



Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

Explainable Machine Learning for Cervical Cancer Risk Factors Assessment

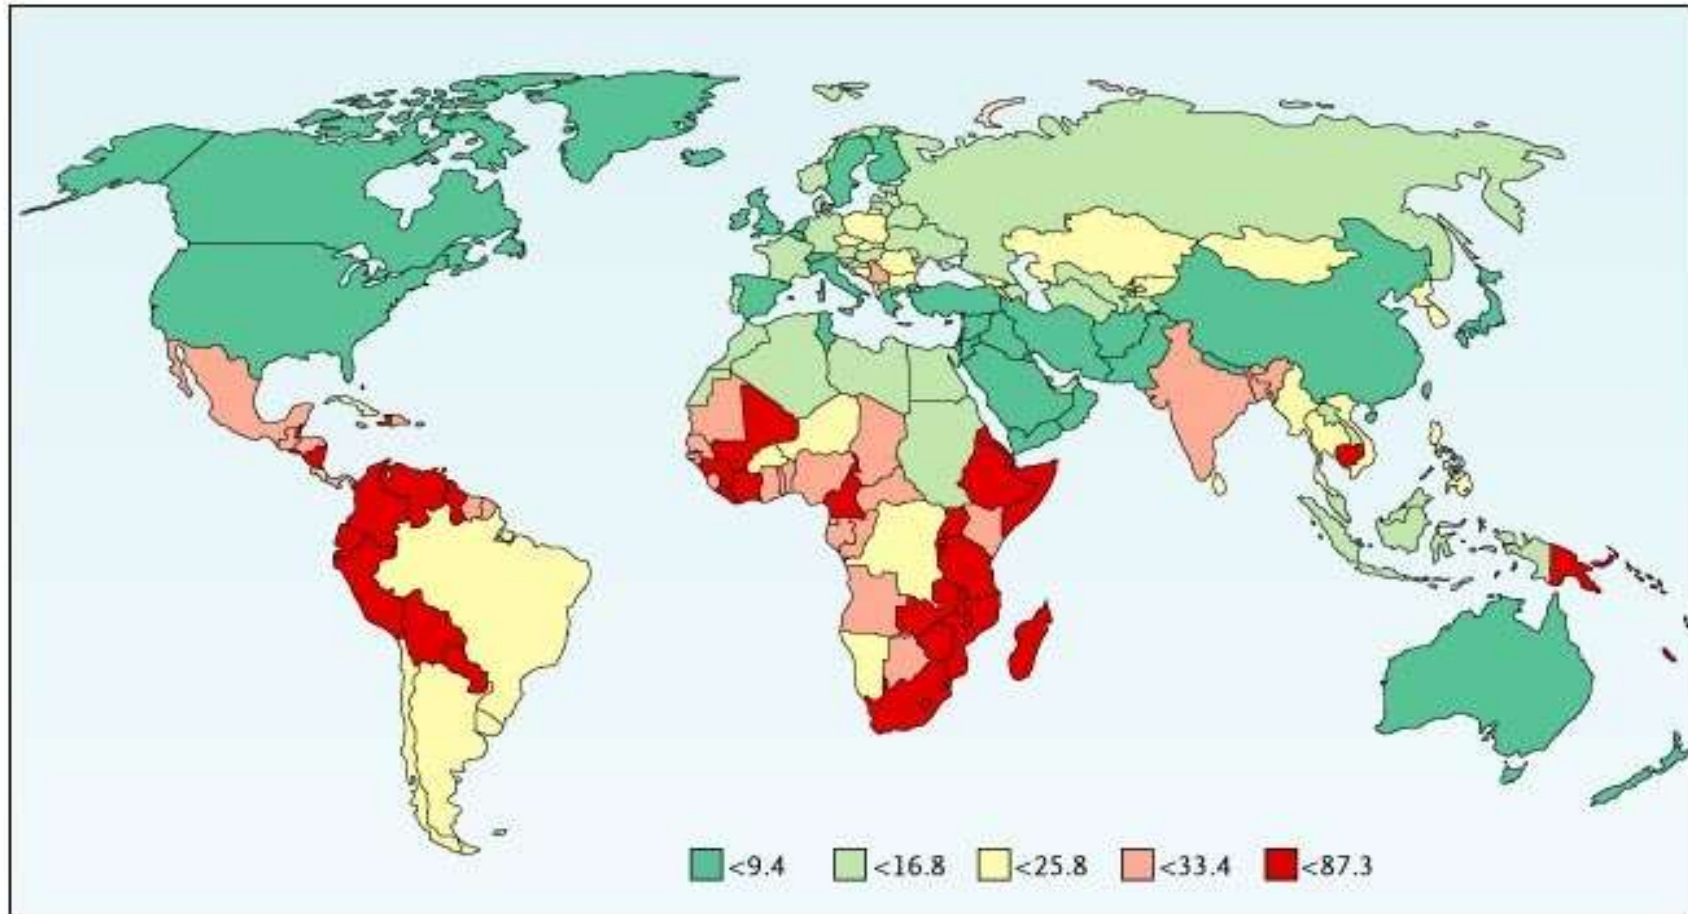
dr. ir. Sultan K. Imangaliyev

Dublin, 25/11/2020

Research question

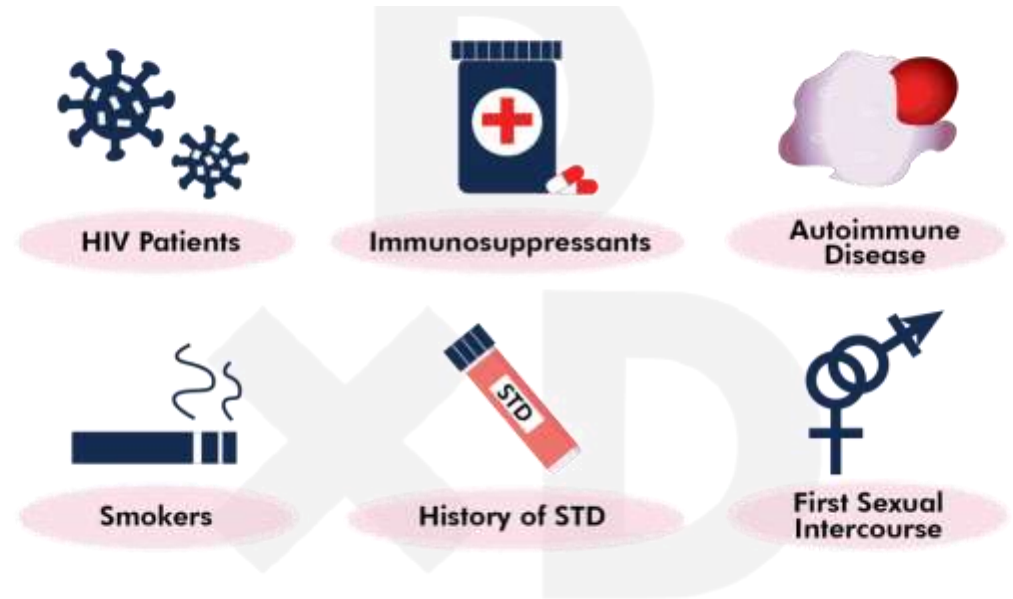

Which epidemiological and/or demographical variables are predictive for cervical cancer screening?

Cervical cancer worldwide



Incidence of Cervical Cancer, cases per 100K population

Cervical cancer risk factors




Case study


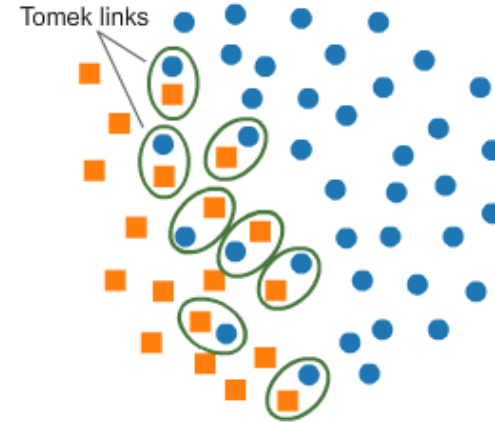

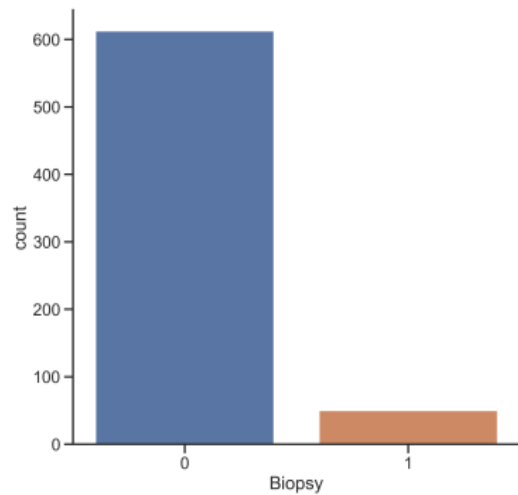
- Publicly available dataset*
 - 858 patients of *Hospital Universitario de Caracas, Venezuela*
- 36 features
 - demographics, lifestyle, medical history *etc.*
- Biopsy as a binary classification label




Missing values map




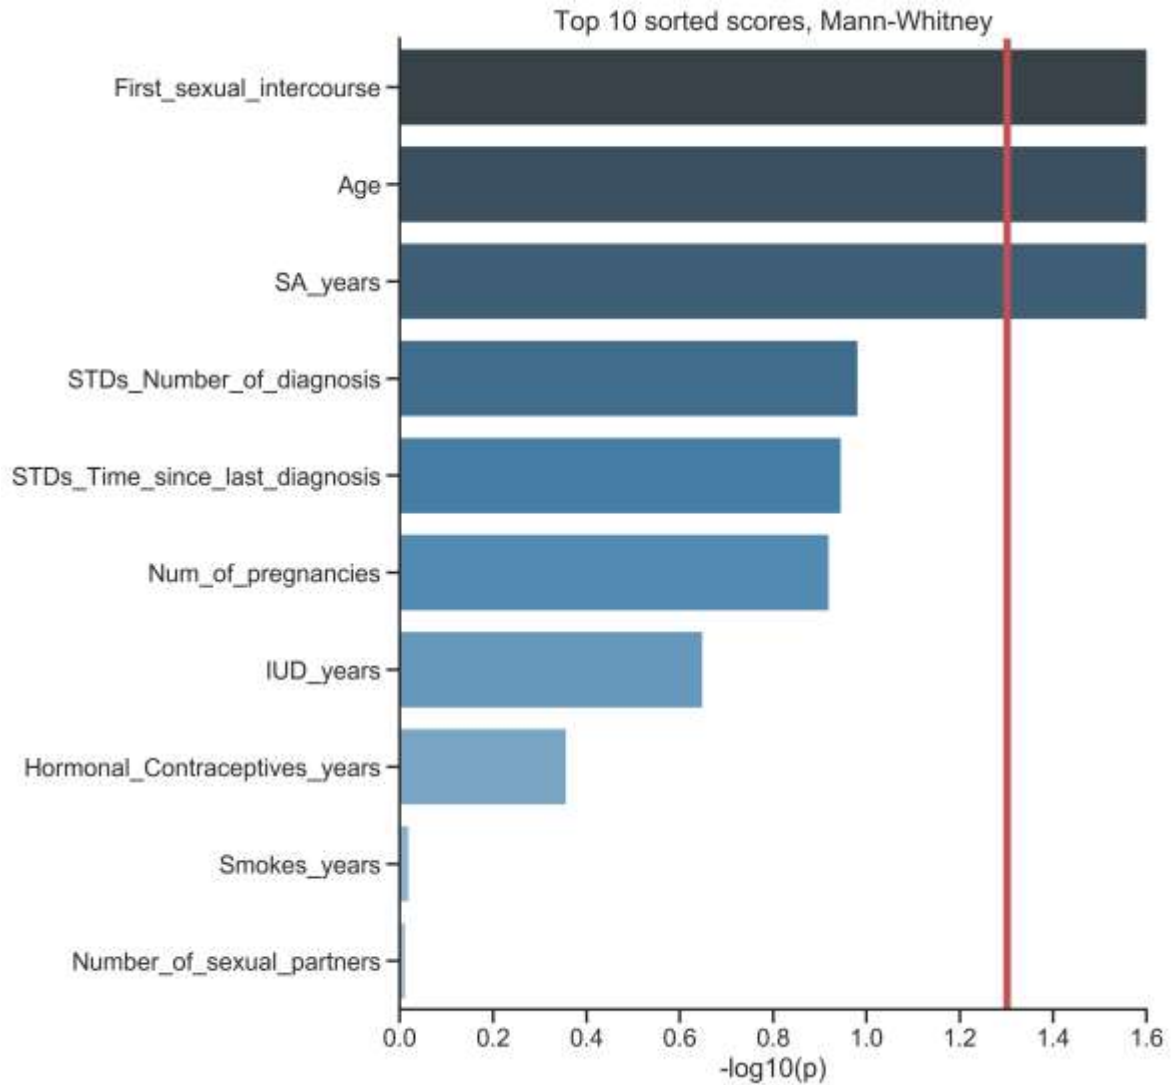
Class imbalance and data subsampling



Majority class subsampling

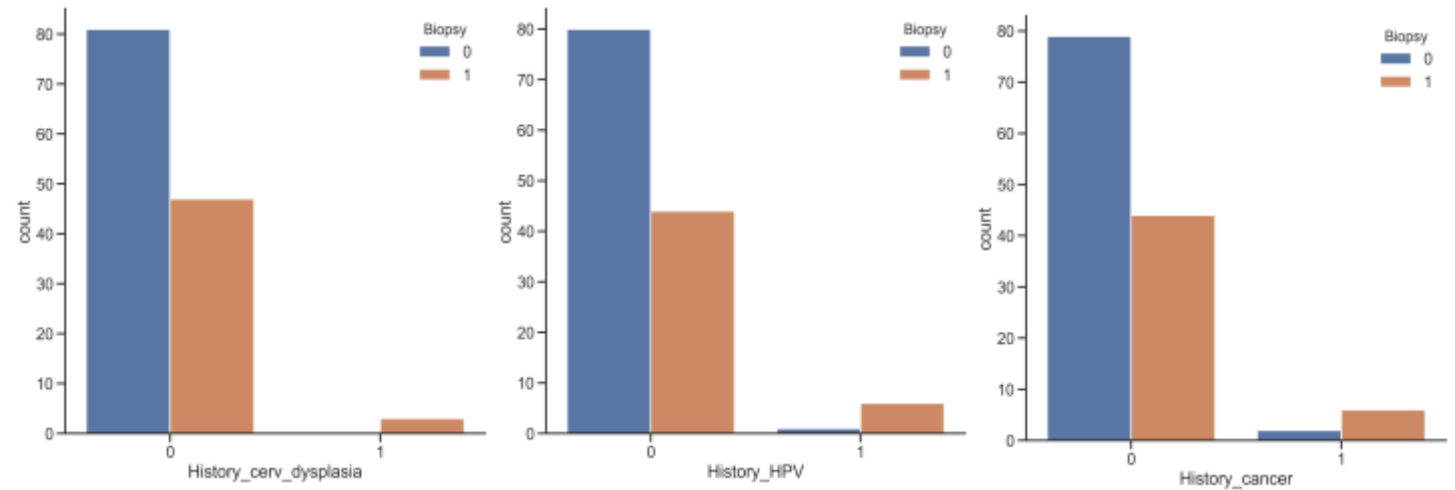
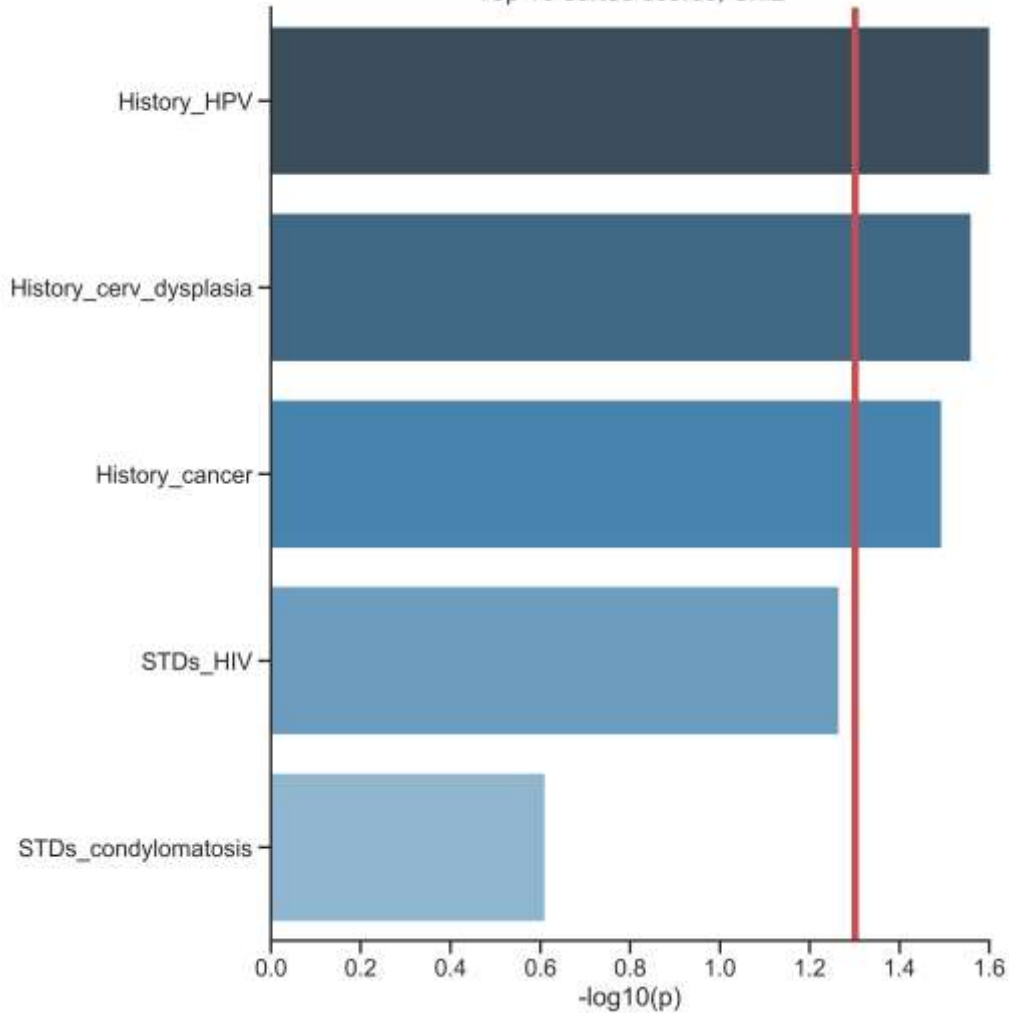


Statistical tests, Mann-Whitney U-test



Statistical tests, Chi-squared test

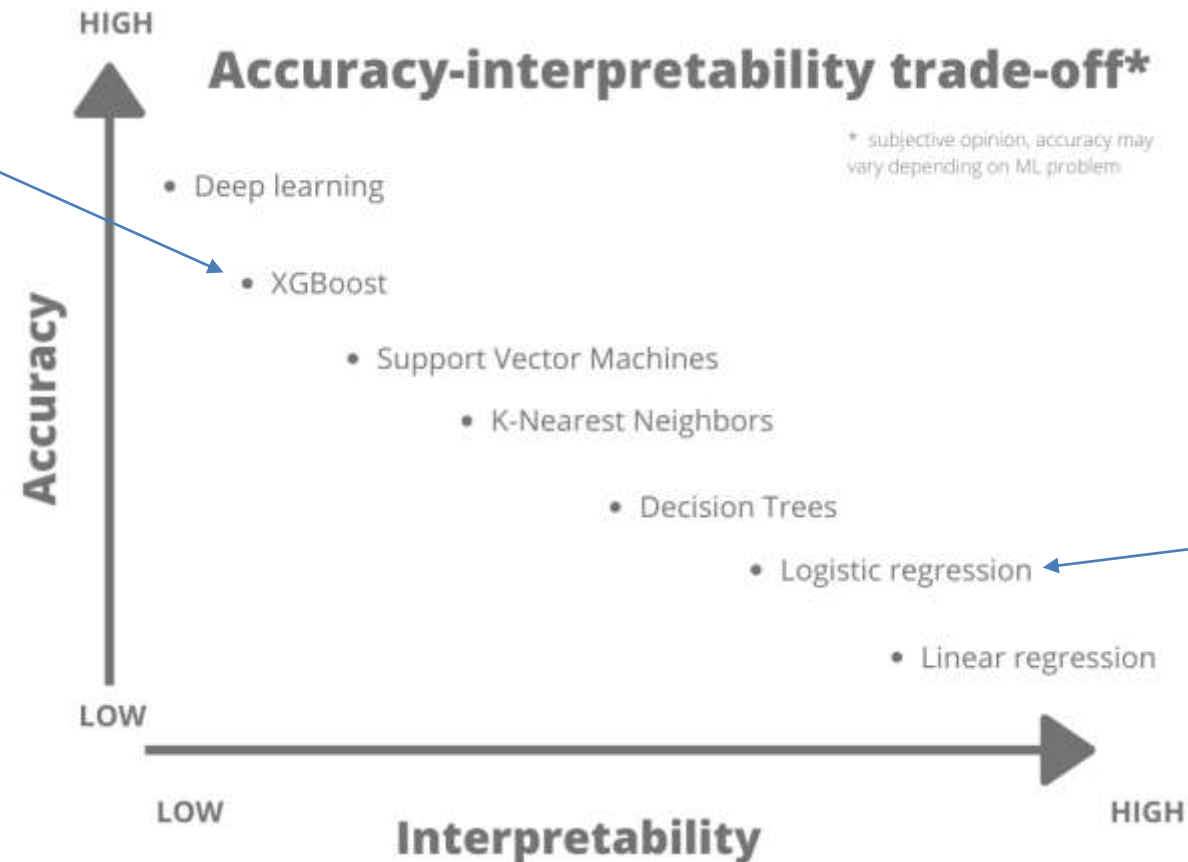
Top 10 sorted scores, Chi2



Motivation for Explainable ML



Published in 2016



Published in 1845



Linear models

The Logistic Function

The diagram shows the logistic function equation: $\text{Log} \left[\frac{Y}{(1-Y)} \right] = b_0 + b_1X_1 + b_2X_2 + b_3X_3 + \dots + b_nX_n$. The left side of the equation is enclosed in a blue oval, with a blue arrow pointing from the text "Log(Likelihood)" below it to the oval. The right side of the equation is annotated with three blue arrows pointing to the coefficients b_1 , b_2 , and b_3 . Below b_1 is the text "diet score (0-15)". Below b_2 is the text "age group (0/1)". Below b_3 is the text "sex (0/1)".

$$\text{Log} \left[\frac{Y}{(1-Y)} \right] = b_0 + b_1X_1 + b_2X_2 + b_3X_3 + \dots + b_nX_n$$

Log(Likelihood)

diet score (0-15)

age group (0/1)

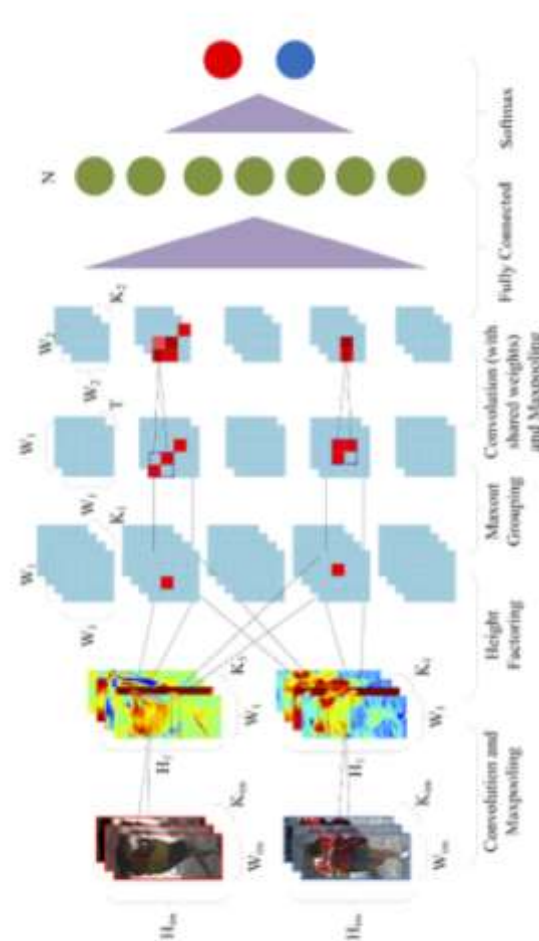
sex (0/1)

Non-linear models




Ah yes, something cool is happening in node 750,345,167...

Maybe we'll see something awesome if we jiggle the inputs?



Local explanations




Complex models are inherently complex!



But a single prediction involves only a small piece of that complexity

Surrogate linear model



Global Interpretation

Local Interpretation

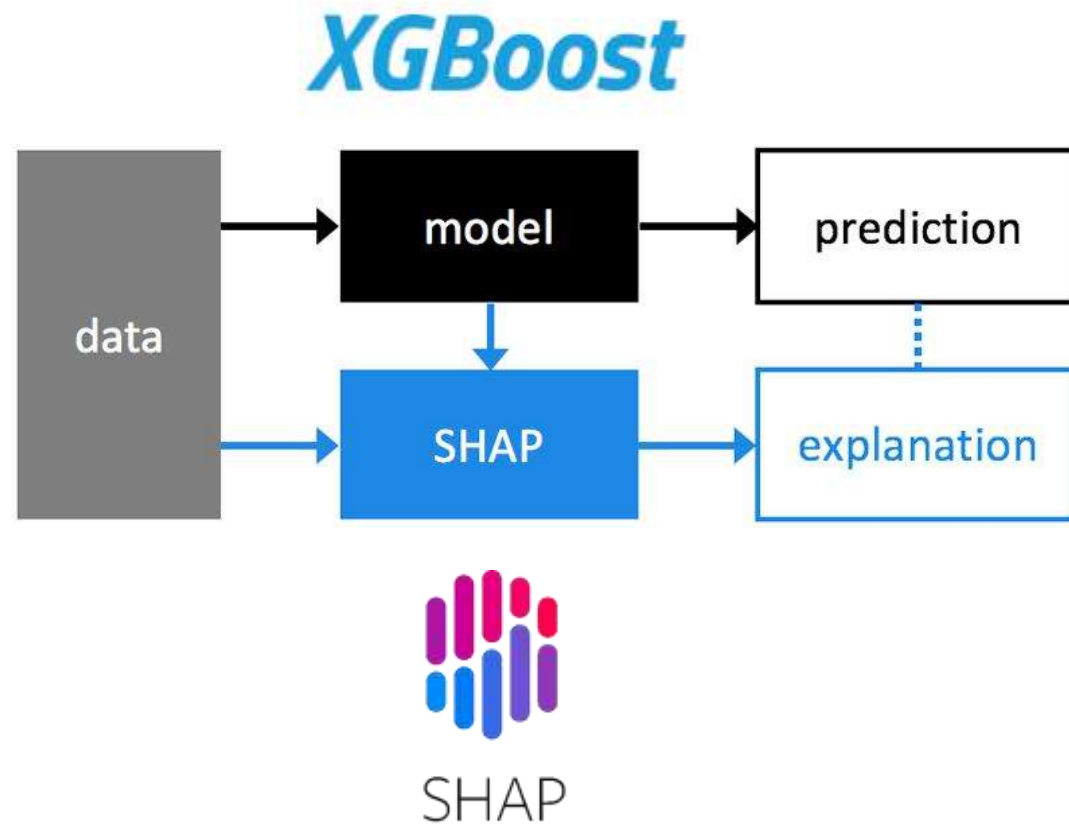


Averages effects over data dimensions



Assesses individual effects

Framework



SHAP Explainer

SHapley Additive exPlanations




Lloyd Shapley (1923 – 2016),
2012 Nobel Prize in Economics

- A method from coalitional game theory
 - Players cooperate in a coalition and receive a certain profit from this cooperation
- Features of the model
 - “*Players*”
- Outcome of the model
 - “*Profit*”

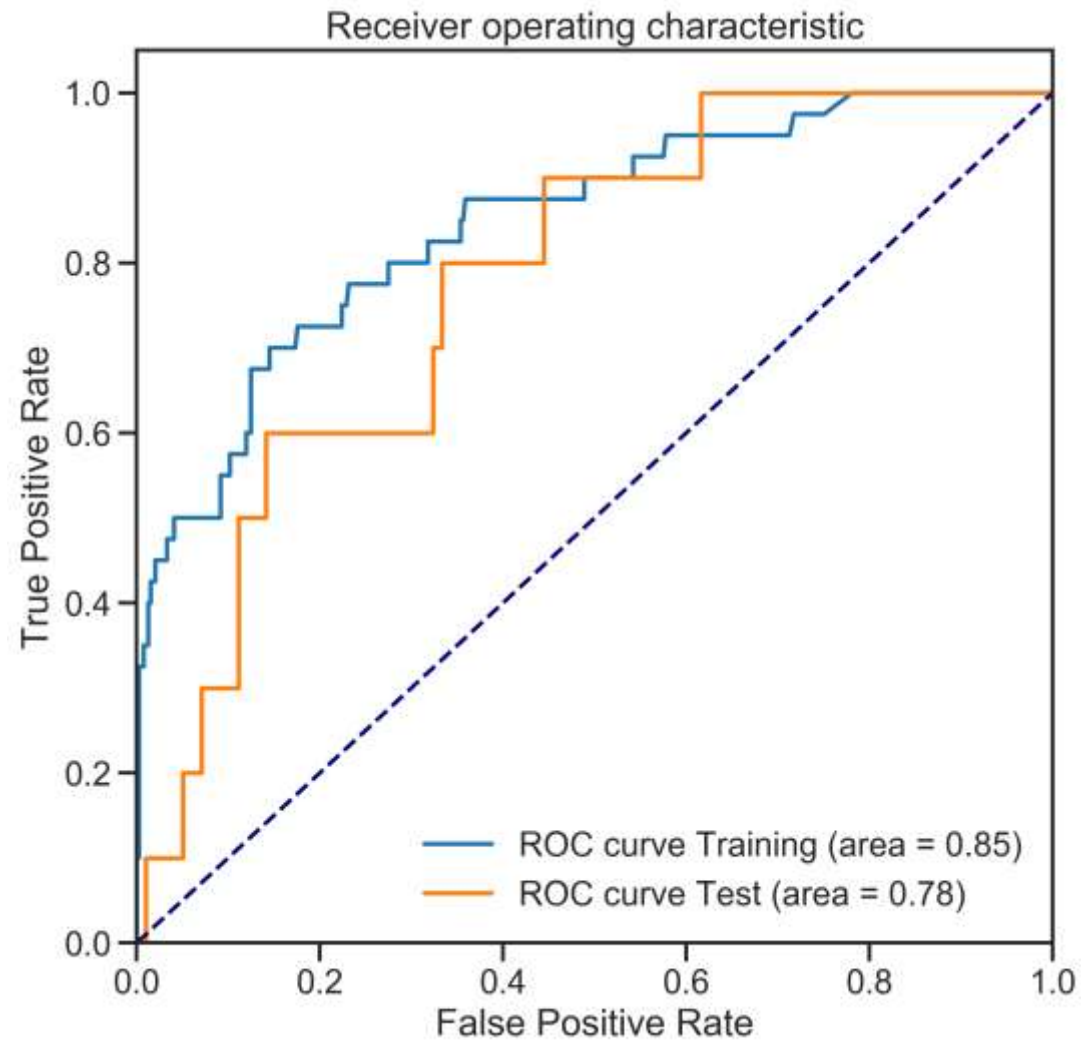
Shapley value in a nutshell

F1	F2	F3	F4
X			~
	X		~
		X	~
X	X		~
	X	X	~
X		X	~
X	X	X	~




How much did the feature contribute to the model's prediction in *all possible* combinations?

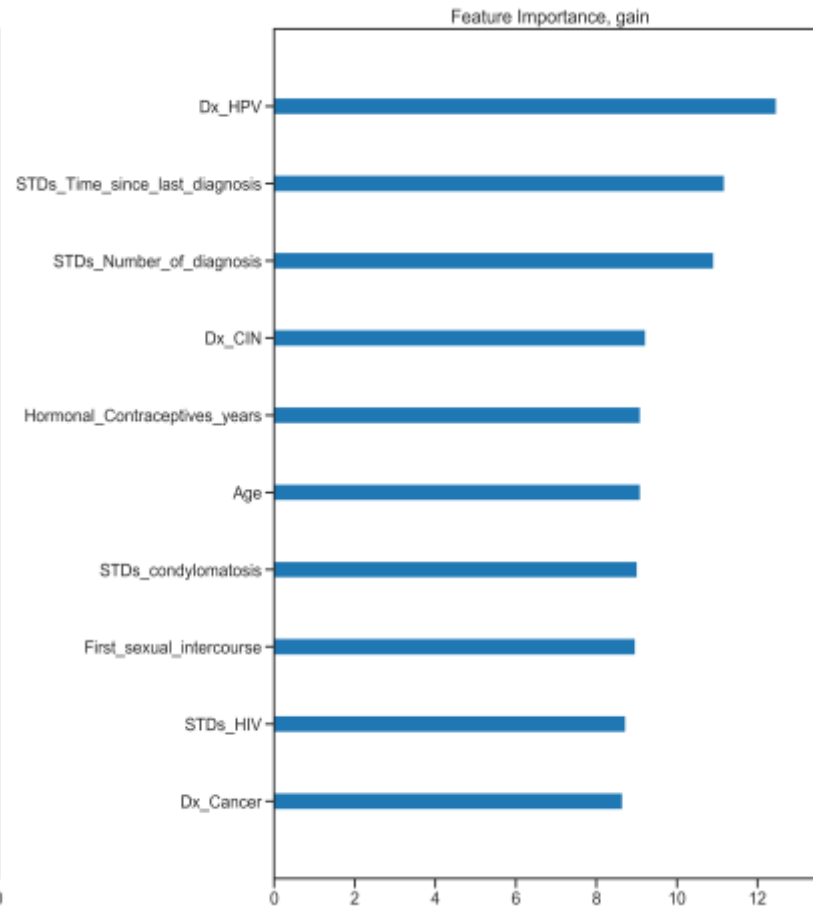
XGBoost performance




Feature importance inconsistency



Cover

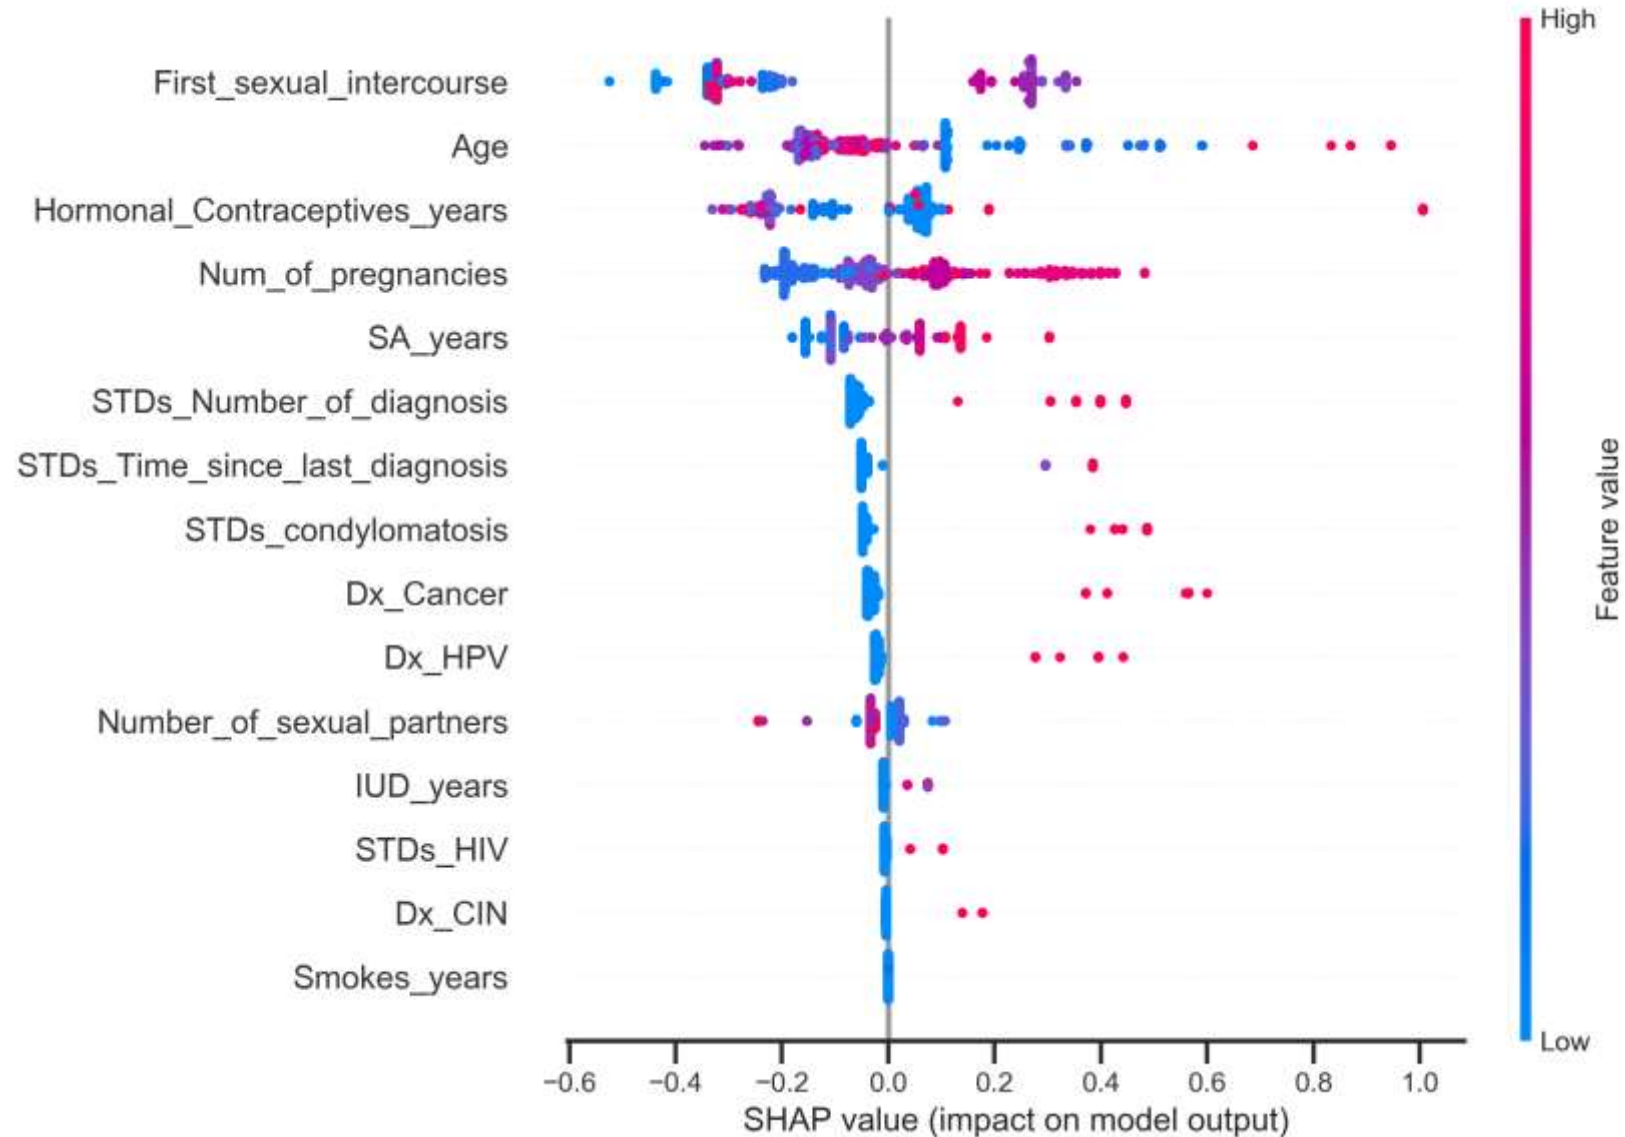


Gain




Weight


SHAP output




Sexual debut age



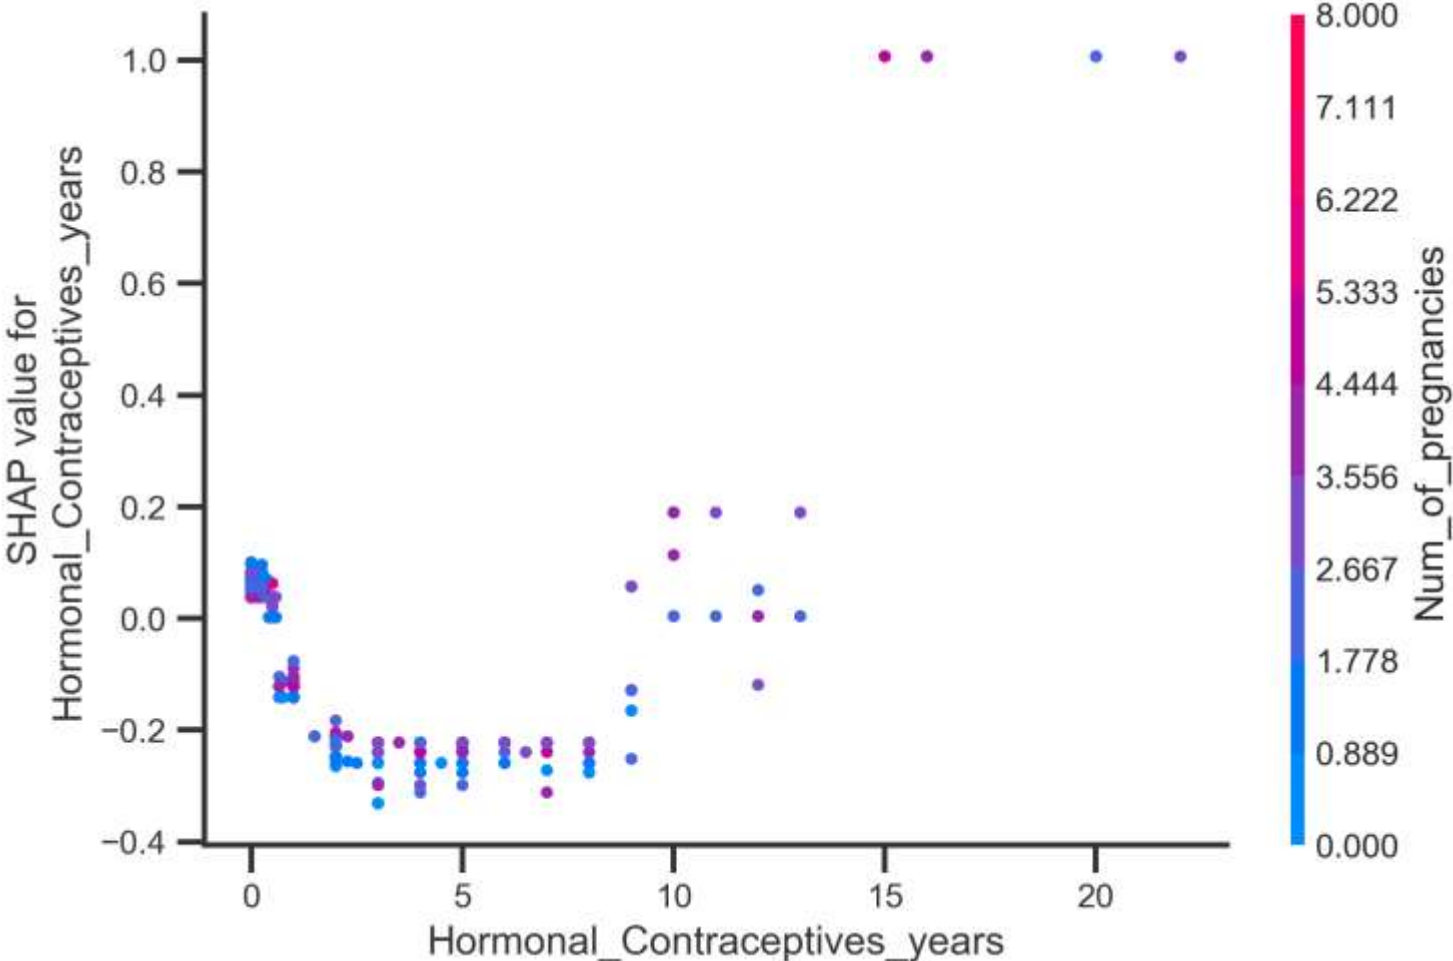
Biological age



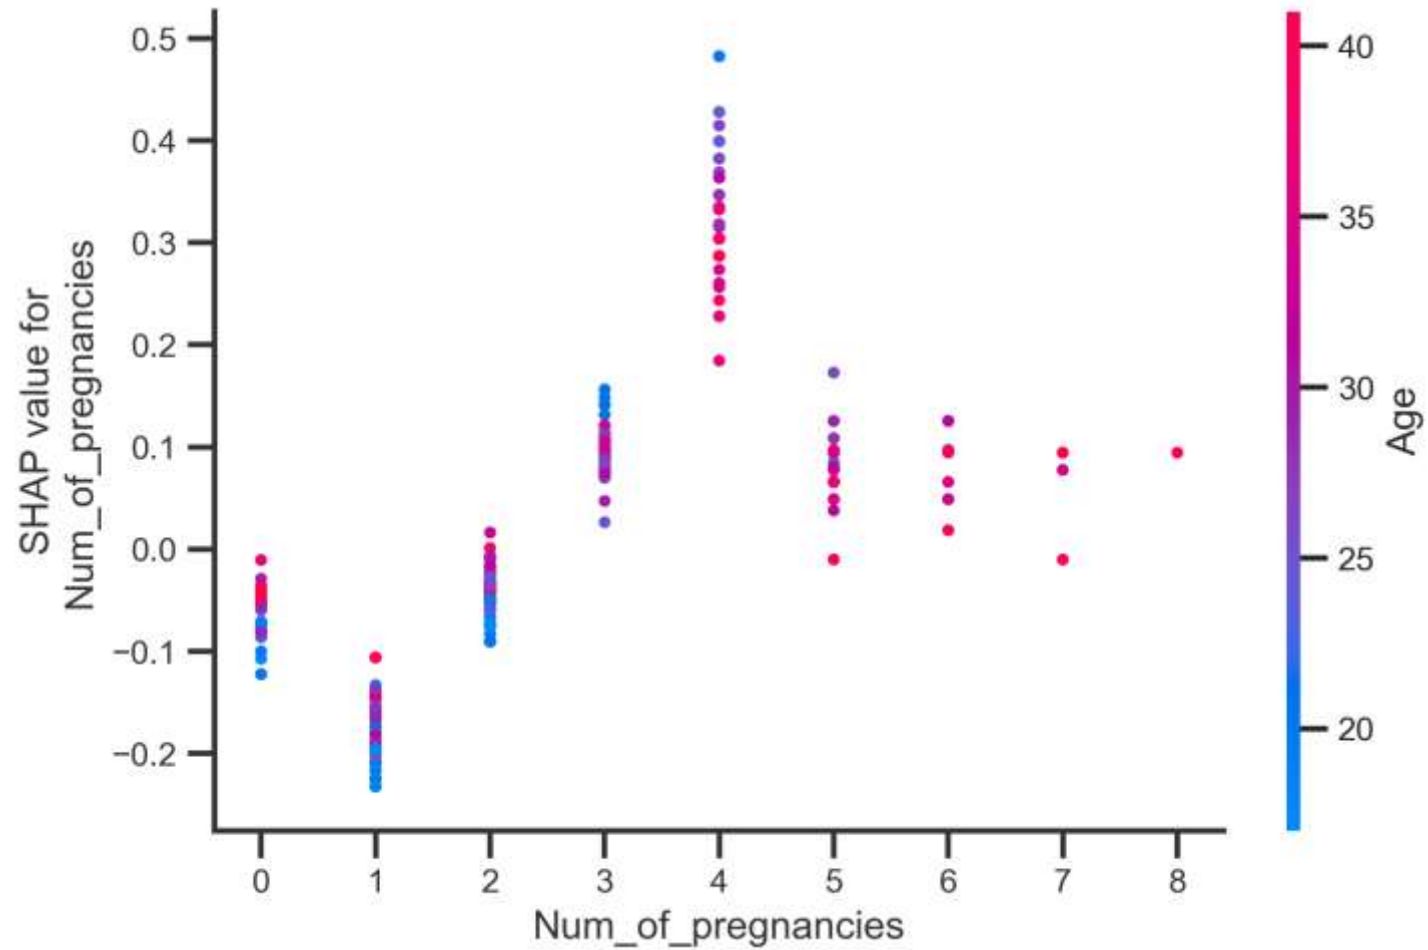
Sexually active age



Hormonal contraception duration



Parity*



*Parity = the number of times that a woman has given birth to a foetus with a gestational age of 24 weeks or more, regardless of whether the child was born alive or was stillborn
Gravidity = number of times that a woman has been pregnant

Examples of positive class predictions

True Positive, $P=0.67$

higher → lower
output value

base value

0.67

0.2

0.3

0.4

0.5

0.6

0.7

Num_of_pregnancies = 0.0

First_sexual_intercourse = 17.0

Age = 21.0

Dx_CIN = 1.0

False Positive, $P=0.63$

higher → lower

output value

base value

0.63

0.2

0.3

0.4

0.5

0.6

0.7

STDs_Time_since_last_diagnosis = 4.0

First_sexual_intercourse = 18.0

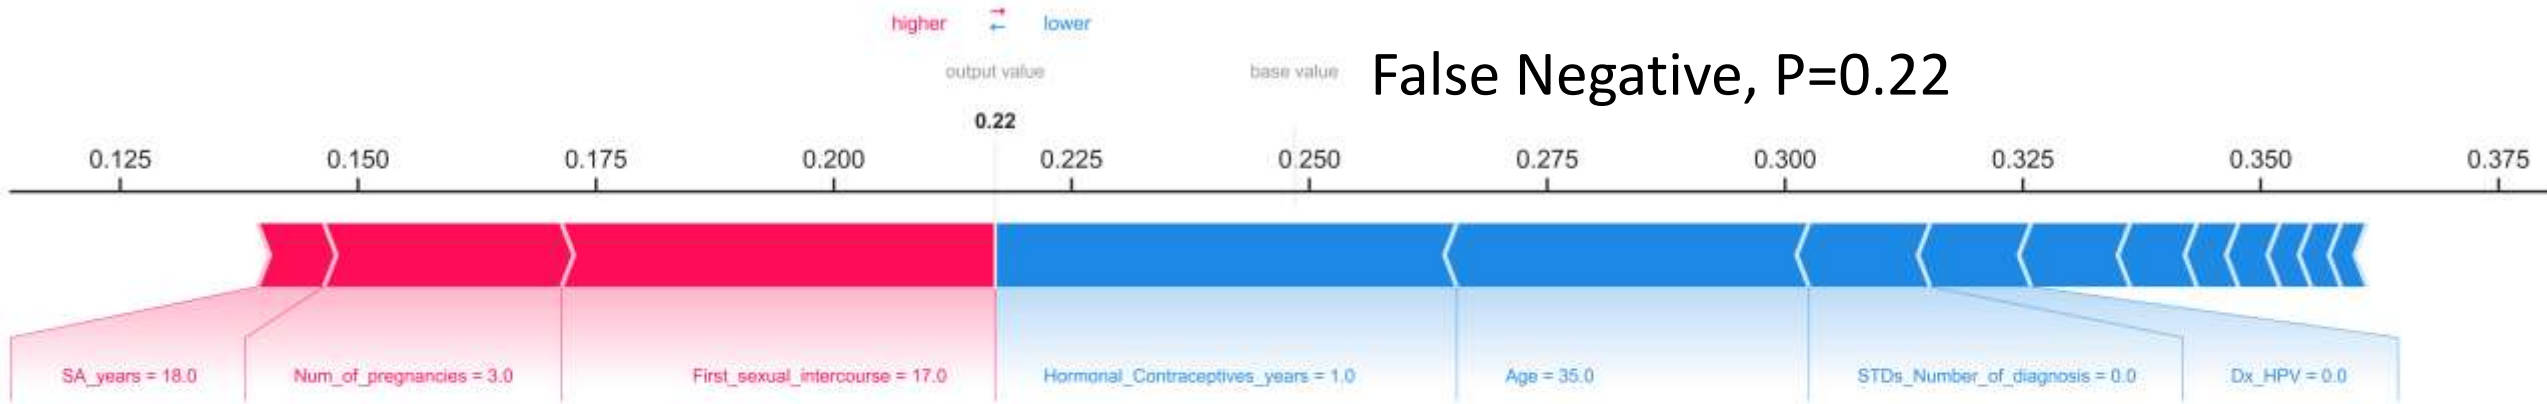

STDs_Number_of_diagnosis = 1.0

Age = 20.0

STDs_condylomatosis = 1.0

Num_of_pregnancies = 1.0

Examples of negative class predictions



Conclusions

- Good predictive performance
 - Training ROC AUC = 0.85, Test ROC AUC = 0.78
- Explainable Machine Learning identified important features
 - Sexually active age
 - Hormonal contraception usage
 - Number of pregnancies
 - History of STDs
- Non-linear dependency between features and predicted cervical cancer risk

Discussion