




Causal structure-based root cause analysis of outliers

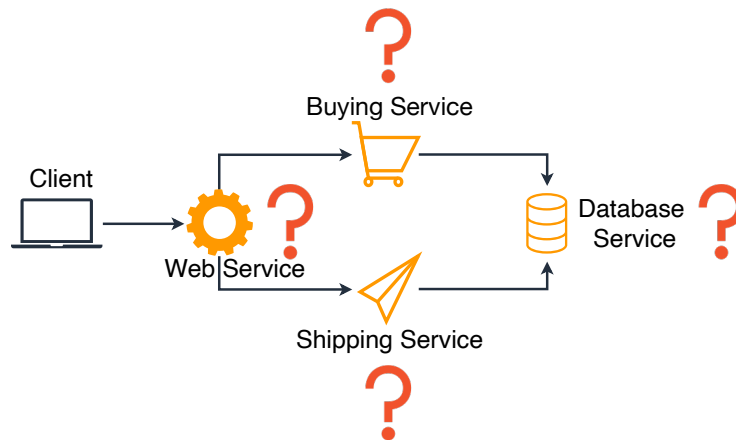

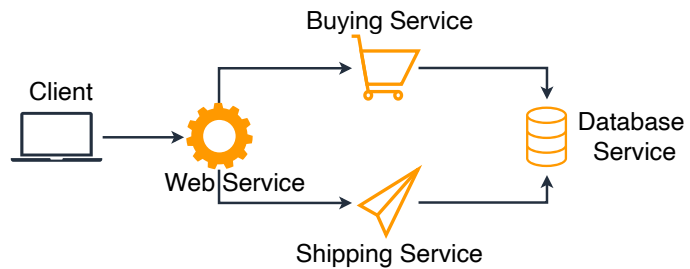
ICML 2022

Kailash Budhathoki, Lenon Minorics, Patrick Blöbaum, Dominik Janzing

July 17, 2021



Motivation



Scenario Event x_n of variable X_n flagged as an outlier!
We jointly observe (x_1, \dots, x_n) of (X_1, \dots, X_n) .

Goal Identify the root causes of outlier x_n
amongst X_1, \dots, X_n



But many outlier scoring algorithms!

IT outlier score

feature map (can be an existing algorithm)

$$S_X^\tau(x) := -\log P_X\{\tau(X) \geq \tau(x)\}$$

↑
event

↑
random variable

IT outlier score: Example

$$\mathbf{z}\text{-score: } \mathbf{z}(x) := \frac{|x - \mu_X|}{\sigma_X}$$

$$S_X^{\mathbf{z}}(x) := -\log P_X\{\mathbf{z}(X) \geq \mathbf{z}(x)\}$$

Formal framework to study
outliers ✓

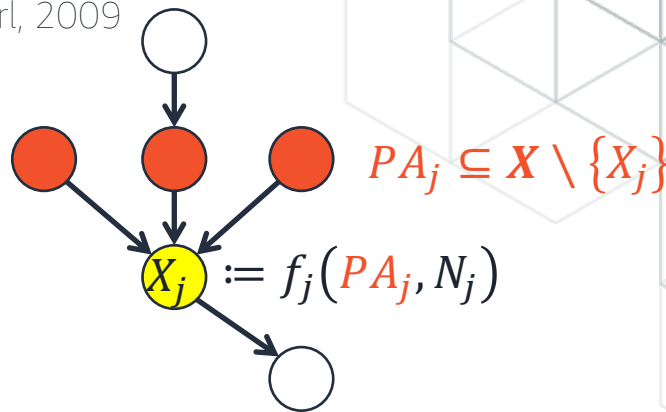
How can we identify the
root causes of outliers?

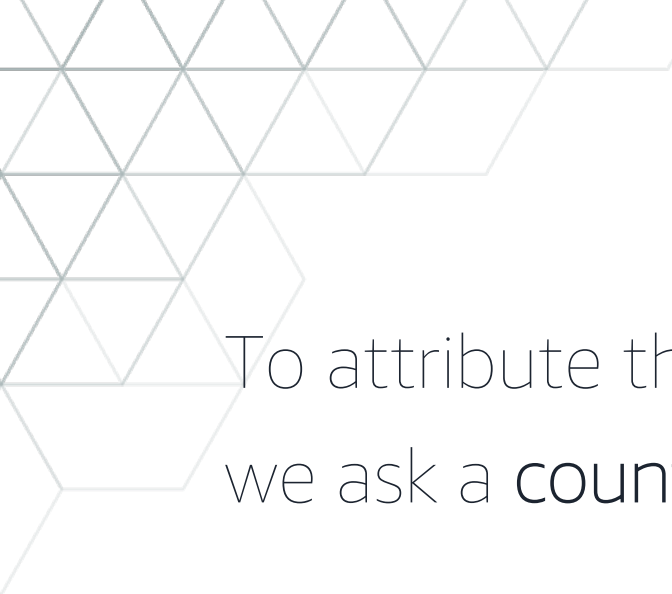


Structural Causal Model (SCM) Pearl, 2009

a quadruple $\langle \mathbf{X}, \mathbf{N}, \mathcal{F}, P_{\mathbf{N}} \rangle$, where

- \mathbf{X} : observed variables (X_1, \dots, X_n)
- \mathbf{N} : unobserved "noise" variables (N_1, \dots, N_n)
- \mathcal{F} : set of "modular" functions or causal mechanisms $\{f_j\}_{j=1}^n$, describing how each X_j is generated, $\{X_j := f_j(PA_j, N_j)\}_{j=1}^n$
- $P_{\mathbf{N}}$: joint distribution over \mathbf{N}



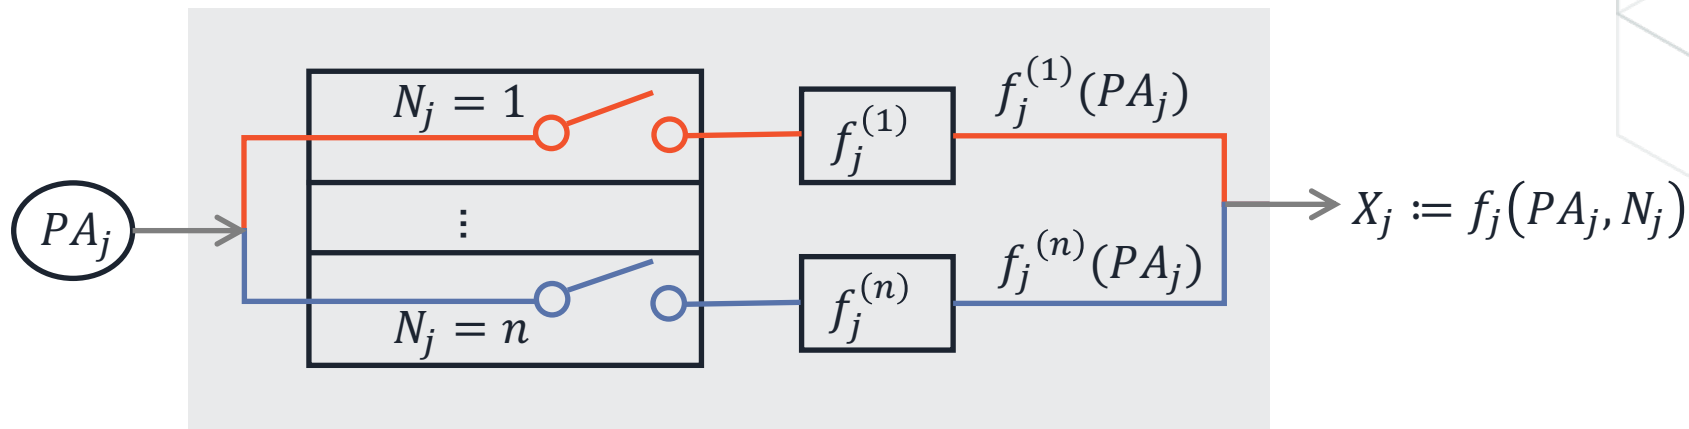


To attribute the outlier x_n to a variable X_j ,
we ask a **counterfactual** question:

Would x_n not have been an outlier had the
causal mechanism at X_j been "normal"?

Canonical Structural Causal Model

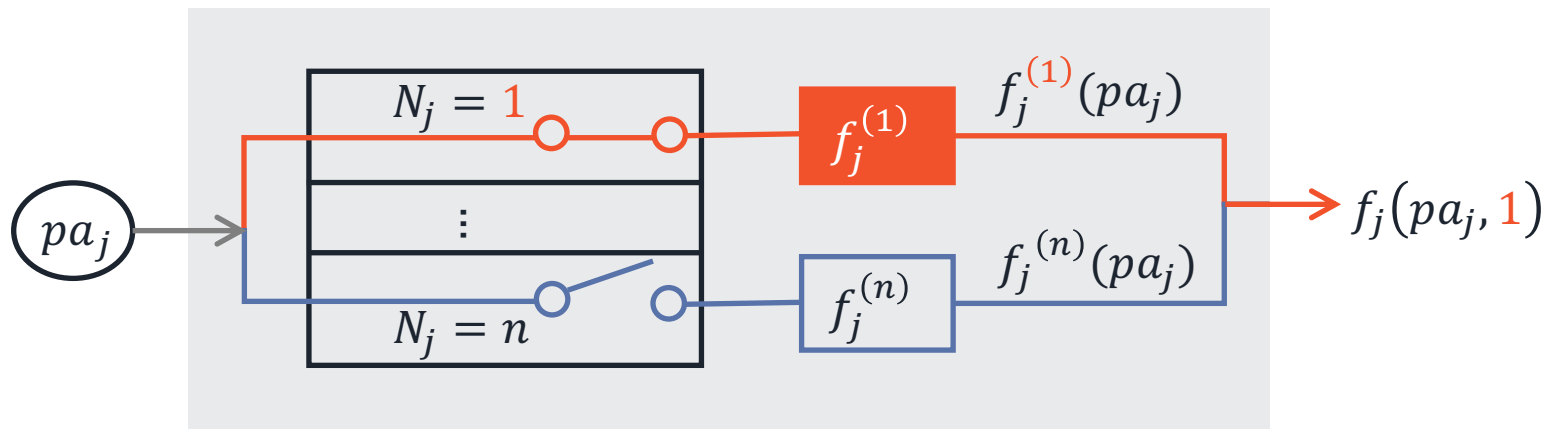
$$X_j := f_j(PA_j, N_j) \text{ for } N_j := n_j$$



P_N defined on the set of mechanisms $\{f_j^{(1)}, \dots, f_j^{(n)}\}$

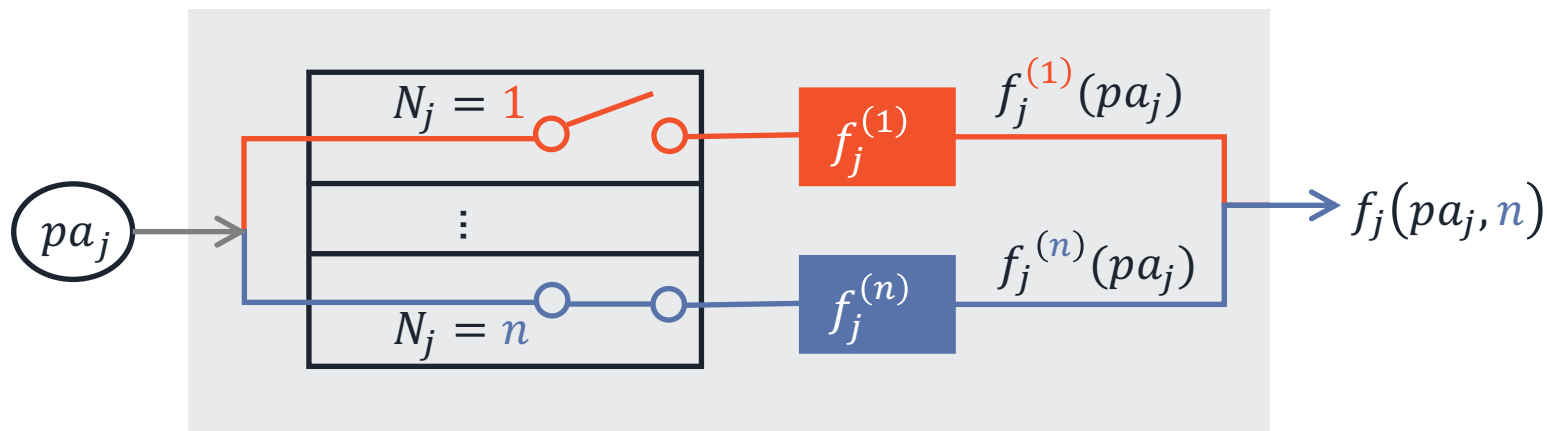
Canonical Structural Causal Model

We observed $x_j := f_j(pa_j, \mathbf{1})$ with $N_j := \mathbf{1}$



Canonical Structural Causal Model

$$\tilde{x}_j := f_j(pa_j, n) \text{ for } N_j := n$$

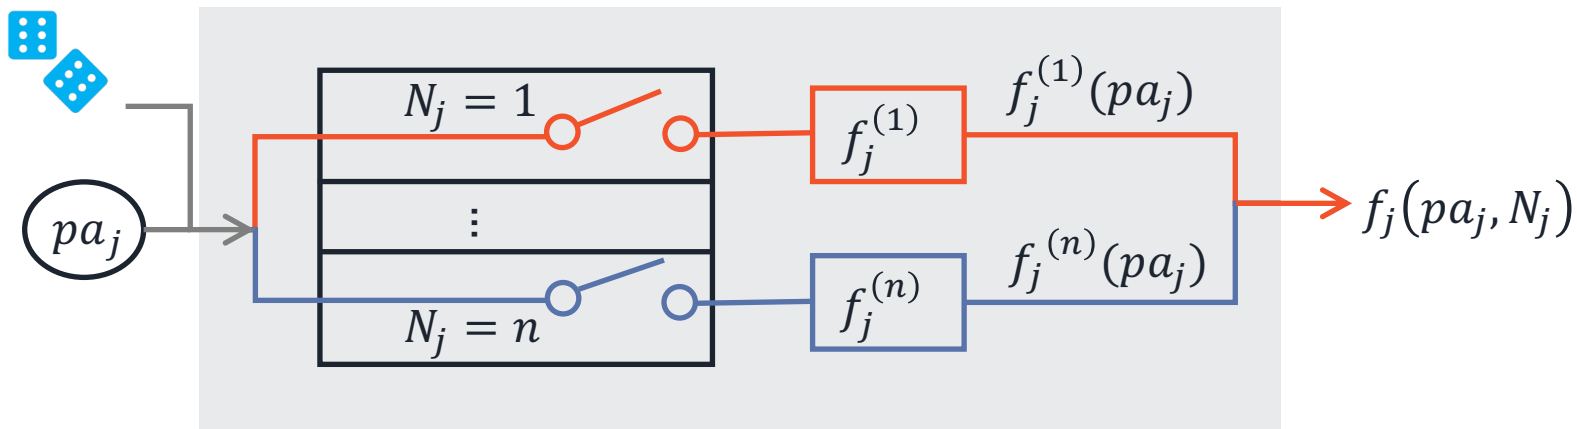


factual mechanism $f_j^{(1)}$

counterfactual mechanism $f_j^{(n)}$

Normal causal mechanisms

Randomize N_j






To attribute the outlier x_n to a variable X_j ,
we ask a **counterfactual** question:

Would x_n not have been an outlier had the
causal mechanism of X_j been "normal"?

Causal Attribution

impact of replacing
causal mechanism at X_j on
the log tail probability
defines the contribution of X_j


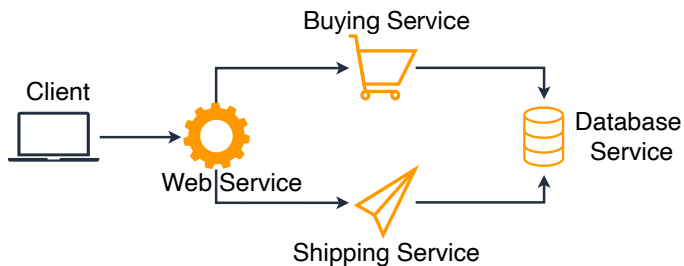


$$C(j \mid 1, \dots, j-1) := \log \frac{P^{RAND(N_1, \dots, N_{j-1})} \{ \tau(X_n) \geq \tau(x_n) \}}{P^{RAND(N_1, \dots, N_j)} \{ \tau(X_n) \geq \tau(x_n) \}}$$



Average contributions over all orderings
to avoid dependence on ordering
(Shapley value approach)

Shapley contributions sum up
to the IT outlier score $S_{X_n}^\tau(x_n)$

Formal framework to identify
root causes of outliers ✓



Structural Causal Model



[Implementation](#) available in Python DoWhy!

<https://github.com/py-why/dowhy>