Introduction to Explainable Artificial Intelligence (XAI)

Szymon Bobek

Jagiellonian University
2023

https://geist.re

Will this class be interesting?

When we use AI, we agree on *something*

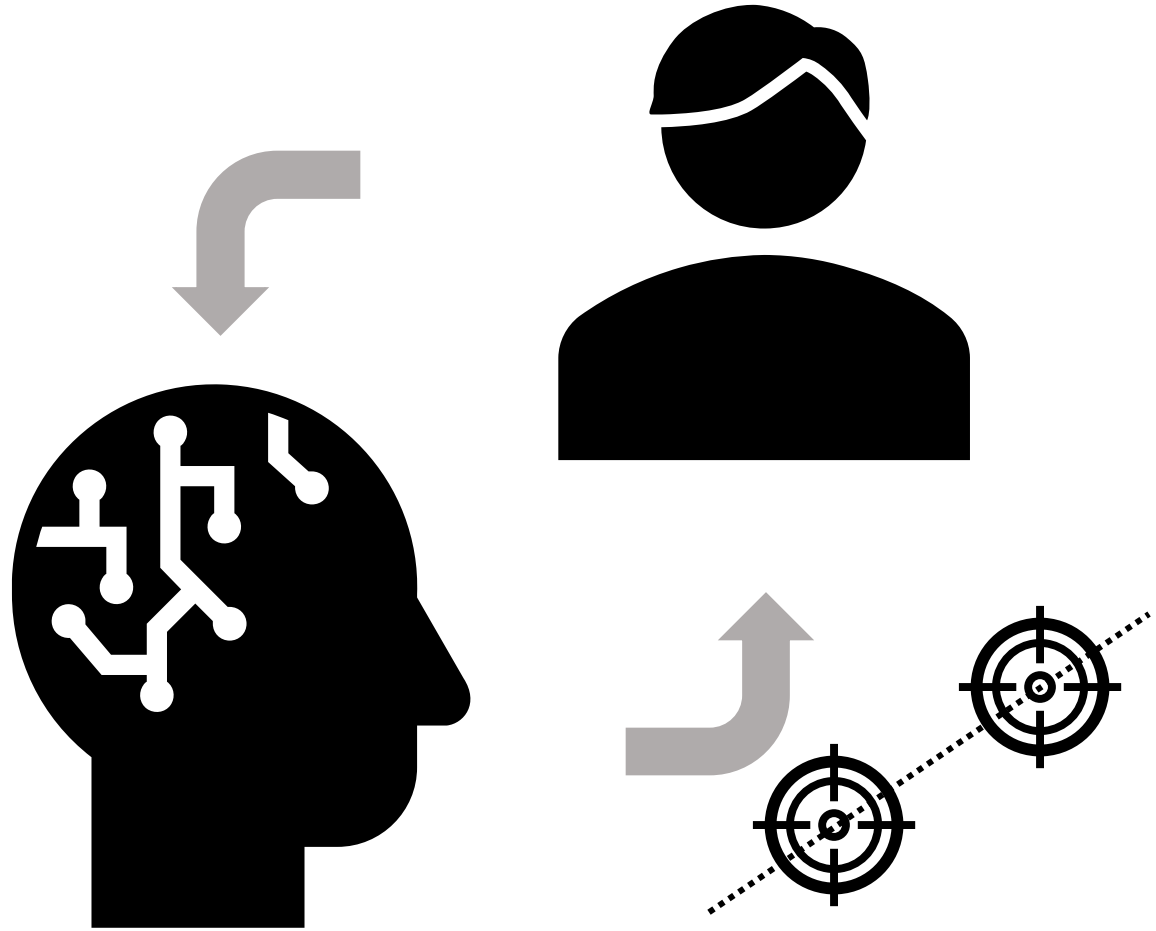

$$\max_{\theta} \ell\ell(\theta) = \max_{\theta} \sum_{i=1}^{N} \left[\mathbb{1}[y = +1] ln \frac{1}{1 + e^{-\theta x^{(i)}}} + \mathbb{1}[y = -1] ln \frac{e^{-\theta x^{(i)}}}{1 + e^{-\theta x^{(i)}}} \right]$$

When we use AI, we agree on *something*

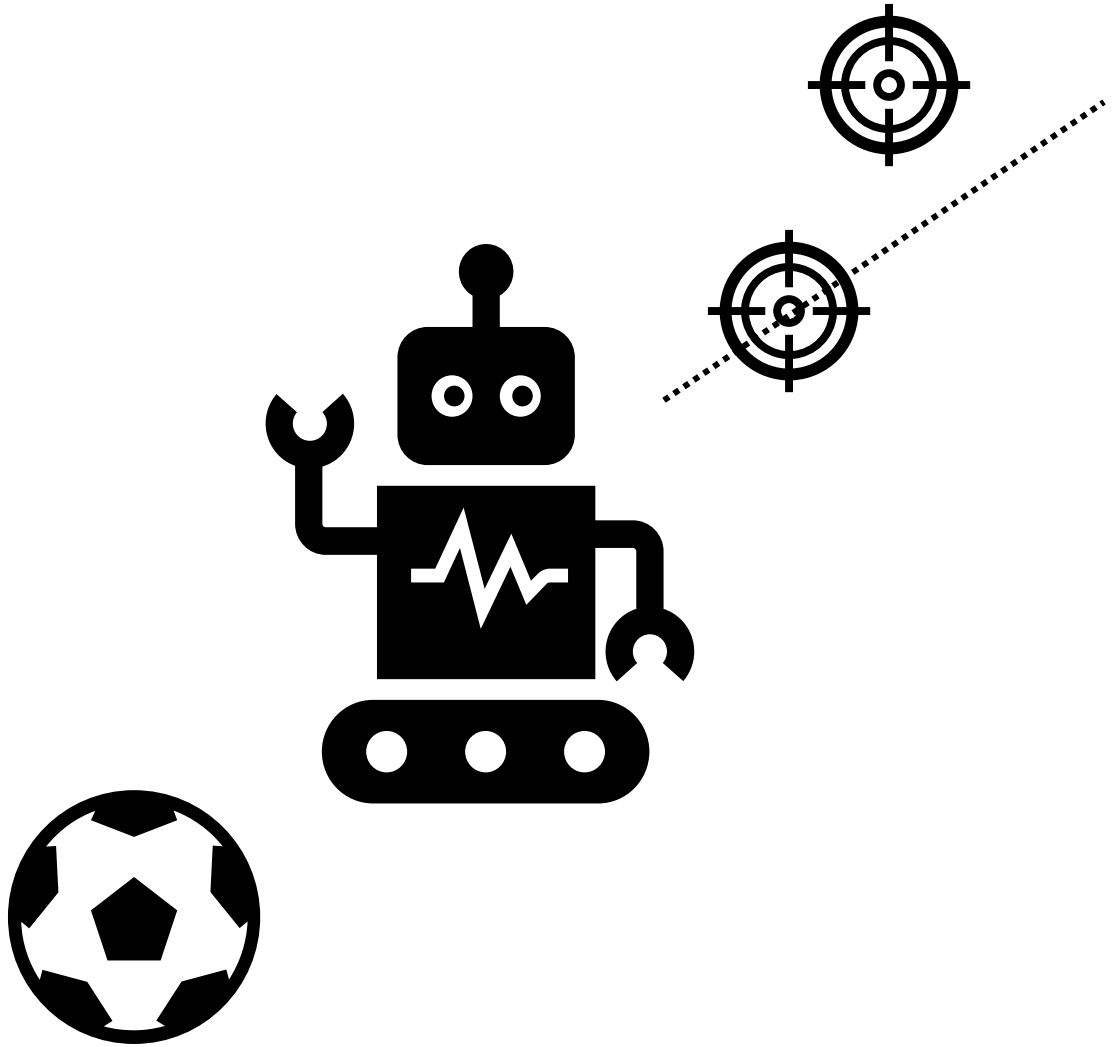

- Large number of examples: https://vkrakovna.wordpress.com/ai-safety-resources/
- Stuart Russel: *Human Compatible: Artificial Intelligence and the Problem of Control*
- Biran Christian: The Alignment Problem: Machine Learning and Human Values
- *Max Tegmark: Life 3.0*

What can go wrong?

Paperclip Theory -- If you instructed a machine to optimize its paperclip production, it would eventually resort to dismantling objects such as computers, refrigerators, or any metal-based items once it depletes alternative sources of metal. This phenomenon is referred to as instrumental convergence.

https://nickbostrom.com/ethics/ai

Source: https://openai.com/research/faulty-reward-functions, 2016

What can go wrong?

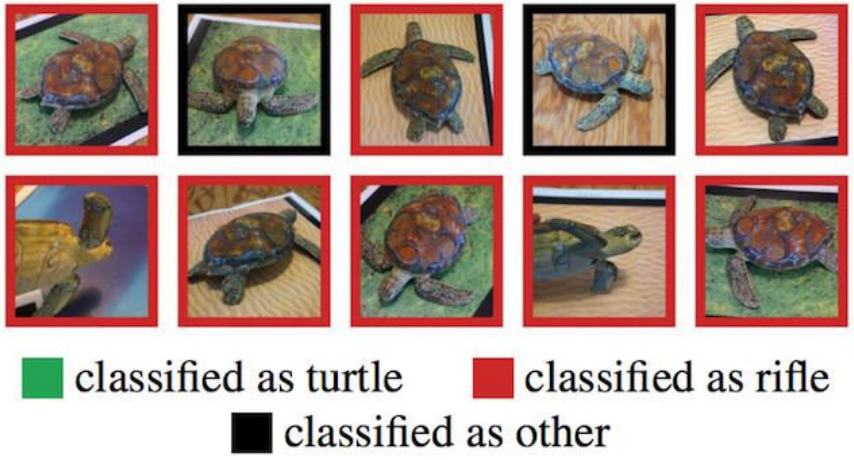

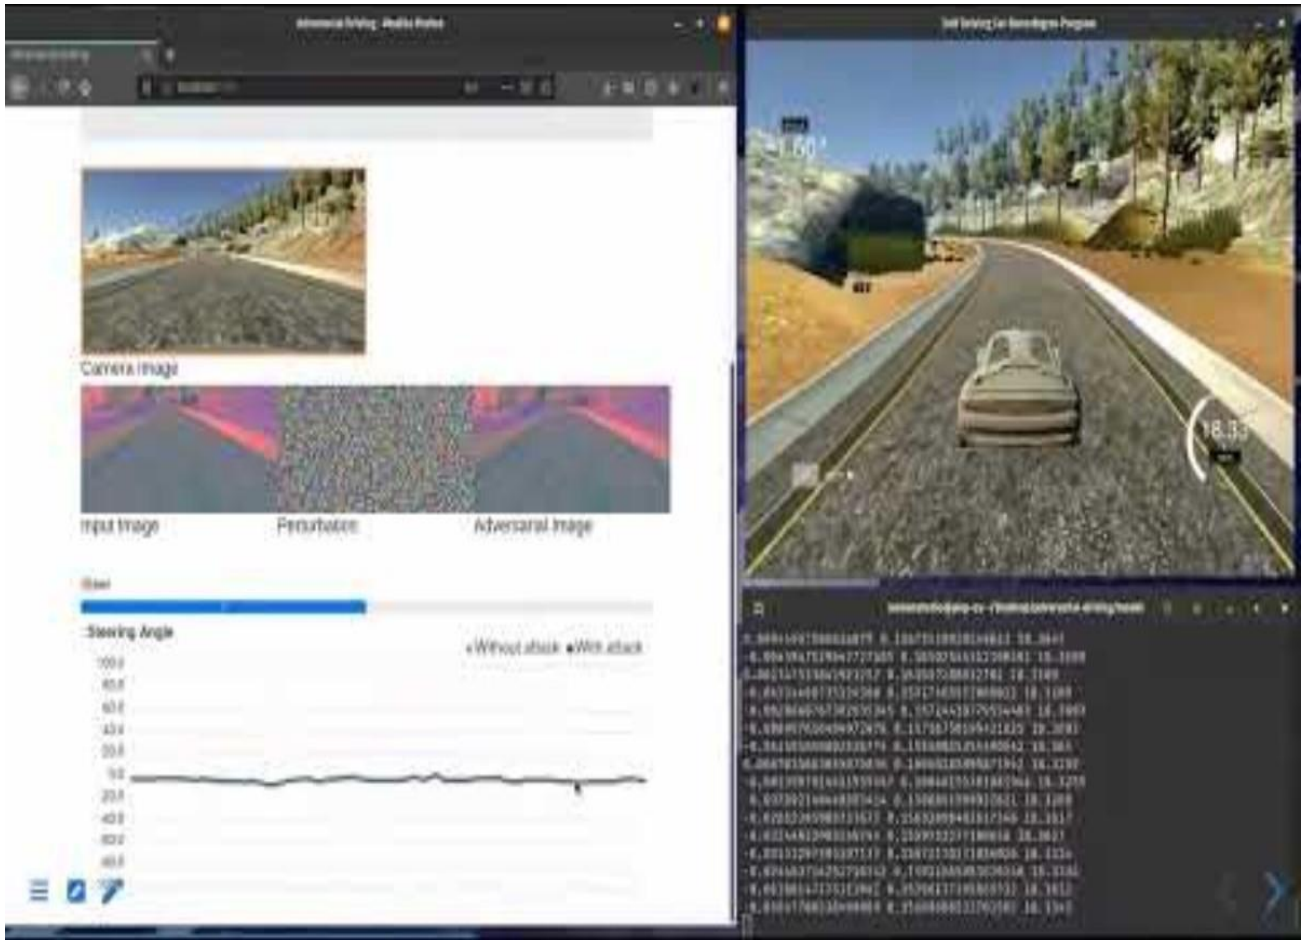

H. Wu, S. Yunas, S. Rowlands, W. Ruan and J. Wahlström, "Adversarial Driving: Attacking End-to-End Autonomous Driving," 2023 IEEE Intelligent Vehicles Symposium (IV), Anchorage, AK, USA, 2023, pp. 1-7, doi: 10.1109/IV55152.2023.10186386.

"But if machines are more intelligent than humans, then giving them the wrong objective would basically be setting upa kind of a chess match between humanity and a machine [...].

And **we wouldn't win** that chess match."

Stuart Russel

"The development of full artificial intelligence could **spell the end** of the human race"

Stephen Hawking

"If we build these devices to take care of everything for us, eventually they'll think faster than us and they'll **get rid of the slow humans** to run companies more efficiently."

Steve Wozniak

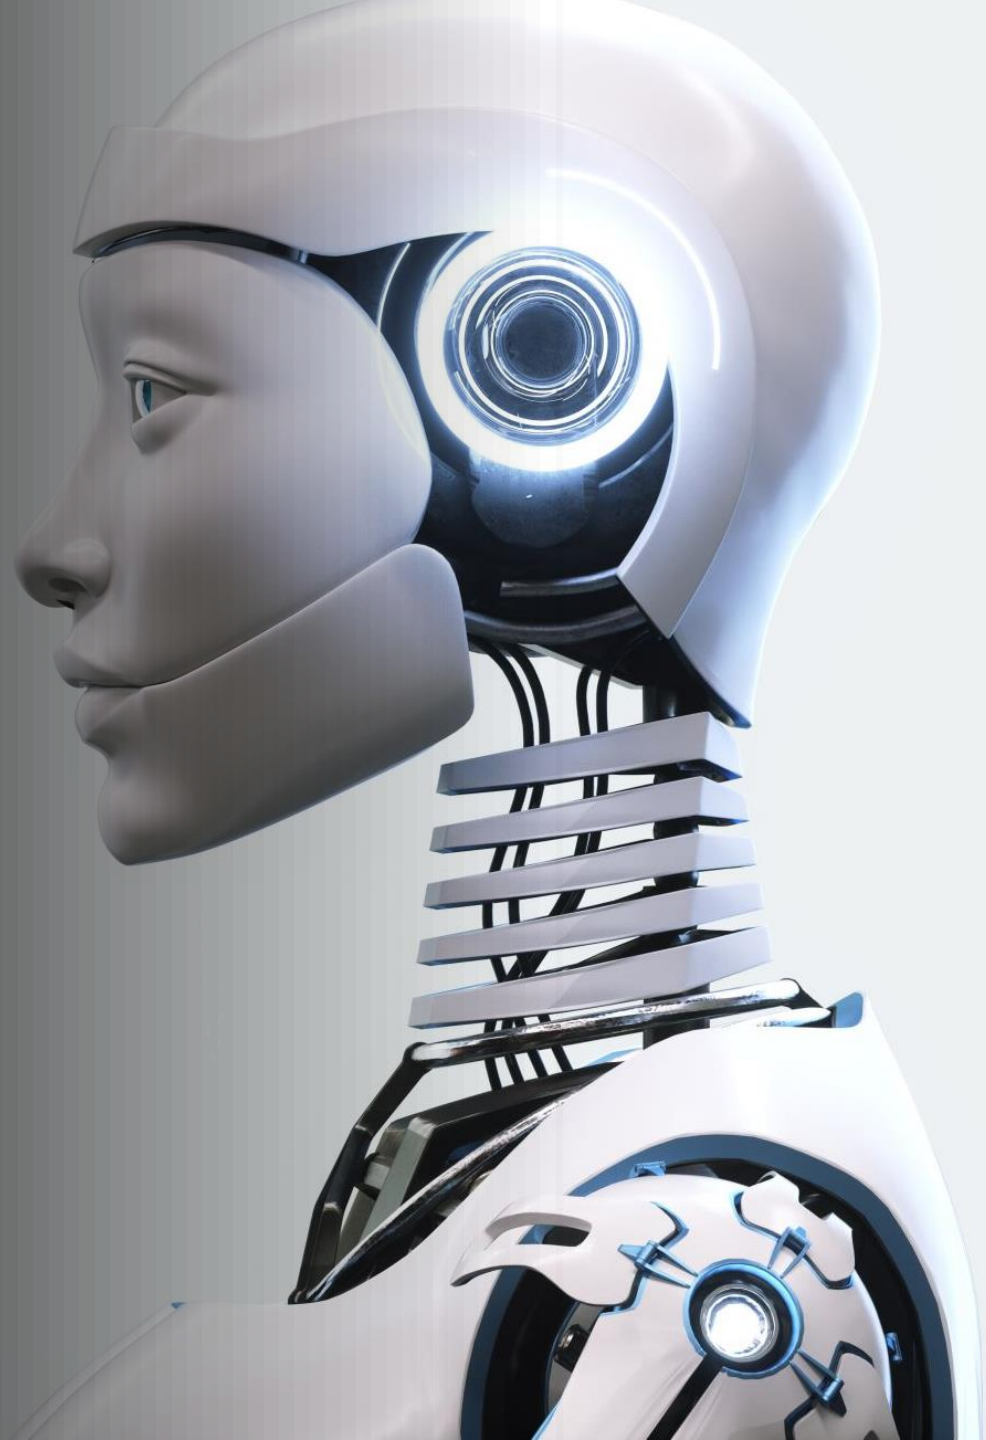

We need to know why…

- Why model made such decisions?
- What influenced the model decision mostly?
- What should be changed to change the model?
- What should be changed to change the model decision?

Otherwise, we will not be able to efficiently monitor and control their behaviour!

Where the true threat and opportunities lie?

- We should worry, but not panic
- Until 2018, it was enough to hide behind the door
- The threat is in the decision-making area, which is unrelated to the robots' motor skills or their ability to open doors.
- General AI – the last invention of humanity. And when will we invent it?
  - John McCarthy answered it very precisely in 1977 ;)
  - Stuart Russell – a few more breakthroughs in AI are needed
- Let's use AI to understand more - then there is a chance that we will make better AI in the future.

Source: https://en.wikipedia.org/wiki/STM_Kargu (CC BY 4.0)
STM Kargum, First flight 2017 – image recognition module to attack individuals

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Source: https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

High-stake decissions taken based on ML model predictions.
No dork-knob ability required to do huge harm

(2015 rok) Outdated, but still funny
https://youtu.be/g0TaYhjpOfo

XAI and ML and bias

- COMPAS -- Correctional Offender Management Profiling for Alternative Sanctions system
- It was actuaslly deployed in US and used by judges
- It was a black-box model

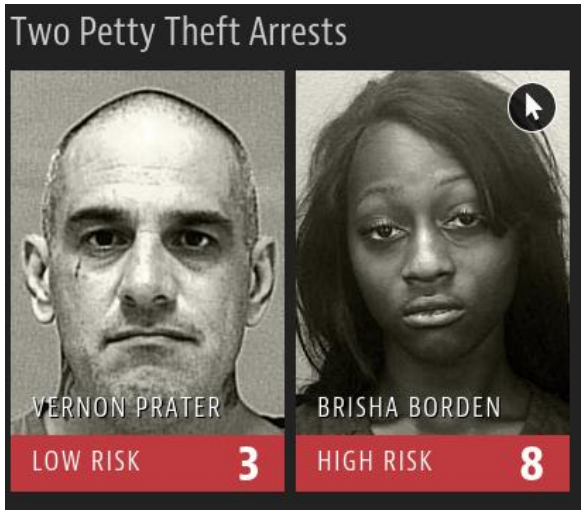

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

XAI and ML and bias

- COMPAS -- Correctional Offender Management Profiling for Alternative Sanctions system
- It was actuaslly deployed in US and used by judges
- It was a black-box model

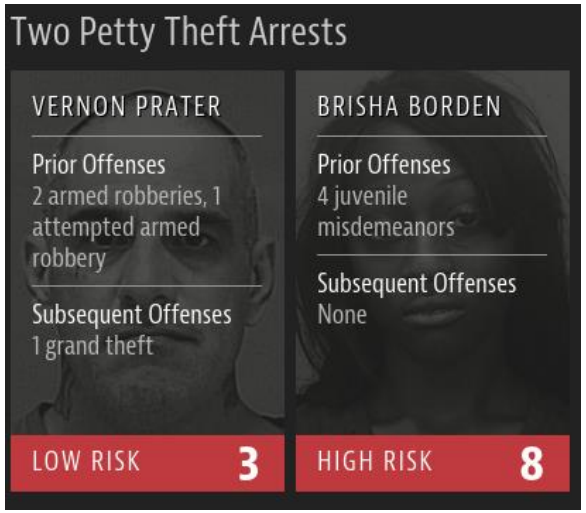

https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

XAI and ML and bias

- Microsoft racist chatbot Tay lasted 16 hours and being shut down for racist comments
- Google's still cannot find gorillas (neither MS or apple) after issue reported by Jacky Alciné in 2015 (as of May 2023)
- Jigsaw (part of Alphabet) released dataset which purpose was comment toxicity prediction. There is unintended bias in the data posed by humans
- What else can go wrong?

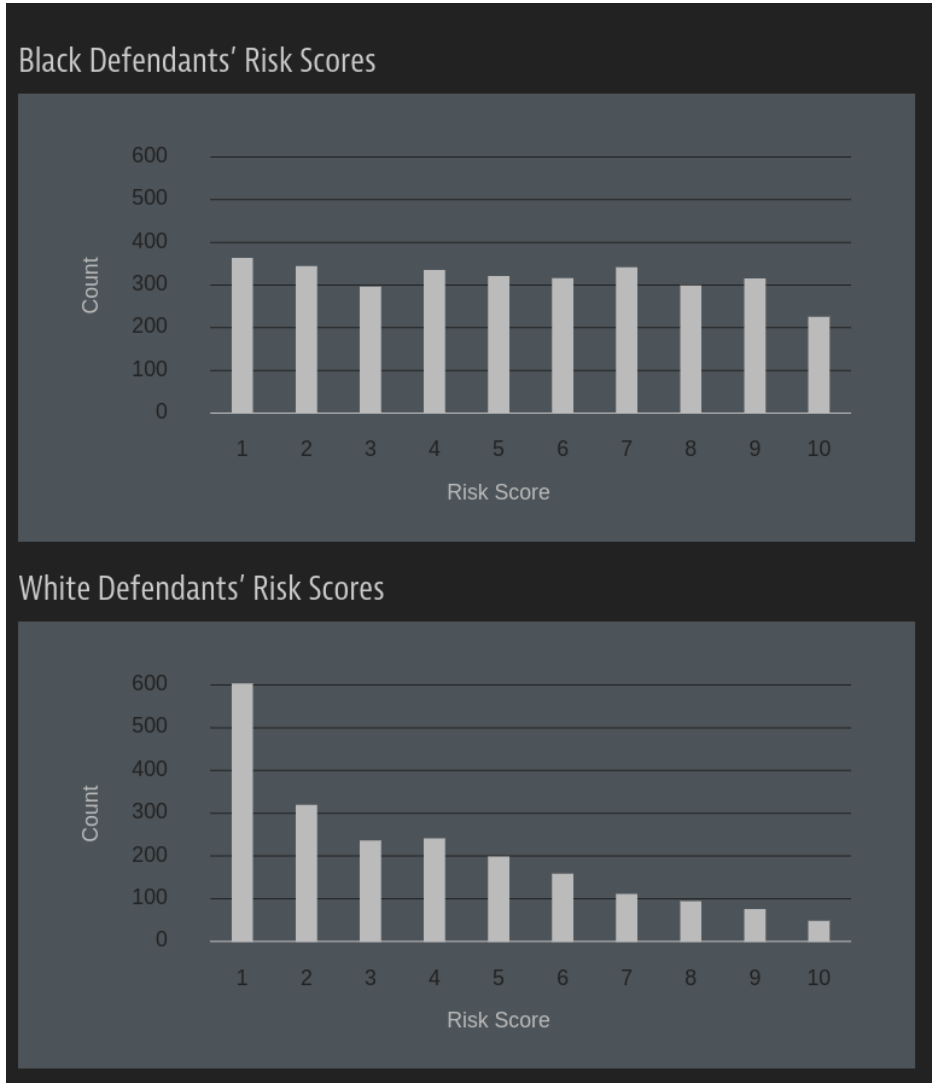

Brief (not full) AI history

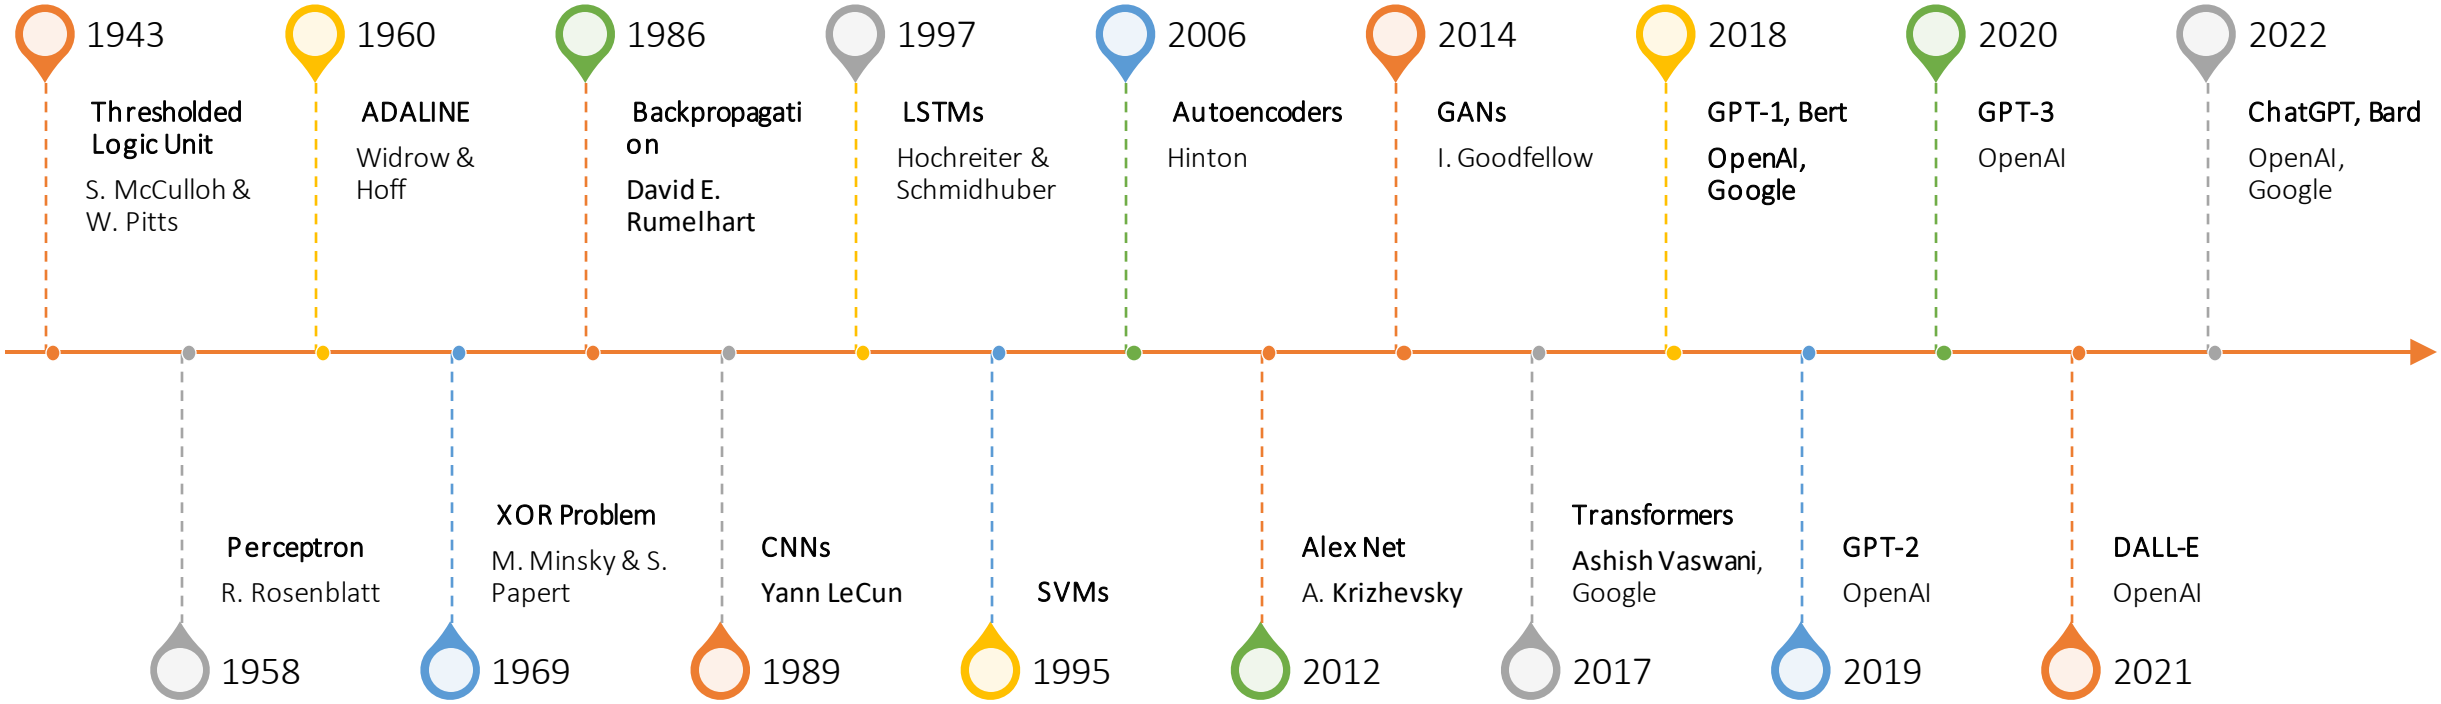

Open letter, DARPA, RODO, AI Act

OPEN-LETTER

Research Priorities for Robust and Beneficial Artificial Intelligence: An Open Letter

There is now a broad consensus that AI research is progressing steadily, and that its impact on society is likely to increase. The potential benefits are huge, since everything that civilization has to offer is a product of human intelligence. Because of the great potential of AI, it is important to research how to reap its benefits while avoiding potential pitfalls.

Signatures
11251

Add your signature

PUBLISHED
October 28, 2015

https://futureoflife.org/open-letter/ai-open-letter/

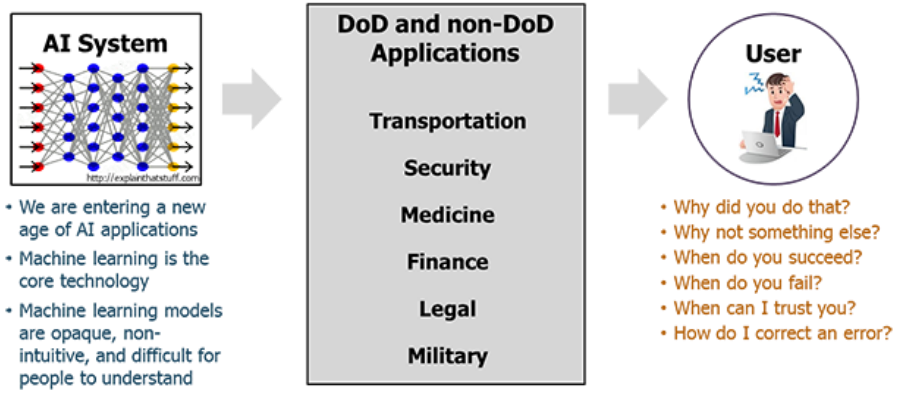

(2016)
https://www.darpa.mil/program/explainable-artificial-intelligence

RODO

"...the data subject and the right to obtain human intervention, to express his or her point of view, to obtain an explanation of the decision reached after such assessment and to challenge the decision"

(Released) 2016

AI Act

Legal framework for building AI systems (under developpment

This is not a new idea

To explain an event is to provide some information about its causal history.
In an act of explaining, someone who is in possession of some information about the causal history of some event - explanatory information, I shall call it - tries to convey it to someone else. - David Lewis

Different approaches

- Intelligibility of the system
- Interpretability of models
- Explainability of ML models

Old topic

- Expert systems
- Recommender systems
- Context-aware systems
- Machine learning

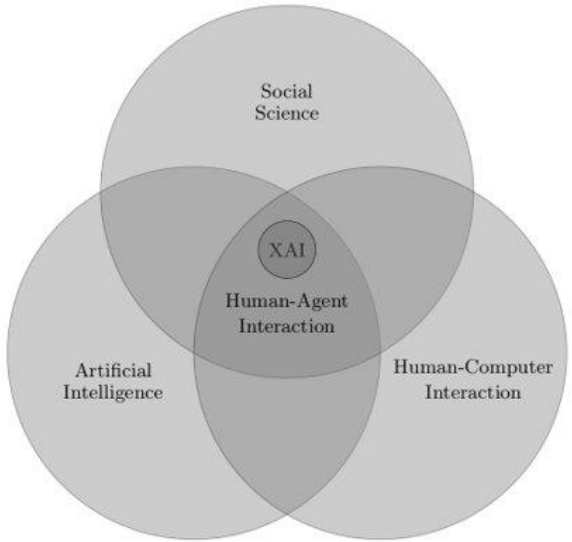

XAI and ML

XAI

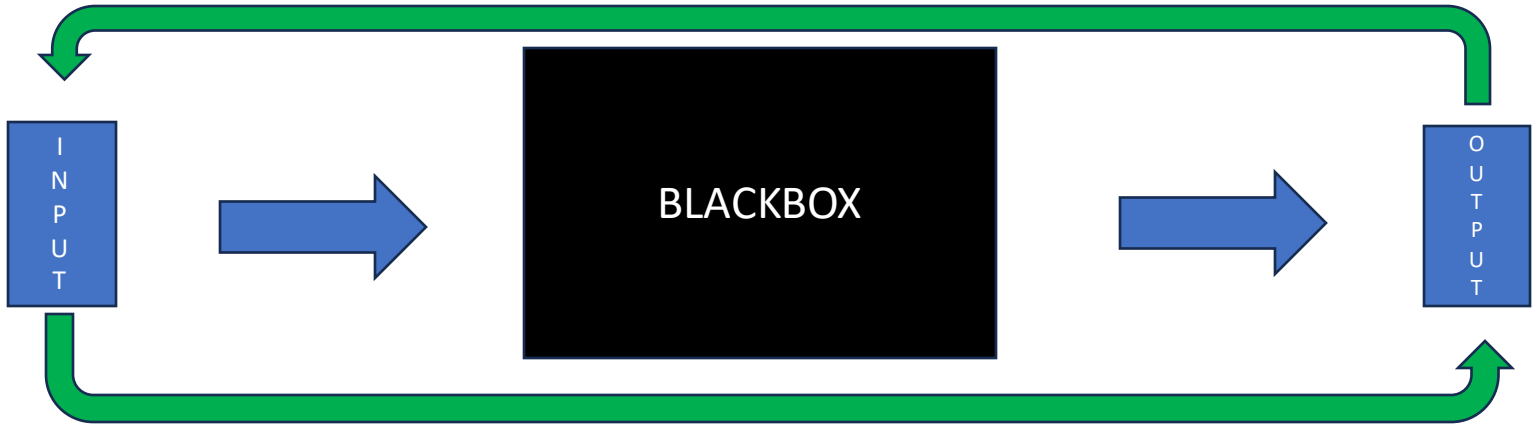

Glassbox

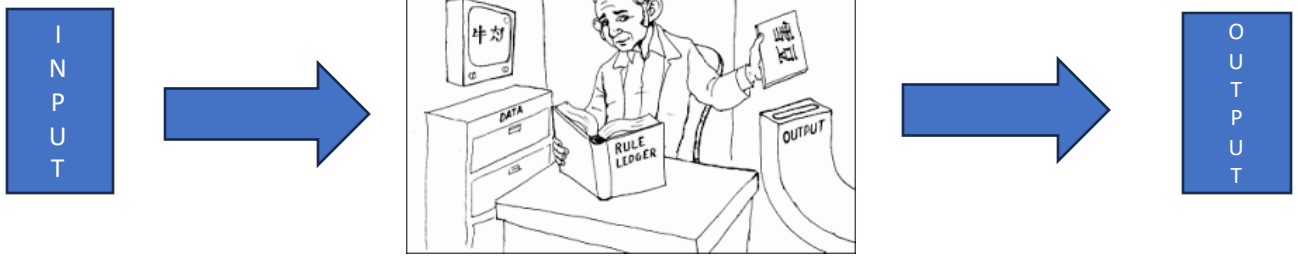

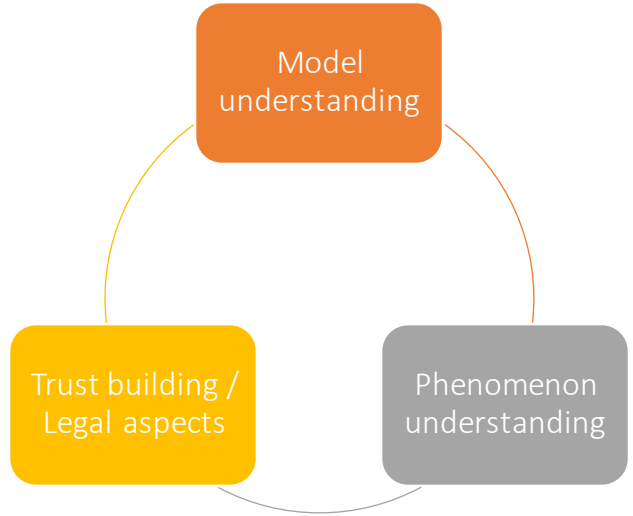

Why XAI is non trivial

*In an **act** of explaining, **someone** who is in possession of **some information***

Artificial intelligence / XAI

Most often feature importance

*about the **causal history of some event** - explanatory information,*

Why input to the model generated such output

*I shall call it - tries to **convey it to someone else**.*

Human

XAI and ML

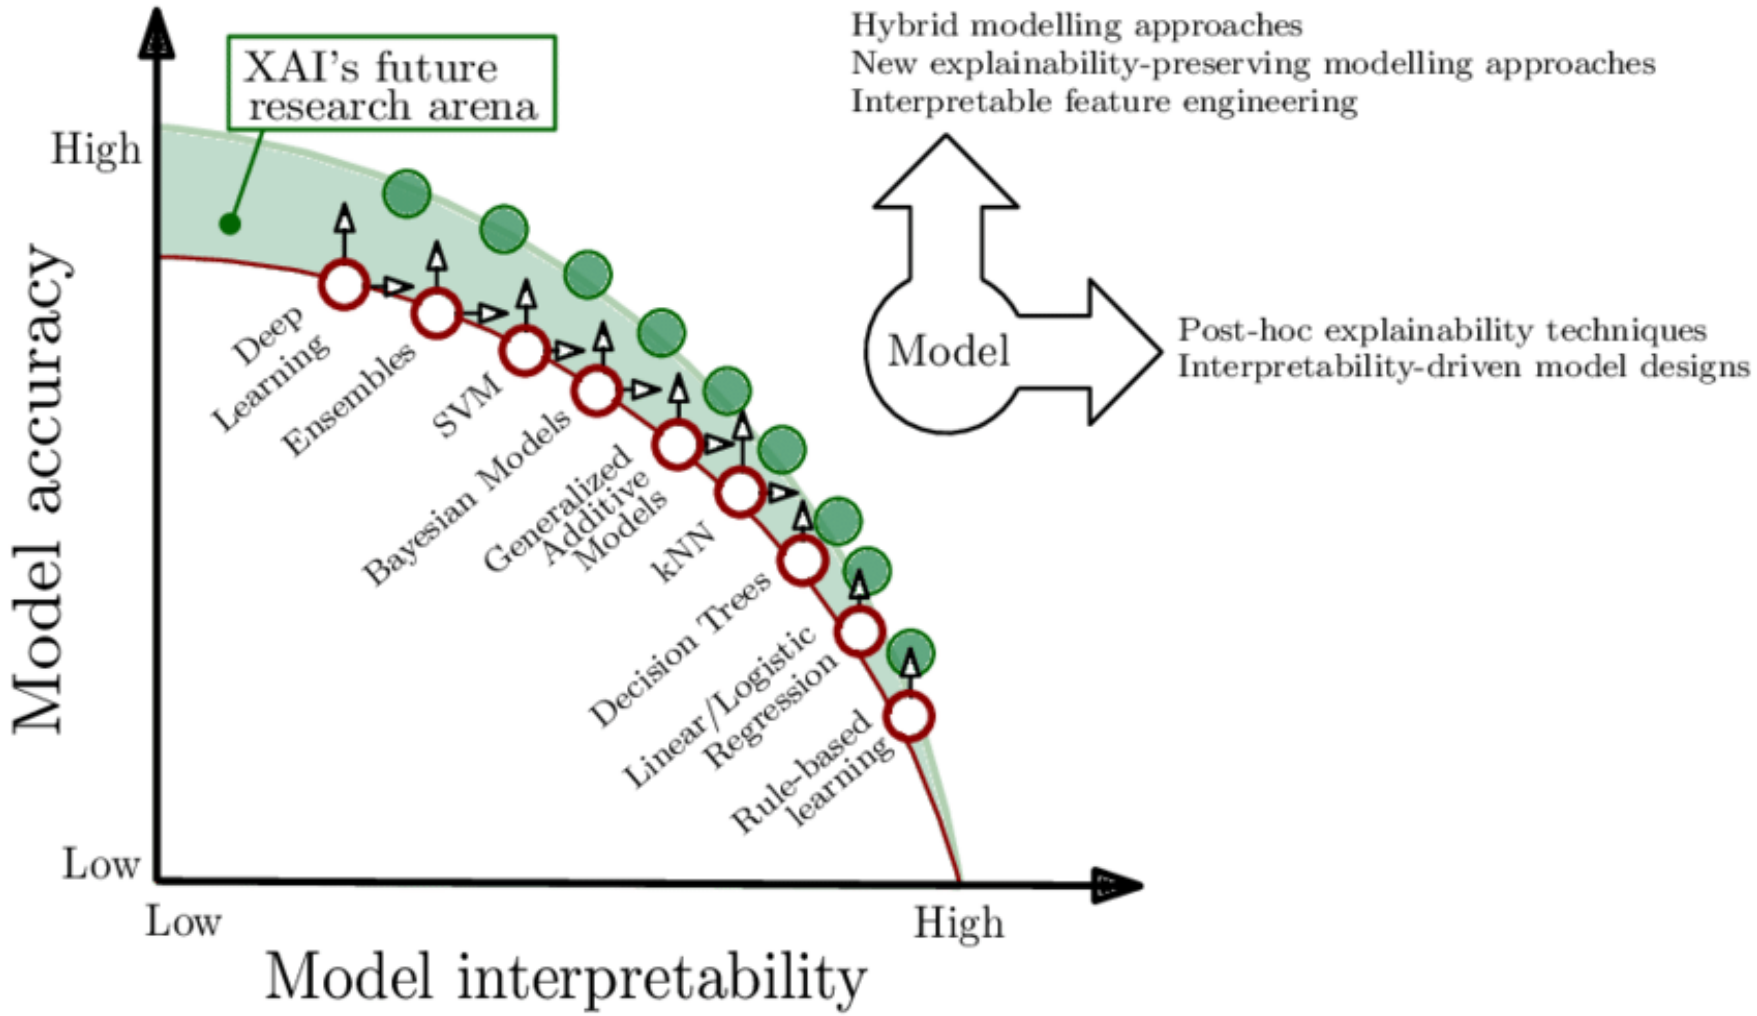
XAI's future research arena
High
Model accuracy
Low
Low
High
Model interpretability
Deep Learning
Ensembles
SVM
Bayesian Models
Generalized Additive Models
kNN
Decision Trees
Linear/Logistic Regression
Rule-based learning
Hybrid modelling approaches
New explainability-preserving modelling approaches
Interpretable feature engineering
Model
Post-hoc explainability techniques
Interpretability-driven model designs

XAI and ML and bias

Historical bias

- Historical bias occurs when the state of the world in which the data was generated is flawed.

Representation bias

- Representation bias occurs when building datasets for training a model, if those datasets poorly represent the people that the model will serve.

Measurement bias

- Measurement bias occurs when the accuracy of the data varies across groups. This can happen when working with proxy variables (variables that take the place of a variable that cannot be directly measured), if the quality of the proxy varies in different groups.

Aggregation bias

- Aggregation bias occurs when groups are inappropriately combined, resulting in a model that does not perform well for any group or only performs well for the majority group. (This is often not an issue, but most commonly arises in medical applications.)

Evaluation bias

- Evaluation bias occurs when evaluating a model, if the benchmark data (used to compare the model to other models that perform similar tasks) does not represent the population that the model will serve.

Deployment bias

- Deployment bias occurs when the problem the model is intended to solve is different from the way it is actually used. If the end users don't use the model in the way it is intended, there is no guarantee that the model will perform well.

Harini Suresh and John Guttag. 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). Association for Computing Machinery, New York, NY, USA, Article 17, 1–9. https://doi.org/10.1145/3465416.3483305

XAI and ML and bias

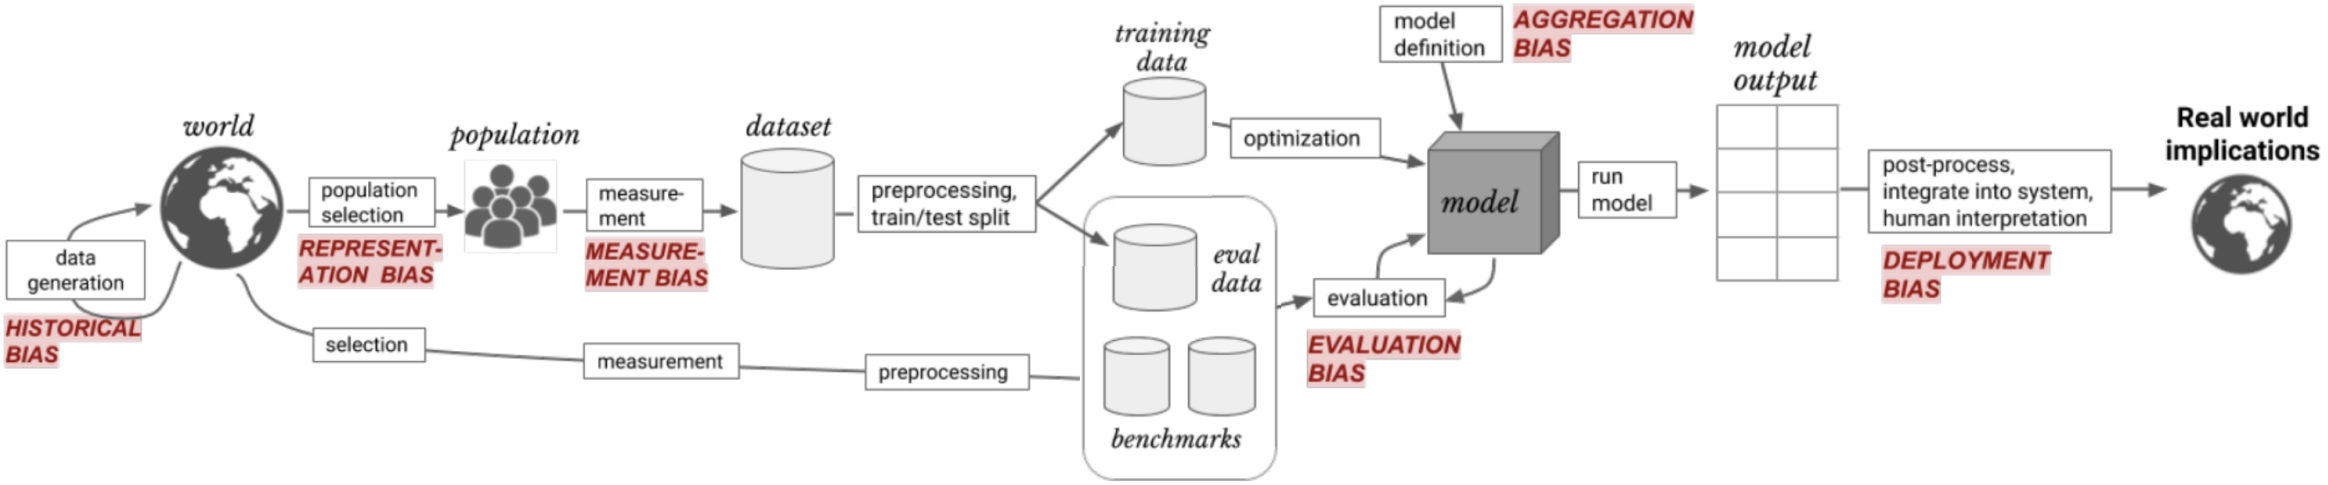

Harini Suresh and John Guttag. 2021. A Framework for Understanding Sources of Harm throughout the Machine Learning Life Cycle. In Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO '21). Association for Computing Machinery, New York, NY, USA, Article 17, 1–9. https://doi.org/10.1145/3465416.3483305

Examples of XAI

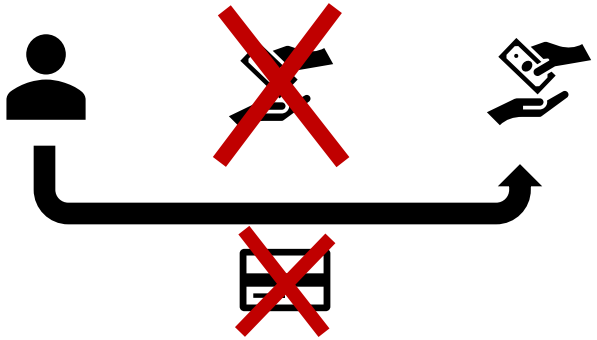

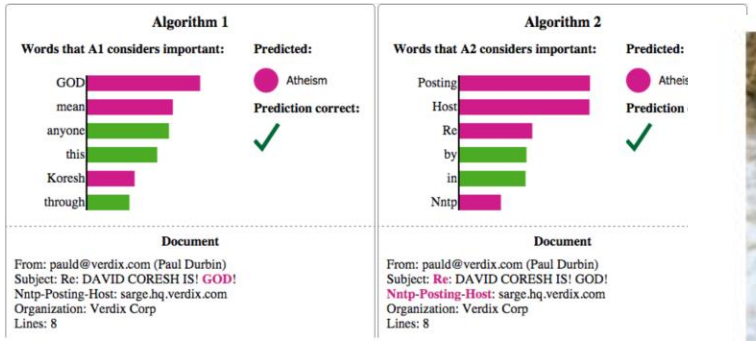

M. T. Ribeiro, S. Singh, and C. Guestrin, '"Why Should I Trust You?": Explaining the Predictions of Any Classifier', in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA, Aug. 2016, pp. 1135–1144. doi: 10.1145/2939672.2939778.

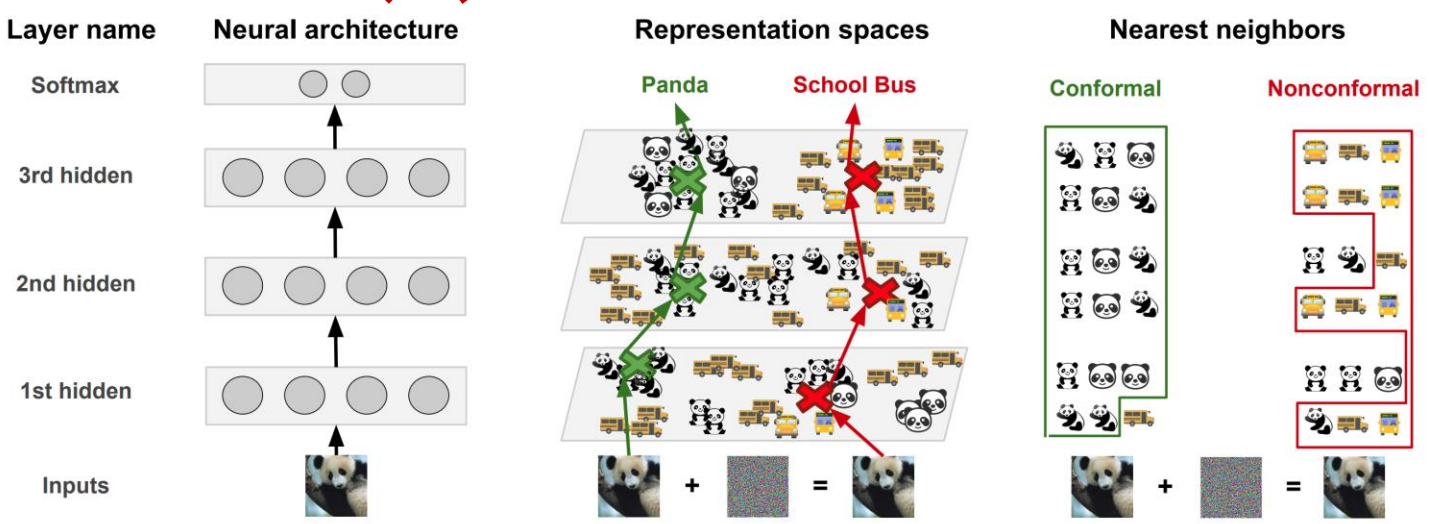

Papernot, N., & Mcdaniel, P. (2018). Deep k-Nearest Neighbors: Towards Confident, Interpretable and Robust Deep Learning. ArXiv, abs/1803.04765.

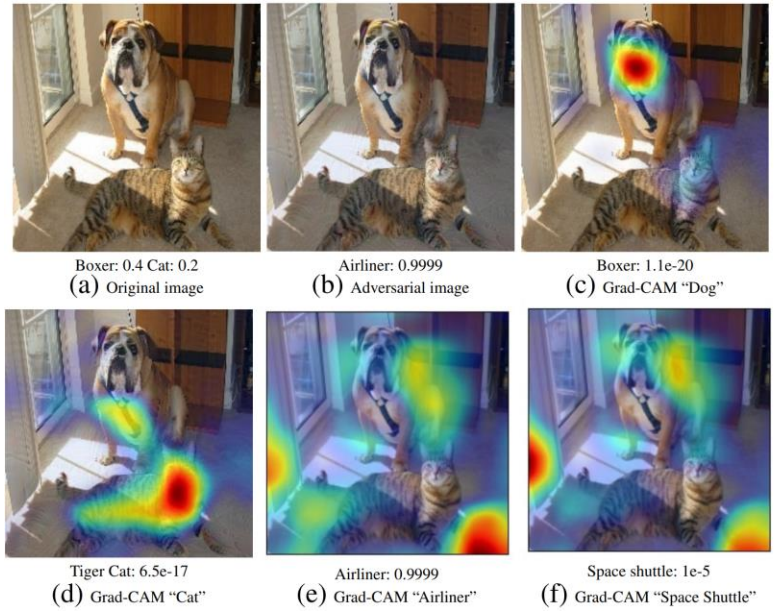

R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, i D. Batra, „Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization", w 2017 IEEE International Conference on Computer Vision (ICCV), paź. 2017, s. 618–626. doi: 10.1109/ICCV.2017.74.

Examples of XAI

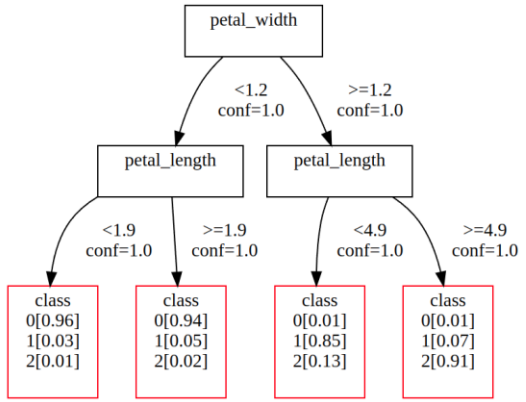

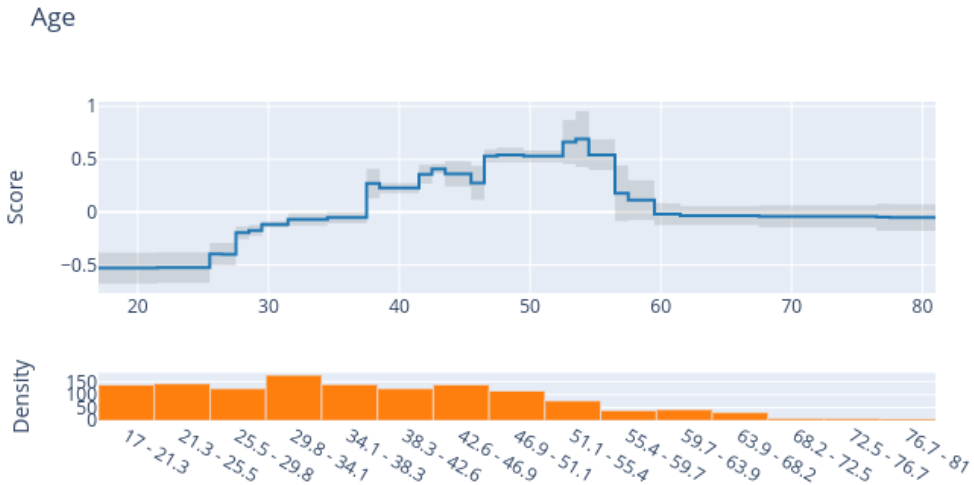

Yin Lou, Rich Caruana, Johannes Gehrke, and Giles Hooker. Accurate intelligible models with pairwise interactions. In Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining, 623–631. 2013.

IF	age between 18-20 and sex is male	THEN predict arrest (within 2 years)
ELSE IF	age between 21-23 and 2-3 prior offenses	THEN predict arrest
ELSE IF	more than three priors	THEN predict arrest
ELSE	predict no arrest.	

Rudin, C. **Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead**. Nat Mach Intell 1, 206–215 (2019). https://doi.org/10.1038/s42256-019-0048-x

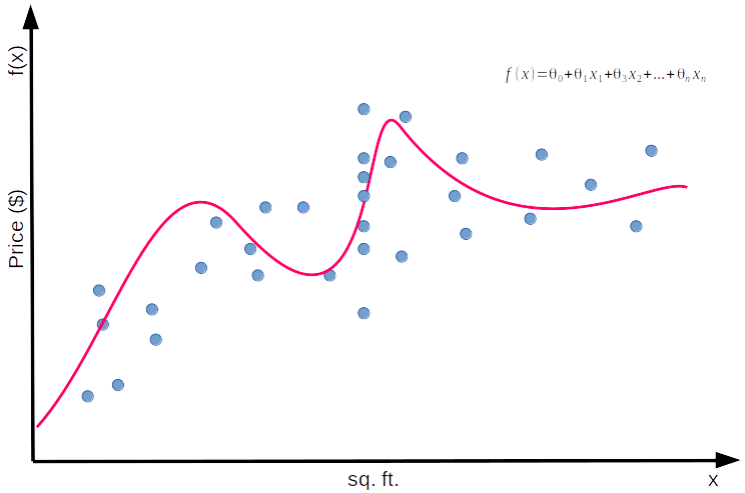

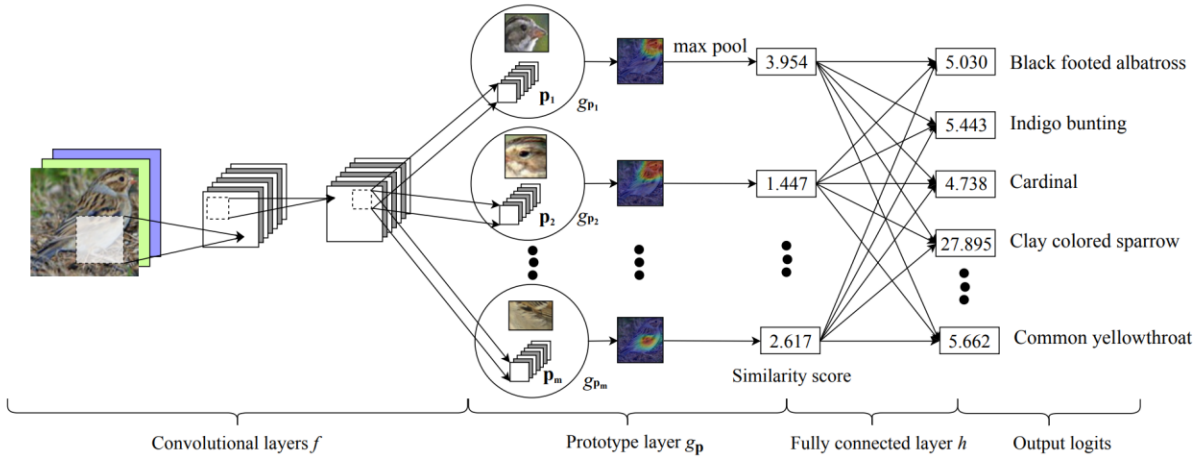

C. Chen, O. Li, C. Tao, A. J. Barnett, J. Su, i C. Rudin, „This Looks Like That: Deep Learning for Interpretable Image Recognition". arXiv, 28 grudzień 2019. doi: 10.48550/arXiv.1806.10574.

What is so difficult?

- An explanation that is not understandable
- An explanation that is intended for someone else
- An explanation that is incorrect
- An explanation that is correct, but not true!

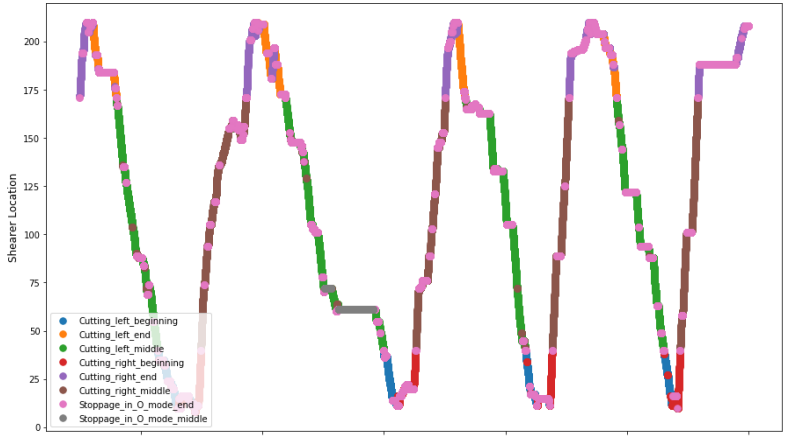

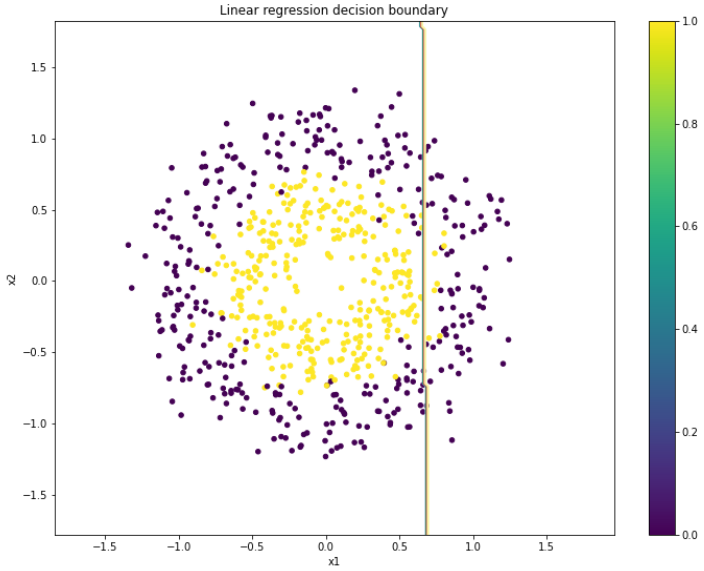

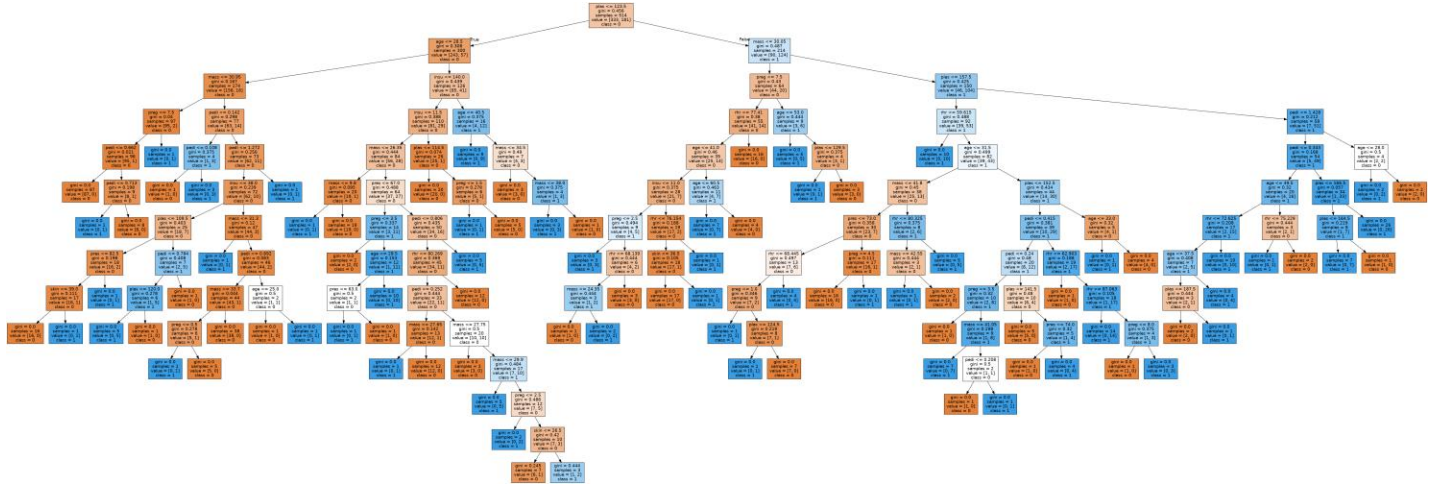

Local vs Global explanations

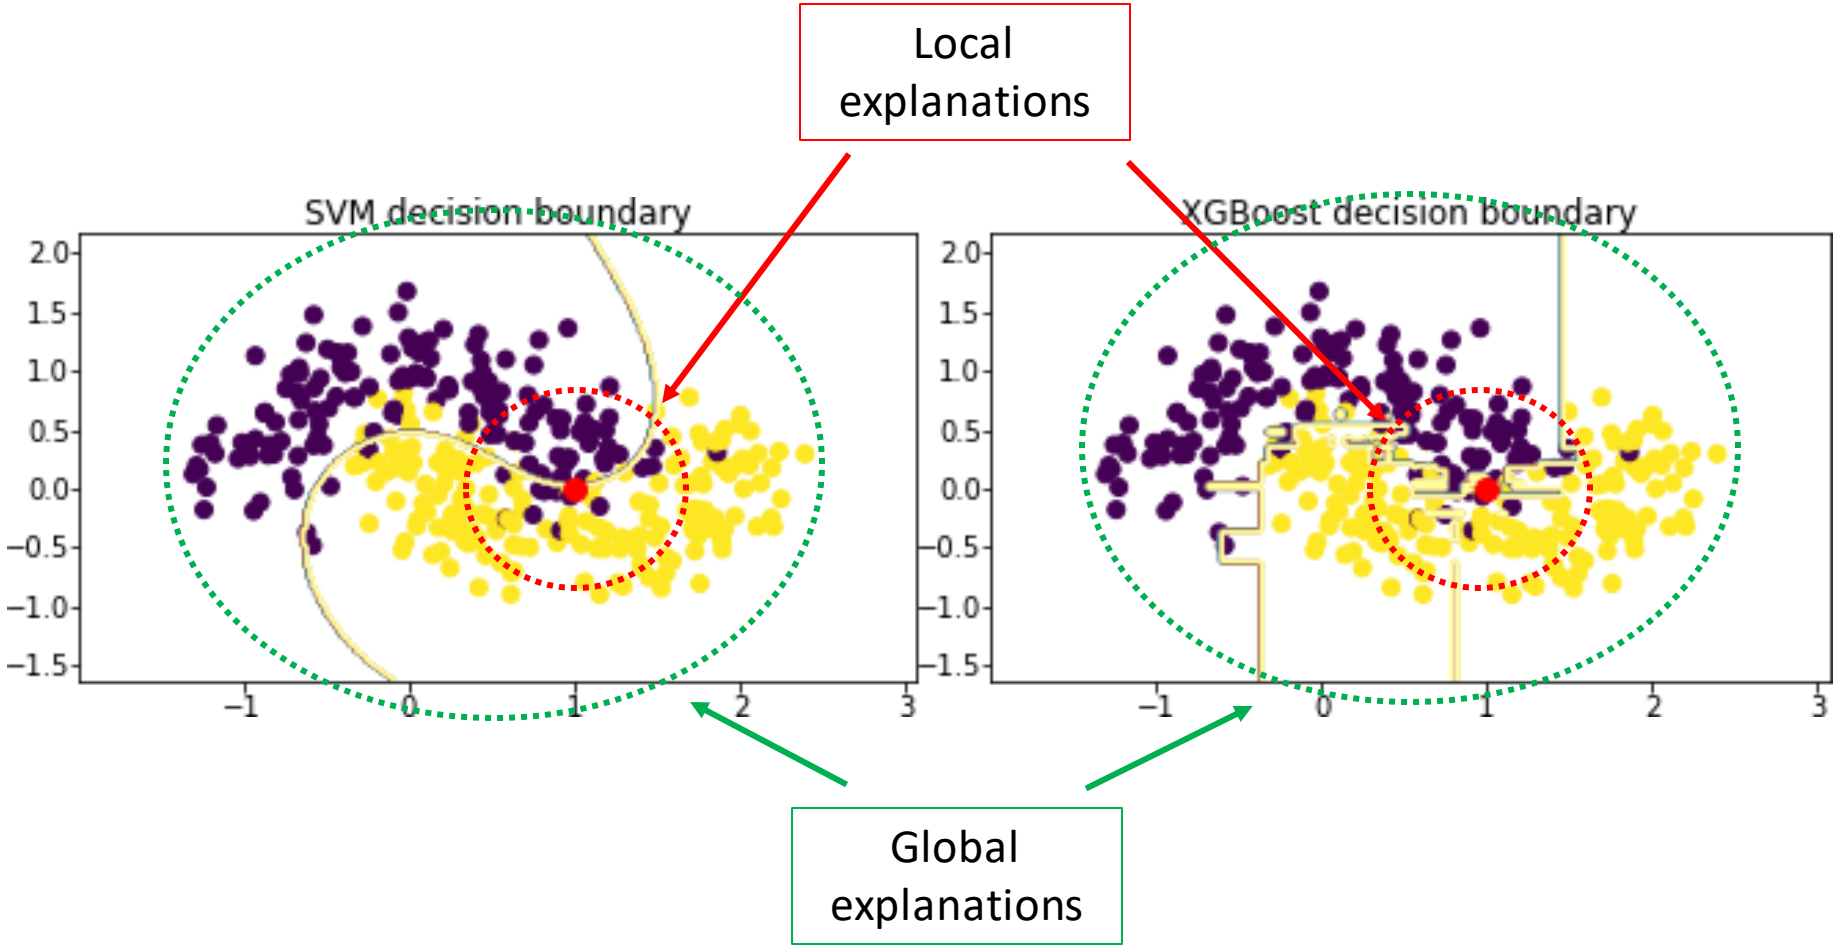

Model-agnostiv vs. Model specific

Model-agnostic

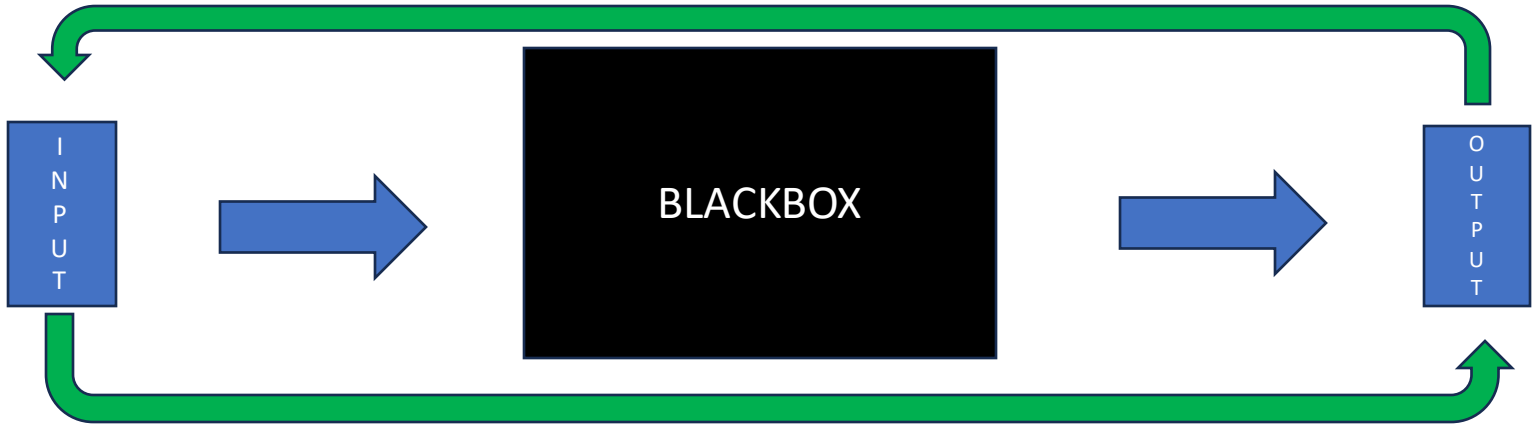

Model-specific

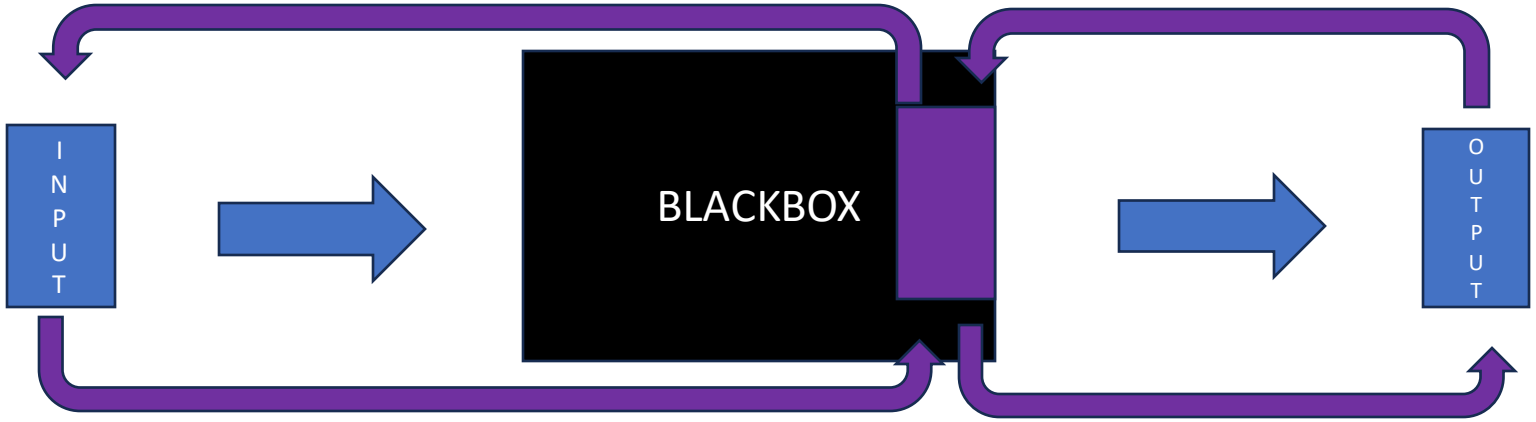

Classification of XAI methods

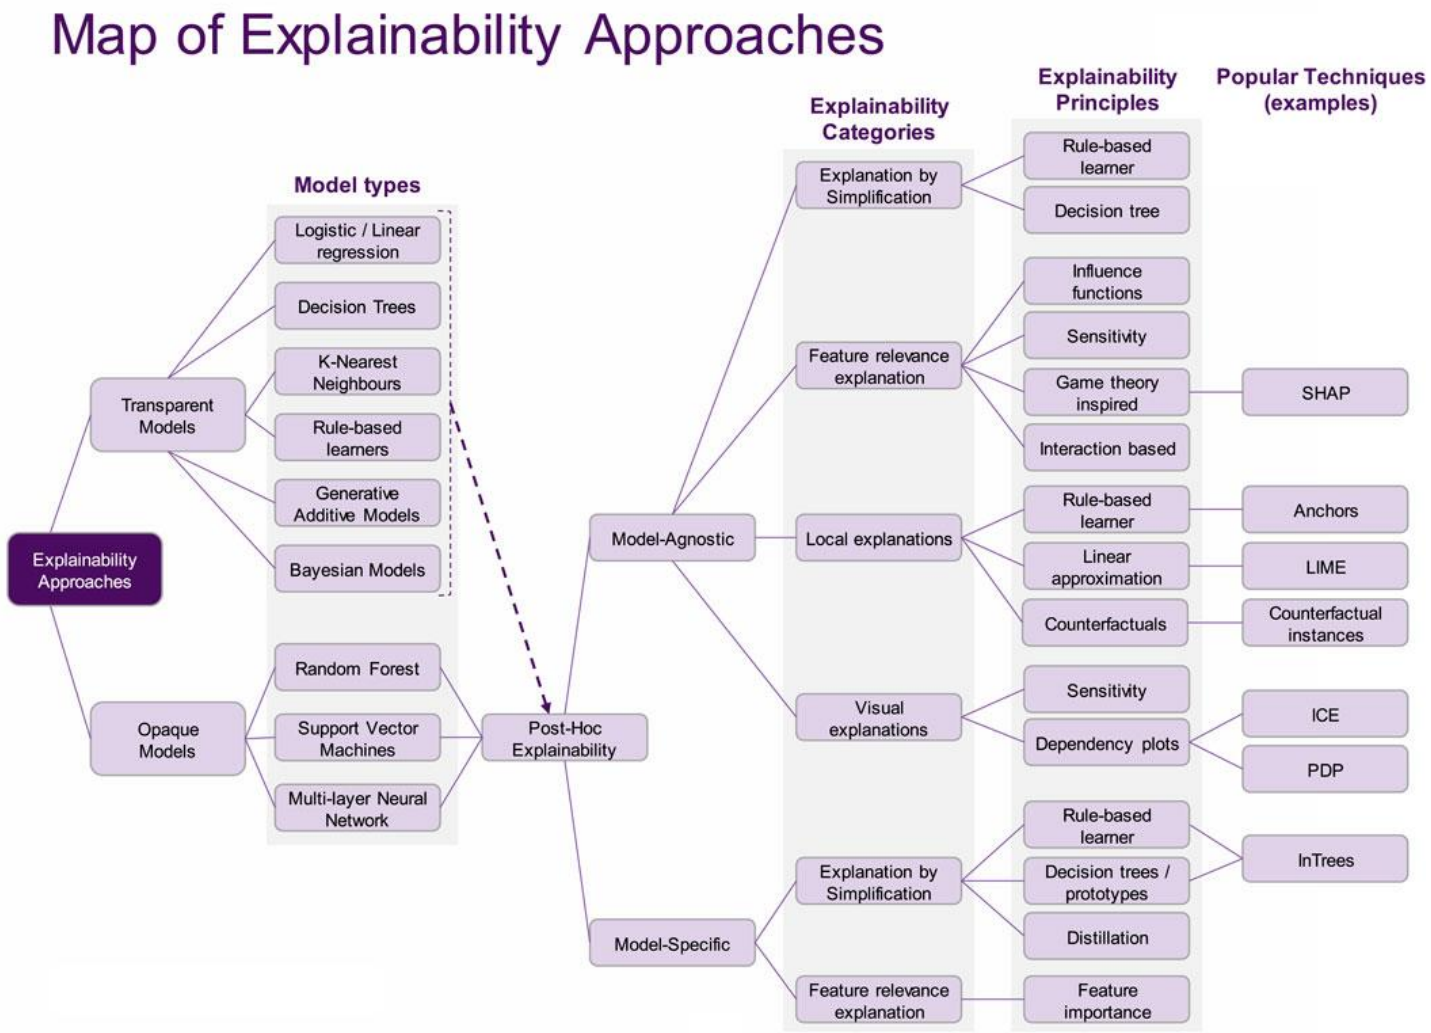

Evaluation of XAI methods

- Types of evaluation approaches
  - **Human-grounded**
  - **Application-grounded**
  - **Functional**
- Popular Quality measures
  - Fidelity (local and global)
  - Stability
  - Consistency
  - Coverage
  - Certainty
  - Representativeness
  - Simplicity/Comprehensability
- Ready to use frameworks
  - Quantus

https://github.com/understandable-machine-intelligence-lab/Quantus

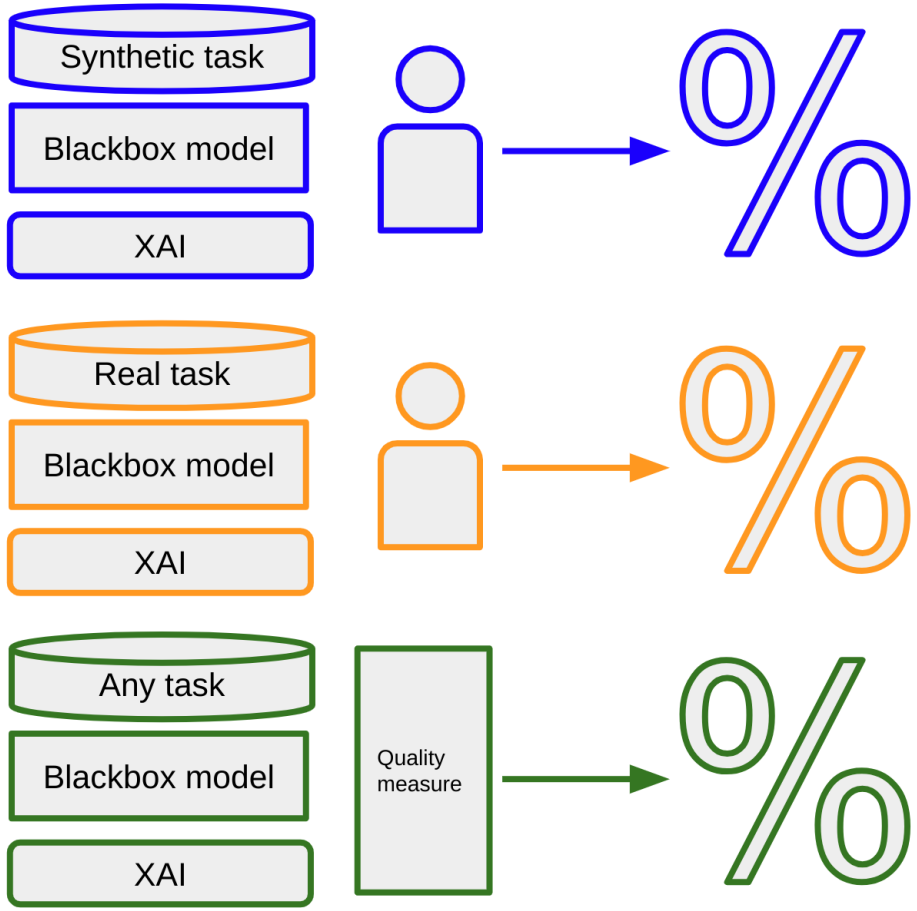

All in all, everything starts from the data

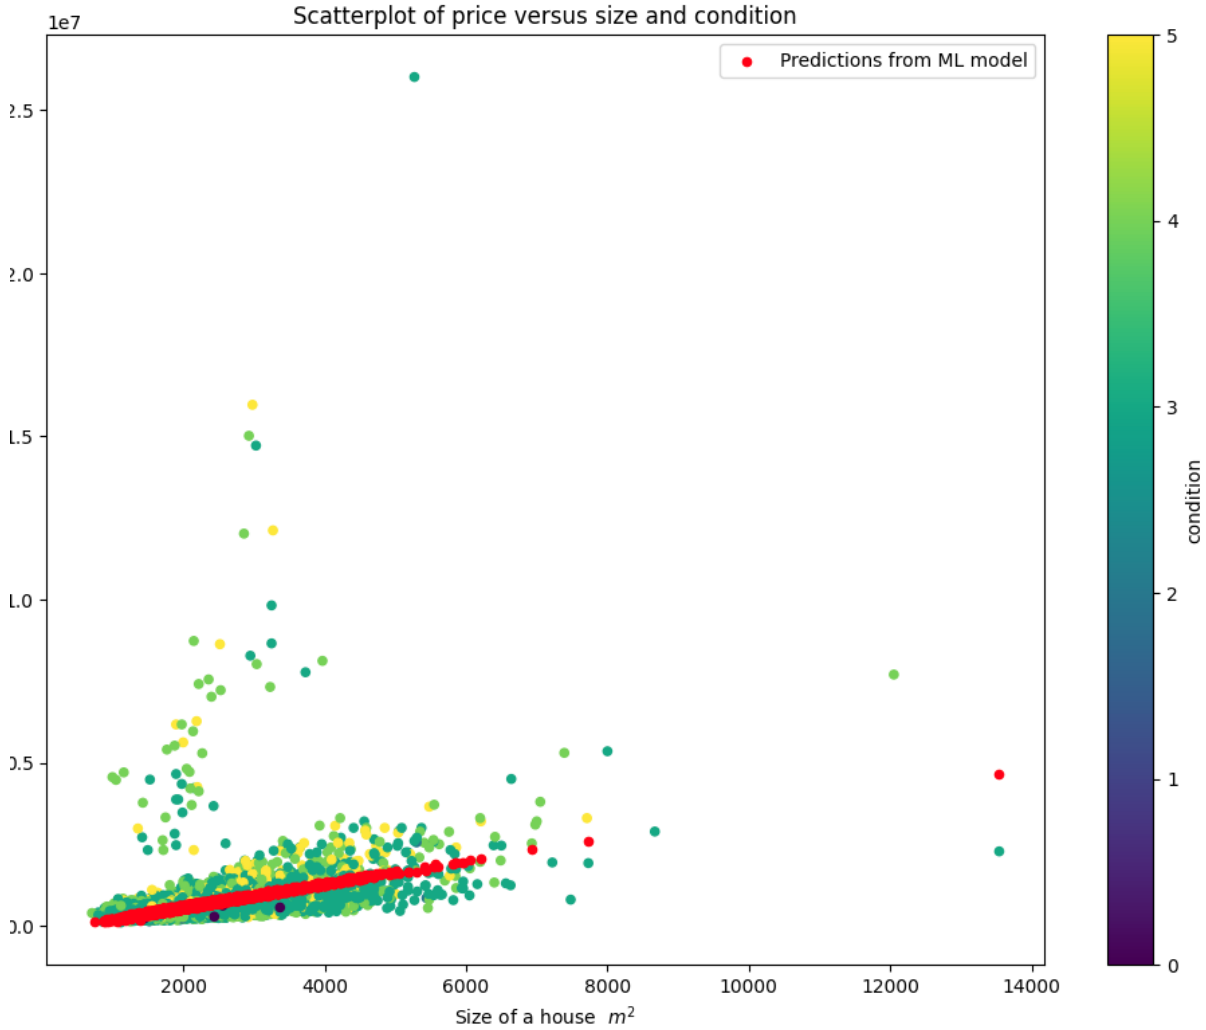

Thank you for your attention!

JAGIELLONIAN UNIVERSITY
IN KRAKÓW

https://geist.re