Prescriptive Neural Networks

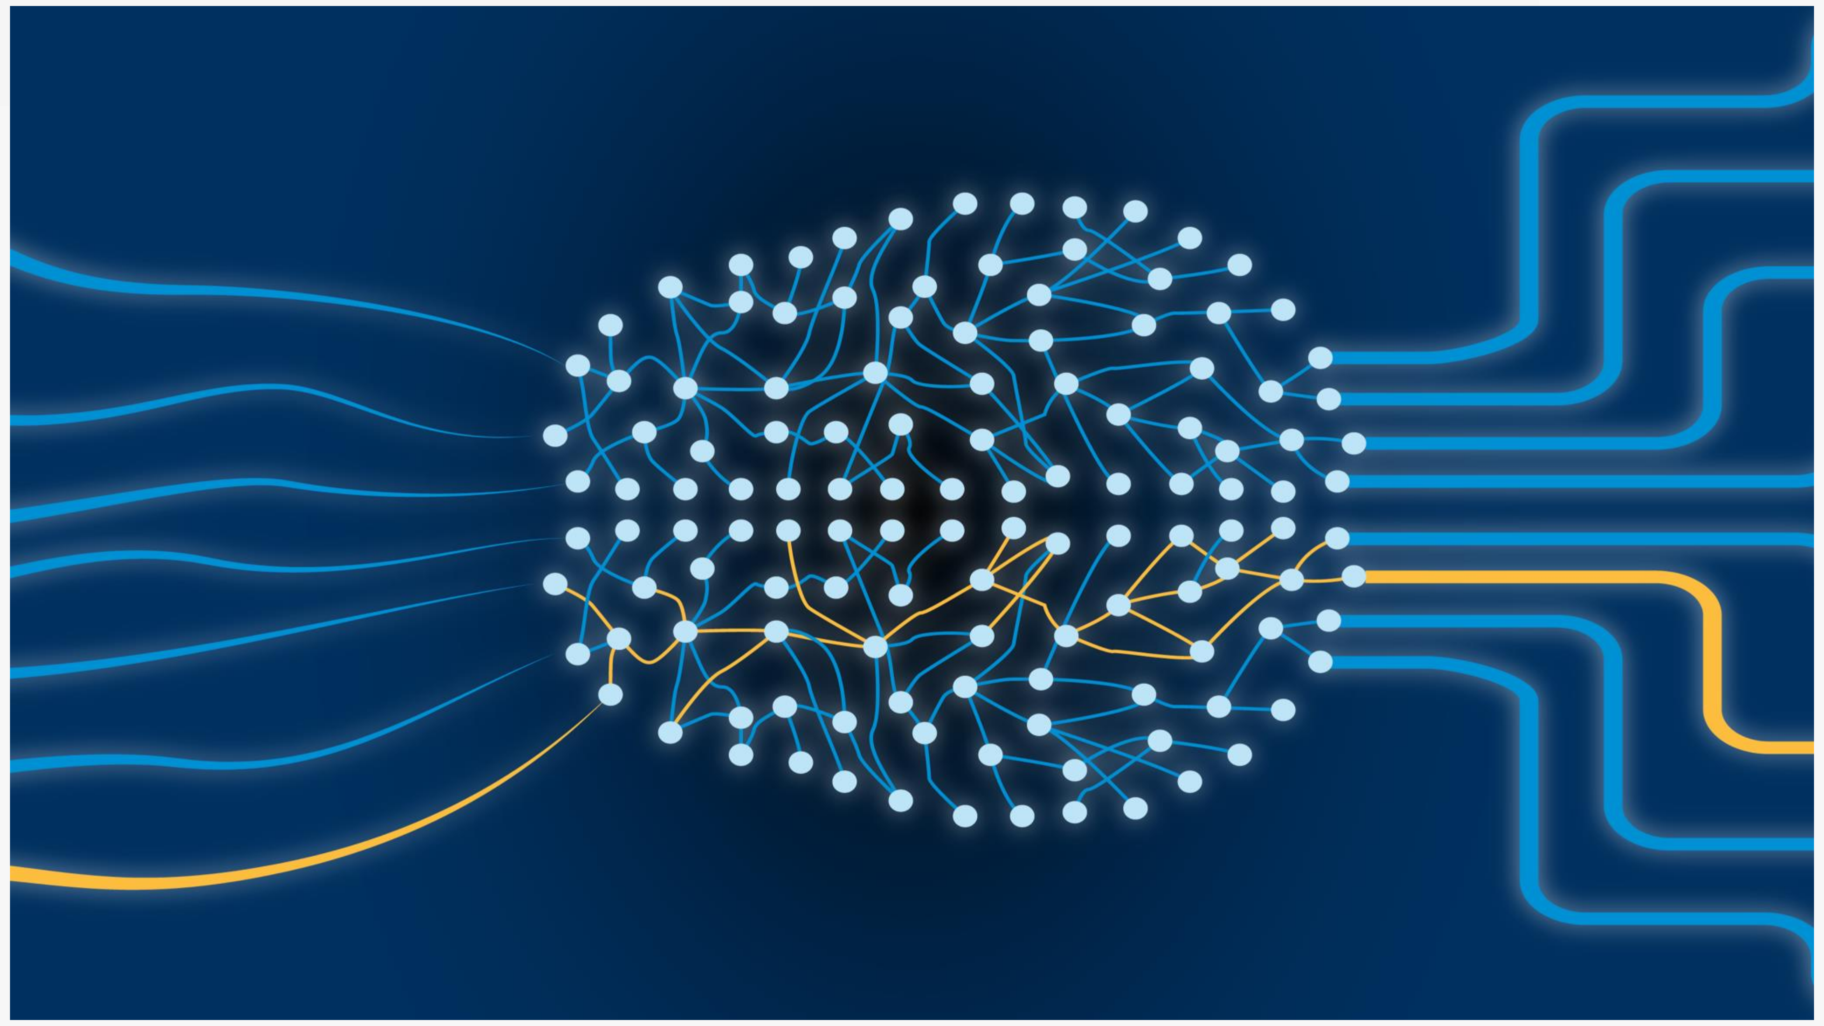

https://www.baslerweb.com/en/learning/deep-learning/

Prescribing optimal treatments for optimal outcomes

Dimitris Bertsimas, Lisa Everest, Vasiliki Stoumpou

Motivation for personalized, data-driven healthcare

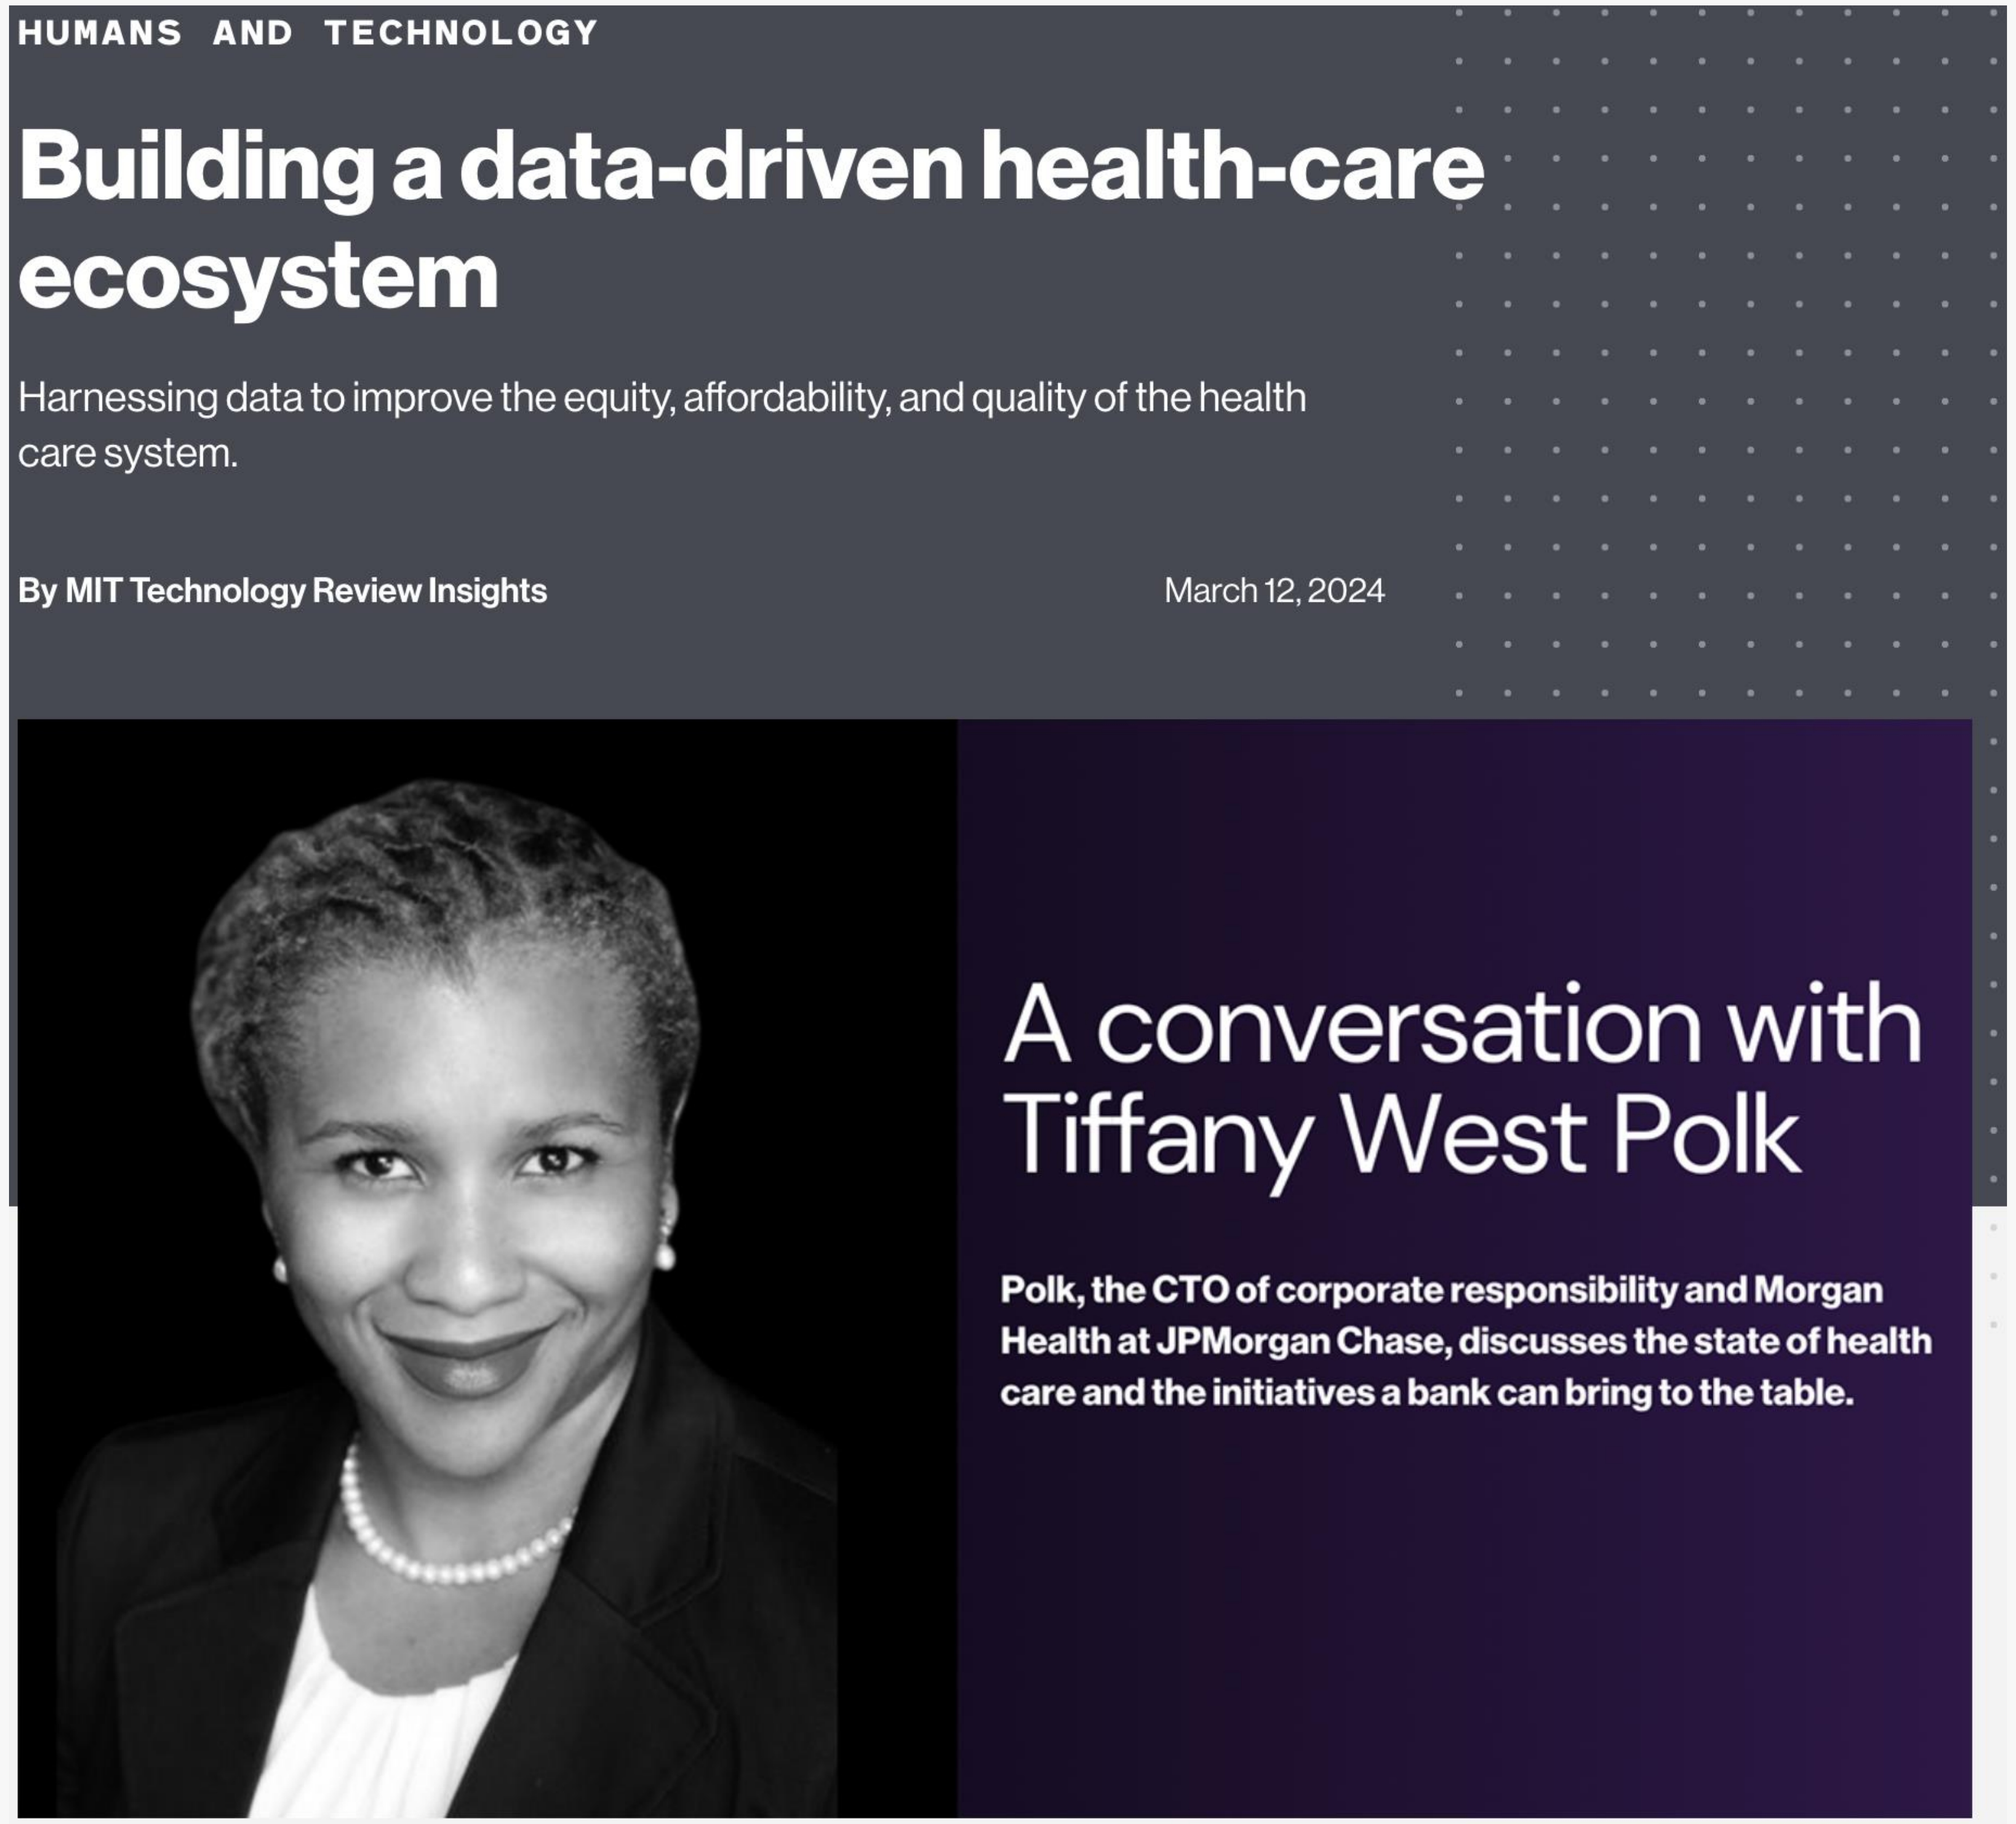
HUMANS AND TECHNOLOGY

Building a data-driven health-care ecosystem

Harnessing data to improve the equity, affordability, and quality of the health care system.

By MIT Technology Review Insights

March 12, 2024

A conversation with Tiffany West Polk

Polk, the CTO of corporate responsibility and Morgan Health at JPMorgan Chase, discusses the state of health care and the initiatives a bank can bring to the table.

Which medication?
What dosage?
Should we intervene?

Prescriptive Neural Networks

What is the “optimal prescription problem”?

In a nutshell, the “optimal prescription problem” is the following question: what treatment do I prescribe a given individual to optimize a specific outcome?

Observational data
$\{ \boldsymbol{x}_i \}_{i=1,\ldots,n}$

Treatments
$t \in \boldsymbol{T}$

Single discrete (binary) | *Multiple discrete* | *Single continuous* | *Multiple continuous*

Outcomes
$\{ y_i \}_{i=1,\ldots,n}$

Patient demographics, current prescriptions, medical data, notes, images, time series

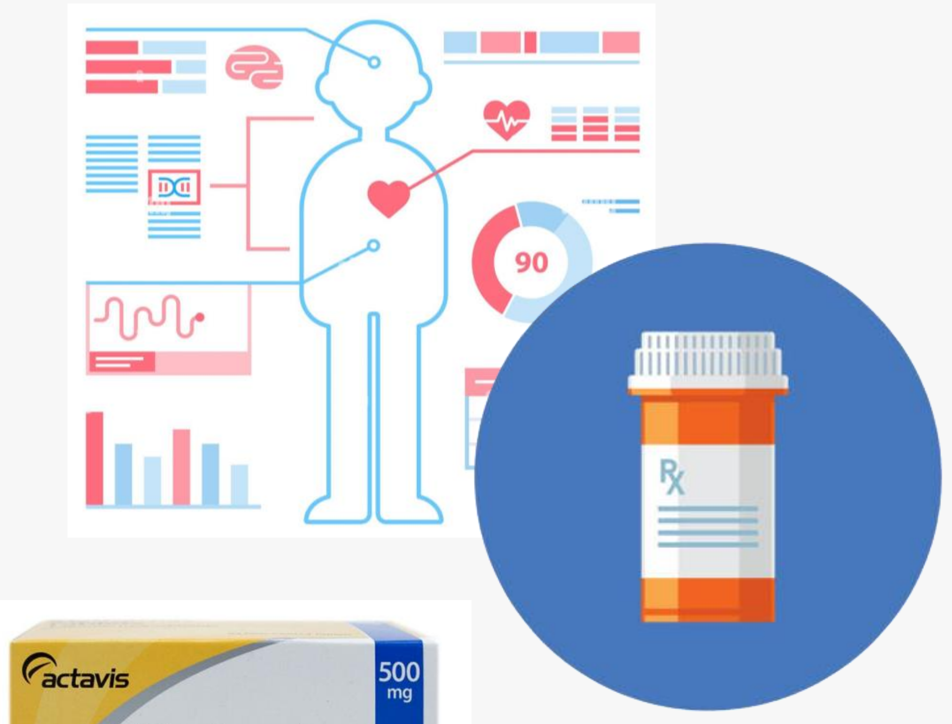

Drug doses of metformin, insulin, and oral blood glucose regulation agents

Patient blood glucose level (HBA_{1c})

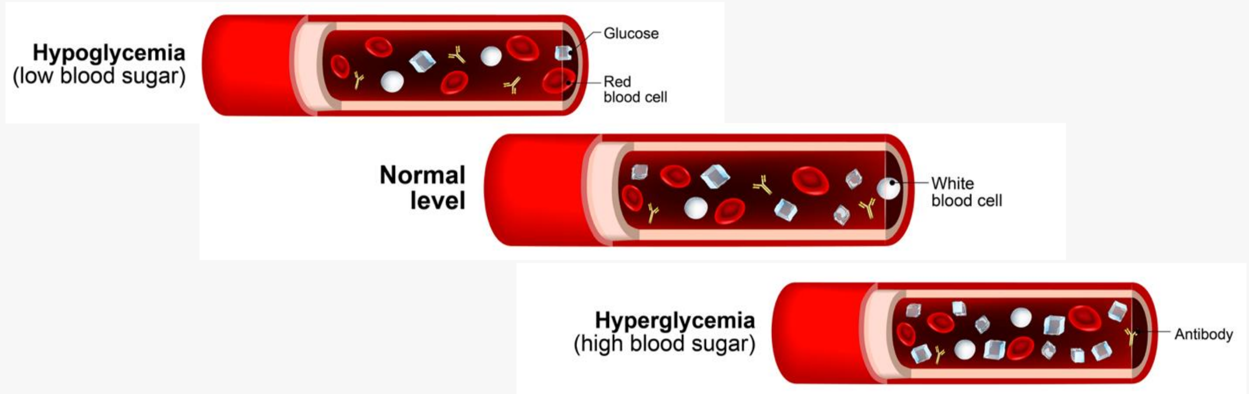

[Minimization Problem]

For a given observation $\boldsymbol{x}_i$, what is the optimal treatment t that minimizes the outcome y_i?

What are some ways to solve this problem?

- Regress and compare[1,2,3], where a regression model is trained to predict the outcomes under all treatments and the best one is chosen
- Causal Forests[4], Causal deep learning models[5,6]
- Tree-based methods, including Optimal Policy Trees[7] and Optimal Prescriptive Trees[8]
- Prescriptive Neural Networks

1. Egil Ferkingstad, Anders L.land, and Mathilde Wilhelmsen. Causal modeling and inference for electricity markets. Energy Economics, 33(3), 2011.
2. Dimitris Bertsimas and Nathan Kallus. From predictive to prescriptive analytics. Management Science, 66(3):1025-1044, 2020.
3. Arman Alam Siddique, Mireille E Schnitzer, Asma Bahamyirou, Guanbo Wang, Timothy H Holtz, Giovanni B Migliori, Giovanni Sotgiu, Neel R Gandhi, Mario H Vargas, Dick Menzies, and Andrea Benedetti. Causal inference with multiple concurrent medications: A comparison of methods and an application in multidrug-resistant tuberculosis. Stat Methods Med Research, 28(12):3534–3549, 2019.
4. Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, 113(523):1228–1242, 2018. doi: 10.1080/01621459., 2017.1319839. URL https://doi.org/10.1080/01621459.2017.1319839.
5. Uri Shalit, Fredrik D. Johansson, and David Sontag. Estimating individual treatment effect: generalization bounds and algorithms, 2017. URL https://arxiv.org/abs/1606.03976.
6. Claudia Shi, David M. Blei, and Victor Veitch. Adapting neural networks for the estimation of treatment effects, 2019. URL https://arxiv.org/abs/1906.02120.
7. Maxime Amram, Jack Dunn, and Ying Daisy Zhuo. Optimal policy trees. Machine Learning, 11:2741–2768, 2022.
8. Dimitris Bertsimas, Jack Dunn, and Nishanth Mundru. Optimal prescriptive trees. INFORMS Journal on Optimization, 1(2):164–183, 2019.

Why Prescriptive Neural Networks (PNN)?

1. Prescriptive deep learning can leverage multimodal data
2. Trees can learn non-linear functions, but neural networks are empirically shown to have more modeling power
3. Neural networks scale better with large number of treatments

The goal is to use Neural Networks to minimize a differentiable loss function, which represents the estimated average outcome after applying the treatment, similarly to the Optimal Policy Trees

Multimodal deep learning approach: Prescriptive Neural Network

1. Embedding extraction

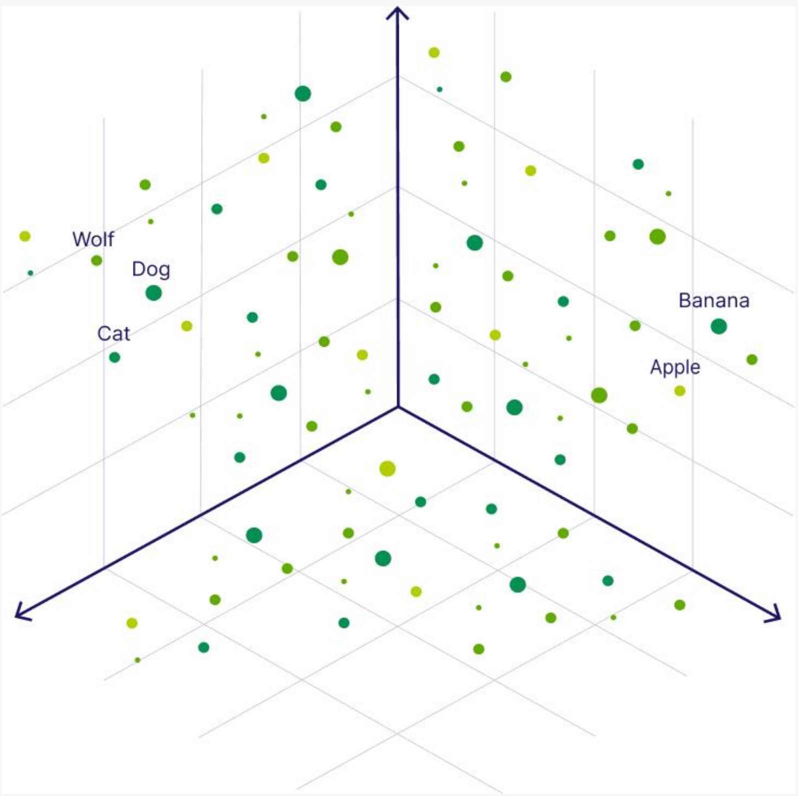

2. Counterfactual estimation

Patient i	Treatment t_1	Treatment t_2	Treatment t_3
$i = 1$	0.82	?	?
$i = 2$	?	0.36	?
…	…	…	…
$i = n - 1$	0.38	?	?
$i = n$	?	?	0.45

3. PNN: Train feedforward neural network with prescriptive objective function

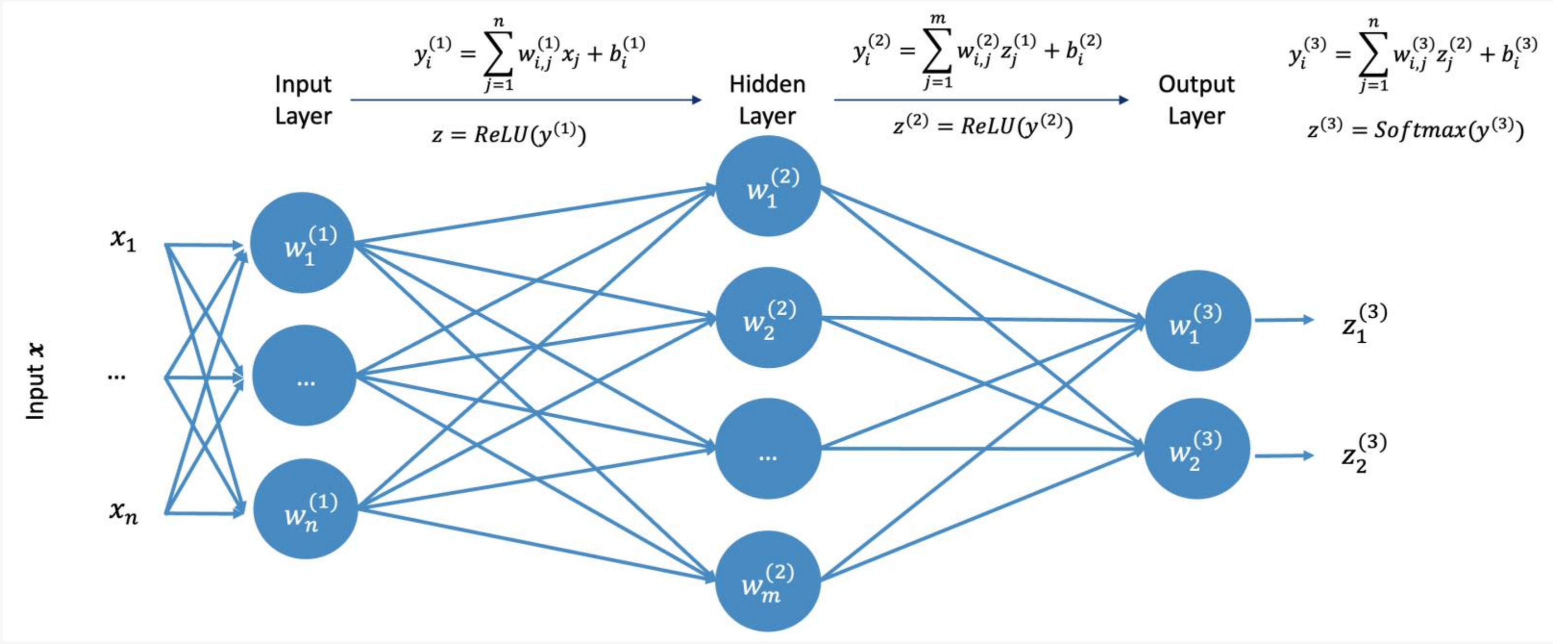

4. Mirrored OCT: Train Optimal Classification Tree (OCT) on the PNN's prescriptions as ground-truth targets

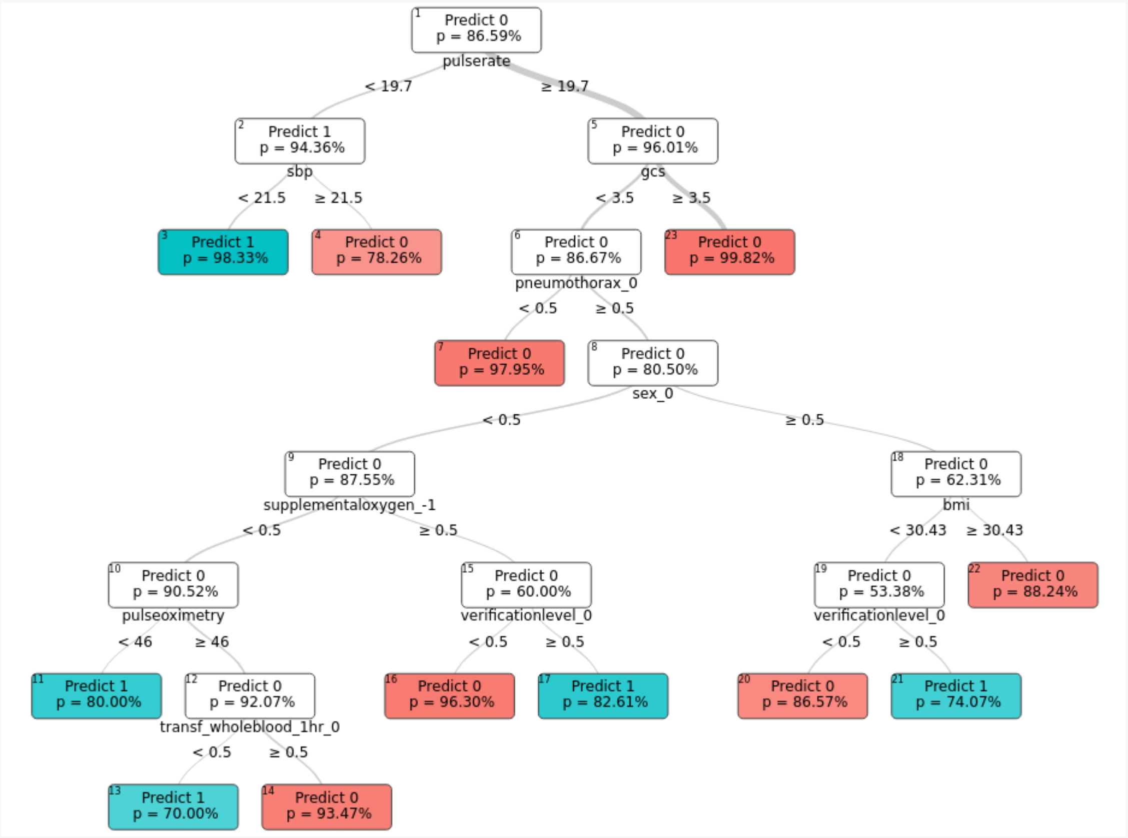

First step: embedding extraction

Structured data

- **Tabular data**: handle with appropriate processing (normalization etc.)

Unstructured data

- **Tabular data**: handle as structured data
- **Language data**: extract embeddings from pretrained models (e.g. Bert, Clinical Longformer)
- **Image data**: extract embeddings from pretrained models (e.g. DenseNet)

Second step: counterfactual estimation

Counterfactuals: unknown outcomes for an observation given a theoretical treatment

Patient i	Treatment t_1	Treatment t_2	Treatment t_3
$i = 1$	0.82	?	?
$i = 2$	?	0.36	?
…	…	…	…
$i = n - 1$	0.38	?	?
$i = n$	?	?	0.45

Second step: counterfactual estimation

Counterfactuals: unknown outcomes for an observation given a theoretical treatment

Counterfactual Estimation: use observed data to infer the unobserved outcomes so that proposed prescription policies could be evaluated

1. Direct method: one regression model for each treatment type (treatment assignment bias)
2. Doubly-robust method:
   a) Classification model to estimate propensity score (treatment assignment probabilities)
   b) Use propensity score to re-weight estimated outcomes from the direct method

Patient i	Treatment t_1	Treatment t_2	Treatment t_3
$i = 1$	0.82	0.76	0.78
$i = 2$	0.45	0.36	0.40
…	…	…	…
$i = n - 1$	0.38	0.39	0.38
$i = n$	0.53	0.44	0.45

Second step: counterfactual estimation

Counterfactuals: unknown outcomes for an observation given a theoretical treatment

Counterfactual Estimation: use observed data to infer the unobserved outcomes so that proposed prescription policies could be evaluated

1. Direct method: one regression model for each treatment type (treatment assignment bias)
2. Doubly-robust method:
   a) Classification model to estimate propensity score (treatment assignment probabilities)
   b) Use propensity score to re-weight estimated outcomes from the direct method

Patient i	Treatment t_1	Treatment t_2	Treatment t_3
$i = 1$	0.82	0.76	0.78
$i = 2$	0.45	0.36	0.40
…	…	…	…
$i = n - 1$	0.38	0.39	0.38
$i = n$	0.53	0.44	0.45

Rewards matrix Γ: where each element $\Gamma_{i,t}$ is the estimated outcome of applying treatment t to the patient i

Prescriptive Neural Networks

Third step: neural network

Input Layer → Hidden Layer → Output Layer

Input Layer → Hidden Layer:

$$y_i^{(1)} = \sum_{j=1}^{n} w_{i,j}^{(1)} x_j + b_i^{(1)}$$

$$z = ReLU(y^{(1)})$$

Hidden Layer → Output Layer:

$$y_i^{(2)} = \sum_{j=1}^{m} w_{i,j}^{(2)} z_j^{(1)} + b_i^{(2)}$$

$$z^{(2)} = ReLU(y^{(2)})$$

Output Layer:

$$y_i^{(3)} = \sum_{j=1}^{n} w_{i,j}^{(3)} z_j^{(2)} + b_i^{(3)}$$

$$z^{(3)} = Softmax(y^{(3)})$$

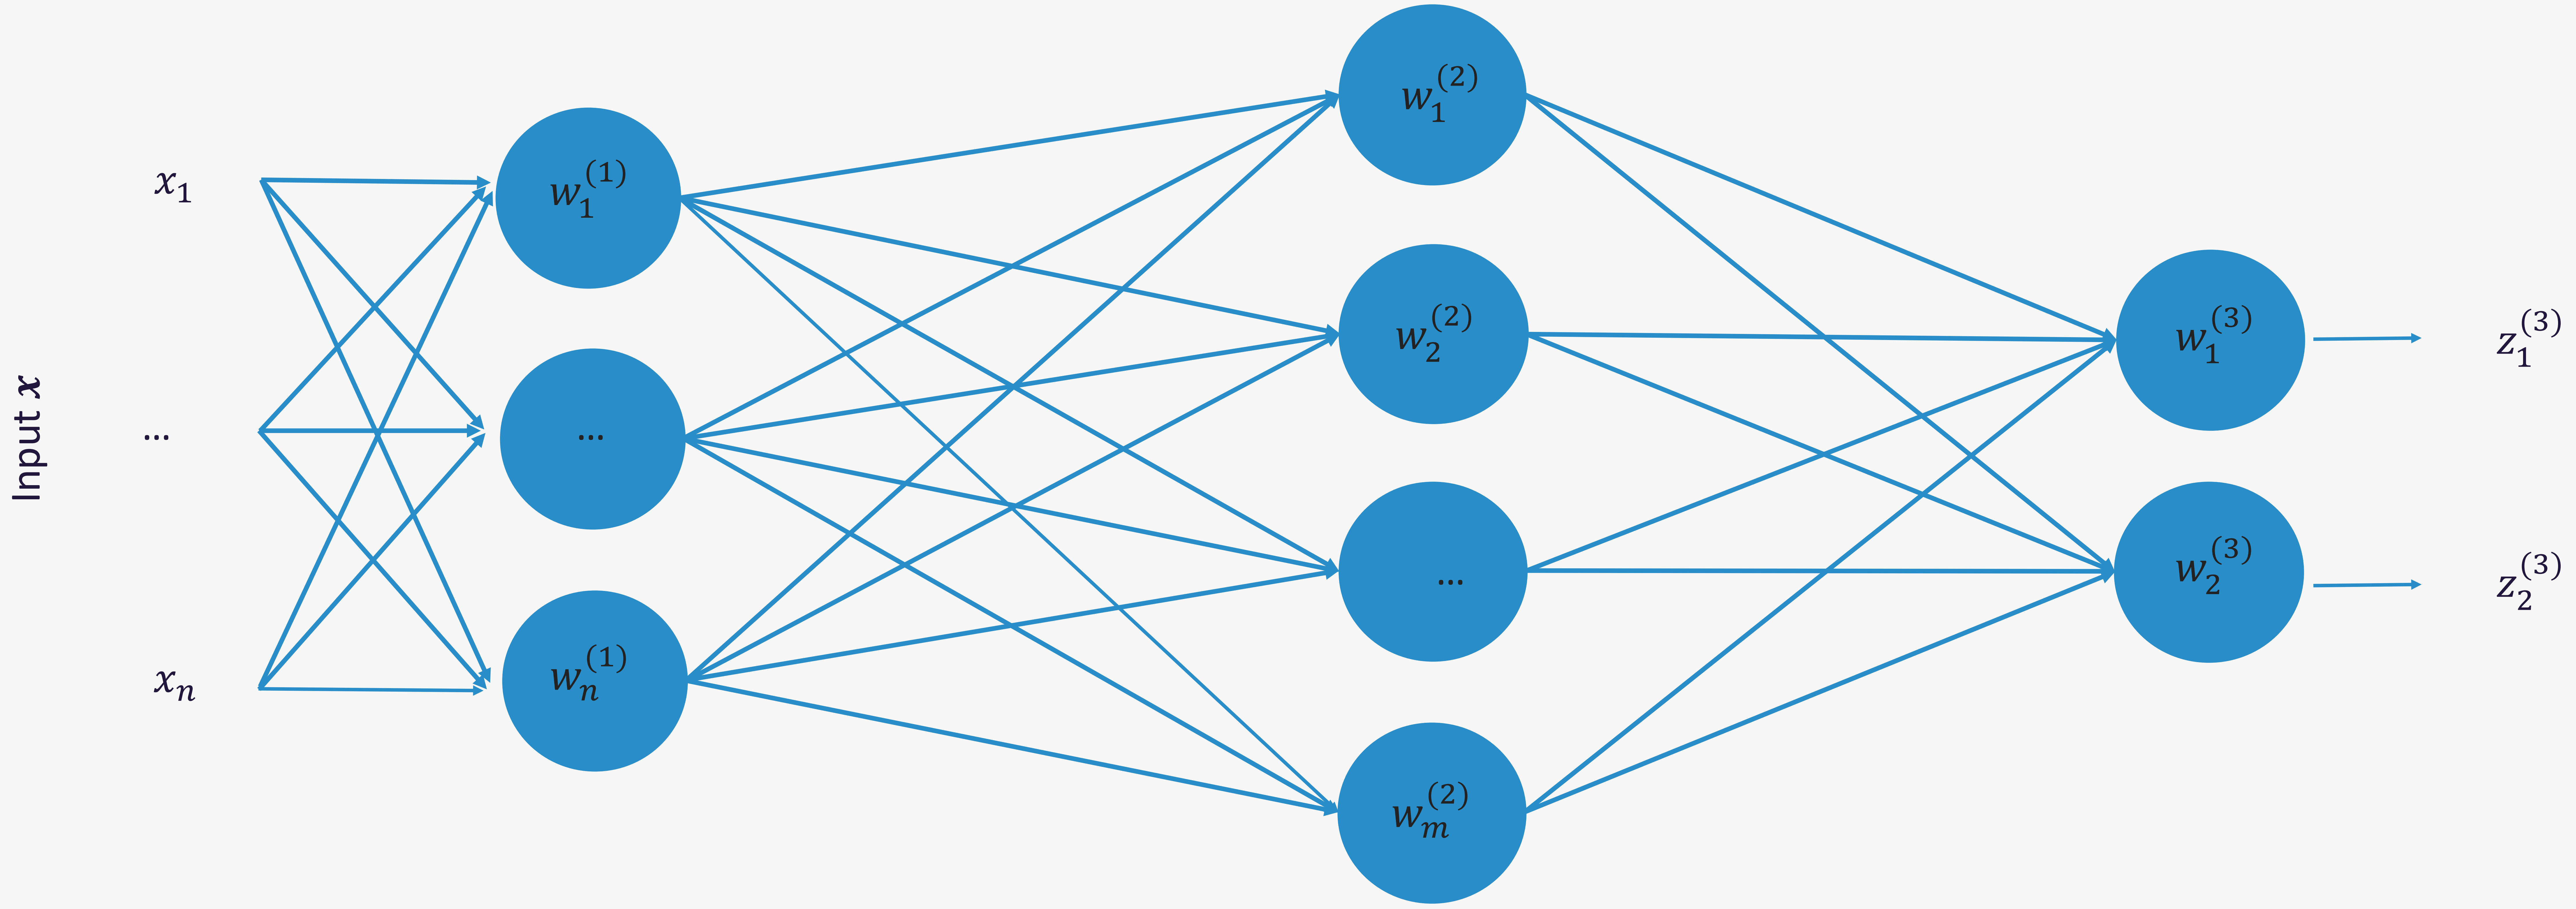

Third step: neural network

Let $\tau(\boldsymbol{x}_i)$= the treatment assigned by the network to patient i

$\Gamma_{i,t}$= the estimated outcome of patient i if treatment t is selected

N= the number of samples

N_t= the number of distinct treatments

Loss Function[1]:

$$min_{\tau(\cdot)} \frac{1}{N} \sum_{i=1}^{N} \sum_{t=1}^{N_t} \mathbb{I}\{\tau(\boldsymbol{x}_i) = t\} \cdot \Gamma_{i,t}$$

1. Maxime Amram, Jack Dunn, and Ying Daisy Zhuo. Optimal policy trees. Machine Learning, 11:2741–2768, 2022.

Third step: neural network

Let $\tau(\boldsymbol{x}_i)$= the treatment assigned by the network to patient i

$\Gamma_{i,t}$= the estimated outcome of patient i if treatment t is selected

N= the number of samples

N_t= the number of distinct treatments

z = the vector output of the network

(PNN's softmax output)

$$\mathbb{P}[\tau(x_i) = t] = \sigma_t(z) = \frac{\exp(z_t)}{\sum_{t=1}^{T} \exp(z_t)}$$

Loss Function:

$$min_{\tau(\cdot)} \frac{1}{N} \sum_{i=1}^{N} \sum_{t=1}^{N_t} \mathbb{P}[\tau(x_i) = t] \cdot \Gamma_{i,t}$$

Example

Counterfactuals			
Patient i	**Treatment t_1**	**Treatment t_2**	**Treatment t_3**
$i = 1$	**0.82**	**0.76**	**0.78**
$i = 2$	**0.45**	**0.36**	**0.40**
$i = 3$	**0.38**	**0.39**	**0.30**
$i = 4$	**0.53**	**0.44**	**0.45**

Prescriptions		
Patient i	**Optimal Policy Tree** $\mathbb{I}\{\tau(\boldsymbol{x}_i) = t\}$	**Prescriptive Neural Network** $\mathbb{P}[\tau(x_i) = t]$
$i = 1$	[0,1,0]	[0.2, 0.45, 0.35]
$i = 2$	[0,1,0]	[0.1, 0.6, 0.3]
$i = 3$	[0,0,1]	[0.45, 0.4, 0.15]
$i = 4$	[0,1,0]	[0.1, 0.5. 0.4]

Objective function (Optimal Policy Tree): 1*0.76 + 1*0.36 + 1*0.3 + 1*0.44

Objective function (Prescriptive Neural Network): (0.2*0.82 + 0.45*0.76 + 0.35*0.78) + (0.1*0.45 + 0.6*0.36 + 0.3*0.4) +

(0.45*0.38 + 0.4*0.39 + 0.15*0.3) + (0.1*0.53 + 0.5*0.44 + 0.4*0.45)

Fourth step: interpretability via Optimal Classification Trees

Neural networks are black-box models lacking interpretability

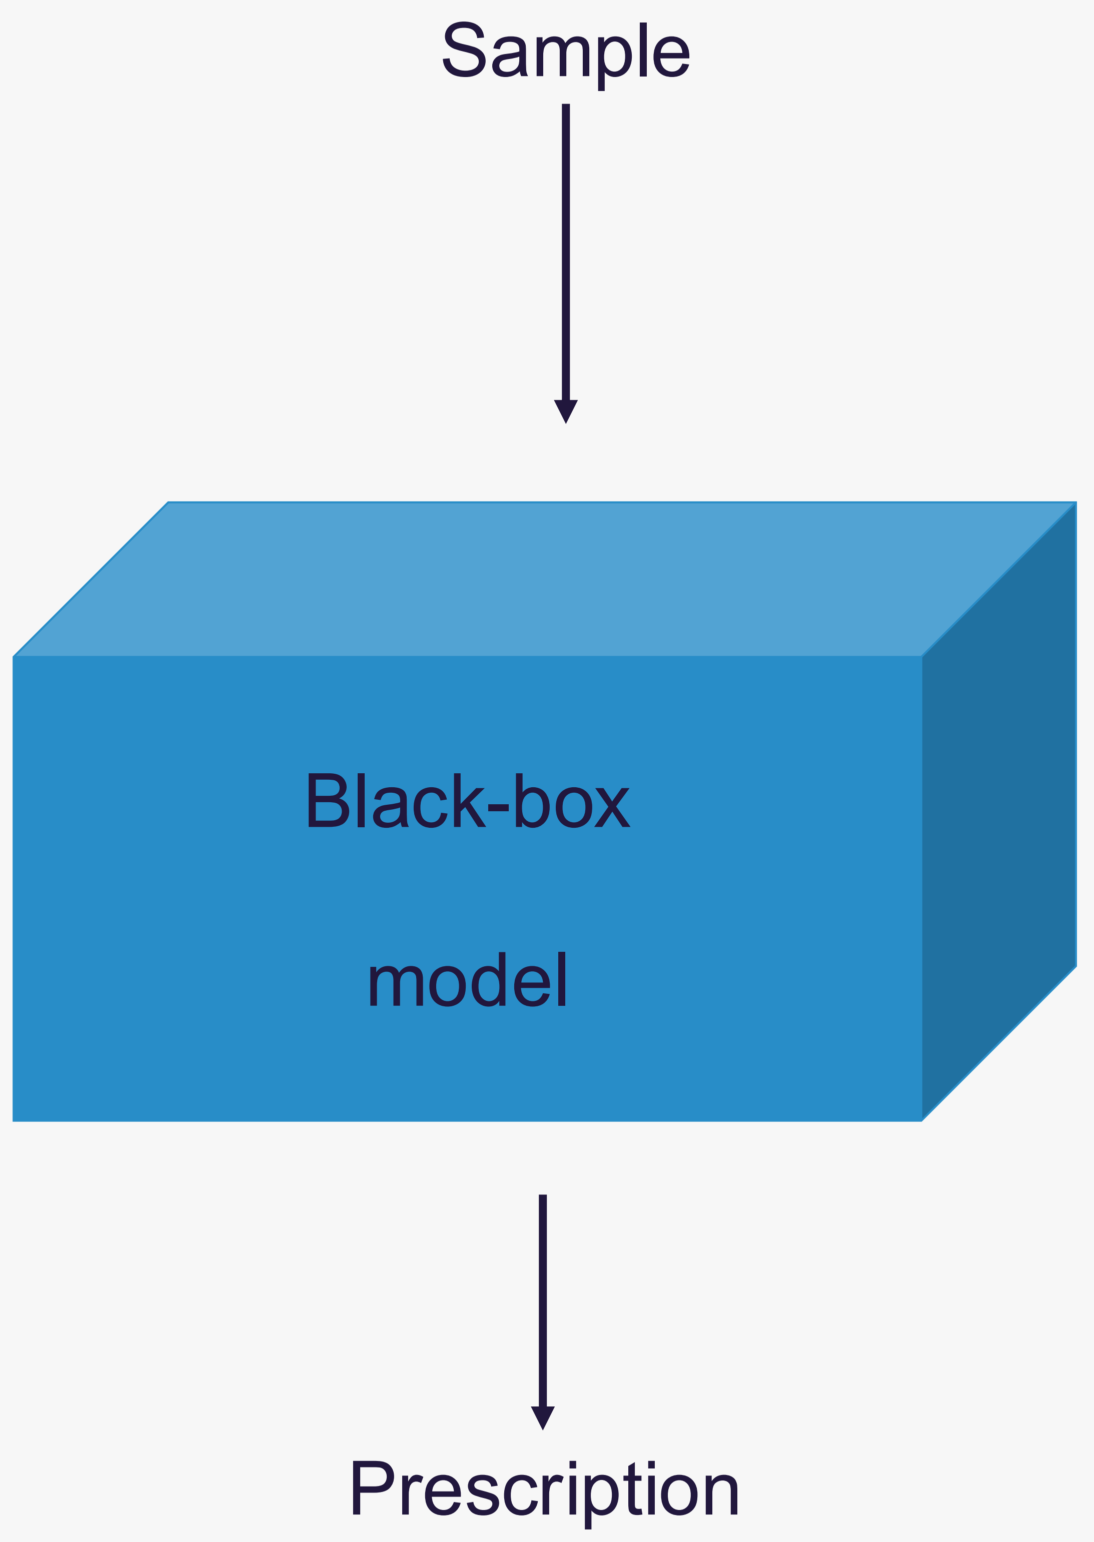

Fourth step: interpretability via Optimal Classification Trees

Neural networks are black-box models lacking interpretability

Tree models are extremely interpretable

Approach: Train “Optimal Classification Tree” (OCT)[1] using the original data’s features and the PNN’s prescriptions as the target classes (mirroring)

Resulting model is a Mirrored OCT

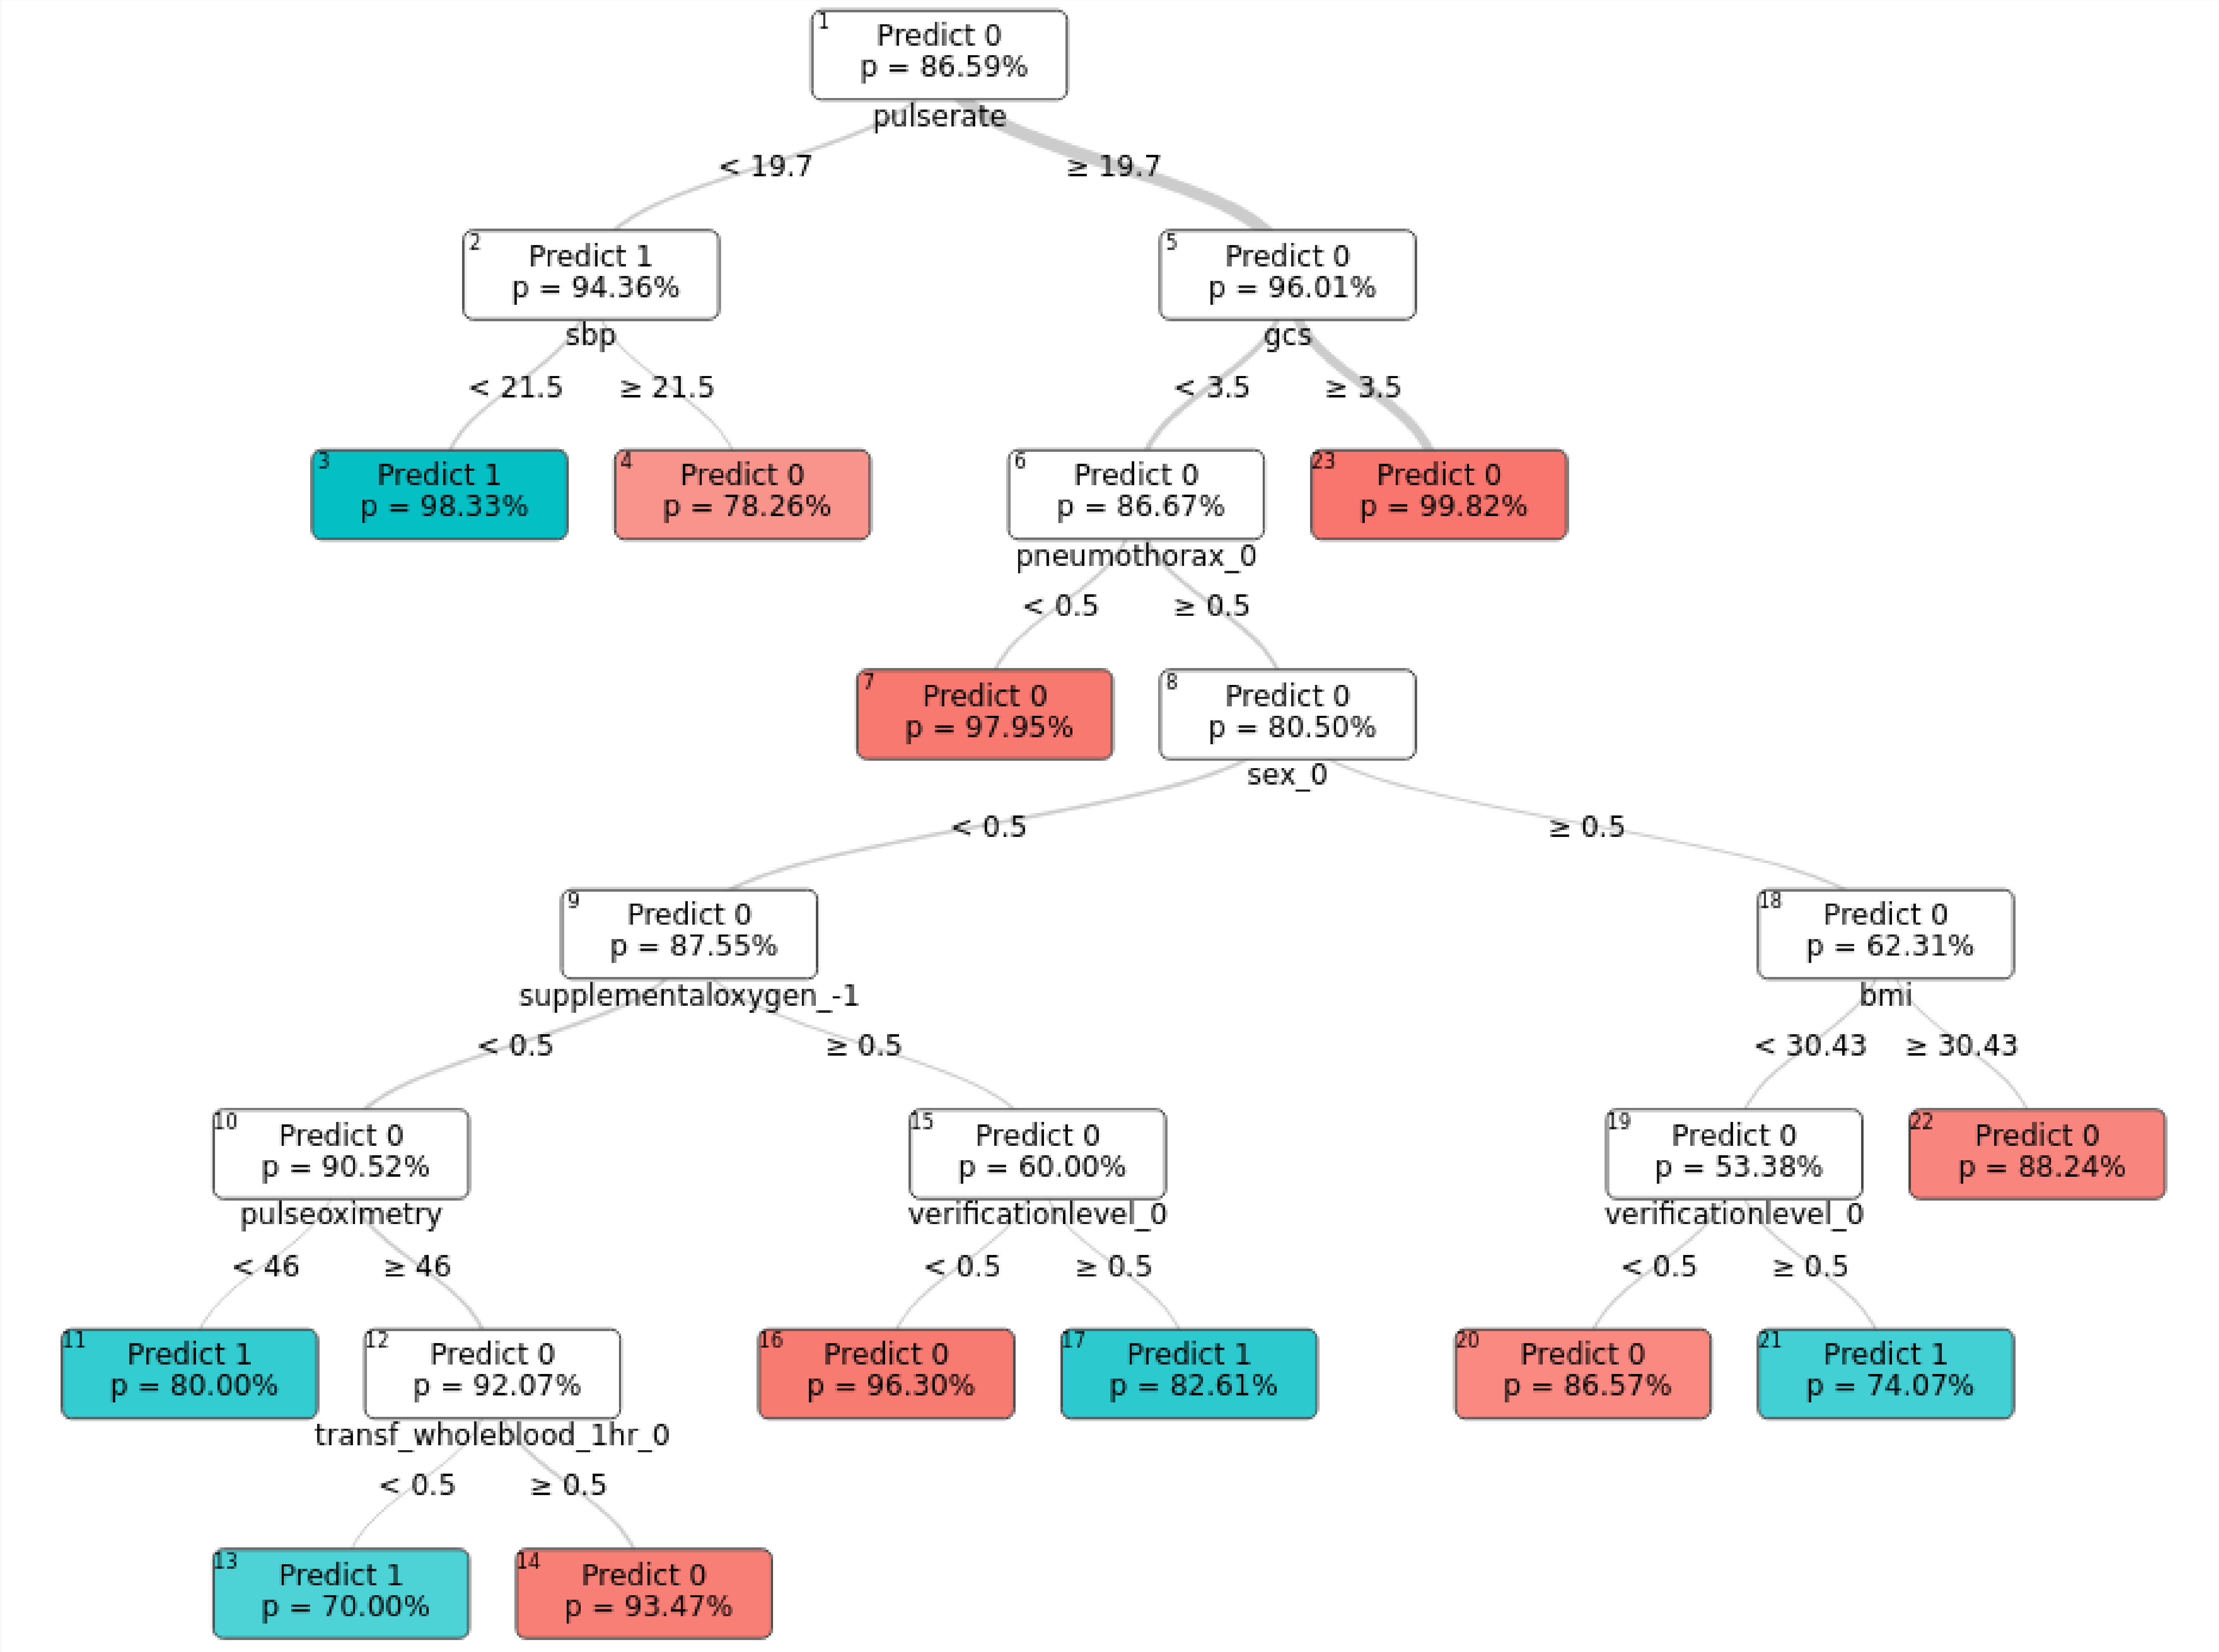

1. Dimitris Bertsimas and Jack Dunn. Optimal classification trees. Machine Learning, 106:1039–1082, 2017.

Experiments: how do we evaluate prescriptive models?

Models

- PNN
- Mirrored OCT
- Regress & Compare (structured only)
- Causal Forests (structured only)
- Optimal Policy Trees (structured only)

Procedure

- Split into many randomized 50% training/50% test sets
- Train different models for each model class, for each data split
- Report relative improvement

Evaluation metric:

Average relative improvement of the outcome on the test set:

$$\bar{I} = \frac{\sum_{i=1}^{n}(\Gamma_{i,\hat{t}_i} - \Gamma_{i,t_i})}{\sum_{i=1}^{n}\Gamma_{i,t_i}}$$

Experiments: TAVR optimal valve prescription

Dataset

Demographic and medical information, as well as radiology reports of echocardiograms and CT scans from 2,127 patients that are having the TAVR procedure

Binary treatment

EVOLUT PRO or SAPIEN 3 valve selection

Goal

Minimize PPI (Permanent Pacemaker Implantation) after the procedure

		Improvement (%)	
Counterfactual Estimator	Model Class	Tabular	Multimodal
Tabular	PNN	5.05 (±2.60)	17.87 (±6.24)
	Mirrored OCT	7.67 (±3.45)	17.14 (±7.41)
Multimodal	PNN	21.09 (±1.08)	**42.89 (±5.64)**
	Mirrored OCT	**22.58 (±2.50)**	41.66 (±7.00)

Experiments: Surgery prescription for liver injury patients

Dataset

Electronic medical records of 722 liver injury patients including patient demographics, history of illness, lab results, reports of liver CT scans

Binary treatment

Surgery or no surgery

Goal

Minimize mortality after the injury

		Improvement (%)	
Estimator	Model Class	Tabular	Multimodal
Tabular	PNN	14.85 (±4.39)	21.74 (±1.96)
	Mirrored OCT	26.77 (±1.61)	26.46 (±1.77)
Multimodal	PNN	23.14 (±1.66)	25.25 (±3.00)
	Mirrored OCT	**29.14 (±2.27)**	**29.15 (±2.17)**

Experiments: Diabetes Management

Dataset

58,200 patients with type 2 diabetes from the Boston Medical Center (1999 to 2014) – demographic data, medical information, etc.

Multiple Continuous Treatments

Continuous dosages of metformin, insulin, and oral blood glucose regulation agents

Goal

Minimize patient's Hemoglobin A1C

Model Class	Improvement (%)
Regress & Compare	2.90 (±0.46)
Causal Forest	1.60 (±0.47)
Optimal Policy Tree	2.55 (±0.52)
PNN	**3.15 (±0.51)**
Mirrored OCT	3.06 (±0.53)

Experiments: Groceries pricing

Dataset
97,295 rows of transaction-, product- and household-level data that include strawberries purchases

Single Continuous Treatment
Price of strawberries

Goal
Maximize revenue
(continuous outcome)

Model Class	Improvement (%)
Regress & Compare	94.17 (±6.25)
Causal Forest	98.68 (±2.67)
Optimal Policy Tree	106.58 (±2.38)
PNN	**110.88 (±1.18)**
Mirrored OCT	110.22 (±3.10)

We report improvement w.r.t. revenue: $p_r = \frac{1}{n}\sum_{i=1}^{n} \Gamma_{i,t_i} \cdot t_i$

Experiments: Spleen injuries treatment

Dataset

35,954 rows of patient data with spleen surgical operations, including demographic data and other medical data

Multiple Discrete Treatments

Splenectomy (spleen removal), angioembolization (minimally invasive), or observation (no treatment)

Goal

Minimize mortality

(binary outcome)

Model Class	Improvement (%)
Regress & Compare	8.46 (±2.06)
Causal Forest	2.43 (±4.57)
Optimal Policy Tree	12.98 (±1.23)
PNN	**13.52 (±1.74)**
Mirrored OCT	9.47 (±1.91)

Experiments: REBOA for blunt trauma patients

Dataset
9,998 patients with noncompressible torso hemorrhage, using demographic and medical features

Binary treatment
REBOA or not

Goal
Minimize mortality
(binary outcome)

Model Class	Improvement (%)
Regress & Compare	-19.69 (±16.04)
Causal Forest	-19.31 (±5.16)
Optimal Policy Tree	17.17 (±3.68)
PNN	17.87 (±3.88)
Mirrored OCT	**18.09 (±3.18)**

Example of a Mirrored OCT

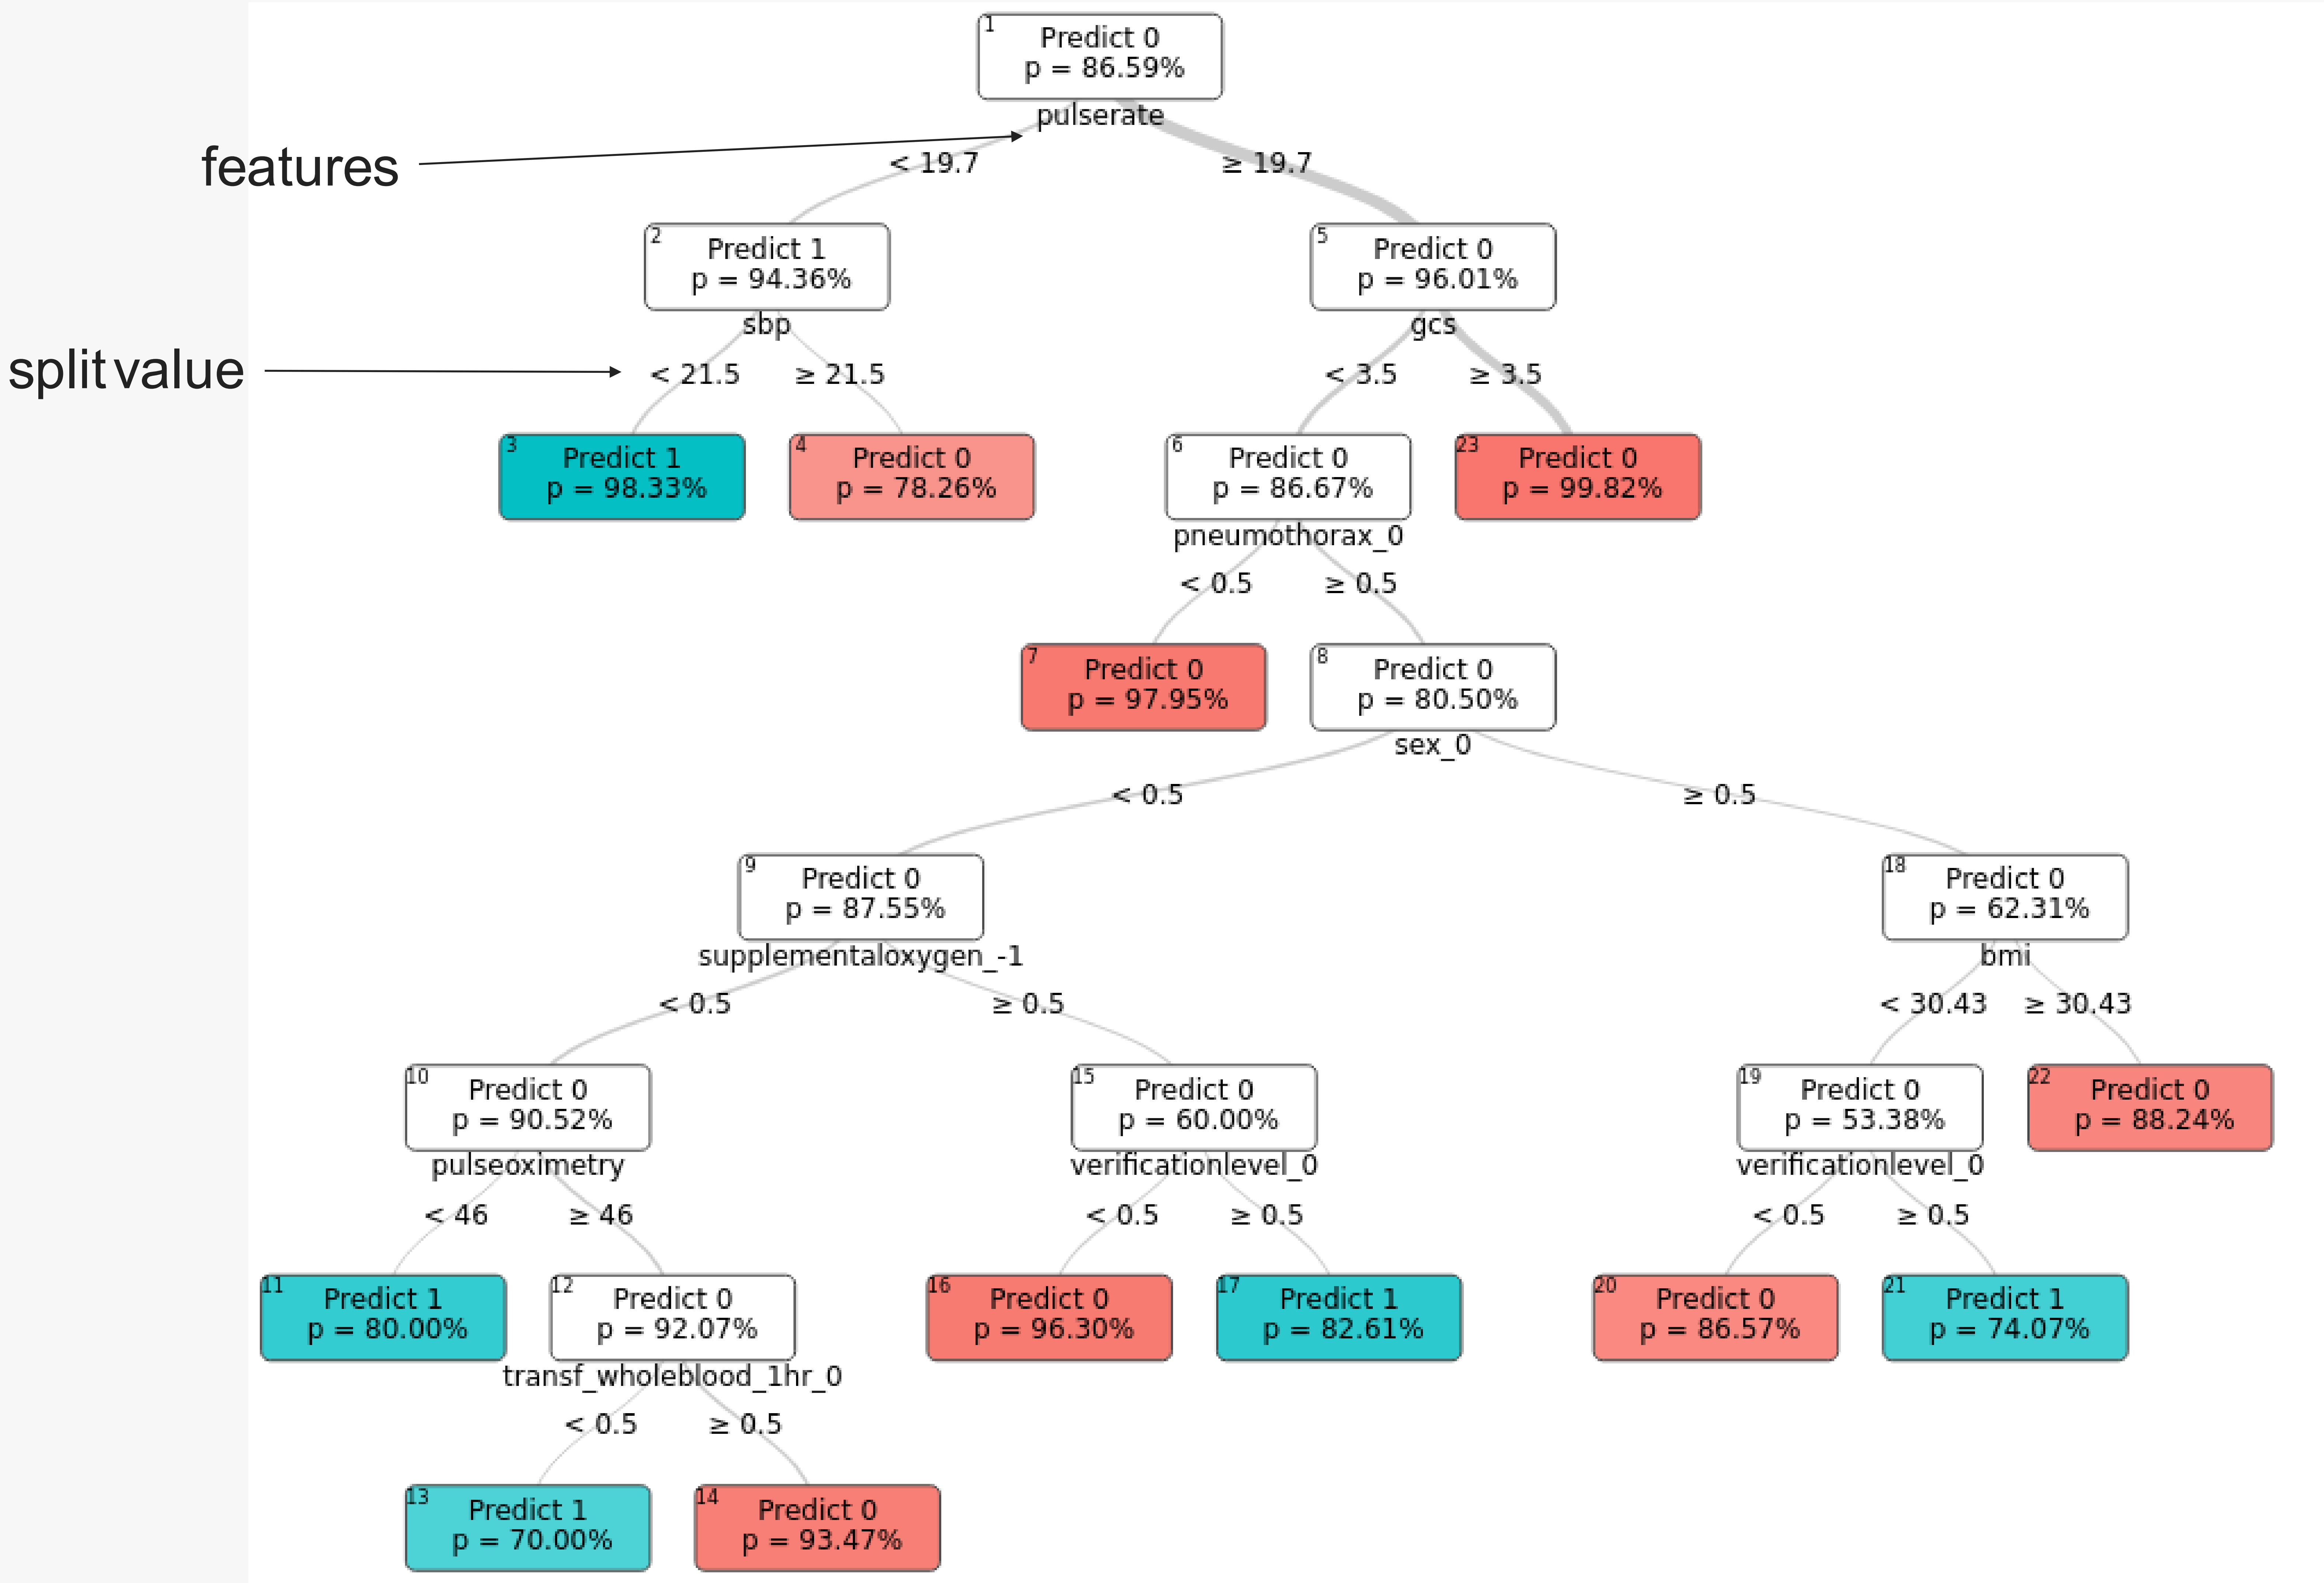

What did we learn today?

- The PNN framework can be viewed as an extension of Optimal Policy Trees for multimodal data and consists of 4 different steps:
  - **Embeddings calculation**
  - **Counterfactual estimation**
  - **PNN training**
  - **Mirrored OCT training**
- With its classification-like feedforward neural network architecture, the PNN framework flexibly handles multimodal data, by easily enabling the incorporation of multiple data sources.
- PNNs and Mirrored OCTs perform better or comparably with existing prescriptive approaches.
- Mirrored OCTs can help recover interpretability without significantly sacrificing PNNs' performance.