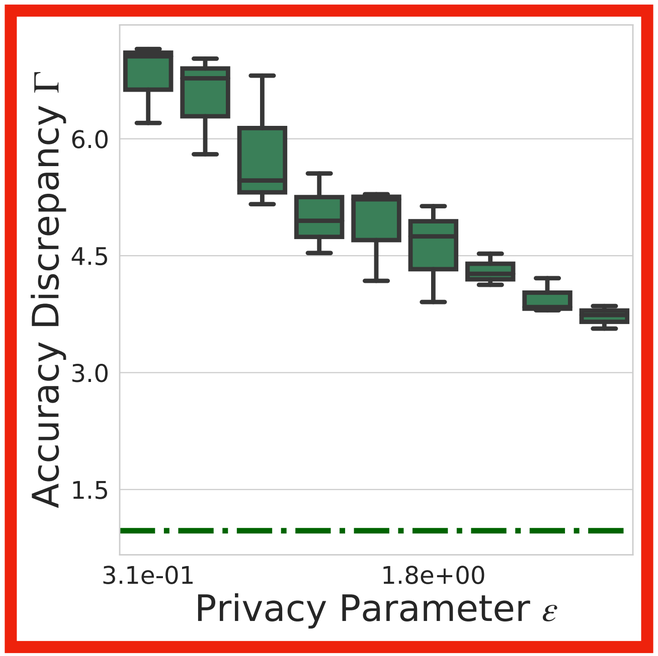

Is this trend systematic?

Main Contribution:
We prove this trend in a model-agnostic setting

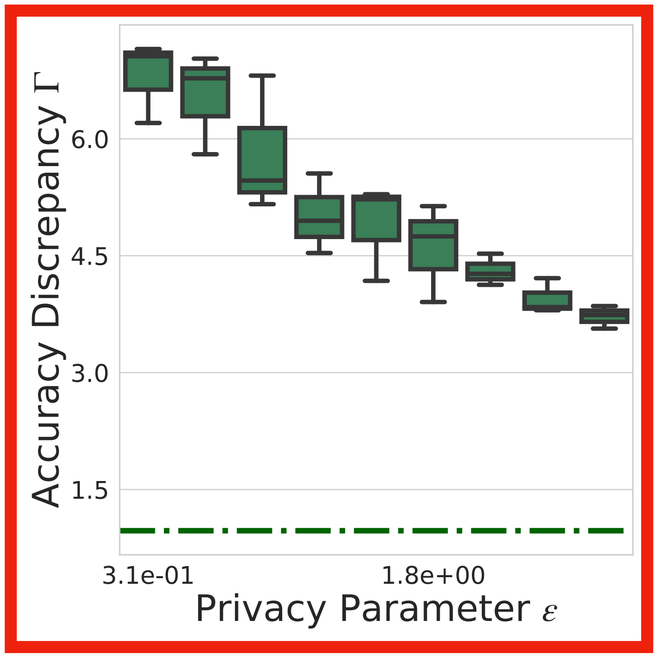

Is this trend systematic?

Main Contribution:
We prove this trend in a model-agnostic setting for long-tailed distribution.

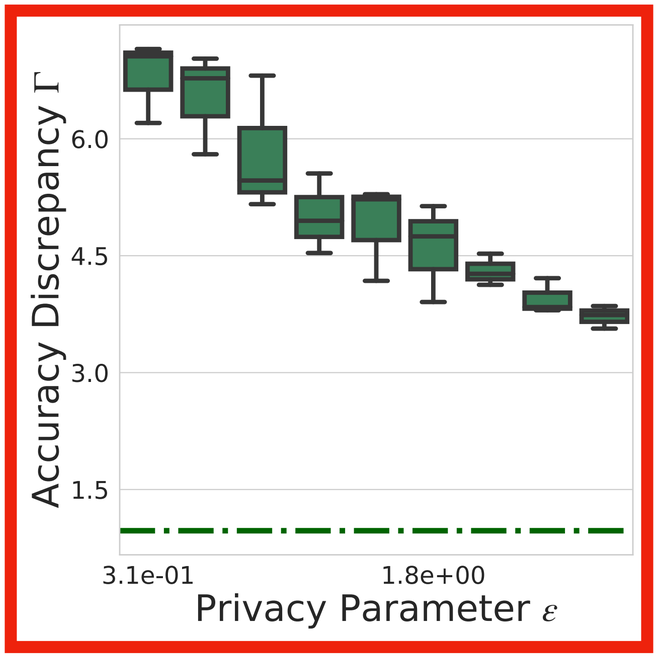

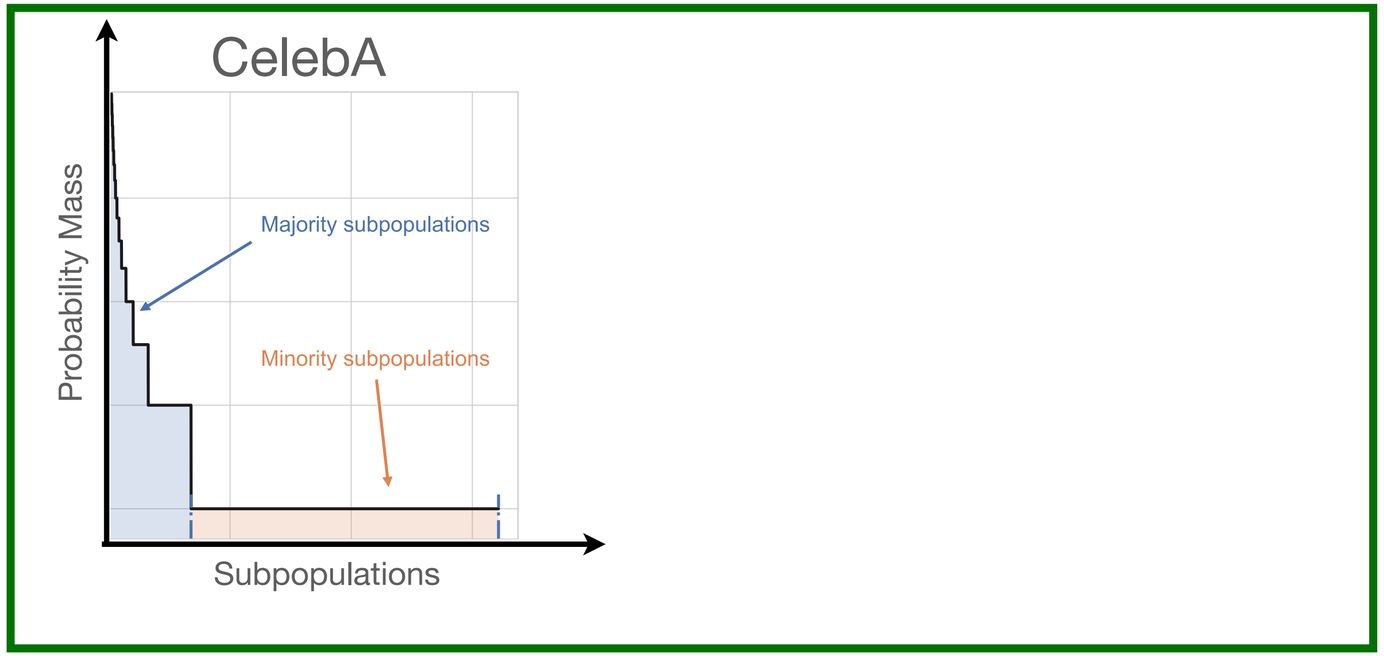

Is this trend systematic?

Main Contribution:
We prove this trend in a model-agnostic setting for long-tailed distribution.

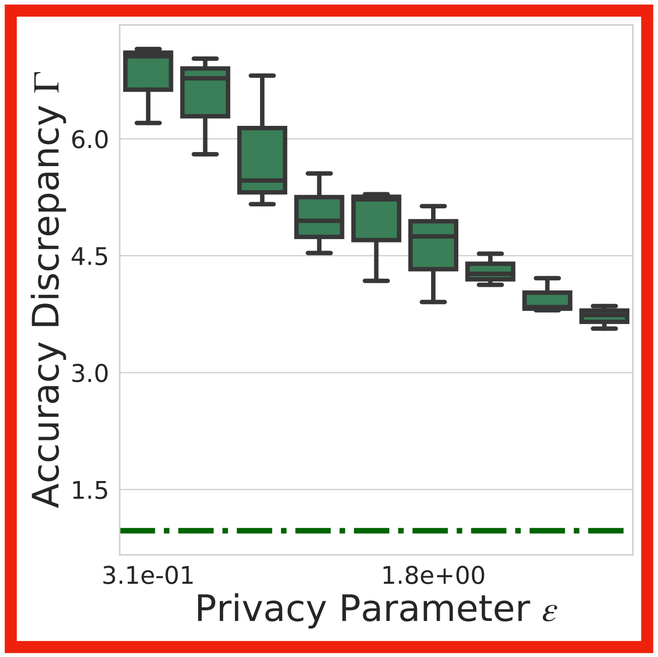

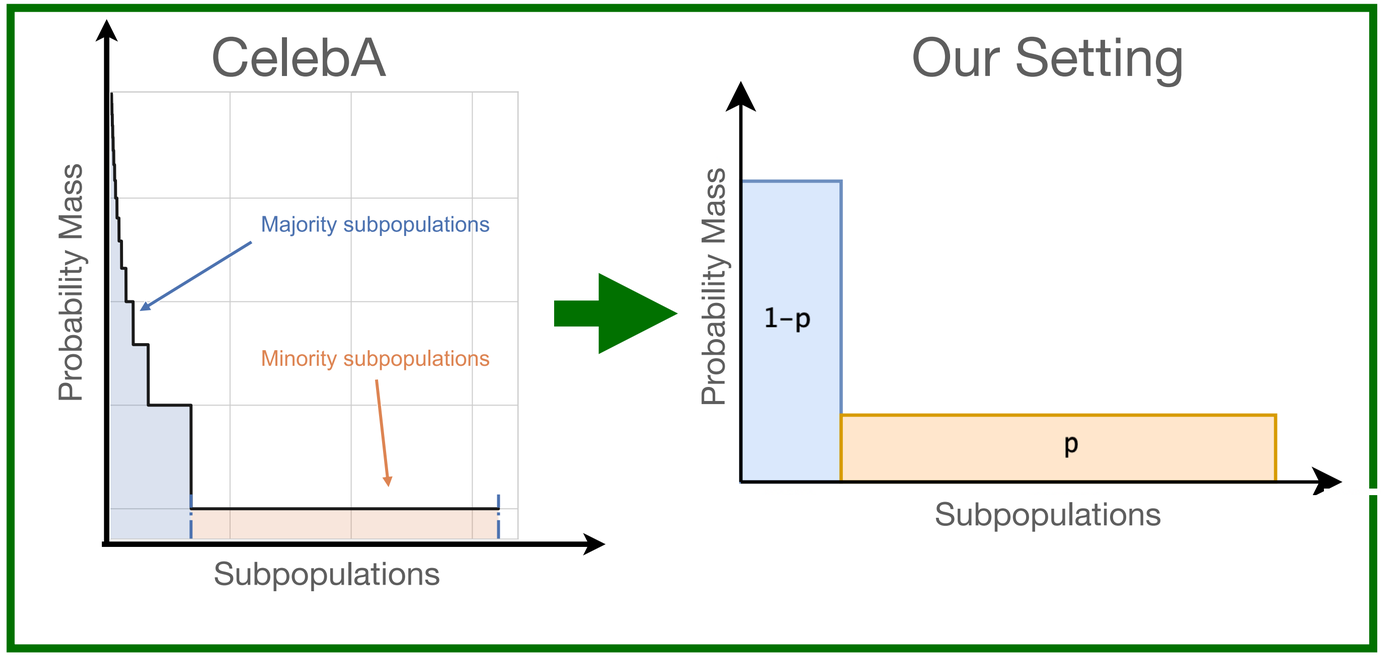

Definitions of error and fairness

Definitions of error and fairness

- Error

Definitions of error and fairness

- Error $\operatorname{err}(A, \Pi, F) =$

Definitions of error and fairness

- Error

$$\mathrm{err}\,(A, \Pi, F) =$$

Learning Algorithm

Definitions of error and fairness

- Error

$$\mathrm{err}\,(A, \Pi, F) =$$

Learning Algorithm

Data Distribution

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) =$$

Learning Algorithm

Data Distribution

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Data Distribution

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Data Distribution

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Data Distribution

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Data Distribution

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Data Distribution

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Data Distribution

Definitions of error and fairness

- Error

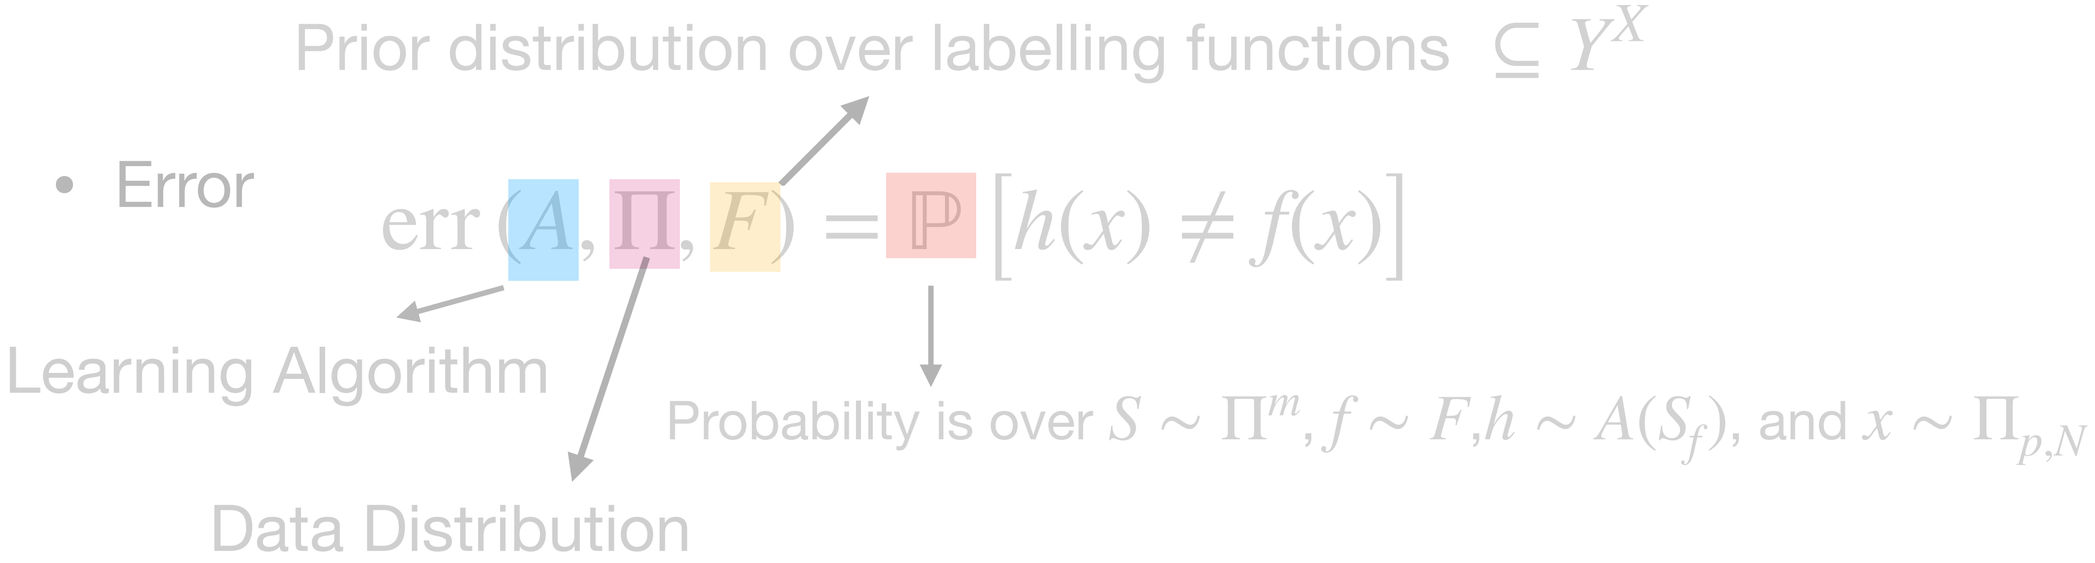

- Accuracy Discrepancy

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\operatorname{err}(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Data Distribution

- Accuracy Discrepancy

$$\Gamma(A, \Pi, F) = \operatorname{err}_{\text{Minority}}(A, \Pi, F) - \operatorname{err}(A, \Pi, F)$$

Definitions of error and fairness

Prior distribution over labelling functions $\subseteq Y^X$

- Error

$$\mathrm{err}\,(A, \Pi, F) = \mathbb{P}\left[h(x) \neq f(x)\right]$$

Learning Algorithm

Probability is over $S \sim \Pi^m, f \sim F, h \sim A(S_f)$, and $x \sim \Pi_{p,N}$

Data Distribution

- Accuracy Discrepancy

Marginalised over minority subpopulations

$$\Gamma\,(A, \Pi, F) = \mathrm{err}_{\mathrm{Minority}}\,(A, \Pi, F) - \mathrm{err}\,(A, \Pi, F)$$

Privacy at the cost of fairness

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

N: # Minority subpopulations
m: # Training points

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

(Minority Subpopulations) Let $\frac{N}{m} \to c$ as $N, m \to \infty$.

N: # Minority subpopulations
m: # Training points

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

(Minority Subpopulations) Let $\frac{N}{m} \to c$ as $N, m \to \infty$.

N: # Minority subpopulations
m: # Training points

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.
- (Long-tailed) Increases with (relative) # of minority subpopulations c.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

(Minority Subpopulations) Let $\frac{N}{m} \to c$ as $N, m \to \infty$.

N: # Minority subpopulations
m: # Training points
F: Label prior

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.
- (Long-tailed) Increases with (relative) # of minority subpopulations c.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

(Minority Subpopulations) Let $\frac{N}{m} \to c$ as $N, m \to \infty$.

(Label prior Entropy) Define $\|F\|_\infty = \max_{x,y} \mathbb{P}_{f \sim F}\left[f(x) = y\right]$

N: # Minority subpopulations
m: # Training points
F: Label prior

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.
- (Long-tailed) Increases with (relative) # of minority subpopulations c.

Privacy at the cost of fairness

Consider any (ϵ, δ)-**DP algorithm** that obtains low error on a long-tailed distribution.

(Minority Subpopulations) Let $\frac{N}{m} \to c$ as $N, m \to \infty$.

(Label prior Entropy) Define $\|F\|_\infty = \max_{x,y} \mathbb{P}_{f \sim F}\left[f(x) = y\right]$

N: # Minority subpopulations
m: # Training points
F: Label prior

(Informal Theorem A) We prove an asymptotic lower bound for accuracy discrepancy which

- (Privacy) Increases with privacy parameter ϵ.
- (Long-tailed) Increases with (relative) # of minority subpopulations c.
- (Label prior) Increases with entropy of the label prior.

Thank you

Subpopulations

Long-tailed data

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Subpopulations

Long-tailed data

- Most datasets follow long-tailed distribution.

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Subpopulations

Long-tailed data

- Most datasets follow long-tailed distribution.
  - With multiple small subpopulations.

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Subpopulations

Long-tailed data

- Most datasets follow long-tailed distribution.
  - With multiple small subpopulations.
  - Individually: small mass.

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Subpopulations

Long-tailed data

- Most datasets follow long-tailed distribution.
  - With multiple small subpopulations.
  - Individually: small mass.
  - Together: sizeable mass.

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Subpopulations

Long-tailed data

- Most datasets follow long-tailed distribution.
  - With multiple small subpopulations.
  - Individually: small mass.
  - Together: sizeable mass.

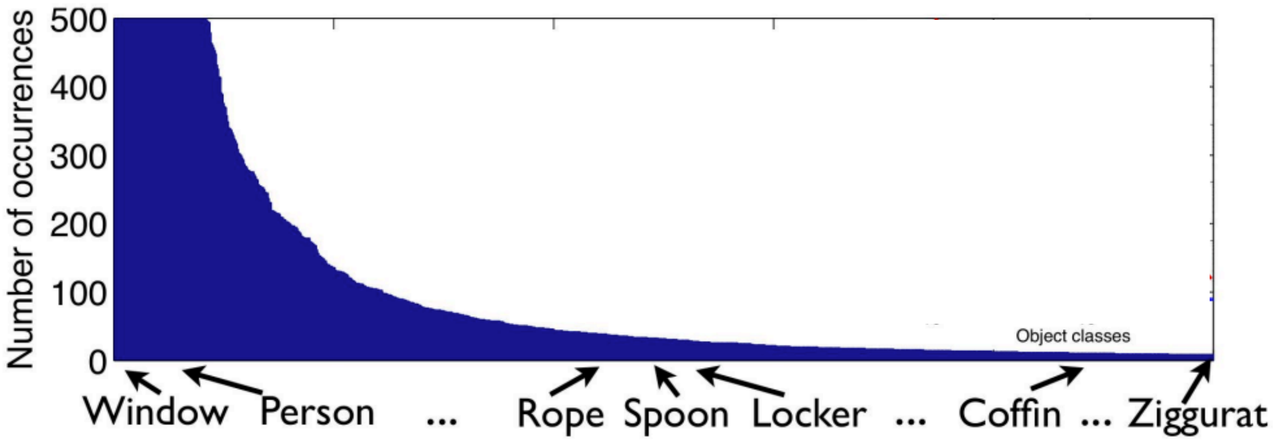

The number of examples by object class in the SUN dataset

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Subpopulations

Long-tailed data

- Most datasets follow long-tailed distribution.
  - With multiple small subpopulations.
  - Individually: small mass.
  - Together: sizeable mass.

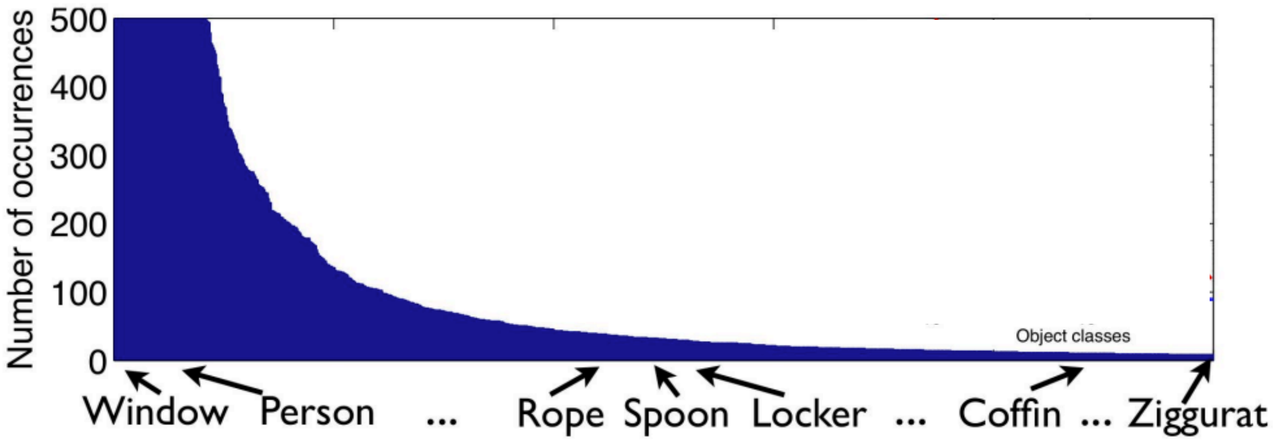

The number of examples by object class in the SUN dataset

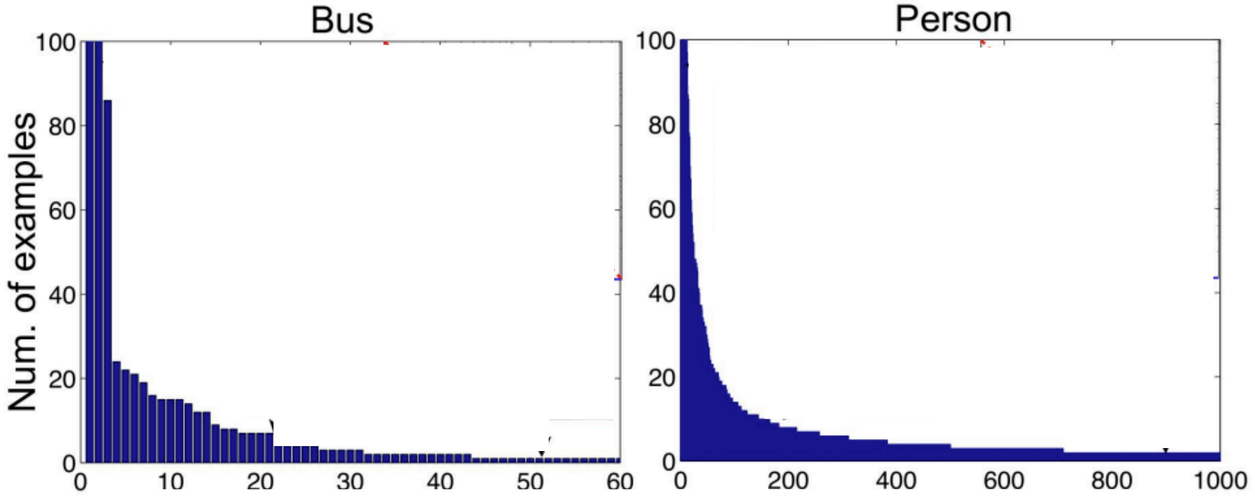

The number of patterns in the PASCAL dataset

X. Zhu, D. Anguelov, and D. Ramanan. “Capturing long-tail distributions of object subcategories”. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition.

Differential Privacy

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

for all subsets Q of the output space,

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

for all subsets Q of the output space,
for all neighboring datasets S_1, S_2.

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

for all subsets Q of the output space,
for all neighboring datasets S_1, S_2.

Neighboring datasets differ in only one entry.

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

- Privacy parameter $\epsilon \in (0, \infty)$

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

for all subsets Q of the output space,

for all neighboring datasets S_1, S_2.

Neighboring datasets differ in only one entry.

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

- Privacy parameter $\epsilon \in (0, \infty)$
- Smaller is more private

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

for all subsets Q of the output space,

for all neighboring datasets S_1, S_2.

Neighboring datasets differ in only one entry.

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

- Privacy parameter $\epsilon \in (0, \infty)$
- Smaller is more private

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

- Probability of catastrophic exposure

for all subsets Q of the output space,

for all neighboring datasets S_1, S_2.

Neighboring datasets differ in only one entry.

Differential Privacy

An algorithm A is (ϵ, δ)-DP if

$$\mathbb{P}\left[A(S_1) \in Q\right] \leq e^{\epsilon}\mathbb{P}\left[A(S_2) \in Q\right] + \delta$$

- Privacy parameter $\epsilon \in (0, \infty)$
- Smaller is more private

- Probability of catastrophic exposure
- Smaller is safer

for all subsets Q of the output space,

for all neighboring datasets S_1, S_2.

Neighboring datasets differ in only one entry.

CIFAR10

With ResNet-18

CIFAR10

With ResNet-18

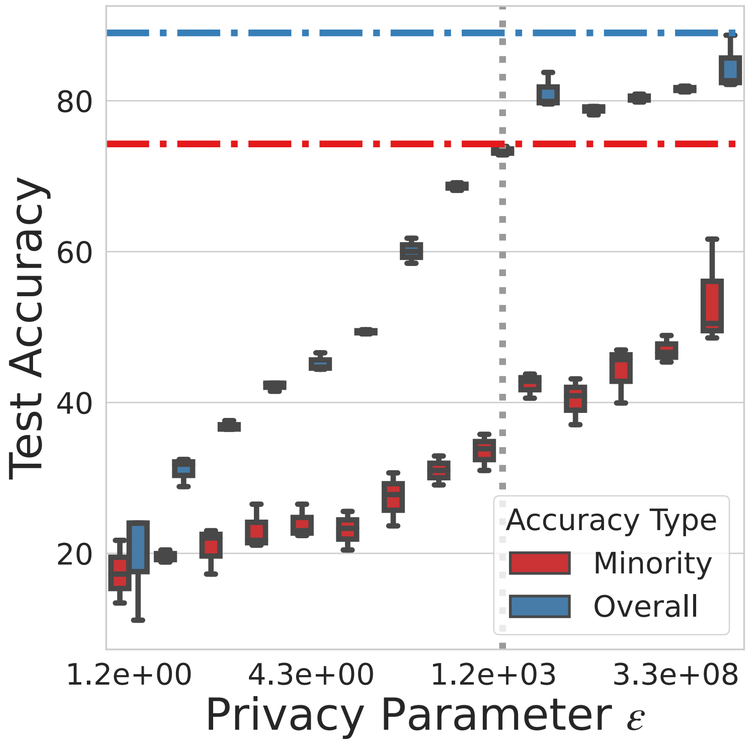

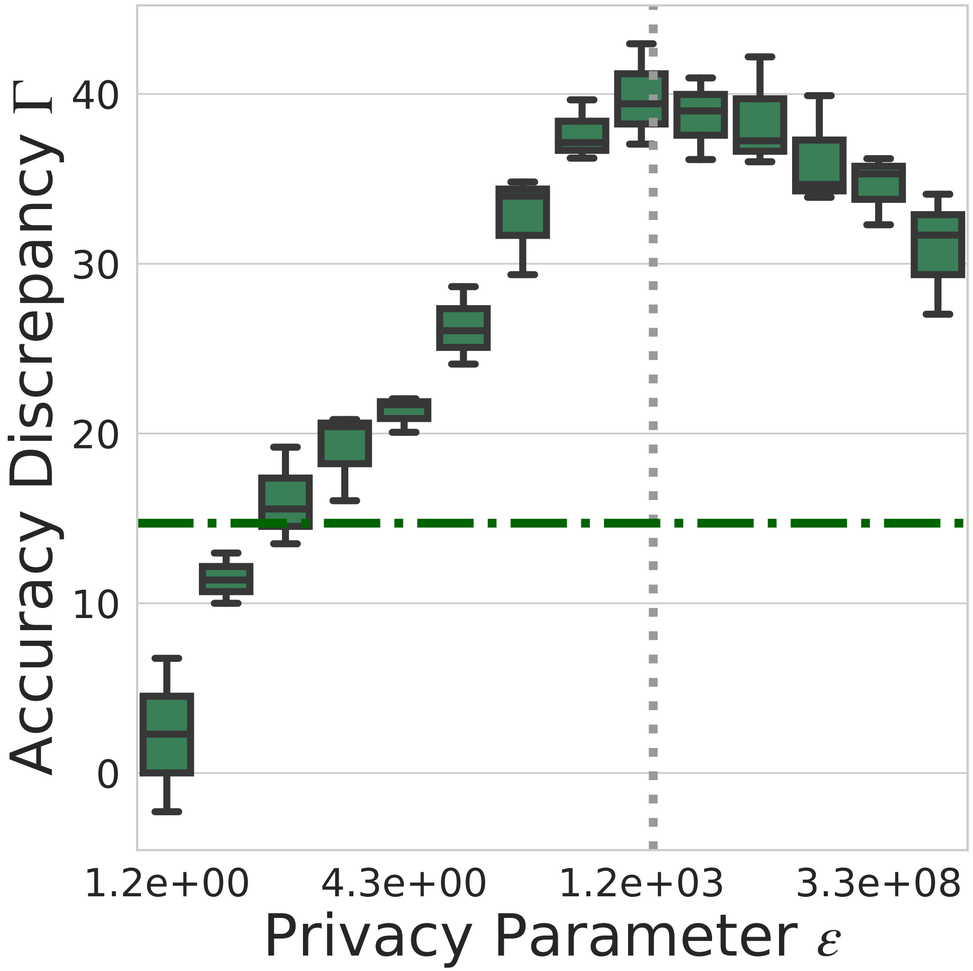

CIFAR10

With ResNet-18

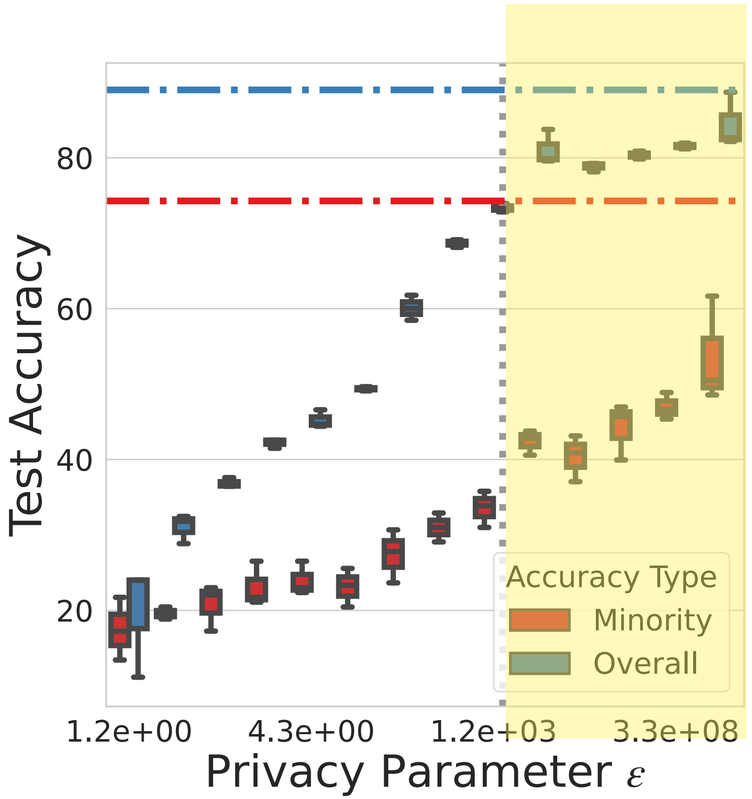

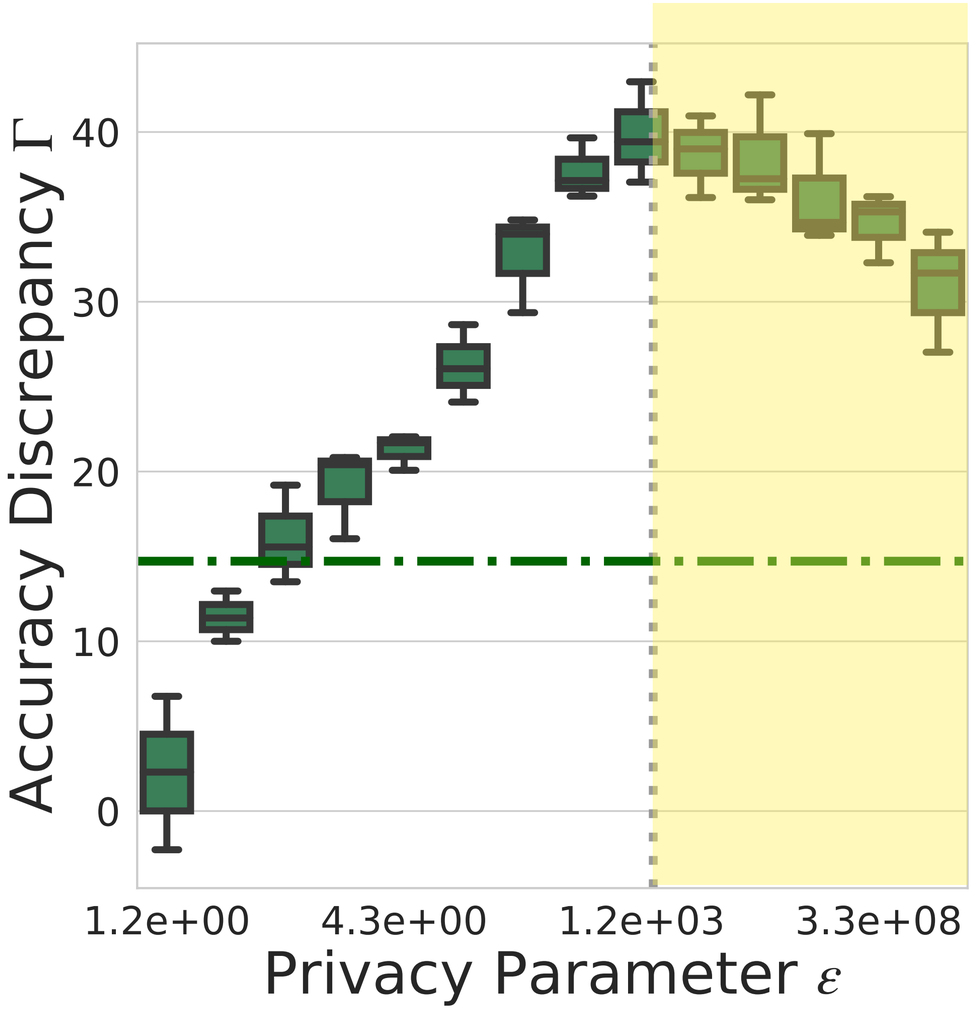

CIFAR10

With ResNet-18

Theorem A

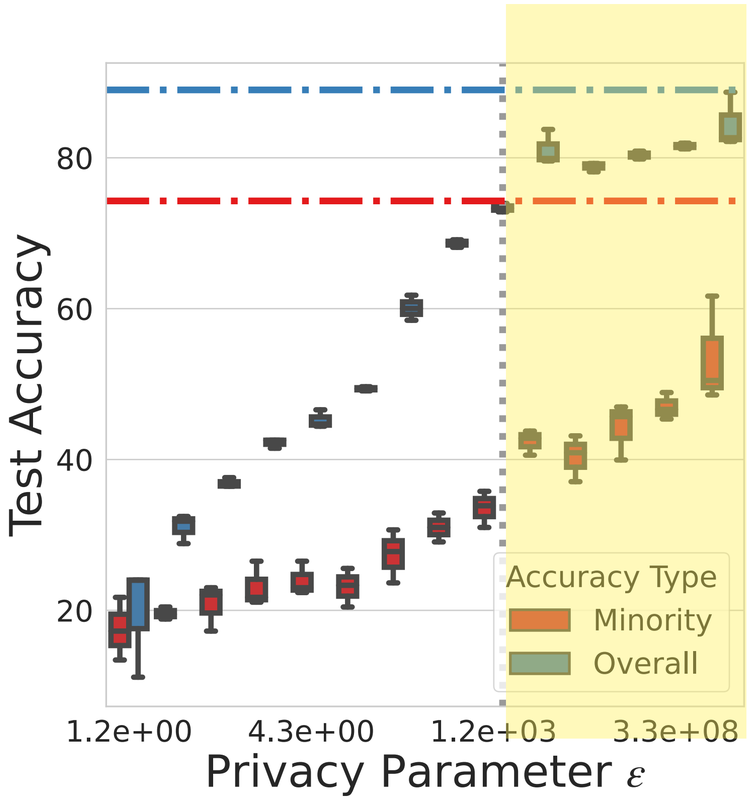

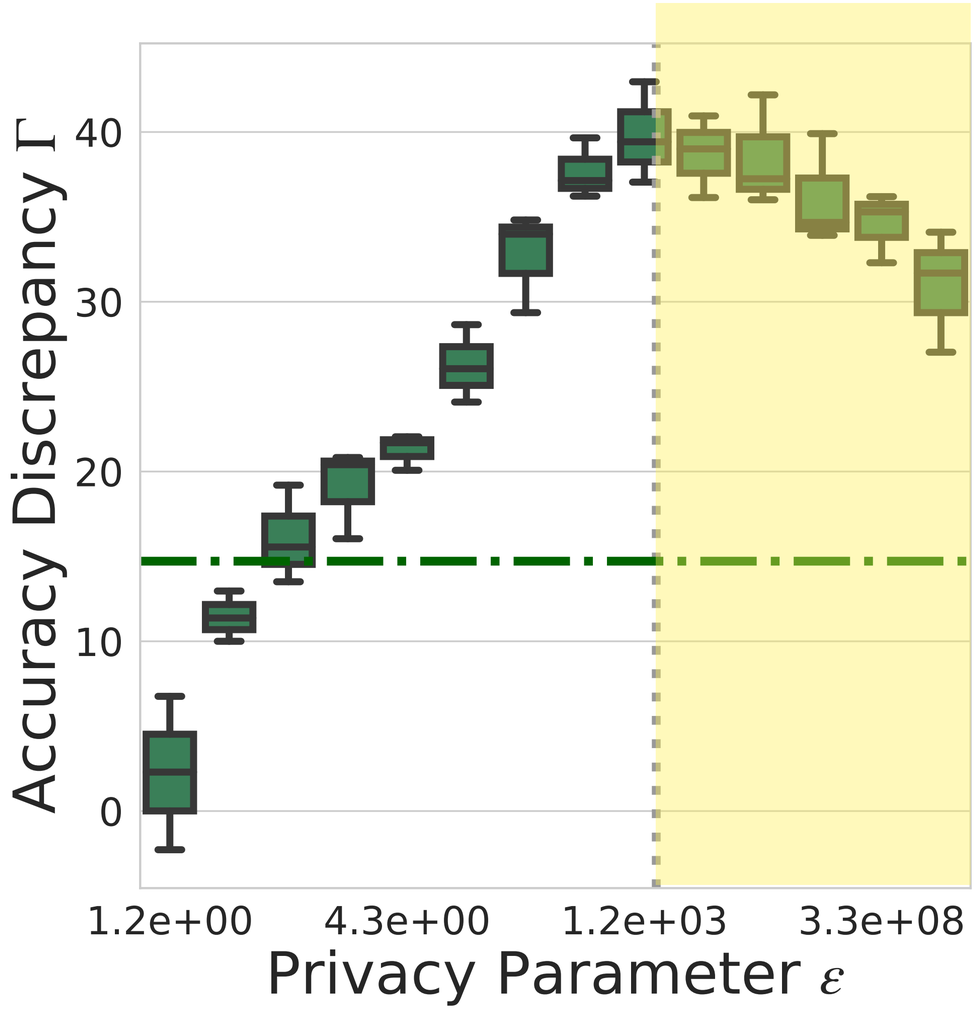

CIFAR10

With ResNet-18

Theorem A

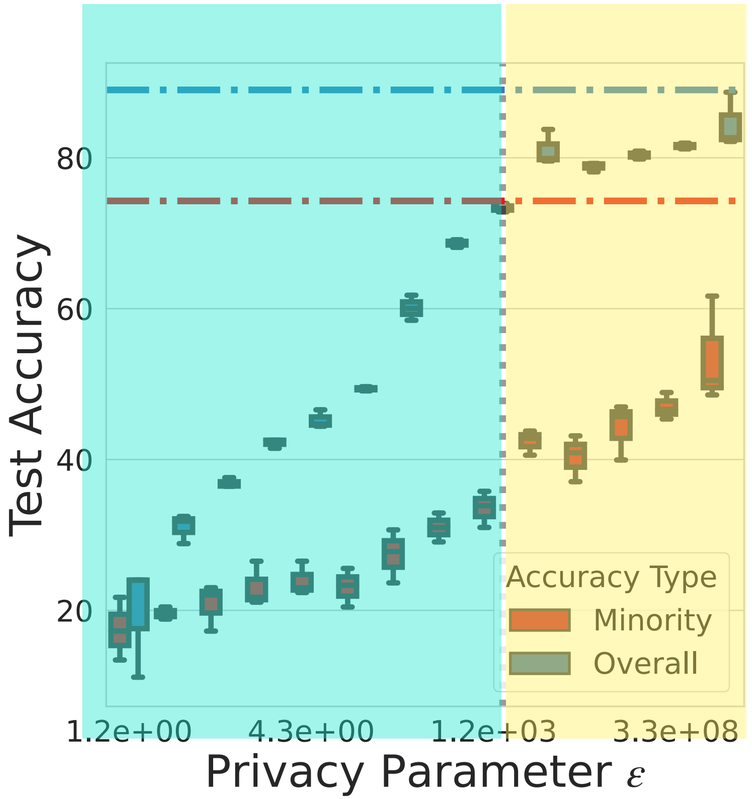

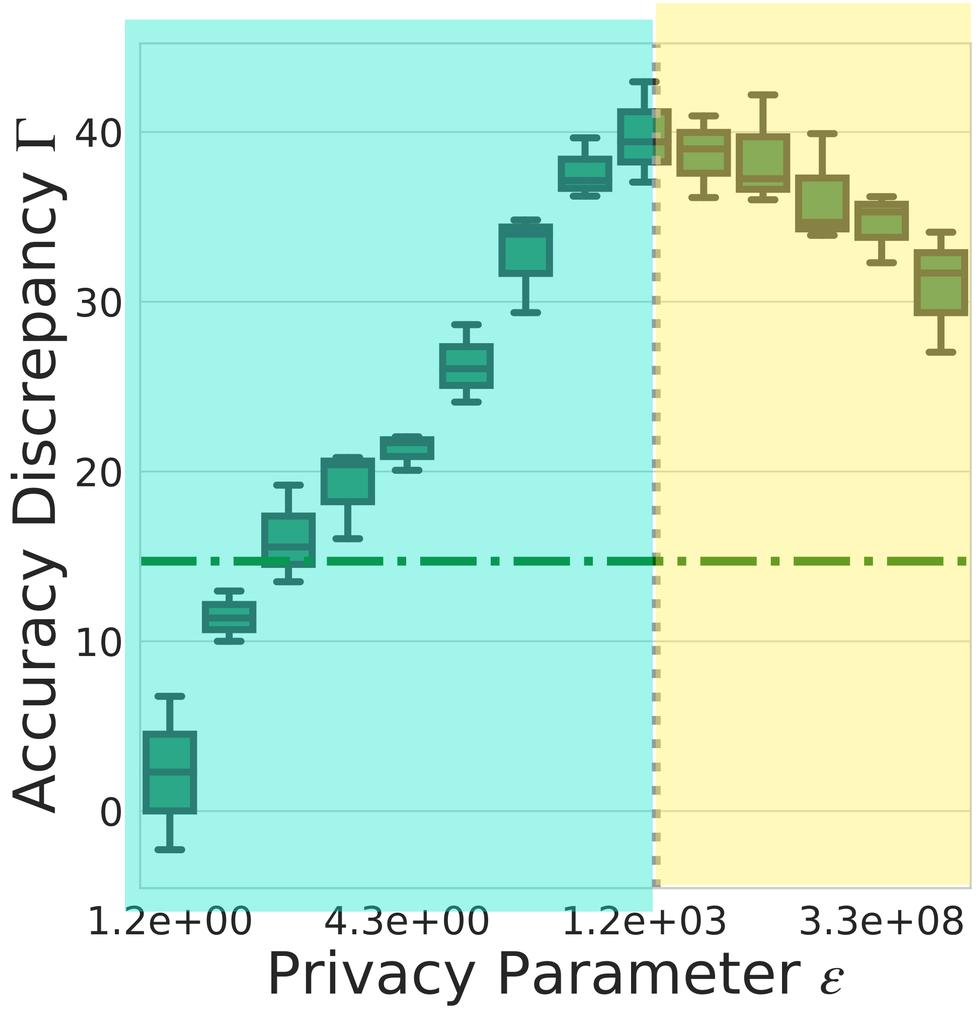

CIFAR10

With ResNet-18

Theorem A

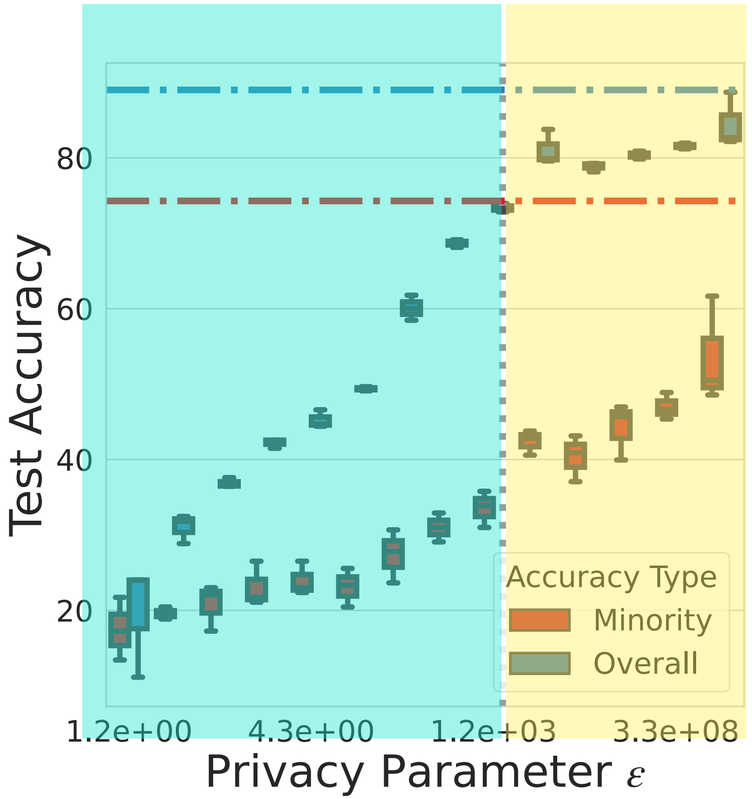

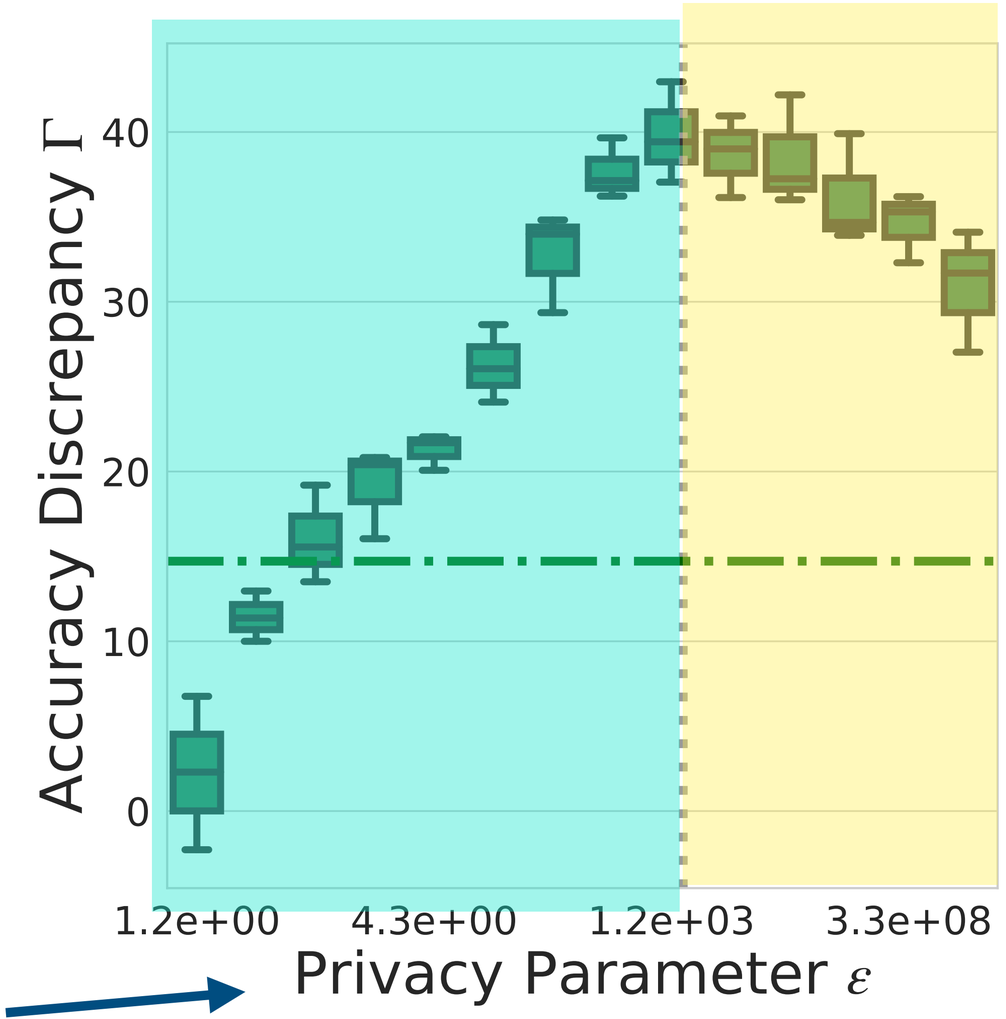

Theorem B

Law School (Tabular)

With Random Forest

Theorem A

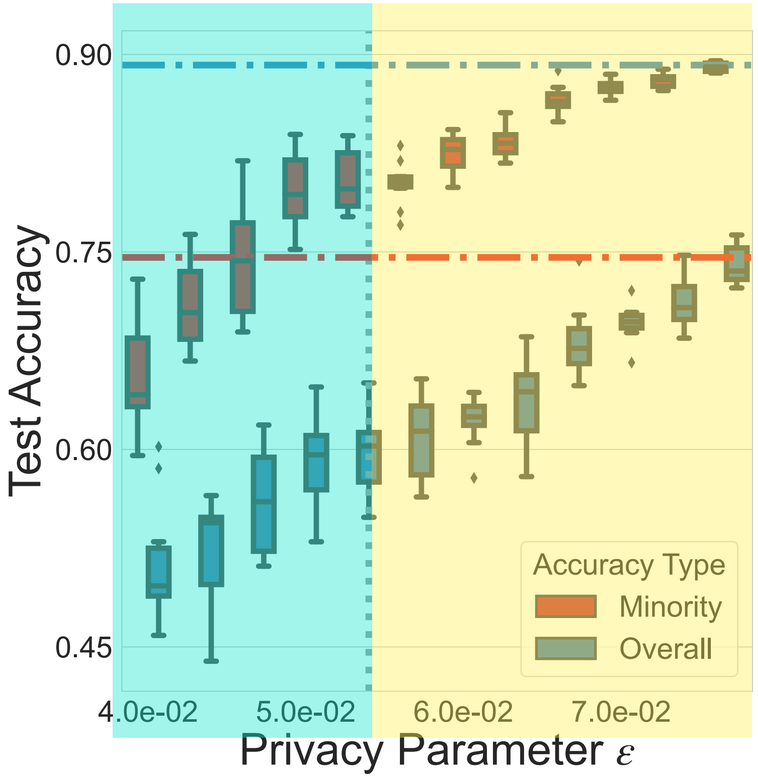

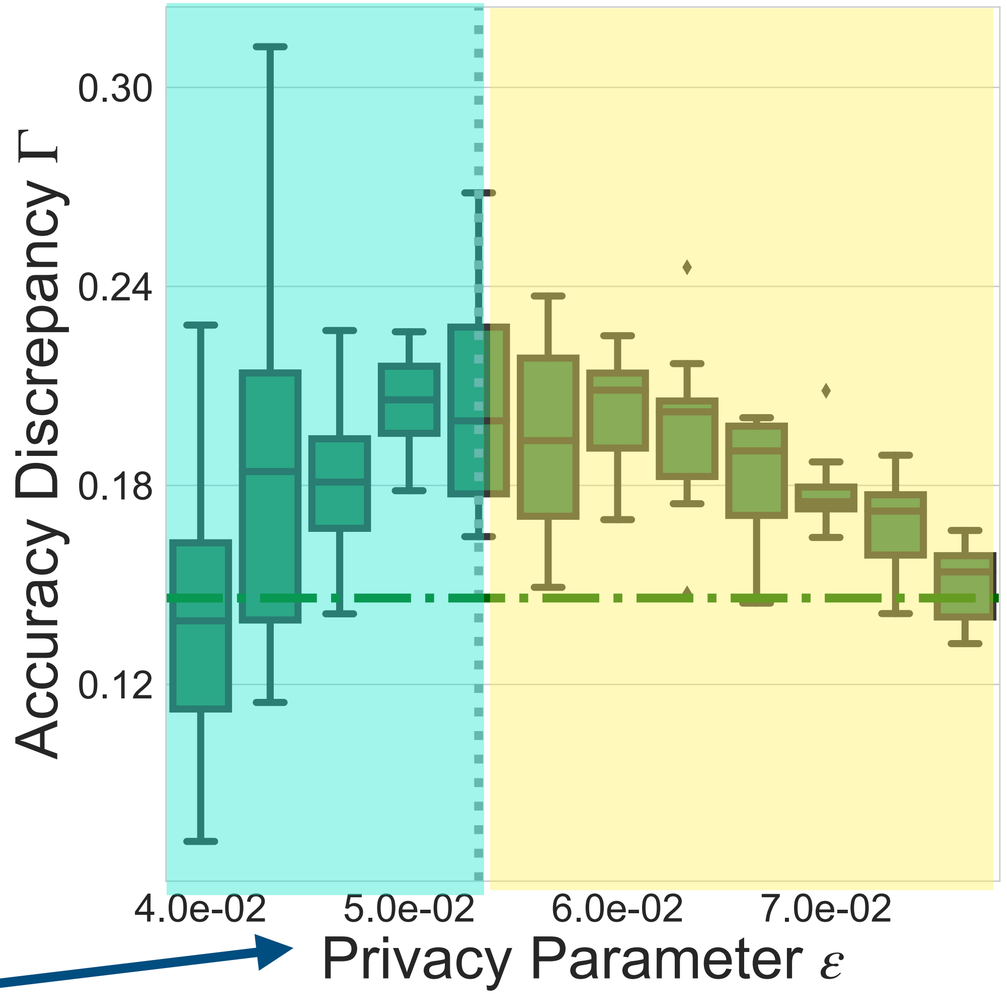

Theorem B

Defining subpopulations