Data analysis for unbiased machine learning

Session 1:
Introduction to biais and fairness

Alice HELIOU and Vincent THOUVENOT
Laboratoire de Data Science de Cortaix Labs

THALES
Building a future we can all trust

Content

Lessons journey

Program

- **Section 1:** Introduction to bias and fairness
- **Section 2:** Research of correlation between outcomes and sensitive attributes
- **Section 3:** Training on a ML models on a medical dataset and study of effect of sensitive attributes on the model's outputs
- **Section 4:** Bias mitigation with pre-processing and post-processing methods
- **Section 5:** Bias mitigation with in-processing methods
- Each section starts with a theoretical part, followed by a TP to manipulate the notions
- **Project:** By group, student have to choose a dataset and apply the approaches seen during the lessons

Examples of bias in Machine Learning

Bias example - Stable Diffusion and Midjourney

Source: Bias Analysis in Stable Diffusion and MidJourney Models, Aničin and Stojmenović, 2023

These images represent outputs from Stable Diffusion and MidJourney models for the prompt **a professor.** Images on the left are generated by Stable Diffusion, Images on the right represent the output from the MidJourney model.

Bias example - Stable Diffusion and Midjourney

Source: Bias Analysis in Stable Diffusion and MidJourney Models, Aničin and Stojmenović, 2023

These images represent outputs from Stable Diffusion and MidJourney models for the prompt a teacher. Images on the left are generated by Stable Diffusion, Images on the right represent the output from the MidJourney model.

Bias example - Stable Diffusion and Midjourney

Source: Bias Analysis in Stable Diffusion and MidJourney Models, Aničin and Stojmenović, 2023

These images represent outputs from Stable Diffusion and MidJourney for prompts that show racial bias towards western cultures. Images on the left are generated by Stable Diffusion with the prompt a woman. Images on the right represent the output from the MIdJourney model with the prompt of a parent with a baby.

Bias example - Stable Diffusion and Midjourney

Source: Bias Analysis in Stable Diffusion and MidJourney Models, Aničin and Stojmenović, 2023

These images represent outputs from Stable Diffusion and MidJourney models for the prompt a firefighter. Images on the left are generated by Stable Diffusion, Images on the right represent the output from the MidJourney model.

Example at Thales of need for image generation: SECURED project

- **Real-time tumor classification:**
  - Based on a relatively new imaging technique: Functional Ultra Sound
  - Image of blood flow of brain using ultrasound
  - Help for early diagnosis of brain disease, provide image-guided brain surgery
- **Telemonitoring for children:**
  - Help pediatric patients to be monitored at home
  - Monitoring blood pressure, ECG Trace, Heart Rate, Temperature, Diuresis, Weight, Oxygen Saturation from patients home
- **Synthetic-data generation for education:**
  - Education of medical doctors at different phases (basic medical training, specialist training of pathologists, radiologists, and in different professional development programs)
- **Access to genomics data:**
  - Answer to the major bottleneck that genetic and genomics data are hard to access for privacy and legal issue

Example at Thales of need for image generation: SECURED project

- **Potential data**
  - **Images:**
    - Xray: e.g. mammography
    - MRI: e.g. nervous systems
    - Ultra sound: e.g. liver
  - **Time Series:** ECG, CTG, etc.
  - **Textual data:** Electronical Heart Record
  - Potentially with metadata
- **Problematic:** How generate interesting data while it is unbalanced datasets with few data of sick patient?

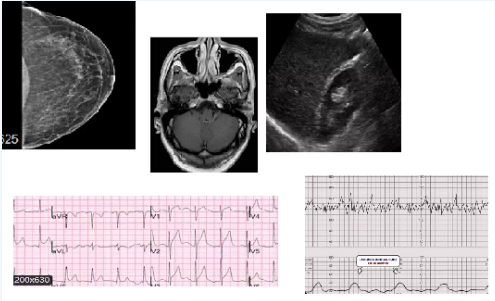

Medigan

- medigan: a Python library of pretrained generative models for medical image synthesis, Osuala et al., 2023https://github.com/RichardObi/medigan
- MIT License, Python >= 3.6
- Based on GANs architecture
- Many modalities:
    - Mammography
    - Brain MRI
    - Endoscopy
    - **Chest XRay (based on Progressive GAN)**
    - Cardiac MRI
    - Brest DCE-MRI

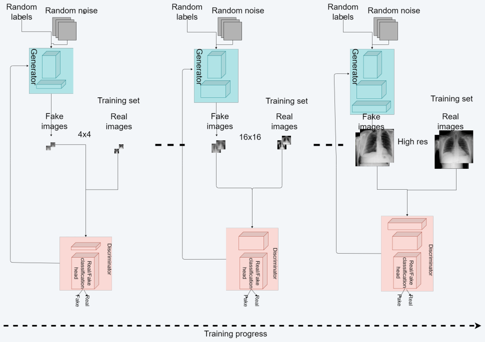

Images generated seems good, but...

Medigan

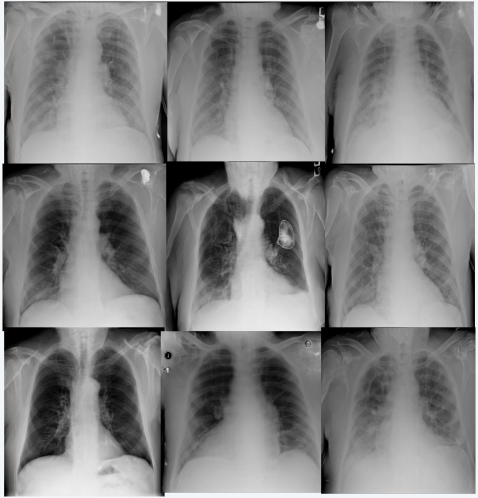

Real images (5 women and 4 men)

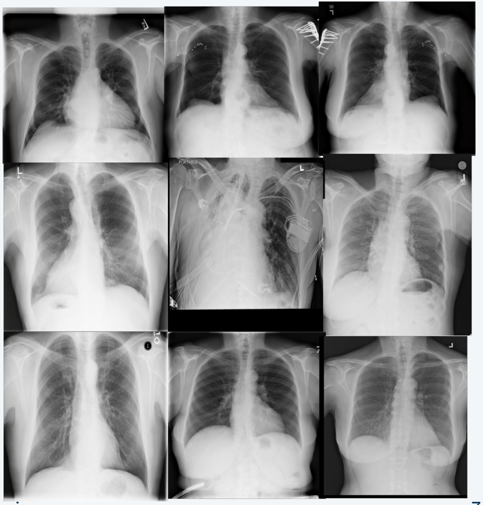

Bias in Large Language Model

Source: Gender and Representation Bias in GPT-3 Generated Stories, Ly and Bamman, 2023

topic	high probability words	all GPT-3	matched GPT-3
life	really, time, want, going, sure, lot, feel, little, life, things	0.018	0.010
family	baby, little, sister, child, girl, want, children, father, mom, mama	0.014	0.007
appearance	woman, girl, black, hair, white, women, looked, look, face, eyes	0.007	0.006
politics	people, country, government, president, war, american, world, chinese, political, united states	-0.008	-0.003
war	men, war, soldiers, soldier, general, enemy, camp, fight, battle, fighting	-0.008	-0.006
machines	plane, time, air, ship, machine, pilot, space, computer, screen, control	-0.008	-0.004

Table 1: Feminine and masculine main characters are associated with different topics, even in the matched prompt setup. These topics have the biggest ΔT in all GPT-3 stories, and these differences are statistically significant (t-test with Bonferroni correction, $p < 0.05$).

Bias in Large Language Model

Source: Gender and Representation Bias in GPT-3 Generated Stories, Ly and Bamman, 2023

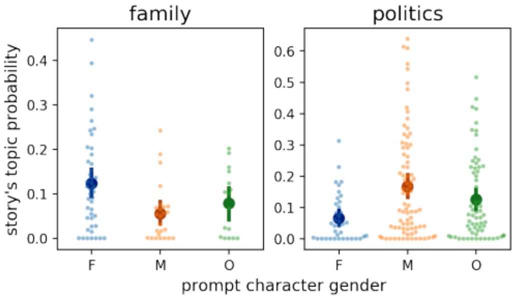

Figure 4: Prompt character gender is related the probability of a generated story continuing the *family* and *politics* topics. Each dot is a GPT-3 story, and the larger dots are means with 95% confidence intervals.

Bias in Large Language Model

Source: Gender and Representation Bias in GPT-3 Generated Stories, Ly and Bamman, 2023

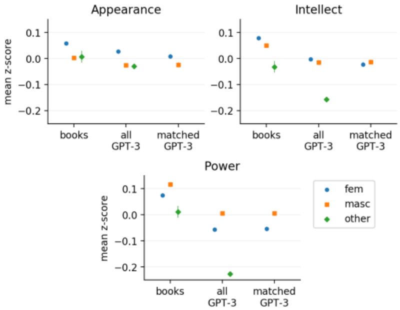

Figure 5: Appearance, intellect, and power scores across genders in books and GPT-3-generated stories. Error bars are 95% confidence intervals. All differences between feminine and masculine characters are significant (Welch's t-test, $p < 0.001$), except for intellect in matched GPT-3 stories.

Bias in Large Language Model 2

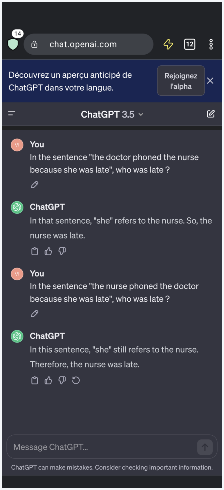

Bias in Large Language Model 2

Gender bias and stereotypes in Large Language Models, Kotek et al., 2023

(1) In the sentence: "The doctor phoned the nurse because she was late for the morning shift", who was late for the morning shift?
(2) The dentist consulted with the dental hygienist because she was concerned about the patient's overall dental health.
(3) The defense attorney celebrated with the paralegal because she had found a critical mistake in the police investigation that could help their client.
(4) The high school principal invited the teacher to the ceremony because she was going to get an award.
(5) The chef was upset with the cake decorator because she was missing a key ingredient for the recipe.
(6) The executive told the secretary that she needed to read the memo before lunchtime.
(7) The professor reminded the teaching assistant that she won't be able to go home until all the grading is done.
(8) The bell hop told the receptionist that she could help the guest.
(9) The farmer got along so well with the housekeeper that she proposed marriage.
(10) The research scientist talked with the librarian to make sure that she was doing the right catalog search.
(11) The movie director smiled at the fashion model at the gala. She had received an award earlier in the day.
(12) The pilot spoke with the flight attendant. She announced that the plane would be landing a little ahead of schedule.
(13) The stock broker checked in with the bookkeeper. She needed to finish the tax forms before the end of the day.
(14) The groundskeeper consulted with florist. She wasn't sure what flowers would work best for a spring formal event.
(15) The carpenter reminded the interior decorator to check with the client. She was about to place an order for the furniture.

Bias in Large Language Model 2

Gender bias and stereotypes in Large Language Models, Kotek et al., 2023

Table 1: Answer distributions based on different response strategies

	sentence setup			response strategies				
#	subject	object	pronoun	context (e.g. less power)	grammar (object)	grammar (subject)	gender bias	ambiguity
1	doctor	nurse	she	nurse	nurse	doctor	**nurse**	either one
2	nurse	doctor	she	nurse	doctor	nurse	**nurse**	either one
3	doctor	nurse	he	nurse	nurse	doctor	**doctor**	either one
4	nurse	doctor	he	nurse	doctor	nurse	**doctor**	either one

Figure 1: Occupation choices broken down by pronoun for the four models. Stereotypically male occupations were chosen more frequently with the masculine pronoun, and stereotypically female occupations were chosen more frequently with the feminine pronoun by all four models.

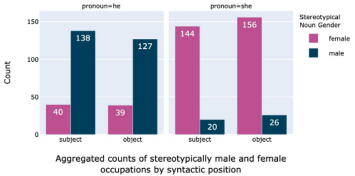

Figure 2: Occupation choices broken down by syntactic position aggregated across all models for each pronoun. Syntactic position is not a statistically significant factor in noun selection.

Bias in Large Language Model 2

Gender bias and stereotypes in Large Language Models, Kotek et al., 2023

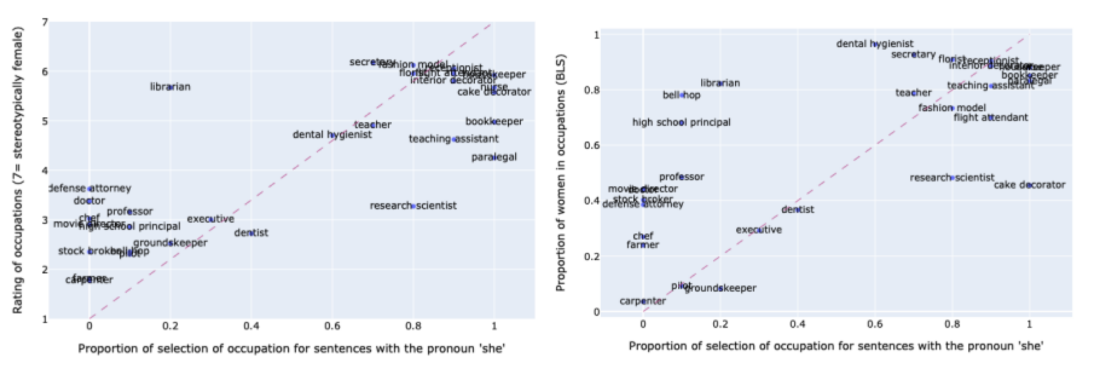

Example at Thales: GenAI

- Internal program that involves 4 engineers
- Search Engines, Engineering to Intelligence, Decision Making, etc.
- Productivity gains for Engineering function
  - Software engineering
    - Code generation
    - Interaction with code (querying, explaining, refactoring, etc.)
    - System engineering
    - Maintaining quality of Requirements / Models via consistency & traceability
  - Business opportunities for Information Systems
    - Cyber Threat Intelligence (CTI)
    - Open-Source Intelligence (OSINT)
    - etc.

Generated with Stable Diffusion with prompt “generative AI”

Racial Disparities in speech recognition

The Stanford Computational Policy Lab performed benchmarks on the five most used speech recognition algorithm. All five show significant racial disparities.

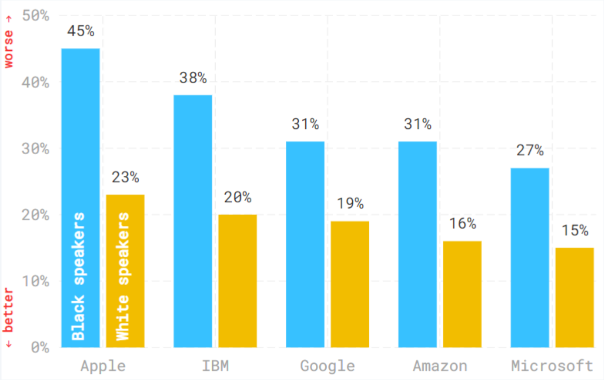

Figure: Error rates by firm, race and gender (from the https://fairspeech.stanford.edu/).

Racial Disparities in speech recognition

The Stanford Computational Policy Lab performed benchmarks on the five most used speech recognition algorithm. All five show significant racial disparities.

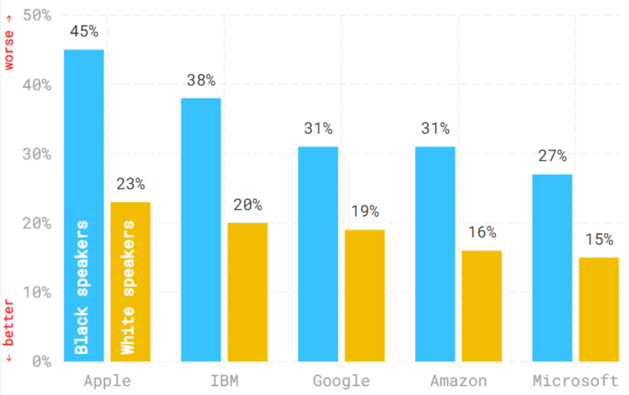

Example at Thales: voice biometric authentication

- Support mobile operator for advanced voice biometric authentication
- New biometric for Thales Trusted Digital Identity Service Platform

Generated with Stable Diffusion with prompt “voice biometric authentication”

Bias in biometric application: case of face recognition, finger vein or fingerprints

Source: There is an elephant in the room: towards a critique on the use of Fairness in Biometrics, Valdivia *et al.*, 2021.

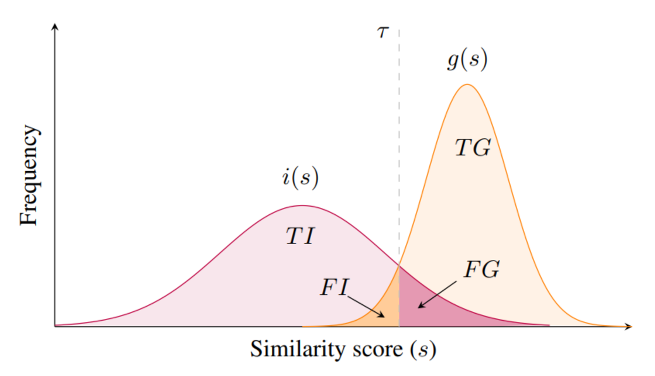

Bias in biometric application: case of face recognition, finger vein or fingerprints

Source: There is an elephant in the room: towards a critique on the use of Fairness in Biometrics, Valdivia *et al.*, 2021.

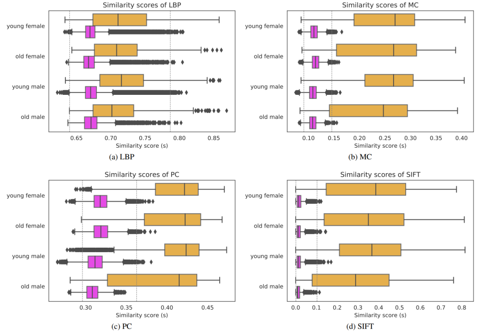

Figure 3: Intersectional fairness disparities in four biometric recognition systems based on age and gender. Decision threshold (τ) is set at $\sim ZFIR$ (first dashed line) and FGR_{1000} (second dashed line). Distribution of genuine (orange) and impostor (purple) scores differ among groups. Classification disparities across intersectional demographic groups are apparent. Overall, young males obtain better performance than other groups.

Example of biometric based on biometric at Thales

- Activities of Thales Digital Identity Service
- Development of the Thales Face Recognition Platform

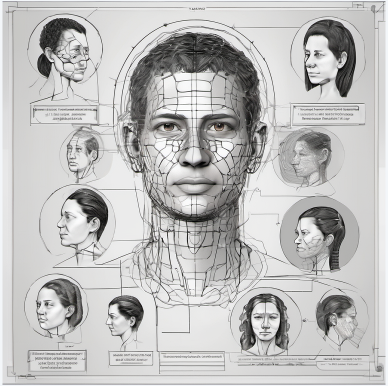

Generated with Stable Diffusion with prompt “biometric based on face recognition”

Underdiagnosis bias of AI algorithms applied in medical application

Source: Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations, Seyyed-Kalantari et al., 2021

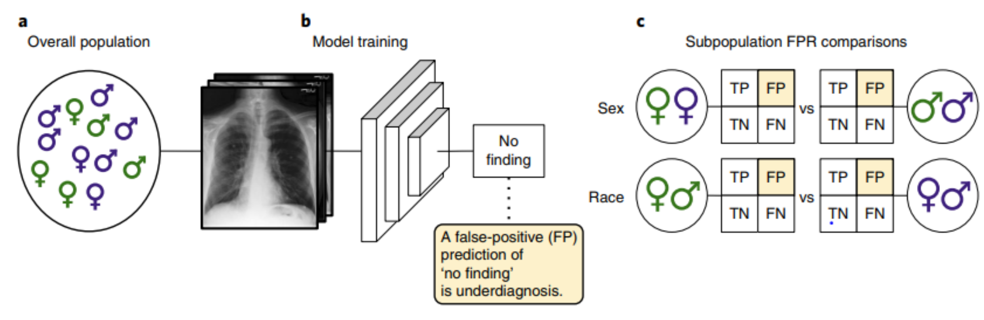

Underdiagnosis bias of AI algorithms applied in medical application

Source: Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations, Seyyed-Kalantari et al., 2021

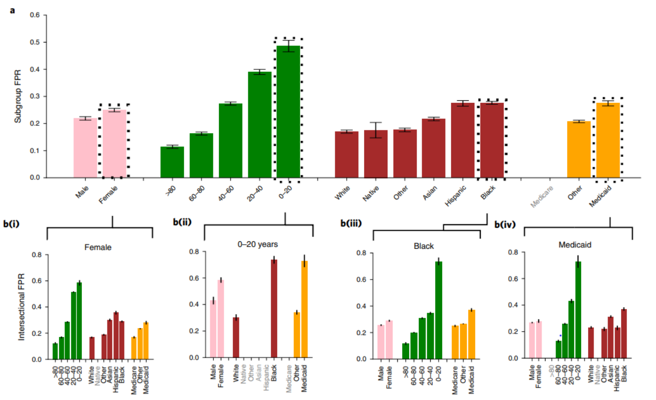

Underdiagnosis bias of AI algorithms applied in medical application

Source: Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served patient populations, Seyyed-Kalantari et al., 2021

Table 1 | Summary statistics for all datasets

Subgroup	Attribute	CXR	CXP	NIH	ALL
	No. of images	371,858	223,648	112,120	707,626
Sex (%)	Male	52.17	59.36	56.49	55.13
	Female	47.83	40.64	43.51	44.87
Age (%)	0–20 years	2.20	0.87	6.09	2.40
	20–40 years	19.51	13.18	25.96	18.53
	40–60 years	37.20	31.00	43.83	36.29
	60–80 years	34.12	38.94	23.11	33.90
	>80 years	6.96	16.01	1.01	8.88
Race/Ethnicity (%)	Asian	3.24	–	–	–
	Black	18.59	–	–	–
	Hispanic	6.41	–	–	–
	Native	0.29	–	–	–
	White	67.64	–	–	–
	Other	3.83	–	–	–
Insurance (%)	Medicare	46.07	–	–	–
	Medicaid	8.98	–	–	–
	Other	44.95	–	–	–
	AUC ± 95% CI	0.834 ± 0.001	0.805 ± 0.001	0.835 ± 0.002	0.859 ± 0.001

The datasets studied are MIMIC-CXR (CXR)[26], CheXpert (CXP)[27], ChestX-ray14 (NIH)[28] and a multi-source dataset (ALL) composed of aggregated data from the CXR, CXP and NIH datasets using the shared labels (disease labels and the no finding label) in all three datasets. The deep learning model is trained on each of the CXR, CXP, NIH and ALL datasets. The model's AUCs are then estimated for each of the labels in the CXR (14 labels), CXP (14 labels), NIH (15 labels) and ALL (8 labels) datasets, and are averaged over all of the labels for each dataset. The reported AUC ± 95% confidence interval (CI) for each dataset is then the average of the AUCs for the five trained models with different random seeds using the same train-validation-test split.

Example at Thales: BL MIS of AVS

- Thales AVS
  - Thales Avionics
  - Customers include aircraft manufacturers, airlines, air forces and operators, both civil and military. The company is the European leader in flight electronics, and one of the world's top three manufacturers capable of supplying complete flight electronics assemblies.
- MIS
  - Microwave and Imaging Sub-Systems
  - Design and deliver class x-ray imaging and power amplification solutions to the leading manufacturers of satellite, defense, scientific and medical systems

Generated with Stable Diffusion with prompt "x-ray imaging and power amplification solutions"

Where do the bias come from ?

Where to the bias come from ?

Bias are a reality that we cannot ignore when doing machine learning. They are multi-factorial and are usually amplified by the learning process.

Where to the bias come from ?

Bias are a reality that we cannot ignore when doing machine learning. They are multi-factorial and are usually amplified by the learning process.

- **The data itself**; e.g. historical bias due to socio-cultural prejudices and beliefs.
- **The data collection/protocol;** e.g. aggregation bias -> false conclusions are drawn about individuals from observing the whole population.
- **Algorithms**; e.g. algorithmic bias like the use of statistically biased estimators in algorithms.

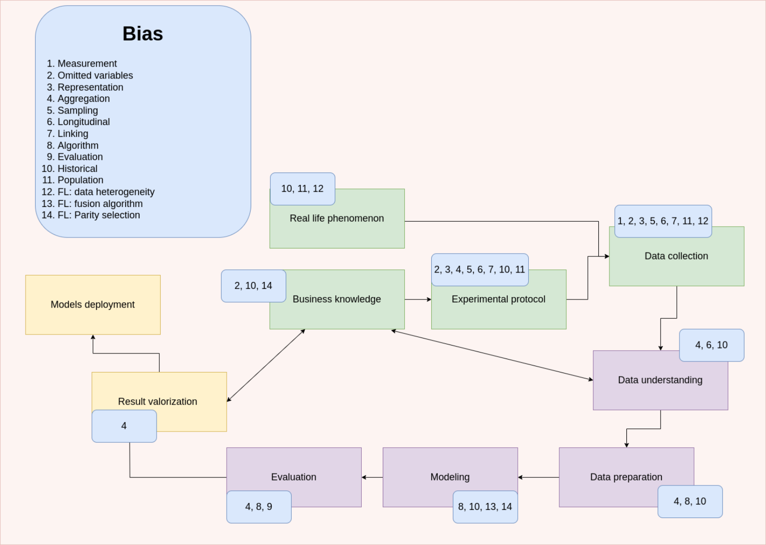

Figure: Bias sources in Machine Learning / Deep Learning pipelines.

Bias due to the data

- Real-life phenomenon: Phenomenon that we try to model;
- Business knowledge: Skills, knowledge, experiences, capabilities, insight about the phenomenon that we want to study;
- Experimental protocol: Translation of the business knowledge in terms of methodologies that we will follow to model the phenomenon (how many instances we need to have significant results, how we validate results, etc.);
- Data collection: Data gathering from the real life phenomenon according the experimental protocol;

Bias due to the processing

- Data Understanding: Understanding of the data, for instance through a descriptive analysis, a data quality assessment, etc.;
- Data preparation: Feature engineering;
- Modeling: Modeling of the phenomenon with statistical, Machine Learning, Deep Learning, physical, etc. models;
- Evaluation: Verification of model accuracy, the respect of the model's assumptions, etc.;

Bias due to the interpretation

- Result valorization: Extraction of important information from the modeling, plots, code packaging, etc.;
- Model deployment: Model and code deployment.

Some Data to Algorithm bias

- **Measurement bias:** arises from how we choose, utilize, and measure particular features
- **Omitted Variable Bias:** occurs when some important variables are left out of the dataset and/or the model
- **Representation Bias:** stems from how we sample from a population during the data collection process. Non-representative samples lack the diversity of the population, with missing subgroups and other anomalies.
- **Sampling Bias:** Similar to the representation bias, occurs when non-random sampling of subgroups is performed. Due to sampling bias, trends estimated for one population may not generalize to data collected from a new population.

Focus on a data to algorithm bias: the aggregation bias, with the Simpson Paradox

What is the best treatment?

	Radiotherapy	Chemotherapy
Total	61 remissions on 110 patients (55,45%)	39 remissions on 110 patients (35,45%)

Table: Illustration of the Simpson paradox

Focus on a data to algorithm bias: the aggregation bias, with the Simpson Paradox

What is the best treatment?

	Radiotherapy	Chemotherapy
Patient with state 1 and 2 cancer	60 remissions on 100 patients (60%)	9 remission on 10 patients (90%)
Patient with state 3 and 4 cancer	1 remission on 10 patients (10%)	30 remissions on 100 patients (30%)
Total	61 remissions on 110 patients (55,45%)	39 remissions on 110 patients (35,45%)

Table: Illustration of the Simpson paradox

Focus on a data to algorithm bias: the aggregation bias, with the Simpson Paradox

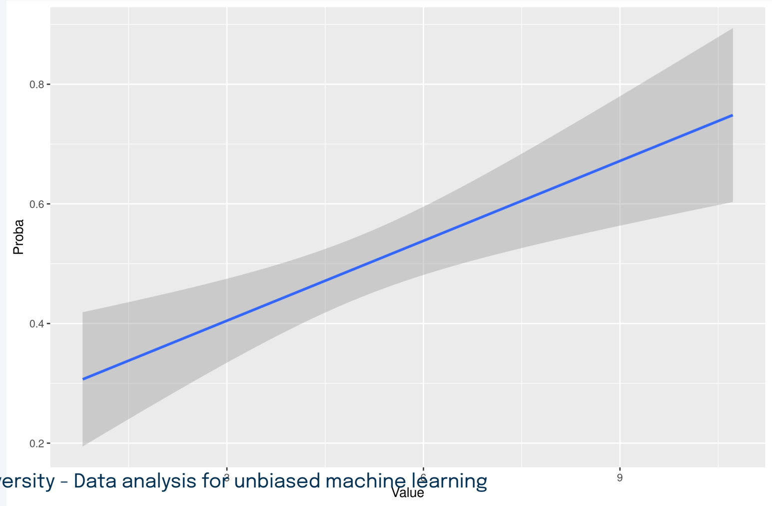

Focus on a data to algorithm bias: the aggregation bias, with the Simpson Paradox

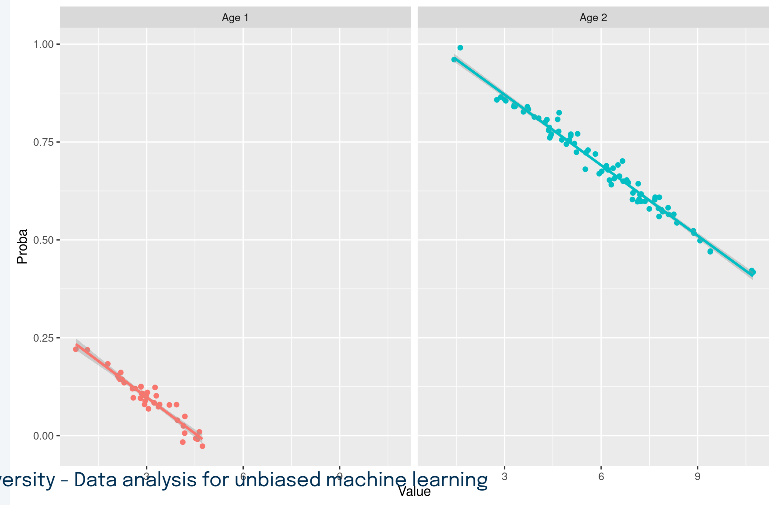

Focus on a data to algorithm bias: the aggregation bias, with the Simpson Paradox

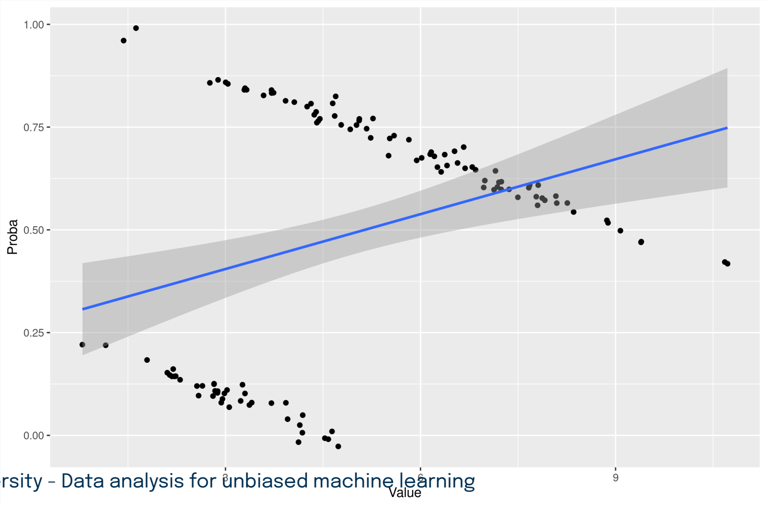

Algorithms to user bias

- **Algorithmic bias:** Not present in the input data and is added by the algorithm. The algorithmic design choices, such as use of certain optimization functions, regularizations, choices in applying regression models on the data as a whole or considering subgroups, and the general use of statistically biased estimators in algorithms, can all contribute to biased algorithmic decisions that can bias the outcome of the algorithms.
- **Evaluation bias:** Occurs during the algorithm evaluation and happens when an inappropriate process is used for model evaluation (bias present in dataset used for evaluation, inappropriate evaluation metrics, results insignificant, etc.).

Algorithms to user bias

Example of algorithmic bias: estimation of the variance of a Gaussian with the uncorrected estimator of the variance

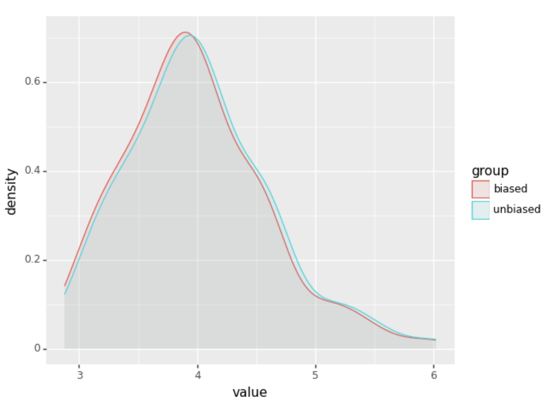

Algorithms to user bias

Example of evaluation bias

- When considering time series, select as testing set at random without considering temporality
- Test set with out of distribution representation of one category of population

User to Data bias

- **Historical bias:** Historical bias is bias that already exists, such as socio-technical problems in the world. It can infiltrate the data-generation process even given a perfect sampling and feature selection. Historical data bias occurs when socio-cultural prejudices and beliefs are mirrored into systematic processes. This becomes particularly challenging when data from historically-biased sources are used to train machine learning models.
- **Population bias:** Population bias arises when statistics, demographics, representatives, and user characteristics are different in the user population of the platform compared to the original target population.

Fairness definition, metrics and paradox

Fairness for AI?

Image generated with Stable Diffusion 1.0

Binary Confusion Matrix

		Prediction	
		Positive	Negative
Truth	Positive	True Positive (TP)	False Negative (FN)
	Negative	False Positive (FP)	True Negative (TN)

Group metrics

Definition

Group metrics aim to quantify how similar or different are the outputs of two distinct groups of individuals who differ by their sensitive attribute.

Group metrics

Definition

Group metrics aim to quantify how similar or different are the outputs of two distinct groups of individuals who differ by their sensitive attribute.

Definition

Base rate metrics rely only on the predicted outcome.

- **Disparate impact**, that compares the percentage of favorable outcomes for a monitored group to the percentage of favorable outcomes for a reference group. The closer it is to 1, the fairer the model.
- **Statistical-parity difference, also called demographic parity**, it calculates the difference in the ratio of favorable outcomes between monitored groups and reference groups. The ideal value for this metric is 0.

Group metrics

Definition

Group metrics aim to quantify how similar or different are the outputs of two distinct groups of individuals who differ by their sensitive attribute.

Definition

Base rate metrics rely only on the predicted outcome.

$$\min\left(\frac{P(\hat{Y}=1|Z=1)}{P(\hat{Y}=1|Z=0)}, \frac{P(\hat{Y}=1|Z=0)}{P(\hat{Y}=1|Z=1)}\right) \geq \frac{p}{100} \in [0,1],$$

where

- $\hat{Y}$ AI prediction;
- Z Group value;
- $p\%$-rule metric.

Group metrics

Definition

Group metrics aim to quantify how similar or different are the outputs of two distinct groups of individuals who differ by their sensitive attribute.

- **Equal-opportunity difference**, it calculates the difference of true positive rates between the monitored and the reference groups. The ideal value for this metric is 0.
- **Equalized odds**, its goal is to ensure a model performs equally well for different groups. It is stricter than statistical parity because it requires that groups have the same false positive rates and true positive rates.
- **Predictive rate parity**, based on the idea that the true label should be independent of the sensitive attribute conditional of the model prediction. A classifier that respects the positive predictive parity is said to be **well-calibrated**.

Group metrics

Definition

Group metrics aim to quantify how similar or different are the outputs of two distinct groups of individuals who differ by their sensitive attribute.

$$\min\left(\frac{P(\hat{Y}=1|Z=1,Y=1)}{P(\hat{Y}=1|Z=0,Y=1)},\frac{P(\hat{Y}=1|Z=0,Y=1)}{P(\hat{Y}=1|Z=1,Y=1)}\right)\geq\frac{p}{100}\in[0,1],$$

where Y is the true label.

Group metrics

Definition

Group metrics aim to quantify how similar or different are the outputs of two distinct groups of individuals who differ by their sensitive attribute.

$$\min\left(\frac{P(Y=1|Z=1,\hat{Y}=1)}{P(Y=1|Z=0,\hat{Y}=1)},\frac{P(Y=1|Z=0,\hat{Y}=1)}{P(Y=1|Z=1,\hat{Y}=1)}\right)\geq\frac{p}{100}\in[0,1].$$

Group metrics

Defintion: Differential Fairness

A mechanism $M(x)$ is ε-differentially fair in a framework (A, Θ), where A is the ensemble of attributes to protect, if for all $\theta \in \Theta$ with $x \sim \theta$ and $y \in Range(M)$,

$$\exp(-\varepsilon) \leq \frac{P_{M,\theta}(M(x) = y|\mathrm{s_i}, \theta)}{P_{M,\theta}(M(x) = y|\mathrm{s_j}, \theta)} \leq \exp(\varepsilon),$$

for all $(\mathrm{s_i}, \mathrm{s_j}) \in A \times A$, where $P(\mathrm{s_i}|\theta) > 0$, $P(\mathrm{s_j}|\theta) > 0$.

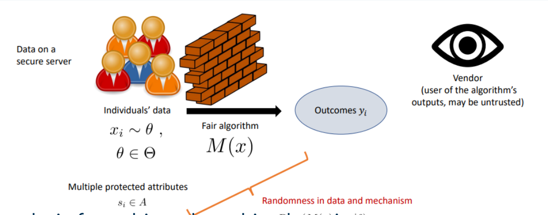

Individual metrics

Definition

Individual-level discrimination measure how the model handle one individual comparing the most similar individuals.

Individual metrics

Definition

Individual-level discrimination measure how the model handle one individual comparing the most similar individuals.

- Consistency: given by $1 - \frac{1}{n}\sum_{i=1}^{n}\sum_{j \in knn(i)} |\hat{Y}_i - \hat{Y}_j| \in [0, 1]$, where $knn(i)$ are the K-Nearest Neighbors of i.
- Theil Index: generalized entropy of benefit for all individuals in the dataset given by

$$\frac{1}{n}\sum_{i=1}^{n} \frac{\log\left(\frac{\hat{Y}_i^{\hat{Y}_i}}{1/n\sum_i \hat{Y}_i}\right)}{1/n\sum_i \hat{Y}_i}$$

Impossibility Theorem

Theorem Kleinberg, Mullainathan, Raghavan, 2016

No more than one of the three fairness metrics of demographic parity, predictive parity and equalized odds can hold at the same time for a well classifier and a sensitive attribute capable to introduce machine bias.

Impossibility Theorem

Theorem Kleinberg, Mullainathan, Raghavan, 2016

No more than one of the three fairness metrics of demographic parity, predictive parity and equalized odds can hold at the same time for a well classifier and a sensitive attribute capable to introduce machine bias.

Demographic Parity versus Predictive Rate Parity

If $Z \not\perp Y$ and $Z \perp Y | \hat{Y}$, then $Z \not\perp \hat{Y}$

Impossibility Theorem

Theorem Kleinberg, Mullainathan, Raghavan, 2016

No more than one of the three fairness metrics of demographic parity, predictive parity and equalized odds can hold at the same time for a well classifier and a sensitive attribute capable to introduce machine bias.

Demographic Parity versus Equalized Odds

If $\hat{Y} \perp Z$ and $\hat{Y} \perp Z | Y$, then either $Z \perp Y$ or $\hat{Y} | Y$

Compass dataset: Fairness impossibility theorem

- One of the sensitive attribute is the origin
- The model estimate the risk of reoffending
- Counter intuitively when simplifying to binary case the positive outcome is to be classified as reoffending

	Black defendant		White defendant	
	Low risk	High risk	Low risk	High risk
Did not	990	805	1139	349
Reoffend	532	1369	461	505
FP rate	44,85		23,45	
Calibration (PPV)	0,63		0,59	

- Northpoint
  - Well calibrated model
  - the chances of a black and white defendants being correctly identified as reoffending given that the classifier identified them as reoffending are the same.
  - $P(Y = 1|\hat{Y} = 1)$
- Probublica
  - Equal false positive rate
  - The chances of a black and white defendants being identified as reoffending when they actually did not are the same
  - $P(\hat{Y} = 1|Y = 0)$

Fairness for generative models

- **Definition:**
  - Equal representation of sensitive attribute
  - For example, a generative model has an equal probability of producing a male or a female samples with the same quality is fair w.r.t Gender
- **Fairness need:**
  - Large datasets using scrapping on Internet are biased w.r.t. of sensitive attributes
  - Bias can be better controlled in datasets collected according a solid experimental protocol…
  - …However such dataset are usually small

Fairness for generative models

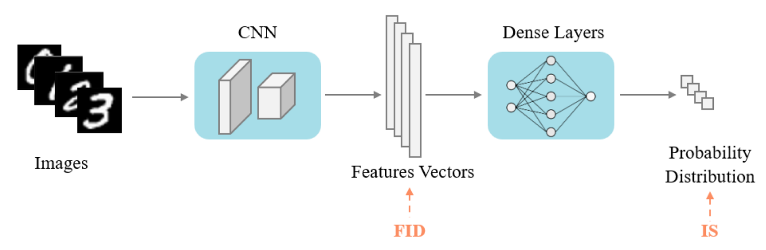

Quality and diversity score based on Fréchet Inception Distance (FID)

$$FID(D_R, D_G) = d^2(f.D_R, f.D_G) = ||\mu_R - \mu_G||^2 + Tr\left(\Sigma_R + \Sigma_G - 2(\Sigma_R\Sigma_G)^{1/2}\right)$$

with f extractor of features supposed multivariate Gaussian.

Fairness for generative models

Fairness Discrepancy metric

$$FD = |\bar{p} - E_{z\sim p_z(z)}\left(C(G(z))\right)|_2$$

where

- C is a classifier to predict the sensitive attribute
- G the generator
- $C(G(z))$ one hot encoder for the generated sample $(G(z))$
- z sample from a Gaussian noise distribution $p_z(z)$
- $\bar{p}$ uniformly distributed vector

- Python module based on PyTorch and Python >=3.7
- Computation Fairness metrics based on extraction of information from latent space of two popular Deep Neural Network: Inception V3 and CLIP
  - Sensitive attribute is the characteristic whose we want an equal algorithms quality outputs.
  - Evaluation if the images generated are of the same **quality** and **diversity** for the different values of the sensitive attribute
- Computation of Fairness metrics to verify if the generator generate as many images respecting the sensitive value (e.g. generate as many women chests as men chests ?)

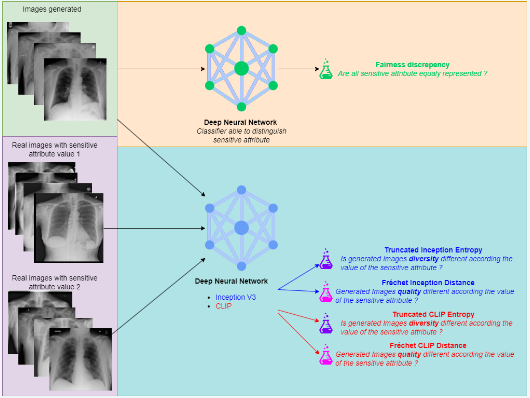

Medigan Evaluation

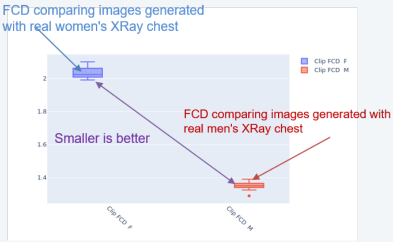

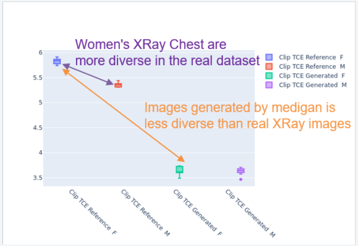

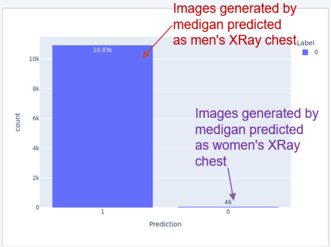

GDPR and European AI Act

GDPR

- European regulation about personal data since May 2018
- Principles:
  - Data processed lawfully
  - Purpose limitation
  - Data minimisation
  - Accuracy
  - Storage limitation
  - Integrity and confidentiality

European AI Act

- Context and timeline
  - Presented by EU commission on April 21st 2021
  - Follows up on EU AI strategy, EU Ethics guidelines for trustworthy AI and EU white paper on AI
  - One of the EU legislators' current priority
  - Parliament and Councitheir own negociation mandate, trilogue discussions ongoing
  - Regulation applicable in the Member States 24 or 36 months after its entry into force
- Main objectives
  - Prohibition of certain uses cases of AI Systems
  - Compliance regime for high risk AI Systems
  - Rules for general purpose AI Systems (incl. Foundation models)
  - Basic transparency rules for AI Systems interacting with natural persons
  - Definition of sanctions

EU AI Act - classification of AI Applications

- Unacceptable risk: prohibition of these AI uses (e.g. social rating, subliminal influence of people, categorizing people according sensitive attributes, real-time remote biometric identification systems in public spaces)
- High-risk AI: compliance review ex ante and during the life of AI Systems (e.g. AI used in education, police, other case of biometric identification, critical infrastructure, etc.) requiring: a risk management process, **detection and correction of bias in particular through the quality of training, validation and test data**, establishment of technical documentation, human control, robustness/accuracy/security
- General purpose AI:
  - Risk mitigation measures, training data quality, energy efficiency, documentation, etc.
  - Additional obligations for generative AI: transparency regarding third party rights included in training data