The Computer Will See You Now: Would You Trust an AI Doctor?

Siân Carey

Supervisors:

CiaránMcInerney, Alwyn Kotzé, Tom Lawton,

Owen Johnson, IbrahimHabli, Marc De Kamps

UK Research and Innovation The Leeds Teaching Hospitals NHS Trust University of York University of Sheffield Bradford Teaching Hospitals NHS Foundation Trust

UNIVERSITY OF LEEDS

UK Research and Innovation

UNIVERSITY OF LEEDS

Fairness in Medical AI

- What is AI?
- How can it help in medicine?
- Why are people worried?
- What are we doing to help?

UK Research and Innovation

UNIVERSITY OF LEEDS

What is AI?

UK Research and Innovation

UNIVERSITY OF LEEDS

What is AI?

AI, or Artificial Intelligence, is the field of computer science focused on creating machines and systems that can perform tasks that typically require human intelligence, such as learning from data, reasoning, problem-solving, and understanding natural language.

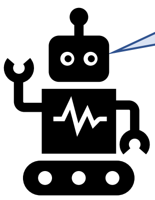

What is AI?

S What is AI?

AI, which stands for "Artificial Intelligence," is a branch of computer science and technology that focuses on creating systems and machines that can perform tasks that typically require human intelligence. These tasks include:

1. Learning: AI systems can acquire knowledge and improve their performance over time through data analysis and experience, often using techniques such as machine learning.
2. Reasoning: AI can make decisions, solve problems, and draw logical conclusions based on the information available to them.
3. Problem-Solving: AI can be used to tackle complex problems and find optimal solutions by searching through large solution spaces.
4. Natural Language Understanding: AI can comprehend and interpret human language, enabling communication between humans and machines.
5. Perception: AI can process and interpret sensory data, such as visual or auditory information, enabling tasks like image recognition and speech recognition.
6. Robotics: AI plays a significant role in creating robots and autonomous systems that can perform tasks in the physical world, including manufacturing, healthcare, and more.

Regenerate

S What is AI in one sentence

AI is the development of computer systems and algorithms that can perform tasks typically requiring human intelligence, such as learning, reasoning, problem-solving, and understanding natural language.

S What is AI? (In one sentence)

AI, or Artificial Intelligence, is the field of computer science focused on creating machines and systems that can perform tasks that typically require human intelligence, such as learning from data, reasoning, problem-solving, and understanding natural language.

UK Research and Innovation
UNIVERSITY OF LEEDS

Different types of AI

Logistic Regression	Decision Tree	NLP	Reinforcement Learning
Support Vector Machine	Linear Regression	Clustering	Neural Network

UK Research and Innovation

UNIVERSITY OF LEEDS

Different types of AI

Logistic Regression	Decision Tree	NLP	Reinforcement Learning
Support Vector Machine	Linear Regression	Clustering	Neural Network

UK Research and Innovation

UNIVERSITY OF LEEDS

Decision Tree

UK Research and Innovation

UNIVERSITY OF LEEDS

Different types of AI

Think of an animal.

UK Research and Innovation

UNIVERSITY OF LEEDS

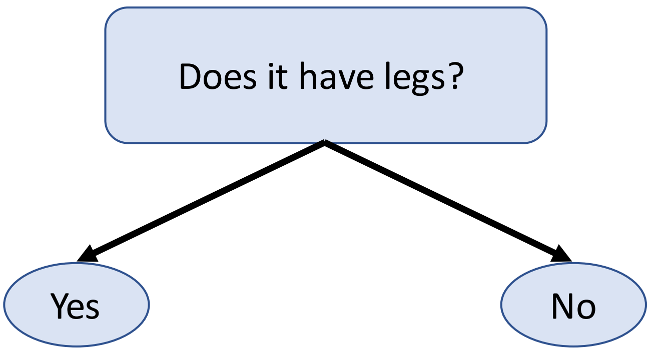

UK Research and Innovation

UNIVERSITY OF LEEDS

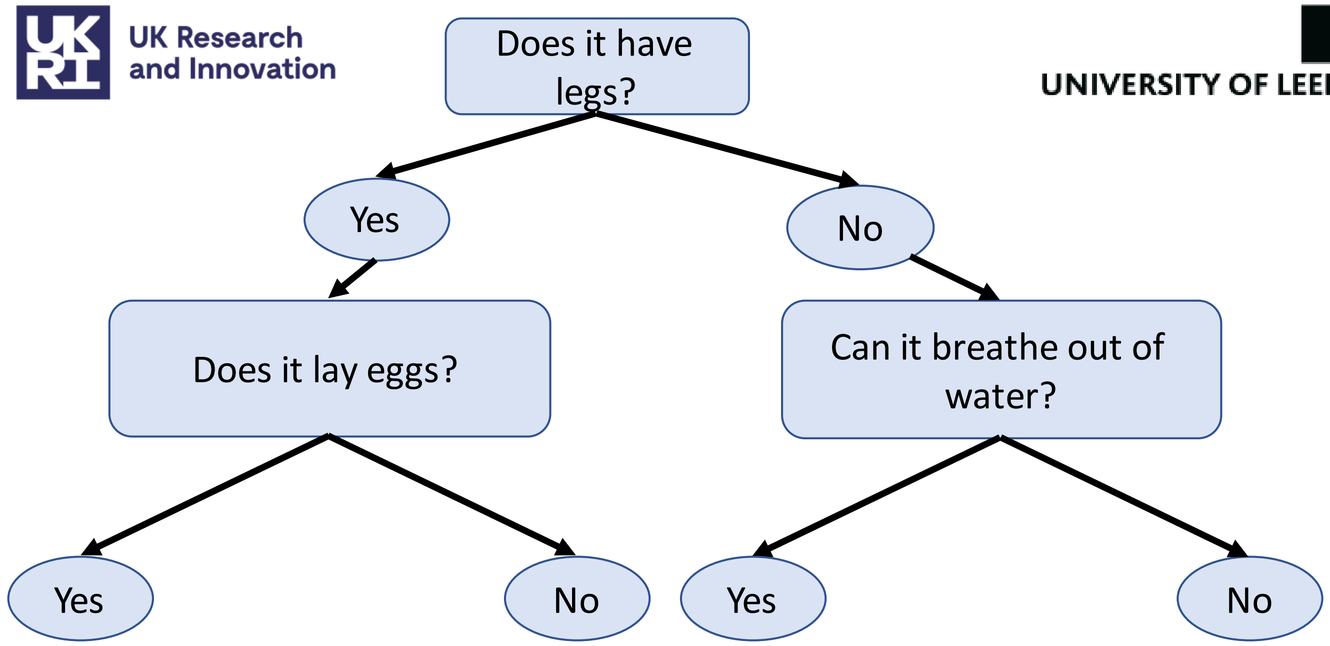

UK Research and Innovation

UNIVERSITY OF LEEDS

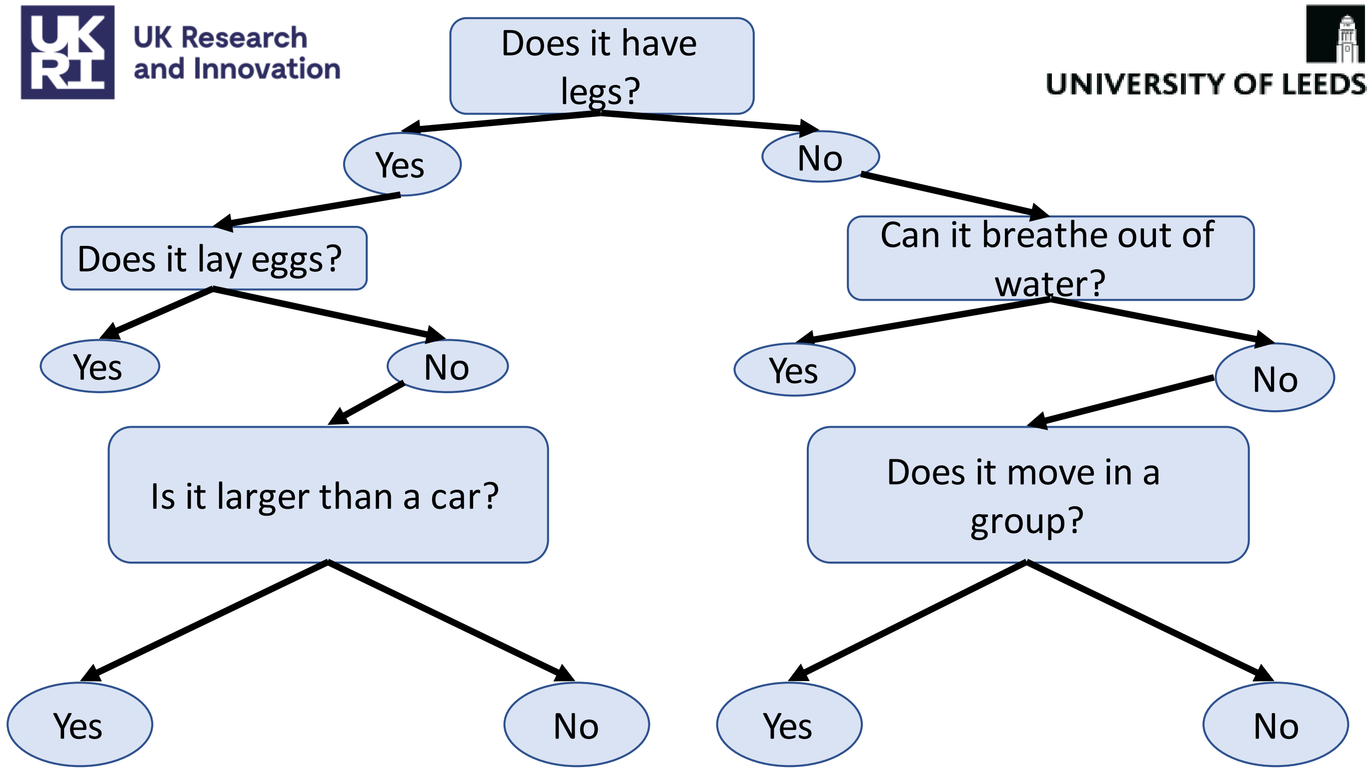

UK Research and Innovation

UNIVERSITY OF LEEDS

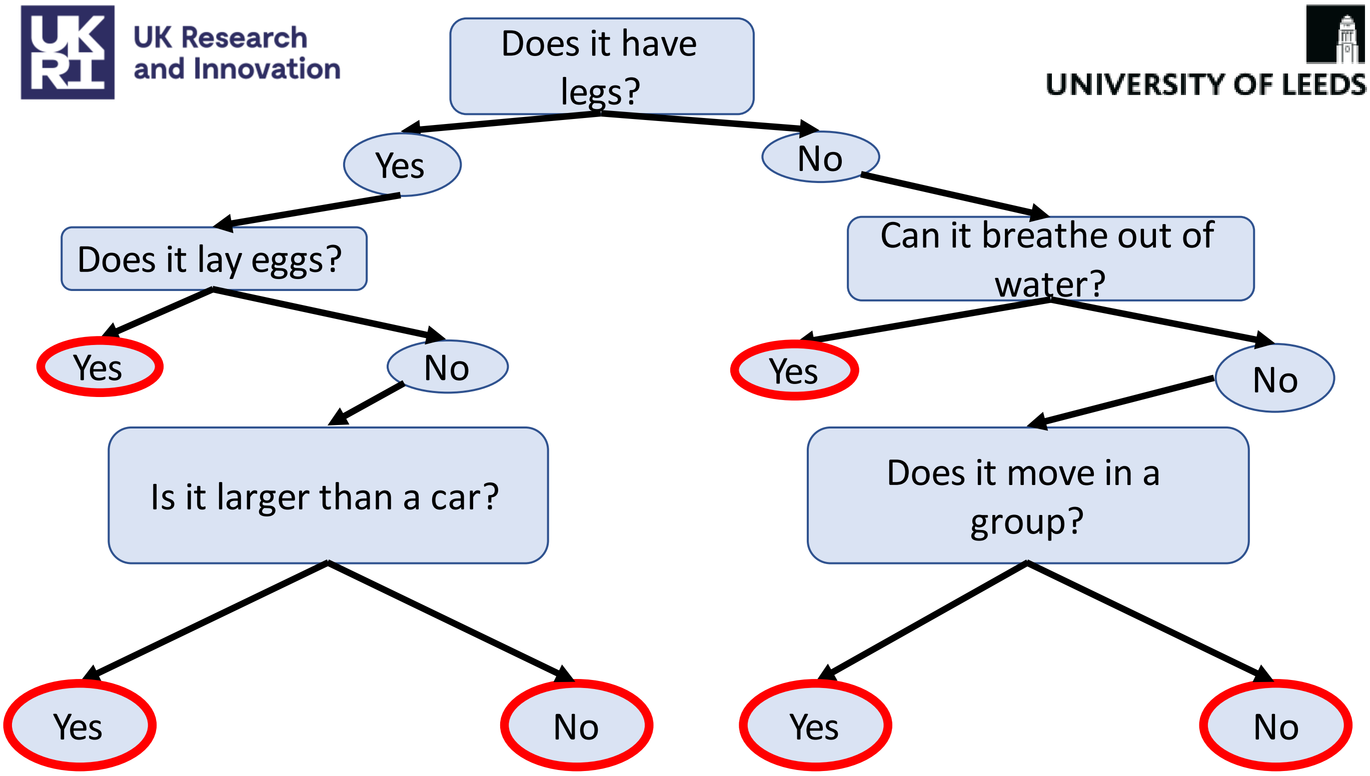

UK Research and Innovation

UNIVERSITY OF LEEDS

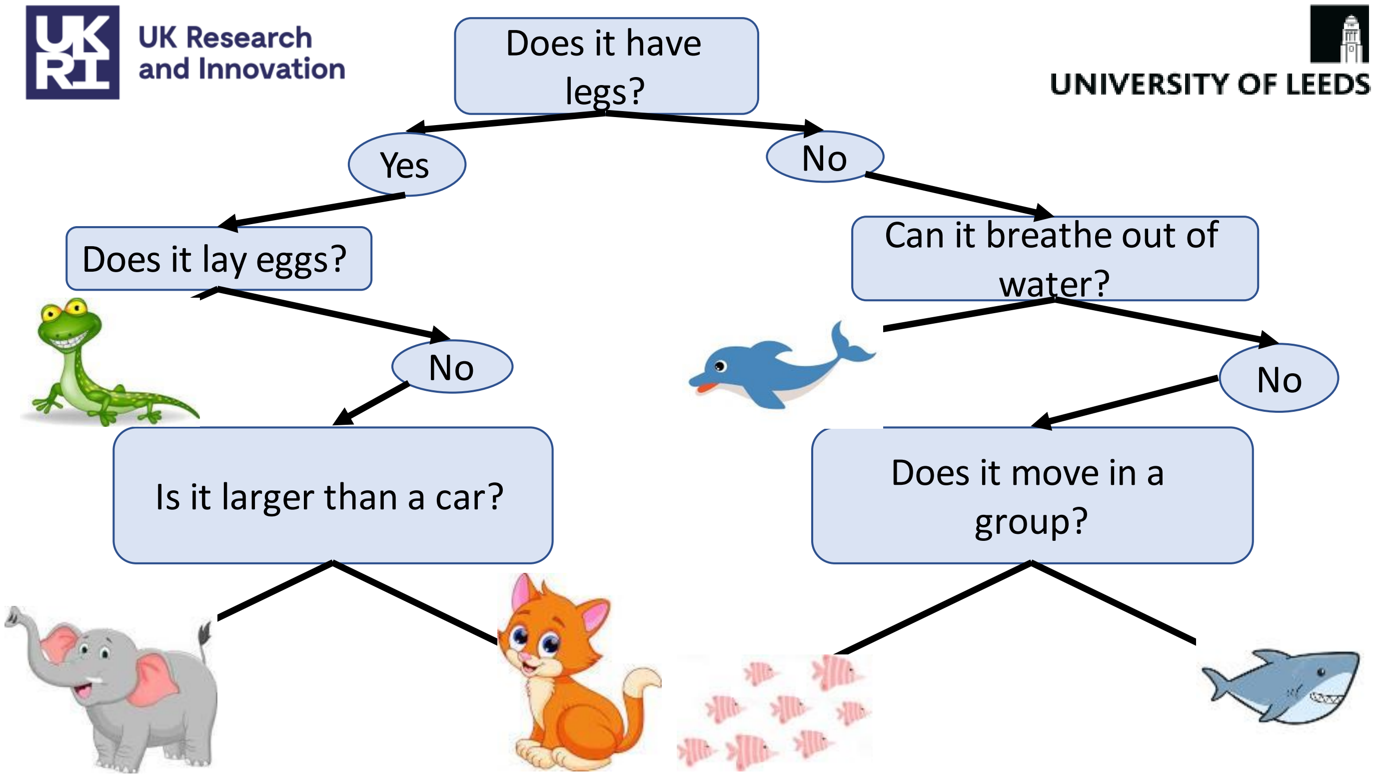

Neural Net

UK Research and Innovation

UNIVERSITY OF LEEDS

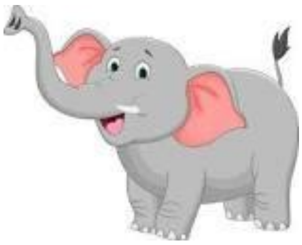

4 Legs | >2,700 kg | Large Ears | Births live young | Africa

UK Research and Innovation

UNIVERSITY OF LEEDS

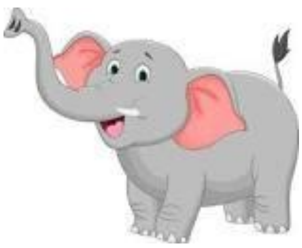

4 Legs | >2,700 kg | Large Ears | Births live young | Africa

Number of legs | Body size | Ear size | Birth type | Home continent

UK Research and Innovation

UNIVERSITY OF LEEDS

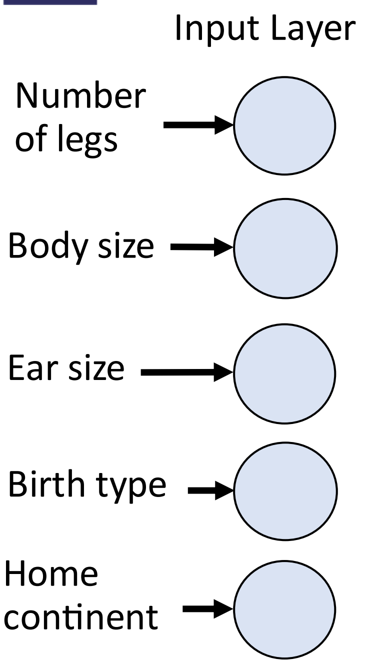

UK Research and Innovation

UNIVERSITY OF LEEDS

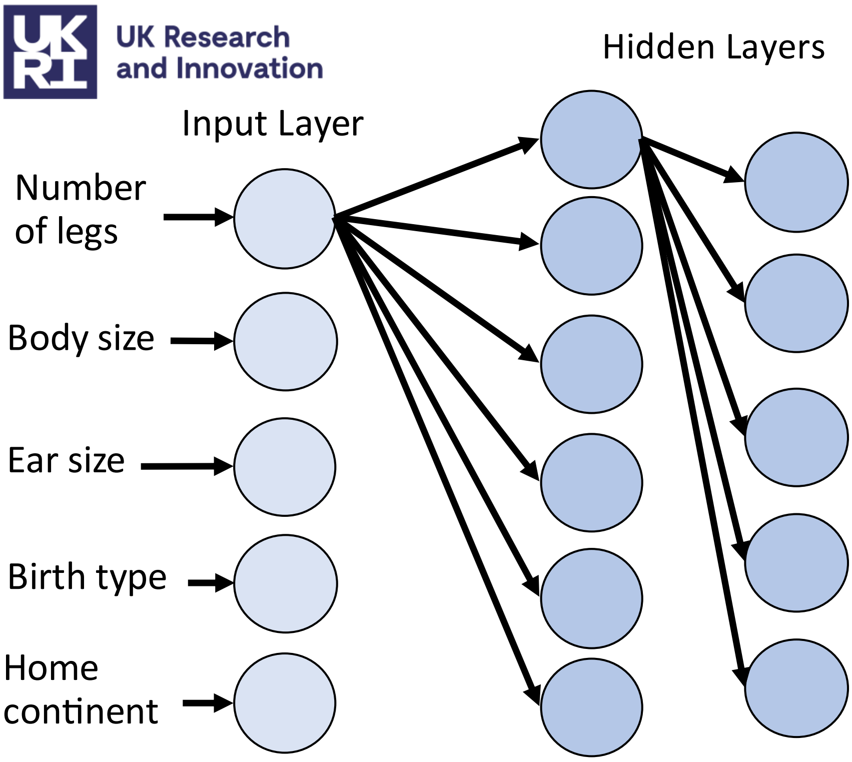

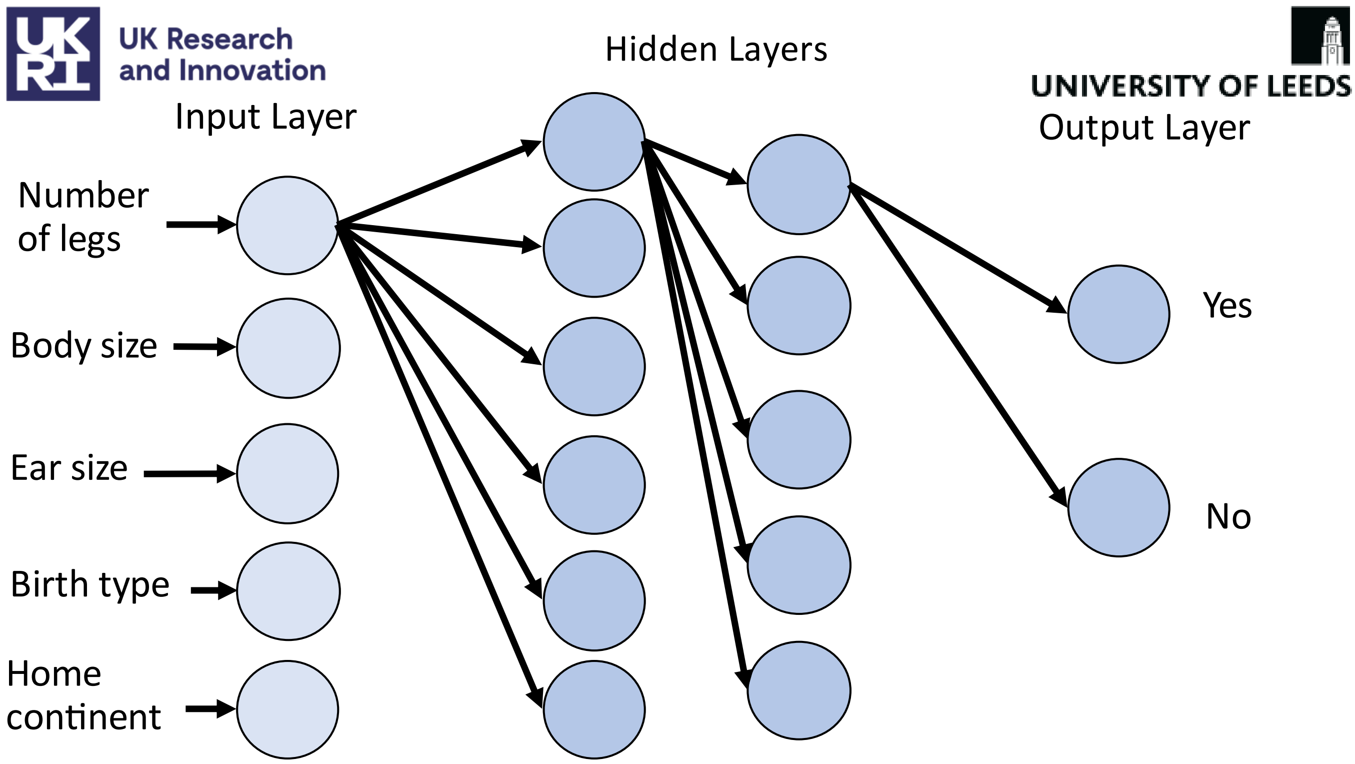
UK Research and Innovation
UNIVERSITY OF LEEDS
Hidden Layers
Input Layer
Output Layer
Number of legs
Body size
Ear size
Birth type
Home continent
Yes
No

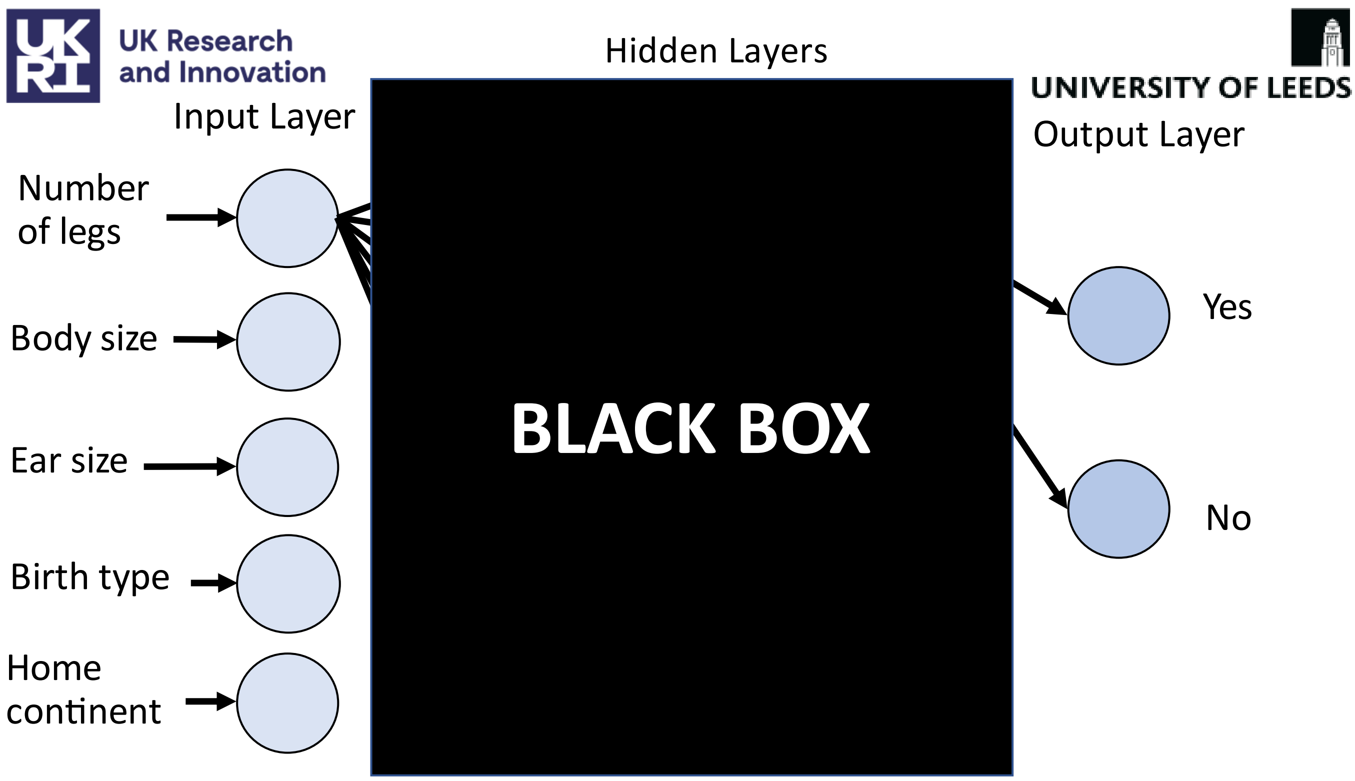
UK Research and Innovation
UNIVERSITY OF LEEDS
Hidden Layers
Input Layer
Output Layer
Number of legs
Body size
Ear size
Birth type
Home continent
BLACK BOX
Yes
No

UK Research and Innovation

UNIVERSITY OF LEEDS

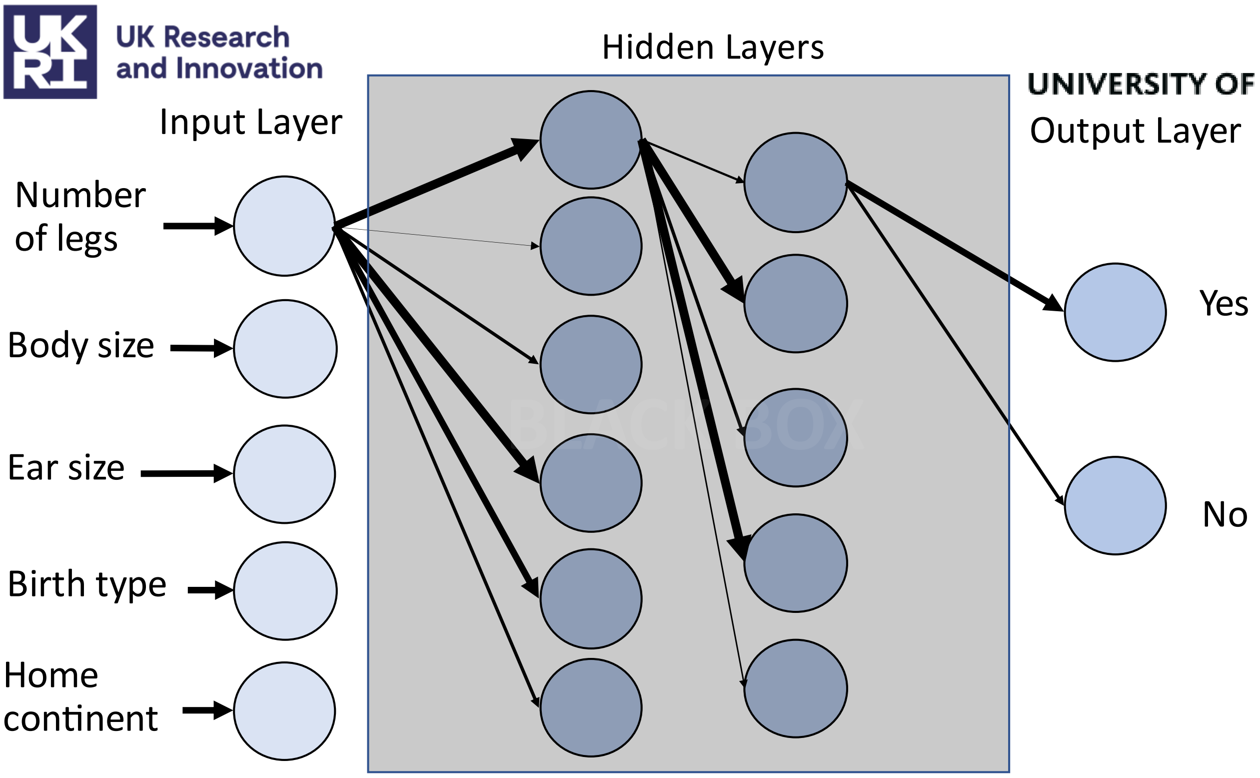

UK Research and Innovation

UNIVERSITY OF LEEDS

AI in Medicine

UK Research and Innovation

UNIVERSITY OF LEEDS

Why?

Shortage of workers

Amount of paperwork

Speed of diagnosis/treatment

UK Research and Innovation

UNIVERSITY OF LEEDS

Use Case 1 - Jack

- Images
- Different types of scanner

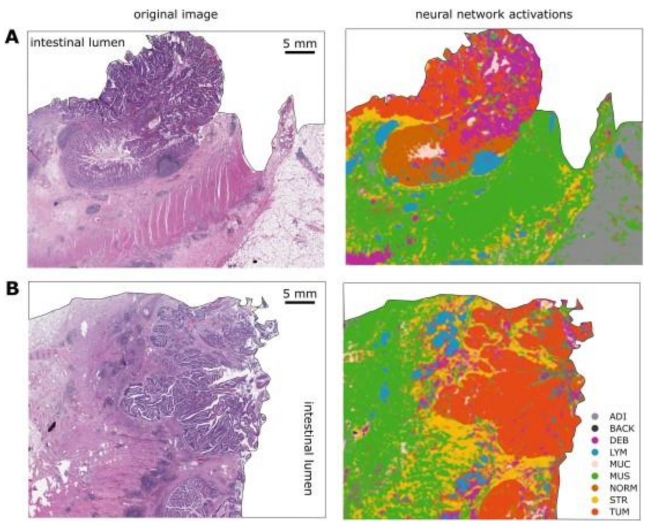

[1] Predicting survival from colorectal cancer histology slides using deep learning: A retrospective multicenter study, Kather et al., 2019

UK Research and Innovation

UNIVERSITY OF LEEDS

Use Case 2 - Emma

- Electronic Healthcare Records

Low haemoglobin	Raised platelets	Constipation	Chest pain	Abdominal pain	Nausea or vomiting	Dyspepsia	Epigastric pain	Reflux	Loss of weight	Dysphagia	
0.2 (0.2–109)	0.5 (0.4–0.5)	0.2 (0.2–0.2)	0.2 (0.2–0.2)	0.3 (0.2–0.3)	0.6 (0.5–0.7)	0.7 (0.6–0.7)	0.9 (0.8–1.0)	0.6 (0.6–0.7)	0.9 (0.7–1.0)	4.8 (4.3–5.9)	PPV as a single symptom
	0.6 (0.6–0.7)	0.4 (0.4–0.5)	0.3 (0.3–0.4)	0.5 (0.4–0.6)	0.9 (0.7–1.1)	1.0 (0.8–1.3)	1.6 (1.1–2.2)	0.9 (0.7–1.2)	1.0 (0.8–1.3)	4.6 (3.4–6.6)	Low haemoglobin
		0.9 (0.6–1.4)	0.8 (0.6–1.2)	0.8 (0.6–1.1)	1.4 (1.0–2.1)	1.4 (0.9–2.2)	1.9 (1.0–3.8)	1.6 (0.9–2.9)	1.8 (1.1–3.0)	6.1 (3.2–13.2)	Raised platelets
			0.4 (0.3–0.5)	0.4 (0.3–0.5)	0.6 (0.4–0.7)	0.8 (0.6–1.1)	1.4 (0.8–2.3)	0.7 (0.5–1.1)	1.1 (0.8–1.7)	4.2 (2.7–7.2)	Constipation
				0.3 (0.3–0.4)	0.6 (0.4–0.8)	0.7 (0.5–0.9)	0.9 (0.6–1.4)	0.6 (0.5–0.9)	1.1 (0.7–1.8)	5.8 (3.5–10.8)	Chest pain
					0.7 (0.5–0.9)	1.0 (0.7–1.3)	0.9 (0.7–1.2)	0.6 (0.5–0.9)	1.4 (0.9–2.2)	6.5 (3.5–13.5)	Abdominal pain
						1.0 (0.8–1.2)	1.3 (0.9–1.8)	1.3 (0.9–2.0)	2.3 (1.5–3.5)	2.8 (1.7–4.8)	7.3 (4.4–13.9)	Nausea or vomiting
							1.2 (1.0–1.5)	1.4 (1.0–2.0)	0.9 (0.7–1.2)	2.1 (1.3–3.5)	9.8 (5.7–20.2)	Dyspepsia
								1.5 (1.0–2.4)	4.2 (1.8–11.0)	9.3 –	Epigastric pain
									3.1 (1.5–6.7)	5.0 (3.3–8.4)	Reflux
										9.2 (4.4–22.7)	Loss of weight
										5.5 (4.2–7.8)	Dysphagia

[2] The risk of oesophago -gastric cancer in symptomatic patients in primary care: a large case –control study using electronic records, Stapley et al., 2013

UK Research and Innovation

UNIVERSITY OF LEEDS

Use Case 3 - Mary

- Sound

Café Scientifique Headingley

Monday 12th February 2024

Dr Alexa: detecting disease from your voice

By: Mary Paterson

UK Research and Innovation

UNIVERSITY OF LEEDS

AI in Medicine - Concerns

UK Research and Innovation

UNIVERSITY OF LEEDS

Why might people be worried? (med AI)

- Can I trust it?
- Will it treat me fairly?
- How does it work?
- Who owns my data?
- What if it gains sentience and kills us all?
- Who do I go to if something is wrong?
- Who is making money from my data?
- Will electrification of doctors kill the planet?
- Will my information be kept private?
- Will doctors lose their jobs?
- Will I be looked after by humans?
- How much will it cost me?

UK Research and Innovation

UNIVERSITY OF LEEDS

Why might people be worried? (med AI)

- Can I trust it?
- Will it treat me fairly?
- How does it work?
- Who owns my data?
- What if it gains sentience and kills us all?
- Who do I go to if something is wrong?
- Who is making money from my data?
- Will electrification of doctors kill the planet?
- Will my information be kept private?
- Will doctors lose their jobs?
- Will I be looked after by humans?
- How much will it cost me?

UK Research and Innovation

UNIVERSITY OF LEEDS

	Assistive AI algorithms		Autonomous AI algorithms		
	Level 1	Level 2	Level 3	Level 4	Level 5
	Data presentation	Clinical decision-support	Conditional automation	High automation	Full automation
Event monitoring	AI	AI	AI	AI	AI
Response execution	Clinician	Clinician and AI	AI	AI	AI
Fallback	Not applicable	Clinician	AI, with a backup clinician available at AI request	AI	AI
Domain, system, and population specificity	Low	Low	Low	Low	High
Liability	Clinician	Clinician	Case dependent	AI developer	AI developer
Example	AI analyses mammogram and highlights high-risk regions	AI analyses mammogram and provides risk score that is interpreted by clinician	AI analyses mammogram and makes recommendation for biopsy, with a clinician always available as backup	AI analyses mammogram and makes biopsy recommendation, without a clinician available as backup	Same as level 4, but intended for use in all populations and systems

[3] Bitterman , Danielle & Aerts, Hugo & Mak, Raymond. (2020). Approaching autonomy in medical artificial intelligence. The Lancet Digital Health.

UK Research and Innovation

UNIVERSITY OF LEEDS

Why might people be worried? (med AI)

- Can I trust it?
- Will it treat me fairly?
- How does it work?
- Who owns my data?
- What if it gains sentience and kills us all?
- Who do I go to if something is wrong?
- Who is making money from my data?
- Will electrification of doctors kill the planet?
- Will my information be kept private?
- Will doctors lose their jobs?
- Will I be looked after by humans?
- How much will it cost me?

UK is 'completely and institutionally ageist'

Care England chief says ageism is national scandal that should be challenged in courts

Amelia Hill

Wed 26 Dec 2018 13.18 GMT

LGBT+ people's health 'seen as less important to NHS', equalities committee warns

NEWS

Review reveals 'vast' ethnic inequalities in NHS services

Women's health

Finally the UK has noticed its rampant sexism in healthcare. What now?

Analysis: acknowledging the shocking female health gap is only a first step - ministers must put money into reversing it

Ministers pledge to reset the dial

Six areas of focus

Alexandra Topping

Thu 23 Dec 2021 00.01 GMT

ORIGINAL ARTICLES: ENDOMETRIOSIS

Attractiveness of women with rectovaginal endometriosis: a case-control study

Paolo Vercellini, M.D.,[a] Laura Buggio, M.D.,[a] Edgardo Somigliana, M.D.,[a] Giussy Barbara, M.D.,[a] Paola Viganò, Ph.D.,[b] and Luigi Fedele, M.D.[a]

Journal of Ovarian Research BMC

J Ovarian Res. 2019; 12: 126. PMCID: PMC6937688
Published online 2019 Dec 30. doi: 10.1186/s13048-019-0600-7 PMID: 31888704

Influence of marital status on overall survival in patients with ovarian serous carcinoma: finding from the surveillance epidemiology and end results (SEER) database

Pei Luo,[#1] Jian-Guo Zhou,[#1,2] Su-Han Jin,[3] Ming-Song Qing,[4] and Hu Ma[1]

Analysis of the Visual Perception of Female Breast Aesthetics and Symmetry: An Eye-Tracking Study

Pietruski, Piotr M.D., Ph.D.; Paskal, Wiktor M.D.; Paskal, Adriana M. M.D.; Jaworowski, Janusz M.D.; Paluch, Łukasz M.D., Ph.D.; Noszczyk, Bartłomiej M.D., Ph.D.

Author Information

r 2019 - Volu

ELSEVIER

Body Image
Volume 8, Issue 2, March 2011, Pages 190-193

Brief research report

Judging the health and attractiveness of female faces: Is the most attractive level of facial adiposity also considered the healthiest?

AMERICAN JOURNAL OF BIOLOGICAL ANTHROPOLOGY

BRIEF COMMUNICATION

Costs of reproduction are reflected in women's faces: Post-menopausal women with fewer children are perceived as more attractive, healthier and younger than women with more children

Urszula M. Marcinkowska, Anthony C. Little, Andrzej Galbarczyk, Ilona Nenko, Magdalena Klimek, Grazyna Jasienska

First published: 13 November 2017 | https://doi.org/10.1002/ajpa.23362 | Citations: 7

UK Research and Innovation

UNIVERSITY OF LEEDS

Fairness in Medical AI

Definitions

Sensitive/protected characteristics

- Age
- Gender reassignment
- Being married or in a civil partnership
- Being pregnant or on maternity leave
- Disability
- Race including colour, nationality, ethnic or national origin
- Religion or belief
- Sex
- Sexual orientation

Proxy Variables

- Variables that are not a sensitive characteristic but are highly correlated to a sensitive characteristic.
- Examples (UK):

  Postcode ~ Ethnicity

  Gaps in career ~ Gender
- Identified by:

a) Testing correlations between variables (if sensitive characteristic available)

b) Field experts thinking/discussing (if sensitive characteristics not available)

UK Research and Innovation

UNIVERSITY OF LEEDS

What is fair?

- Group v Individual
- Difference v Discrimination
- Allocation harm
- Quality-of-service harm

UK Research and Innovation

UNIVERSITY OF LEEDS

Fairness in Medical AI

Identifying Bias

Mitigating Bias

UK Research and Innovation

UNIVERSITY OF LEEDS

Fairness in Medical AI

Identifying Bias

Mitigating Bias

UK Research and Innovation

UNIVERSITY OF LEEDS

Fairness Notions

Selection Criteria

Ground Truth Availability	Equal Base Rate	Reliable Outcome	Presence of Explanatory Variables
Emphasis on Precision v Recall	Emphasis on FP v FN	Cost of Misclassification	Fixed or Floating
Intersectionality	Masking	Source of Bias	Legal Framework

Ground Truth Availability

A ground truth value is the true and correct observed outcome corresponding to given sample in the data.

Predicting laryngeal cancer: The ground truth is the images from the patient's airway.

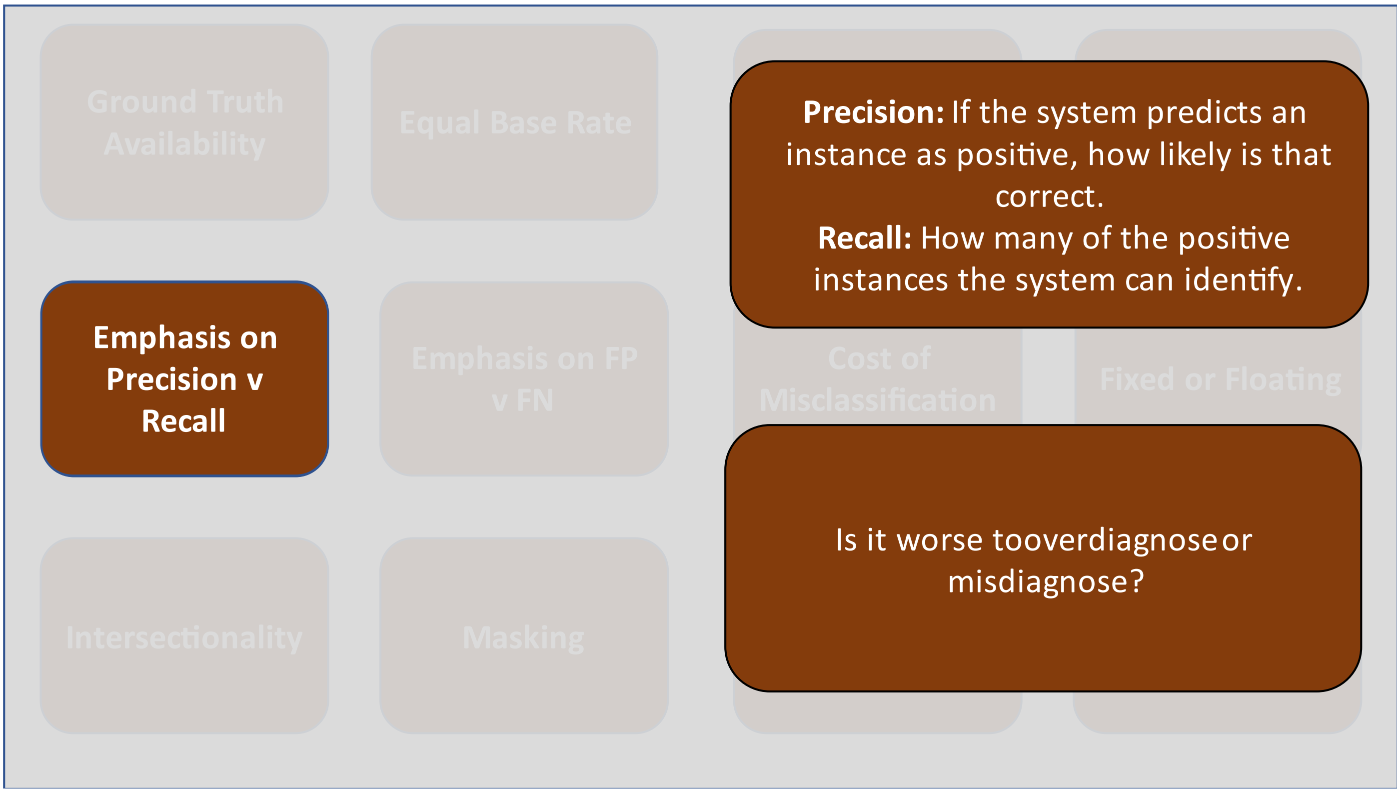
Ground Truth Availability
Equal Base Rate
Precision: If the system predicts an instance as positive, how likely is that correct.
Recall: How many of the positive instances the system can identify.
Emphasis on Precision v Recall
Emphasis on FP v FN
Cost of Misclassification
Fixed or Floating
Is it worse to overdiagnose or misdiagnose?
Intersectionality
Masking

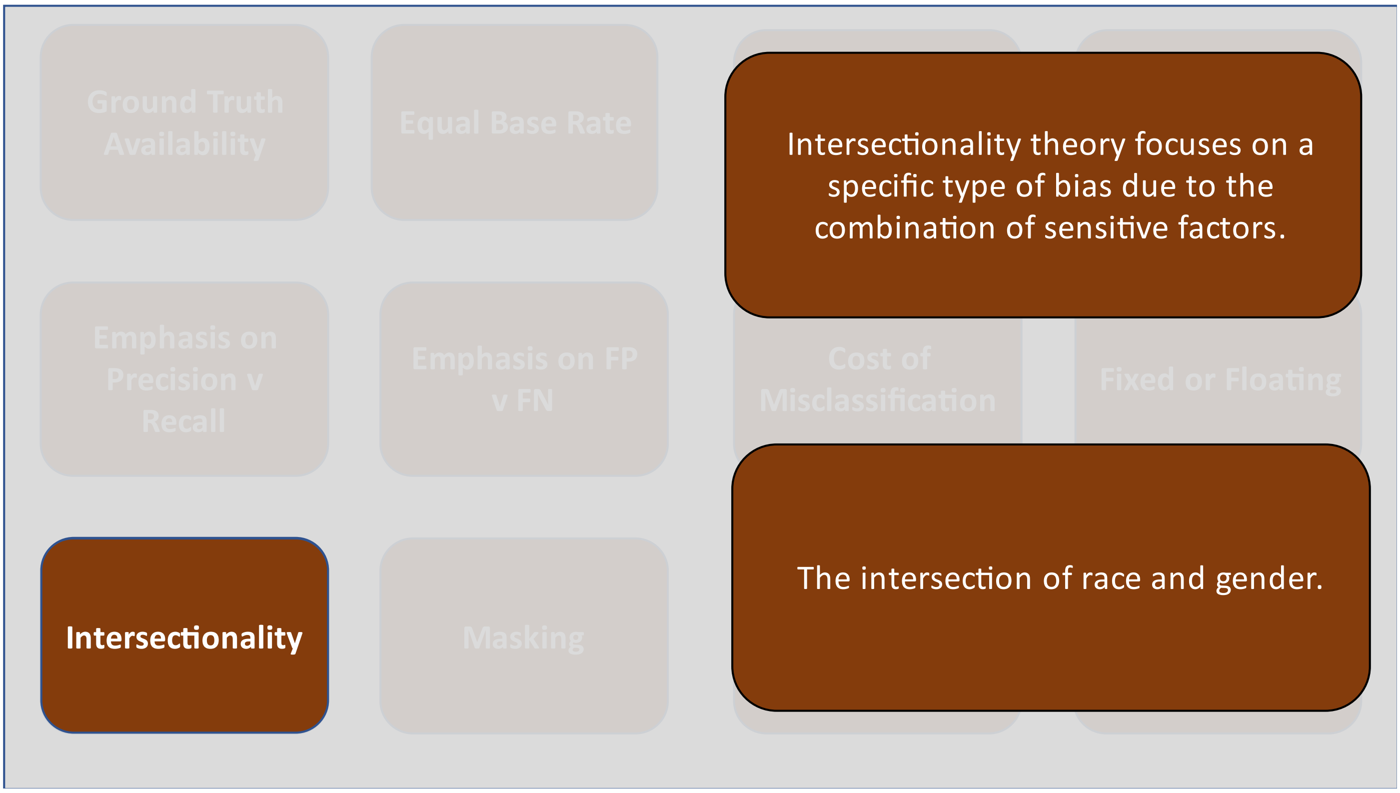
Ground Truth Availability
Equal Base Rate
Intersectionality theory focuses on a specific type of bias due to the combination of sensitive factors.
Emphasis on Precision v Recall
Emphasis on FP v FN
Cost of Misclassification
Fixed or Floating
The intersection of race and gender.
Intersectionality
Masking

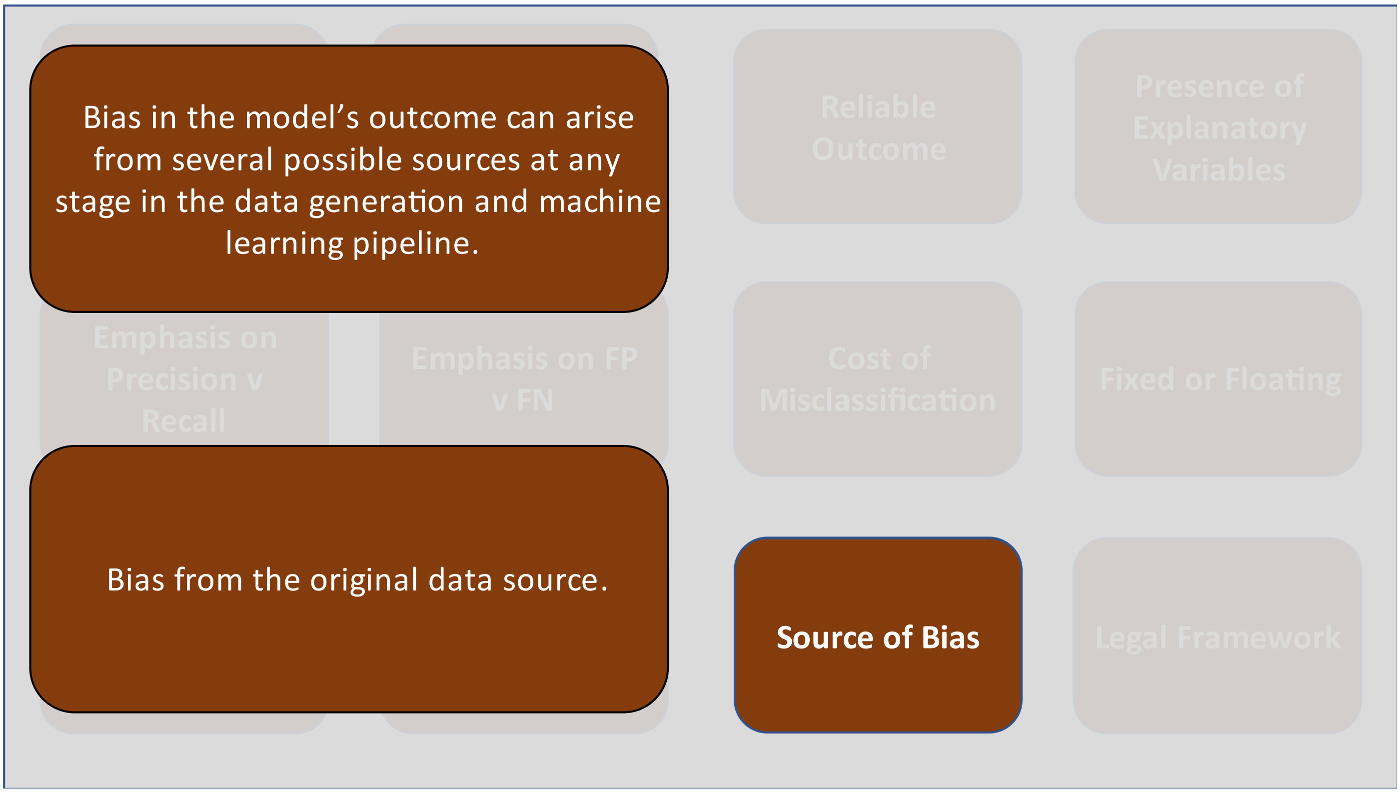

Fairness Notions

Group Notions

Conditional Statistical parity

Statistical Parity

Equalised Odds

Treatment Equality

Calibration

Requires the prediction to be statistically independent of the sensitive attribute

$$P(\hat{Y} \mid A = 0) = P(\hat{Y} \mid A = 1)$$

Statistical Parity

Acceptance rate of hiring male and female applicants is 0.57 (4 out of 7) and 0.4 (2 out of 5), respectively. Thus, this does not satisfy statistical parity.

(a) Dataset						(b) Prediction	
Gender	Education level	Job experience	Age	Marital status	Y	Ŷ	S
Female 1	8	2	39	Single	0	1	0.5
Female 2	8	2	26	Married	1	0	0.1
Female 3	12	8	32	Married	1	1	0.5
Female 4	11	3	35	Single	0	0	0.2
Female 5	9	5	29	Married	1	0	0.3
Male 1	11	3	34	Single	1	1	0.8
Male 2	8	0	48	Married	0	0	0.1
Male 3	7	3	43	Single	1	0	0.1
Male 4	8	2	26	Married	1	1	0.5
Male 5	8	2	41	Single	0	1	0.5
Male 6	12	8	30	Single	1	1	0.8
Male 7	10	2	28	Married	1	0	0.3

(a) Dataset						(b) Prediction	
Gender	Education level	Job experience	Age	Marital status	Y	Ŷ	S
Female 1	8	2	39	Single	0	1	0.5
Female 2	8	2	26	Married	1	0	0.1
Female 3	12	8	32	Married	1	1	0.5
Female 4	11	3	35	Single	0	0	0.2
Female 5	9	5	29	Married	1	0	0.3
Male 1	11	3	34	Single	1	1	0.8
Male 2	8	0	48	Married	0	0	0.1
Male 3	7	3	43	Single	1	0	0.1
Male 4	8	2	26	Married	1	1	0.5
Male 5	8	2	41	Single	0	1	0.5
Male 6	12	8	30	Single	1	1	0.8
Male 7	10	2	28	Married	1	0	0.3

Equalised Odds

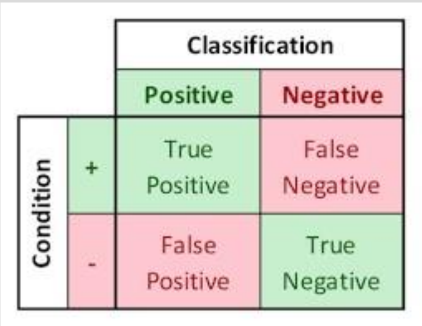

Equalised odds requires both subpopulations to have the same true positive rate (recall) and false positive rate (model error)

$P(\hat{Y} = 1 \mid Y = y, A = 0) = P(\hat{Y} = 1 \mid Y = y, A = 1)\ \forall y \in \{0, 1\}$

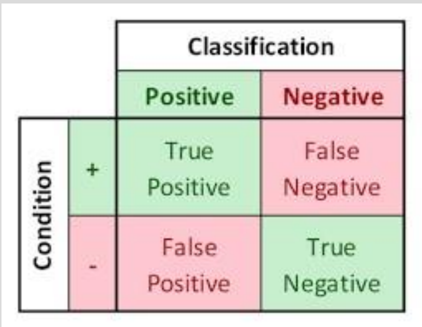

Statistical Parity

Treatment equality is achieved when the ratio of false positives and false negatives is the same for both protected and unprotected groups

Treatment Equality

Calibration

Fairness Notions

Individual Notions

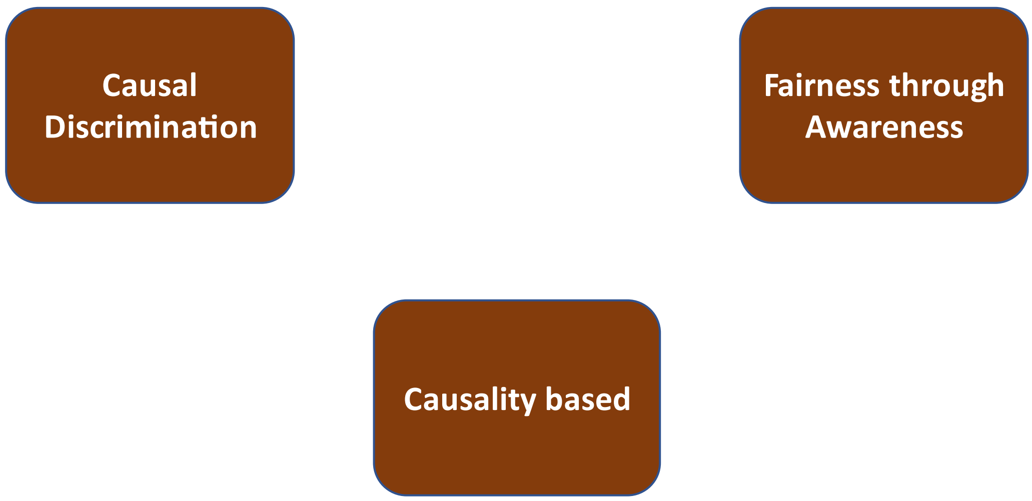

Causal Discrimination

Causal Discrimination implies that a classifier should produce the same prediction for individuals who differ only from gender while possessing identical attributes X.

Causality based

Causality-based fairness notions differ from all aforementioned statistical fairness approaches in that they are not totally based on data but consider additional knowledge about the structure of the world, in the form of a causal model.

Myth Busting

Fairness through Unawareness

Problems with Fairness Notions

They’re very strict. Often will require relaxation.

It is impossible to satisfy all fairness notions simultaneously except in extreme, degenerate, and dump scenarios.

Example

Surgical Outcome Risk Tool

- Nationally mandated
- Used at individual and group level
- Is fair, with regards to sex!

UK Research and Innovation

UNIVERSITY OF LEEDS

Other ways

UK Research and Innovation

UNIVERSITY OF LEEDS

Not Fairness Notions

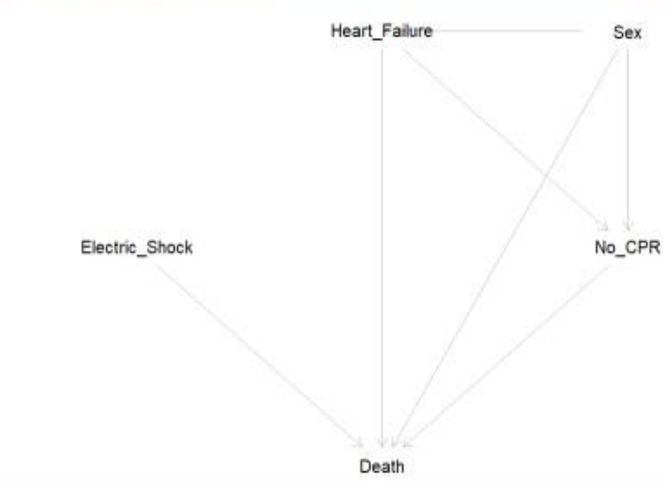

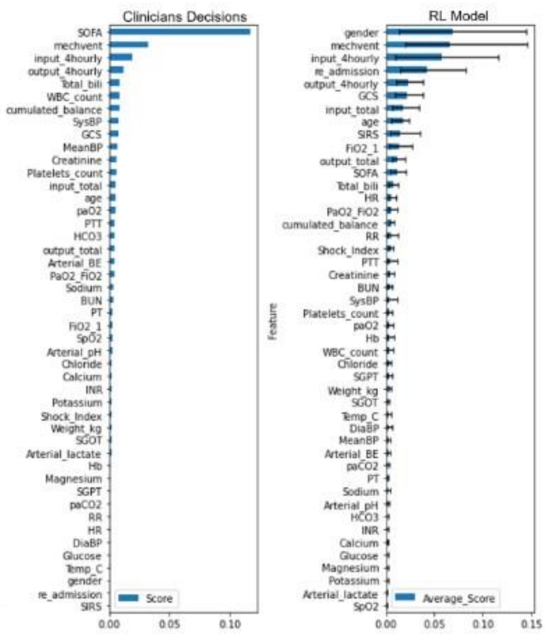

Mitigating Bias

UK Research and Innovation

UNIVERSITY OF LEEDS

Mitigating bias

PRE-PROCESSING

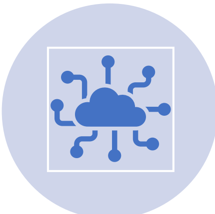

IN-MODEL

POST-PROCESSING

Mitigating bias

PRE-PROCESSING

Mitigating bias

IN-MODAL

Mitigating bias

POST-PROCESSING

PUBLIC