110° Nexa Lunch Seminar

Towards Responsible AI in Banking: Addressing Bias for Fair Decision-Making

Alessandro Castelnovo

Politecnico di Torino

Nexa Center for Internet & Society

Business Worldwide Declares AI as a Strategic Goal

SCALE ADOPTION OF AI-BASED SOLUTIONS LEAD TO MANY BENEFITS…

SUPPORT FOR REVENUE GROWTH

- Improvement of the Customer Experience through personalized offers and experiences
- Improvement of the Customer Satisfaction through dedicated marketing

SIGNIFICANT COSTS REDUCTION

- Optimization of operational processes by reducing manual activities

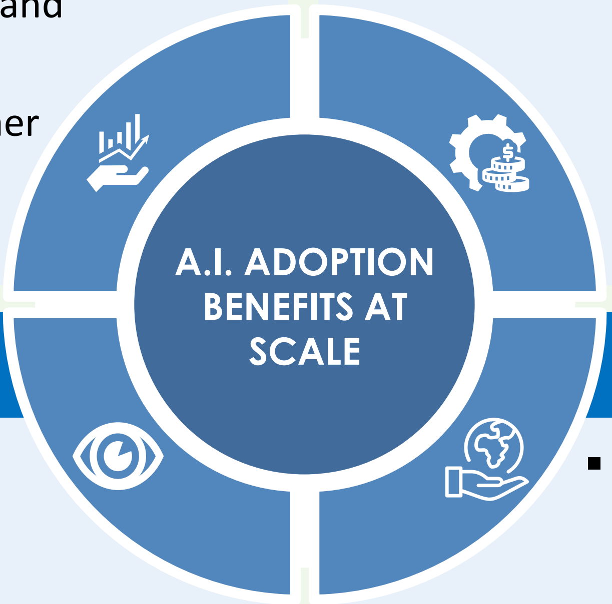

OPTIMIZATION OF RISK PROFILES

- Increase in the accuracy of the Credit Score through enrichment of the database for risk analysis

SUSTAINABILITY AND COMPLIANCE

- Automation of Conformity Checks for compliance purposes
- Optimization of profiling for ESG purposes, also through external data

…WHOSE VALUE IS NOW RECOGNIZED WITHIN THE INDUSTRY

Executives see Artificial Intelligence at the center of business evolution…

84%

of Executives believes in the centrality of A.I. to achieve growth goals[1]

… with a potential incremental value of enormous magnitude for the Banking Industry

$1 Trillion

Potential annual incremental value from AI and analytics for the global banking industry[2]

Sources: 1 Accenture 2021 Banking Conference; 2) «The executive’s AI playbook,» McKinsey.com

AI Implementation Doesn't Always Guarantee Expected Benefits

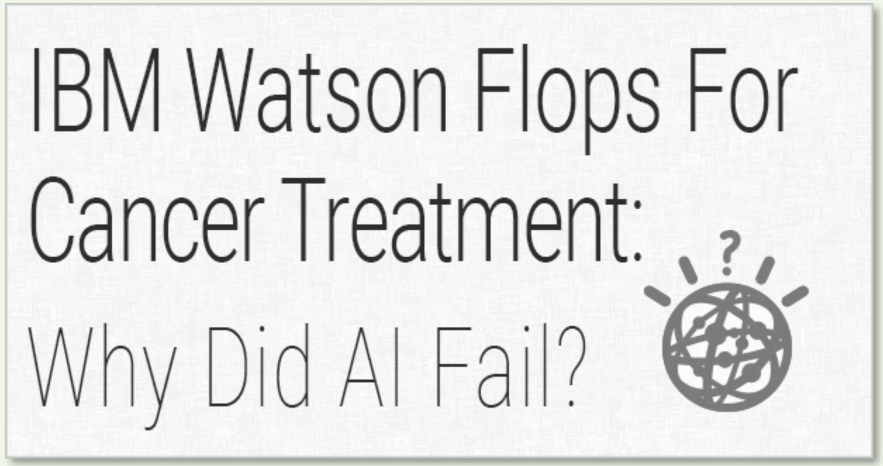

IBM's Watson for Oncology cancelled

Google's AdFishertool served significantly fewer ads for high paid jobs to women than men

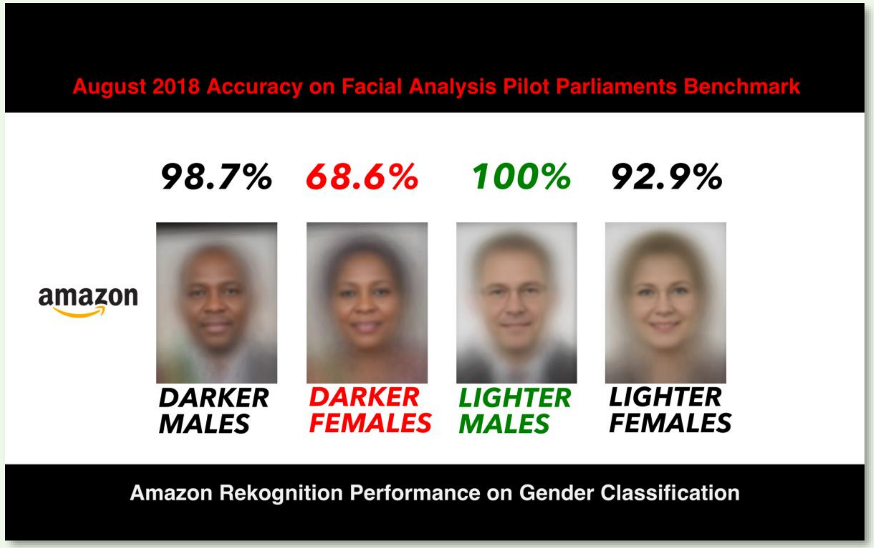

Amazon's Facial recognition works better for white males

Microsoft's bot Taytaken offline after racist tweets

This Could Represent a Risk for People

data as a social mirror

ML could amplify and perpetuate biases already present in data, at large scale

sample size imbalances

ML could disregard minority groups, effectively producing bias even if absent in the data

this can have a huge impact on people's lives
e.g. Recruiting / Loans approval

Regulation Risk for Companies - The New European Regulation on AI

The importance of ***limiting AI risks*** is unequivocally demonstrated by the European Union’s proactive efforts to regulate AI, aiming to create a ***more favorable environment for the development and deployment of AI***

A.I. ACT

- Provides for **different levels** of **risk** based on possible **discrimination** and **impacts** on fundamental **human rights** such as
  - Dignity
  - Freedom
  - Equality
  - Solidarity (including health protection)
- Identifies **cross-sectoral "high risk" A.I. systems** and contains no specific **provisions** for **FSI**

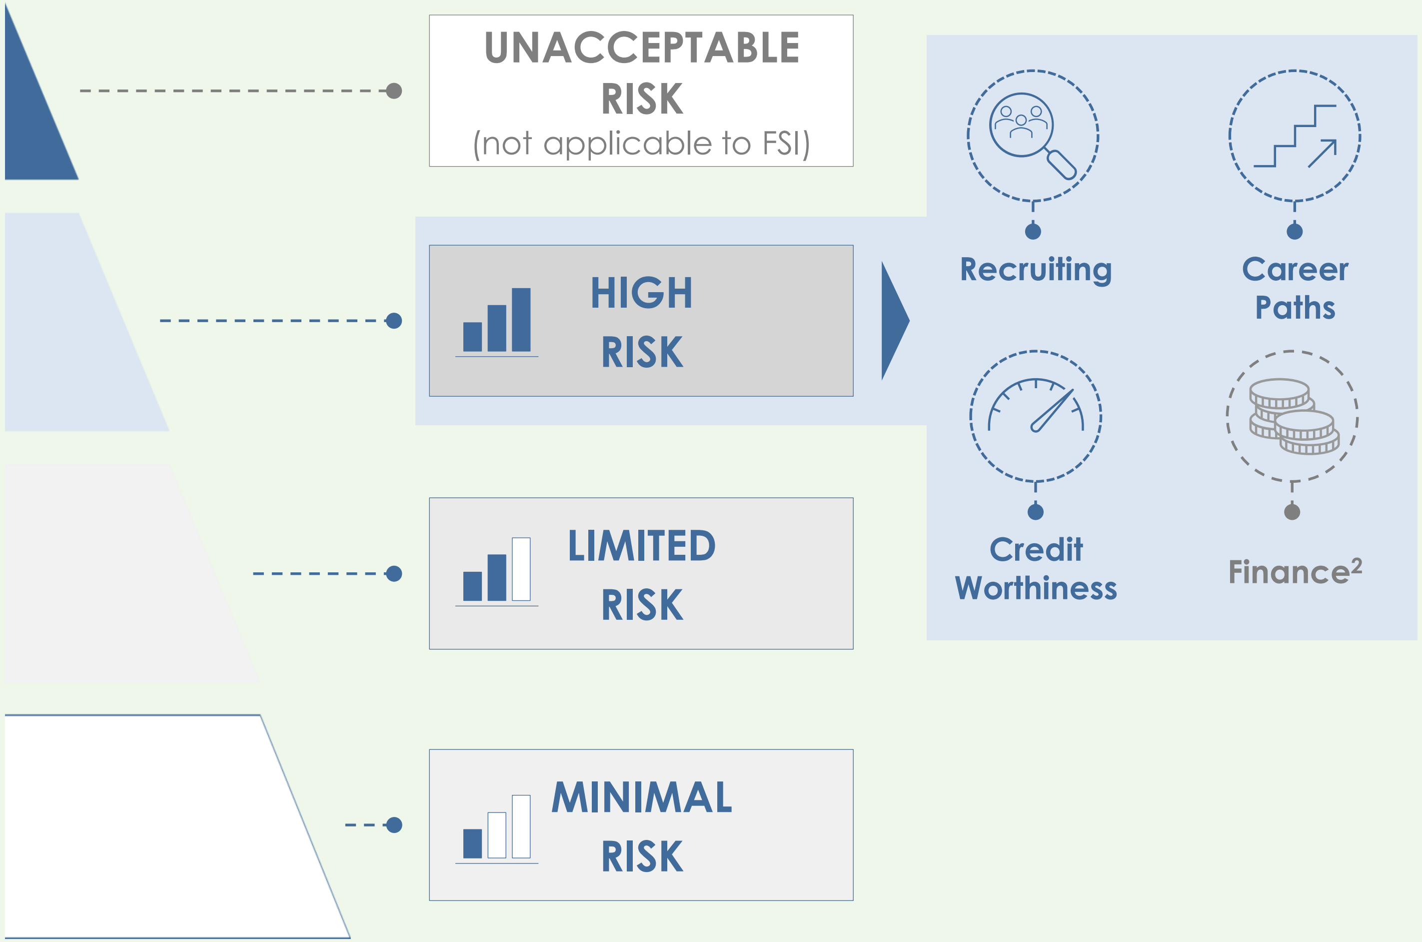

In April 2021, the European Commission introduced the First Proposal for a “Regulation laying down harmonised rules on artificial intelligence” (AI Act) [169]. On December 8, 2023, the European Parliament and Council reached a provisional agreement on the AI Act.

Reputation Risk for Companies

HOME > TECH

Amazon built an AI tool to hire people but had to shut it down because it was discriminating against women

Isobel Asher Hamilton Oct 10, 2018, 11:47 AM

INSIDER

PRO PUBLICA

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

GOOGLE IS POISONING ITS REPUTATION WITH AI RESEARCHERS

The firing of top Google AI ethics researchers has created a significant backlash

By James Vincent | Apr 13, 2021, 9:30am EDT

THE VERGE

WILL KNIGHT BUSINESS 11.19.2019 09:15 AM

WIRED

The Apple Card Didn't 'See' Gender—and That's the Problem

The way its algorithm determines credit lines makes the risk of bias more acute.

Control Measures for promoting trust in A.I. Solutions

The adoption of AI by a company should be contingent upon widespread understanding, not only among data scientists and developers but also within governance and compliance structures

CONTROL MEASURES PROVIDED DEPENDING ON THE RISK LEVEL

FAIRNESS

fairness enables the **measurement** and **mitigation** of **undesired biases,** ensuring that AI systems exhibit **desirable ethical characteristics**

EXPLAINABILITY

explanations for an AI system provide **human-understandable interpretations** of its inner workings and outcomes

HUMAN OVERSIGHT

Definition of **supervision** and **human monitoring mechanisms** on **A.I. systems**, on the outputs **produced** and on the **resulting decisions**

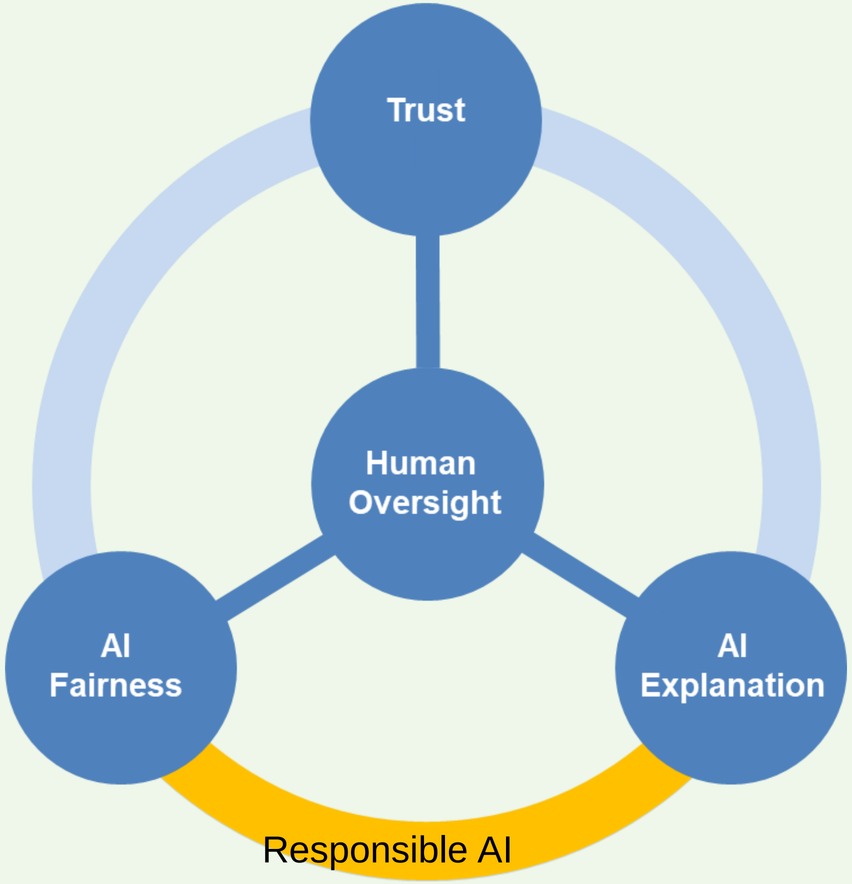

Promoting trust in AI-based decision-making requires the integration of three key elements: fairness, eXplainable Artificial Intelligence (XAI), and human oversight, which must be combined harmoniously

Contributions Landscape of my thesis

Understanding Bias

1. Bias and Moral Framework in AI-based Decision Making
2. A framework for generating synthetic data that emulates Fundamental biases
3. Fairness Metrics Landscape

Mitigating Bias

4. Fairness Mitigation
5. FFTree: A Flexible Tree to Mitigate Multiple Fairness Criteria

Accounting for Bias

6. Criteria for choosing a fairness metric
7. Towards Fairness Through Time
8. Addressing Fairness in the Banking Sector

Part I

Understanding Bias

This part aims to deepen our understanding of how bias is generated, how it manifests in the data, and how it impacts the outcomes of AI systems.

1. Bias and Moral Framework in AI-based Decision Making
2. A framework for generating synthetic data that emulates Fundamental biases
3. Fairness Metrics Landscape

Publications related to this part are:

Baumann, J., Castelnovo, A., Crupi, R., Inverardi, N., and Regoli, D. (2023). **Bias on Demand: A Modelling Framework That Generates Synthetic Data With Bias**. In *2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT '23*, New York

Castelnovo, A., Crupi, R., Greco, G., Regoli, D., Penco, I. G., and Cosentini, A. C. (2022b). **A clarification of the nuances in the fairness metrics landscape**. *Scientific Reports, 12(1):1–21*

Ethical Moral Frameworks for Choosing Fairness in Machine Learning

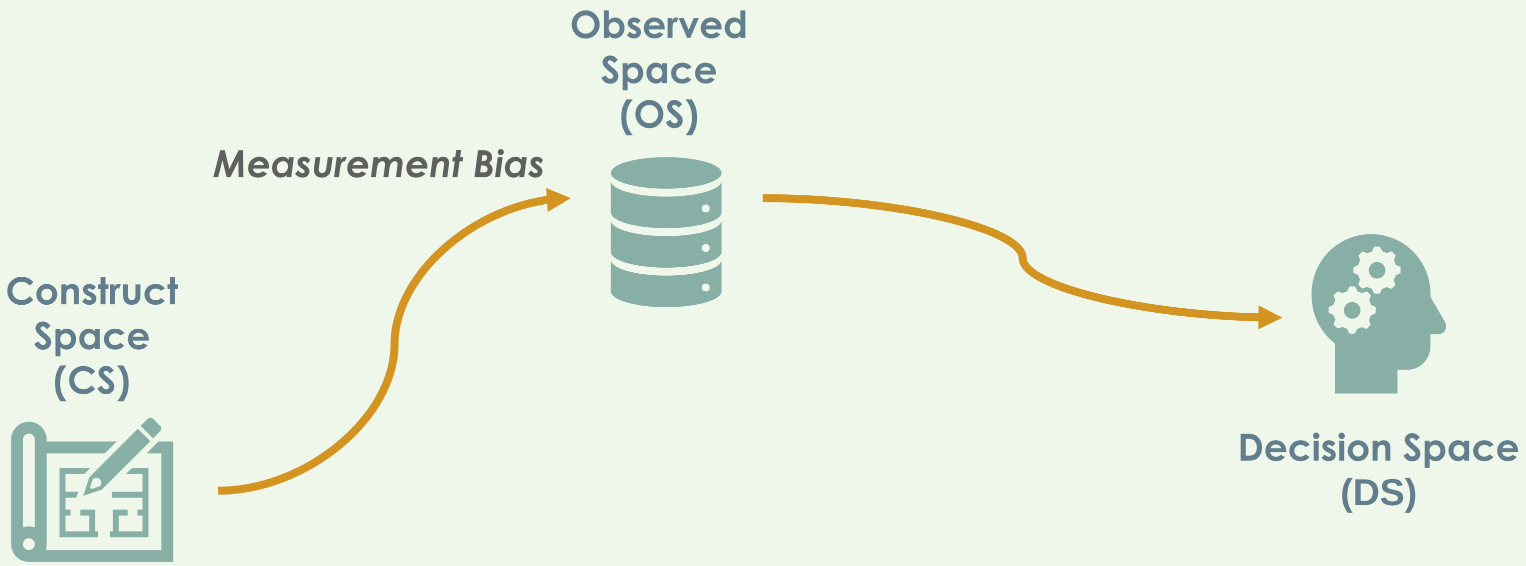

Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. “The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making”. In: Communications of the ACM 64.4 (2021),

Ethical Moral Frameworks for Choosing Fairness in Machine Learning

In The ***What You See Is What You Get* (WYSIWYG)** worldview, CS and OS must be considered equal, and any eventual difference between them is irrelevant to the fairness of the corresponding choice in DS

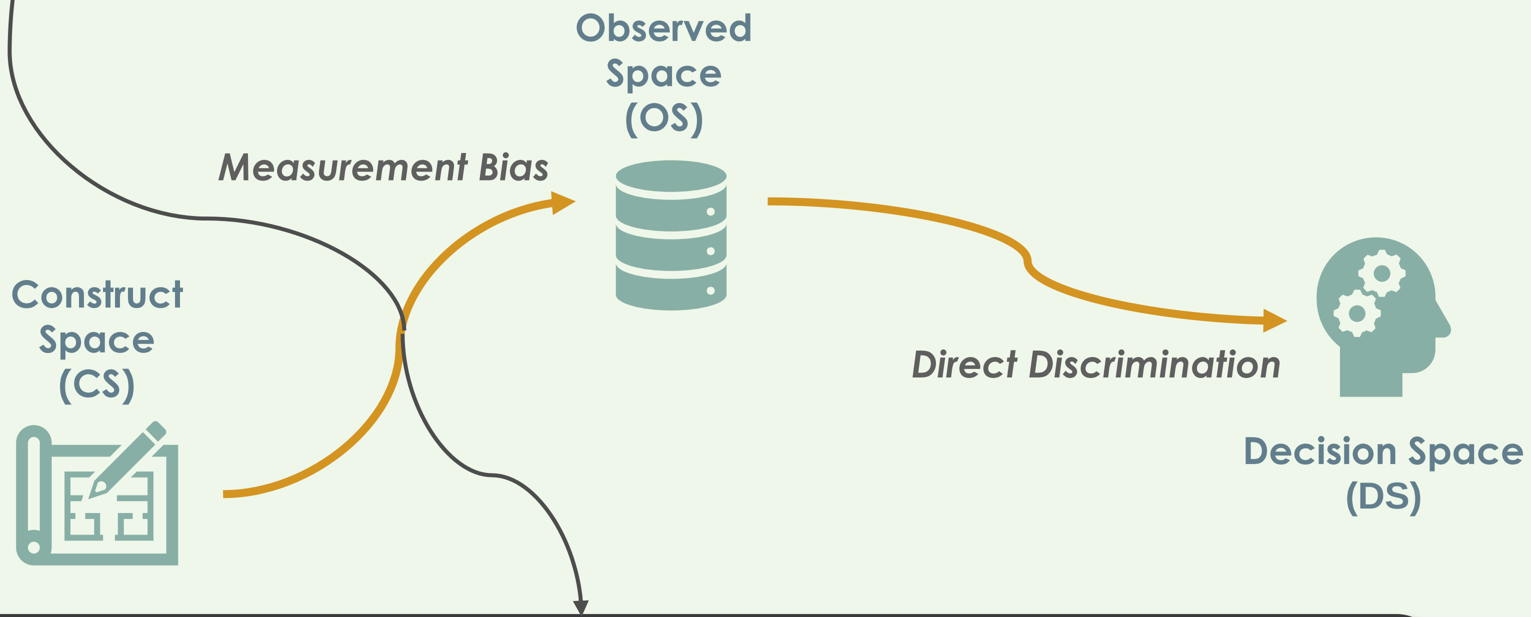

In The ***We Are All Equal* (WAE)** worldview, states that individuals are all equal at a certain point in time in CS. Therefore, in this perspective, any distortion detectable between CS and OS must be interpreted as caused by a biased observation method corresponding to an unfair mapping.

Sorelle A Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. “The (im) possibility of fairness: Different value systems require different mechanisms for fair decision making”. In: Communications of the ACM 64.4 (2021),

1

Ethical Moral Frameworks for Choosing Fairness in Machine Learning

In The ***What You See Is What You Get (WYSIWYG)*** worldview, CS and OS must be considered equal, and any eventual difference between them is irrelevant to the fairness of the corresponding choice in DS

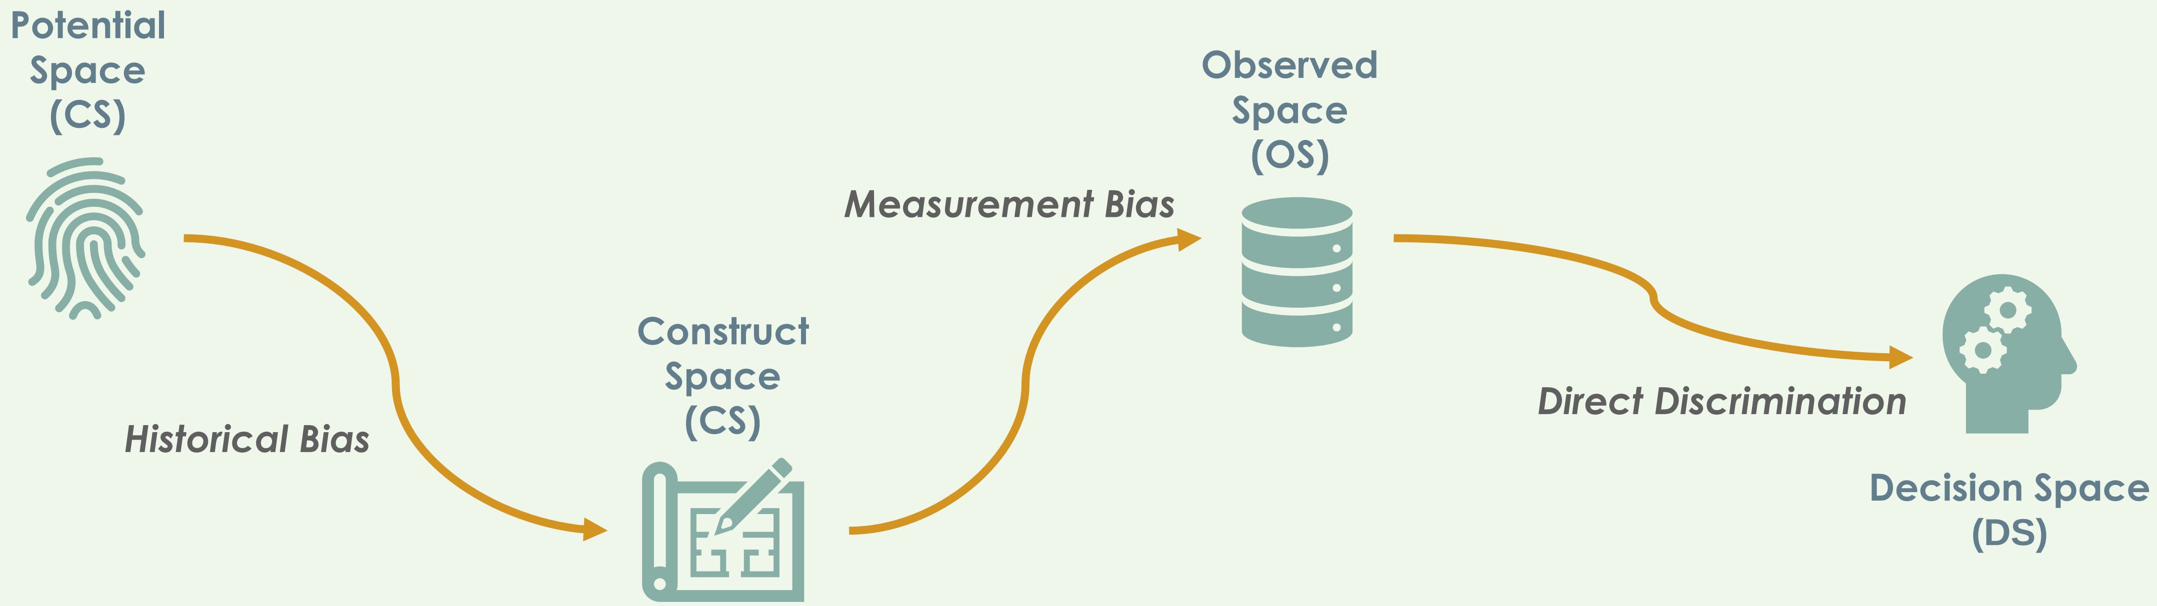

Hertweck, C., Heitz, C., and Loi, M. (2021). On the moral justification of statistical parity. In Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, pages 747–757.

Bias Throughout the ML Life Cycle

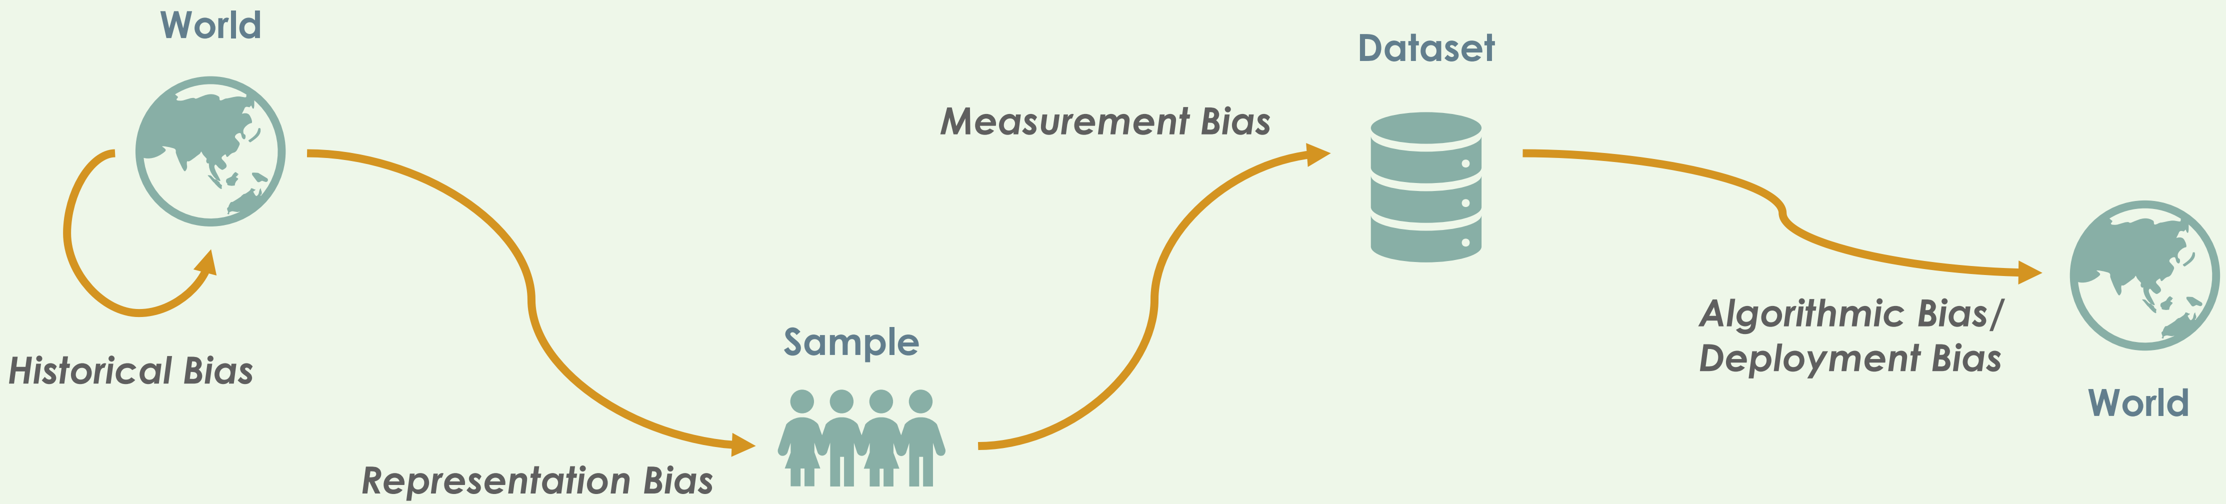

Suresh, H. and Guttag, J. (2021). A framework for understanding sources of harm throughout the machine learning life cycle. In Equity and Access in Algorithms, Mechanisms, and Optimization, EAAMO ’21, New York, NY, USA. Association for Computing Machinery

Family of Biases

Bias From Users to Data

Bias is present in the underline phenomenon that generates the data

$$Y = f(X) + \epsilon$$

Variables needed by the model

Historical/Life Bias

Bias From Data to Algortihm

Bias is due to the data collection mechanism

$$\tilde{X} = g(X);$$
$$\tilde{Y} = h(Y)$$

Variables used by the model

Measurement Bias

Representation/Sampling Bias

Omission Bias

Bias From Algortihm to User

Bias is due to the predictor/classification mechanism

$$\hat{Y} = \hat{f}(\tilde{X})$$

Function learned by the model

Algorithmic Bias (Aggregation bias, Learning Bias, Evaluation Bias)

Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. (2021). A survey on bias and fairness in machine learning. ACM Computing Surveys (CSUR), 54(6):1–35.

Bias On Demand

***Bias On Demand** is a toolkit that permits to generate synthetic dataset with different combination of bias.*

Advantages of this approach are:

Education

Easy Toolkit to show how different biases arise in data

Investigation

Effects of different bias combinations on performance and fairness evaluation

Research

Ad-Hoc scenarios for developing "bias-aware" systems

Baumann, J., Castelnovo, A., Crupi, R., Inverardi, N., and Regoli, D. (2023). Bias on Demand: A Modelling Framework That Generates Synthetic Data With Bias. In 2023 ACM Conference on Fairness, Accountability, and Transparency, FAccT '23, New York, NY, USA.

The Zoo of the Fairness Metrics

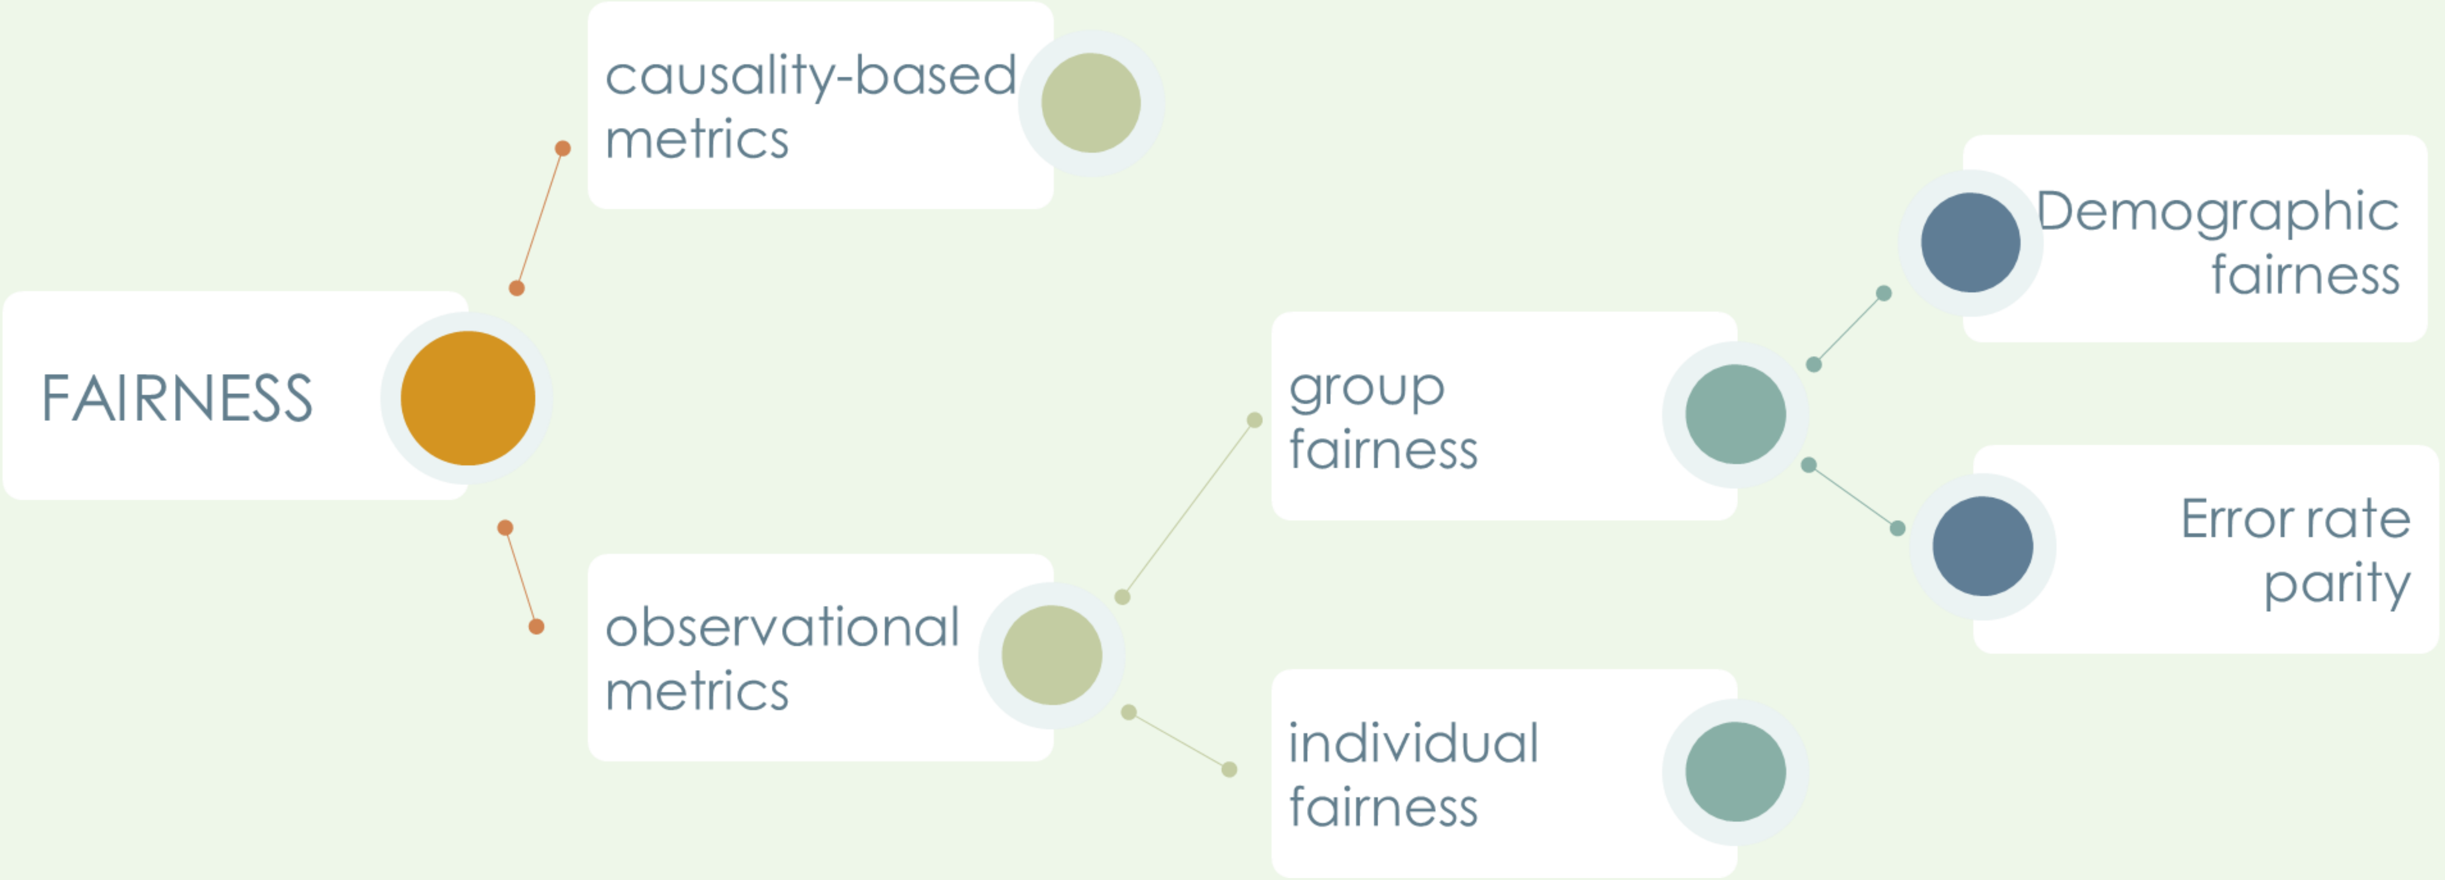

Narayanan, A., 2018. Translation tutorial: 21 fairness definitions and their politics, in: Proc. Conf. Fairness Accountability Transp., New York,

In General, Fairness Metrics are Non-Compatible With One Another

Individual Fairness

"individuals who are similar (with respect to some task) should be treated similarly (with respect to that task)" (Dwork,2012)

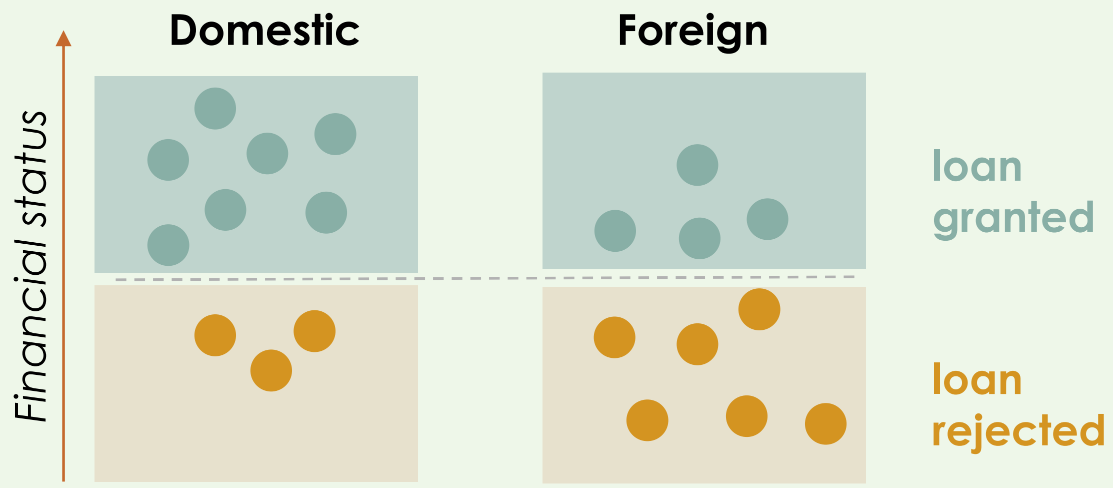

Individual Fairness mantains the **status quo**

In line with WYSIWYG worldview

Statistical Parity

"demographic groups should, *on the whole*, receive equal decisions"
(inspired by civil rights law in the US and UK)

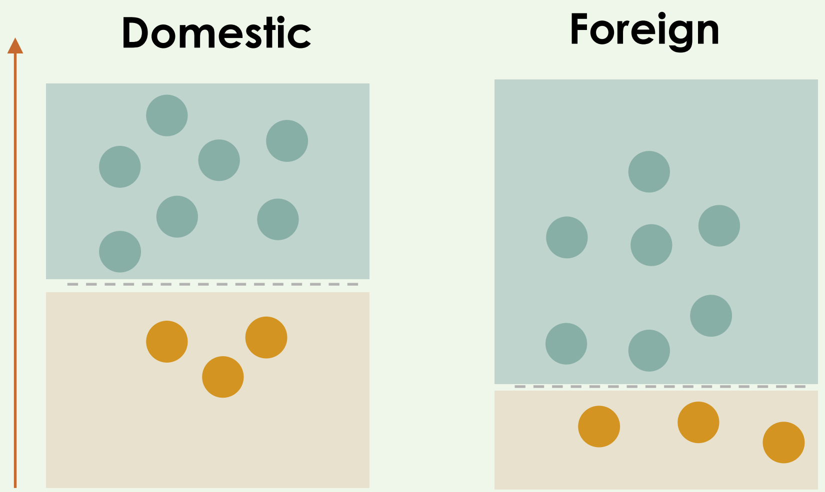

Statistical Parity breaks the **status quo**

In line with WAE worldview

Incompatibility between Error Rate Parity Metric

The Compas debate

n the USA, a software was developed to predict criminal recidivism. This software, with fairness in mind, was developed to meet **Precision Parity**

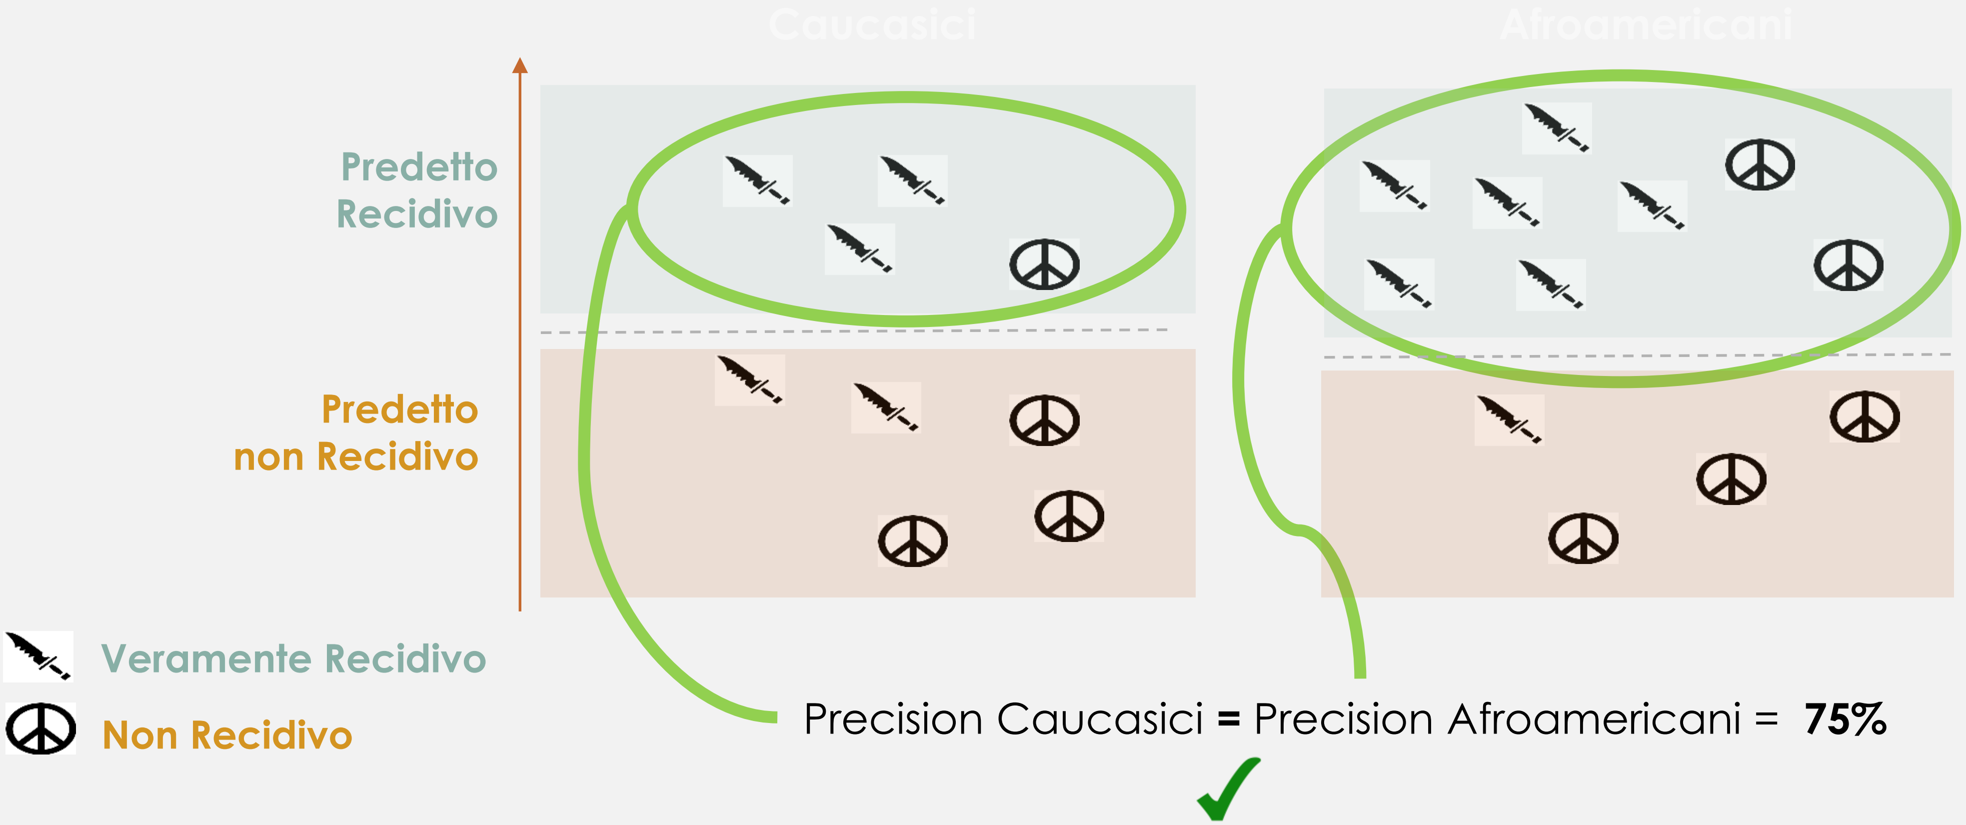

Incompatibility between Error Rate Parity Metric

The Compas debate

COMPAS was criticized in the media for being discriminatory as it did not meet **Equal Opportunity**

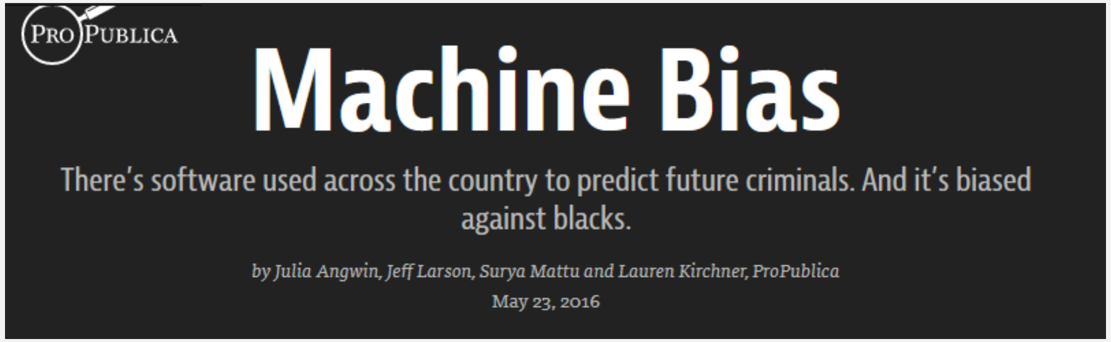
PROPUBLICA

Machine Bias

There's software used across the country to predict future criminals. And it's biased against blacks.

by Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, ProPublica

May 23, 2016

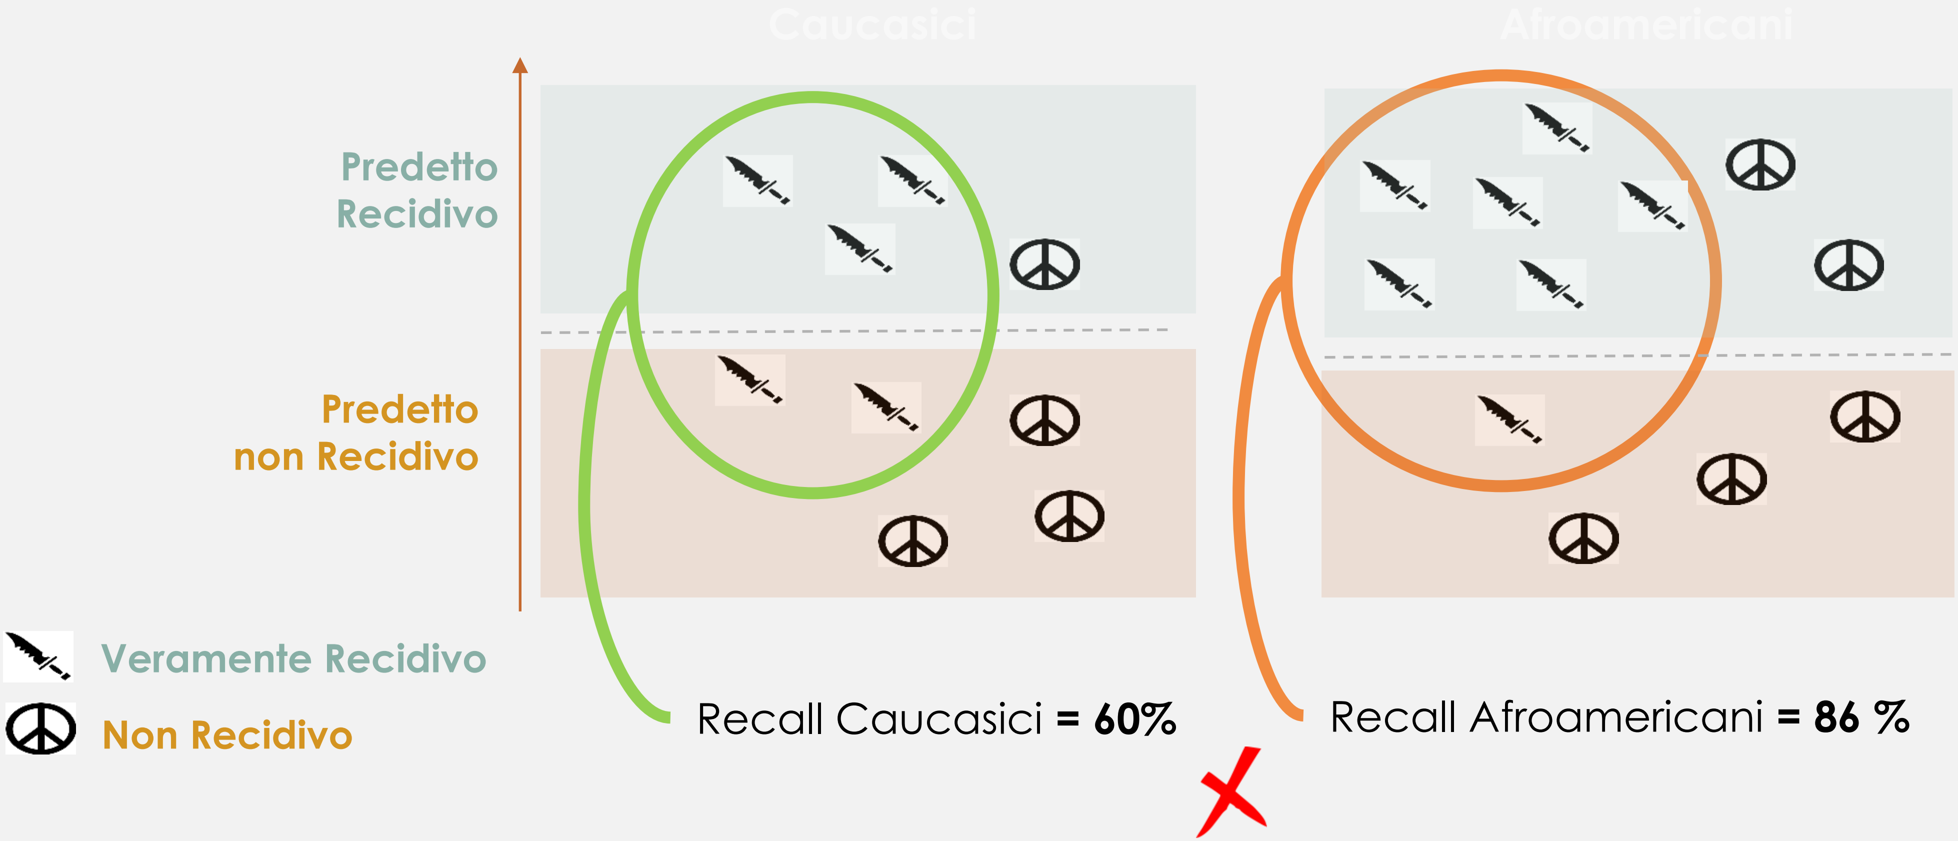

Part II

Mitigating Bias

This part focus on addressing bias at different stages of AI decision-making, such as pre-processing, in-processing, and post-processing. They aim to mitigate bias by carefully handling data inputs, optimizing learning algorithms, and refining model outputs.

Publications related to this part are:

Castelnovo, A., Cosentini, A., Malandri, L., Mercorio, F., and Mezzanzanica, M. (2022a). Fftree: A flexible tree to handle multiple fairness criteria. Information Processing & Management

Castelnovo, A., Crupi, R., Del Gamba, G., Greco, G., Naseer, A., Regoli, D., and Gonzalez, B. S. M. (2020). Befair: Addressing fairness in the banking sector. In 2020 IEEE International Conference on Big Data (Big Data)

Fairness Mitigation Strategies

pre-processing:

F. Kamiran and T. Calders, "Data preprocessing techniques for classification without discrimination," Knowledge and Information Systems, vol. 33, no. 1, pp. 1–33, 2012.

in-processing:

B. H. Zhang, B. Lemoine, and M. Mitchell, "Mitigating unwanted biases with adversarial learning," in Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, 2018, pp. 335–340.

post-processing:

M. Hardt, E. Price, and N. Srebro, "Equality of opportunity in supervised learning," in Advances in neural information processing systems, 2016, pp. 3315–3323.

BeFair: a Fairness Mitigation Toolkit

*To assist data scientists at Intesa Sanpaolo in their efforts to achieve fairness mitigation, we have developed a comprehensive toolkit called **BeFair***

	Mitigation Technique	Demographic Parity	Error Rate Parity	Individual Fairness	Counterfactual Fairness
Pre Processing	FTU			✓	
	Suppression	✓			
	Massaging	✓			
	Sampling	✓			
In Processing	CFF				✓
	AdvDP	✓			
	AdvEO		✓		
	AdvCDP	✓		✓	
	ReductionsGS	✓			
	ReductionsEG	✓			
Post processing	ThreshDP				
	ThreshEO		✓		
	ThreshEopp		✓		
	ThreshCDP	✓		✓	

Castelnovo, A., Crupi, R., Del Gamba, G., Greco, G., Naseer, A., Regoli, D., and Gonzalez, B. S. M. (2020). Befair: Addressing fairness in the banking sector. In 2020 IEEE International Conference on Big Data (Big Data)AIF

BeFair Experimentsv(1/2)

Experiment on Fairness Mitigation using Real-World Data on Credit Lending

~200,000 loan applications
~50 predictors, including financial variables and personal information.
The target is the final decision of a human officer.

Throughout the analysis, we focus on

CITIZENSHIP = {0, 1}

as **sensitive attribute** with respect to which assess fairness.

Bias, measured in terms of Demographic Parity, is negligible in the original target, but amplified by the application of a ML model.

		fairness				performance		
family	**type**	**DP**	**EO**	**EOpp**	**PP**	**AUROC**	**Accuracy**	**F1**
no mitigation	Logistic	0.324	0.272	0.272	**0.032**	0.817	0.761	0.823
	Random forest	0.221	0.202	-0.104	0.068	**0.838**	0.804	0.875
	Neural network	0.219	0.198	0.104	0.072	0.830	0.811	**0.876**
pre-process	FTU	0.164	0.124	0.058	0.095	**0.838**	0.812	**0.876**
	Suppression	0.099	-0.053	0.065	0.152	0.753	0.748	0.840
	Massaging	-0.004	0.062	0.062	0.163	0.818	**0.868**	0.803
	Sampling	0.080	0.012	0.012	0.115	0.835	0.791	0.851
	CFF	0.218	0.192	0.104	0.070	0.832	0.810	0.874
in-process	AdvDP	-0.034	0.073	0.063	0.176	0.823	0.802	0.869
	AdvEO	0.102	0.029	-0.010	0.148	0.819	0.805	0.871
	AdvCDP	0.147	0.101	-0.050	0.112	0.830	0.807	0.872
	ReductionsGS	0.012	0.077	0.049	0.159	0.812	0.794	0.864
	ReductionsEG	0.007	0.084	0.051	0.161	–	0.794	0.864
post-process	ThreshDP	**0.003**	0.099	0.056	0.164	–	0.805	0.872
	ThreshEO	0.082	**0.006**	0.006	0.138	–	0.812	0.873
	ThreshEOpp	0.100	0.048	**0.005**	0.119	–	0.809	0.874
	ThreshCDP	0.186	0.159	0.072	0.083	–	0.810	0.875

Castelnovo, A., Crupi, R., Del Gamba, G., Greco, G., Naseer, A., Regoli, D., and Gonzalez, B. S. M. (2020). Befair: Addressing fairness in the banking sector. In 2020 IEEE International Conference on Big Data (Big Data)AIF

BeFair Experiments (2/2)

Utilizing the BeFair Interface to Facilitate Optimal Mitigation Approach Selection

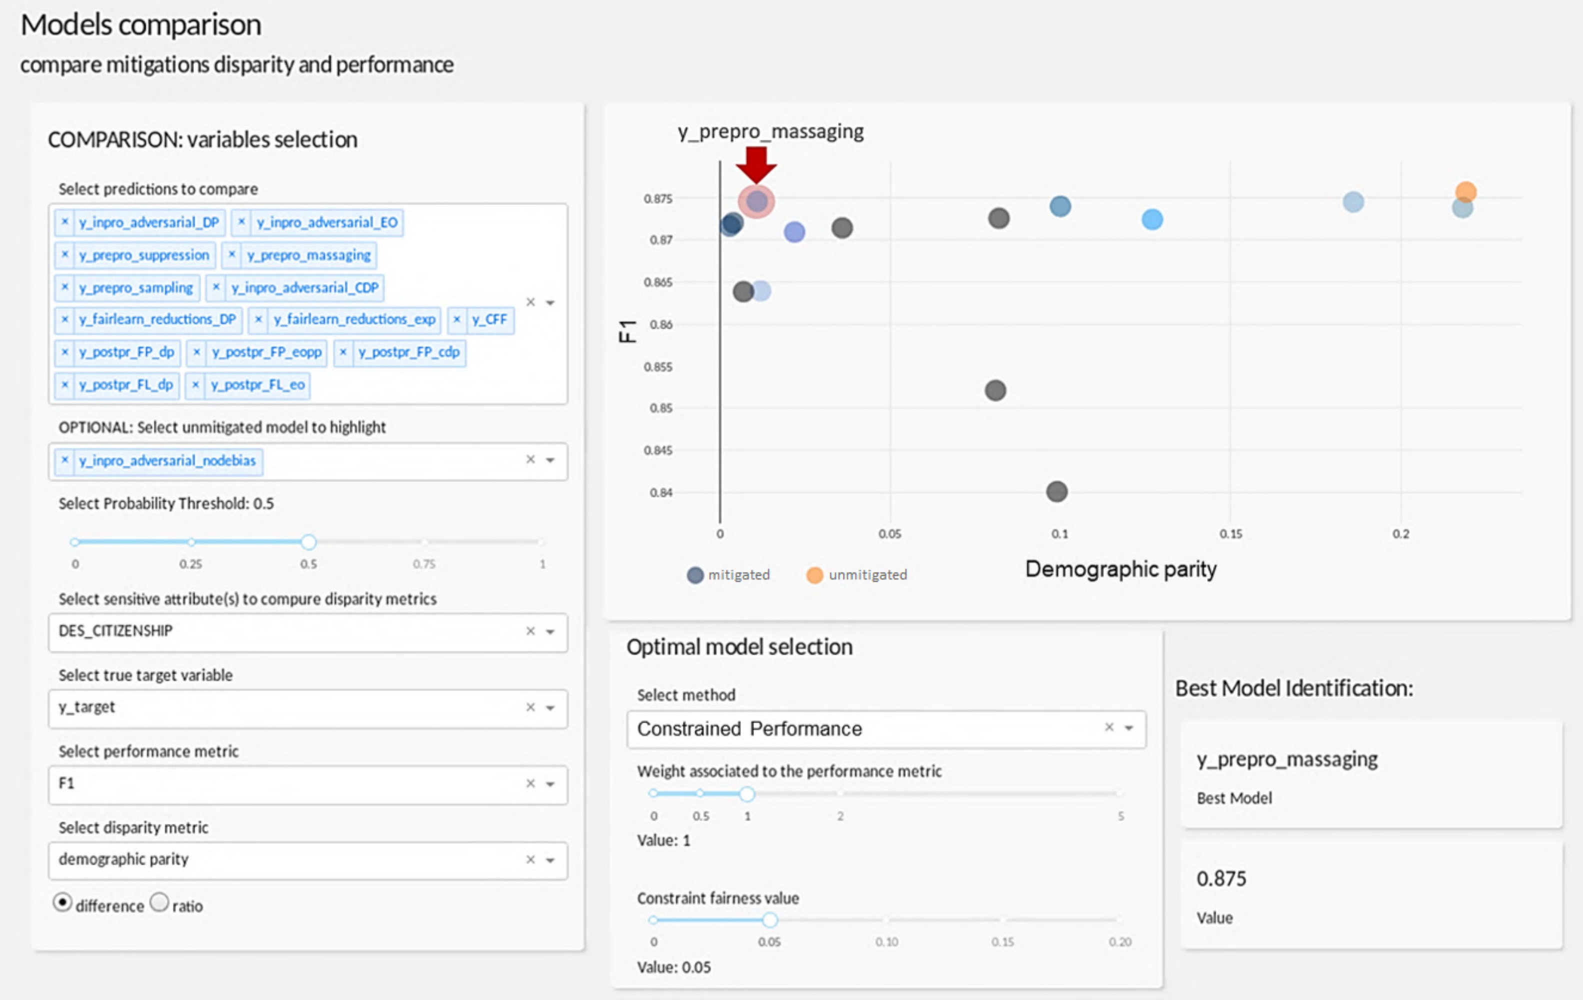

Proposed methods to identify the best perfomance-fairness tradeoff:

Trade-off fairness-performance

$$(1+\ \beta^2)\frac{(1-|\phi|\ )\ *\pi}{\beta^2\ *(1-|\phi|\ )+\pi}$$

Constrained performance

$$\max_{\phi\leq\Phi}\pi$$

π and φ are the preferred performance and fairness metrics, respectively and beta is the weight associated with the performance metric.

Castelnovo, A., Crupi, R., Del Gamba, G., Greco, G., Naseer, A., Regoli, D., and Gonzalez, B. S. M. (2020). Befair: Addressing fairness in the banking sector. In 2020 IEEE International Conference on Big Data (Big Data)AIF

Common Challenges in Fairness Mitigation

As detailed in [2], the mitigation strategies proposed in prior studies typically lack flexibility with respect to the following aspects:

- They are specifically designed for only one fairness criterion, and cannot accommodate more than one simultaneously;
- They cannot ensure fairness with respect to multiple sensitive features simultaneously (e.g., gender and race);
- They are typically designed as a black box, i.e. they are not directly interpretable.

[2] Zafar, Valera, Gomez-Rodriguez, and Gummadi. Fairness constraints: A flexible approach for fair classification. JMLR, 20(1):2737–2778, 2019.

FFTree: A flexible tree to handle multiple fairness criteria

We present ***FFTree***, a new transparent, flexible and fairnessaware classifier. As a novelty, ***FFTree*** enchances the classical approach introduced in [3] with a new approach to find a "fair" split to:

- ✓ Satisfy a fairness constraint selected from a wide range of possible definitions of fairness;
- ✓ Implement more than one fairness criterion;
- ✓ Handle more than one sensitive attribute at the same time;
- ✓ Set the required level of fairness as an input parameter to meet different business needs or regulatory requirements.

State-of-the-art Fair Tree	DI	DT	DM	MD	MS	BN
Kamiran et al.	✓	✗	✗	✗	✗	✗
Zhang and Ntoutsi	✓	✗	✗	✗	✗	✗
Aghaei et al.	✓	✓	✗	✗	✗	✓
FFTree (our method)	✓	✓	✓	✓	✓	✓

[3] Brieman, Friedman, Olshen, and Stone. Classification and regression trees. Wadsworth Inc.

Castelnovo, A., Cosentini, A., Malandri, L., Mercorio, F., and Mezzanzanica, M. (2022a). Fftree: A flexible tree to handle multiple fairness criteria. Information Processing & Management

Part III

Accounting For Bias

This part focus on proposing approaches proactively account for bias by incorporating bias-aware decision-making mechanisms. They also prioritize human involvement, allowing for human intervention and oversight, while ensuring that understandable explanations of AI outcomes are provided.

6 — Criteria for choosing a fairness metric

7 — Towards Fairness Through Time

8 — Addressing Fairness in the Banking Sector

Publications related to this part are:

Castelnovo, A., Malandri, L., Mercorio, F., Mezzanzanica, M., and Cosentini, A. Towards fairness through time. In Machine Learning and Principles and Practice of Knowledge Discovery in Databases: International Workshops of ECML PKDD 2021

Castelnovo, A., Inverardi, N., Malandri, L., Mercorio, F., Mezzanzanica, M., and Seveso, A. (2023b). Leveraging group contrastive explanations for handling fairness. In World Conference on Explainable Artificial Intelligence, pages 332–345. Springer.

Castelnovo, A., Crupi, R., Del Gamba, G., Greco, G., Naseer, A., Regoli, D., and Gonzalez, B. S. M. (2020). Befair: Addressing fairness in the banking sector. In 2020 IEEE International Conference on Big Data (Big Data)

Effects of Fairness Metrics (1/2)

The choice of fairness metric primarily depends on the willingness to change the status quo

Some metrics that maintain the *status quo*:

- Individual Fairness
- Error Rate Parity

"A perfect model is also perfectly fair"

In line with WYSIWYG worldview

Some metrics that change the *status quo*:

- Statistical Parity
- Causality-based Metrics

"A perfect model cannot be fair"

In line with WAE worldview

Effects of Fairness Metrics(2/2)

La scelta di cambiare lo status quo può dipendere sia da ragioni strategiche, sia di beneficio per la popolazione. In alcuni casi, decisioni positive possono creare danni.

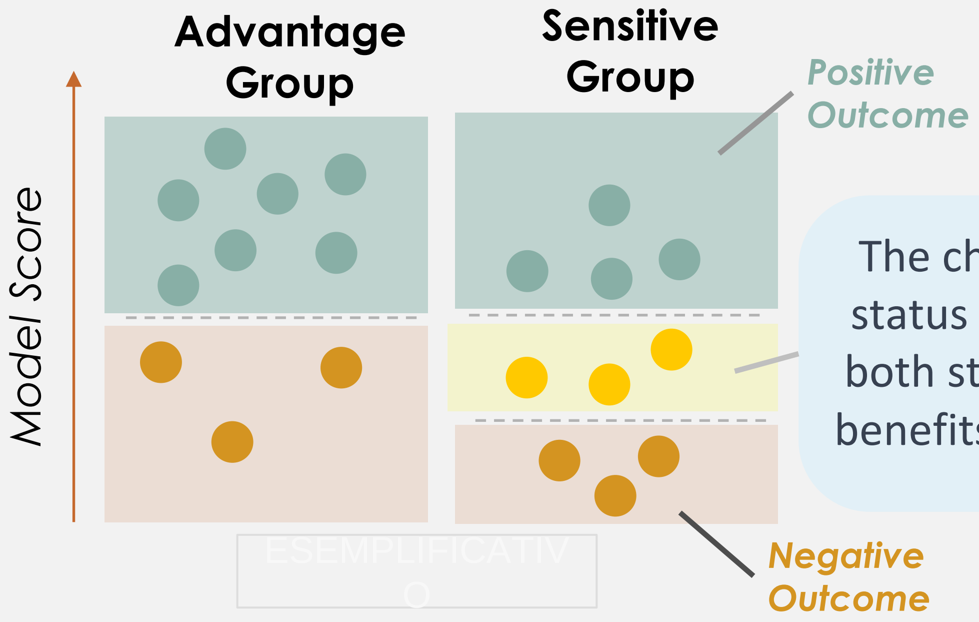

The choice to change the status quo can depend on both strategic reasons and benefits for the population.

They should be classified positively when the decision brings benefits that improve the entire group and there are no reasons to believe that classes are treated differently

They should be classified negatively when the decision can cause harm to the recipient or to the entire society

Monitoring Fairness Through Time

Environment

real-world financial data
800,000 personal loan granted from 2016 to 2019
+ 50 financial features

Target

Human decision on granting the loan

Bias *"in the inputs"*

The sensitive attribute *citizenship* is correlated with many financial features

Bias *"in the outputs"*

The target is correlated with the sensitive attribute *citizenship*

Chosen Mitigation Policy:
Deploy a ML model that ensure Demographic Parity
To lead an improvement to the vulnerable class and reach in **long-term** DP and Individual Fairness simultaneously - ***Optimal Situation***

Challenging Questions

C1
Will the outputs of a mitigation model continue to ensure Demographic Parity over time?

C2
How can XAI techniques be used to verify that the chosen fairness policy (*ensure Demographic Parity*) is helping to reduce individual discrimination over time?

Castelnovo, A., Malandri, L., Mercorio, F., Mezzanzanica, M., and Cosentini, A. Towards fairness through time. In Machine Learning and Principles and Practice of Knowledge Discovery in Databases: International Workshops of ECML PKDD 2021

1st Challenging Questions

Will the outputs of a mitigation model continue to ensure Demographic Parity over time?

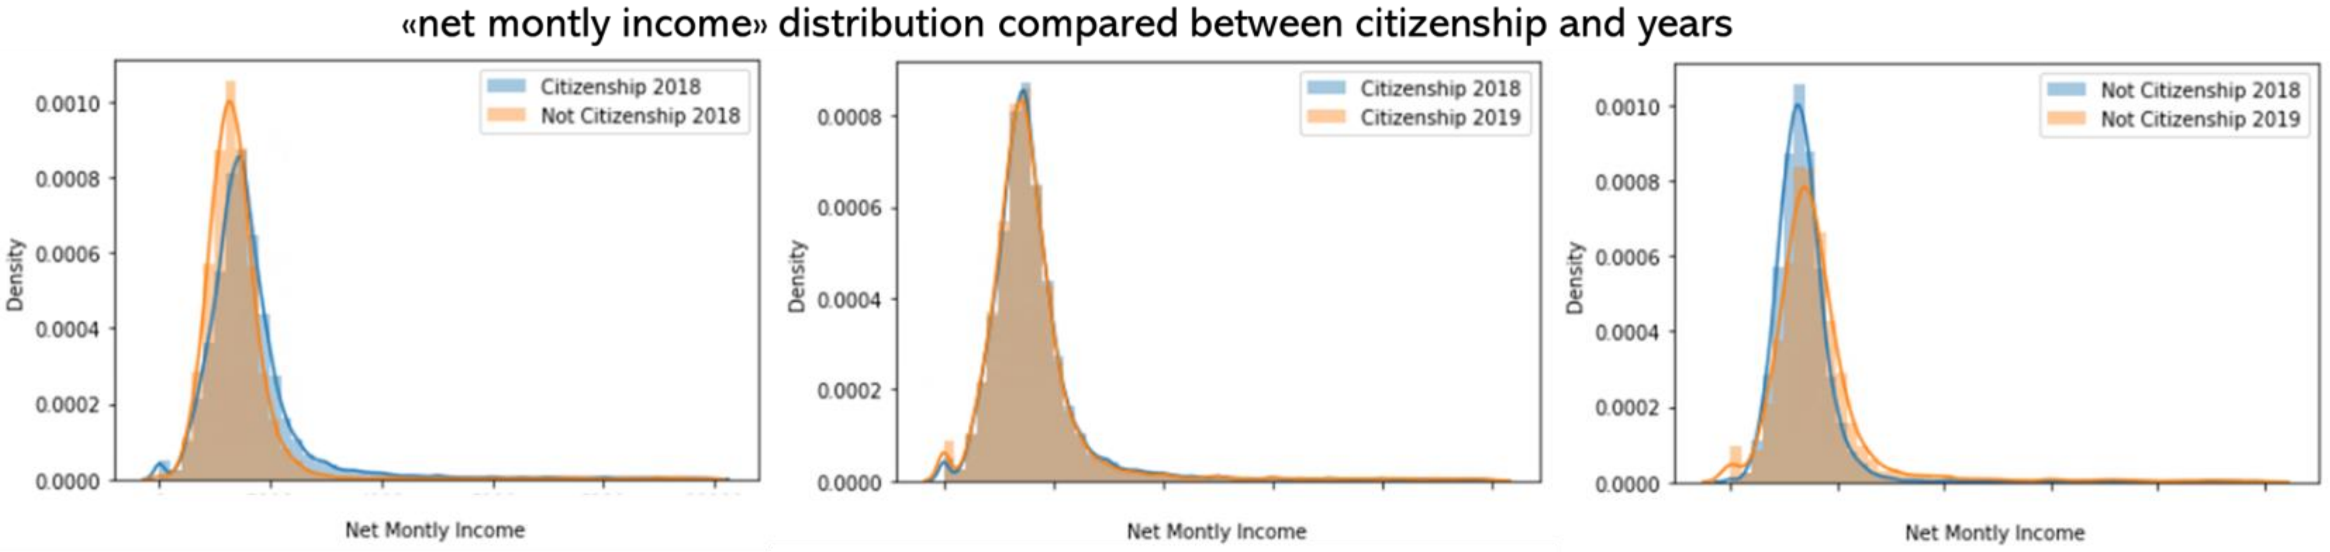

Density plot of the variable net montly income conditioned to vary combination of citizenship and year. Distribution values are blinded for data privacy.

Castelnovo, A., Malandri, L., Mercorio, F., Mezzanzanica, M., and Cosentini, A. Towards fairness through time. In Machine Learning and Principles and Practice of Knowledge Discovery in Databases: International Workshops of ECML PKDD 2021

1st Challenging Questions

Will the outputs of a mitigation model continue to ensure Demographic Parity over time?

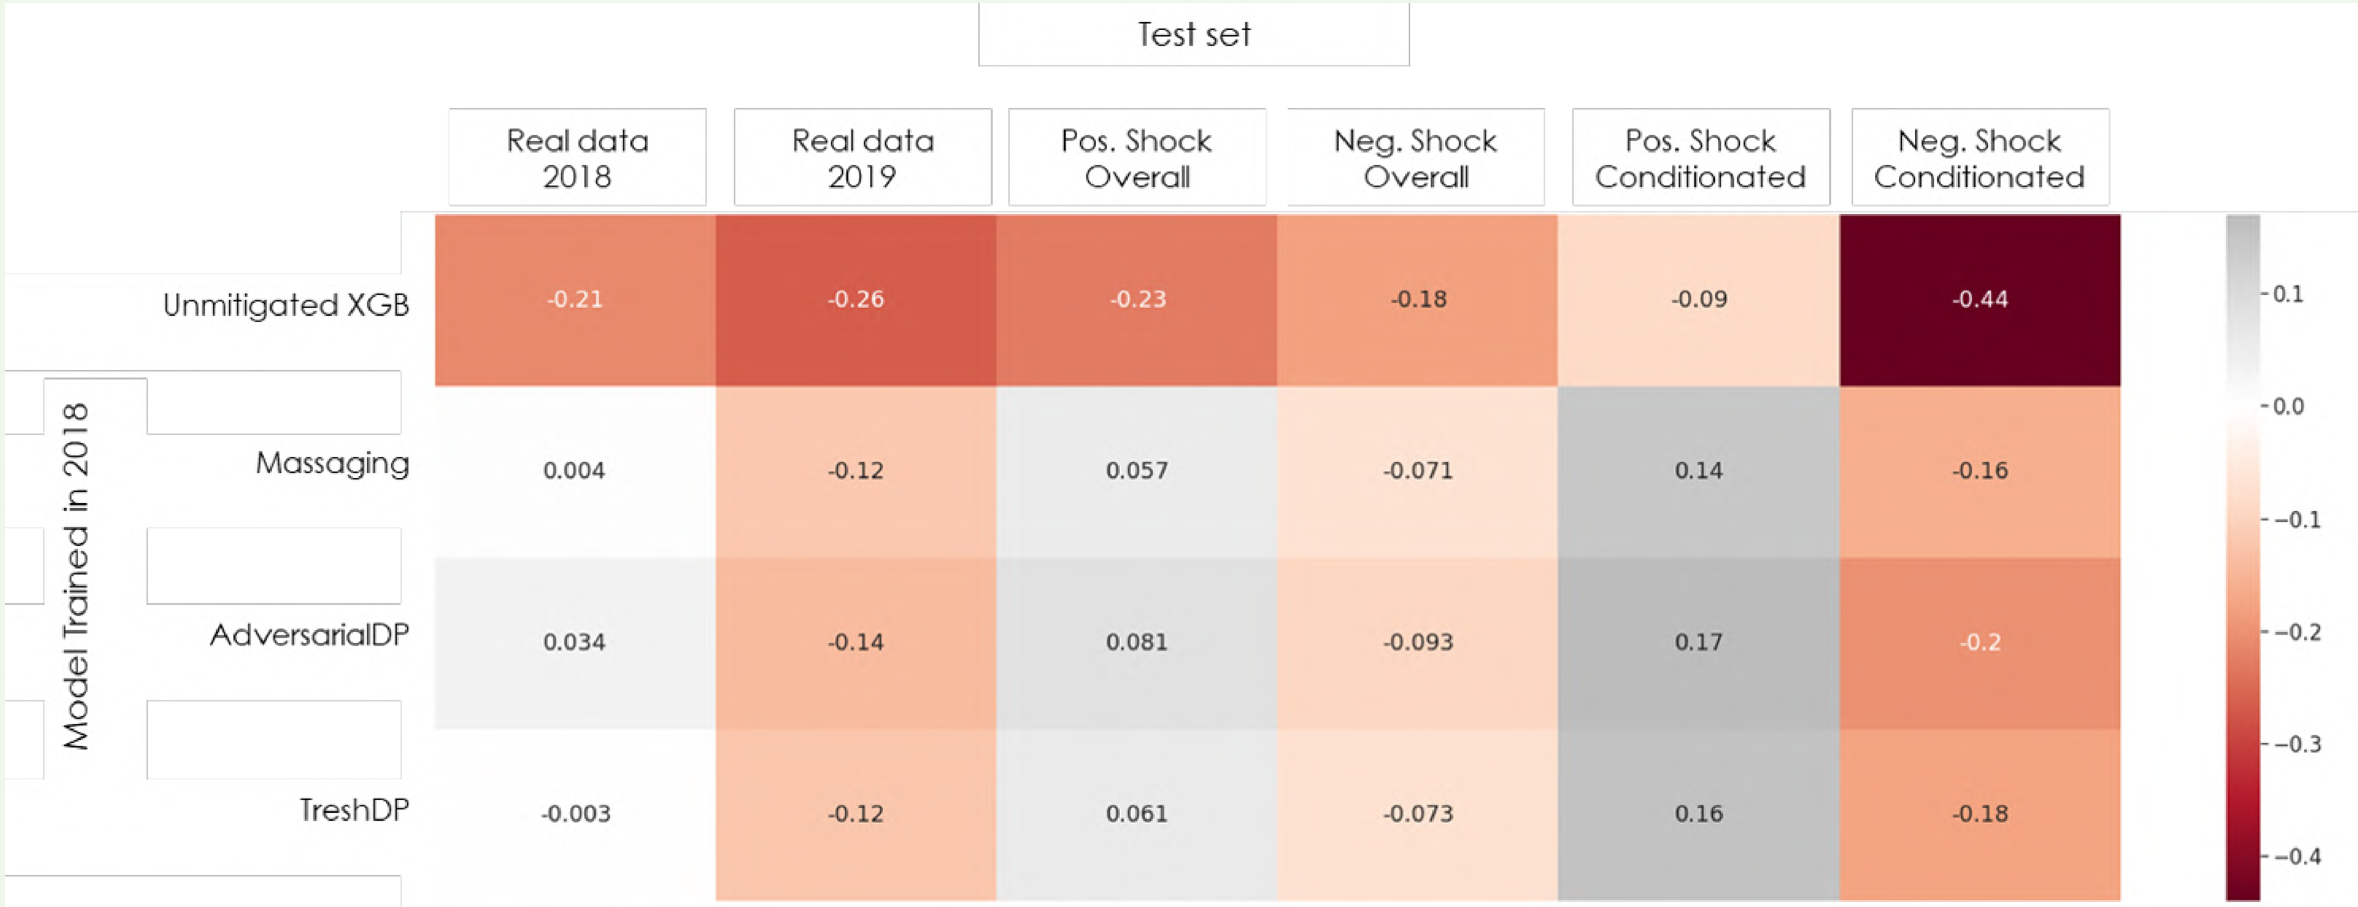

Demographic Parity of the various models tested on different temporal samples or after different stress tests. DP is calculated using citizenship as sensitive feature.

Castelnovo, A., Malandri, L., Mercorio, F., Mezzanzanica, M., and Cosentini, A. Towards fairness through time. In Machine Learning and Principles and Practice of Knowledge Discovery in Databases: International Workshops of ECML PKDD 2021

2nd Challenging Questions

How can XAI techniques be used to verify that the chosen fairness policy (ensure Demographic Parity) is helping to reduce individual discrimination over time?

FairX

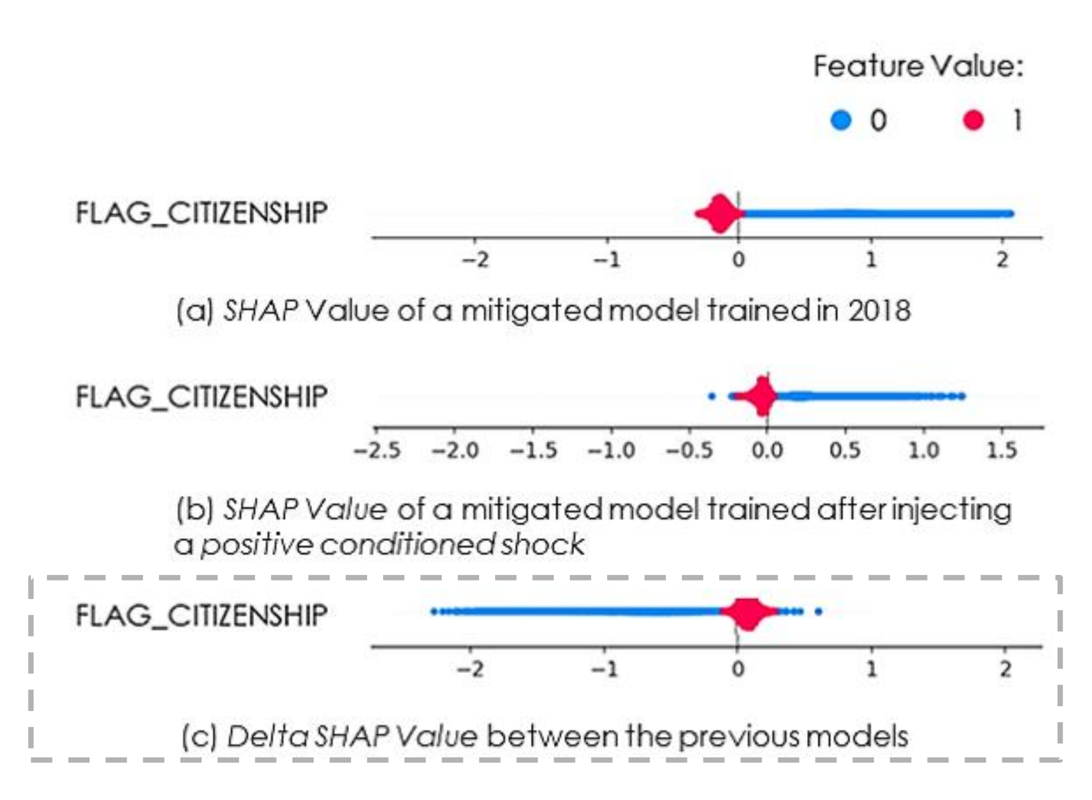

Representation of the Shapley values of two mitigated models trained in 2018, after injecting a positive conditioned shock and the relative differences.

The group mitigation model has to assign a marginal contribution to the vulnerable class to provide demographic parity in the outcome

The marginal contribution on the sensitive variables is a proxy of individual discrimination

SHAP helps to observe the marginal contribution

∆ ***Shapley values*** are reasonable to observe changes in individual discrimination between the two models

A Raodmap for Addressing Fairness in the Banking Sector

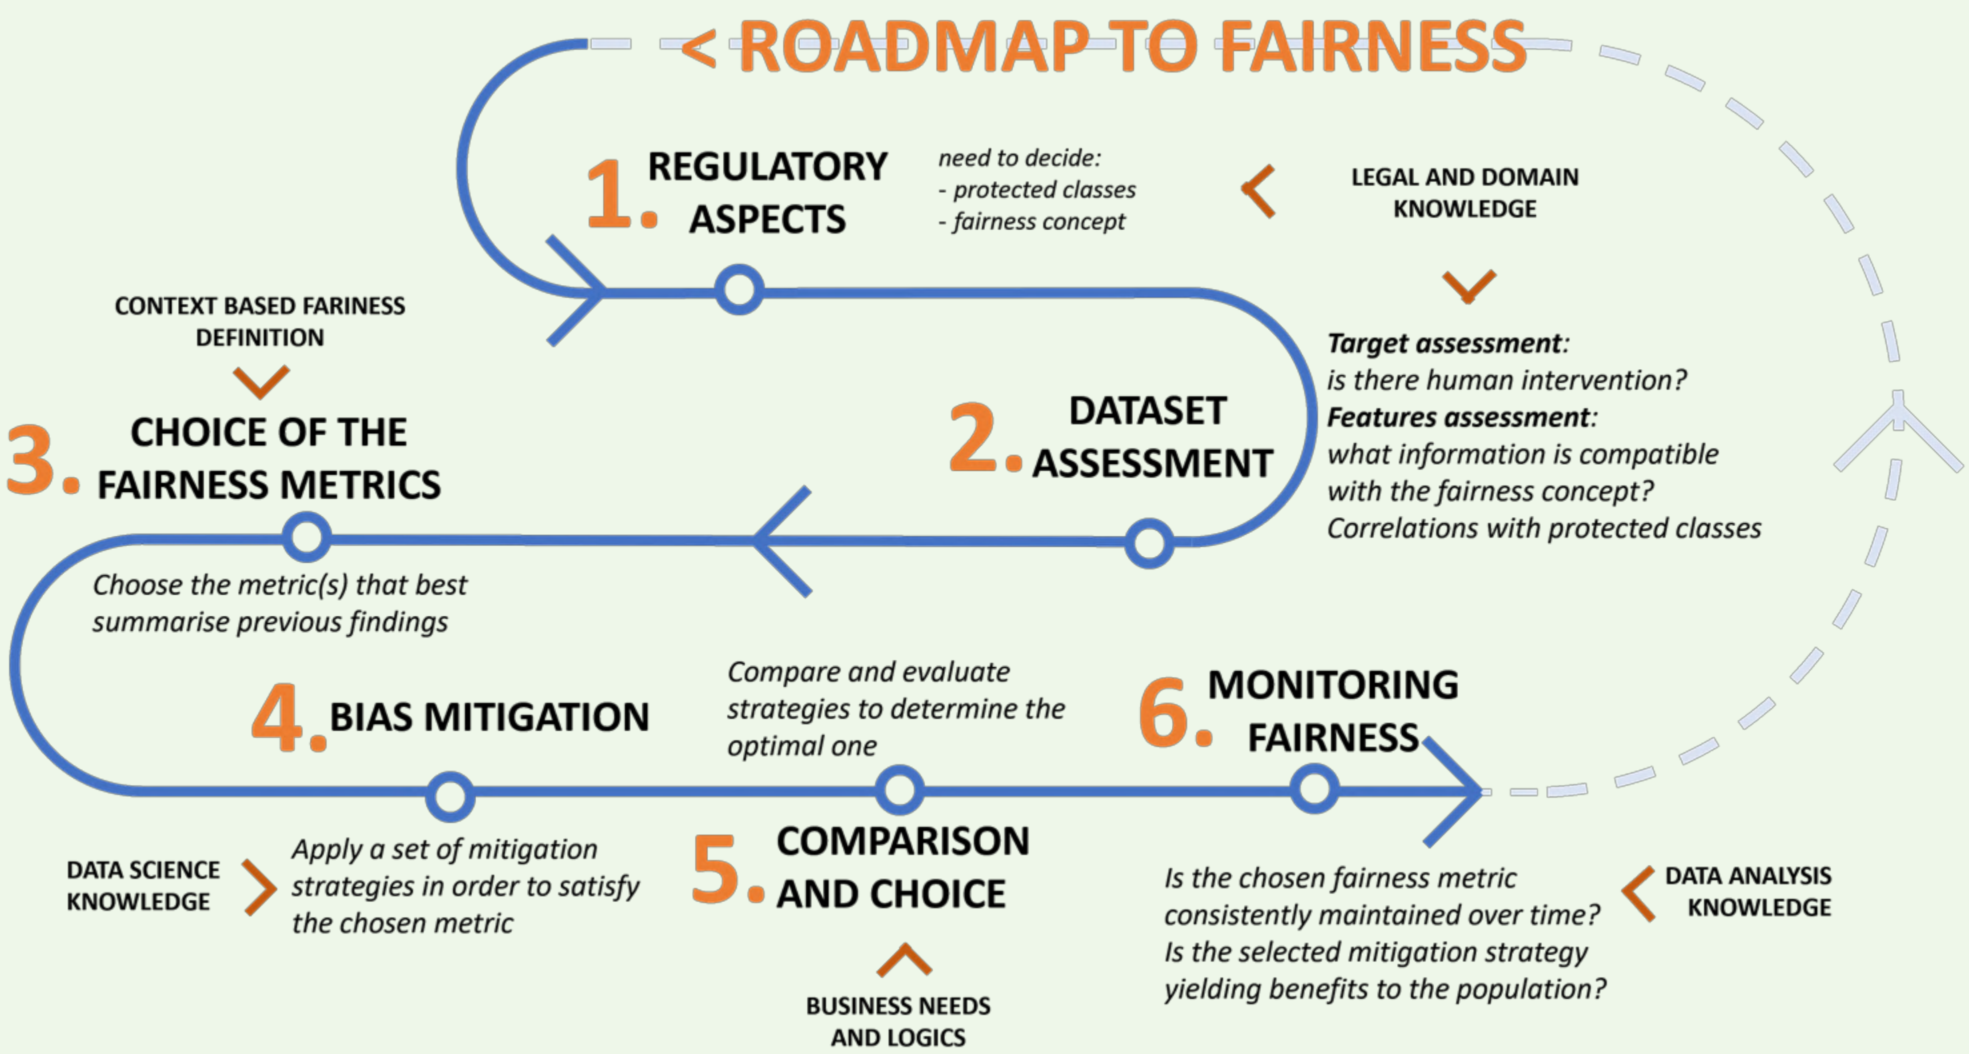

Castelnovo, A., Crupi, R., Del Gamba, G., Greco, G., Naseer, A., Regoli, D., and Gonzalez, B. S. M. (2020). Befair: Addressing fairness in the banking sector. In 2020 IEEE International Conference on Big Data (Big Data)

SCIENTIFIC PUBLICATIONS

[1] **BeFair: Addressing Fairness in the Banking Sector**
Proceedings of the 2020 IEEE International Conference on Big Data (Big Data)

[2] **Towards Fairness Through Time**
Machine Learning and Principles and Practice of Knowledge Discovery in Databases. ECML PKDD 2021.

[3] **Towards Responsible AI: A Design Space Exploration of Human-Centered Artificial Intelligence User Interfaces to Investigate Fairness**
International Journal of Human-Computer Interaction 2022

[4] **A clarification of the nuances in the Fairness metrics landscape**
Scientific Reports 2022

[5] **Counterfactual Explanations as Interventions in Latent Space**
Data Mining & Knowledge Discovery 2022

[6] **FFTree: A flexible tree to handle multiple fairness criteria**
Information Processing & Management 2022

[7] **Bias on Demand: A Modelling Framework that Generates Synthetic Data with Bias.**
Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency

[8] **Declarative Encoding of Fairness in Logic Tensor Networks**
Accepted for publication in the *Proceedings of the 26th European Conference of Artificial Intelligence*

SCIENTIFIC CONFERENCES

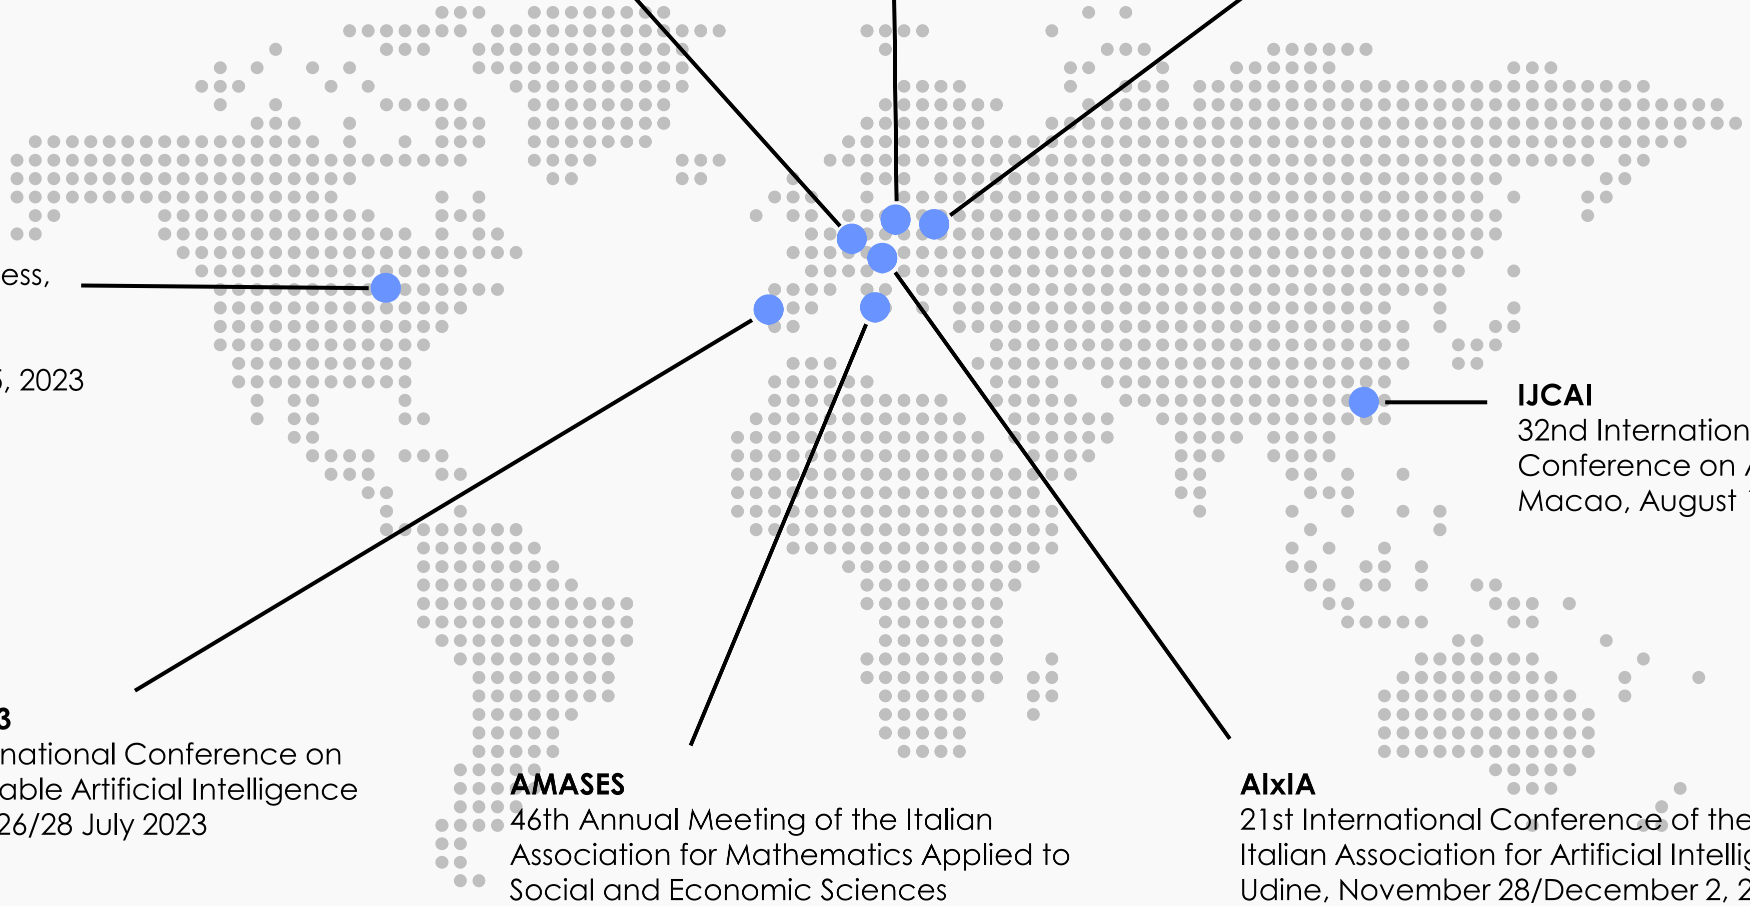

EWAF
European Workshop on Algorithmic Fairness
Winterthur, Switzerland, June 7/9, 2023

IJCAI
31st International Joint Conference on AI
Vienna, July 23/29, 2022

ECAI
26th European Conference on AI
Krakow, Poland, Oct 9/13, 2023

ACM FAccT
Conference on Fairness, Accountability, and Transparency
Chicago, June 12/15, 2023

IJCAI
32nd International Joint Conference on AI
Macao, August 19/25, 2022

xAI 2023
1st International Conference on eXplainable Artificial Intelligence
Lisbon, 26/28 July 2023

AMASES
46th Annual Meeting of the Italian Association for Mathematics Applied to Social and Economic Sciences
Palermo, 22/28 September 2022

AIxIA
21st International Conference of the Italian Association for Artificial Intelligence
Udine, November 28/December 2, 2022

Thank you