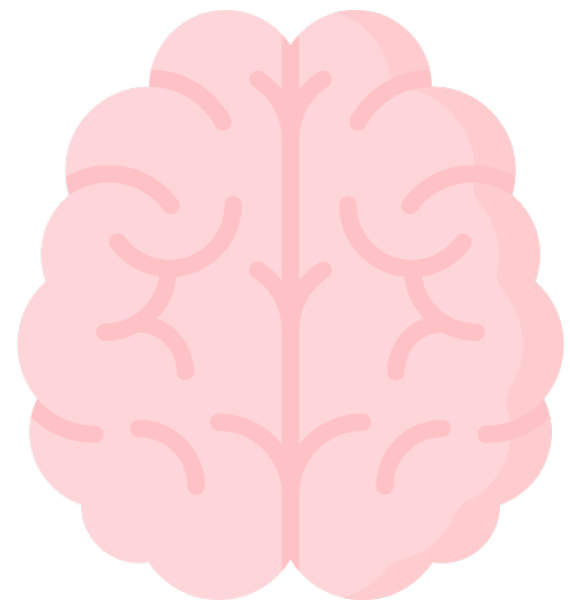
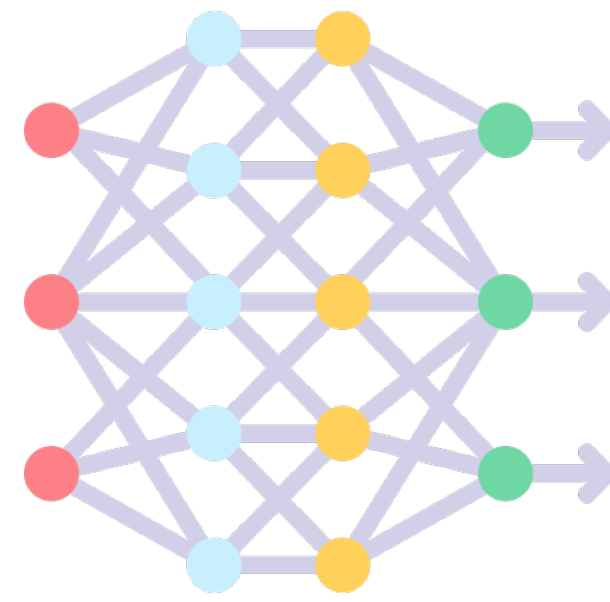


# Outline



Cognitive  
science



Machine  
learning

AaI

Large language  
models

# Causal machine learning

The amount of work at the interface of causality and machine learning, often referred to as **causal machine learning**, has been increasing very rapidly.

Kaddour et al. *"Causal machine learning: A survey and open problems."* arXiv preprint, 2022.

Peters et al. *"Elements of causal inference: foundations and learning algorithms."* The MIT Press, 2017.

# Causal machine learning

The amount of work at the interface of causality and machine learning, often referred to as **causal machine learning**, has been increasing very rapidly.

Causal machine learning operationalizes causal (counterfactual) reasoning about  
the **outputs** of machine learning models,  
the **data** used by these models, and  
the **users** of these models  
using the theoretical framework of **structural causal models (SCMs)**.

Kaddour et al. *"Causal machine learning: A survey and open problems."* arXiv preprint, 2022.

Peters et al. *"Elements of causal inference: foundations and learning algorithms."* The MIT Press, 2017.

# Structural Causal Models (SCMs)

Given a set of random variables  $\mathbf{X} = \{X_1, \dots, X_n\}$ , a SCM defines a **complete data-generating process** via a collection of assignments

$$X_i := f_i(\mathbf{PA}_i, U_i),$$

where  $\mathbf{PA}_i \subseteq \mathbf{X} \setminus X_i$  are the direct causes of  $X_i$ ,

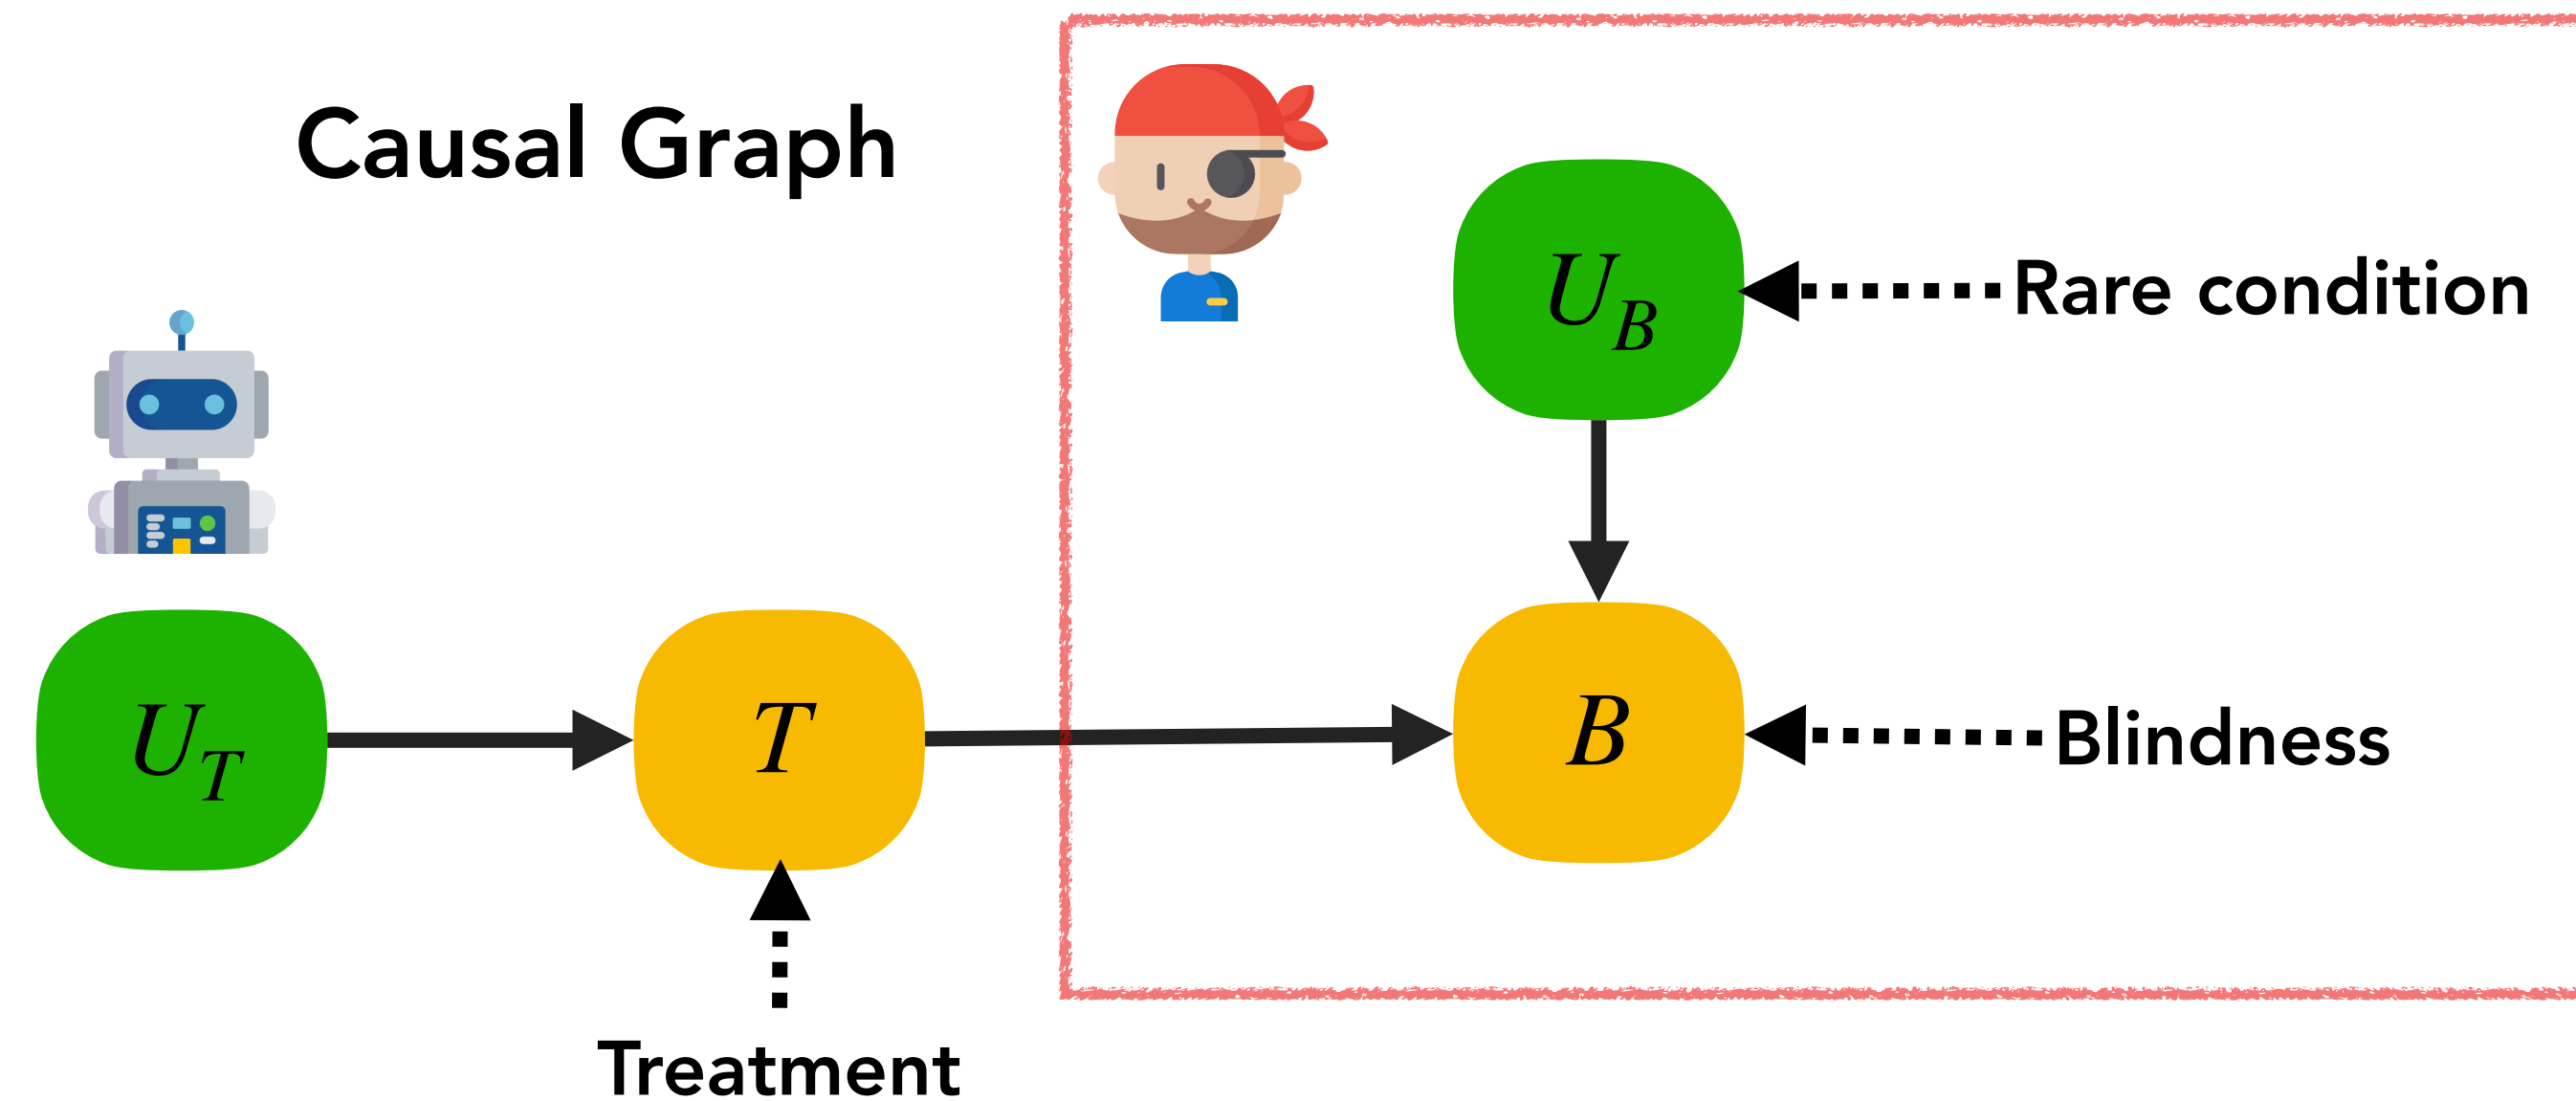
$\mathbf{U} = \{U_1, \dots, U_n\}$  are jointly independent noise variables

$\mathbf{F} = \{f_1, \dots, f_n\}$  are deterministic causal mechanisms, and

$P(\mathbf{U})$  denotes the (prior) distribution of the noise variables.

# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) Interventional and (3) Counterfactual Queries



**Structural Causal Model  $\mathcal{M}$**

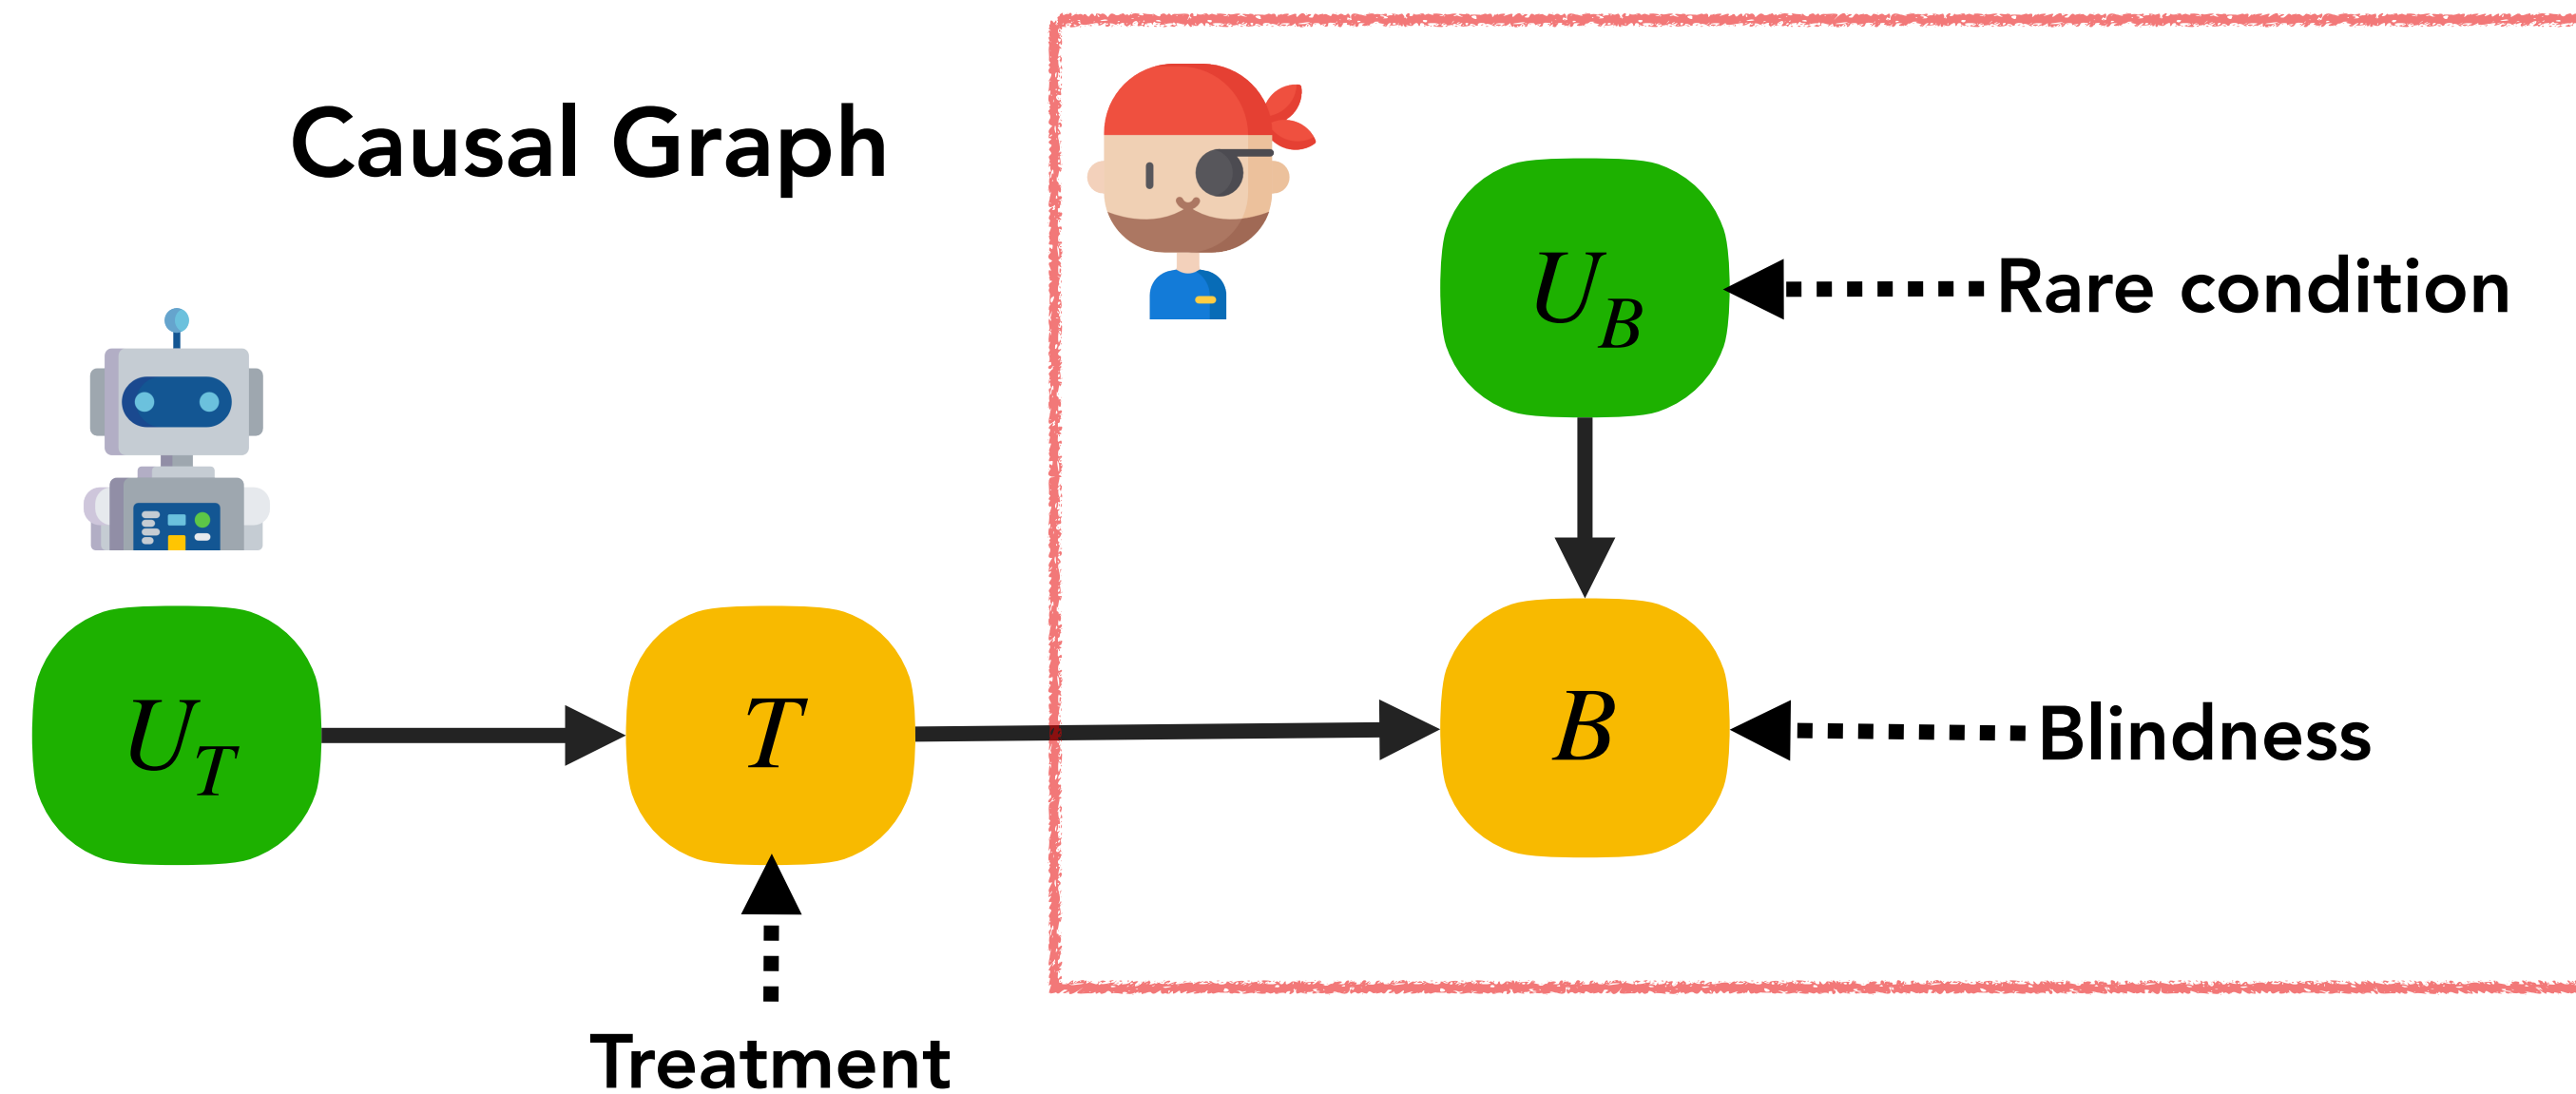
$$T := U_T$$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B)$$

$$U_B \sim \text{Ber}(0.01), \quad U_T \sim \text{Ber}(0.5)$$

# What kind of (causal) questions can we answer with SCMs?

(1) **Observational**, (2) Interventional and (3) Counterfactual Queries



**Structural Causal Model  $\mathcal{M}$**

$$T := U_T$$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B)$$

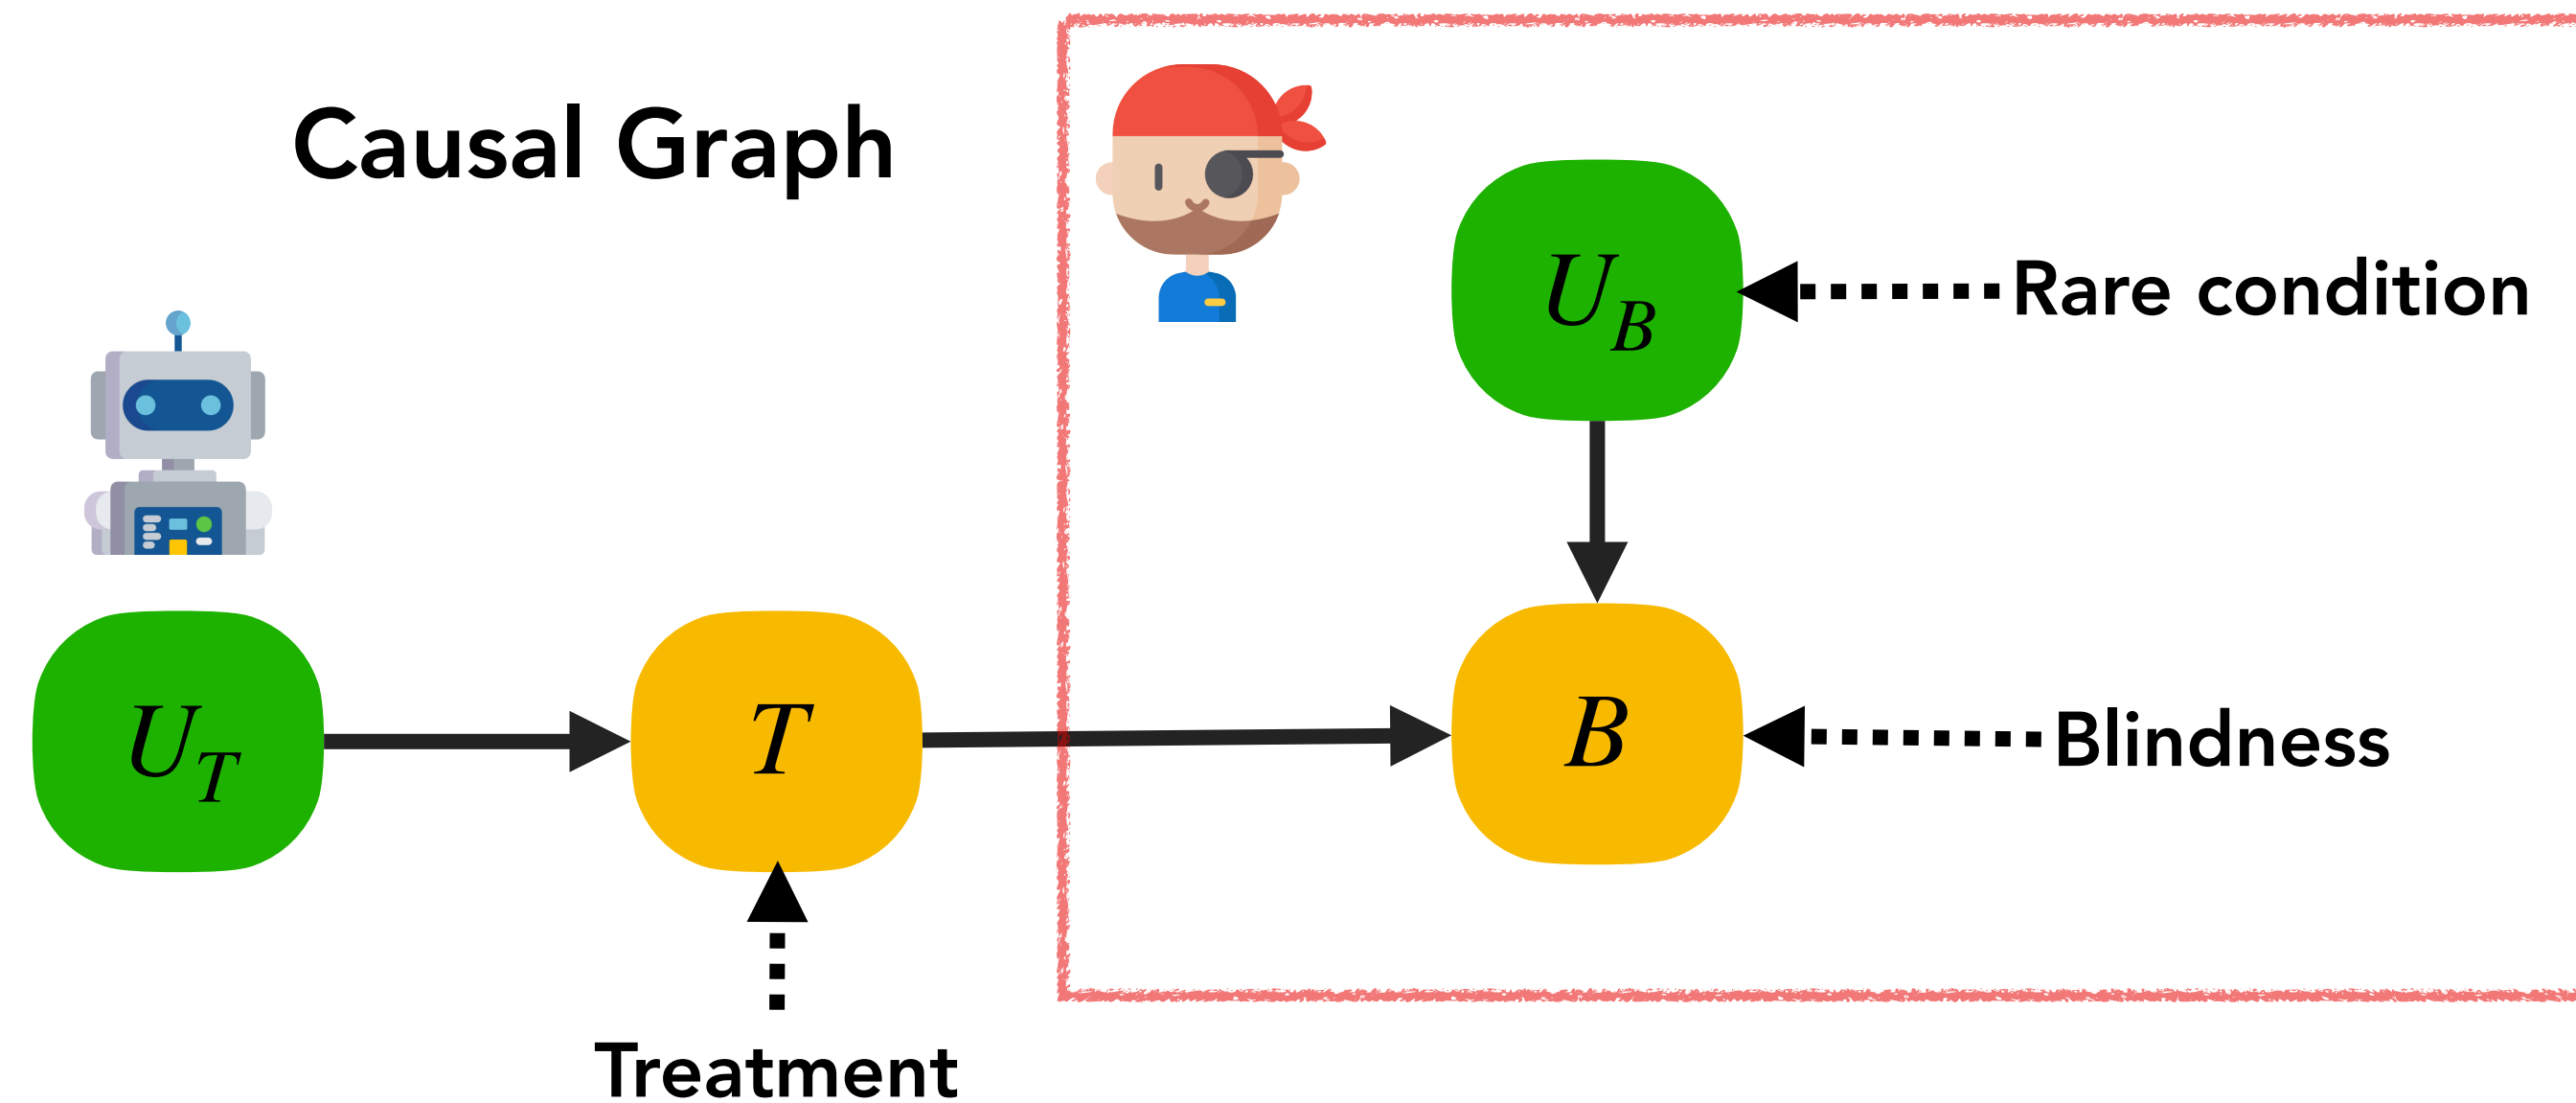
$$U_B \sim \text{Ber}(0.01), \quad U_T \sim \text{Ber}(0.5)$$

**Observational question**

What will happen to the patient?

# What kind of (causal) questions can we answer with SCMs?

(1) **Observational**, (2) Interventional and (3) Counterfactual Queries



**Structural Causal Model  $\mathcal{M}$**

$$T := U_T$$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B) \xrightarrow{\text{"observe"}} \quad$$

$$U_B \sim \text{Ber}(0.01), \quad U_T \sim \text{Ber}(0.5)$$

**Observational question**

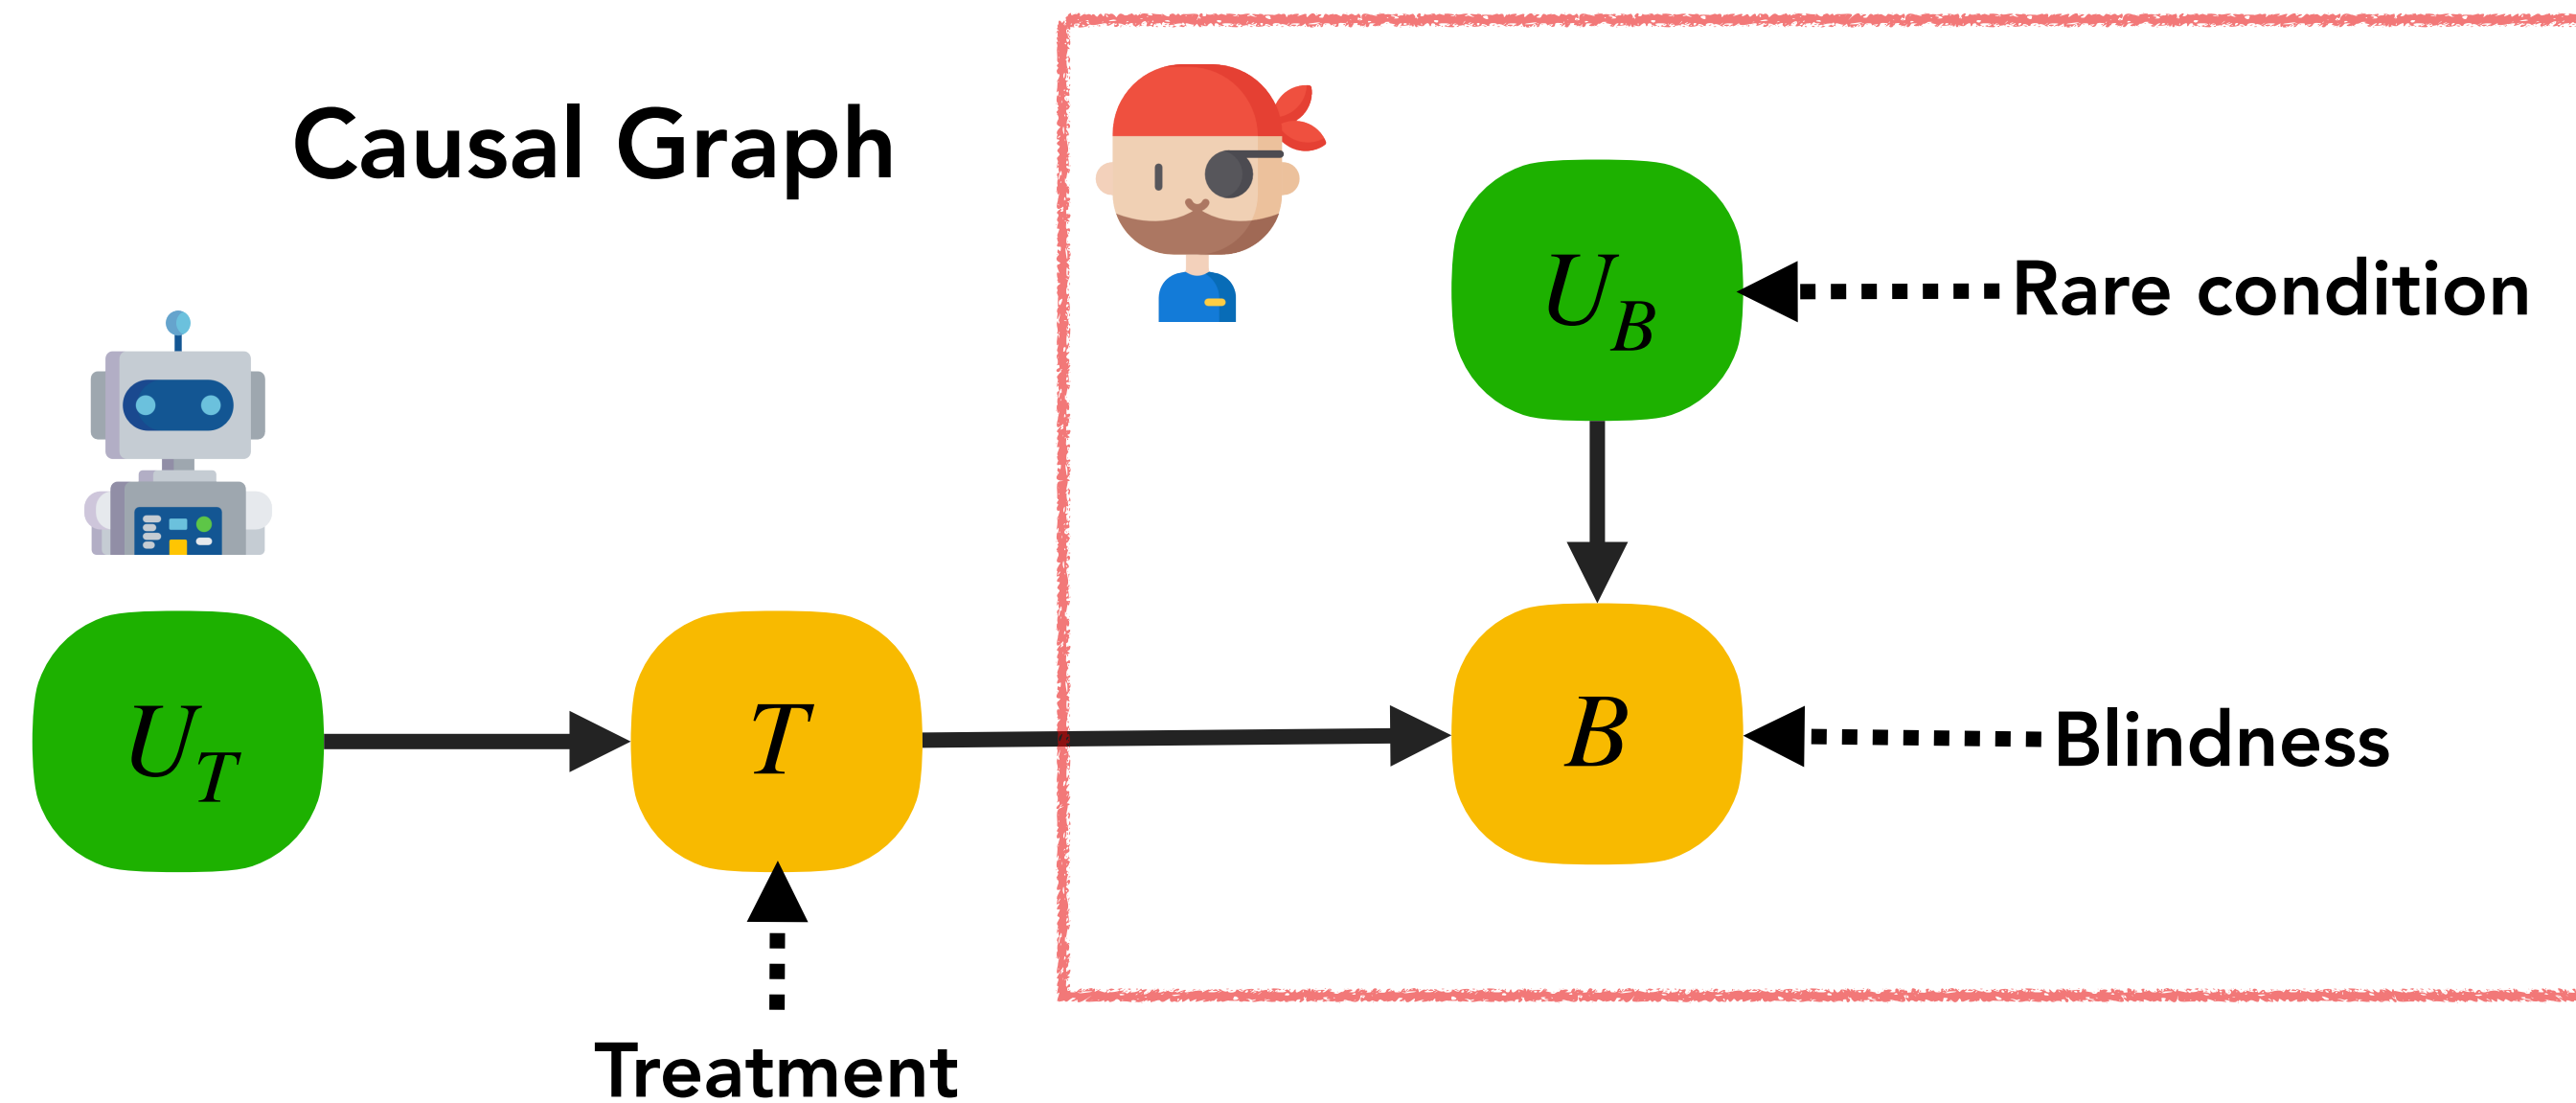
What will happen to the patient?

The patient will get blind ( $B = 1$ ) with prob. 0.5



# What kind of (causal) questions can we answer with SCMs?

(1) **Observational**, (2) Interventional and (3) Counterfactual Queries



**Structural Causal Model  $\mathcal{M}$**

$$T := U_T$$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B) \xrightarrow{\text{"observe"}} \quad$$

$$U_B \sim \text{Ber}(0.01), \quad U_T \sim \text{Ber}(0.5)$$

**Observational question**

What will happen to the patient?

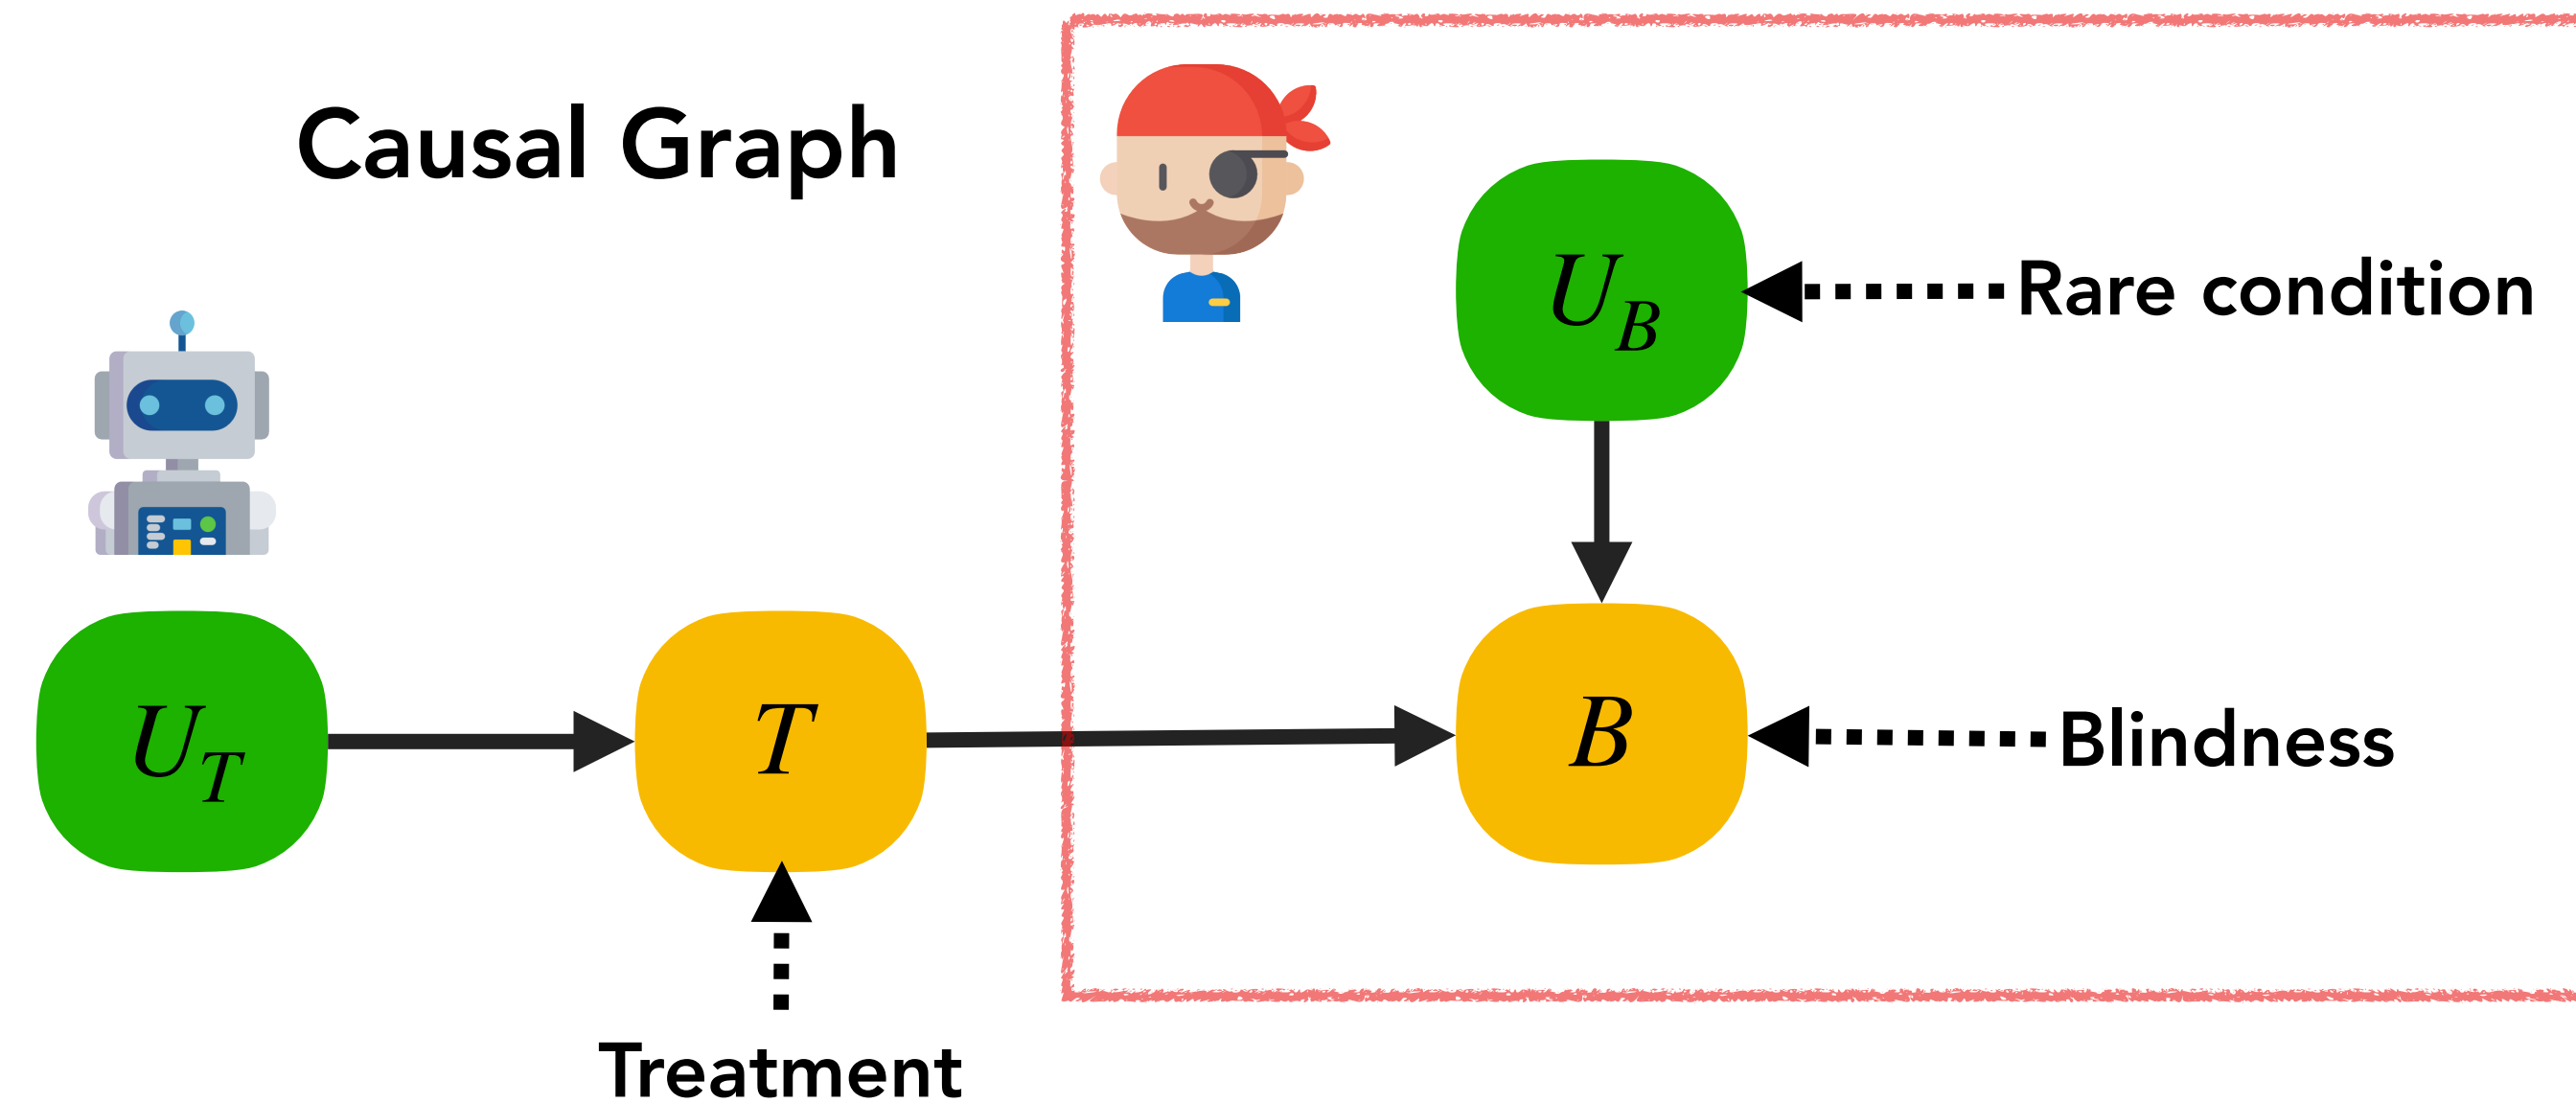
The patient will get blind ( $B = 1$ ) with prob. 0.5

Formally,  $P^{\mathcal{M}}(B = 1) = 0.5$



# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) **Interventional** and (3) Counterfactual Queries



**Structural Causal Model  $\mathcal{M}$**

$$T := U_T$$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B)$$

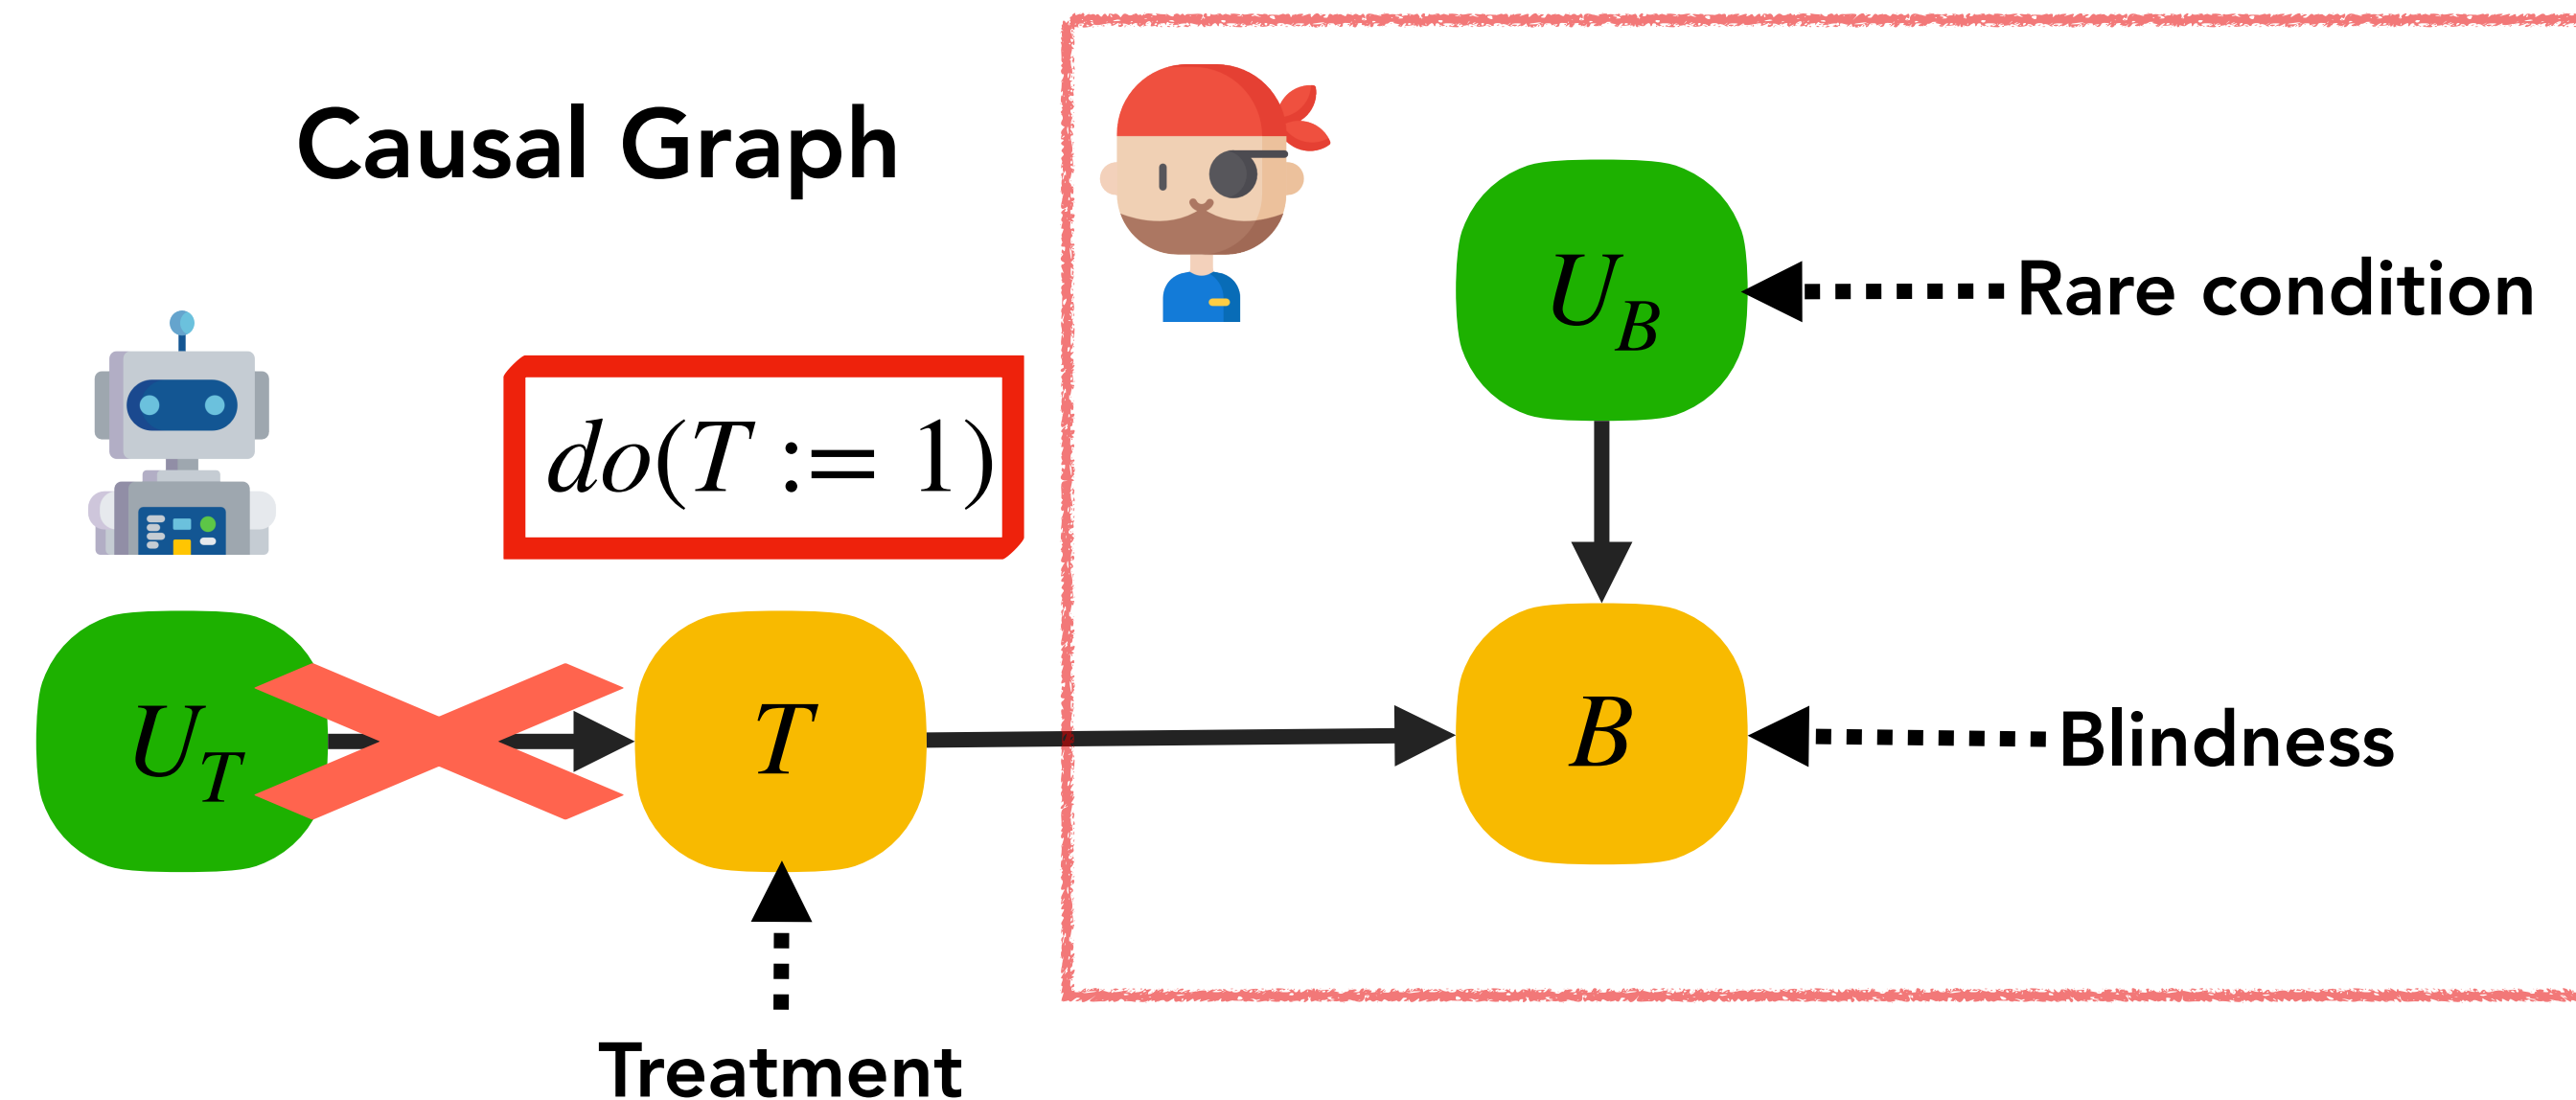
$$U_B \sim \text{Ber}(0.01), \quad U_T \sim \text{Ber}(0.5)$$

**Interventional question**

What will happen to the patient if a doctor breaks the robot and always administers the treatment?

# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) **Interventional** and (3) Counterfactual Queries



## Structural Causal Model $\mathcal{M}$

~~$T := U_T$~~   $T := 1$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B)$$

$U_B \sim \text{Ber}(0.01),$   ~~$U_T \sim \text{Ber}(0.5)$~~

"do"

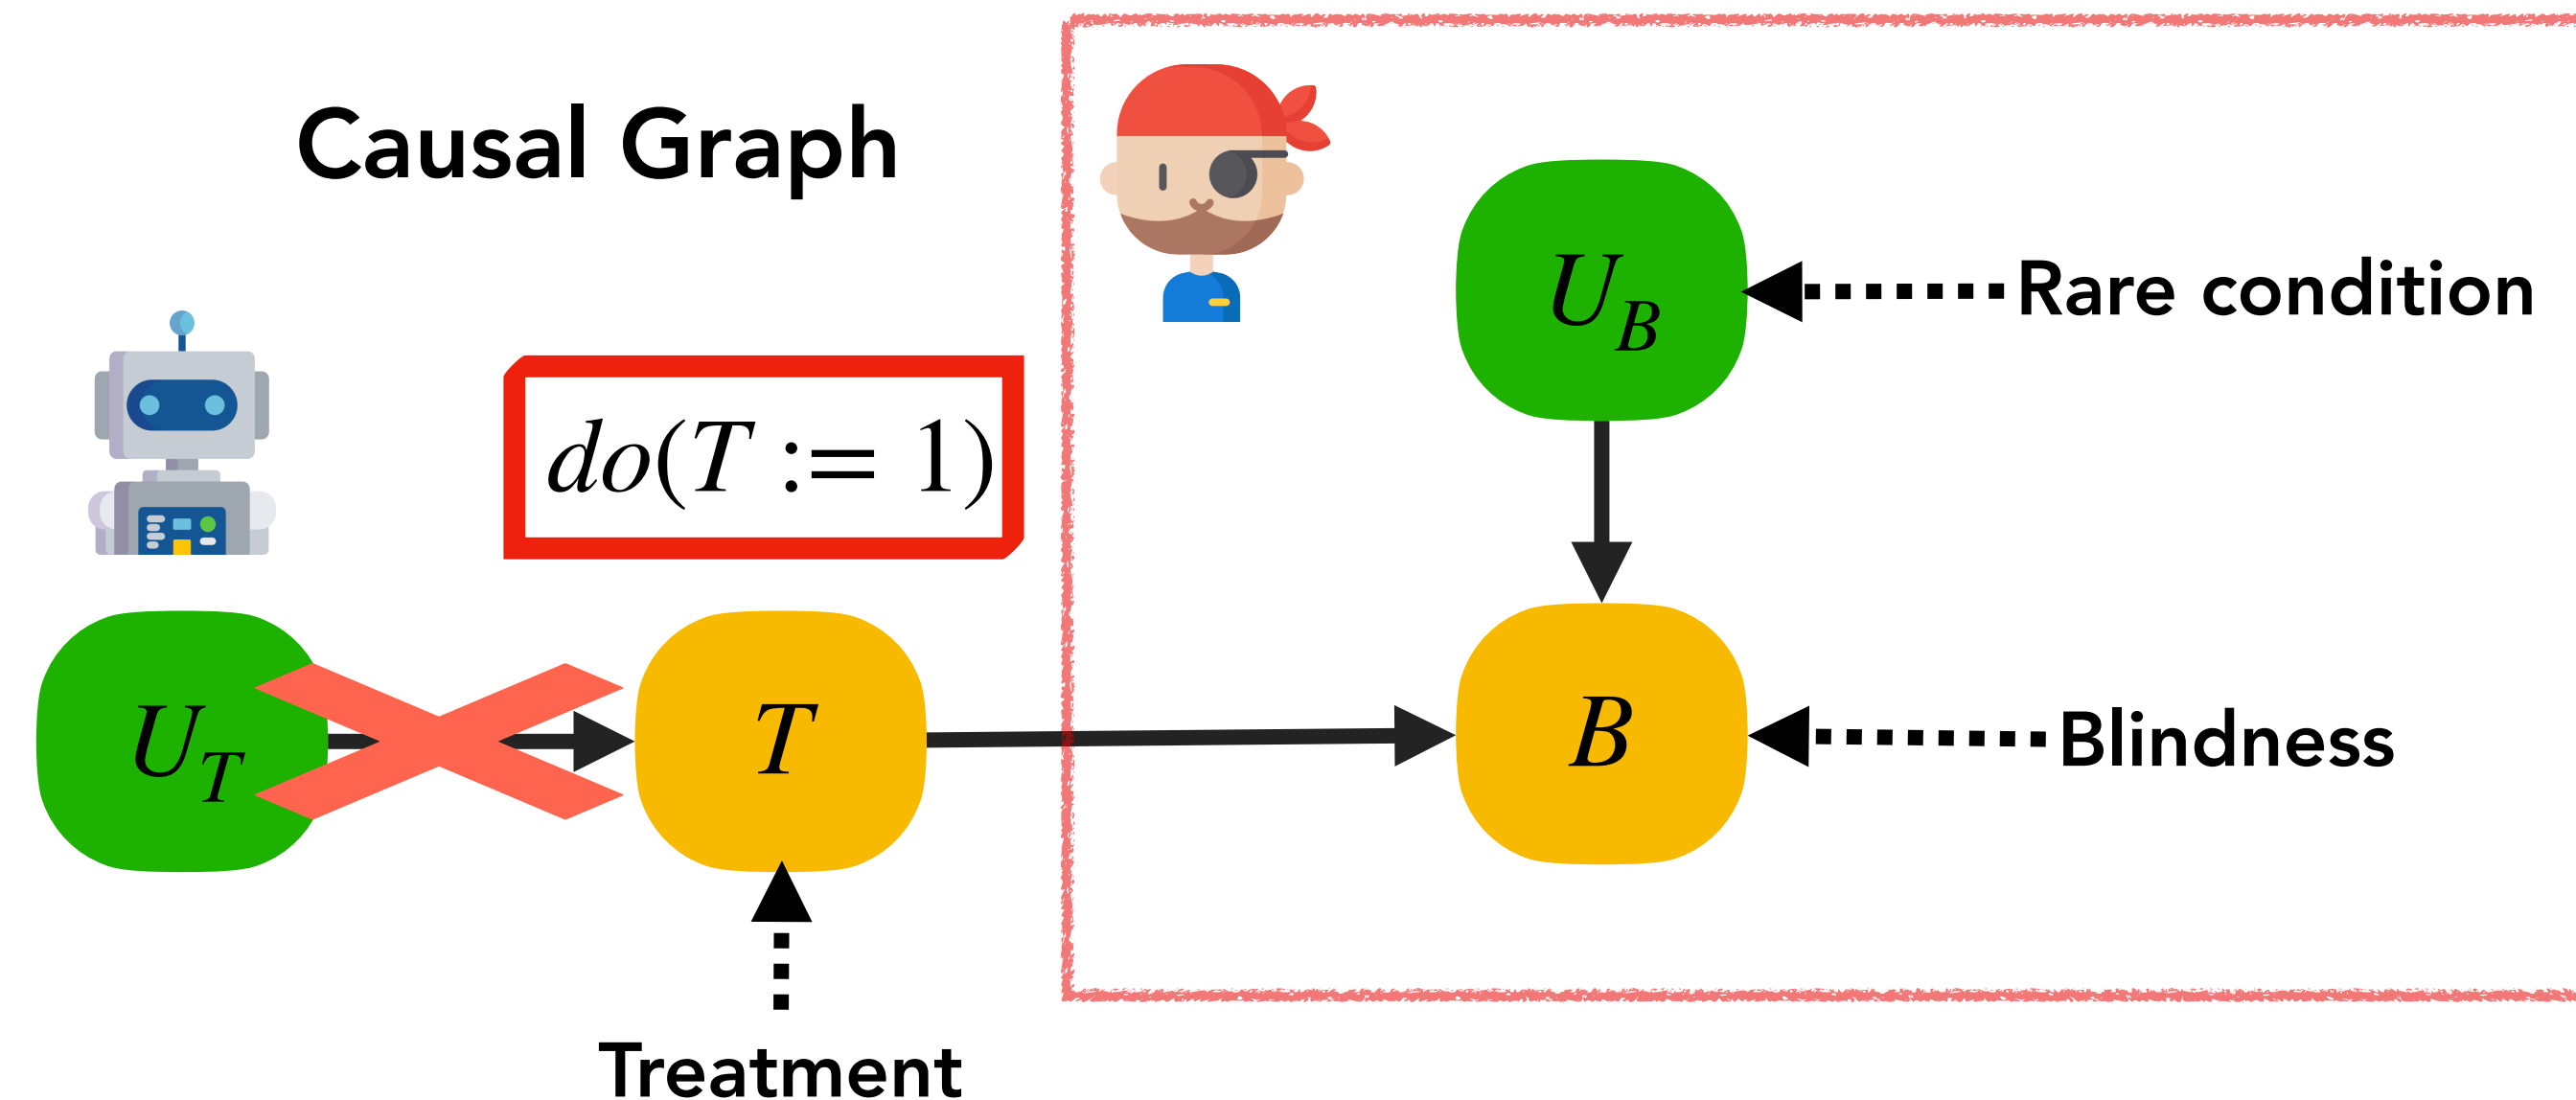
## Interventional question

What will happen to the patient if a doctor breaks the robot and always administers the treatment?

The patient will get blind ( $B = 1$ ) with prob. 0.01

# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) **Interventional** and (3) Counterfactual Queries



## Structural Causal Model $\mathcal{M}$

~~$T := U_T$~~   $T := 1$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B)$$

$U_B \sim \text{Ber}(0.01)$ ,  ~~$U_T \sim \text{Ber}(0.5)$~~

"do"

## Interventional question

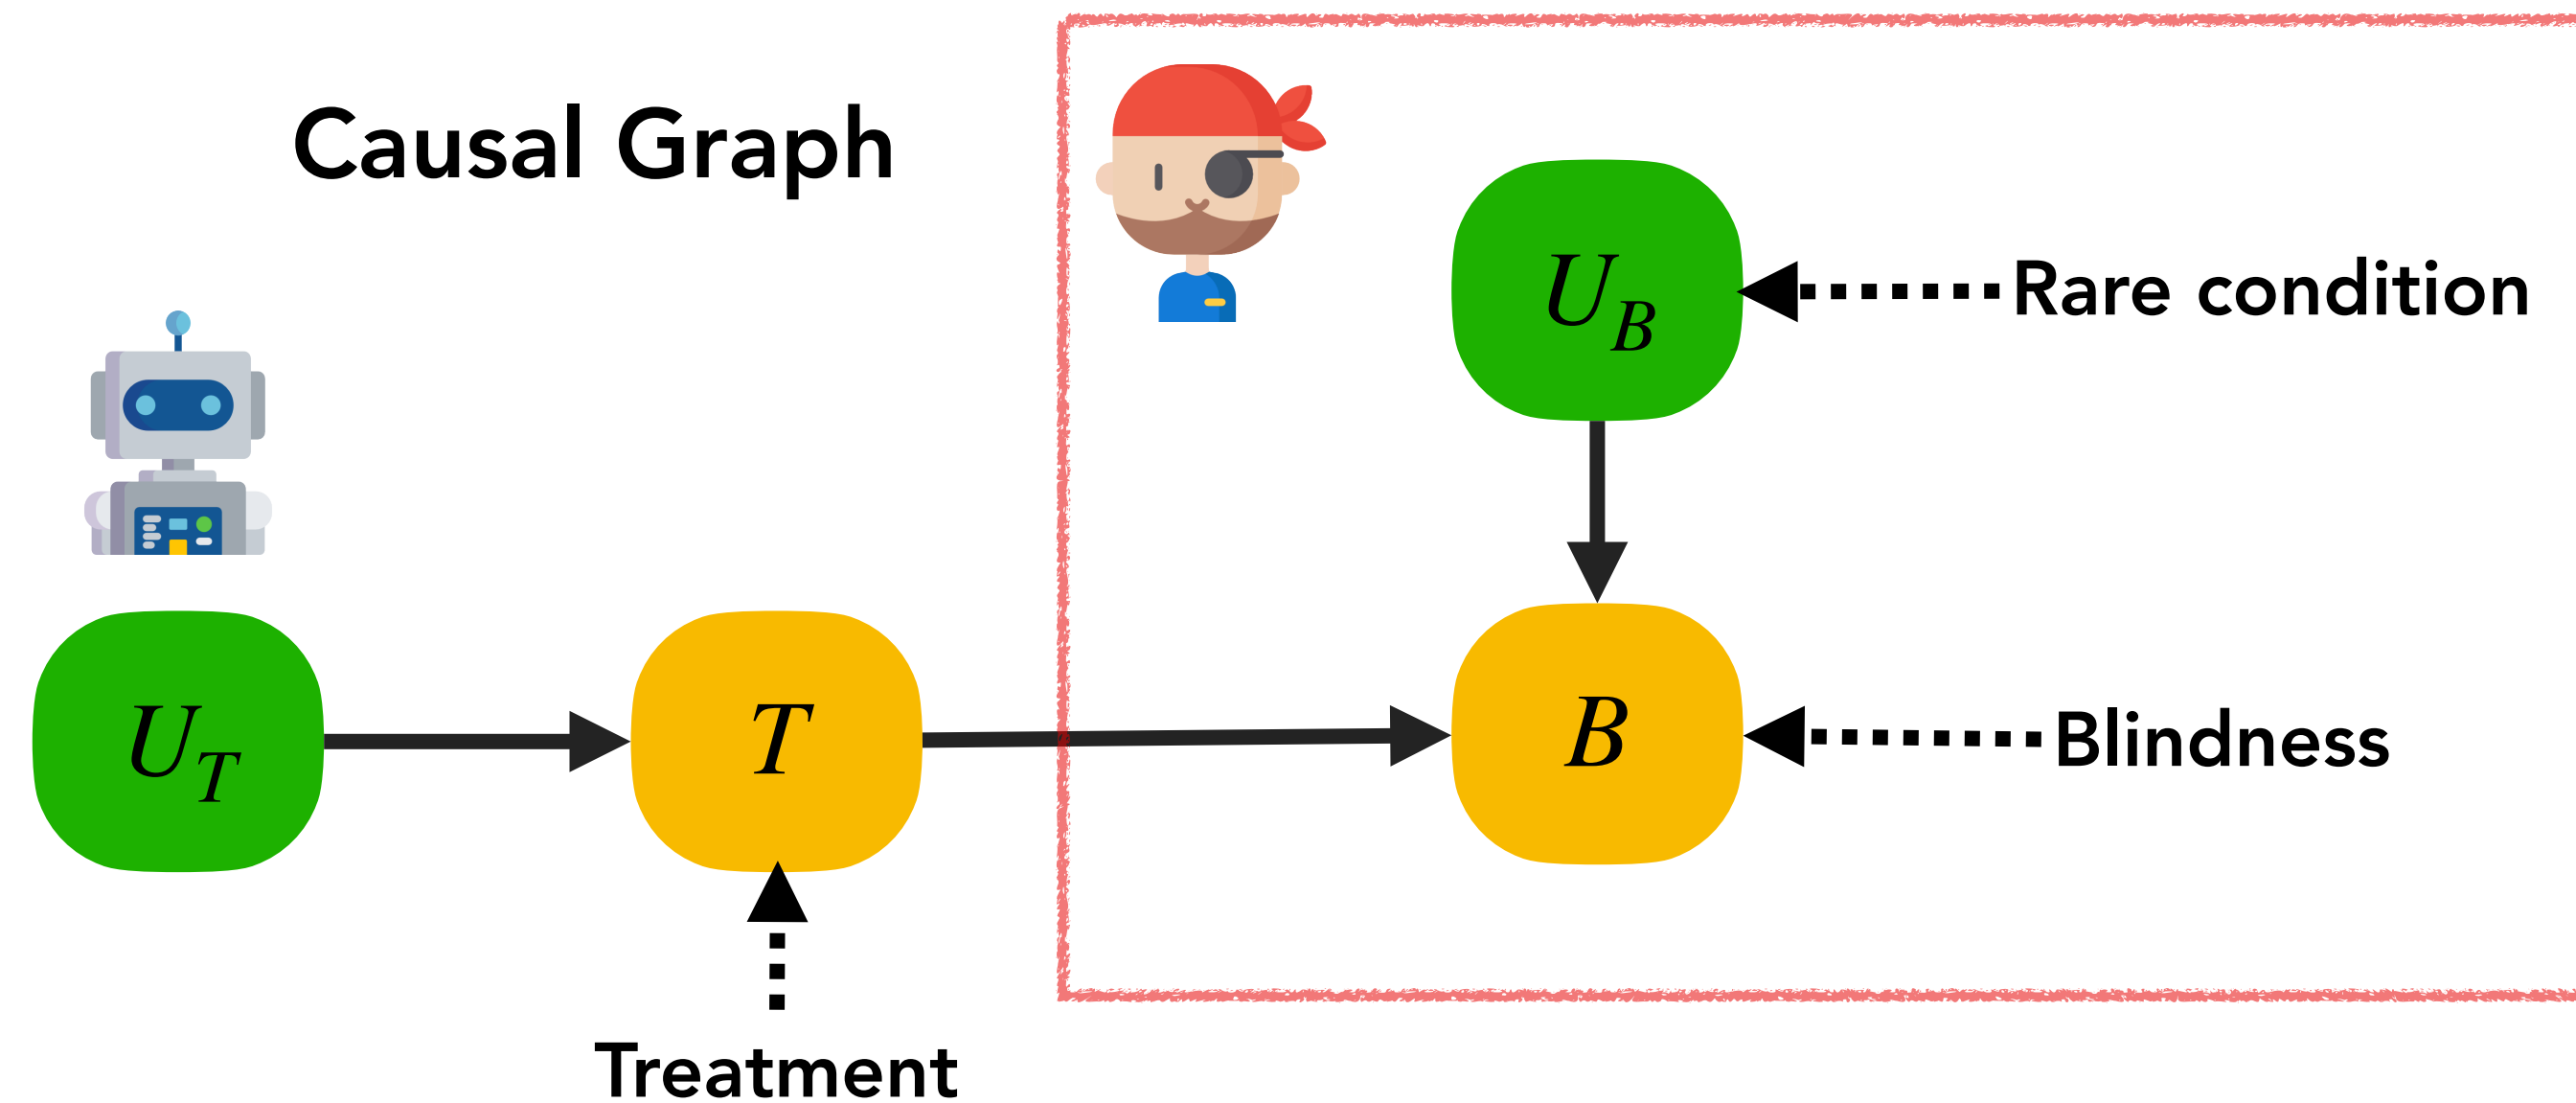
What will happen to the patient if a doctor breaks the robot and always administers the treatment?

The patient will get blind ( $B = 1$ ) with prob. 0.01

Formally,  $P^{\mathcal{M}; do(T=1)}(B = 1) = 0.01$

# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) Interventional and (3) **Counterfactual** Queries



**Structural Causal Model  $\mathcal{M}$**

$$T := U_T$$

$$B := T \cdot U_B + (1 - T) \cdot (1 - U_B)$$

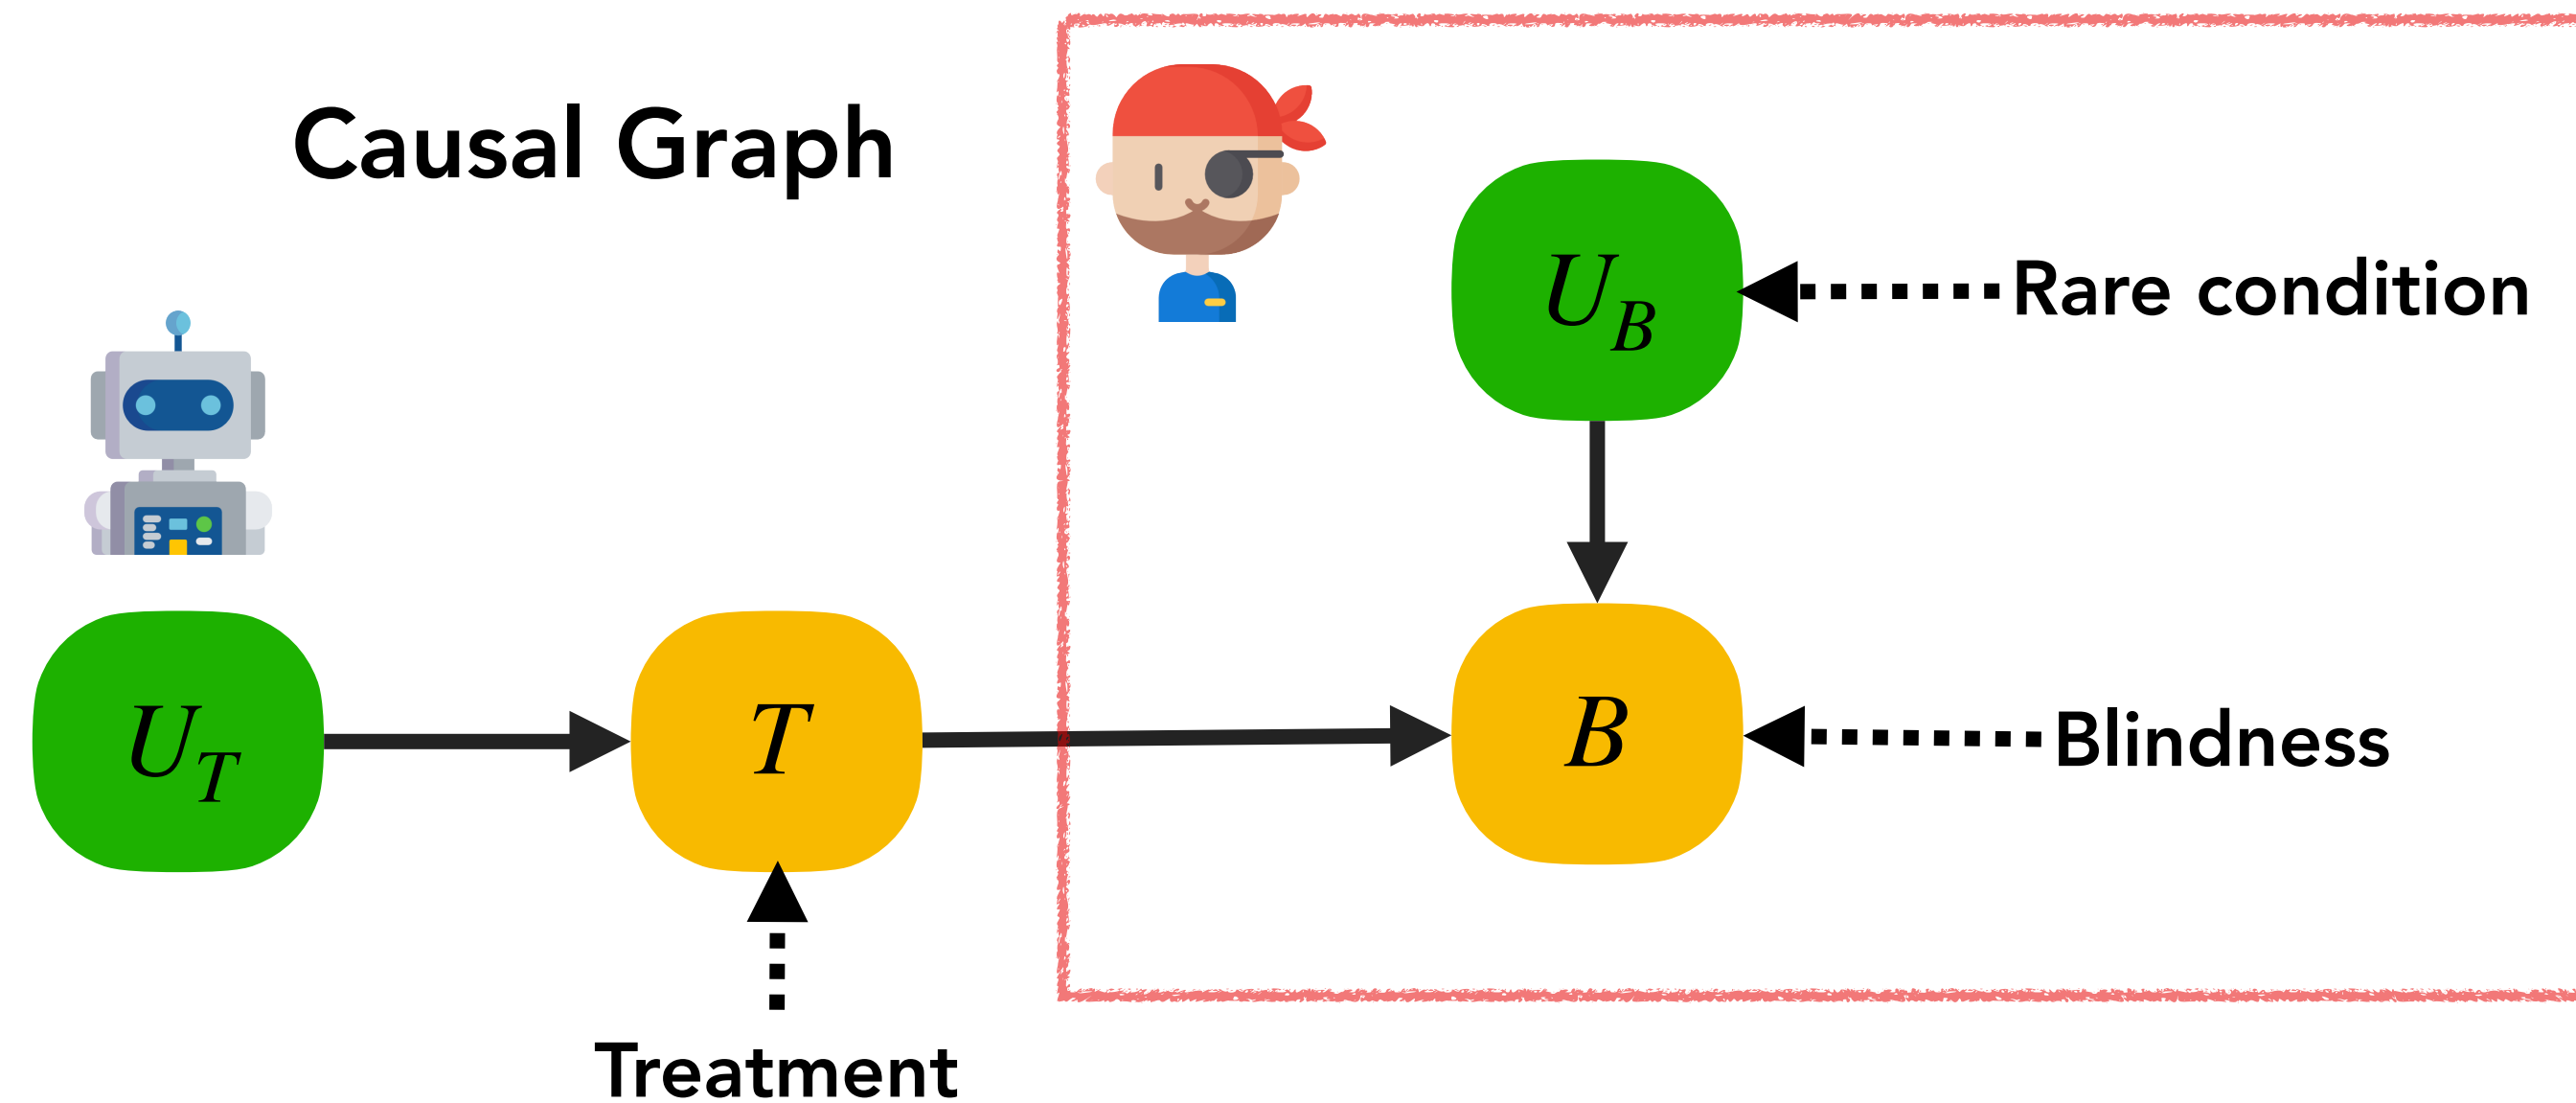
$$U_B \sim \text{Ber}(0.01), \quad U_T \sim \text{Ber}(0.5)$$

**Counterfactual question**

The treatment was administered and the patient got blind. What would have happened if the treatment had not been administered?

# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) Interventional and (3) **Counterfactual** Queries



**Modified Structural Causal Model**  $\mathcal{M}_{T=1, B=1}$

$T := 1$

$B := T$

$U_B = 1$  with prob. 1  $\leftarrow$  Posterior distribution of the noise

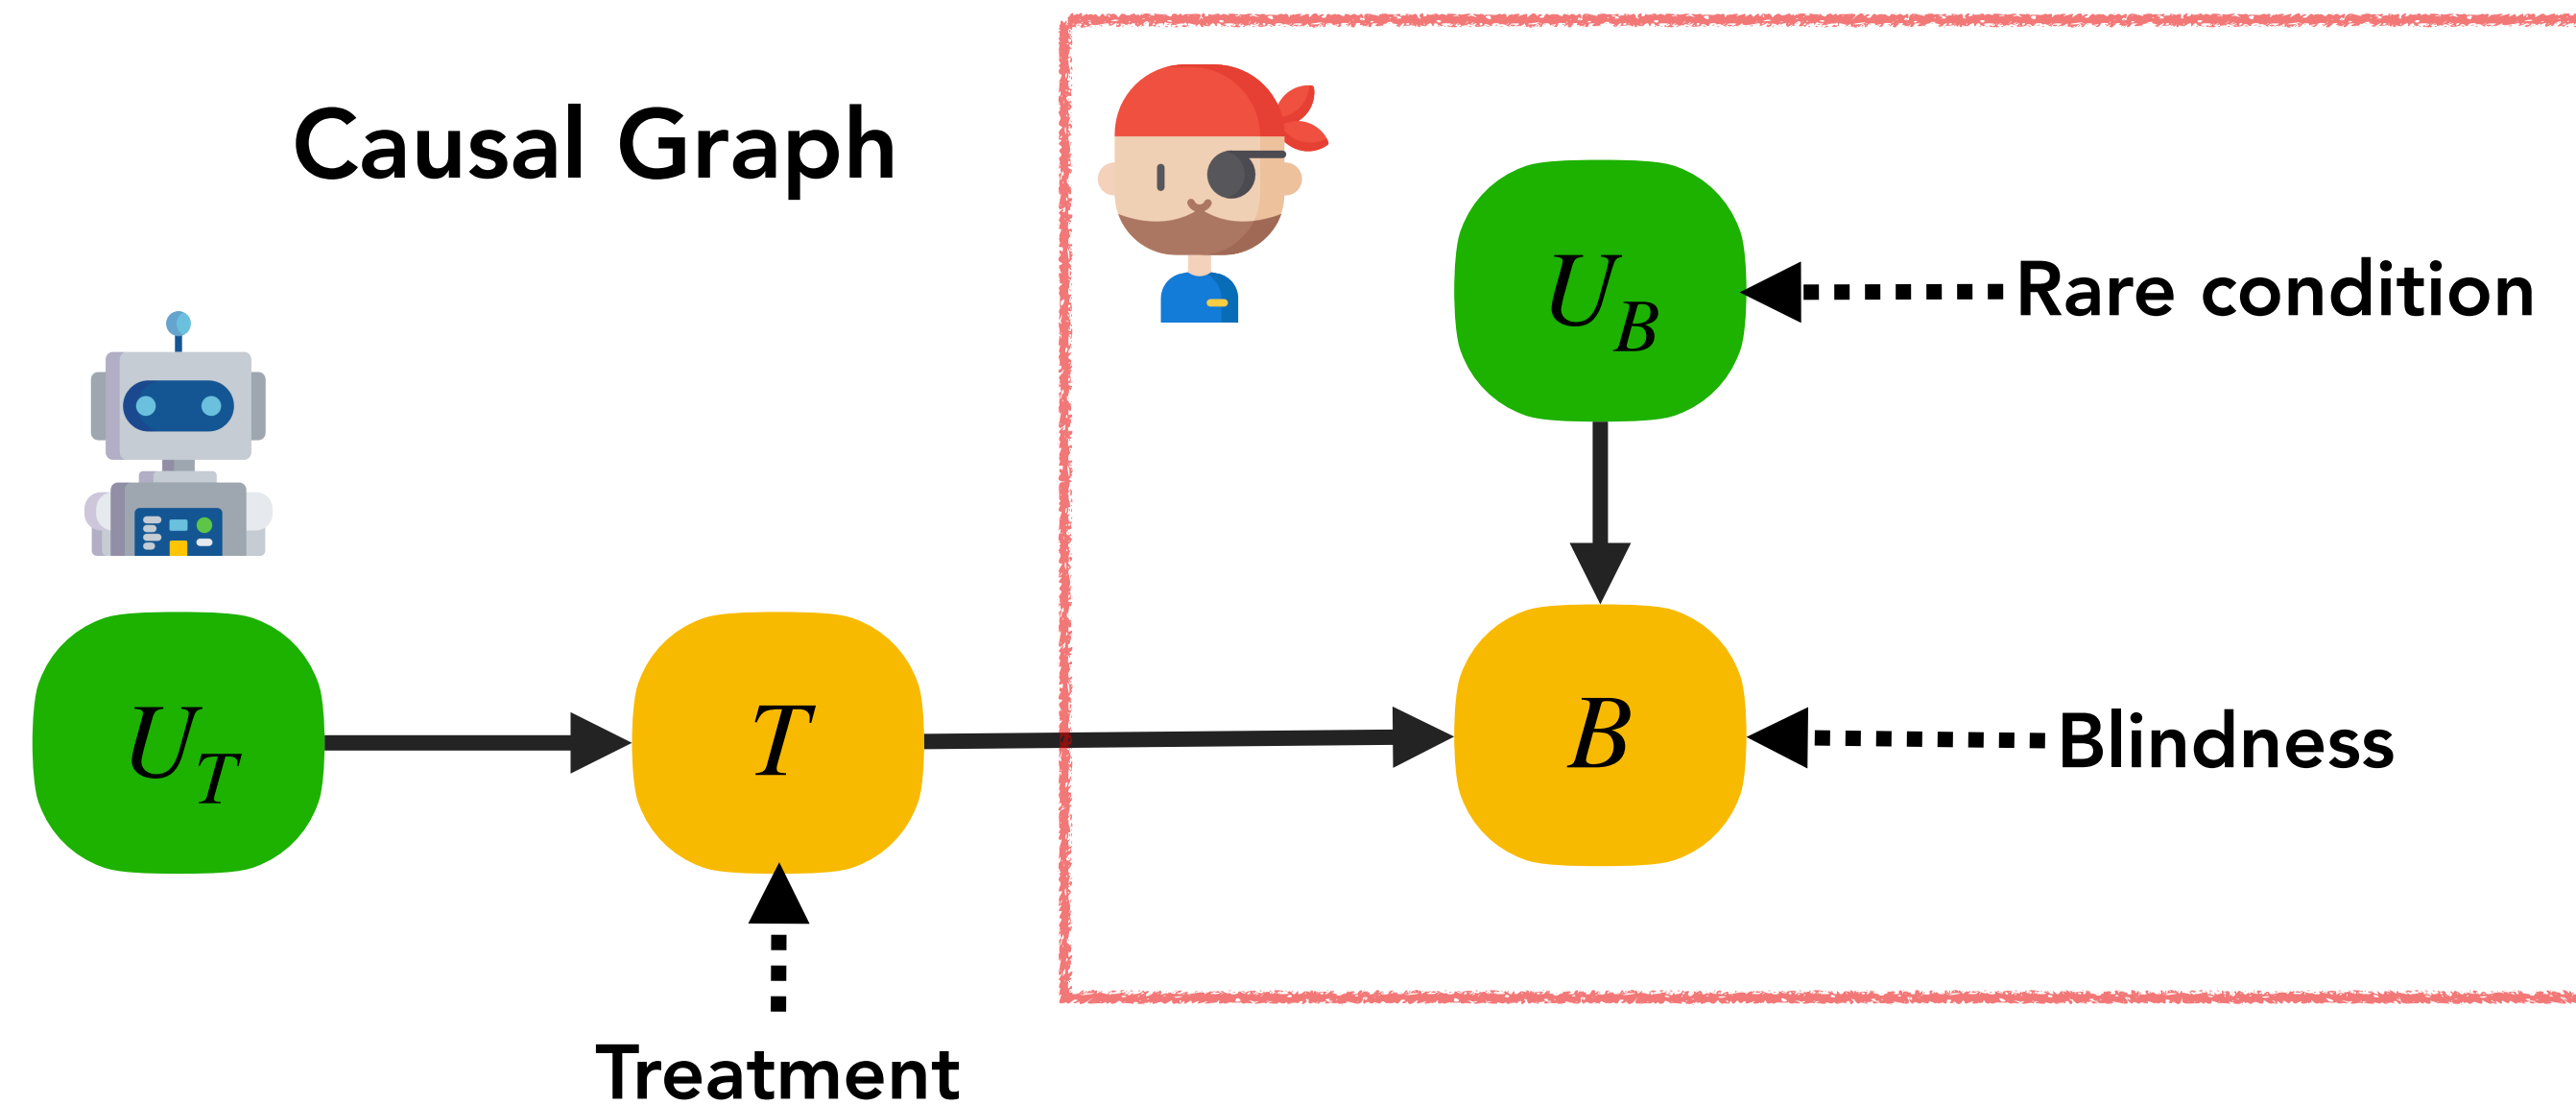
**Counterfactual question**

The treatment was administered and the patient got blind. What would have happened if the treatment had not been administered?



# What kind of (causal) questions can we answer with SCMs?

(1) Observational, (2) Interventional and (3) **Counterfactual** Queries



**Modified Structural Causal Model**  $\mathcal{M}_{T=1, B=1}$

~~$T := 1$~~   $T := 0$

$B := T$

$U_B = 1$  with prob. 1

"imagine"

**Counterfactual question**

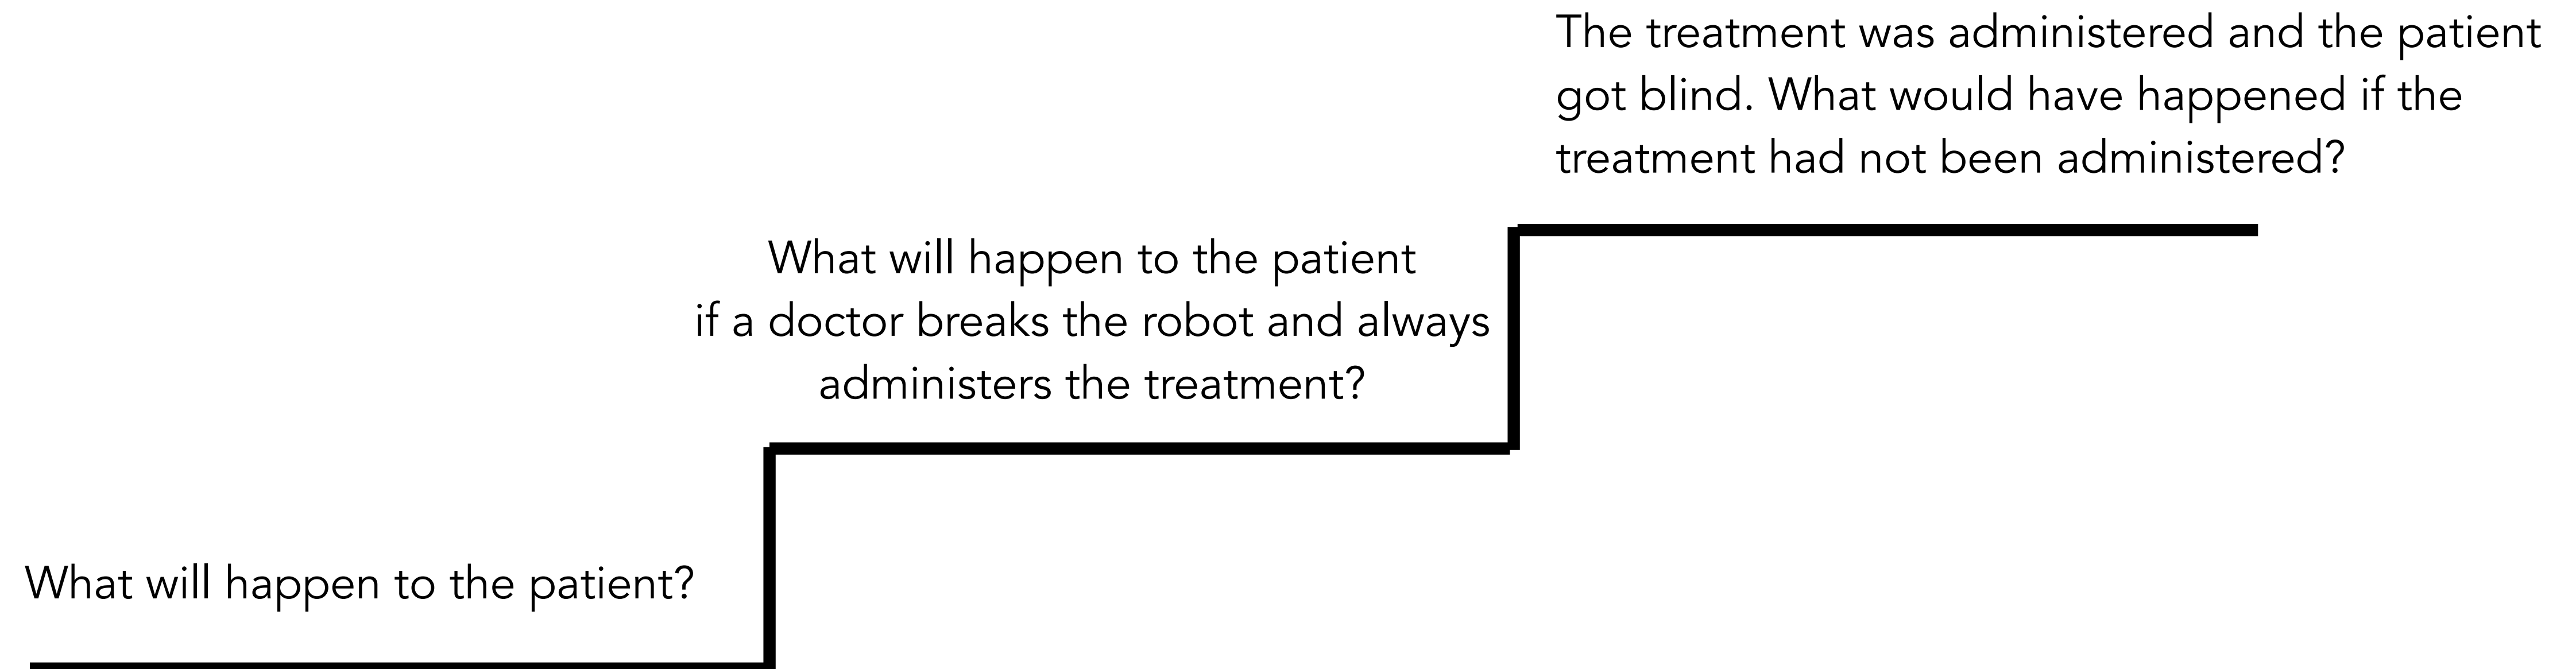
The treatment was administered and the patient got blind. **What would have happened if the treatment had not been administered?**

The patient would not have gotten blind ( $B = 0$ )

Formally,  $P^{\mathcal{M}} | T=1, B=1 ; do(T=1)(B = 1) = 0$

# The ladder of causation

(1) Observational, (2) Interventional and (3) Counterfactual Queries



It is called **ladder of causation** because questions at level  $i \in \{1,2,3\}$  can only be answered if information from level  $j \geq i$  is available. Counterfactuals sit at the top of the ladder!

Pearl. *"Causality."* Cambridge university press, 2009.

Bareinboim et al. *"On Pearl's hierarchy and the foundations of causal inference."* Probabilistic and causal inference: the works of Judea Pearl, 2022.



# Identifiability

Identification of

an interventional probability, e.g.,  $P^{\mathcal{M}}; do(T=1)(B)$ , or

a counterfactual probability, e.g.,  $P^{\mathcal{M}} | T=1, B=1; do(T=1)(B)$

refers to the process of estimating it using (observational) data from  $\mathcal{M}$ .

Shpitser and Pearl. "Complete identification methods for the causal hierarchy." JMLR, 2008.

Perkovic et al. "Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs." JMLR, 2018.

Shalit et al. "Estimating individual treatment effect: generalization bounds and algorithms." ICML, 2017.

Kallus. "Treatment effect risk: Bounds and inference." Management Science, 2023.

# Identifiability

Identification of

an interventional probability, e.g.,  $P^{\mathcal{M}}; do(T=1)(B)$ , or

a counterfactual probability, e.g.,  $P^{\mathcal{M}} | T=1, B=1; do(T=1)(B)$

refers to the process of estimating it using (observational) data from  $\mathcal{M}$ .

If an interventional or counterfactual probability is not identifiable, then regardless of how much data we have, we will not be able to estimate it.

Shpitser and Pearl. "Complete identification methods for the causal hierarchy." JMLR, 2008.

Perkovic et al. "Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs." JMLR, 2018.

Shalit et al. "Estimating individual treatment effect: generalization bounds and algorithms." ICML, 2017.

Kallus. "Treatment effect risk: Bounds and inference." Management Science, 2023.

# Identifiability

Identification of

an interventional probability, e.g.,  $P^{\mathcal{M}}; do(T=1)(B)$ , or

a counterfactual probability, e.g.,  $P^{\mathcal{M}} | T=1, B=1; do(T=1)(B)$

refers to the process of estimating it using (observational) data from  $\mathcal{M}$ .

If an interventional or counterfactual probability is not identifiable, then regardless of how much data we have, we will not be able to estimate it.

There exist methods to

- (i) determine the identifiability of specific interventional and counterfactual probabilities, and
- (ii) estimate (or bound) quantities derived from these probabilities (e.g., individual treatment effects)

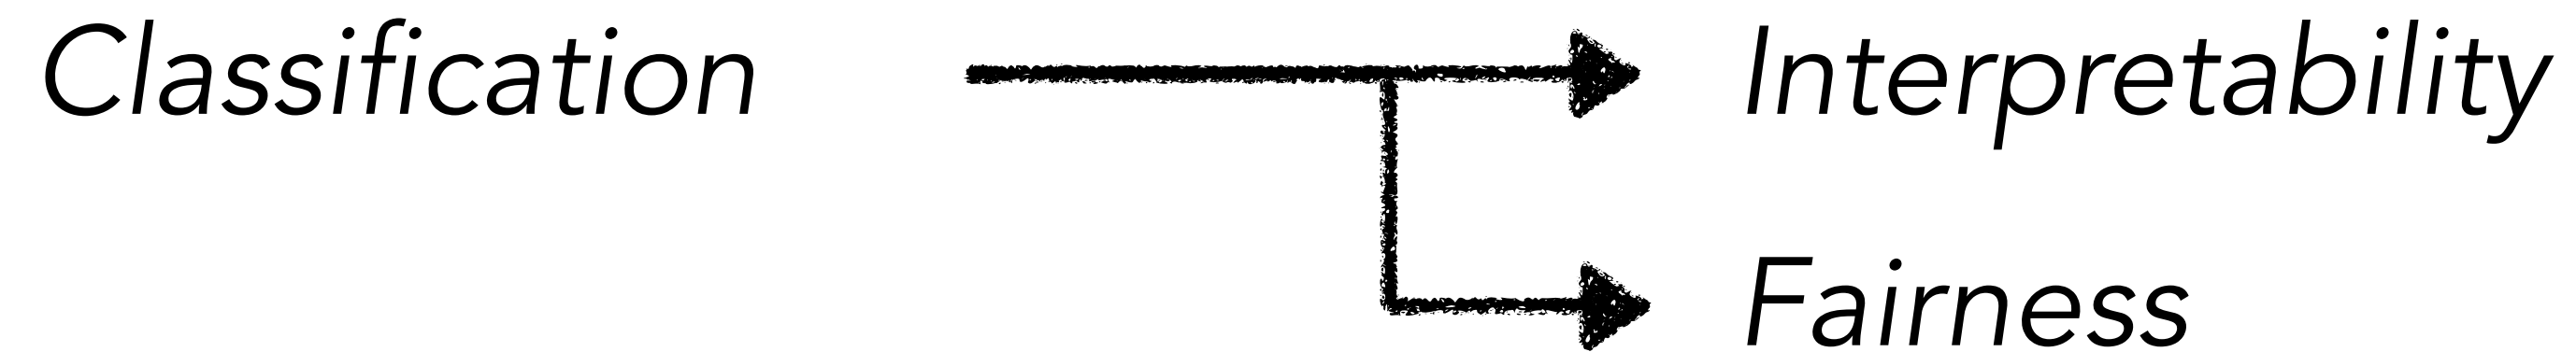
Shpitser and Pearl. "Complete identification methods for the causal hierarchy." JMLR, 2008.

Perkovic et al. "Complete graphical characterization and construction of adjustment sets in markov equivalence classes of ancestral graphs." JMLR, 2018.

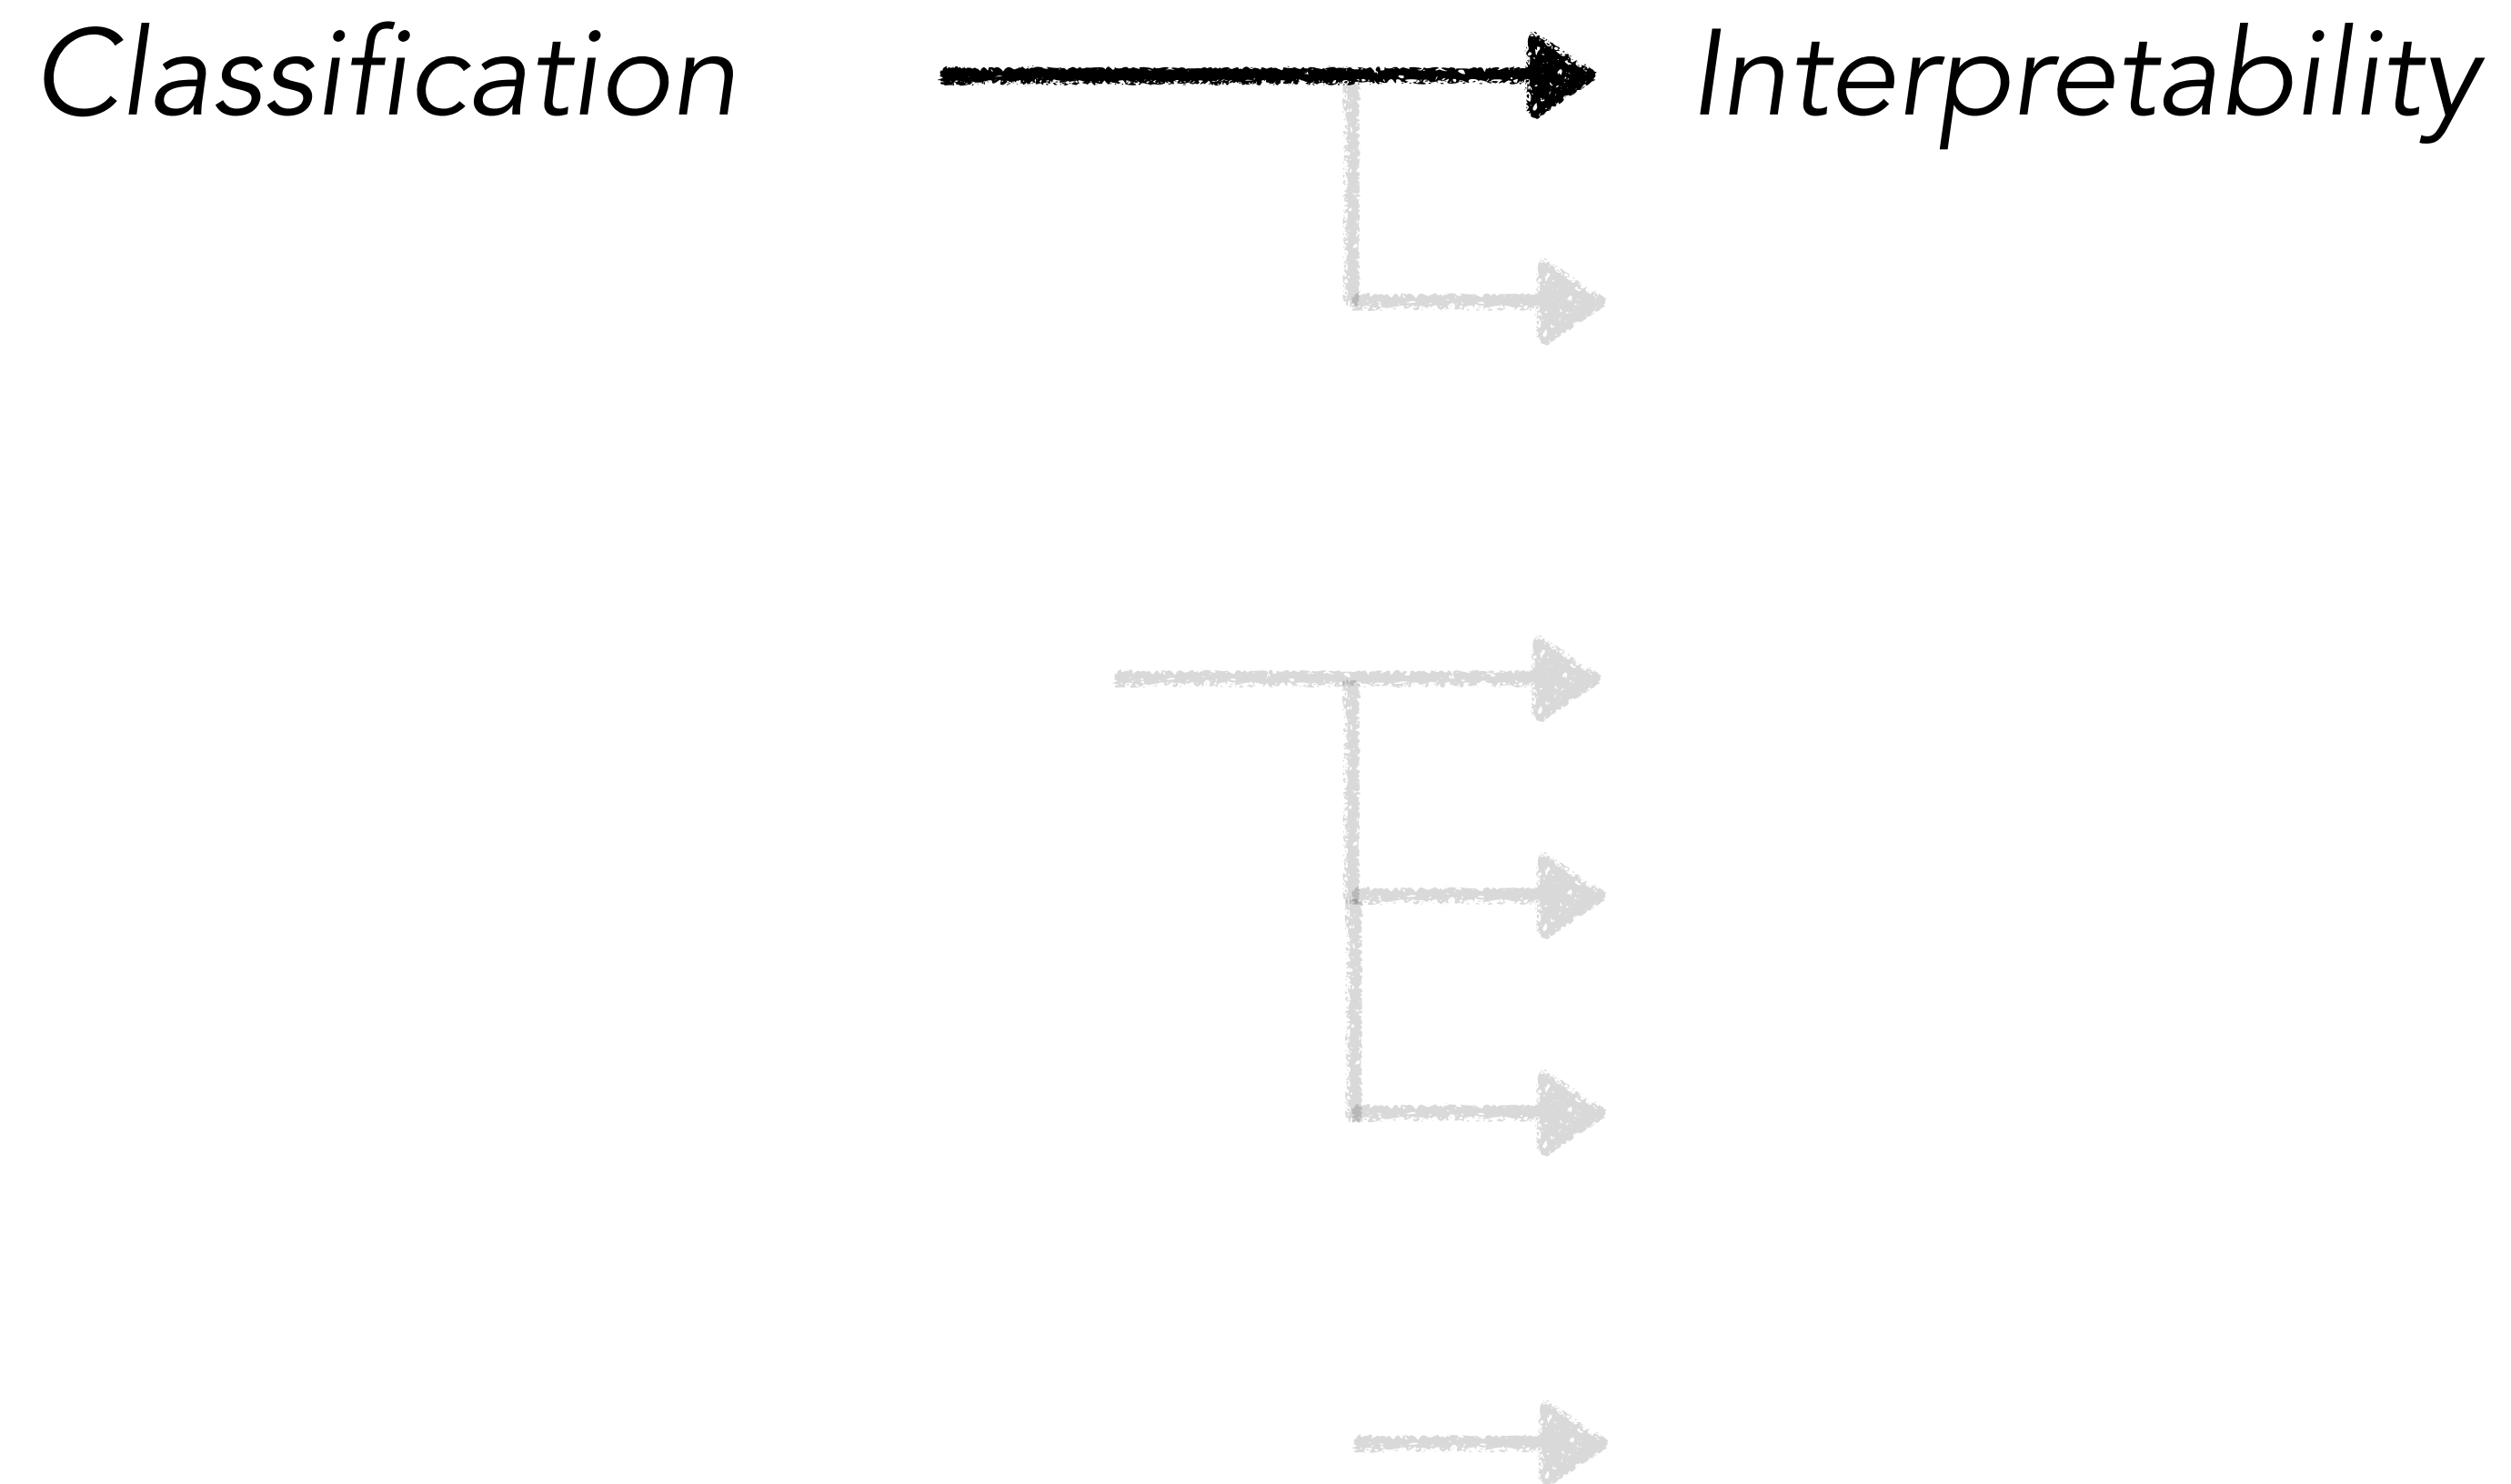
Shalit et al. "Estimating individual treatment effect: generalization bounds and algorithms." ICML, 2017.

Kallus. "Treatment effect risk: Bounds and inference." Management Science, 2023.

# Use cases of counterfactuals in machine learning



# Use cases of counterfactuals in machine learning



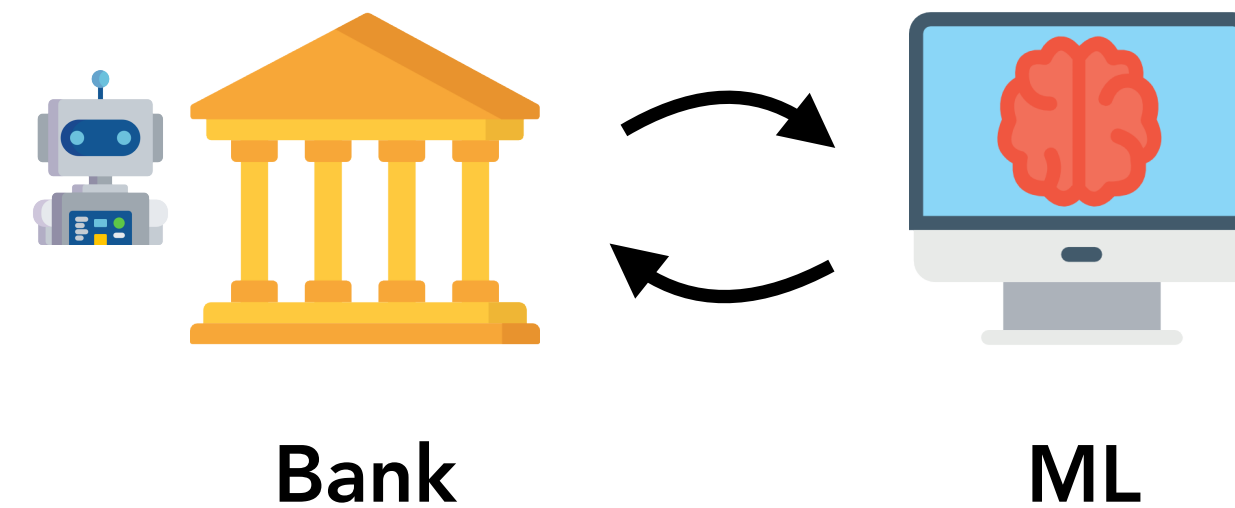
# Counterfactual explanations

The term counterfactual has arguably become mainstream in the field of machine learning after the seminal work on counterfactual explanations by Wachter et al.

Wachter et al. "*Counterfactual explanations without opening the black box: Automated decisions and the GDPR.*" Harv. JL & Tech., 2017.

# Counterfactual explanations

The term counterfactual has arguably become mainstream in the field of machine learning after the seminal work on counterfactual explanations by Wachter et al.



Wachter et al. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR." Harv. JL & Tech., 2017.



# Counterfactual explanations

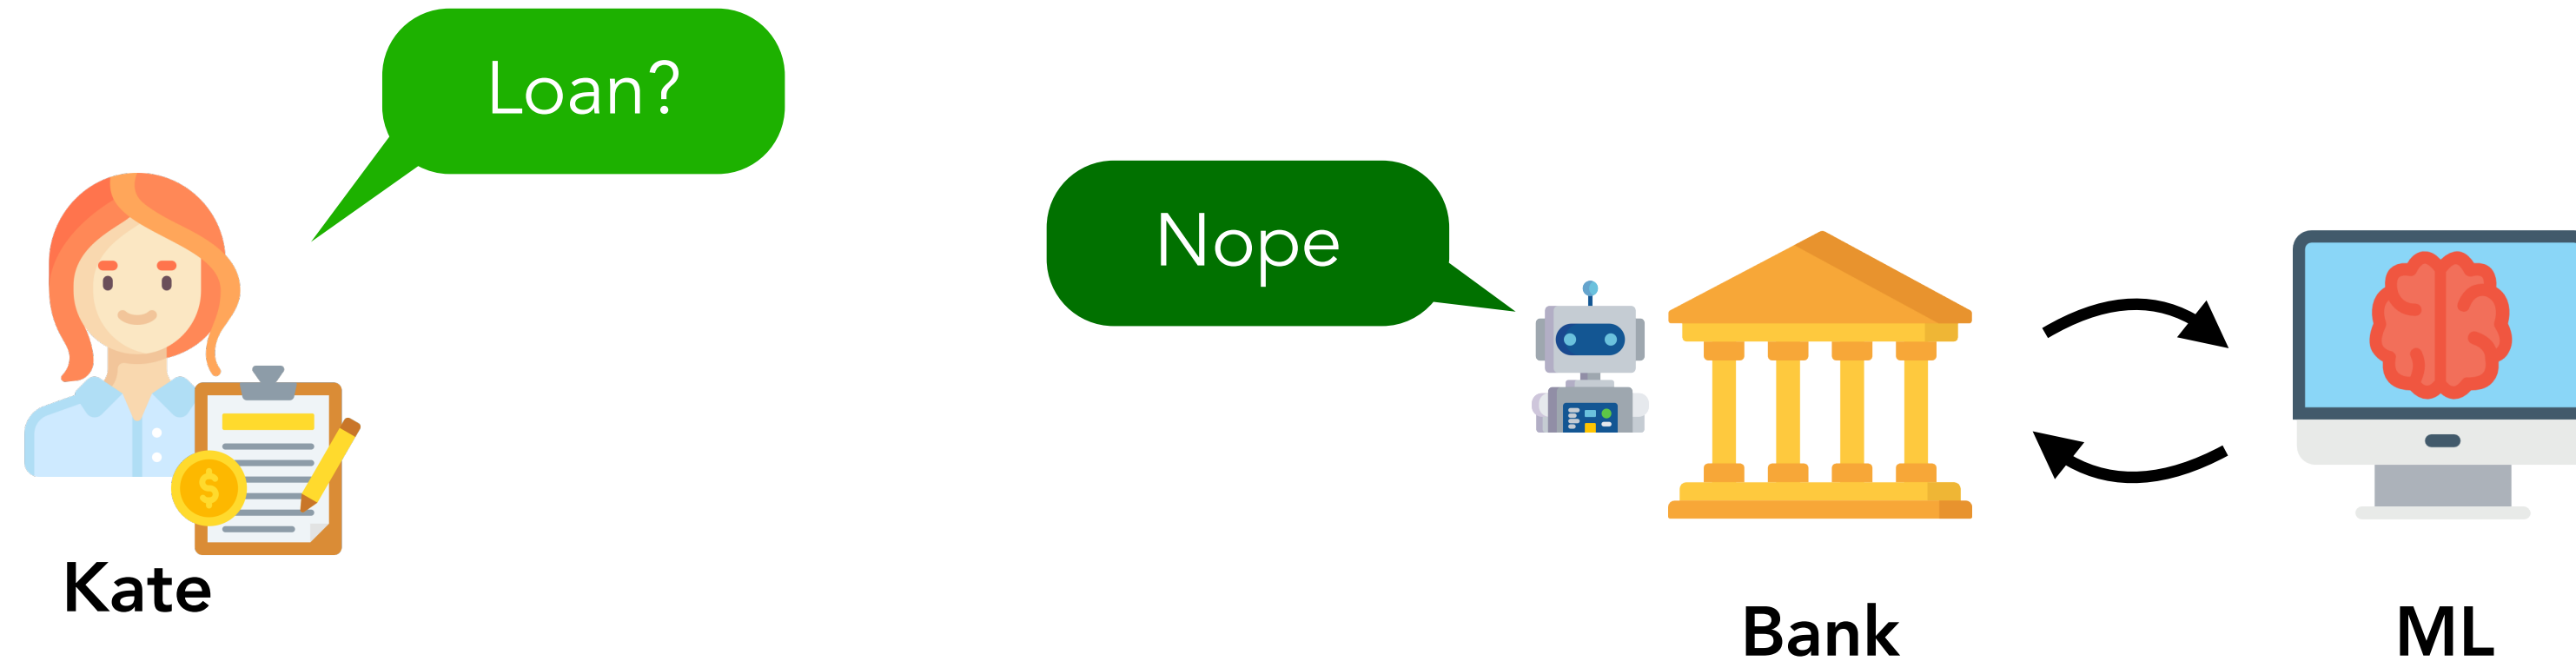
The term counterfactual has arguably become mainstream in the field of machine learning after the seminal work on counterfactual explanations by Wachter et al.



Wachter et al. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR." Harv. JL & Tech., 2017.

# Counterfactual explanations

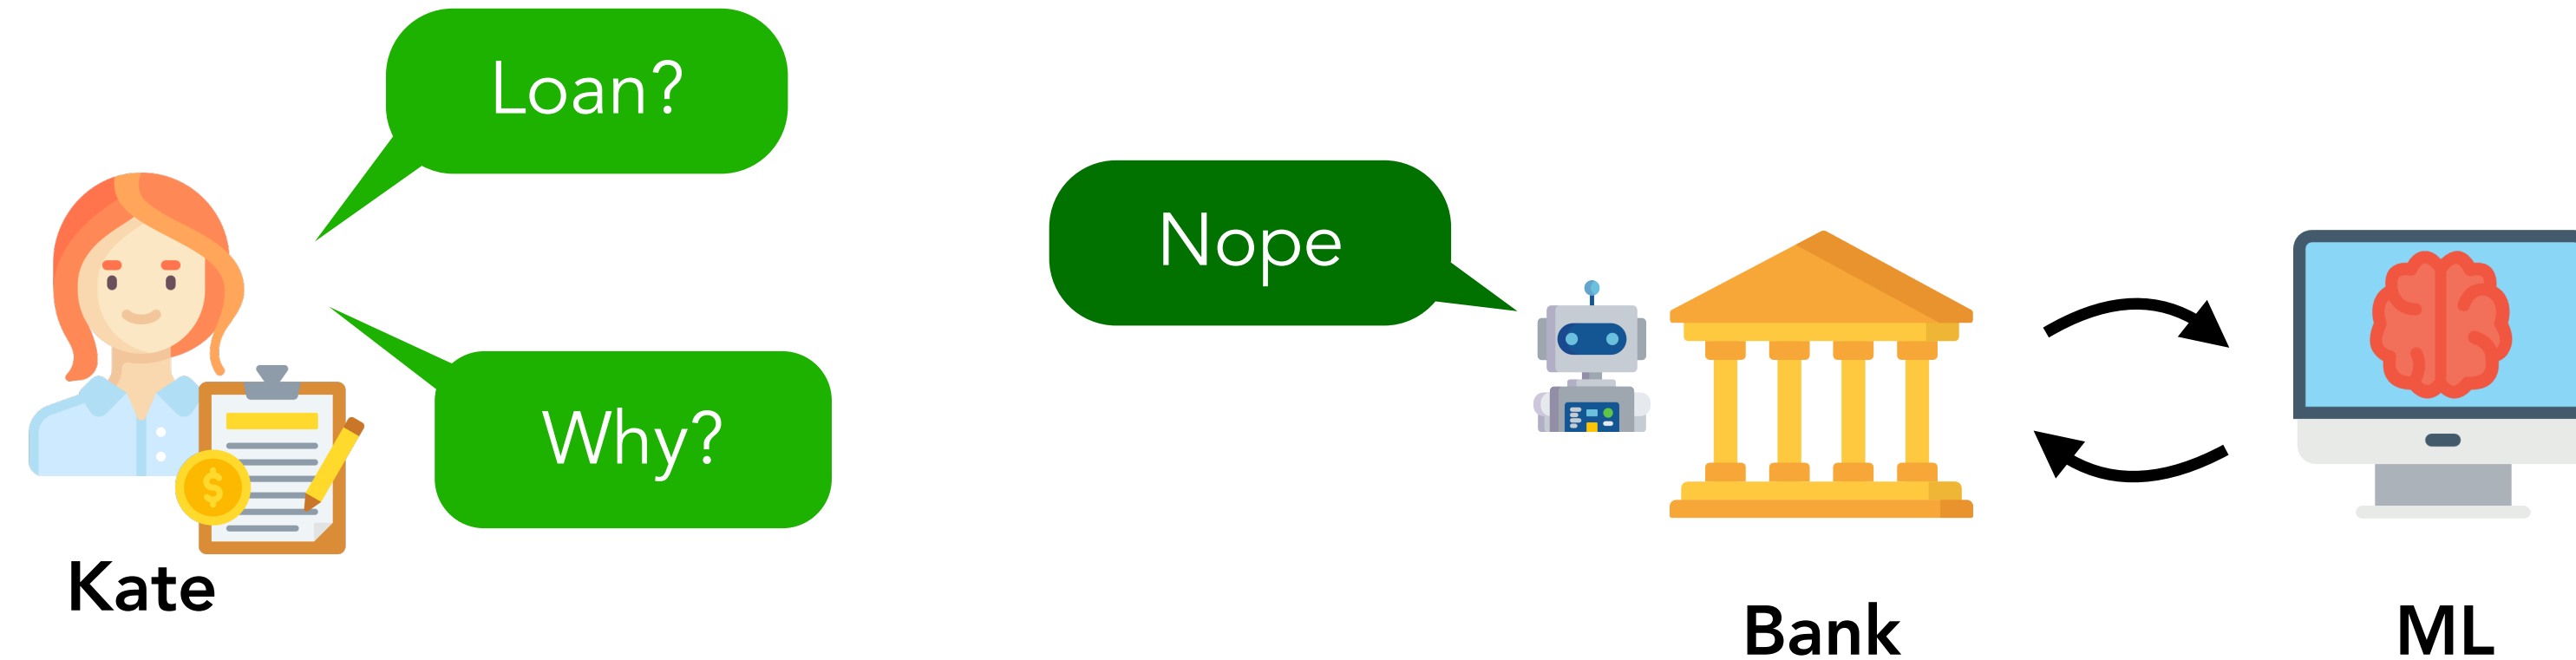
The term counterfactual has arguably become mainstream in the field of machine learning after the seminal work on counterfactual explanations by Wachter et al.



Wachter et al. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR." Harv. JL & Tech., 2017.

# Counterfactual explanations

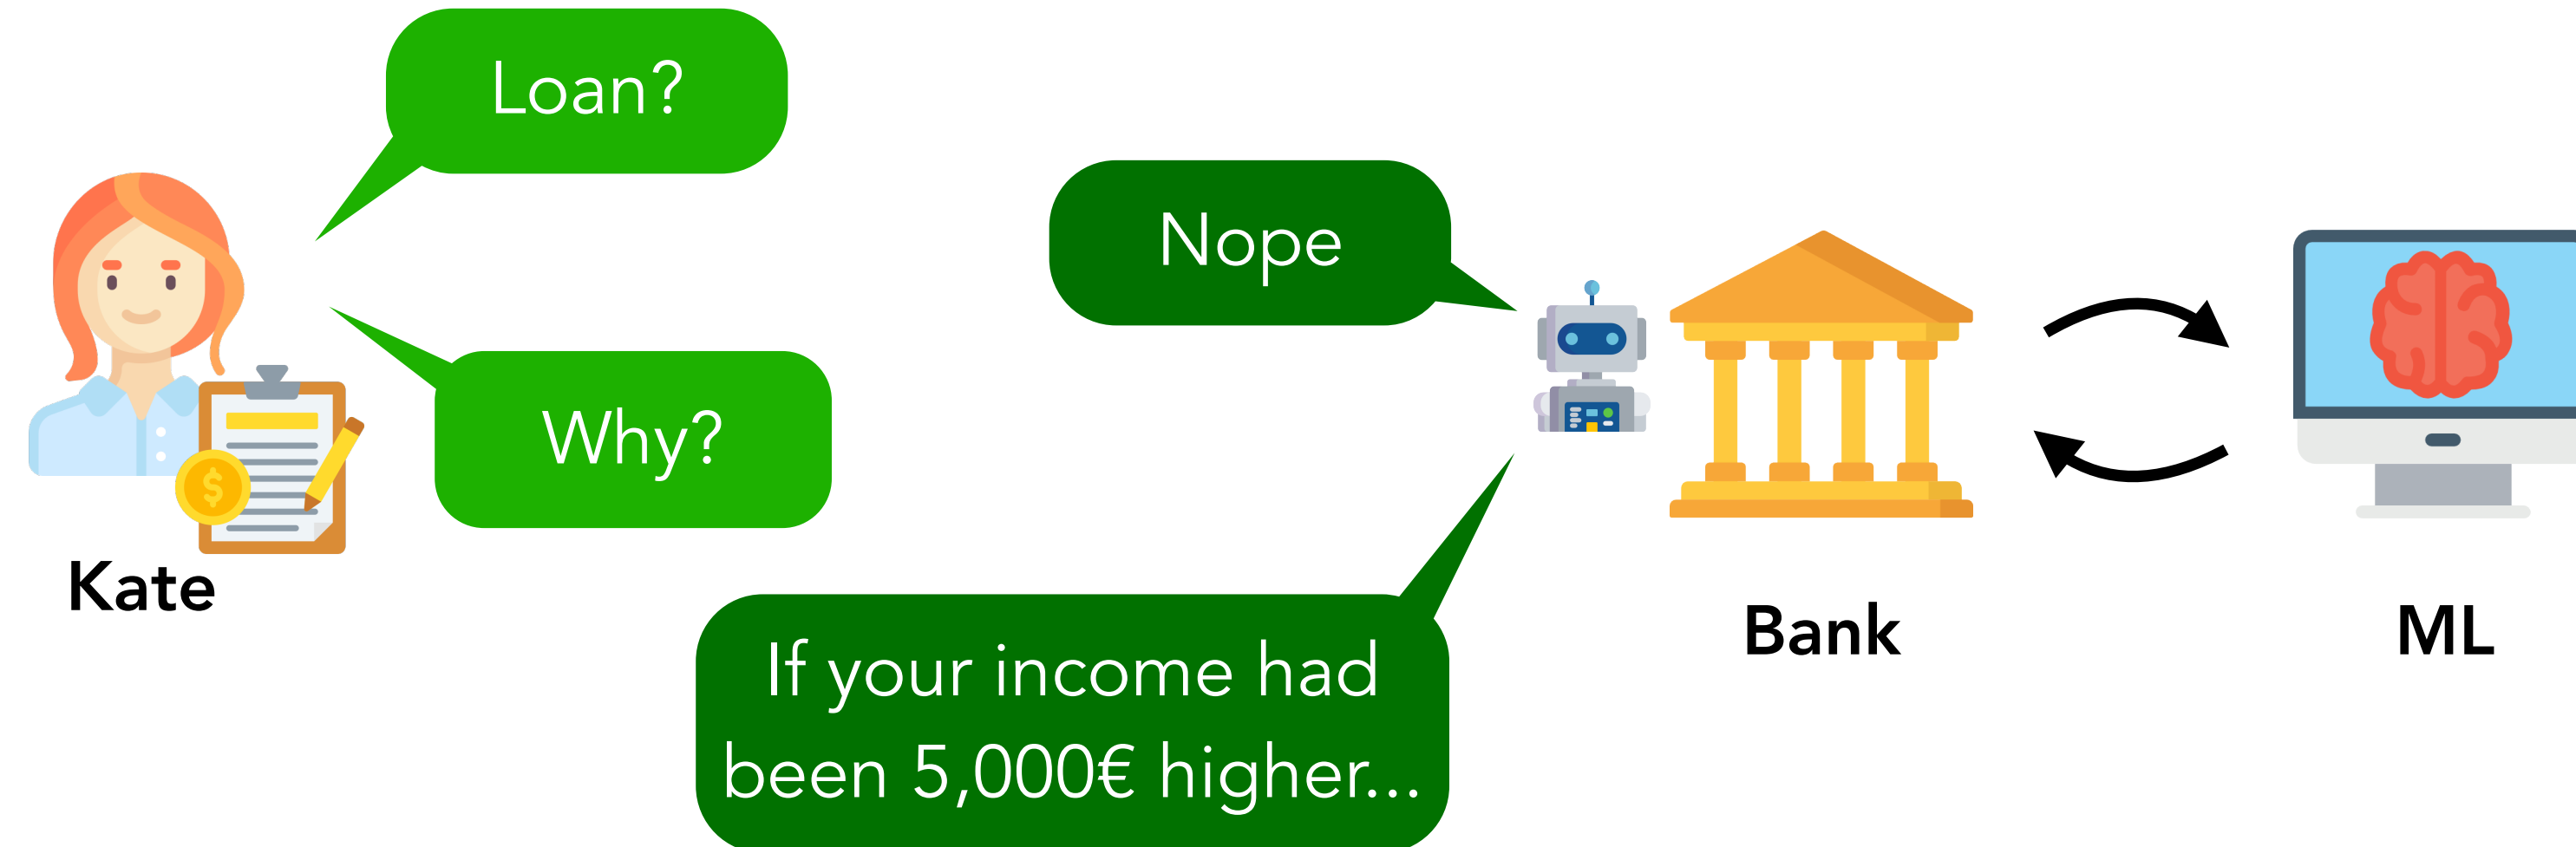
The term counterfactual has arguably become mainstream in the field of machine learning after the seminal work on counterfactual explanations by Wachter et al.



Wachter et al. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR." Harv. JL & Tech., 2017.

# Counterfactual explanations

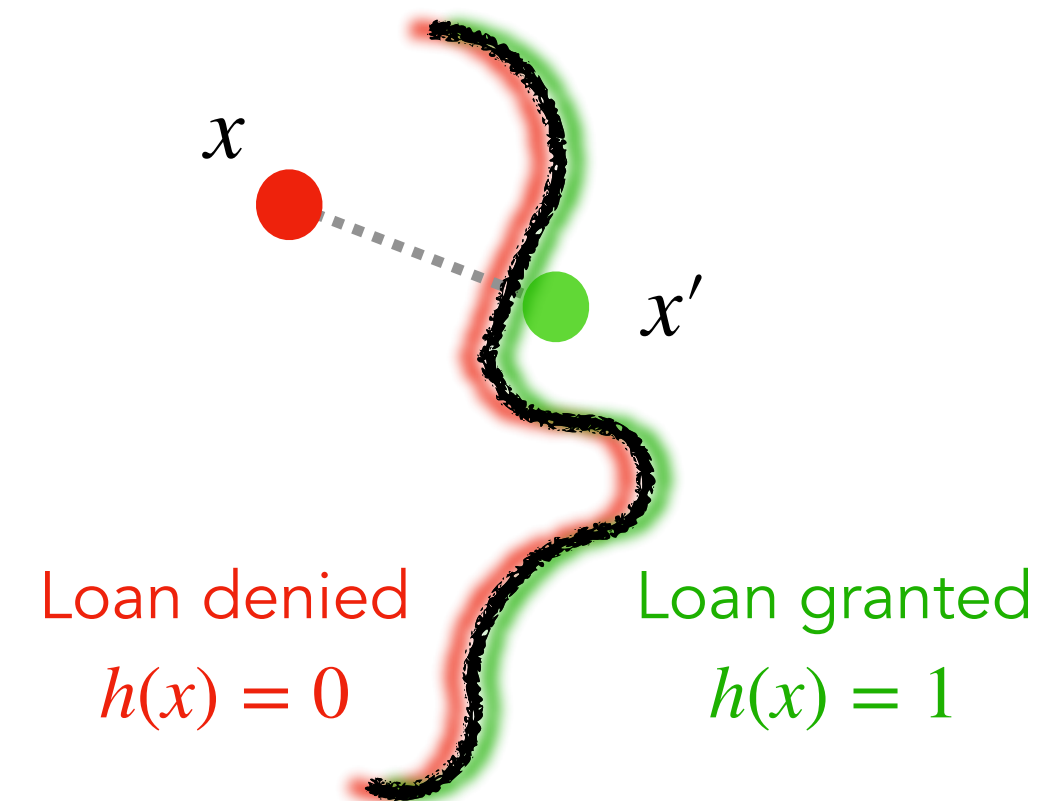
The term counterfactual has arguably become mainstream in the field of machine learning after the seminal work on counterfactual explanations by Wachter et al.



Wachter et al. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR." Harv. JL & Tech., 2017.

# Counterfactual explanations

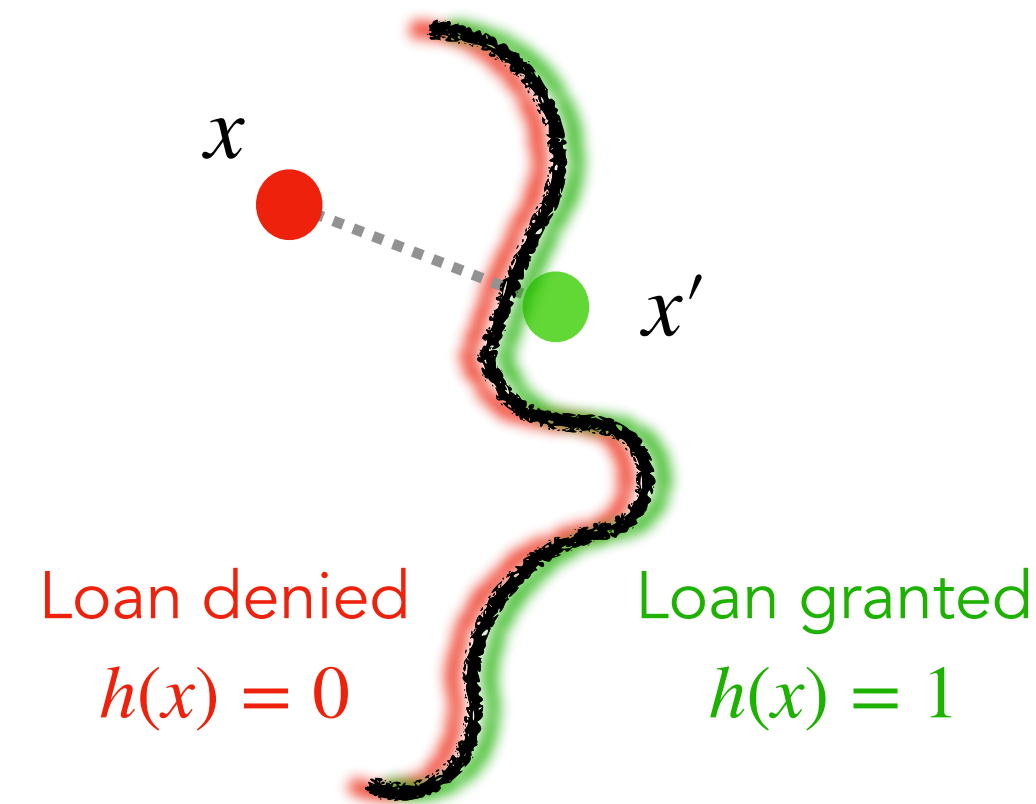
Given a (binary) prediction  $h(x)$  by a machine learning model about an individual with features  $x$ , a counterfactual explanation is given by the closest feature value  $x'$  under which  $h(x') \neq h(x)$



Wachter et al. "Counterfactual explanations without opening the black box: Automated decisions and the GDPR." Harv. JL & Tech., 2017.

# Counterfactual explanations

Given a (binary) prediction  $h(x)$  by a machine learning model about an individual with features  $x$ , a counterfactual explanation is given by the closest feature value  $x'$  under which  $h(x') \neq h(x)$



By showing a feature-perturbed version of an individual, a counterfactual explanation is, in principle, telling the individual what to do to secure a better decision in the future.

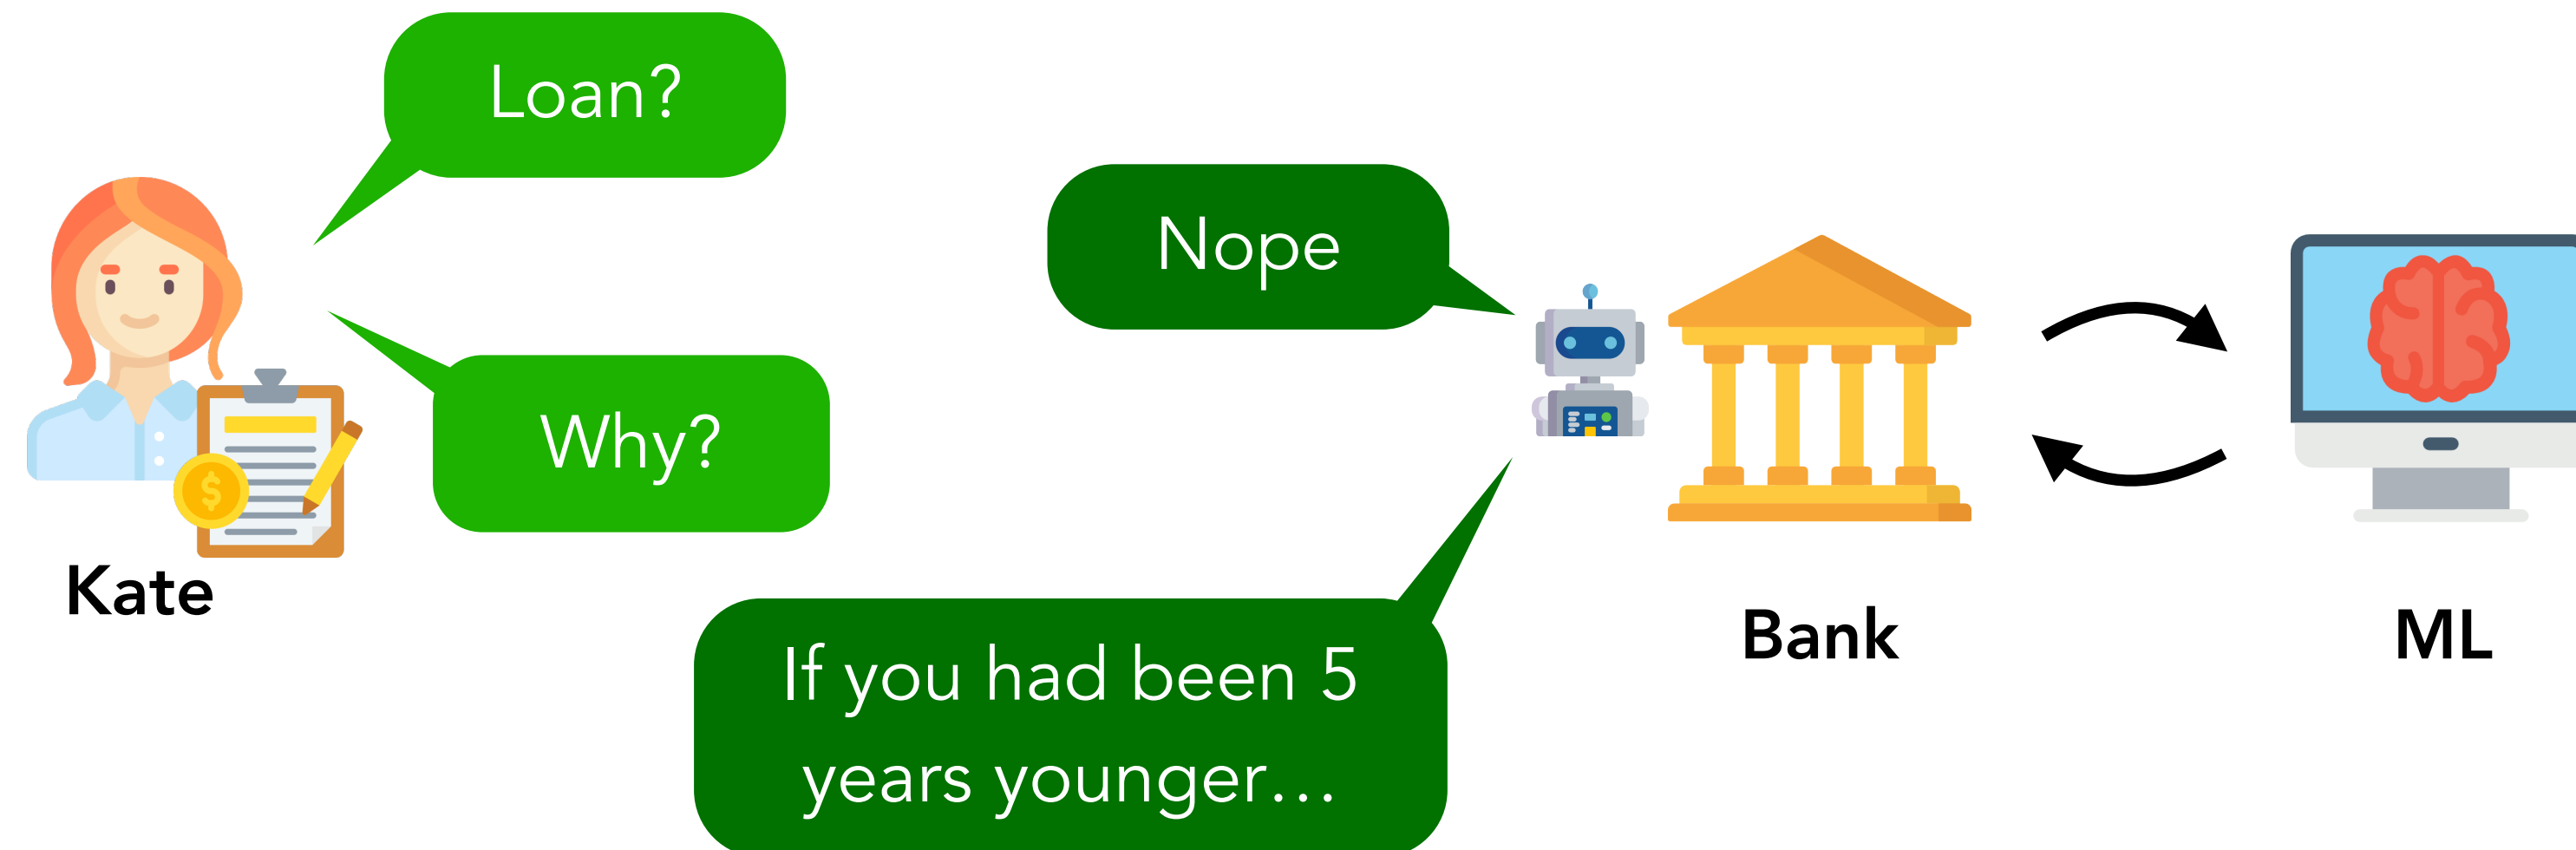
# Counterfactual explanations

However, the closest feature value  $x'$  may not be actionable, and may not even be plausible.



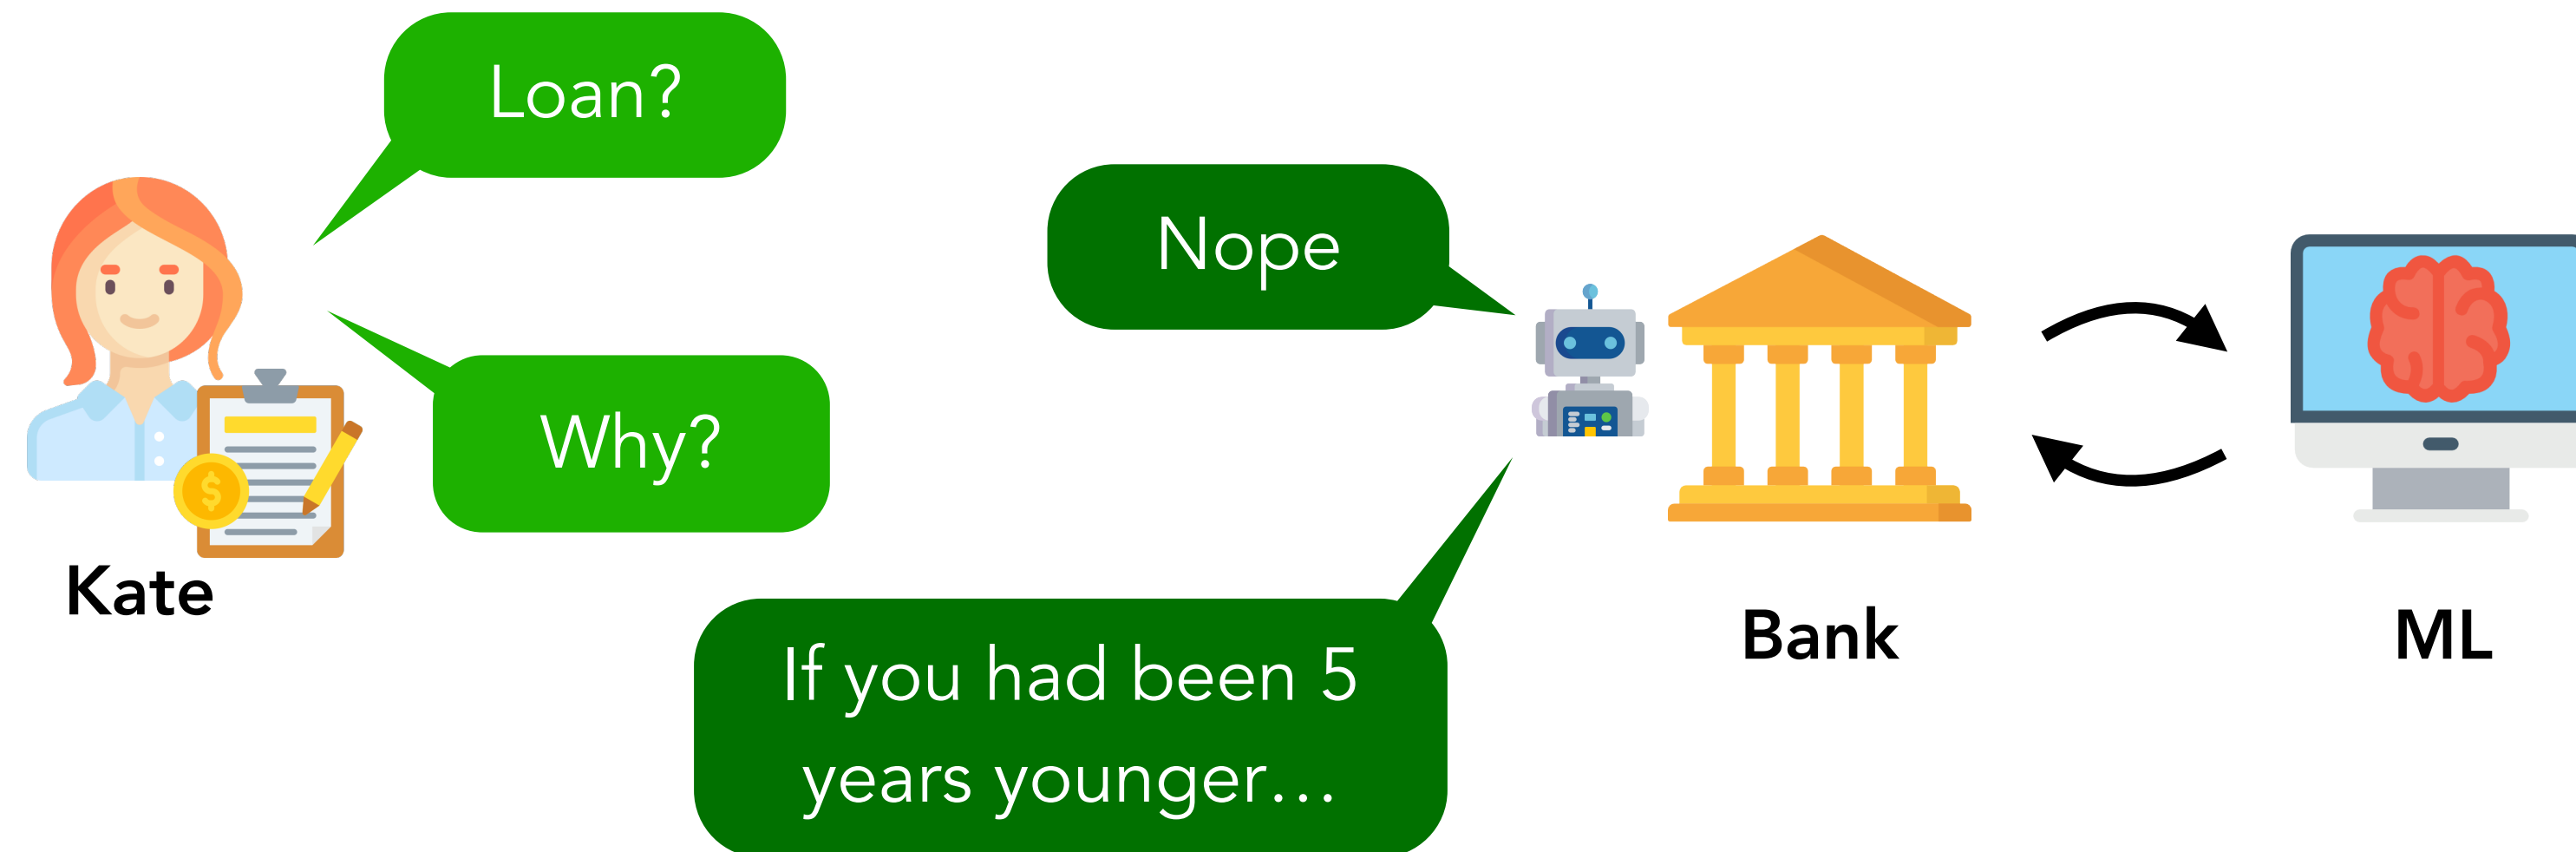
# Counterfactual explanations

However, the closest feature value  $x'$  may not be actionable, and may not even be plausible.



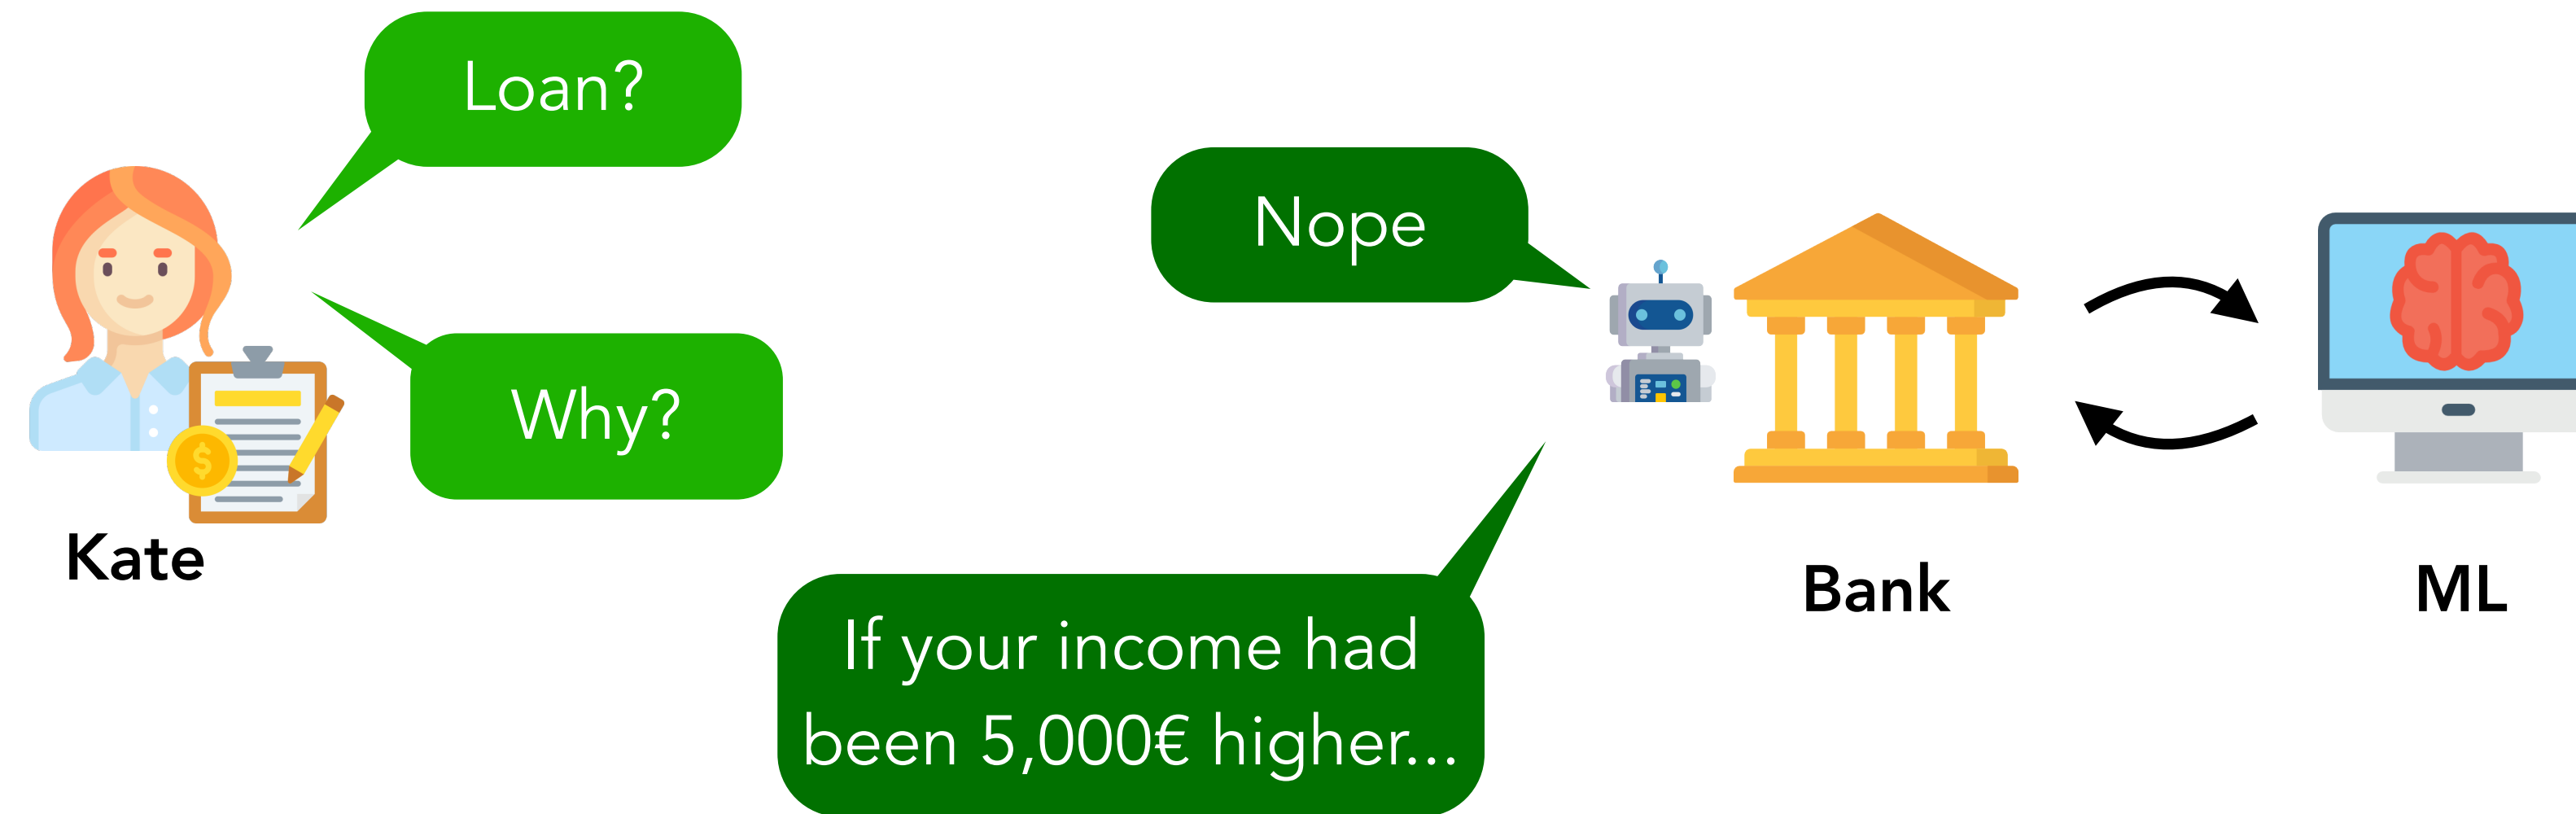
# Counterfactual explanations

However, the closest feature value  $x'$  may not be actionable, and may not even be plausible.



Many follow-up works have addressed this problem by finding the closest feature value subject to a variety of actionability and plausibility constraints.

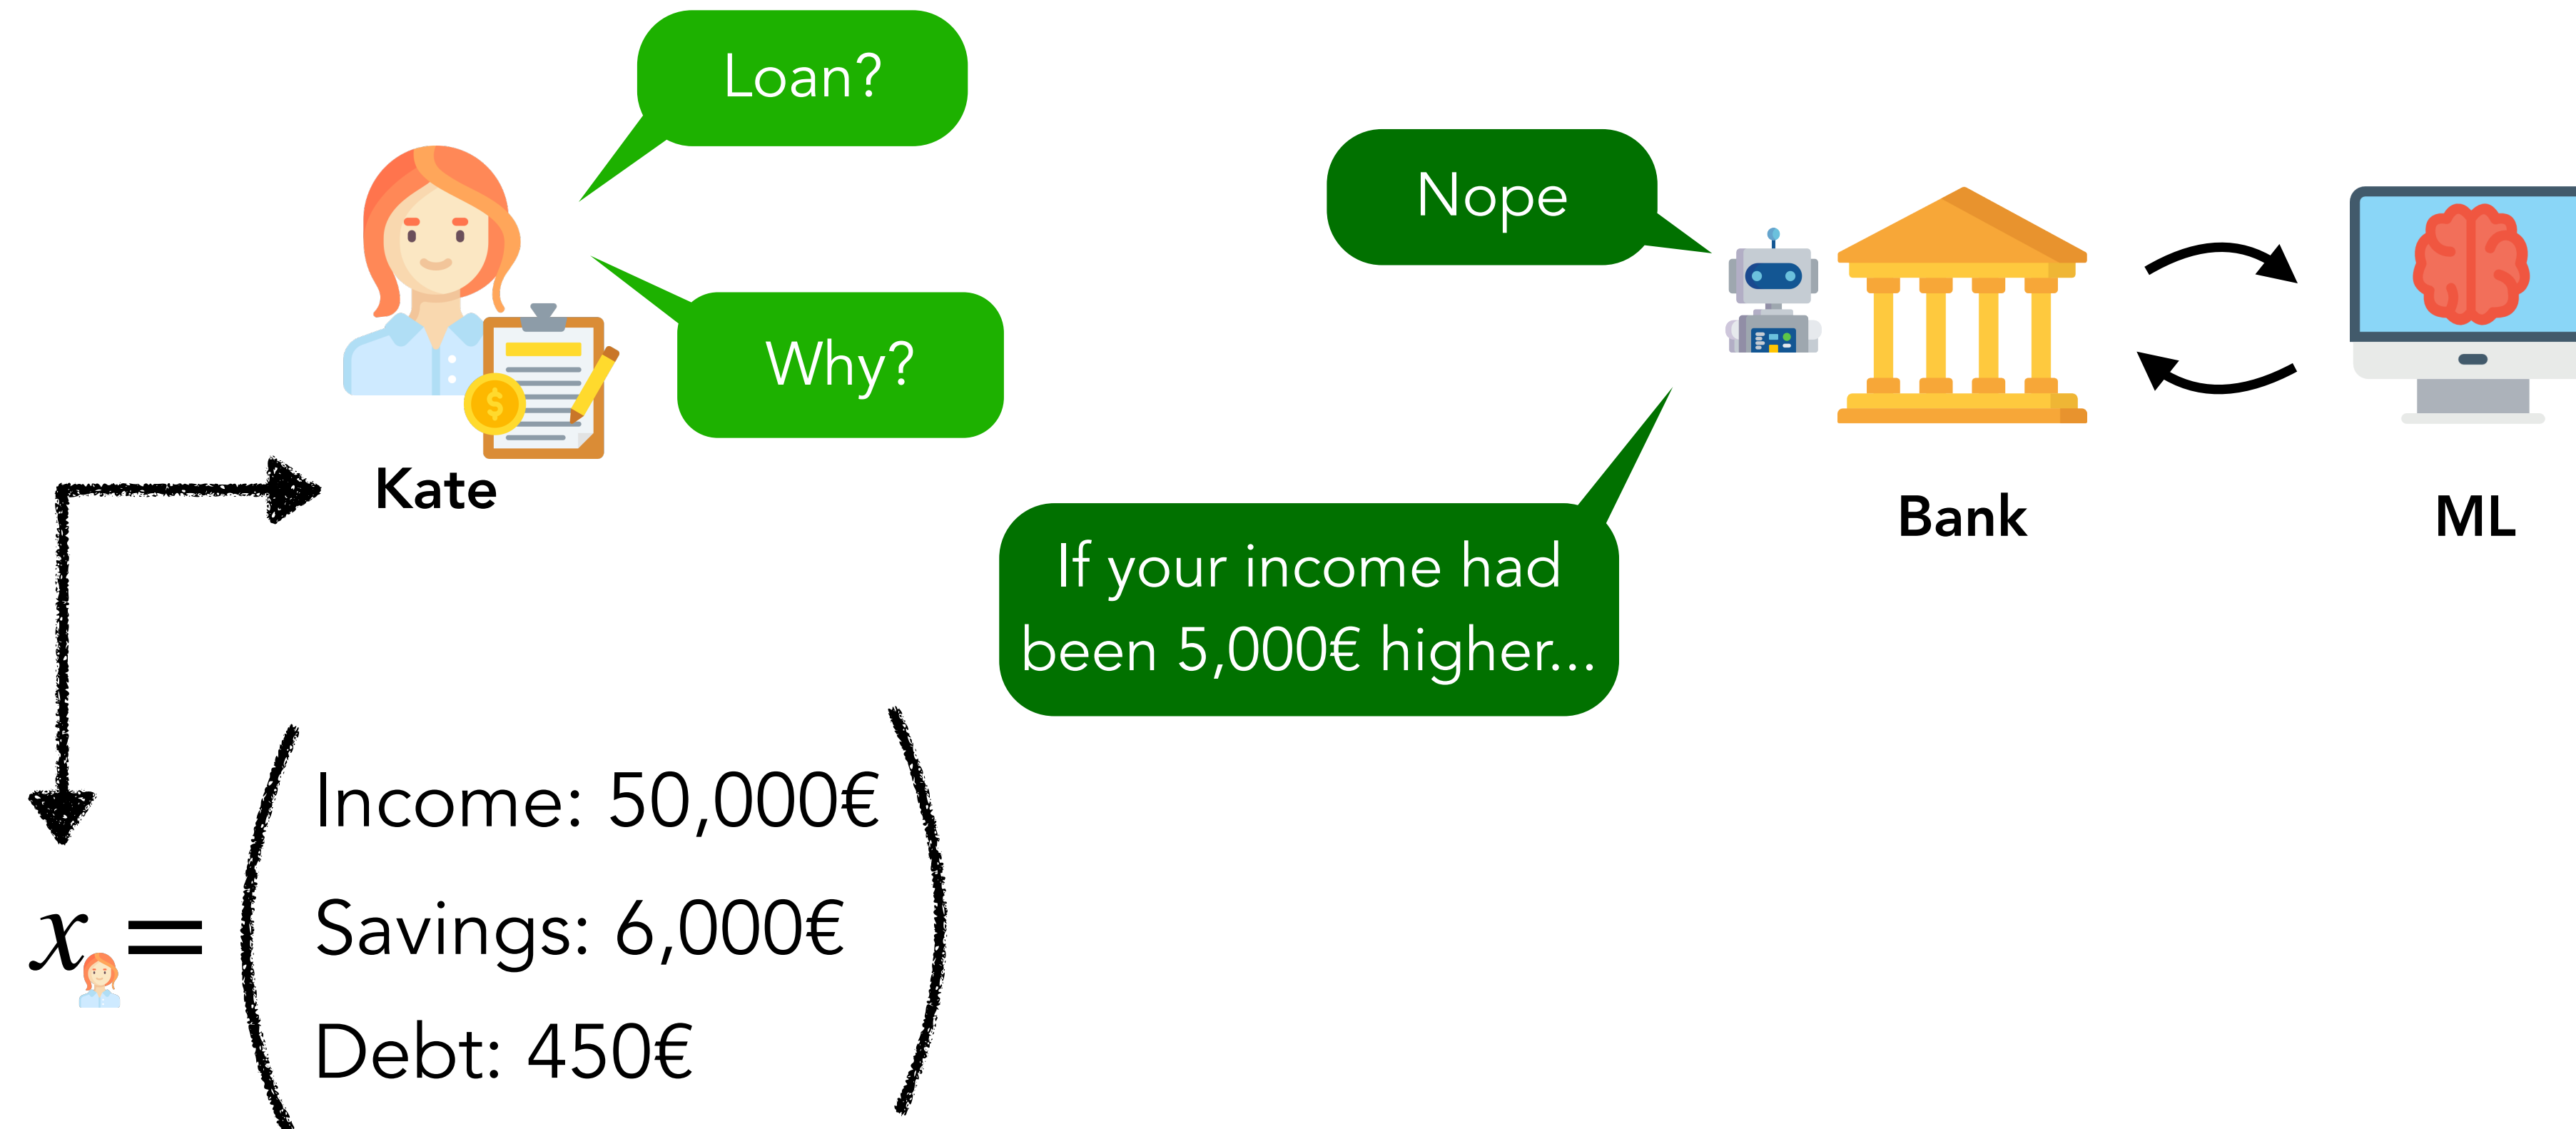
# Counterfactual explanations ignore causal dependencies



Beckers. "Causal explanations and xai." CLeaR, 2022.

Crupi et al. "Counterfactual explanations as interventions in latent space." DMKD, 2022.

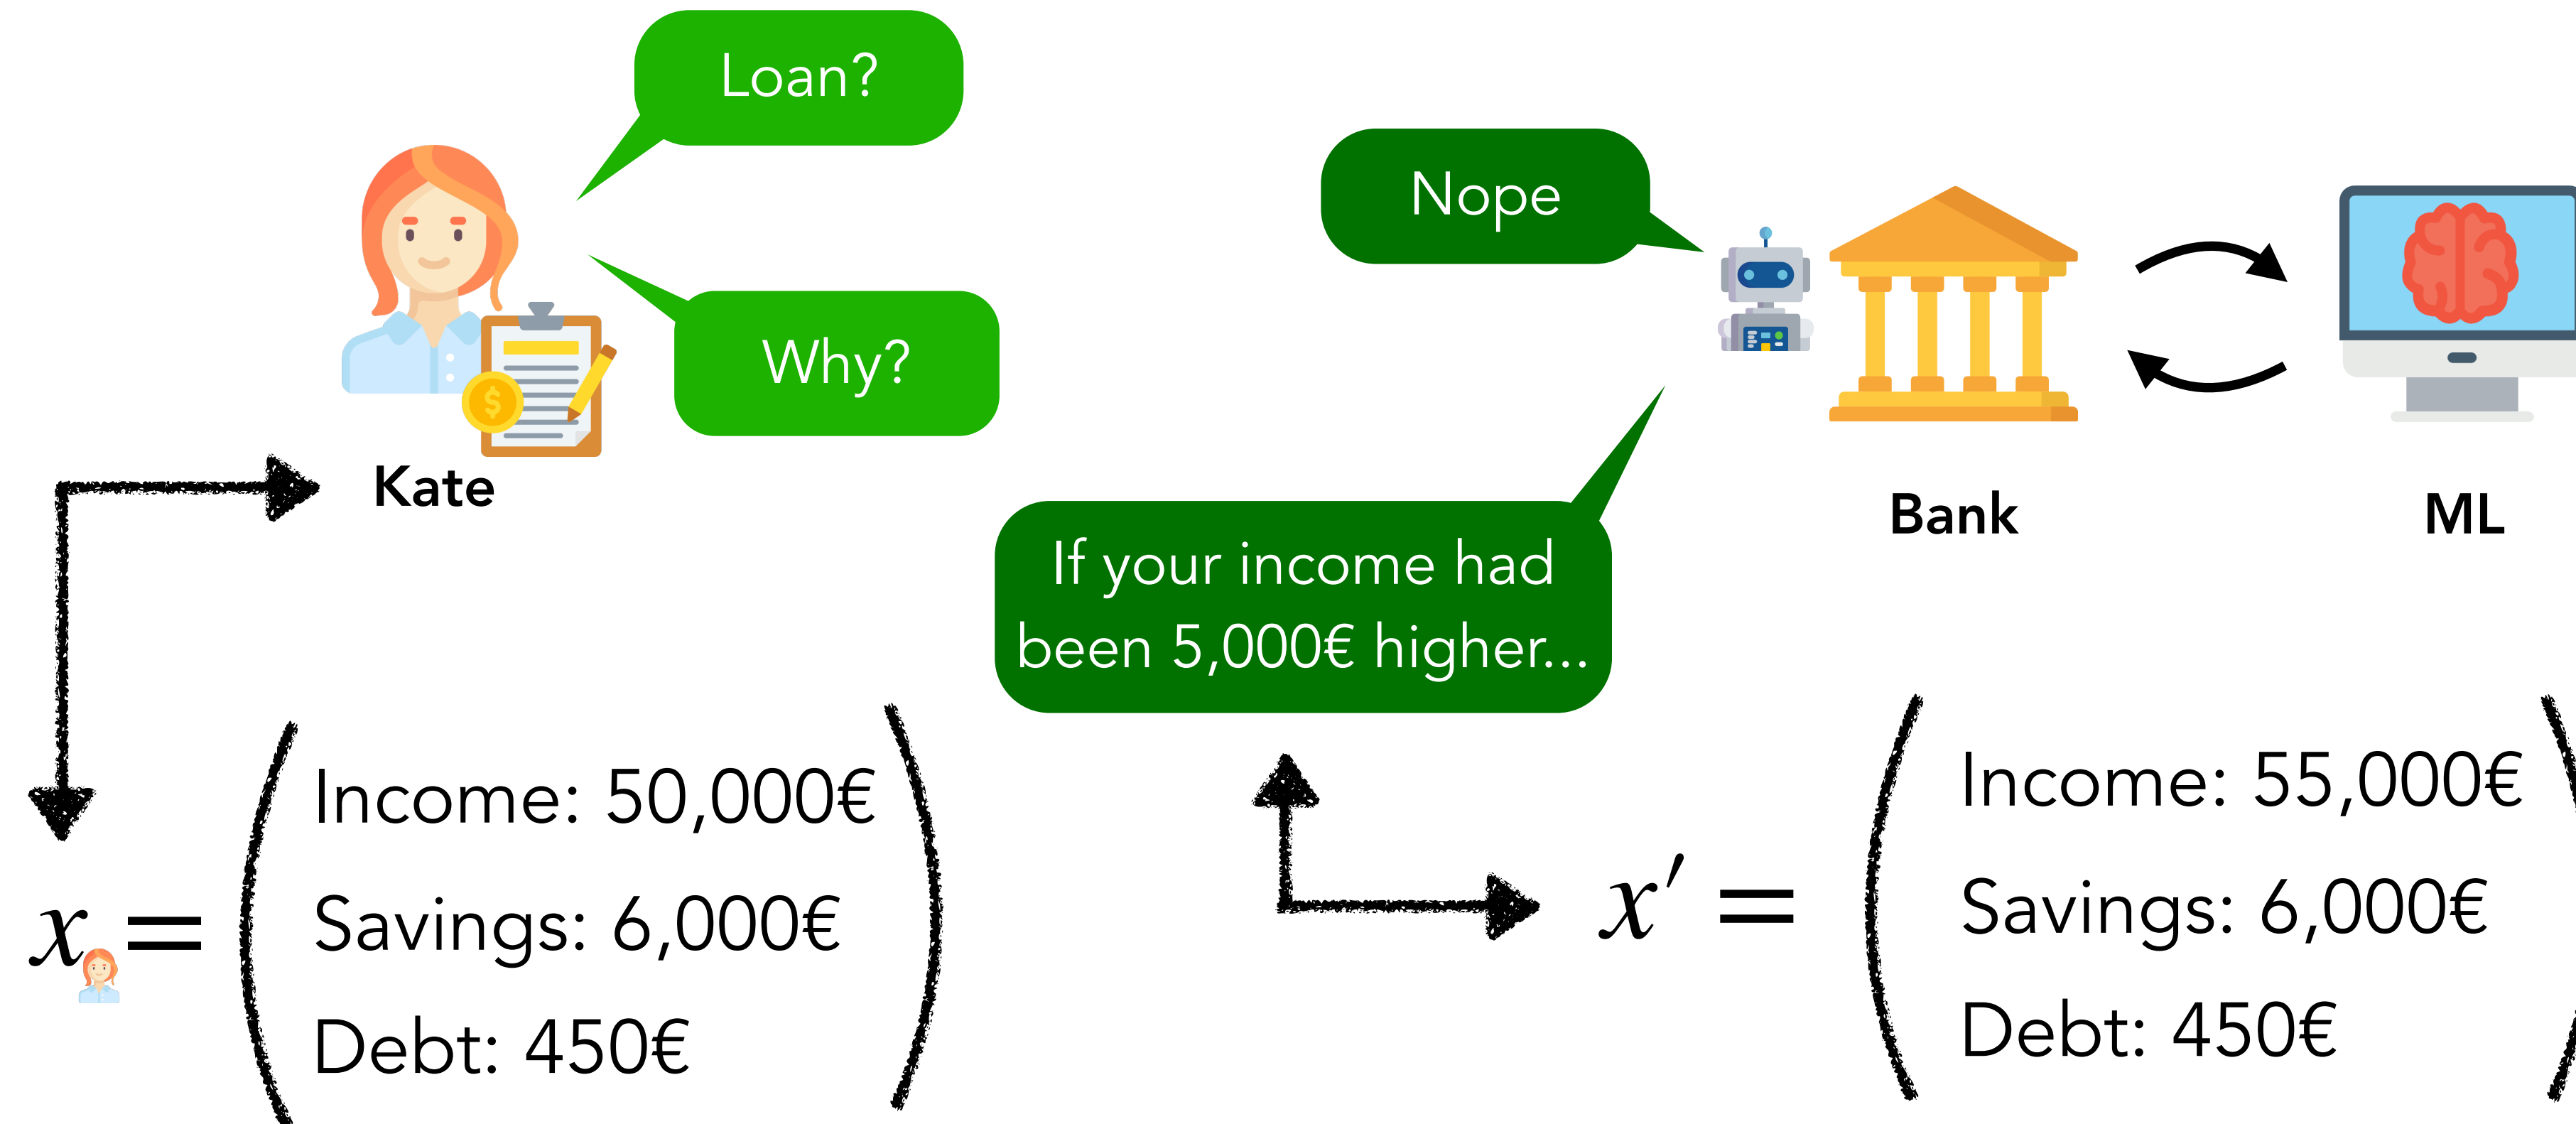
# Counterfactual explanations ignore causal dependencies



Beckers. "Causal explanations and xai." CLeaR, 2022.

Crupi et al. "Counterfactual explanations as interventions in latent space." DMKD, 2022.

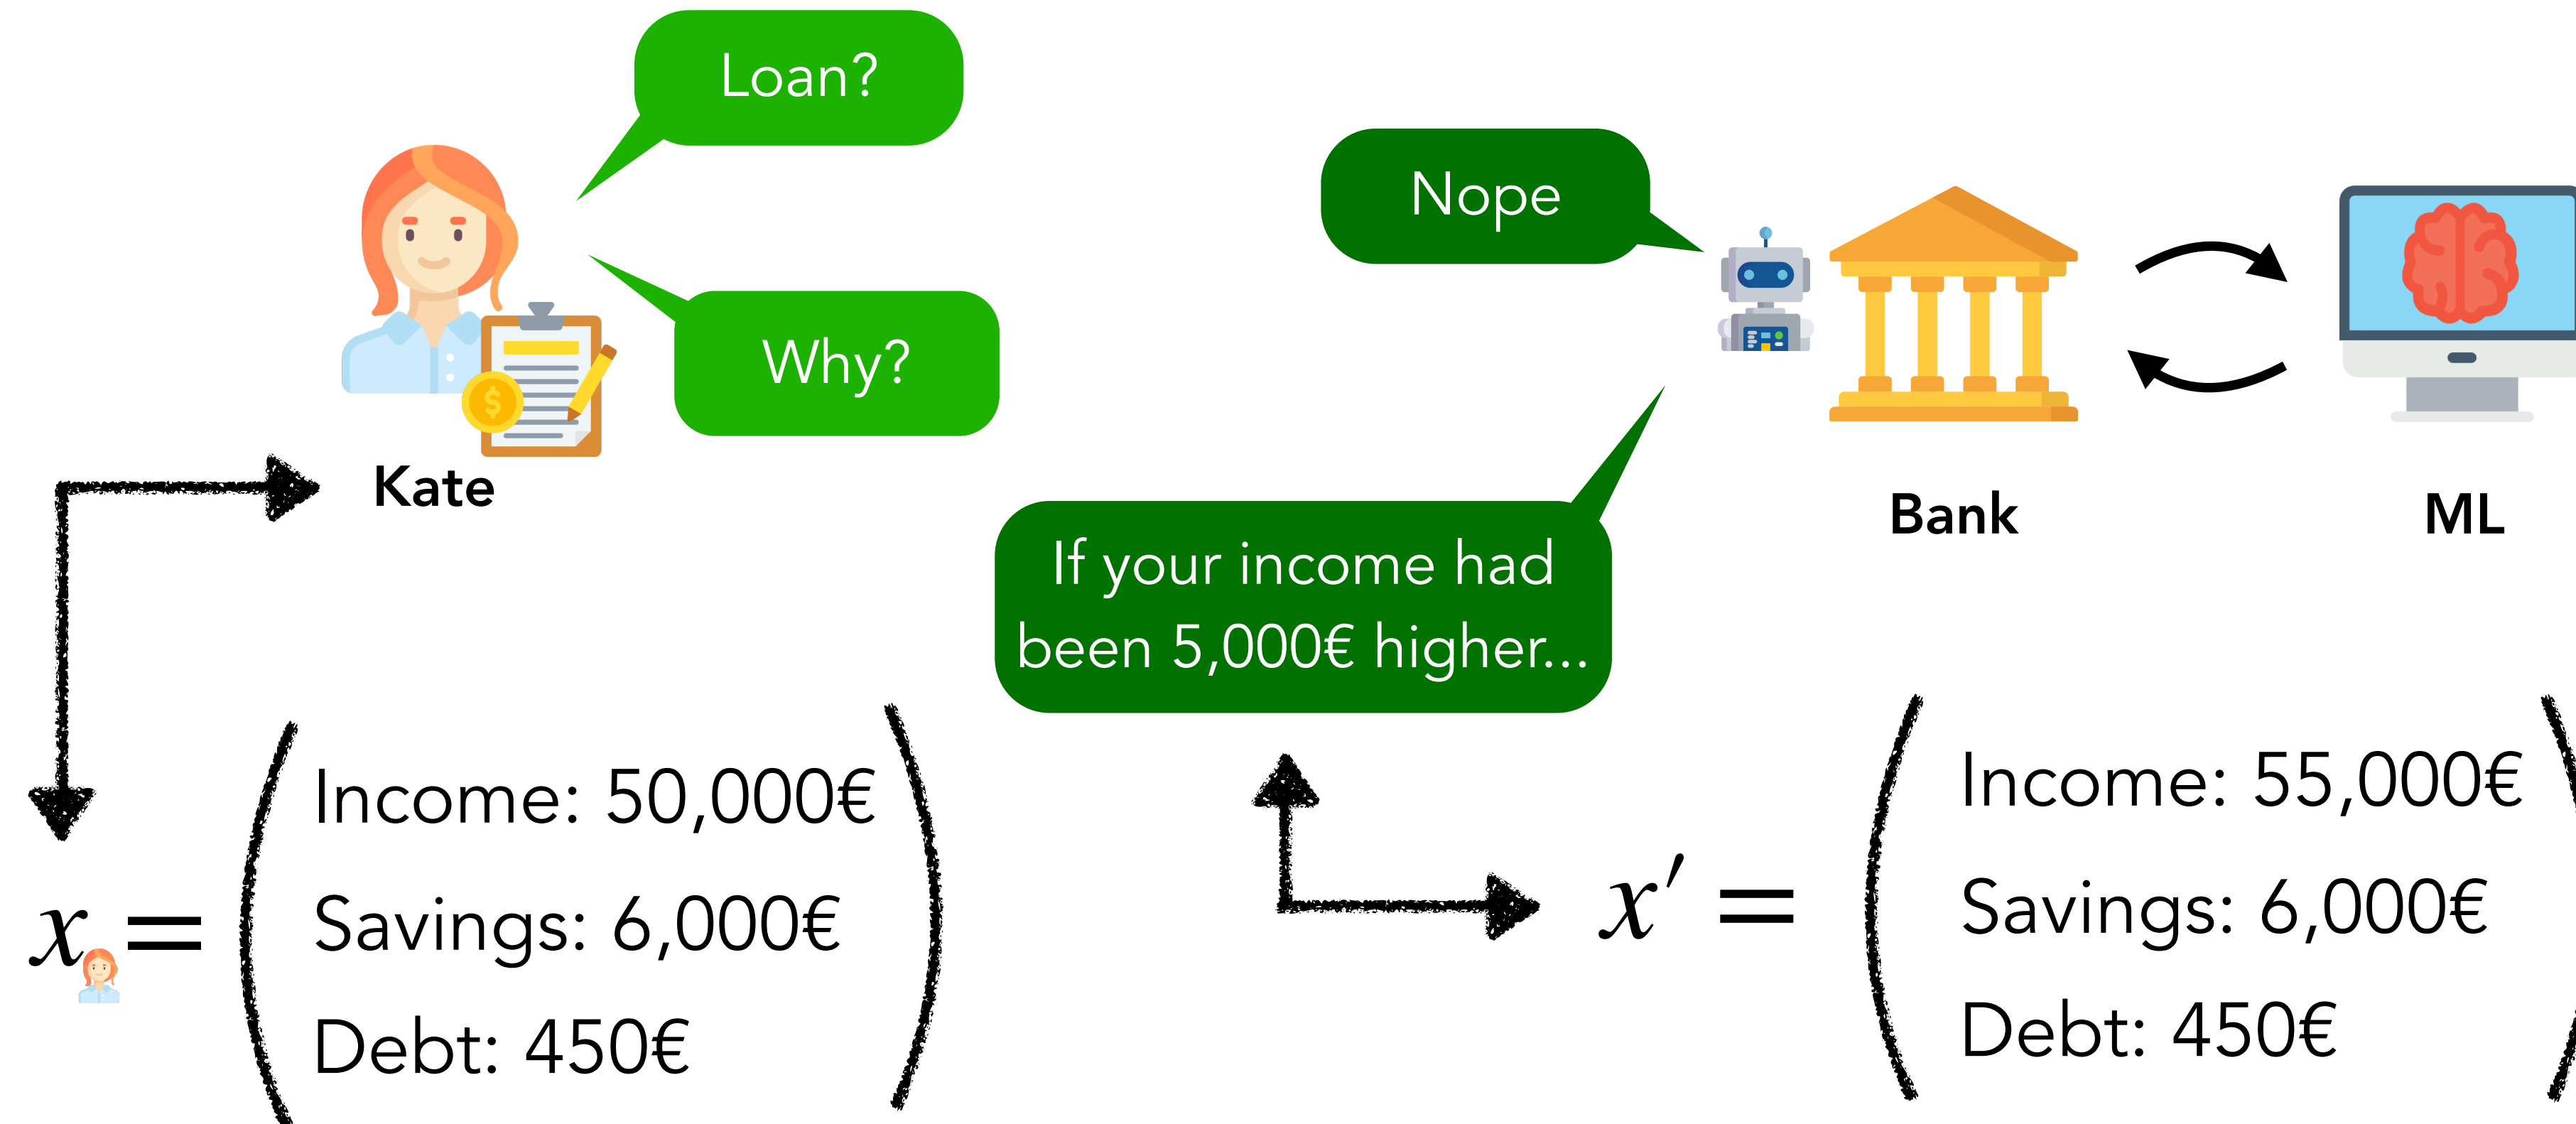
# Counterfactual explanations ignore causal dependencies



Beckers. "Causal explanations and xai." CLeaR, 2022.

Crupi et al. "Counterfactual explanations as interventions in latent space." DMKD, 2022.

# Counterfactual explanations ignore causal dependencies



 **If Kate's income had been 5,000€ higher, Kate's savings would have been more than 6,000€!**

Beckers. "Causal explanations and xai." CLeaR, 2022.

Crupi et al. "Counterfactual explanations as interventions in latent space." DMKD, 2022.

# Counterfactual explanations as interventions

A counterfactual explanation does not answer a counterfactual question but an interventional question.



# Counterfactual explanations as interventions

A counterfactual explanation does not answer a counterfactual question but an interventional question.

## Structural Causal Model $\mathcal{M}$

$$X_1 := f_{X_1}(D)$$

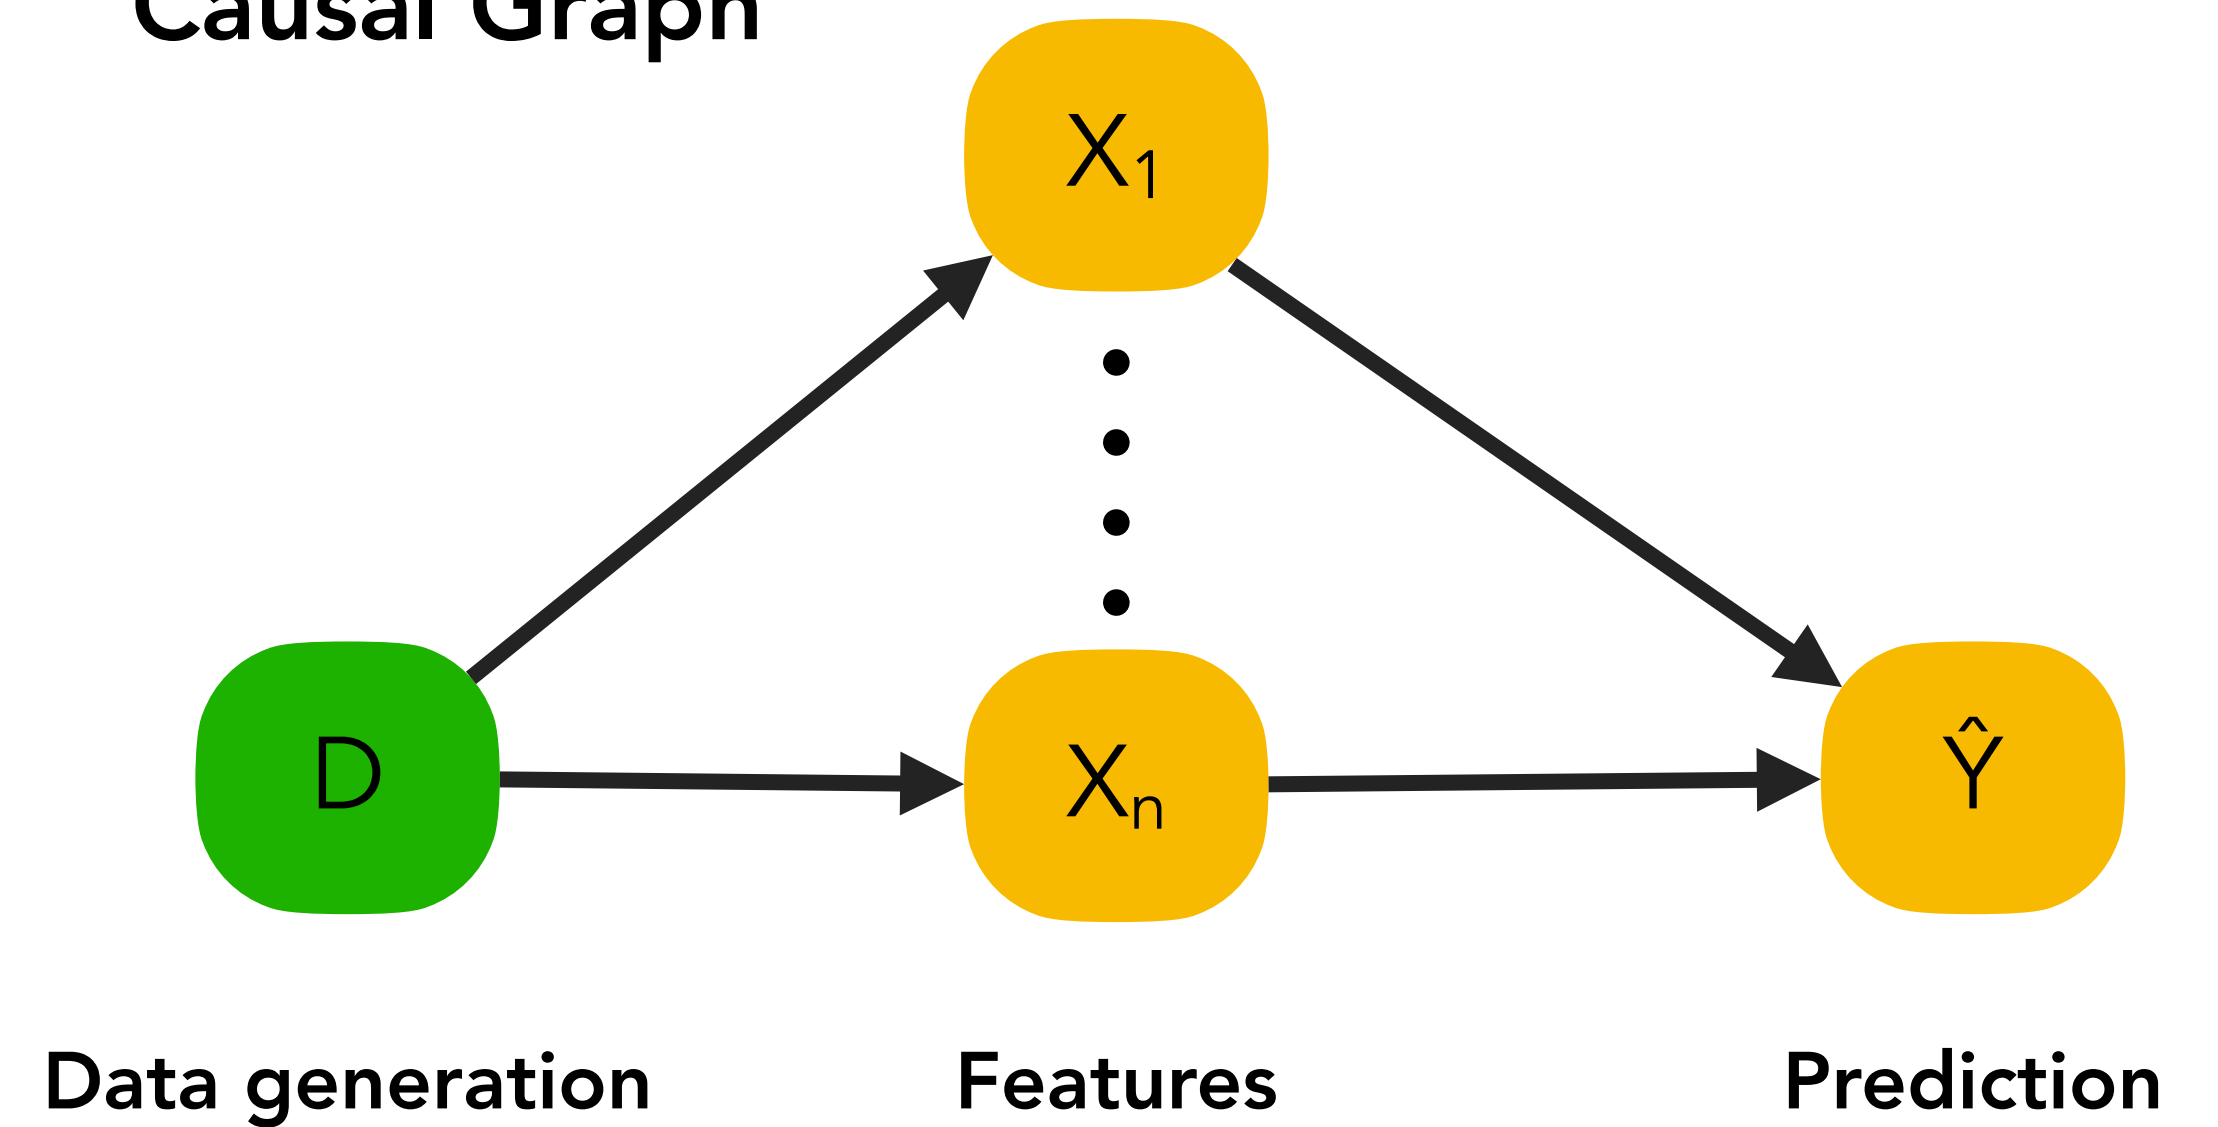
$$\vdots$$

$$X_n := f_{X_n}(D)$$

$$\hat{Y} := h(X)$$

$$D \sim P(D)$$

## Causal Graph



# Counterfactual explanations as interventions

A counterfactual explanation does not answer a counterfactual question but an interventional question.

## Structural Causal Model $\mathcal{M}$

$$\cancel{X_1 := s_{X_1}(D)} \quad X_1 := x'_1$$

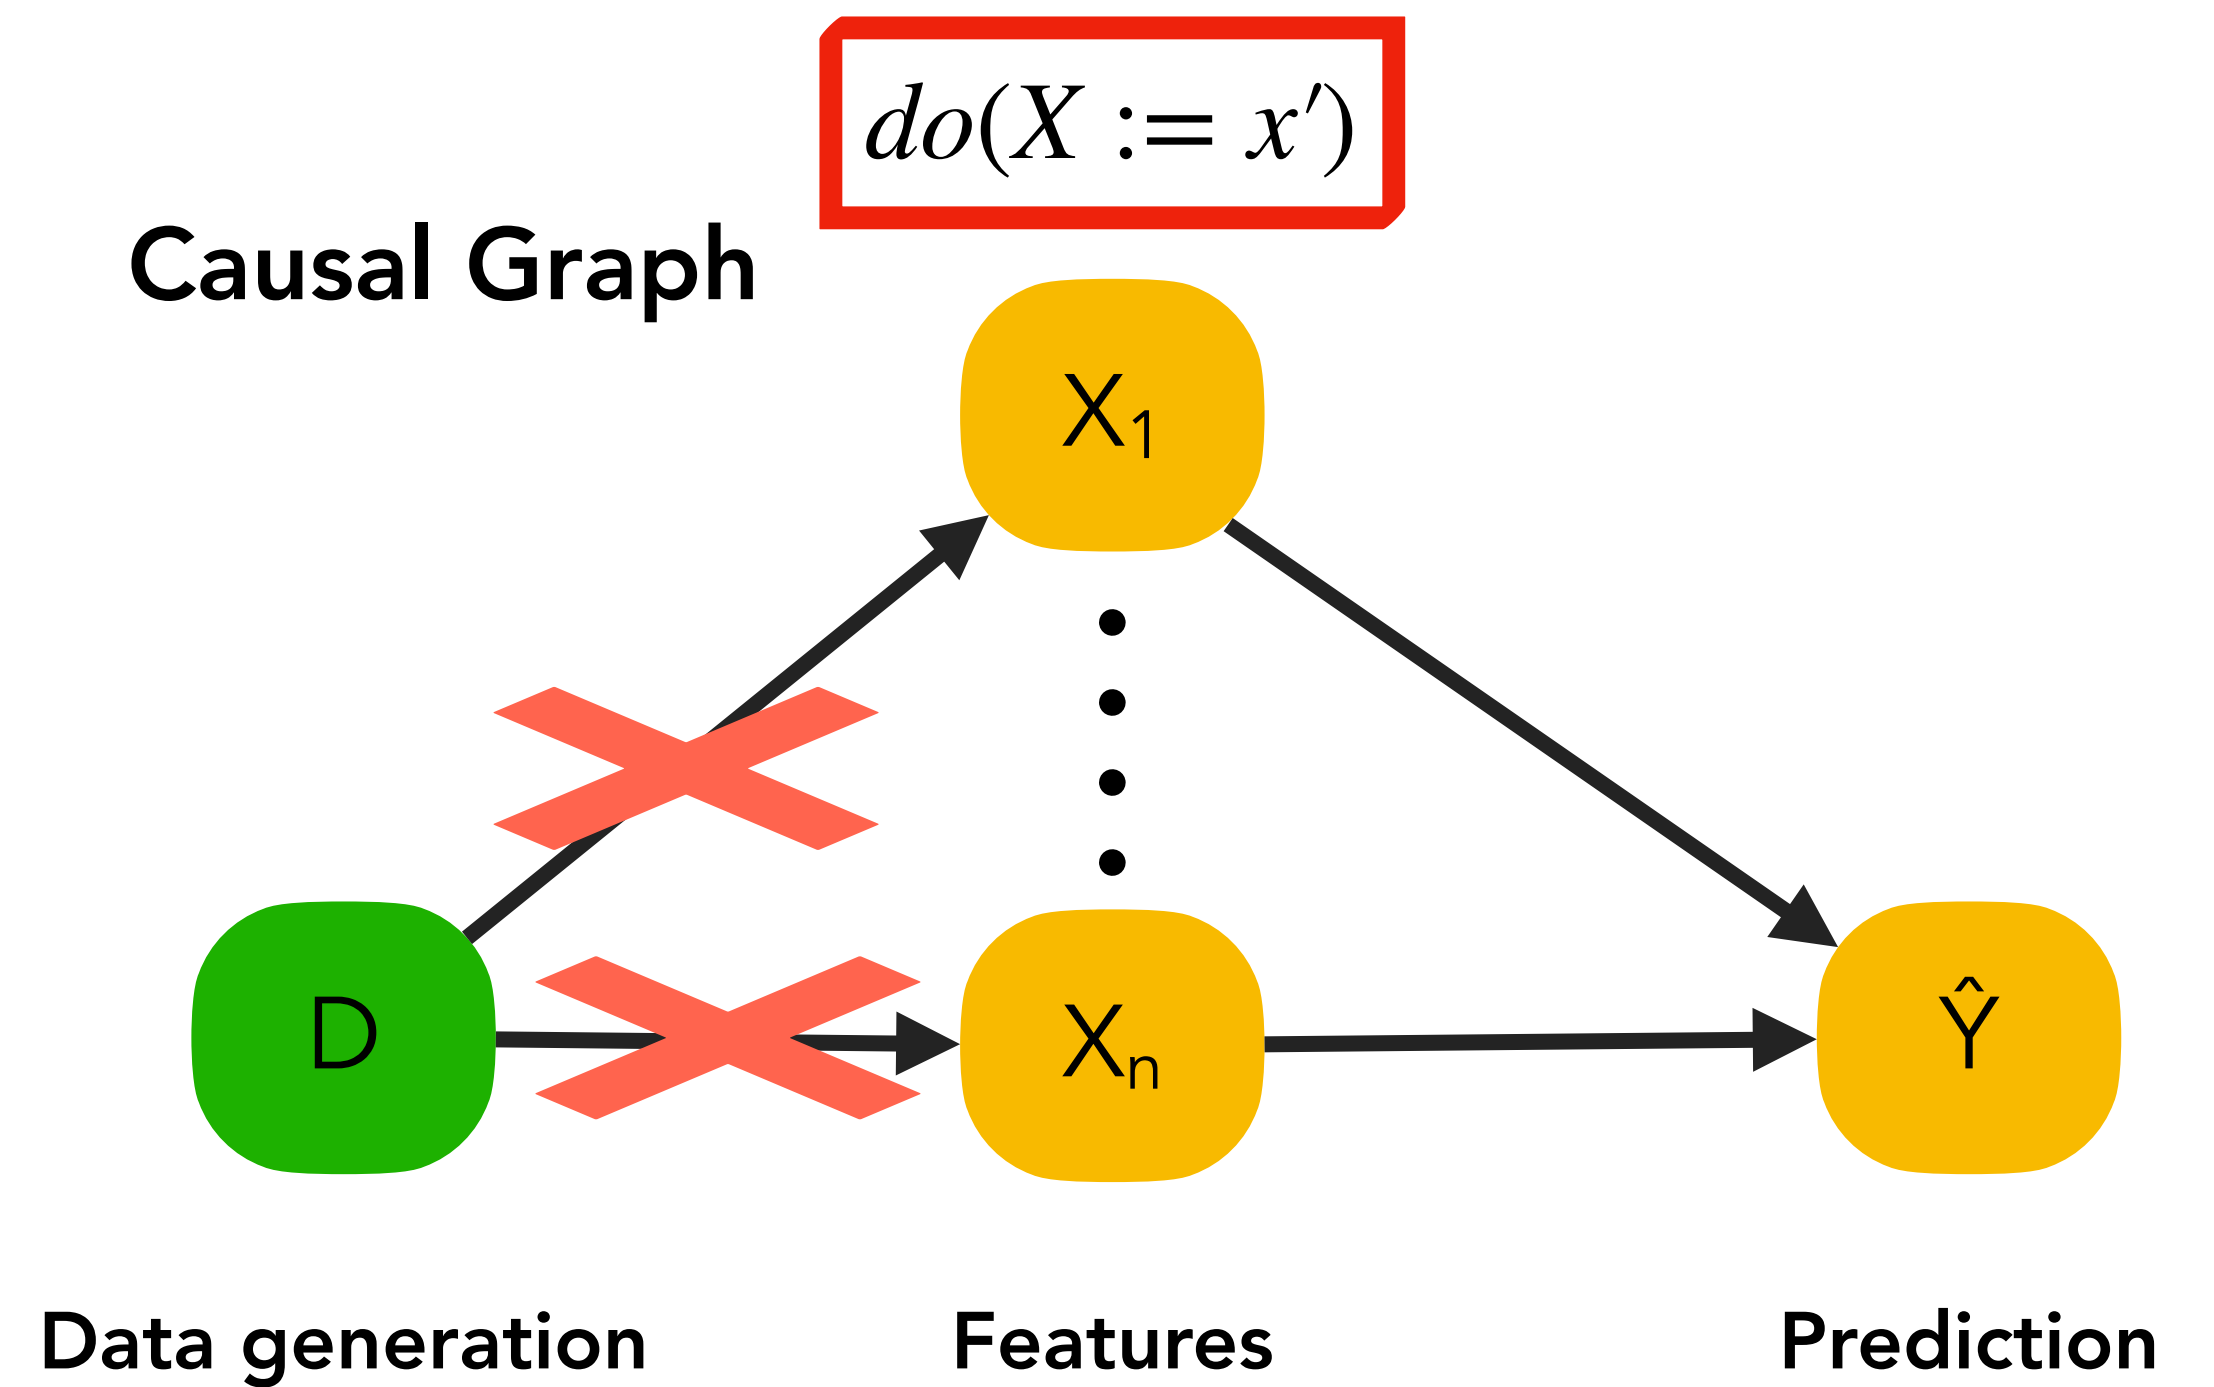
$$\vdots$$

$$\cancel{X_n := s_{X_n}(D)} \quad X_n := x'_n$$

$$\hat{Y} := h(X)$$

$$D \sim P(D)$$

## Causal Graph



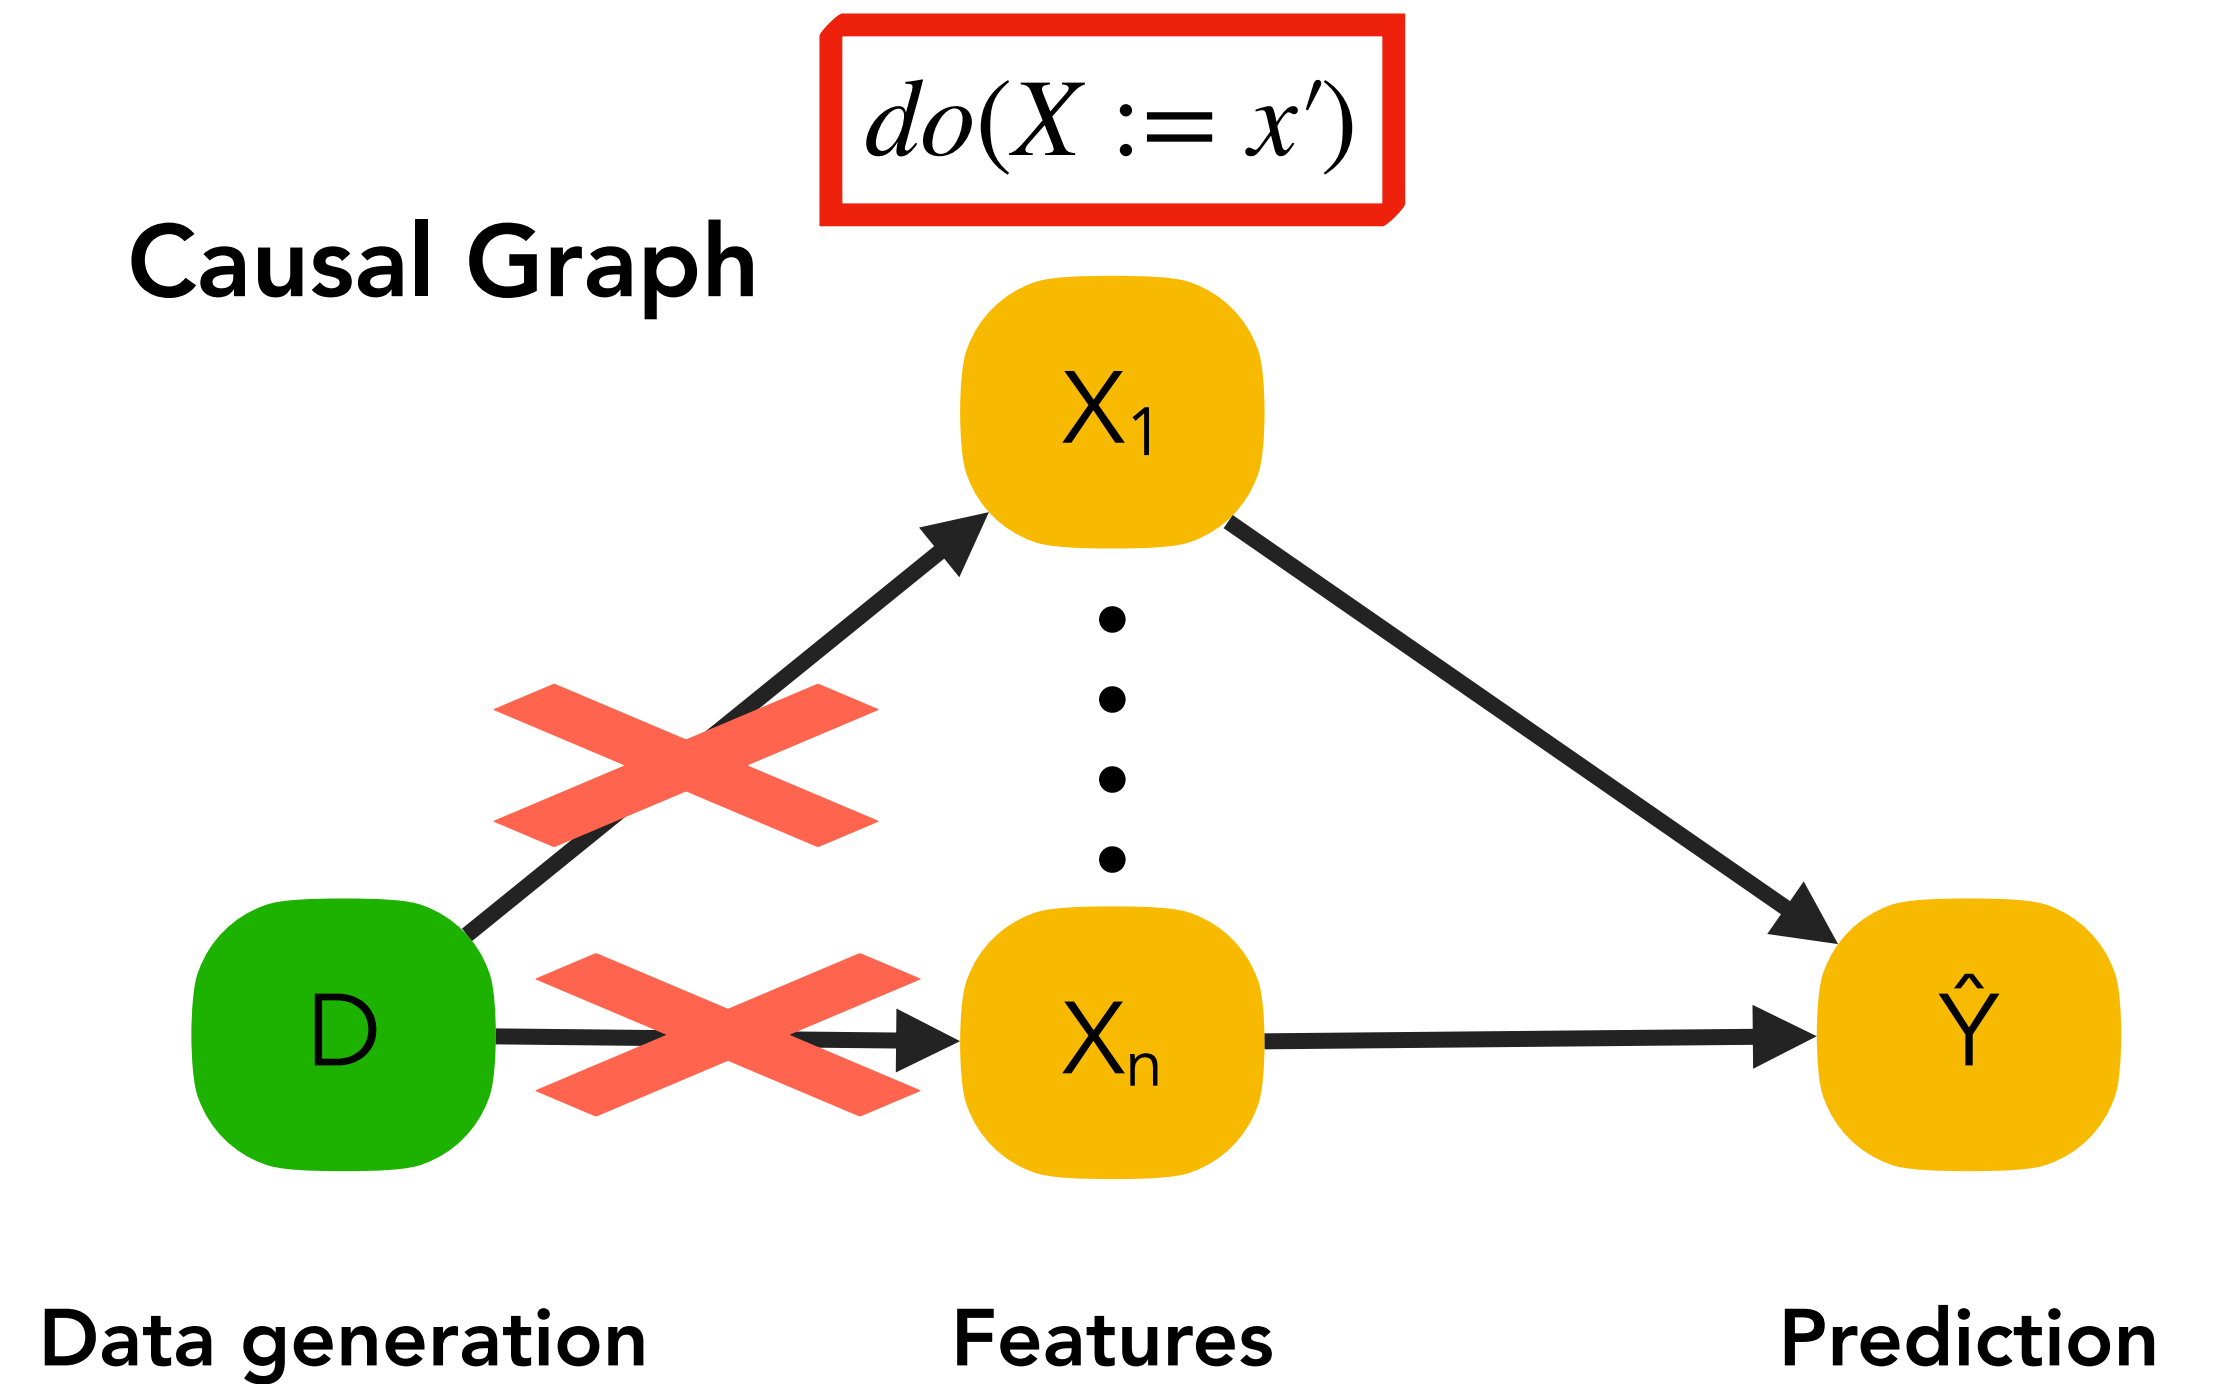
# Counterfactual explanations as interventions

A counterfactual explanation does not answer a counterfactual question but an interventional question.

## Structural Causal Model $\mathcal{M}$

$$\begin{array}{ll} \cancel{X_1 \leftarrow s_1(D)} & X_1 := x'_1 \\ \vdots & \vdots \\ \cancel{X_n \leftarrow s_n(D)} & X_n := x'_n \\ \hat{Y} := h(X) & \\ D \sim P(D) & \end{array}$$

## Causal Graph



A counterfactual explanation encourages an individual to change the value of the features  $x_l$  such that  $x_l \neq x'_l$ . However, it does not take into account that such a change may induce changes in features  $x_l$  such that  $x_l = x'_l$ .

# Algorithmic recourse

Algorithmic recourse seeks to find the minimal intervention  $a$  under which  $h(x + a) \neq h(x)$  while accounting for causal dependencies between features.

Karimi et al. "Algorithmic recourse: from counterfactual explanations to interventions." FAccT, 2021.

Karimi et al. "Algorithmic recourse under imperfect causal knowledge: a probabilistic approach." NeurIPS, 2020.

# Algorithmic recourse

Algorithmic recourse seeks to find the minimal intervention  $a$  under which  $h(x + a) \neq h(x)$  while accounting for causal dependencies between features.

## Structural Causal Model $\mathcal{M}$

$$X_1 := f_{X_1}(D)$$

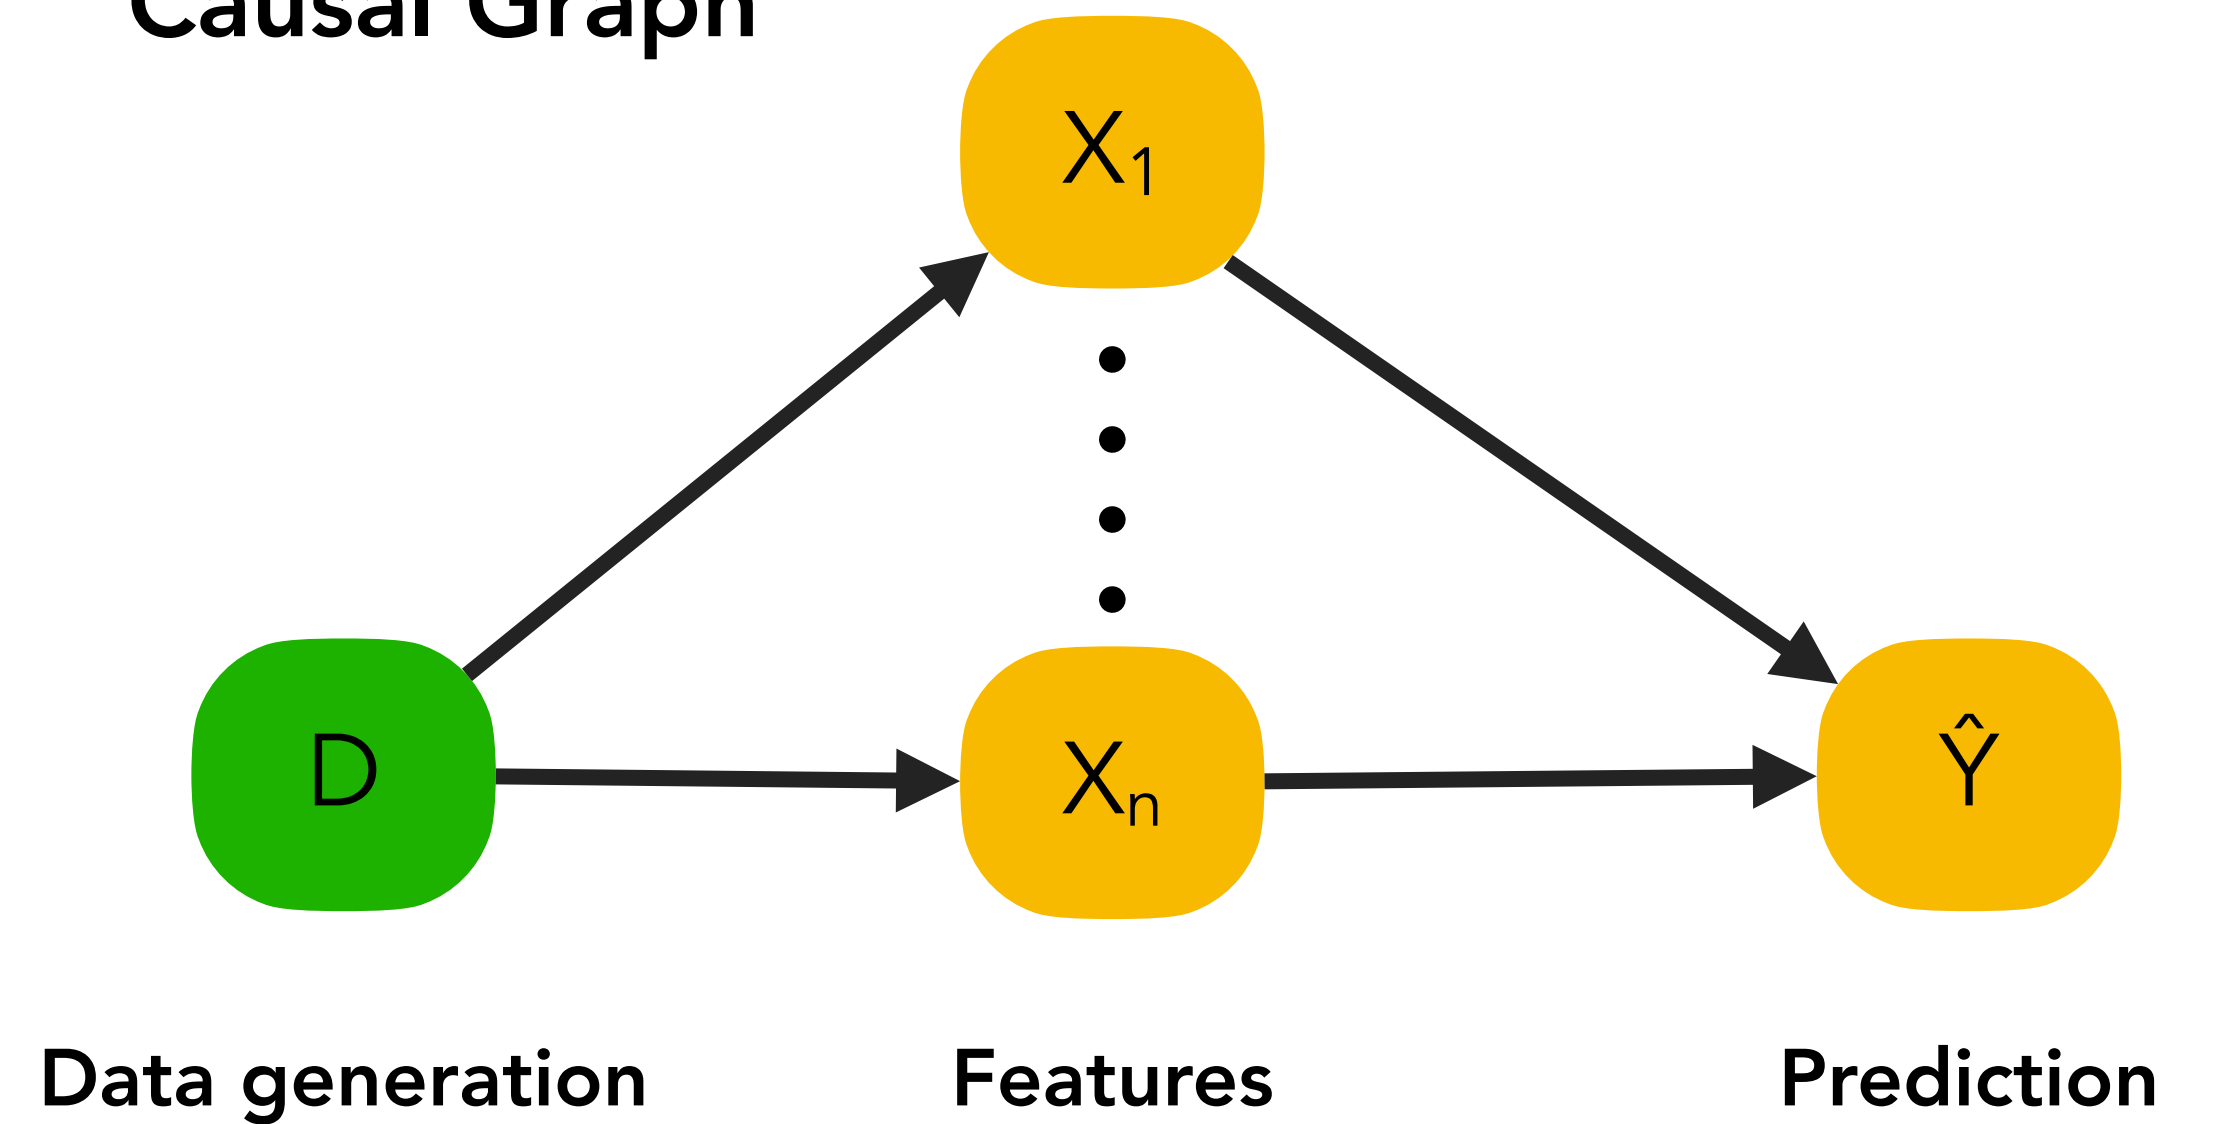
$\vdots$

$$X_n := f_{X_n}(D)$$

$$\hat{Y} := h(X)$$

$$D \sim P(D)$$

## Causal Graph



Karimi et al. "Algorithmic recourse: from counterfactual explanations to interventions." FAccT, 2021.

Karimi et al. "Algorithmic recourse under imperfect causal knowledge: a probabilistic approach." NeurIPS, 2020.

# Algorithmic recourse

Algorithmic recourse seeks to find the minimal intervention  $a$  under which  $h(x + a) \neq h(x)$  while accounting for causal dependencies between features.

**Modified Structural Causal Model  $\mathcal{M}_{X=x}$**

$$X_1 := f_{X_1}(D)$$

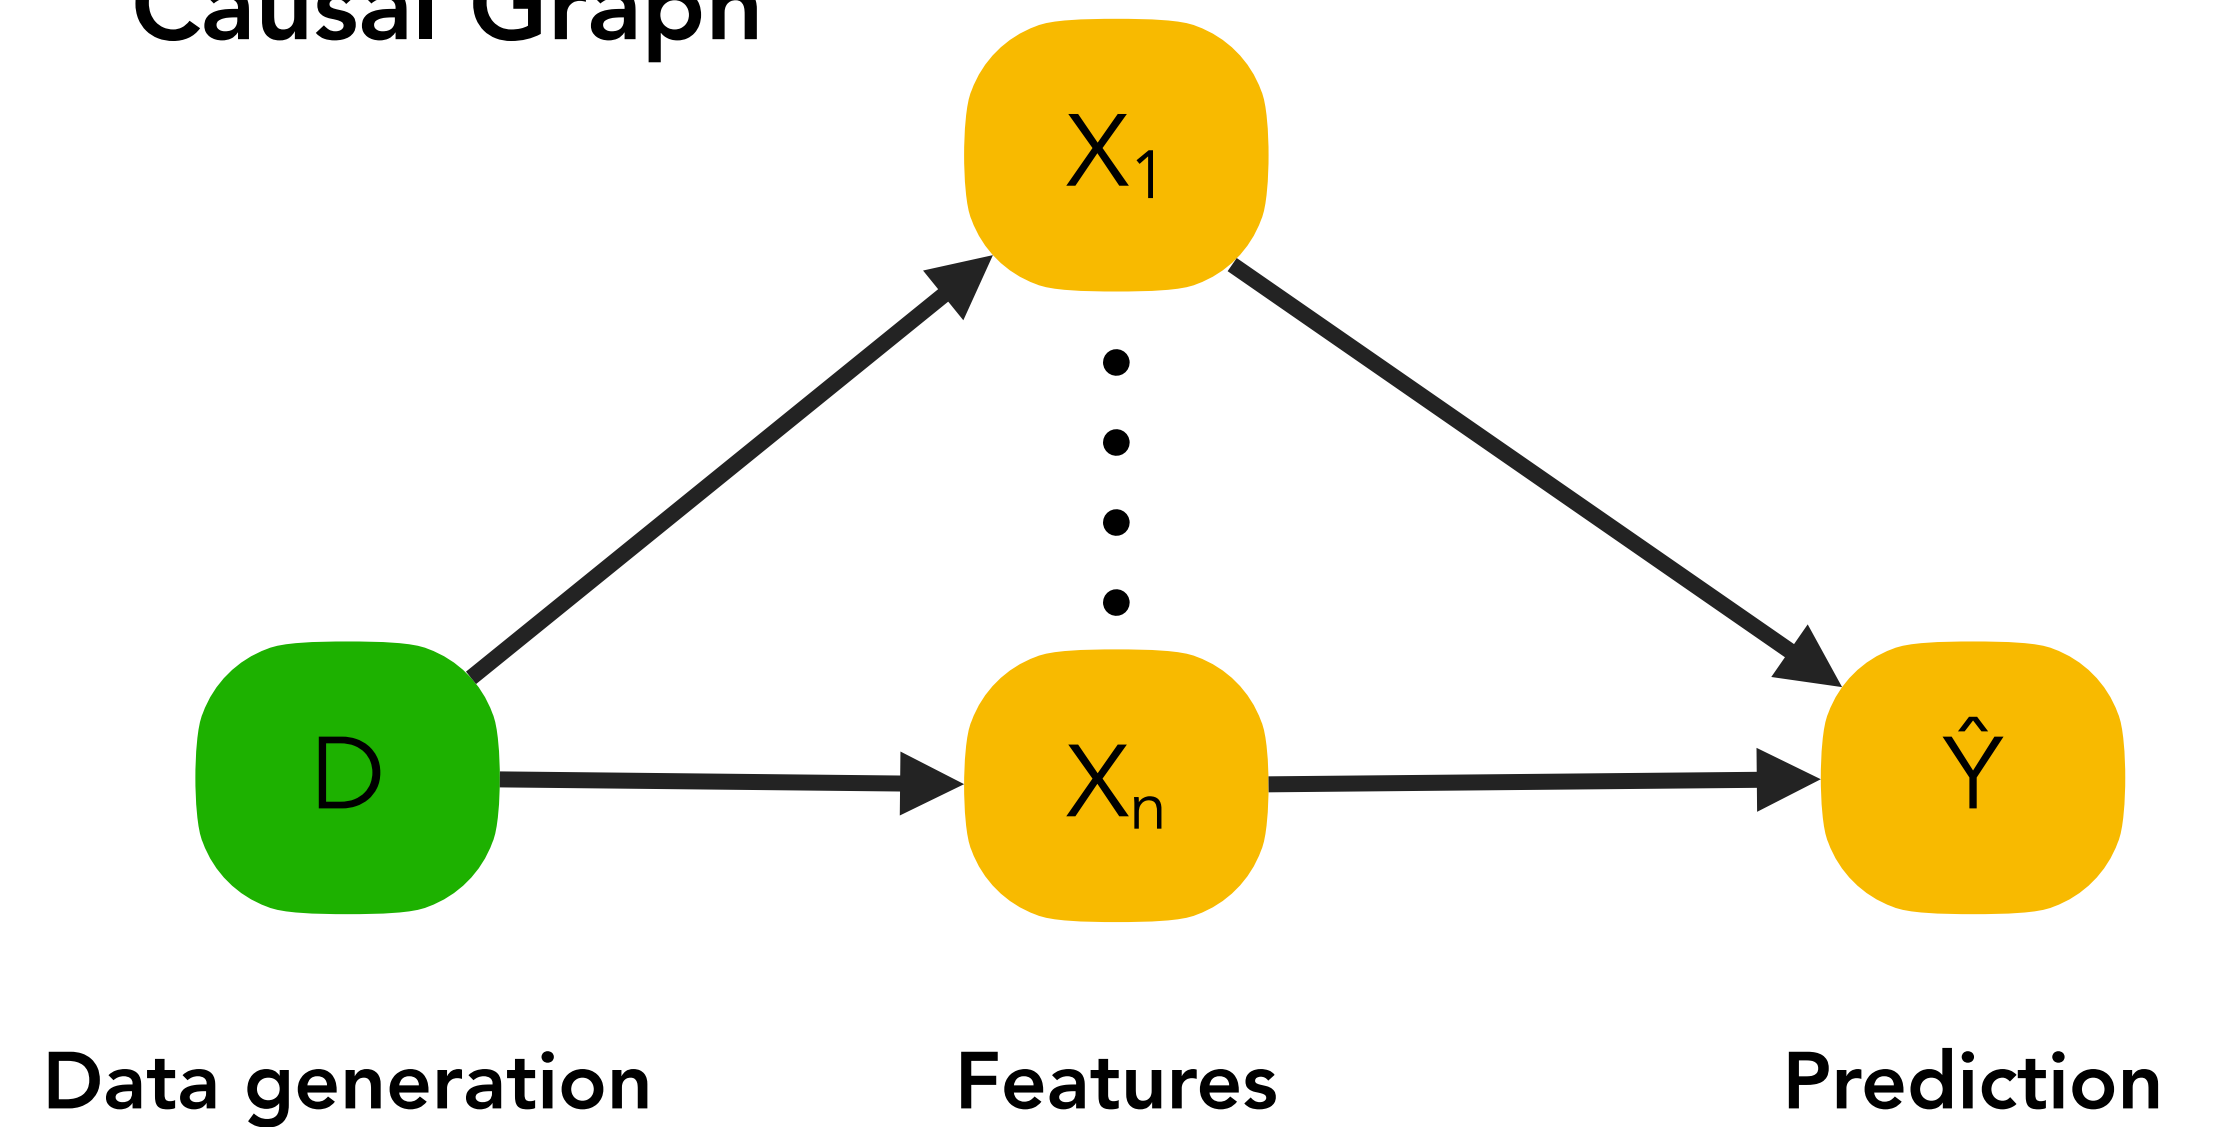
$\vdots$

$$X_n := f_{X_n}(D)$$

$$\hat{Y} := h(X)$$

$$D \sim P(D \mid X = x) \quad \leftarrow \text{Posterior distribution of the noise}$$

**Causal Graph**



Karimi et al. "Algorithmic recourse: from counterfactual explanations to interventions." FAccT, 2021.

Karimi et al. "Algorithmic recourse under imperfect causal knowledge: a probabilistic approach." NeurIPS, 2020.

# Algorithmic recourse

Algorithmic recourse seeks to find the minimal intervention  $a$  under which  $h(x + a) \neq h(x)$  while accounting for causal dependencies between features.

**Modified Structural Causal Model**  $\mathcal{M}_{X=x}$

$$X_1 := f_{X_1}(D) + a_1$$

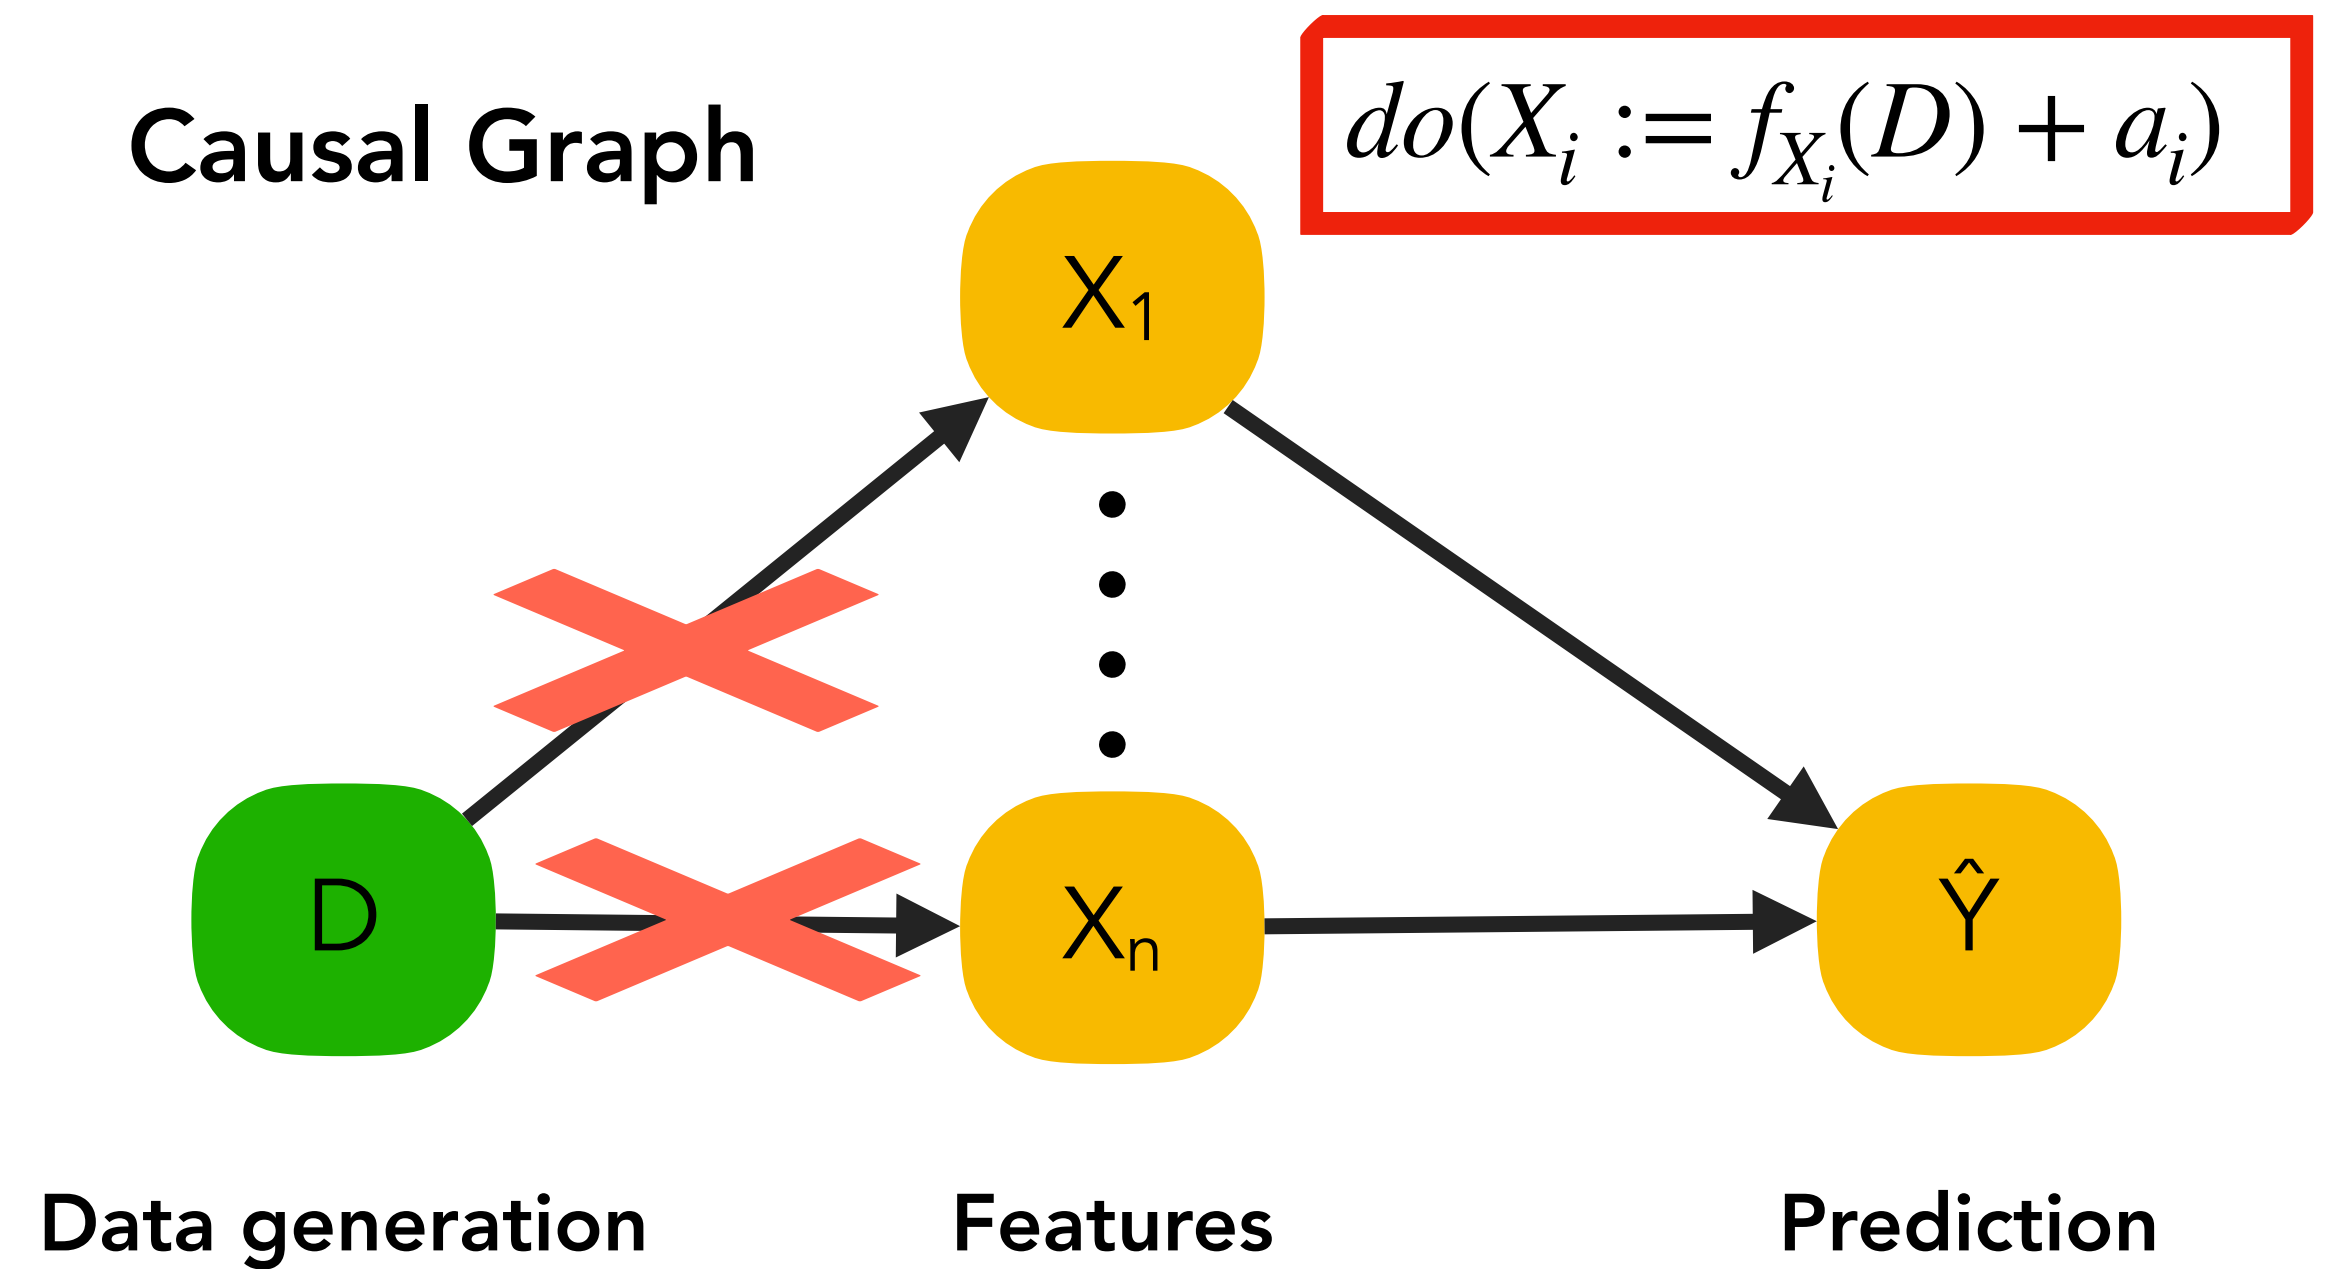
$\vdots$

$$X_n := f_{X_n}(D) + a_n$$

$$\hat{Y} := h(X)$$

$$D \sim P(D \mid X = x) \quad \leftarrow \text{Posterior distribution of the noise}$$

**Causal Graph**



Whenever  $a_i = 0$ , the value of  $X_i$  may still change!

Karimi et al. "Algorithmic recourse: from counterfactual explanations to interventions." FAccT, 2021.

Karimi et al. "Algorithmic recourse under imperfect causal knowledge: a probabilistic approach." NeurIPS, 2020.

# Counterfactual explanations & performativity

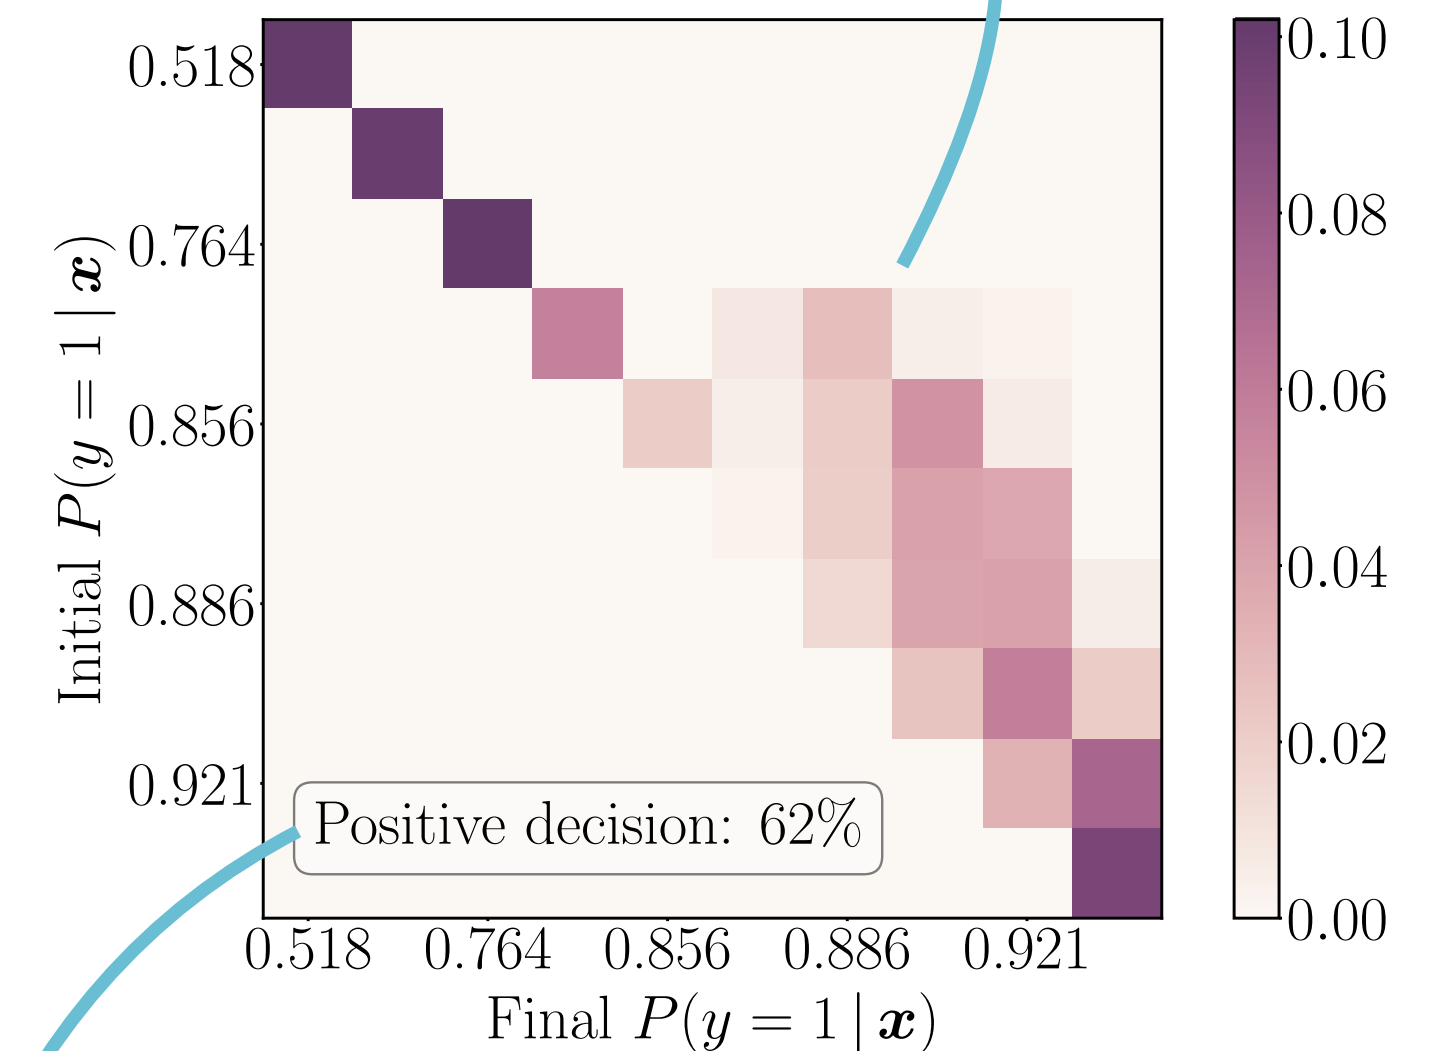
If a sizable number of individuals follow the changes prescribed by counterfactual explanations, the feature distribution  $P(X)$  may change.



# Counterfactual explanations & performativity

If a sizable number of individuals follow the changes prescribed by counterfactual explanations, the feature distribution  $P(X)$  may change.

Chances of repayment would improve  
for large part of the population



More people would  
receive credit

# Counterfactual explanations & performativity

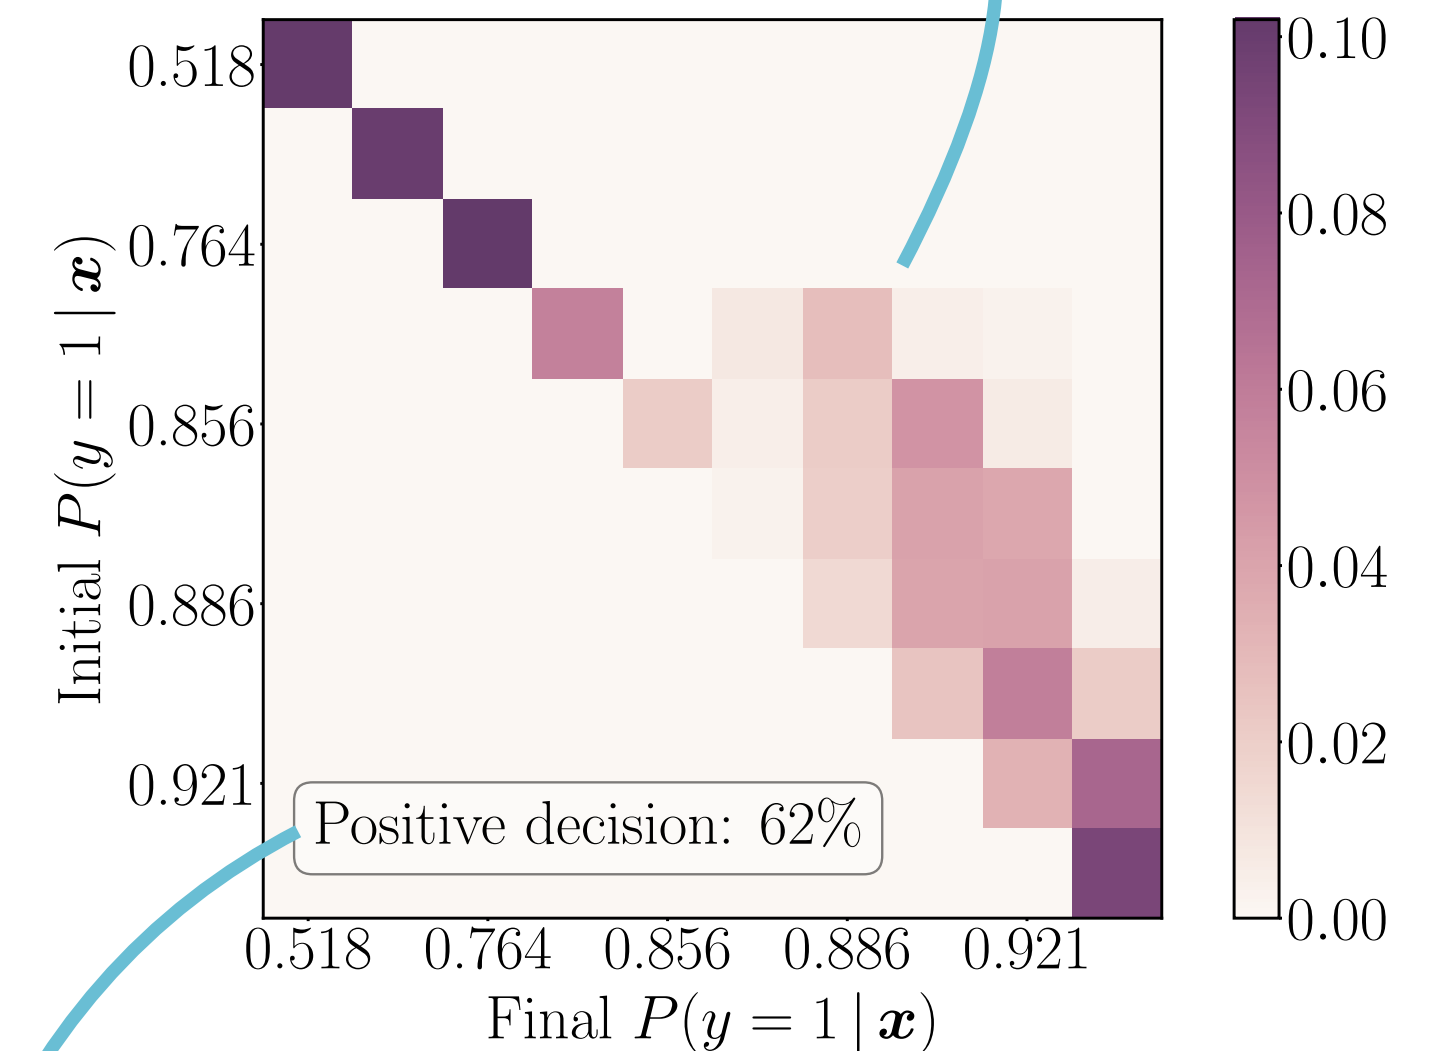
If a sizable number of individuals follow the changes prescribed by counterfactual explanations, the feature distribution  $P(X)$  may change.

*This raises the question of finding decision policies  $\pi$  and counterfactual explanations  $\mathcal{A}$  that are optimal in terms of utility.*

$$\max_{\pi, \mathcal{A}} u(\pi, \mathcal{A}) := \mathbb{E}_{x \sim P(X | \pi, \mathcal{A})} \left[ \pi(x) (P(Y = 1 | x) - \gamma) \right]$$

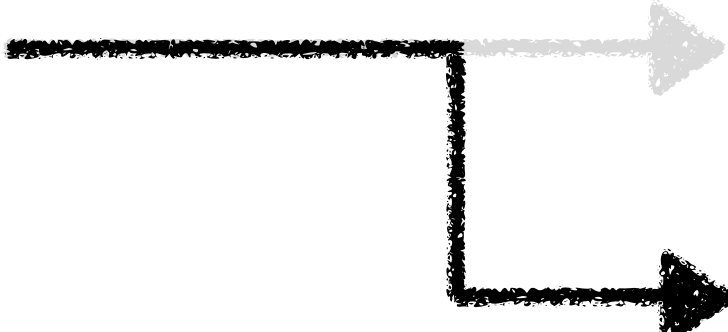
↑  
constant reflecting economic considerations  
of the decision maker

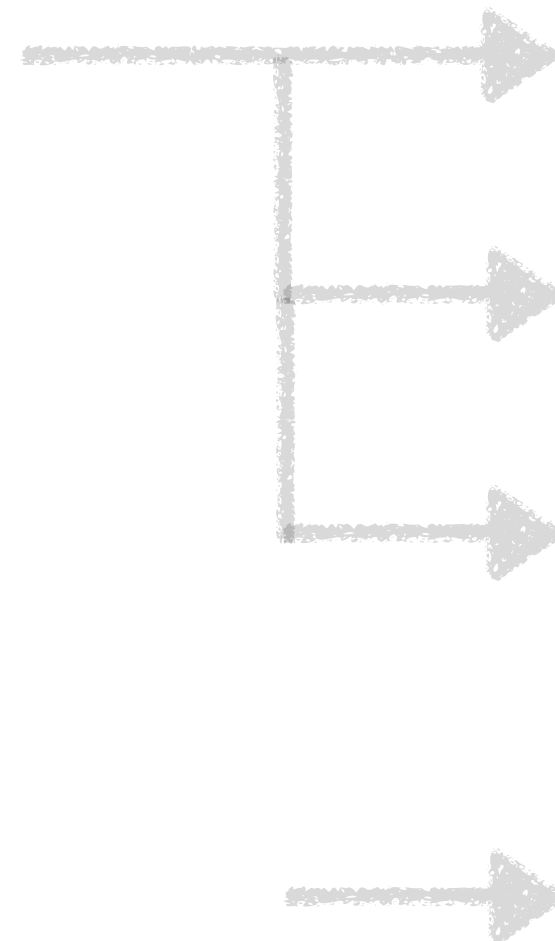
**Chances of repayment would improve  
for large part of the population**



**More people would  
receive credit**

# Use cases of counterfactuals in machine learning

*Classification*  *Fairness*



# Counterfactual fairness

Counterfactual fairness captures the intuition that a prediction by a machine learning model is fair towards an individual who belongs to a demographic group  $A = a$  if it would have been the same had the individual belonged to a different demographic group  $A = a'$ .

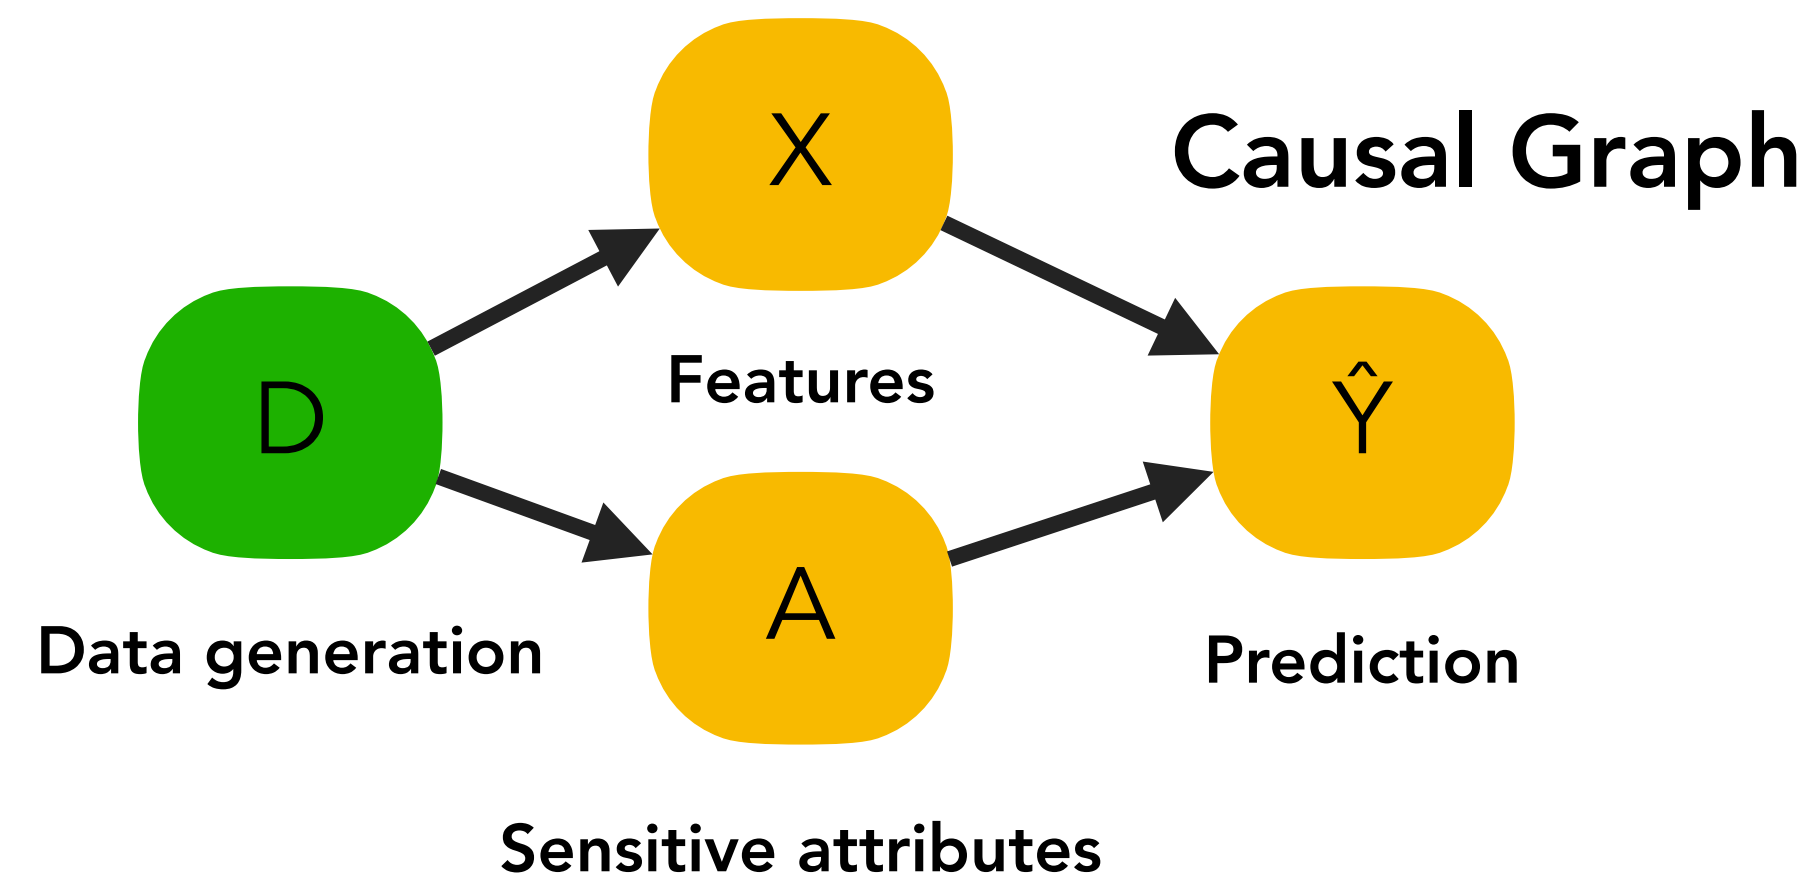
# Counterfactual fairness

Counterfactual fairness captures the intuition that a prediction by a machine learning model is fair towards an individual who belongs to a demographic group  $A = a$  if it would have been the same had the individual belonged to a different demographic group  $A = a'$ .

## Structural Causal Model $\mathcal{M}$

$$X := f_X(D) \quad \hat{Y} := h(X, A)$$

$$A := f_A(D) \quad D \sim P(D)$$



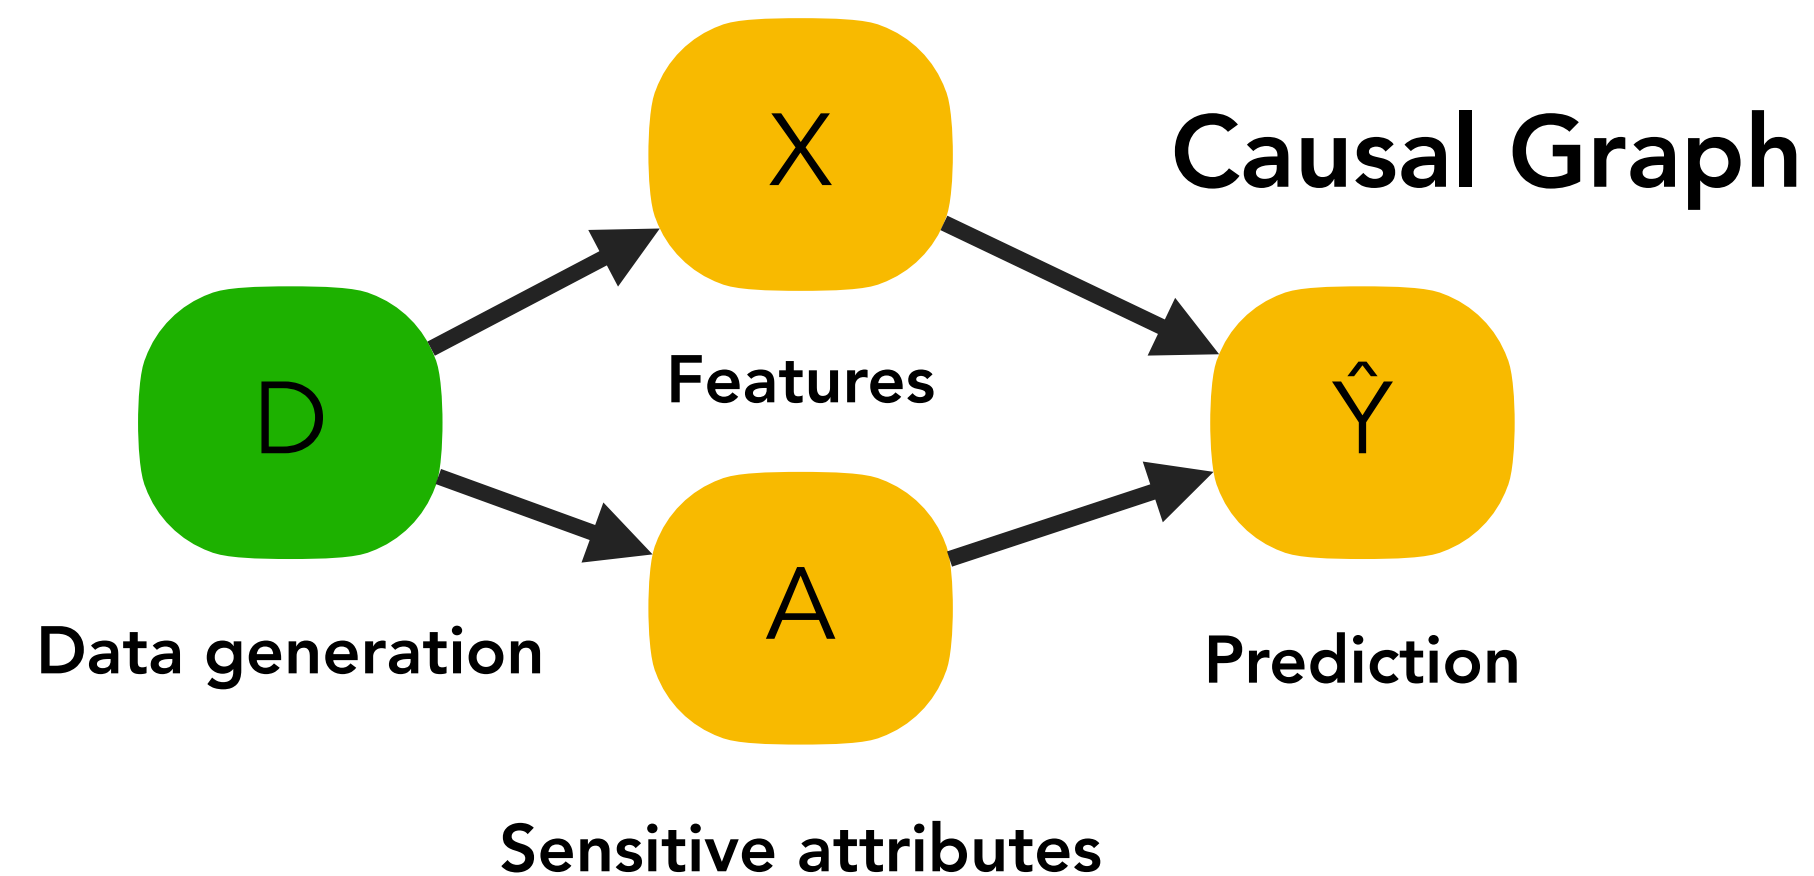
# Counterfactual fairness

Counterfactual fairness captures the intuition that a prediction by a machine learning model is fair towards an individual who belongs to a demographic group  $A = a$  if it would have been the same had the individual belonged to a different demographic group  $A = a'$

## Structural Causal Model $\mathcal{M}$

$$X := f_X(D) \quad \hat{Y} := h(X, A)$$

$$A := f_A(D) \quad D \sim P(D)$$



## Counterfactual fairness

$$P^{\mathcal{M}} | X=x, A=a; do(A=a')(\hat{Y}) = P^{\mathcal{M}} | X=x, A=a(\hat{Y})$$

The diagram shows two circular icons representing individuals: a male (orange) and a female (green). Dashed arrows point from each icon to the  $A=a$  term in the equation. A third circular icon (orange) has a dashed arrow pointing to the  $do(A=a')$  term.

# Counterfactual fairness can be too restrictive

Counterfactual fairness considers the full effect of the demographic group on the prediction as problematic. However, this is not the case in certain scenarios.

# Counterfactual fairness can be too restrictive

Counterfactual fairness considers the full effect of the demographic group on the prediction as problematic. However, this is not the case in certain scenarios.

## ***Alleged gender bias case at Berkeley***



8,442 male applicants for the fall of 1973, 44 percent were admitted,



4,351 female applicants, 35 percent were admitted



# Counterfactual fairness can be too restrictive

Counterfactual fairness considers the full effect of the demographic group on the prediction as problematic. However, this is not the case in certain scenarios.

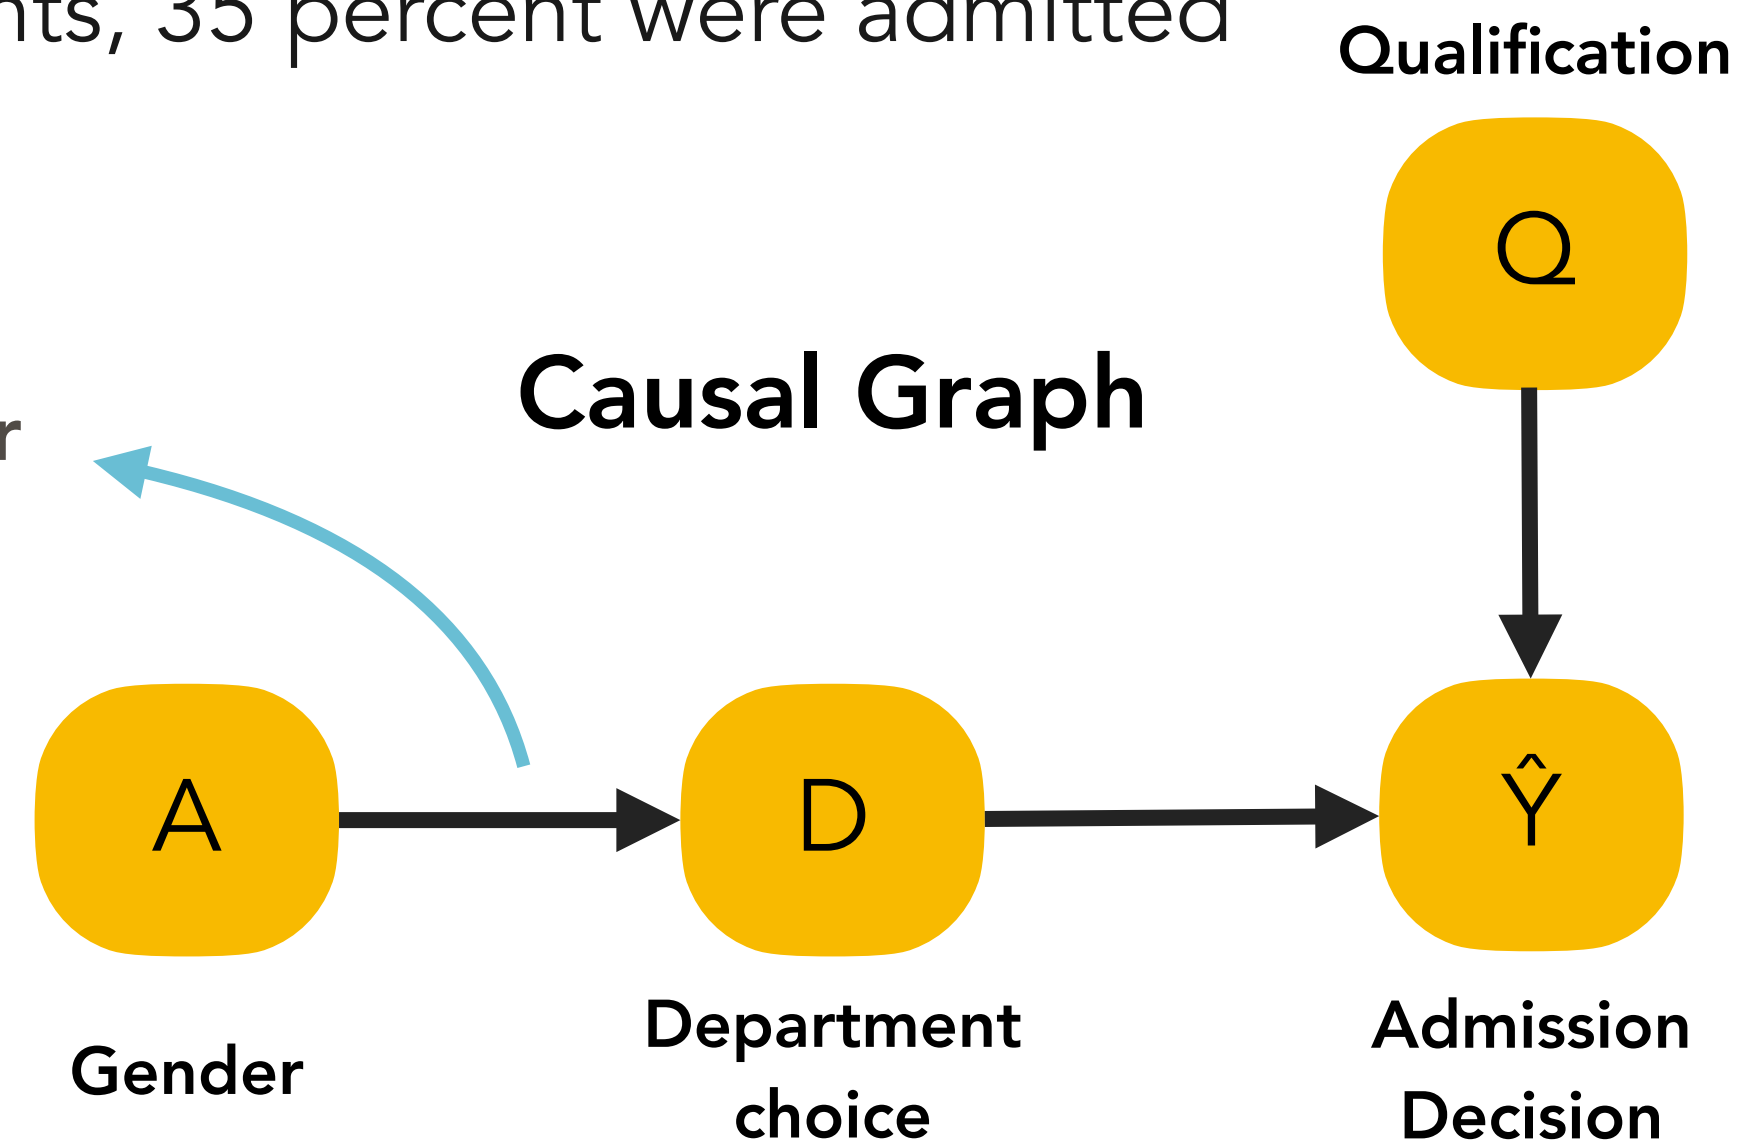
## *Alleged gender bias case at Berkeley*

 8,442 male applicants for the fall of 1973, 44 percent were admitted,

 4,351 female applicants, 35 percent were admitted

Female applied to  
departments with lower  
admission rates

### Causal Graph

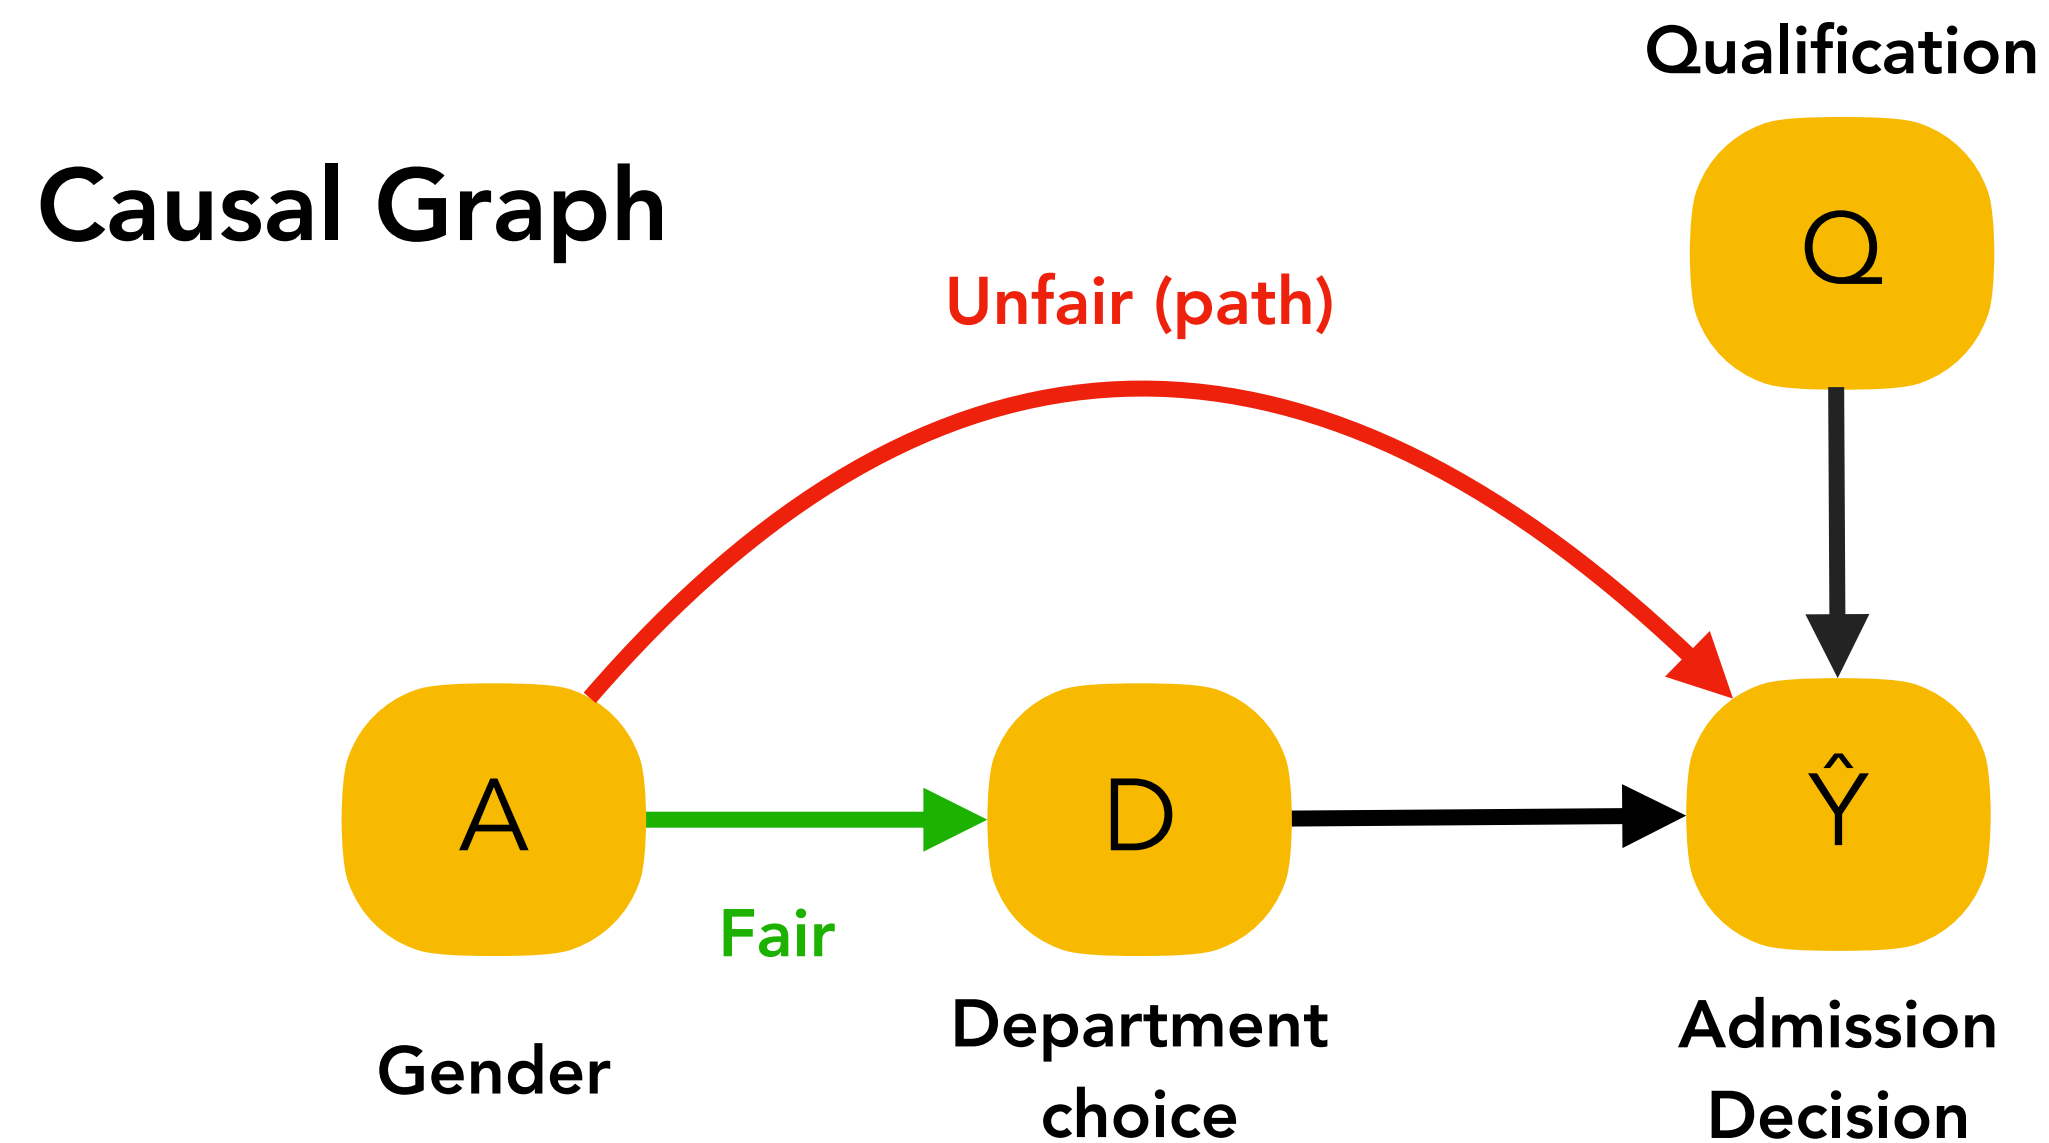


### Counterfactual fairness is violated

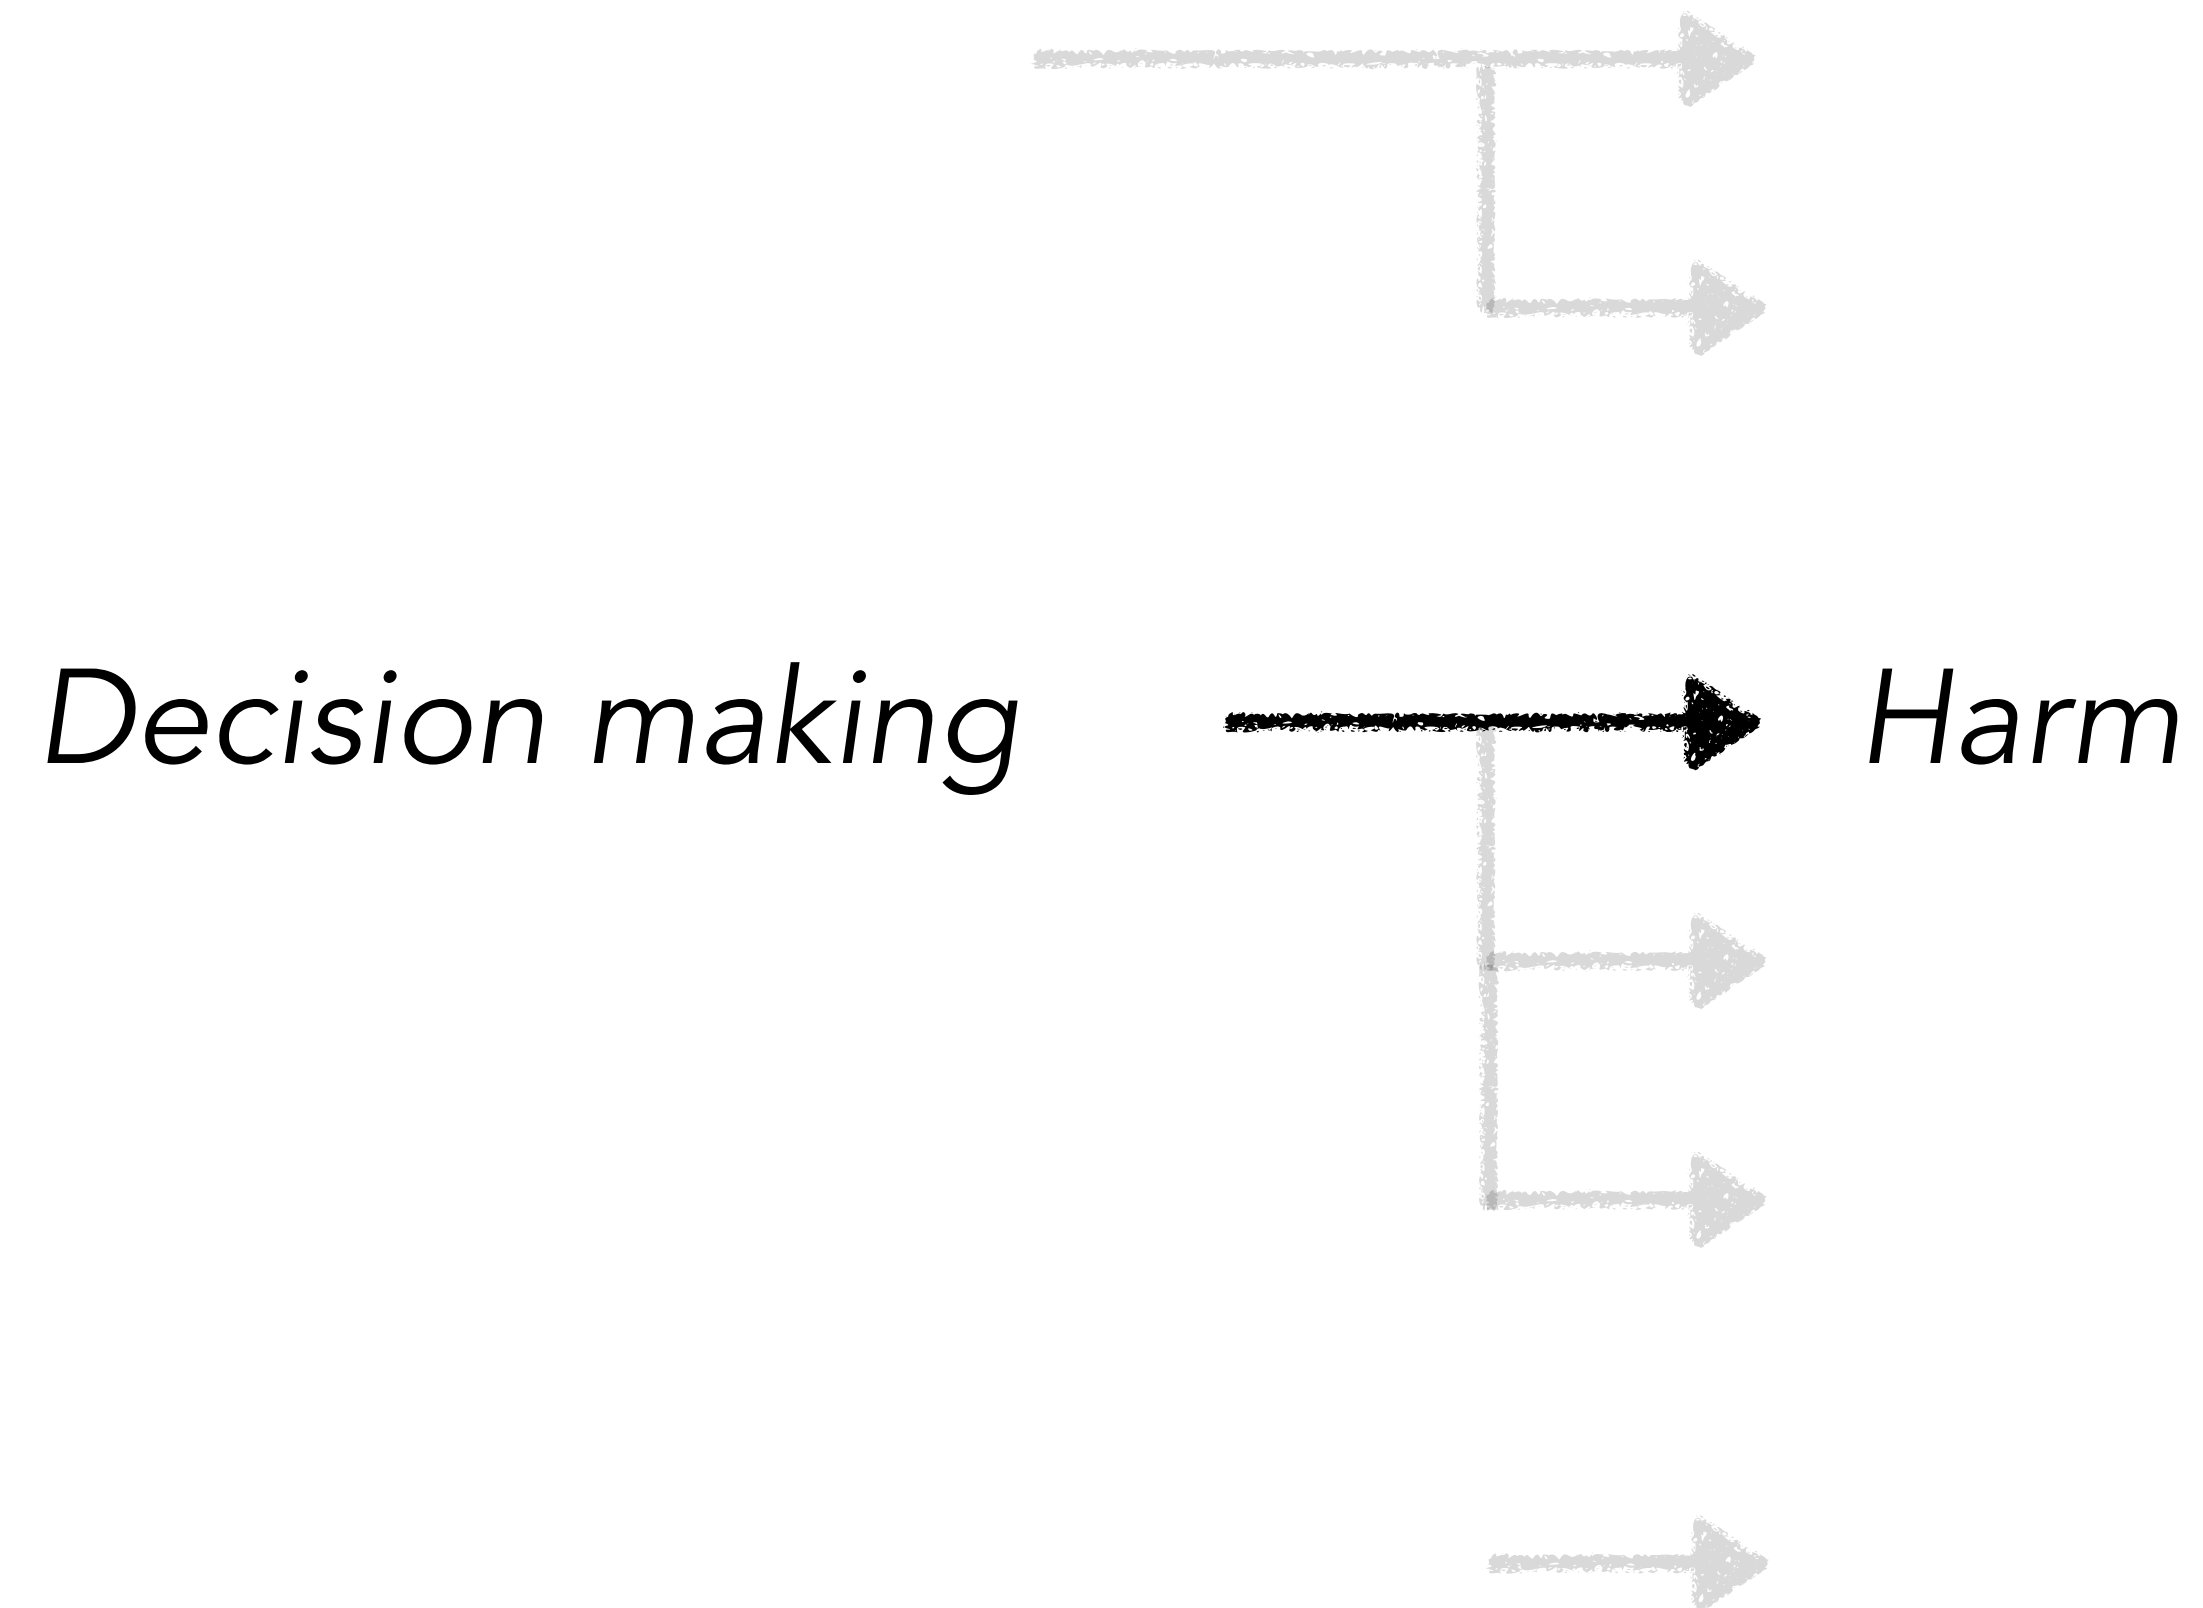
$$P^{\mathcal{M}} | Q=q, A=a; do(A=a')(\hat{Y}) \neq P^{\mathcal{M}} | Q=q, A=a(\hat{Y})$$

# Path-specific counterfactual fairness

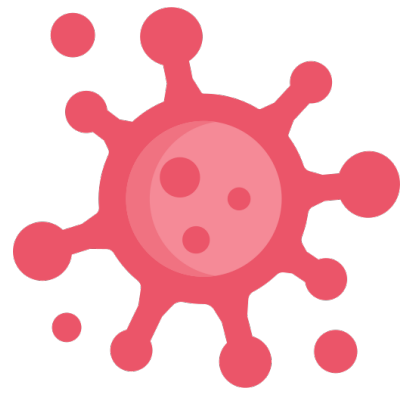
Path-specific counterfactual fairness is a more fine-grained fairness criterion that deals with sensitive attributes affecting the prediction along both fair and unfair pathways.



# Use cases of counterfactuals in machine learning

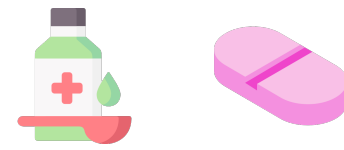


# Counterfactual harm



**Disease**

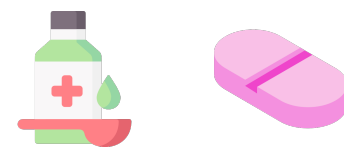
**50% mortality rate**



**Treatment A**

**60% chance of  
curing a patient**

**40% chance of  
having no effect**

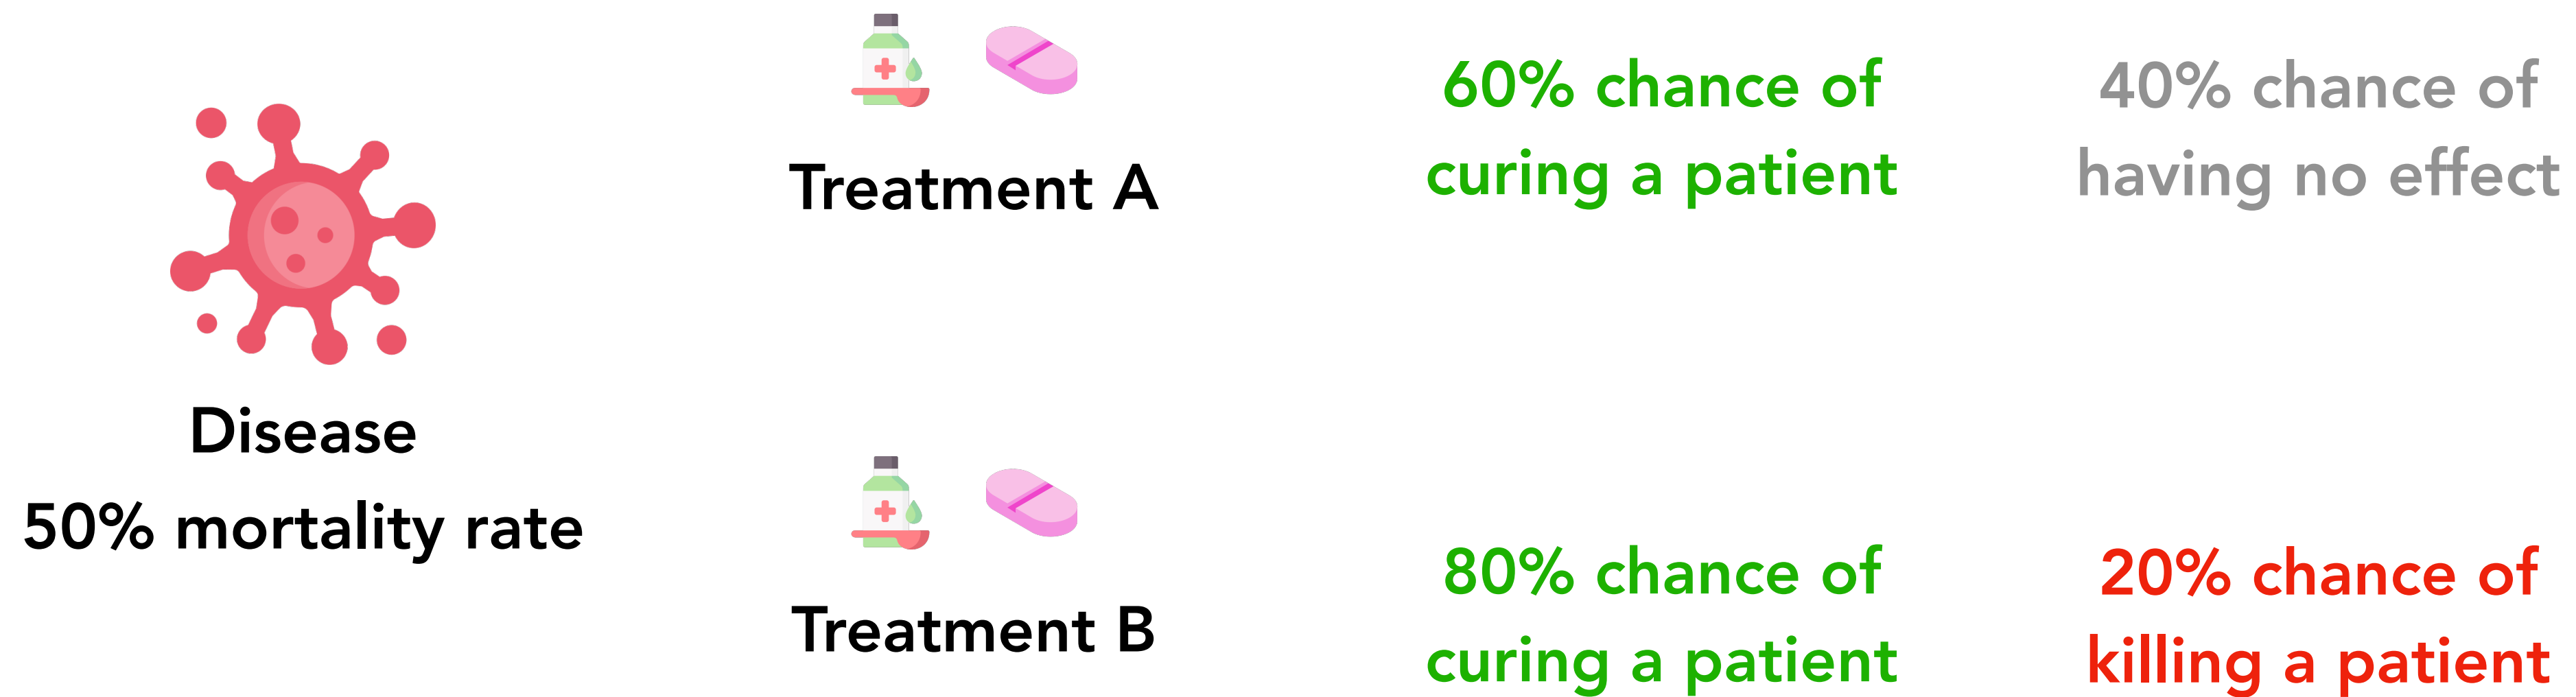


**Treatment B**

**80% chance of  
curing a patient**

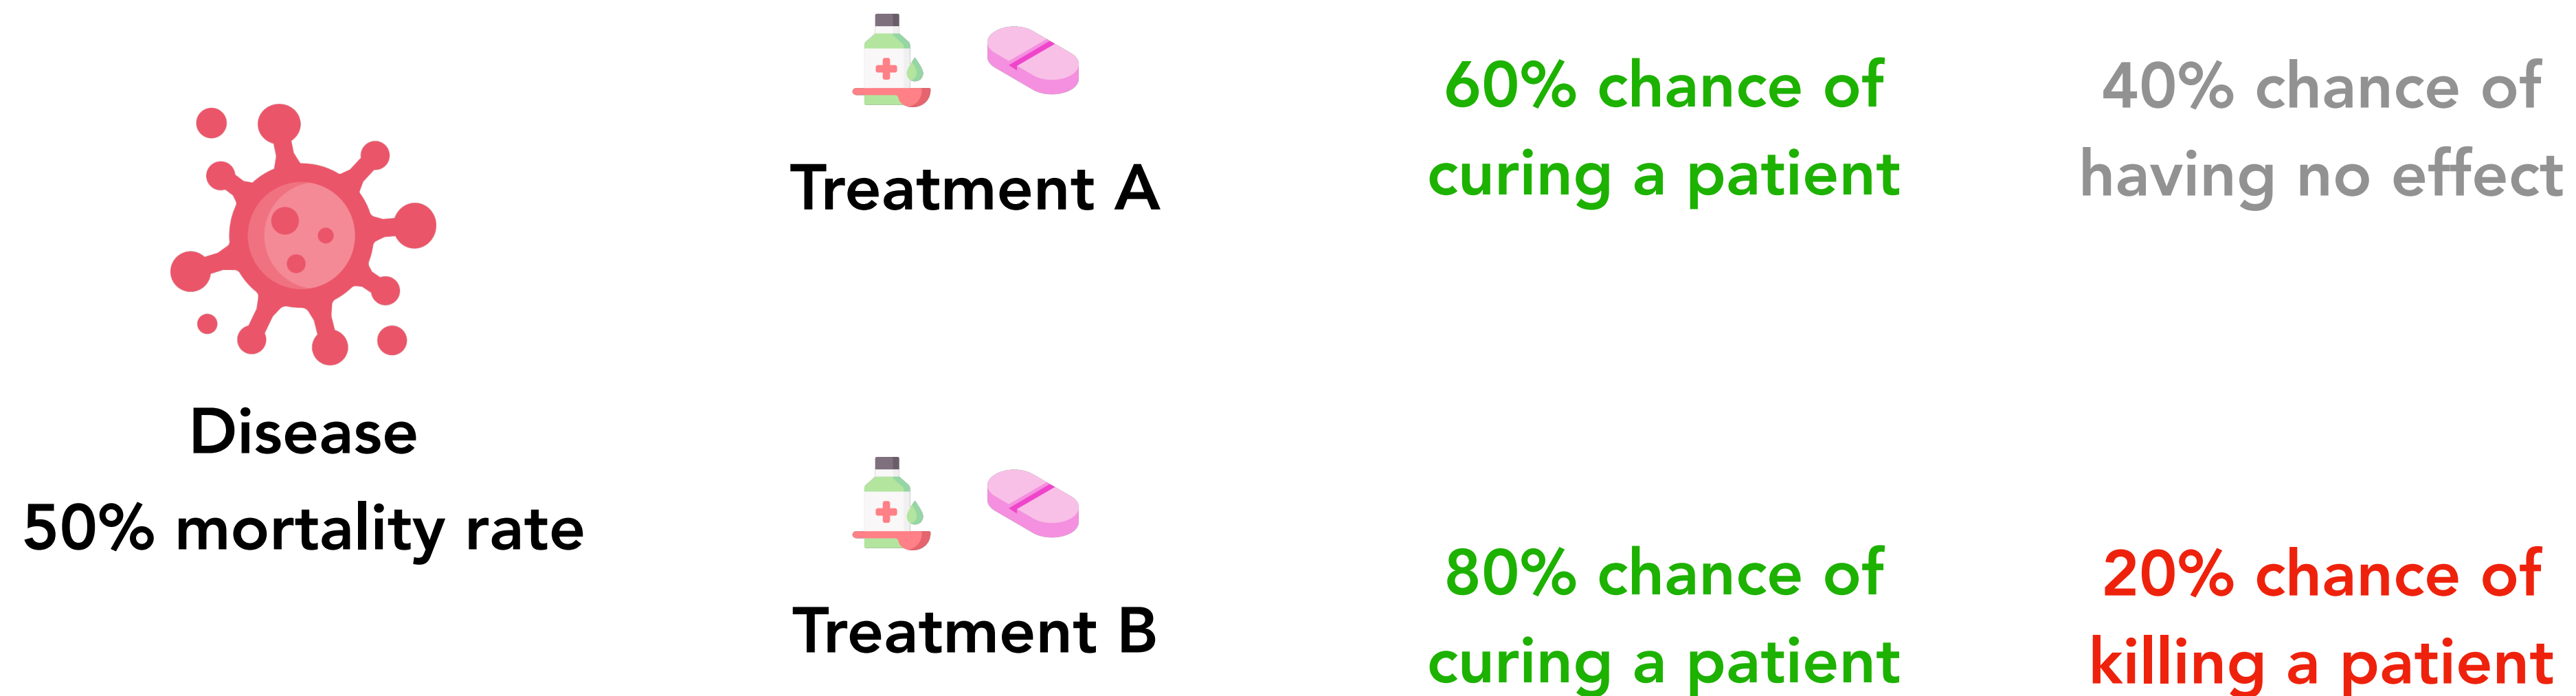
**20% chance of  
killing a patient**

# Counterfactual harm



Treatments A and B have **identical recovery rates**. However, doctors would systematically favor treatment A as it achieves the same recovery rate but never harms the patient.

# Counterfactual harm



Treatments A and B have **identical recovery rates**. However, doctors would systematically favor treatment A as it achieves the same recovery rate but never harms the patient.

- Under treatment A, there are no patients that would have survived had they not been treated.
- Under treatment B, there are patients who die following treatment who would have lived had they not been treated.

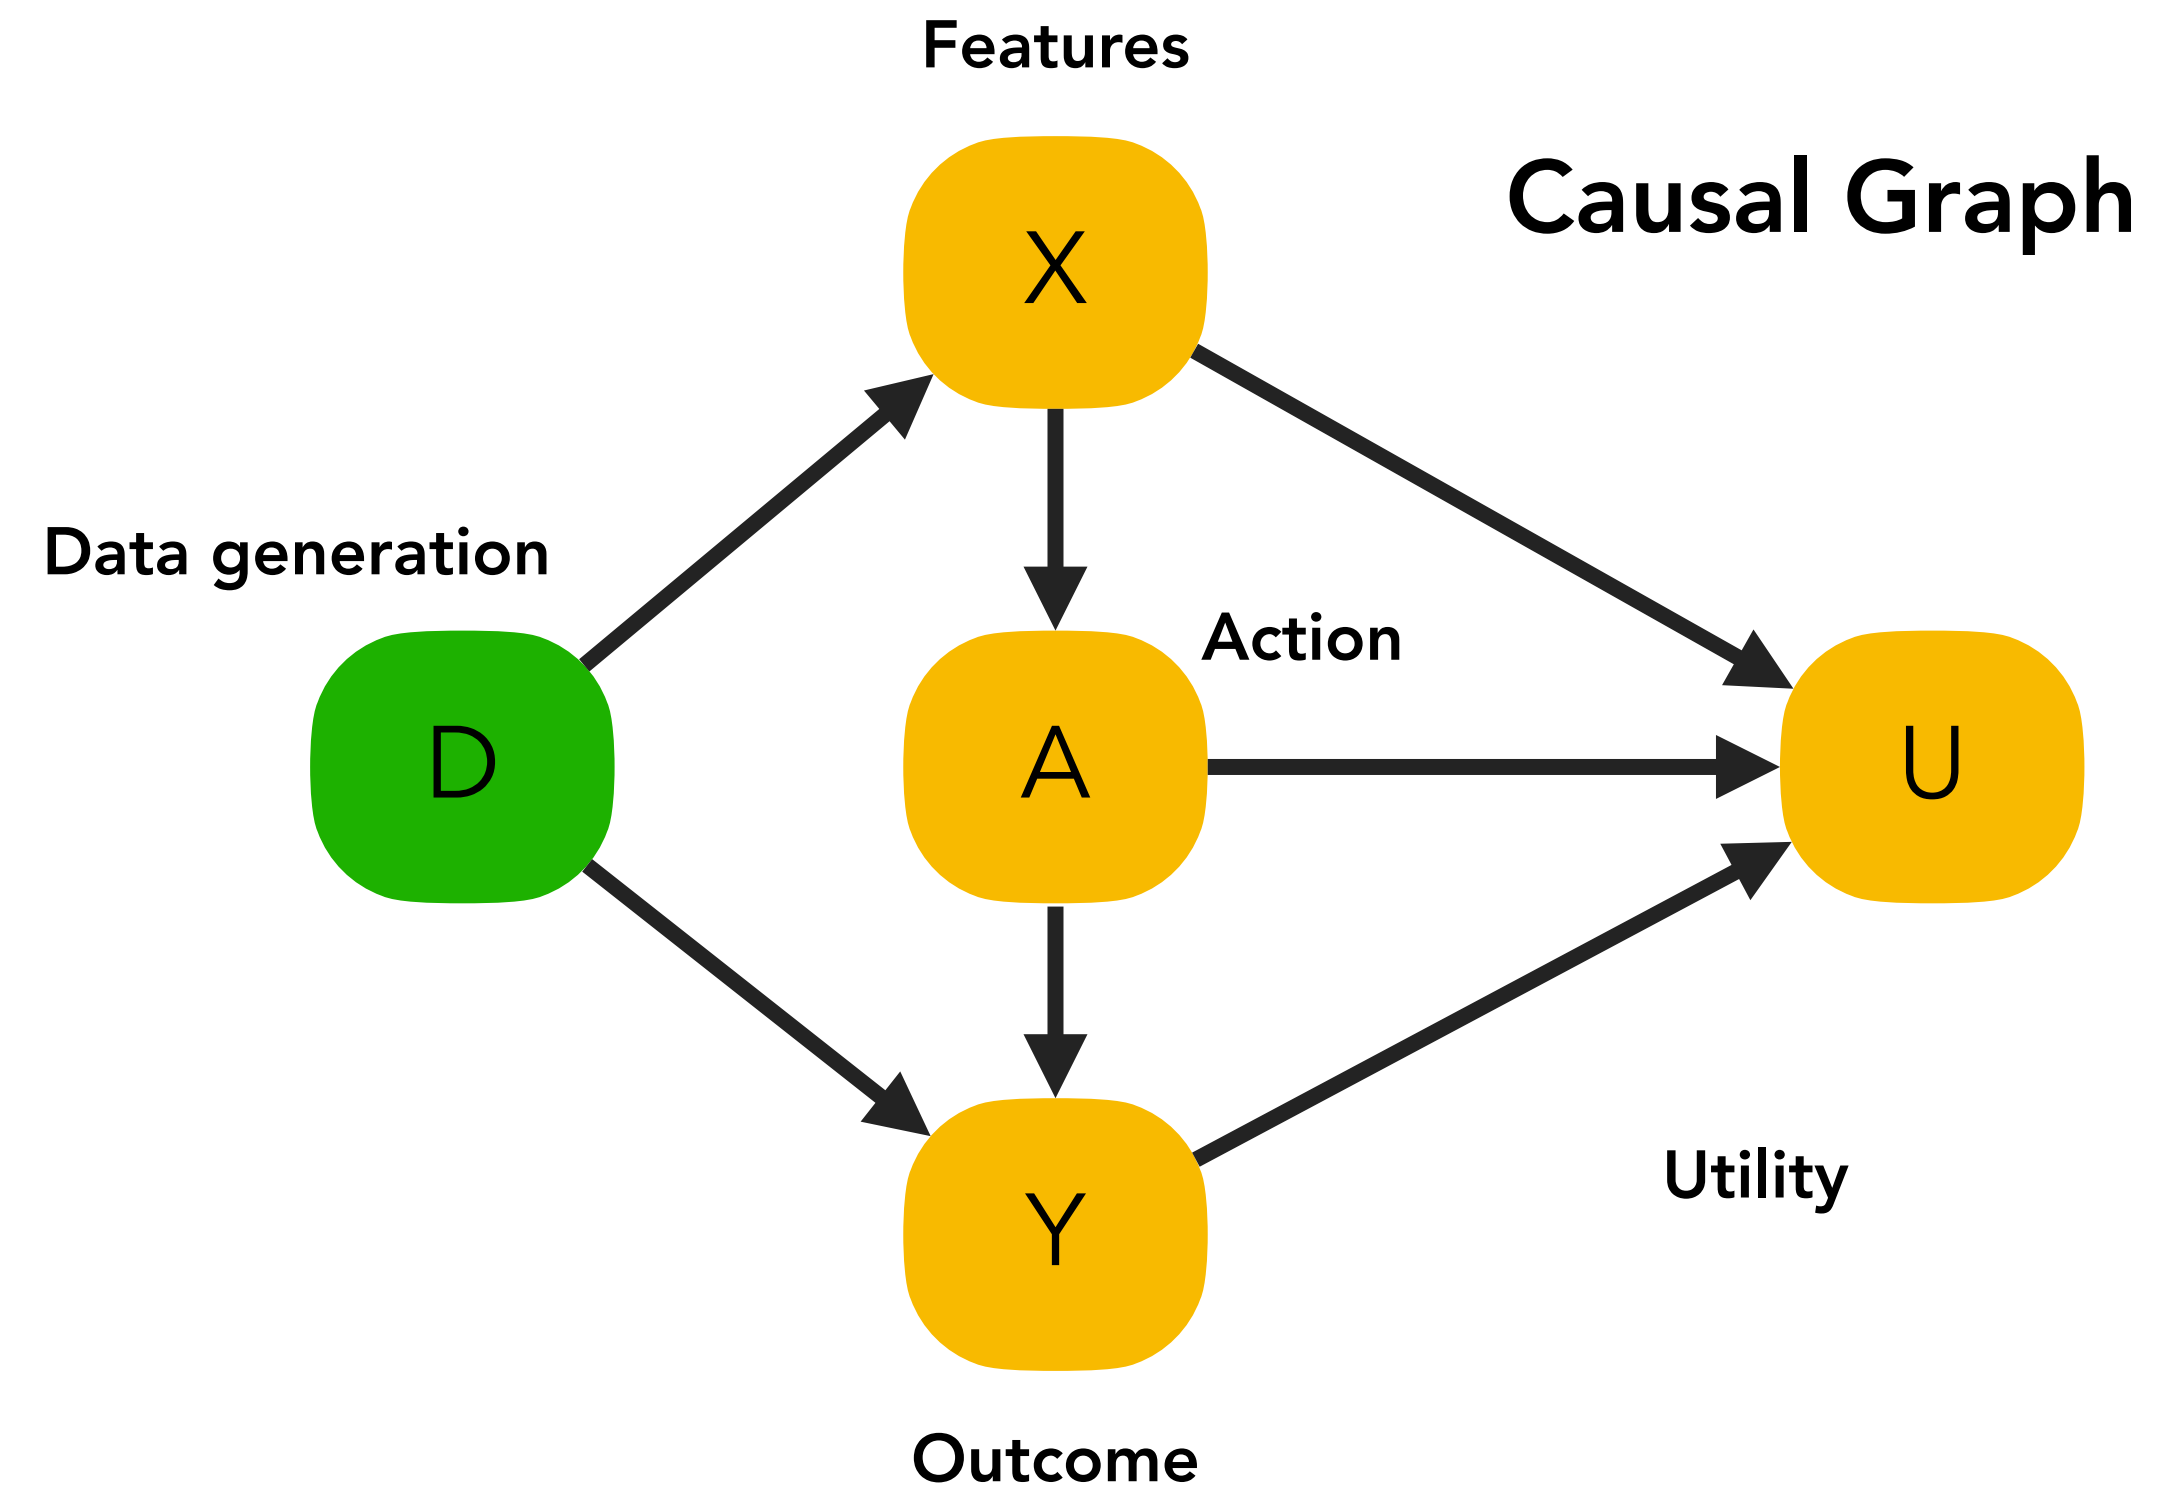
# Formalizing counterfactual harm

## Structural Causal Model $\mathcal{M}$

$$X := f_X(D) \quad Y := f_Y(D) \quad D \sim P(D)$$

$$A := \pi(X) \quad \leftarrow \text{Algorithmic policy} \quad \text{robot icon}$$

$$U := f_U(A, X, Y)$$





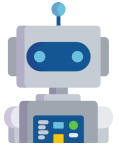
# Formalizing counterfactual harm

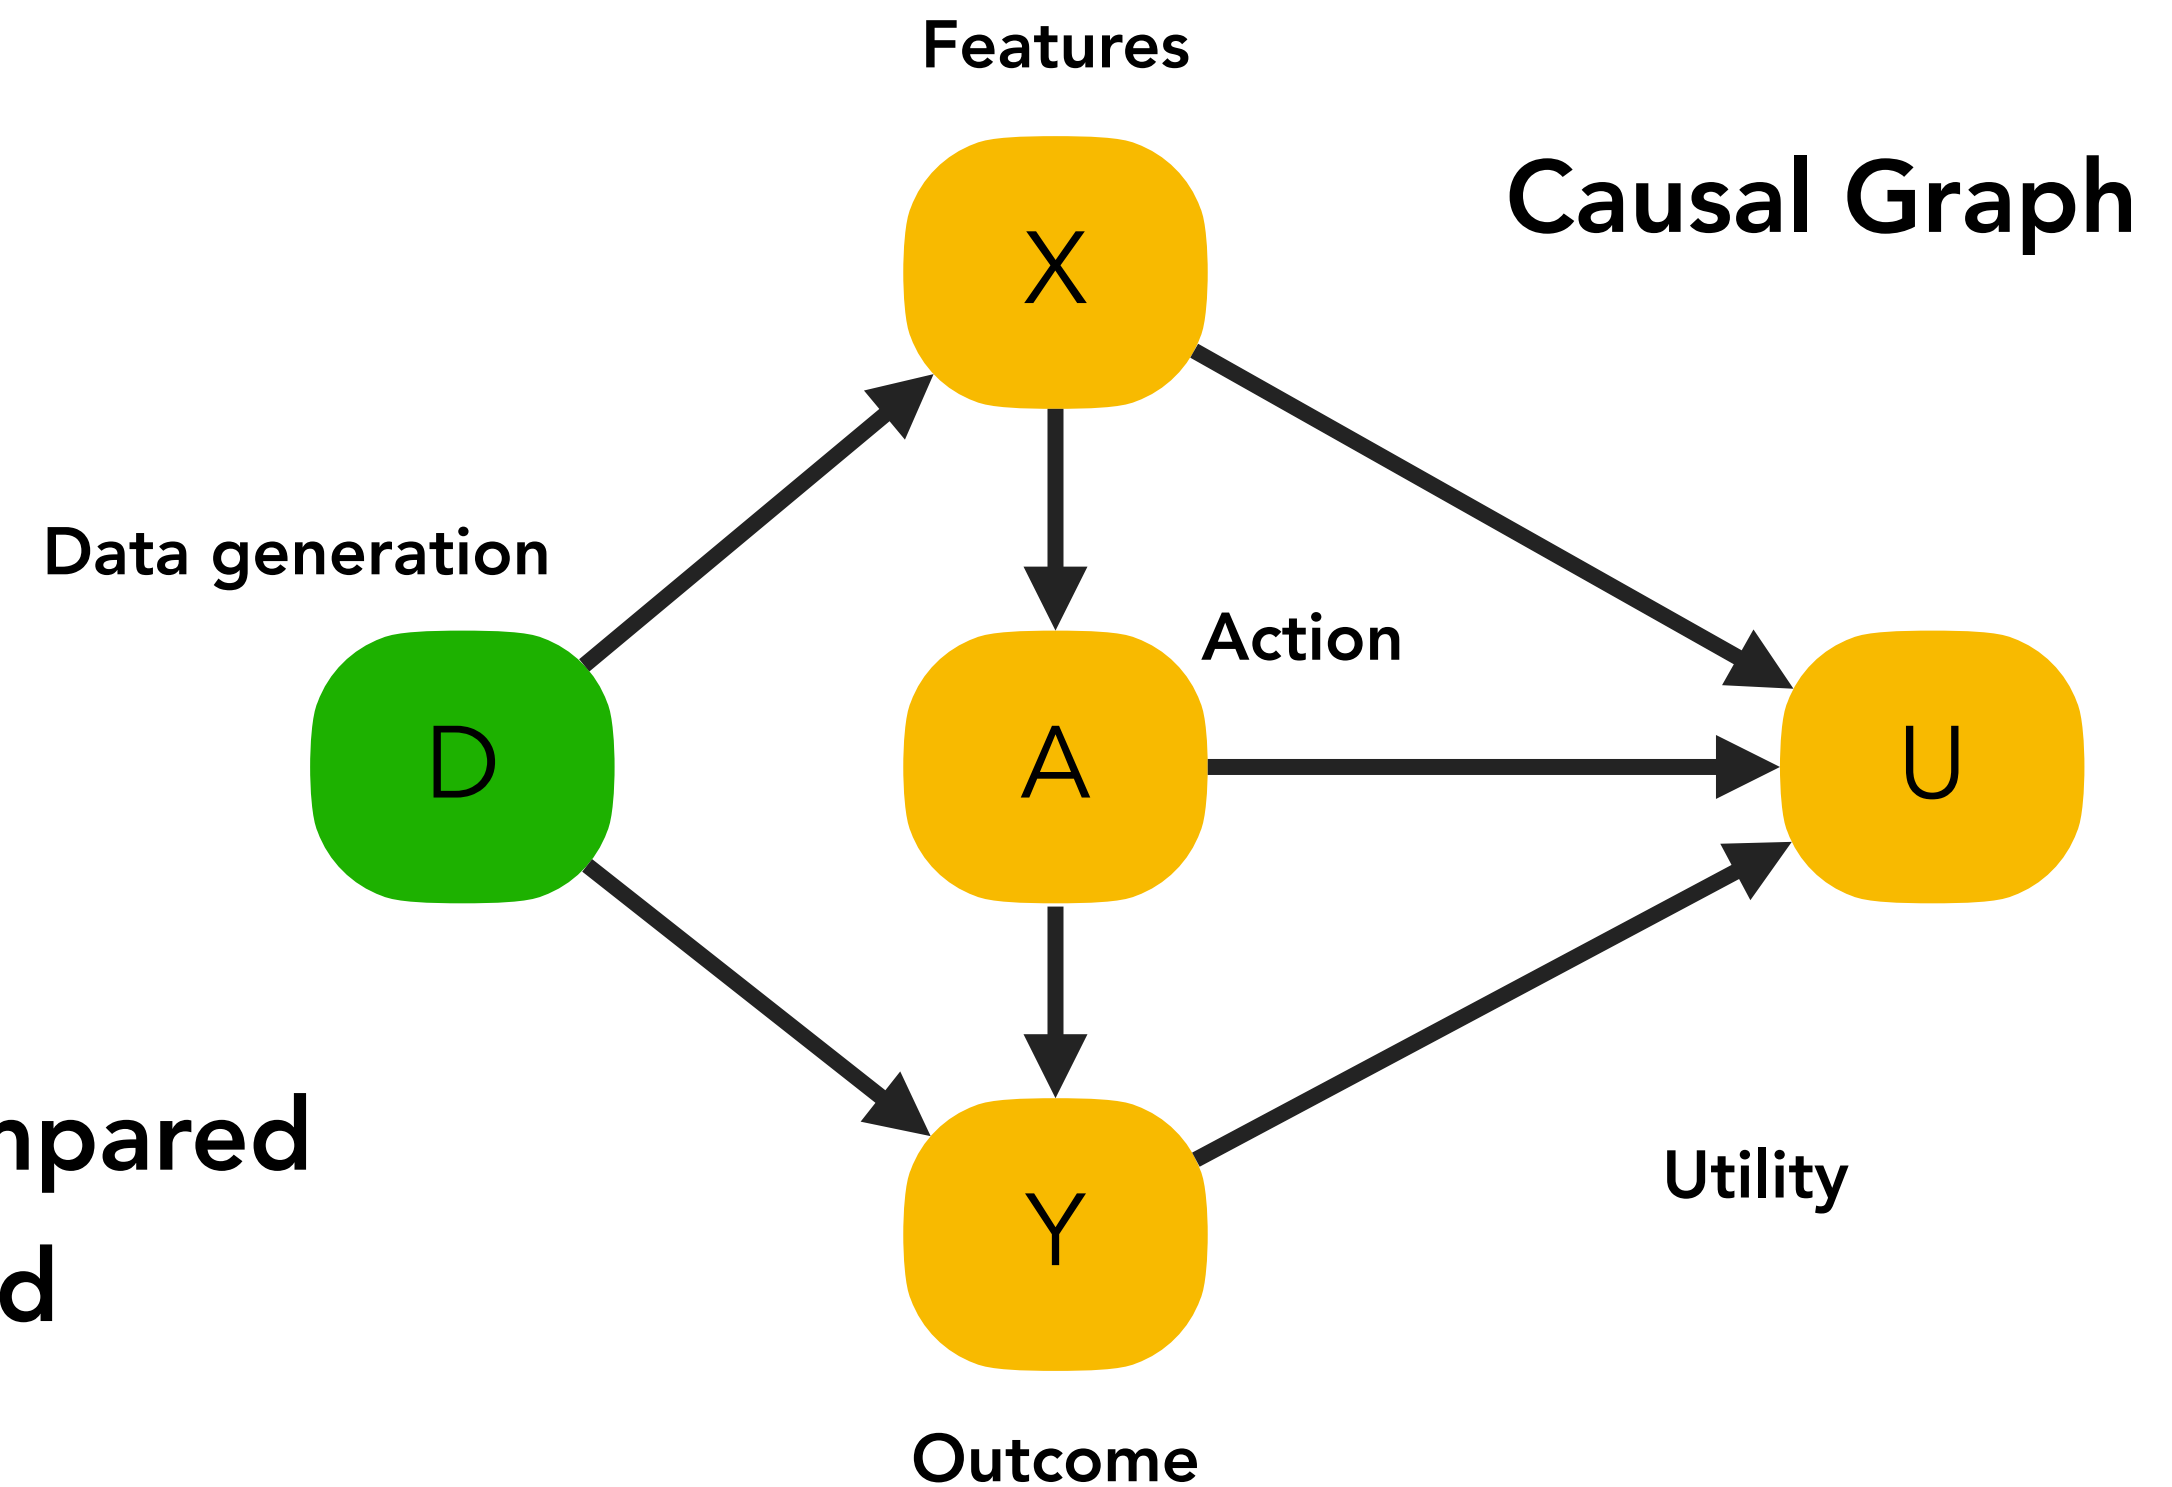
## Structural Causal Model $\mathcal{M}$

$$X := f_X(D) \quad Y := f_Y(D) \quad D \sim P(D)$$

$$A := \pi(X) \quad \leftarrow \text{Algorithmic policy} \quad \text{robot icon}$$

$$U := f_U(A, X, Y)$$

Harm caused by action  $a$  taken by  compared to default action  $\bar{a}$  given context  $X = x$  and outcome  $Y = y$



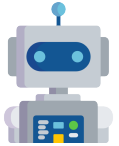
# Formalizing counterfactual harm

## Structural Causal Model $\mathcal{M}$

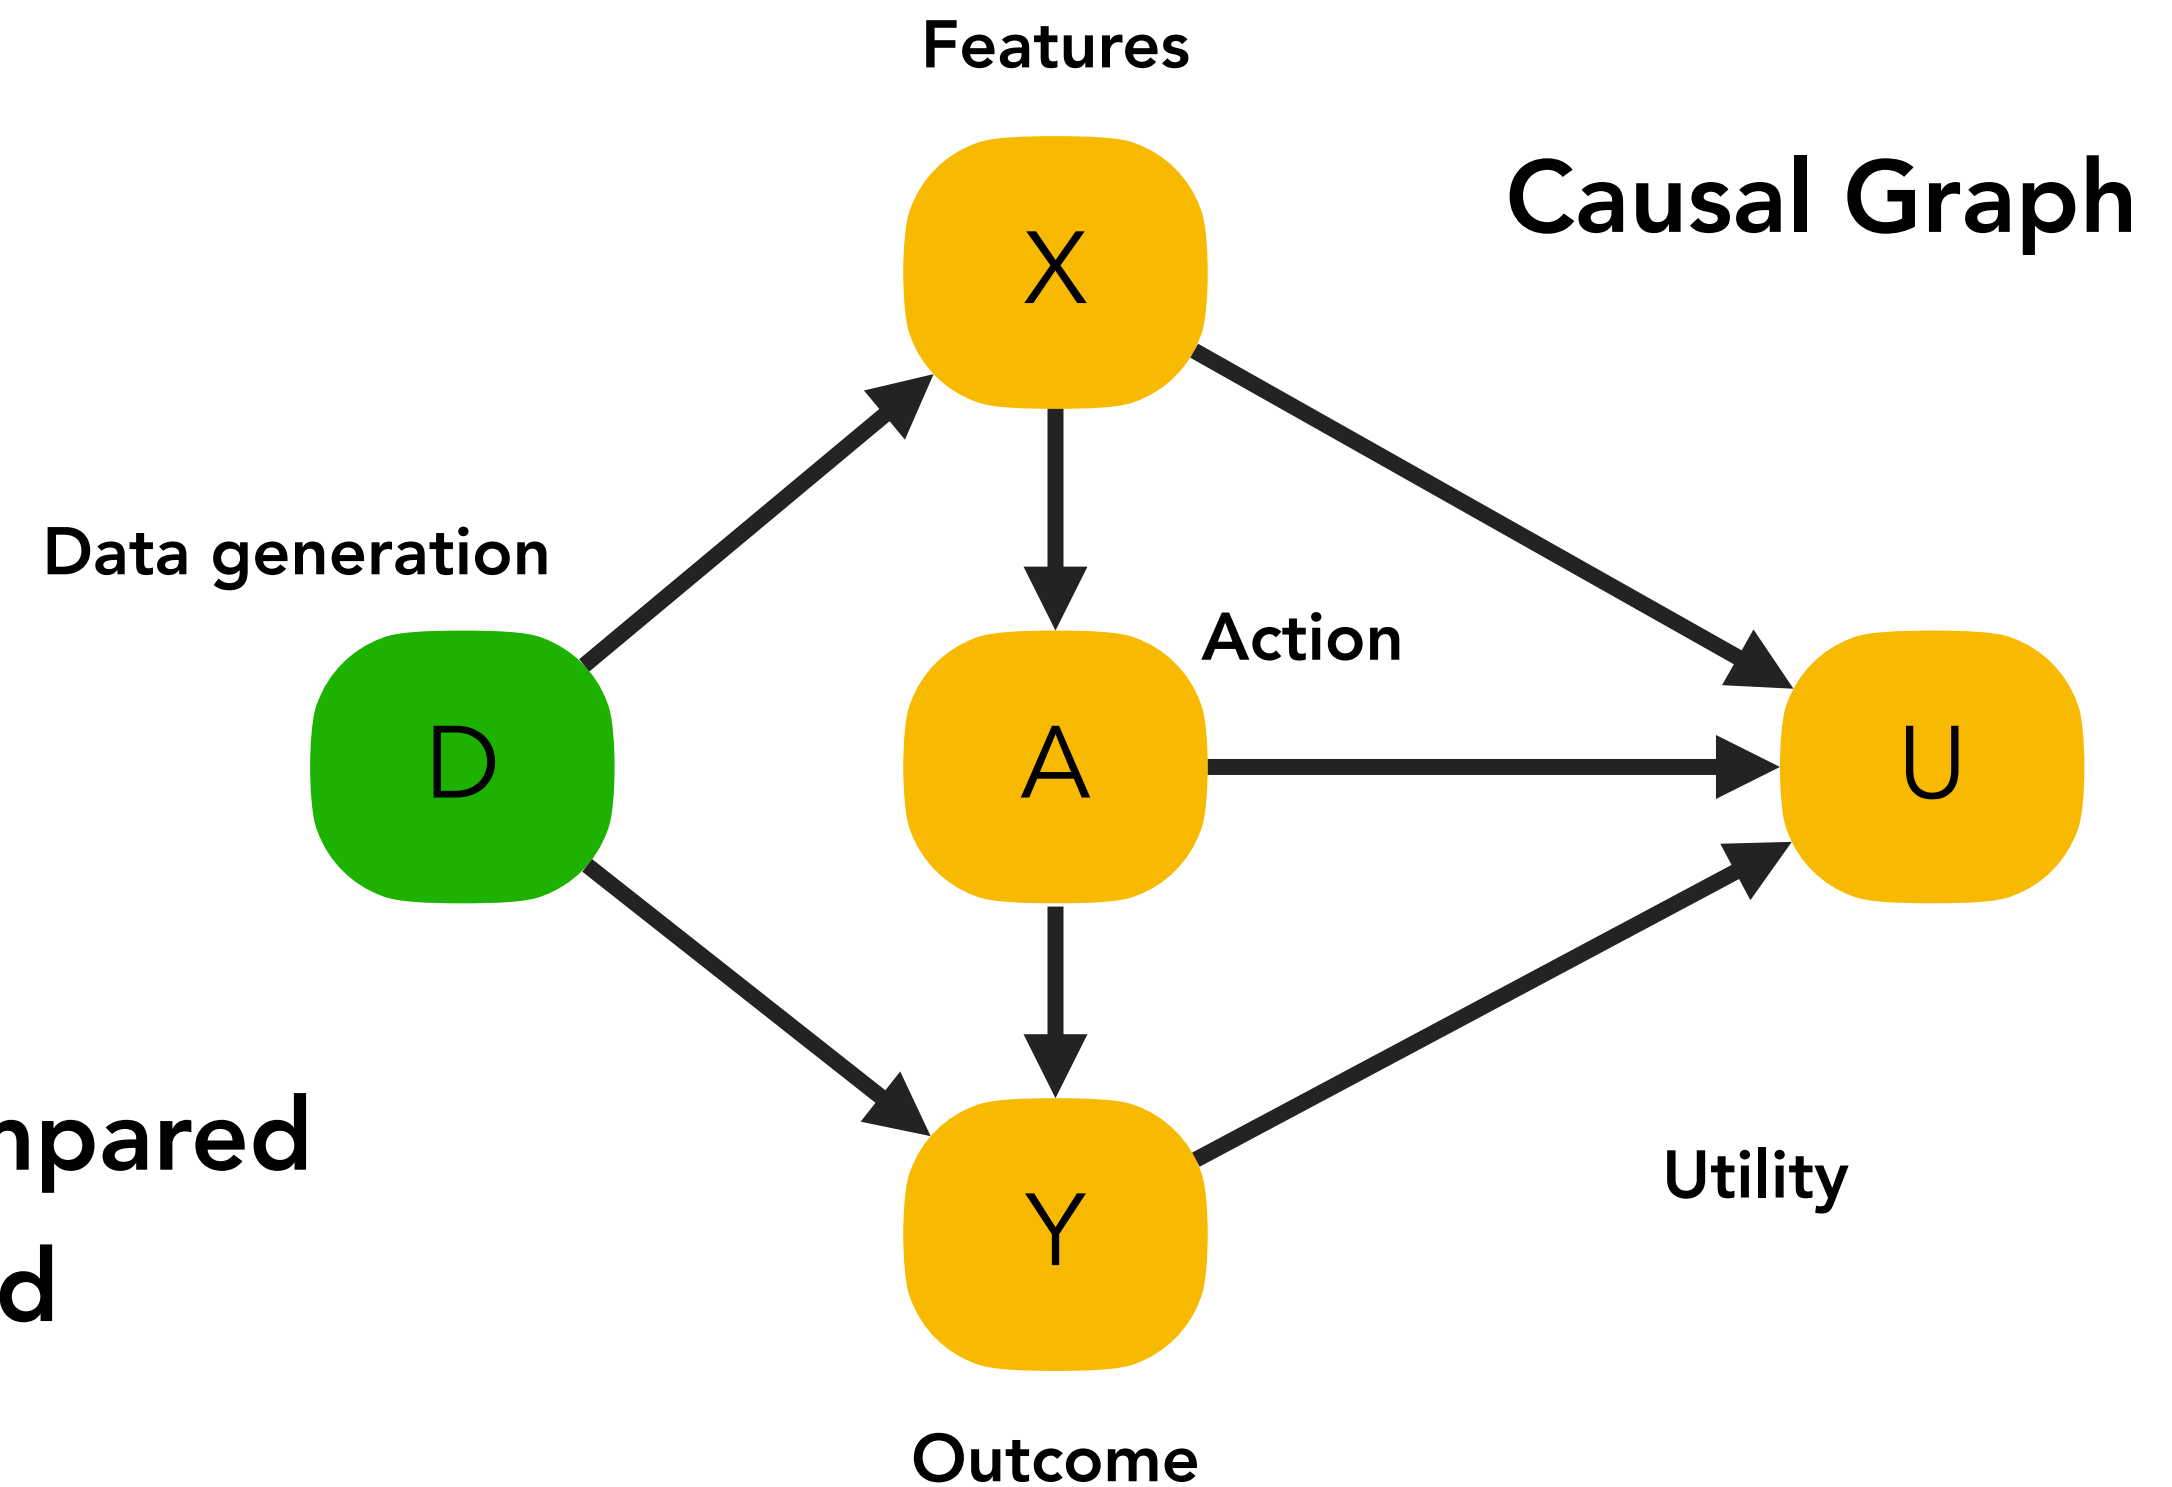
$$X := f_X(D) \quad Y := f_Y(D) \quad D \sim P(D)$$

$$A := \pi(X) \quad \leftarrow \text{Algorithmic policy} \quad \text{robot icon}$$

$$U := f_U(A, X, Y)$$

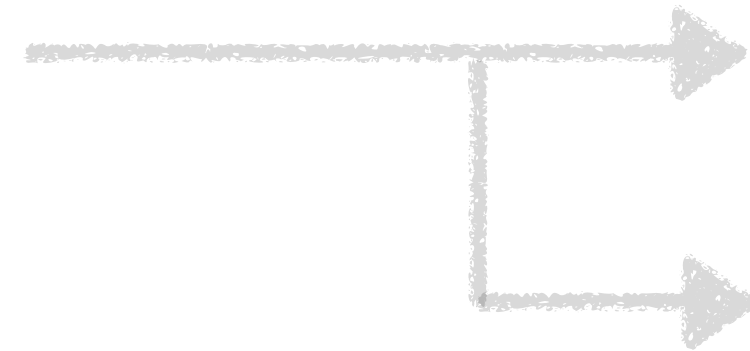
Harm caused by action  $a$  taken by  compared to default action  $\bar{a}$  given context  $X = x$  and outcome  $Y = y$

$$h(a, x, y) = \int_{y'} P^{\mathcal{M}} | X=x, Y=y, A=a ; do(A=\bar{a}) (Y = y') \max \left( 0, \underbrace{U(\bar{a}, x, y')}_{\text{Counterfactual utility}} - \underbrace{U(a, x, y)}_{\text{Utility}} \right) dy'$$

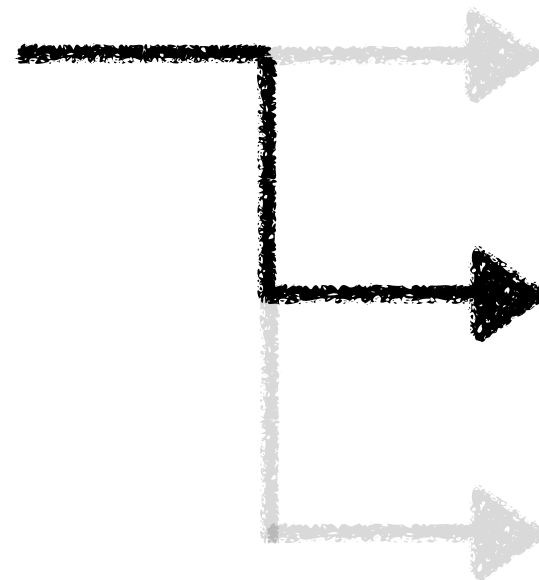


# Use cases of counterfactuals in machine learning

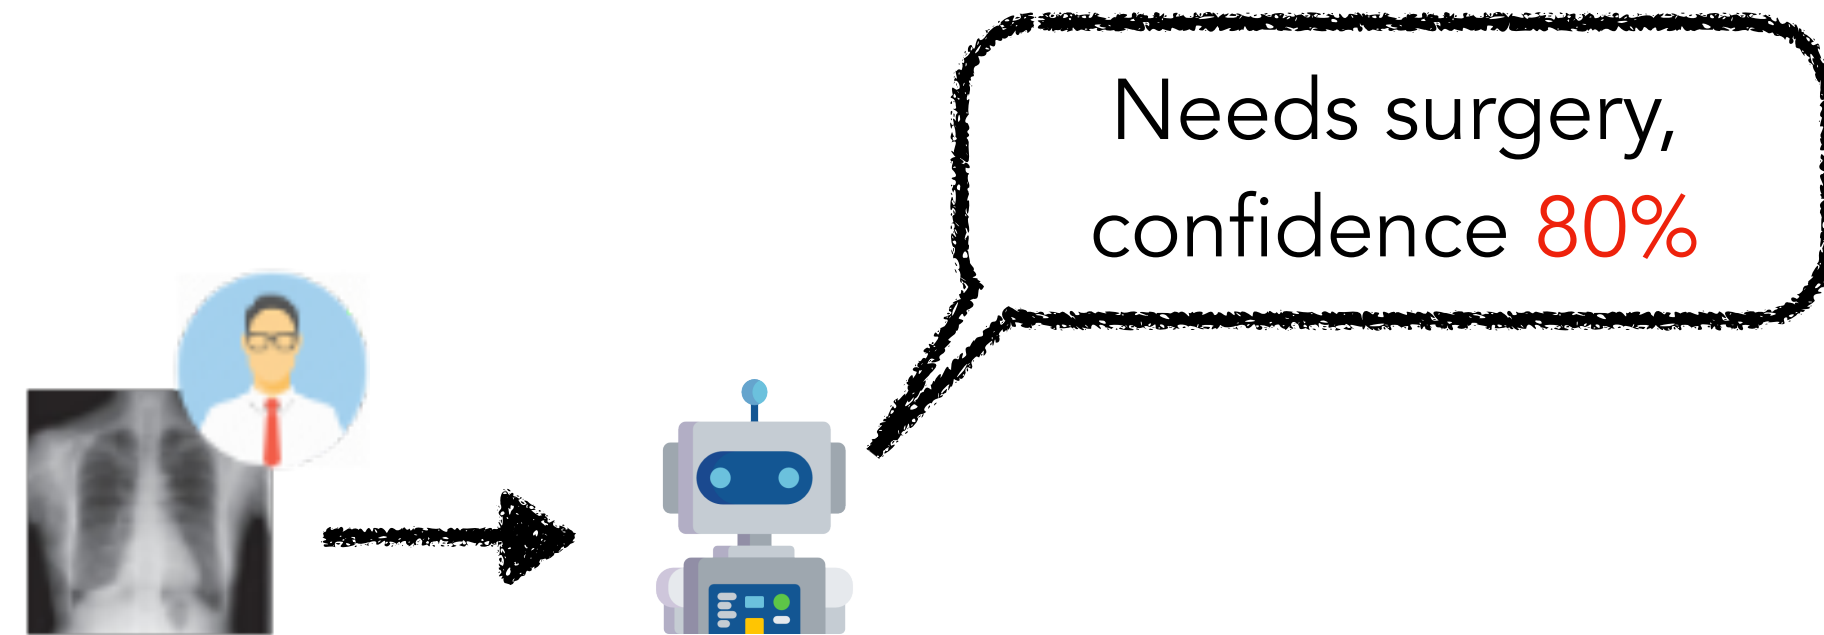
*Decision making*



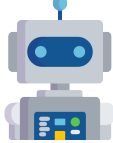
*Calibration*



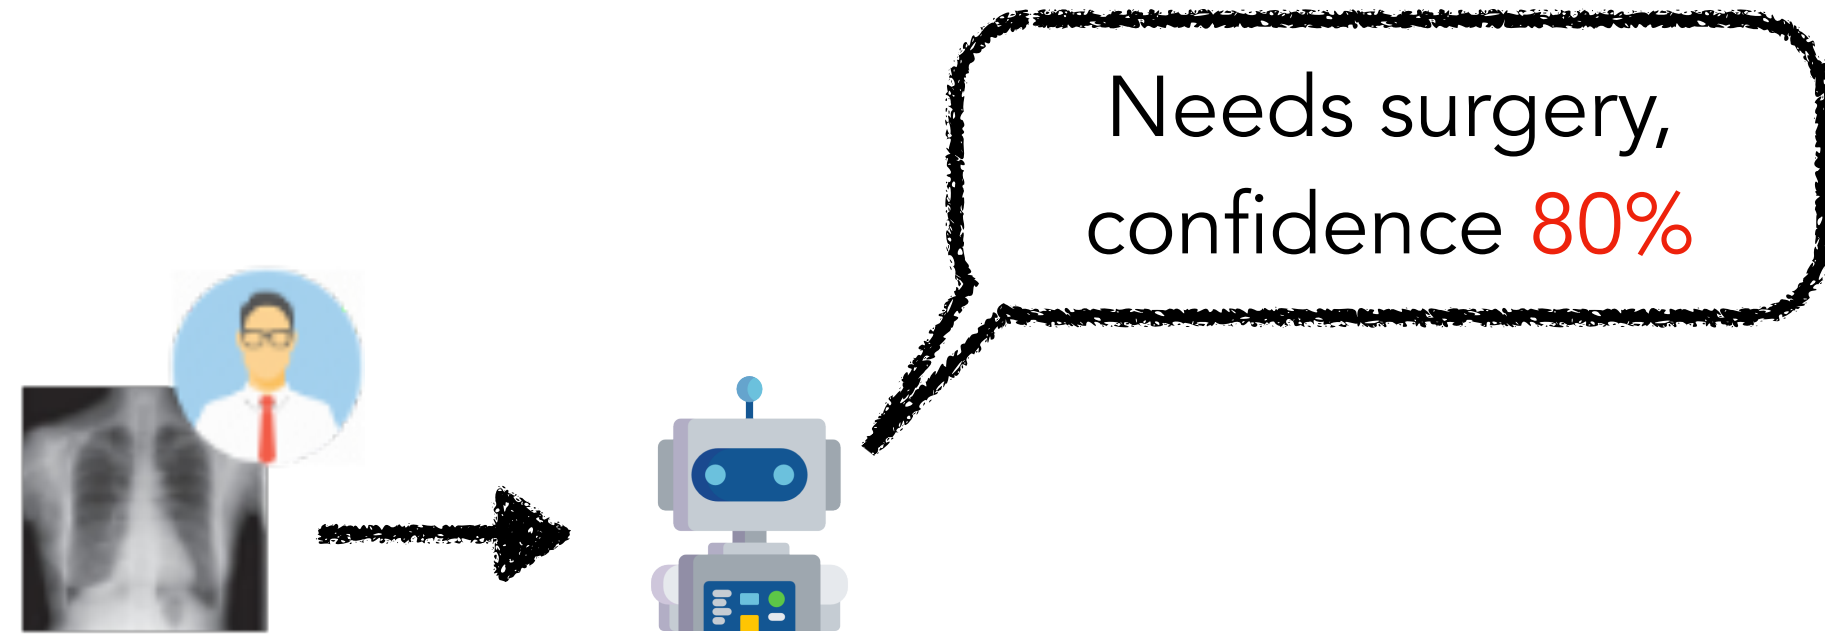
# Calibration



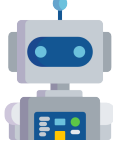
## Calibration:

Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

# Calibration

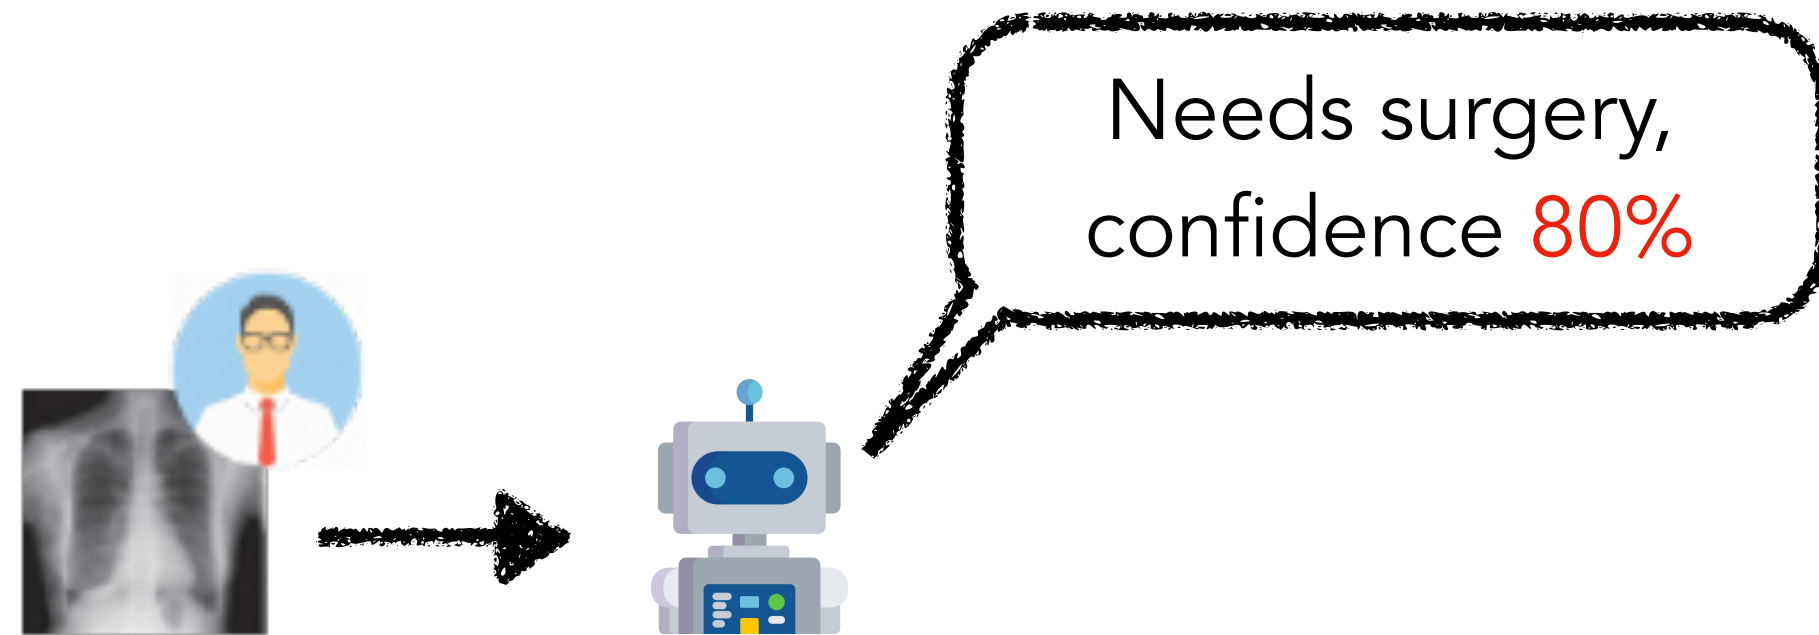


## Calibration:

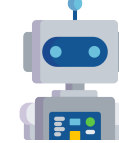
Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

Counterfactual reasoning reveals that the way in which machine learning models compute confidence values today is problematic.

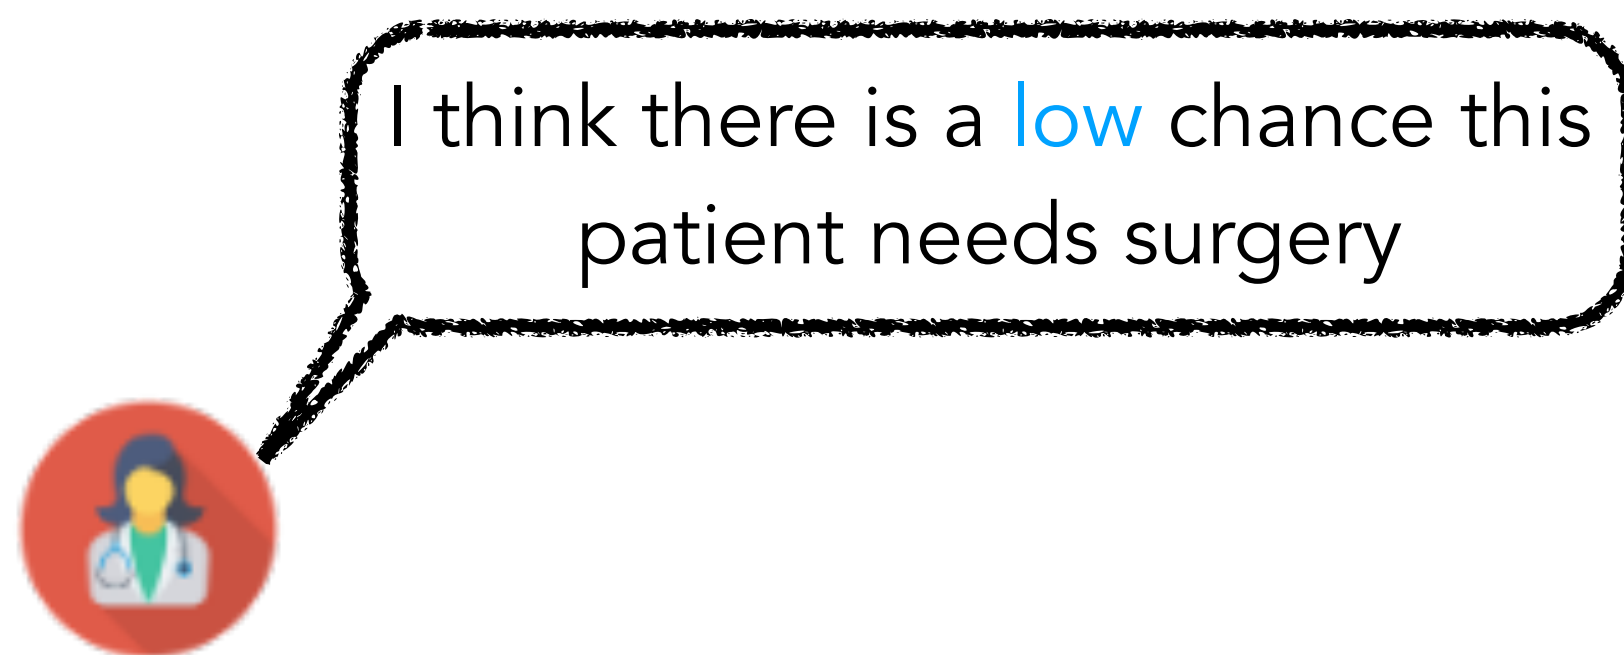
# Calibration



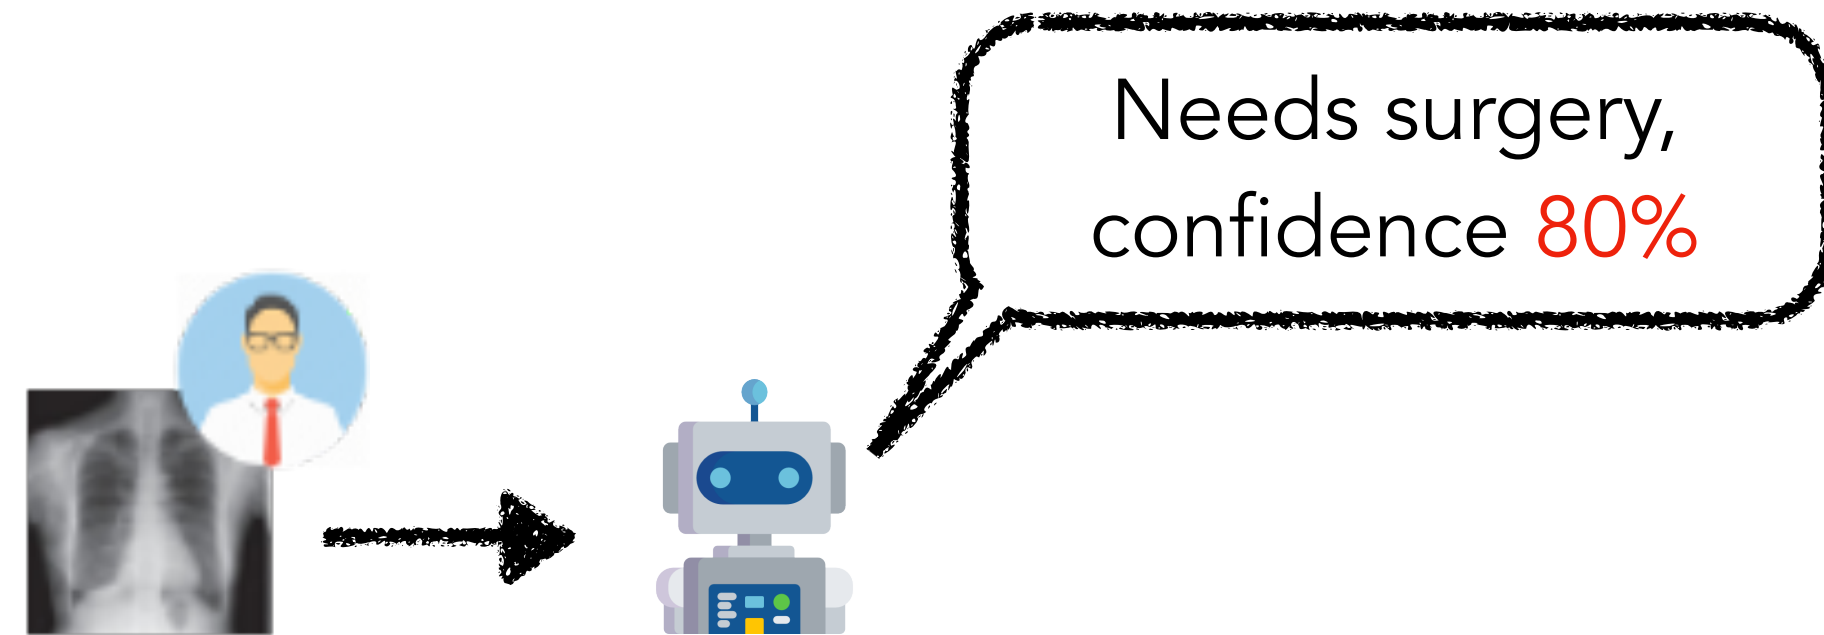
## Calibration:

Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

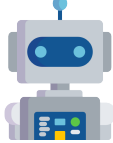
Counterfactual reasoning reveals that the way in which machine learning models compute confidence values today is problematic.



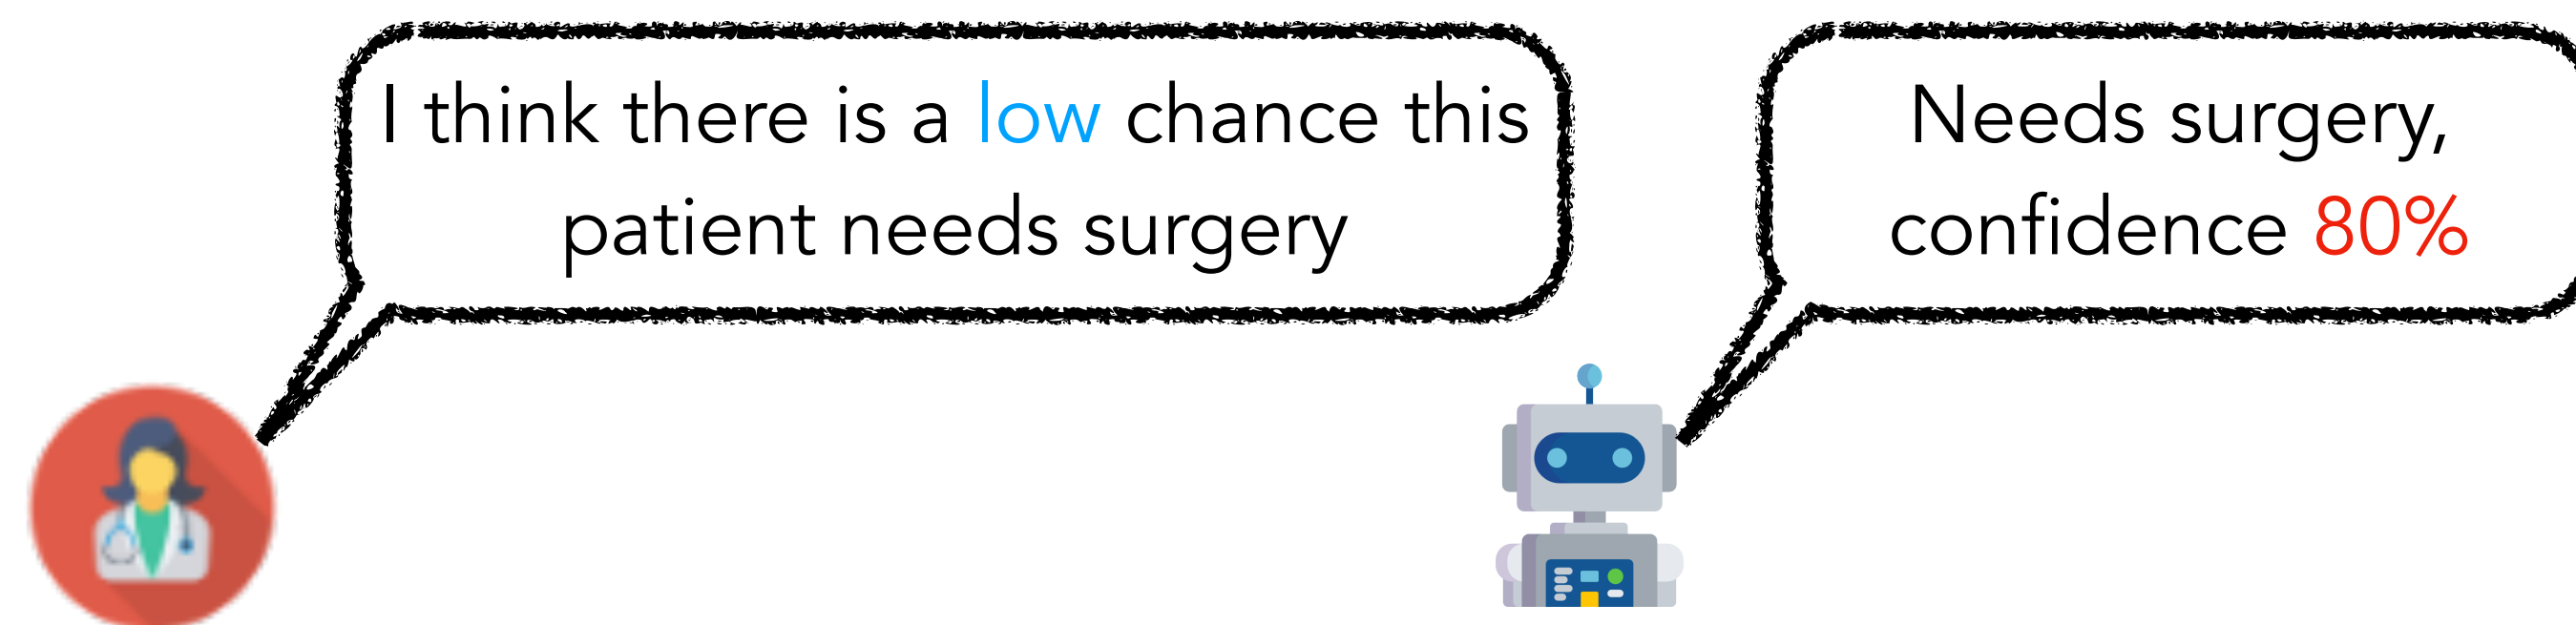
# Calibration



## Calibration:

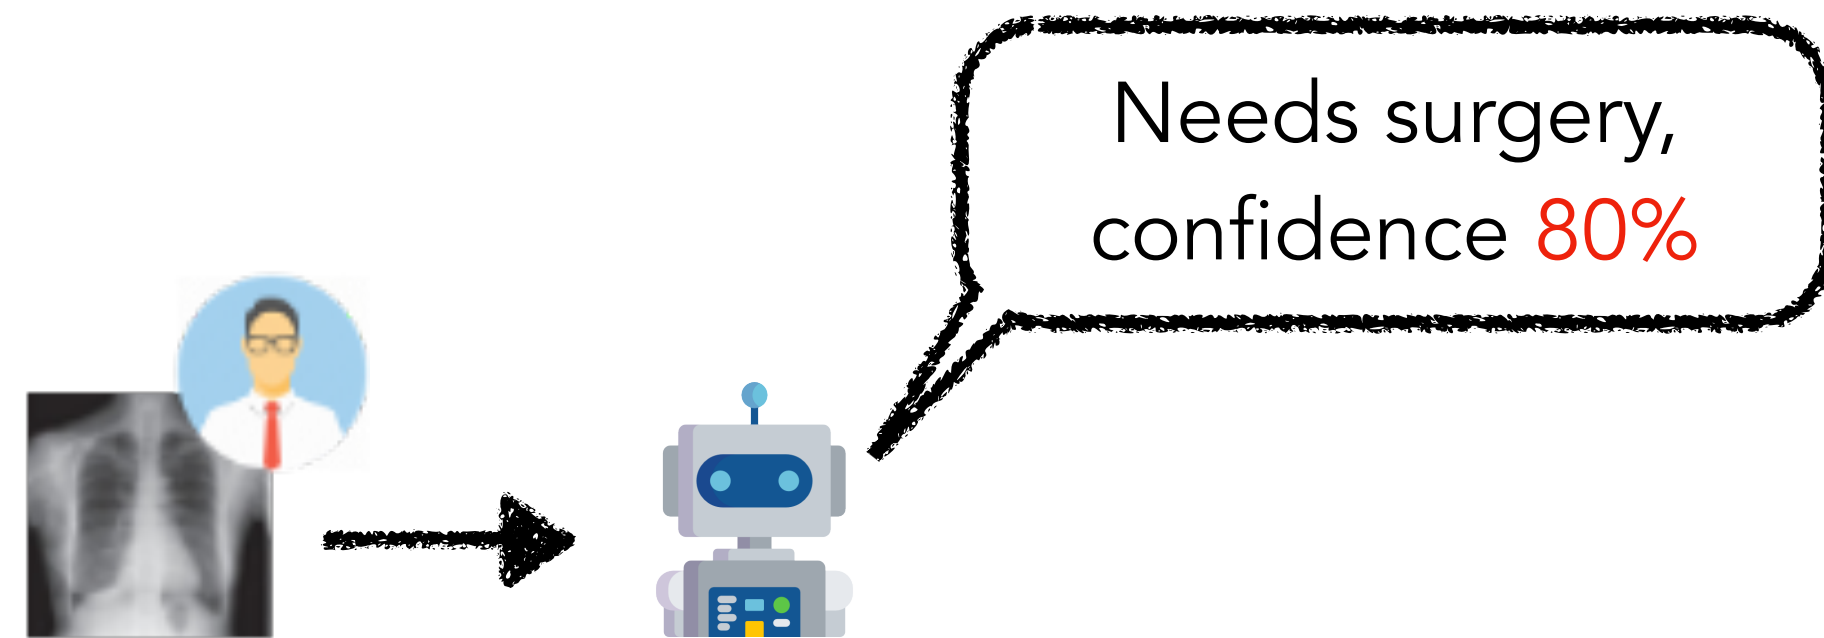
Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

Counterfactual reasoning reveals that the way in which machine learning models compute confidence values today is problematic.

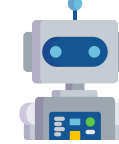




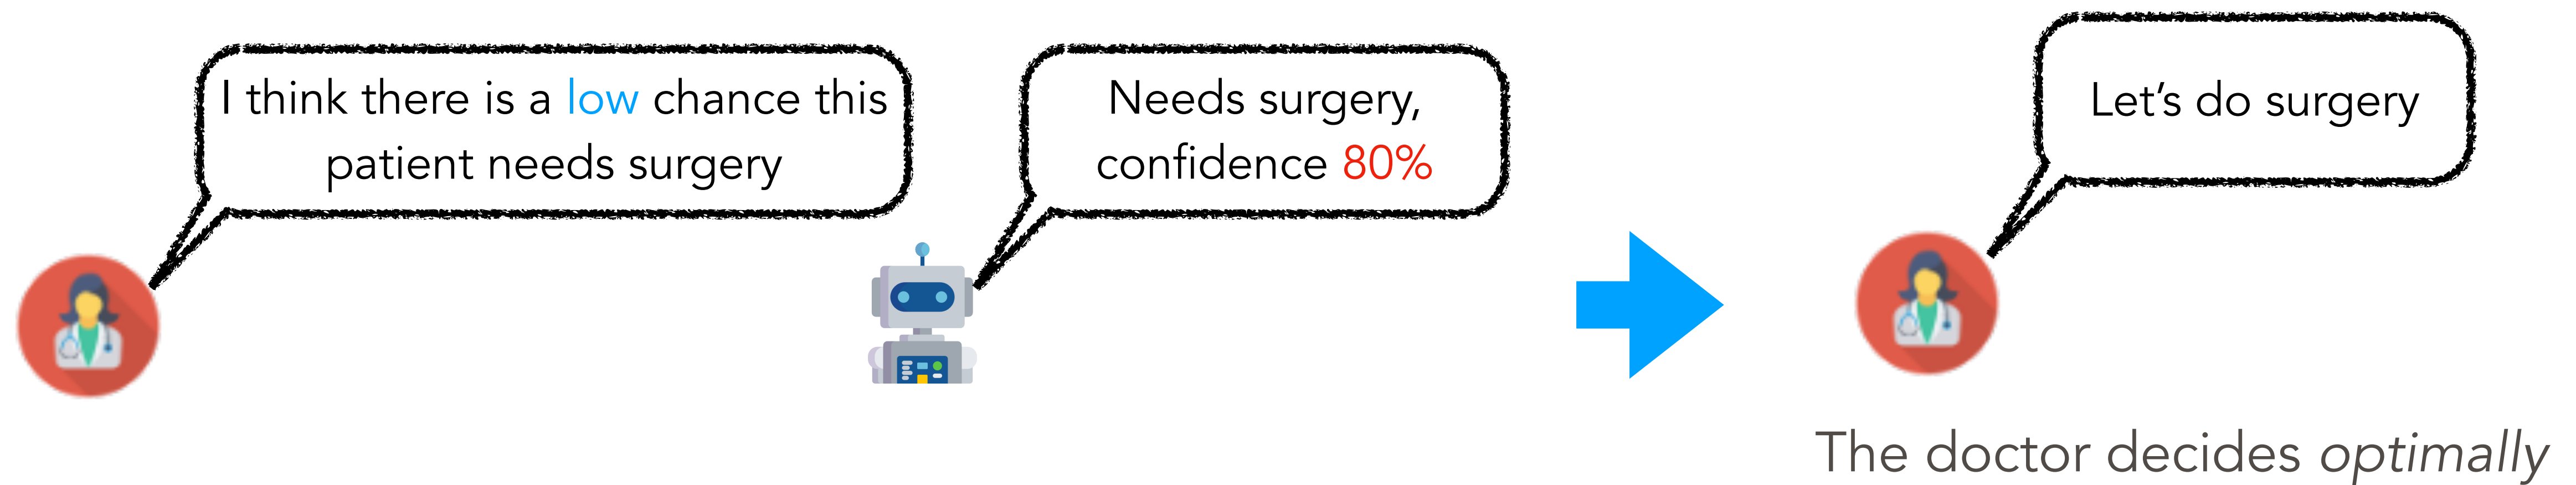
# Calibration



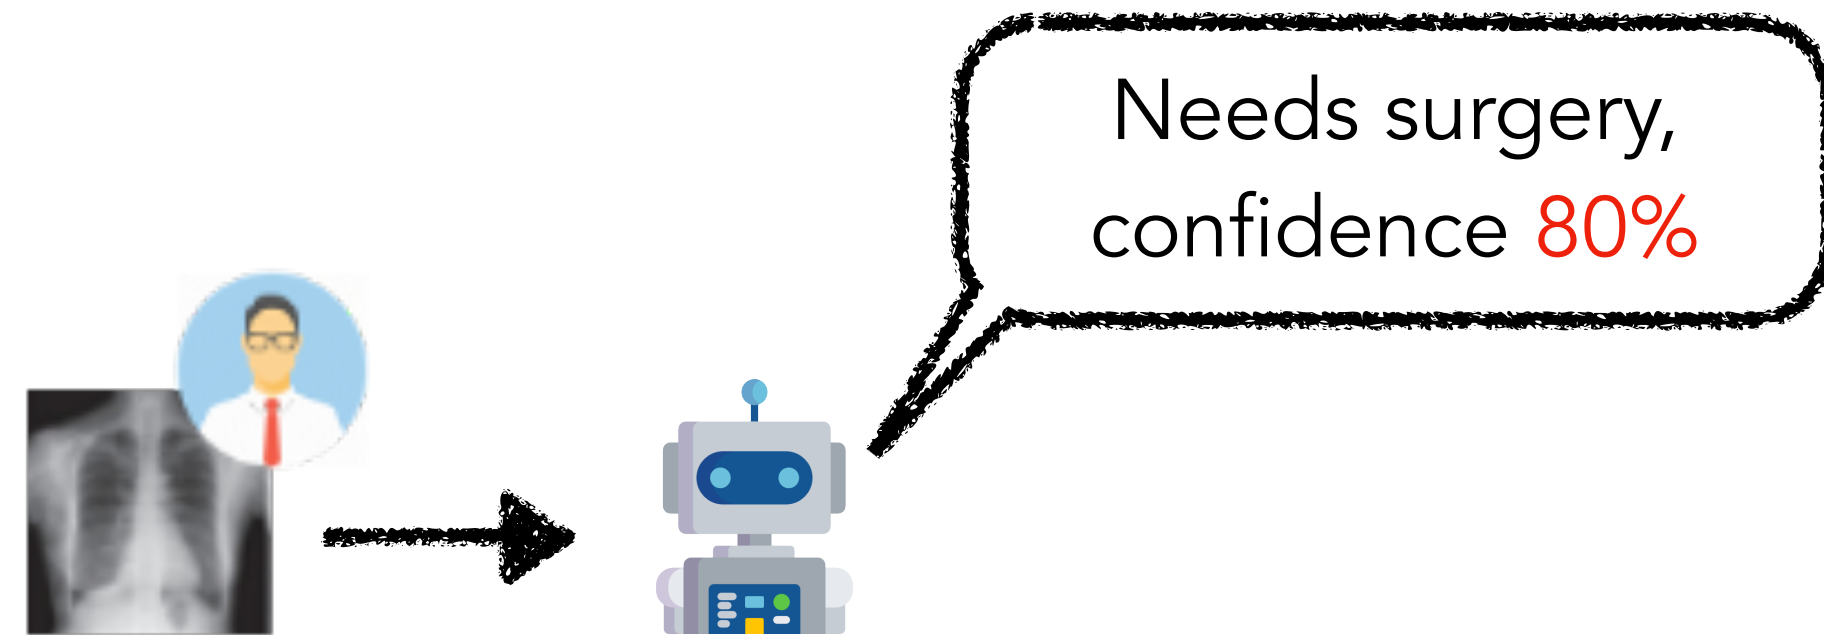
## Calibration:

Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

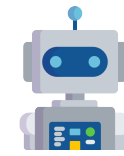
Counterfactual reasoning reveals that the way in which machine learning models compute confidence values today is problematic.



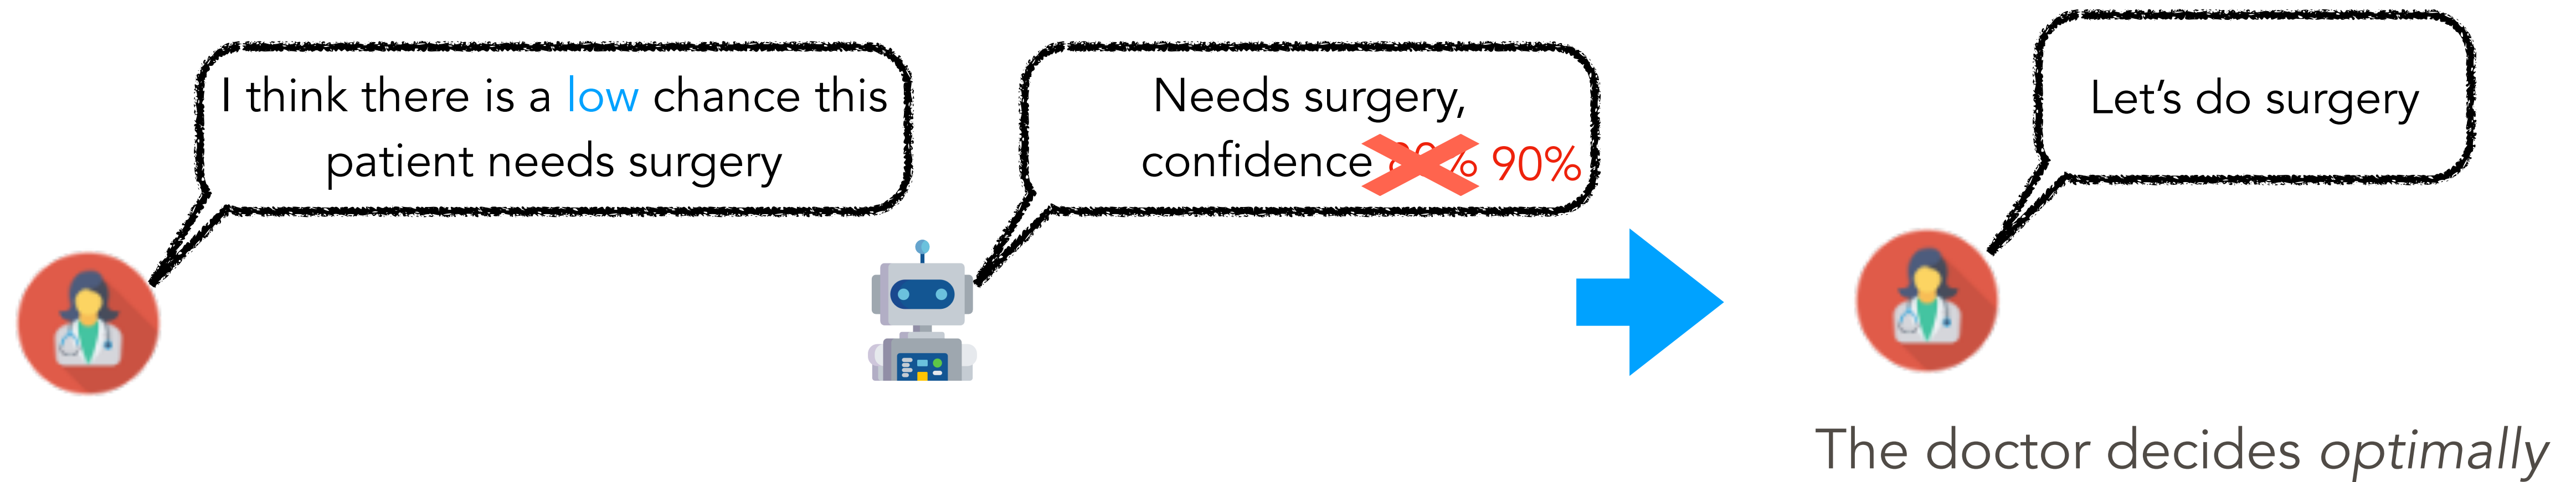
# Calibration



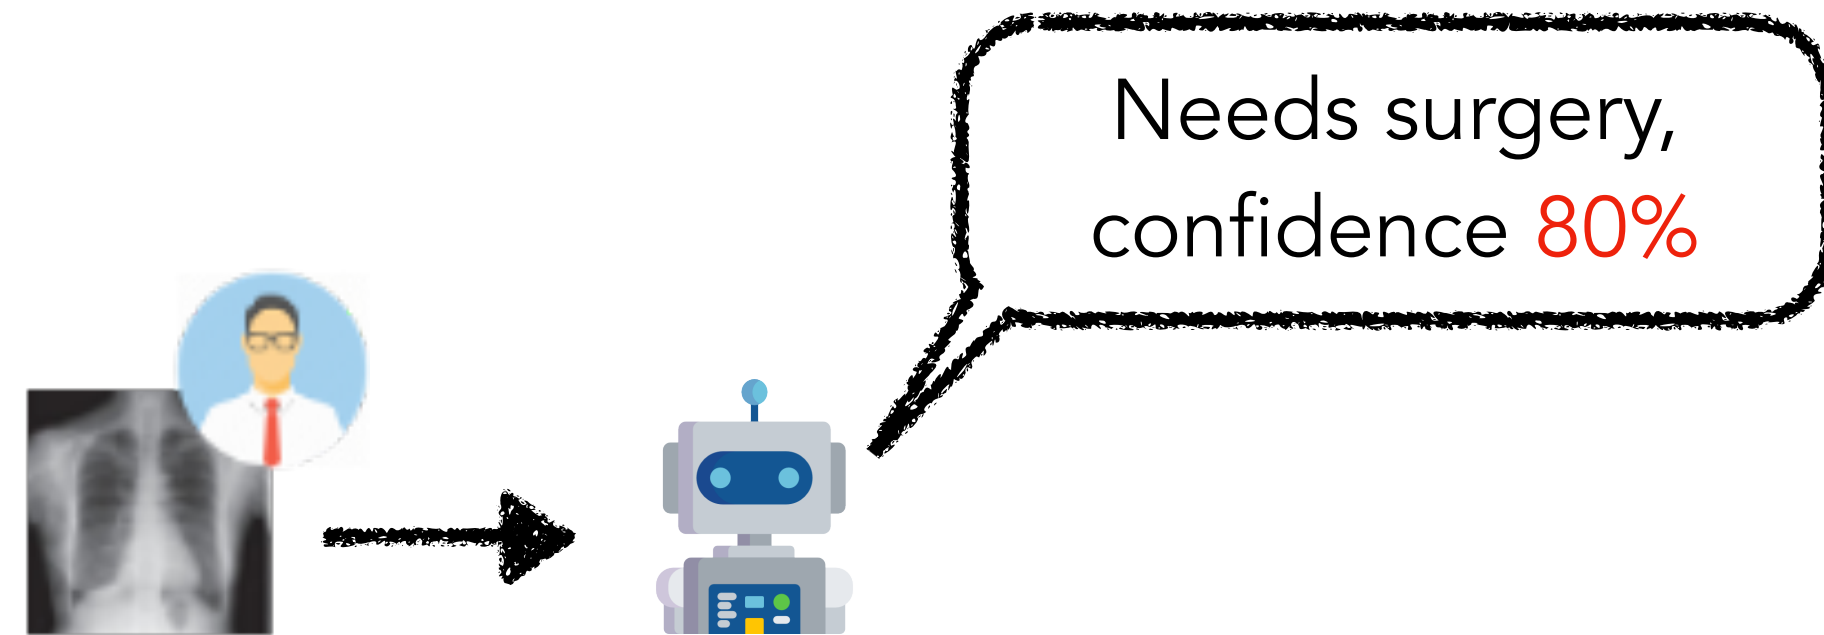
## Calibration:

Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

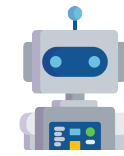
Counterfactual reasoning reveals that the way in which machine learning models compute confidence values today is problematic.



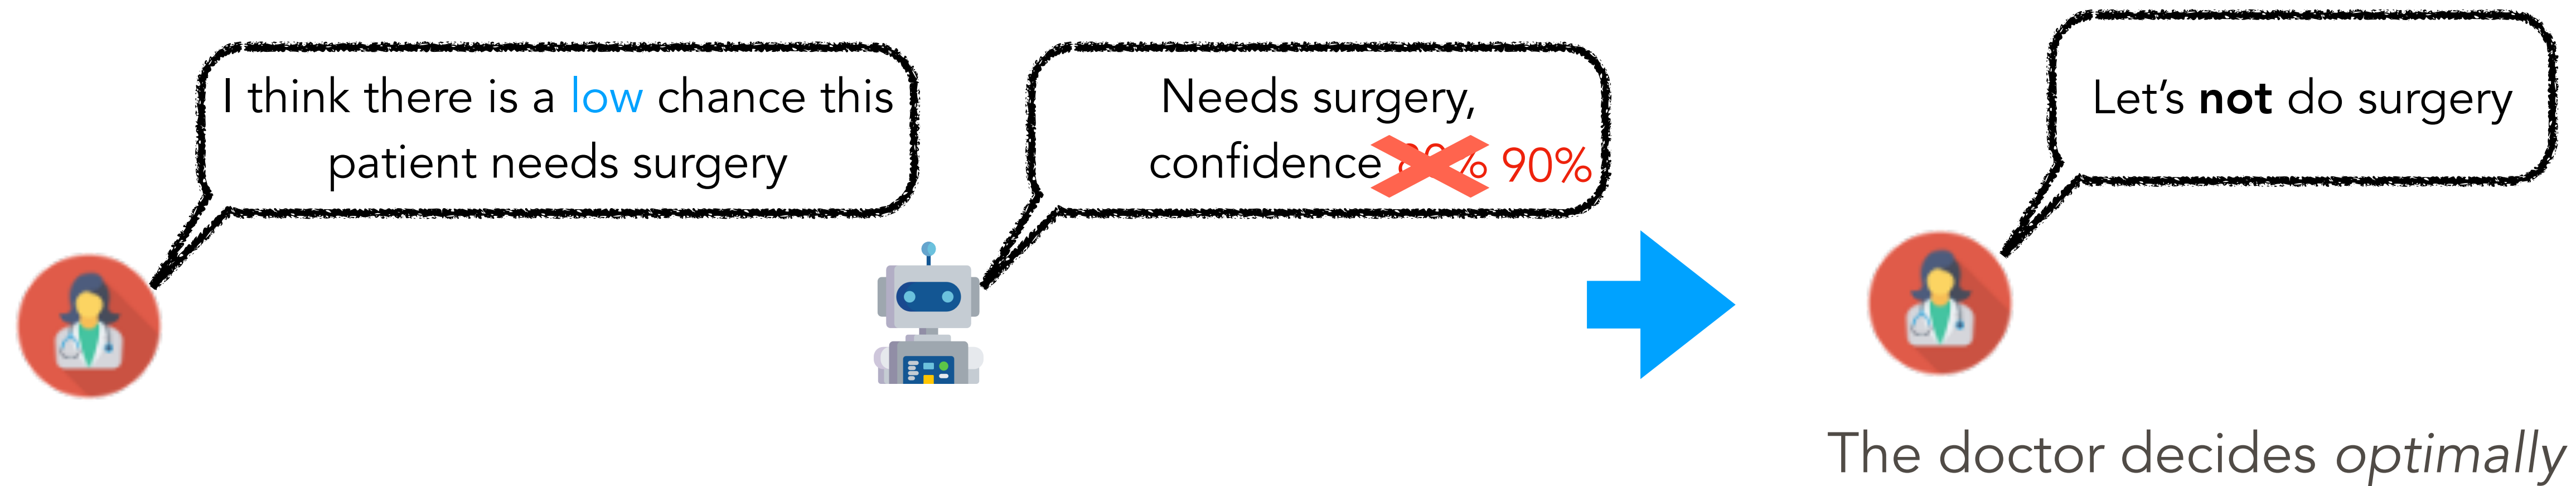
# Calibration



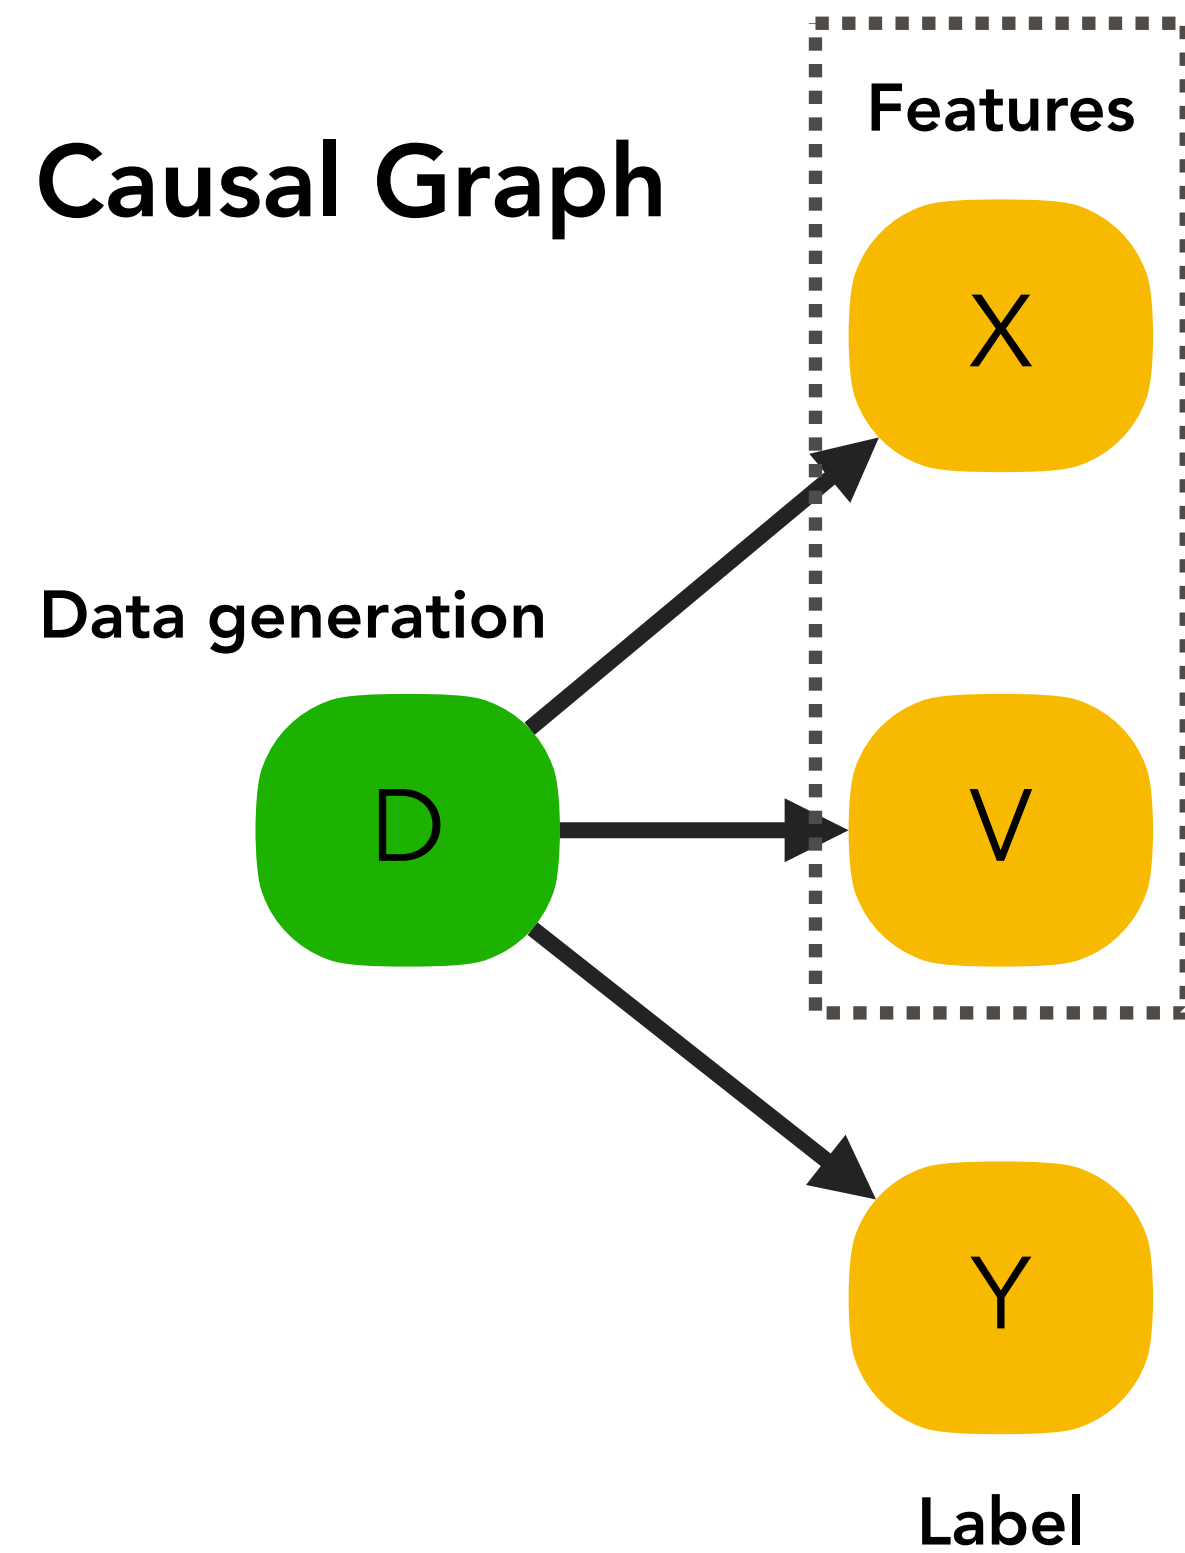
## Calibration:

Across all patients who  predicts there is a 80% chance they need surgery, it truly happens 80% of them needs surgery

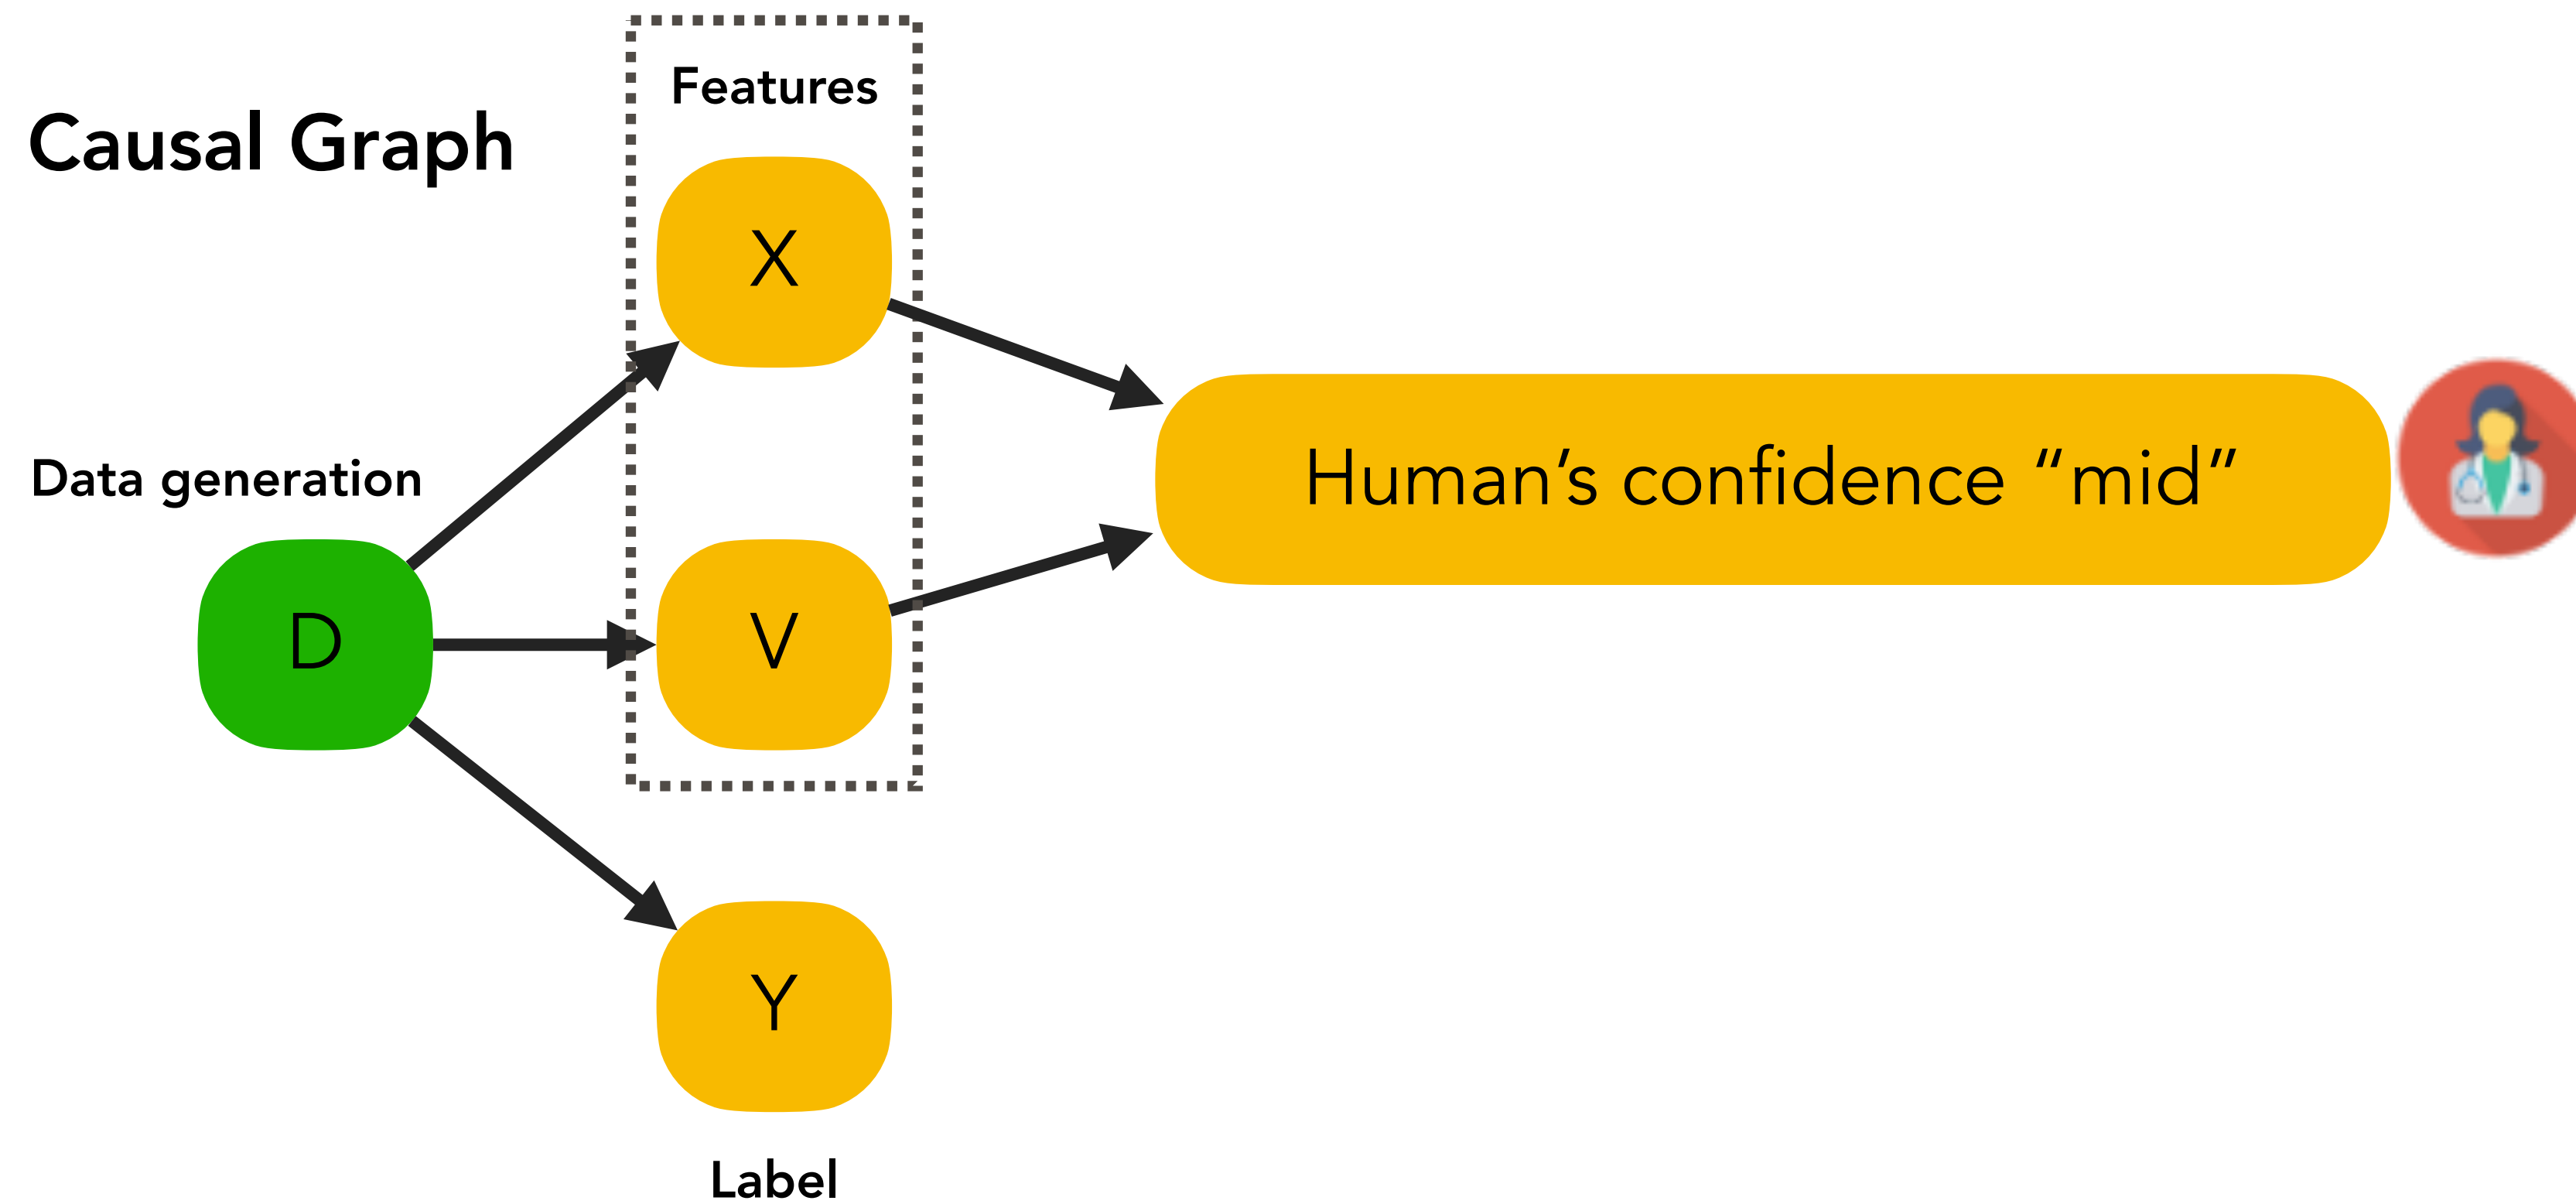
Counterfactual reasoning reveals that the way in which machine learning models compute confidence values today is problematic



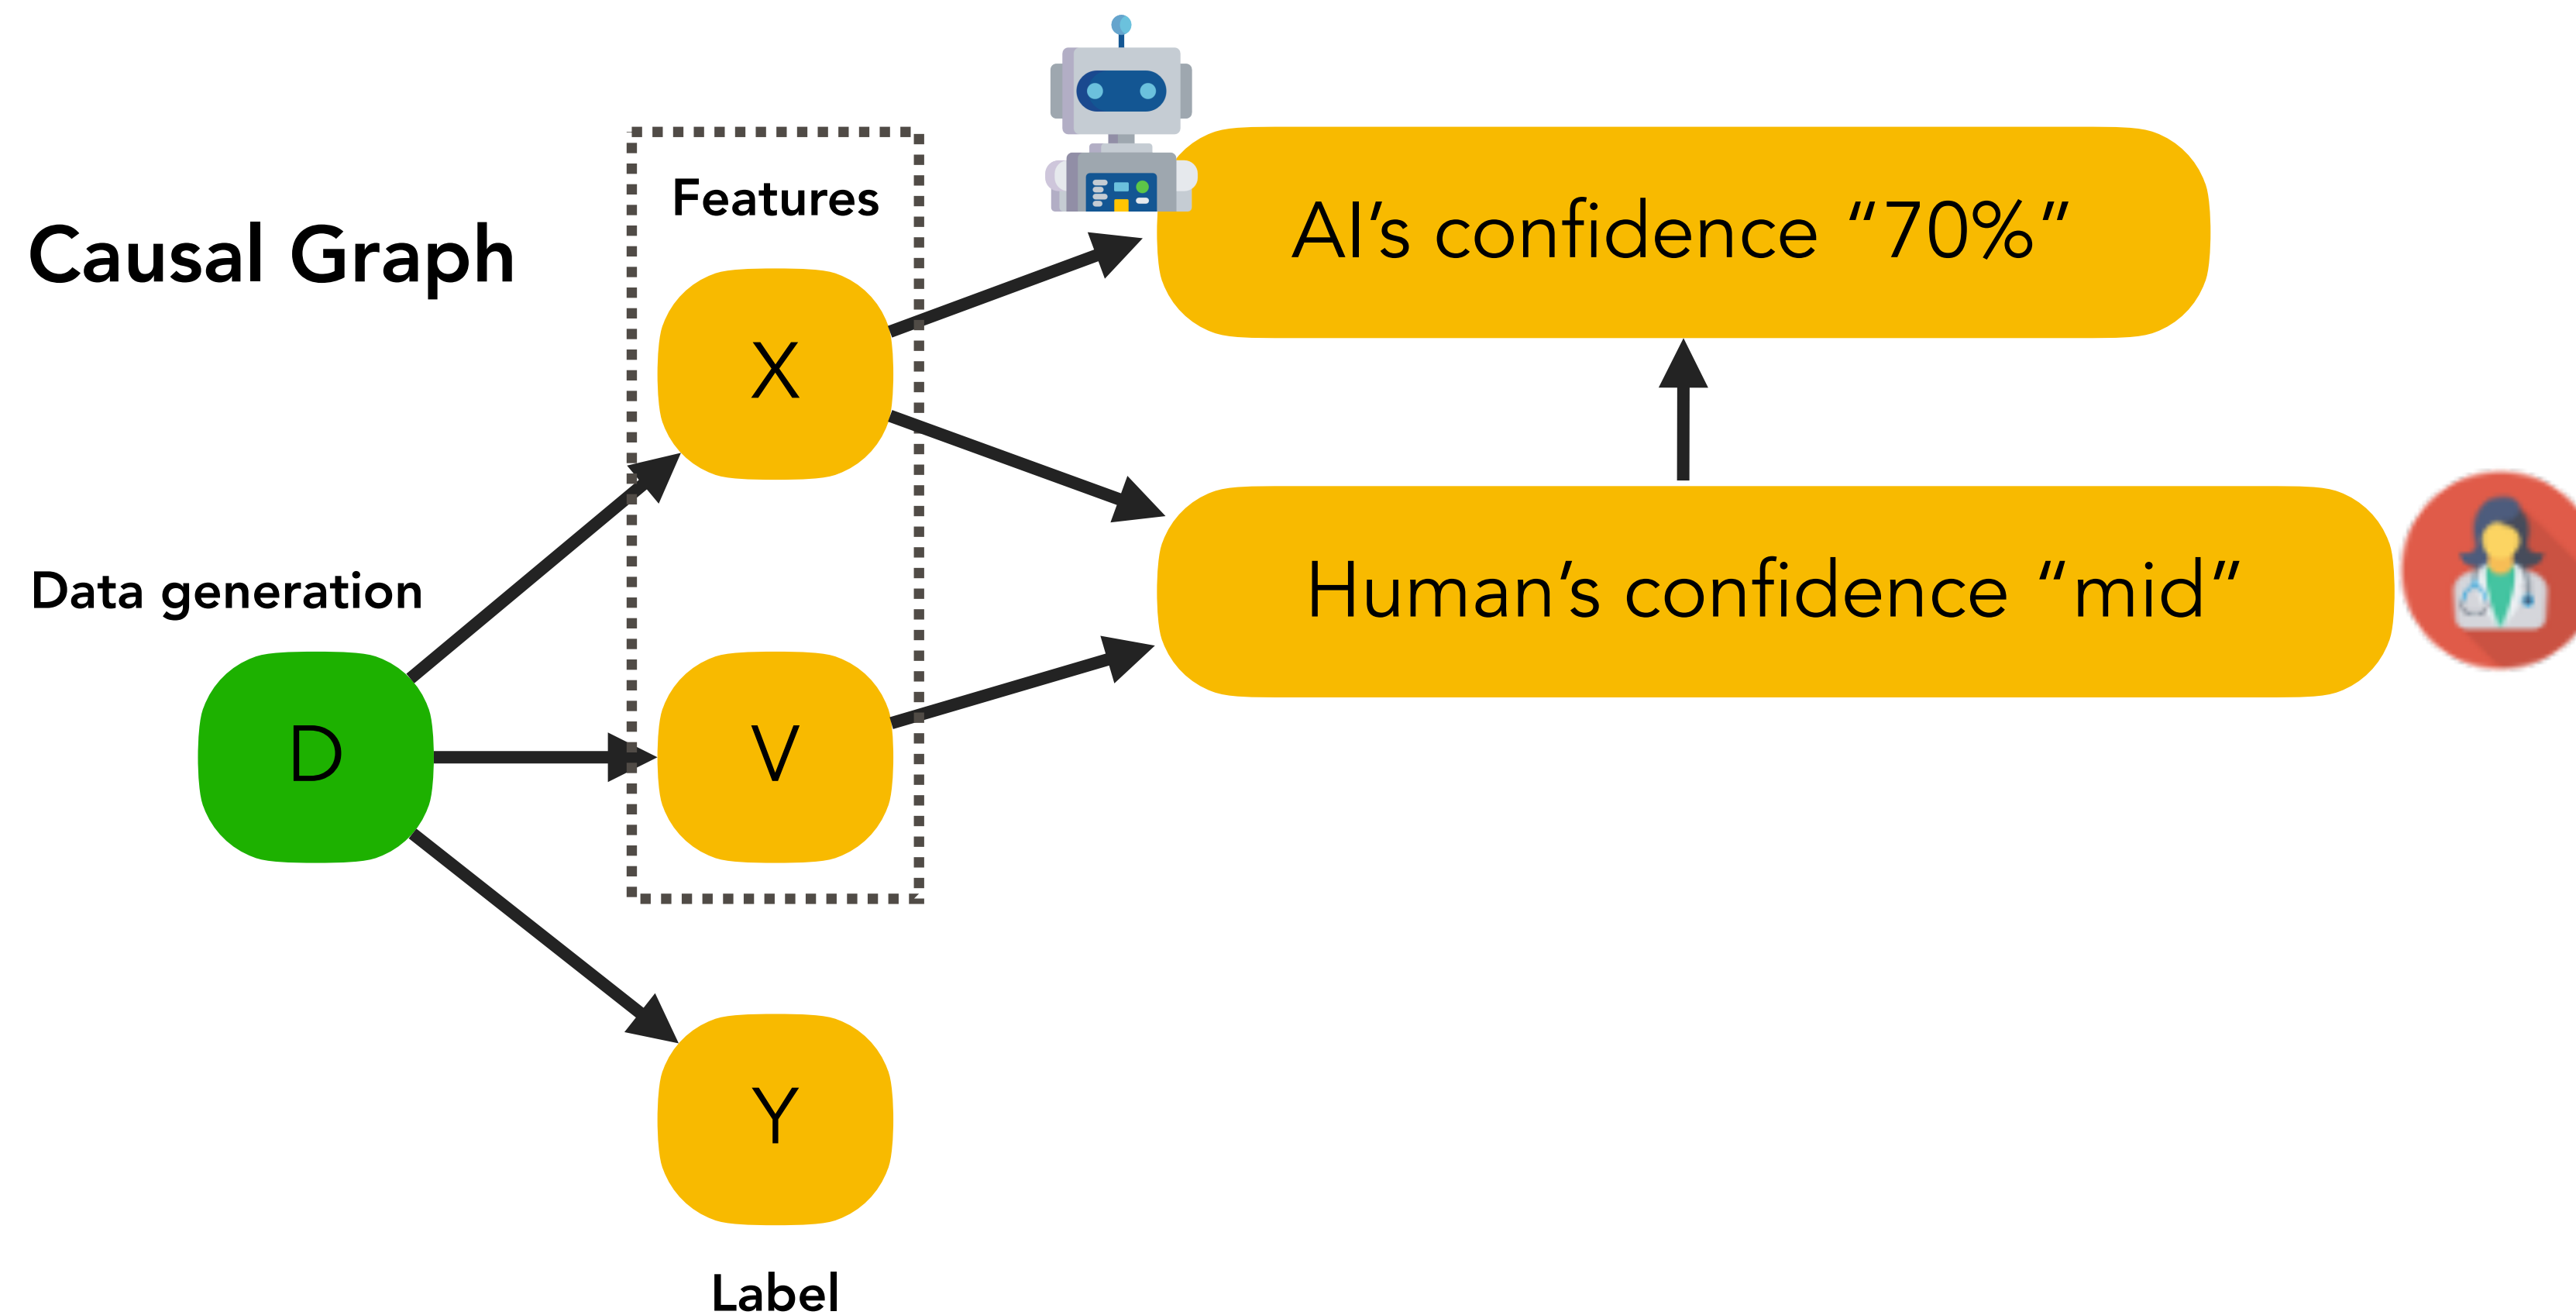
# Calibration



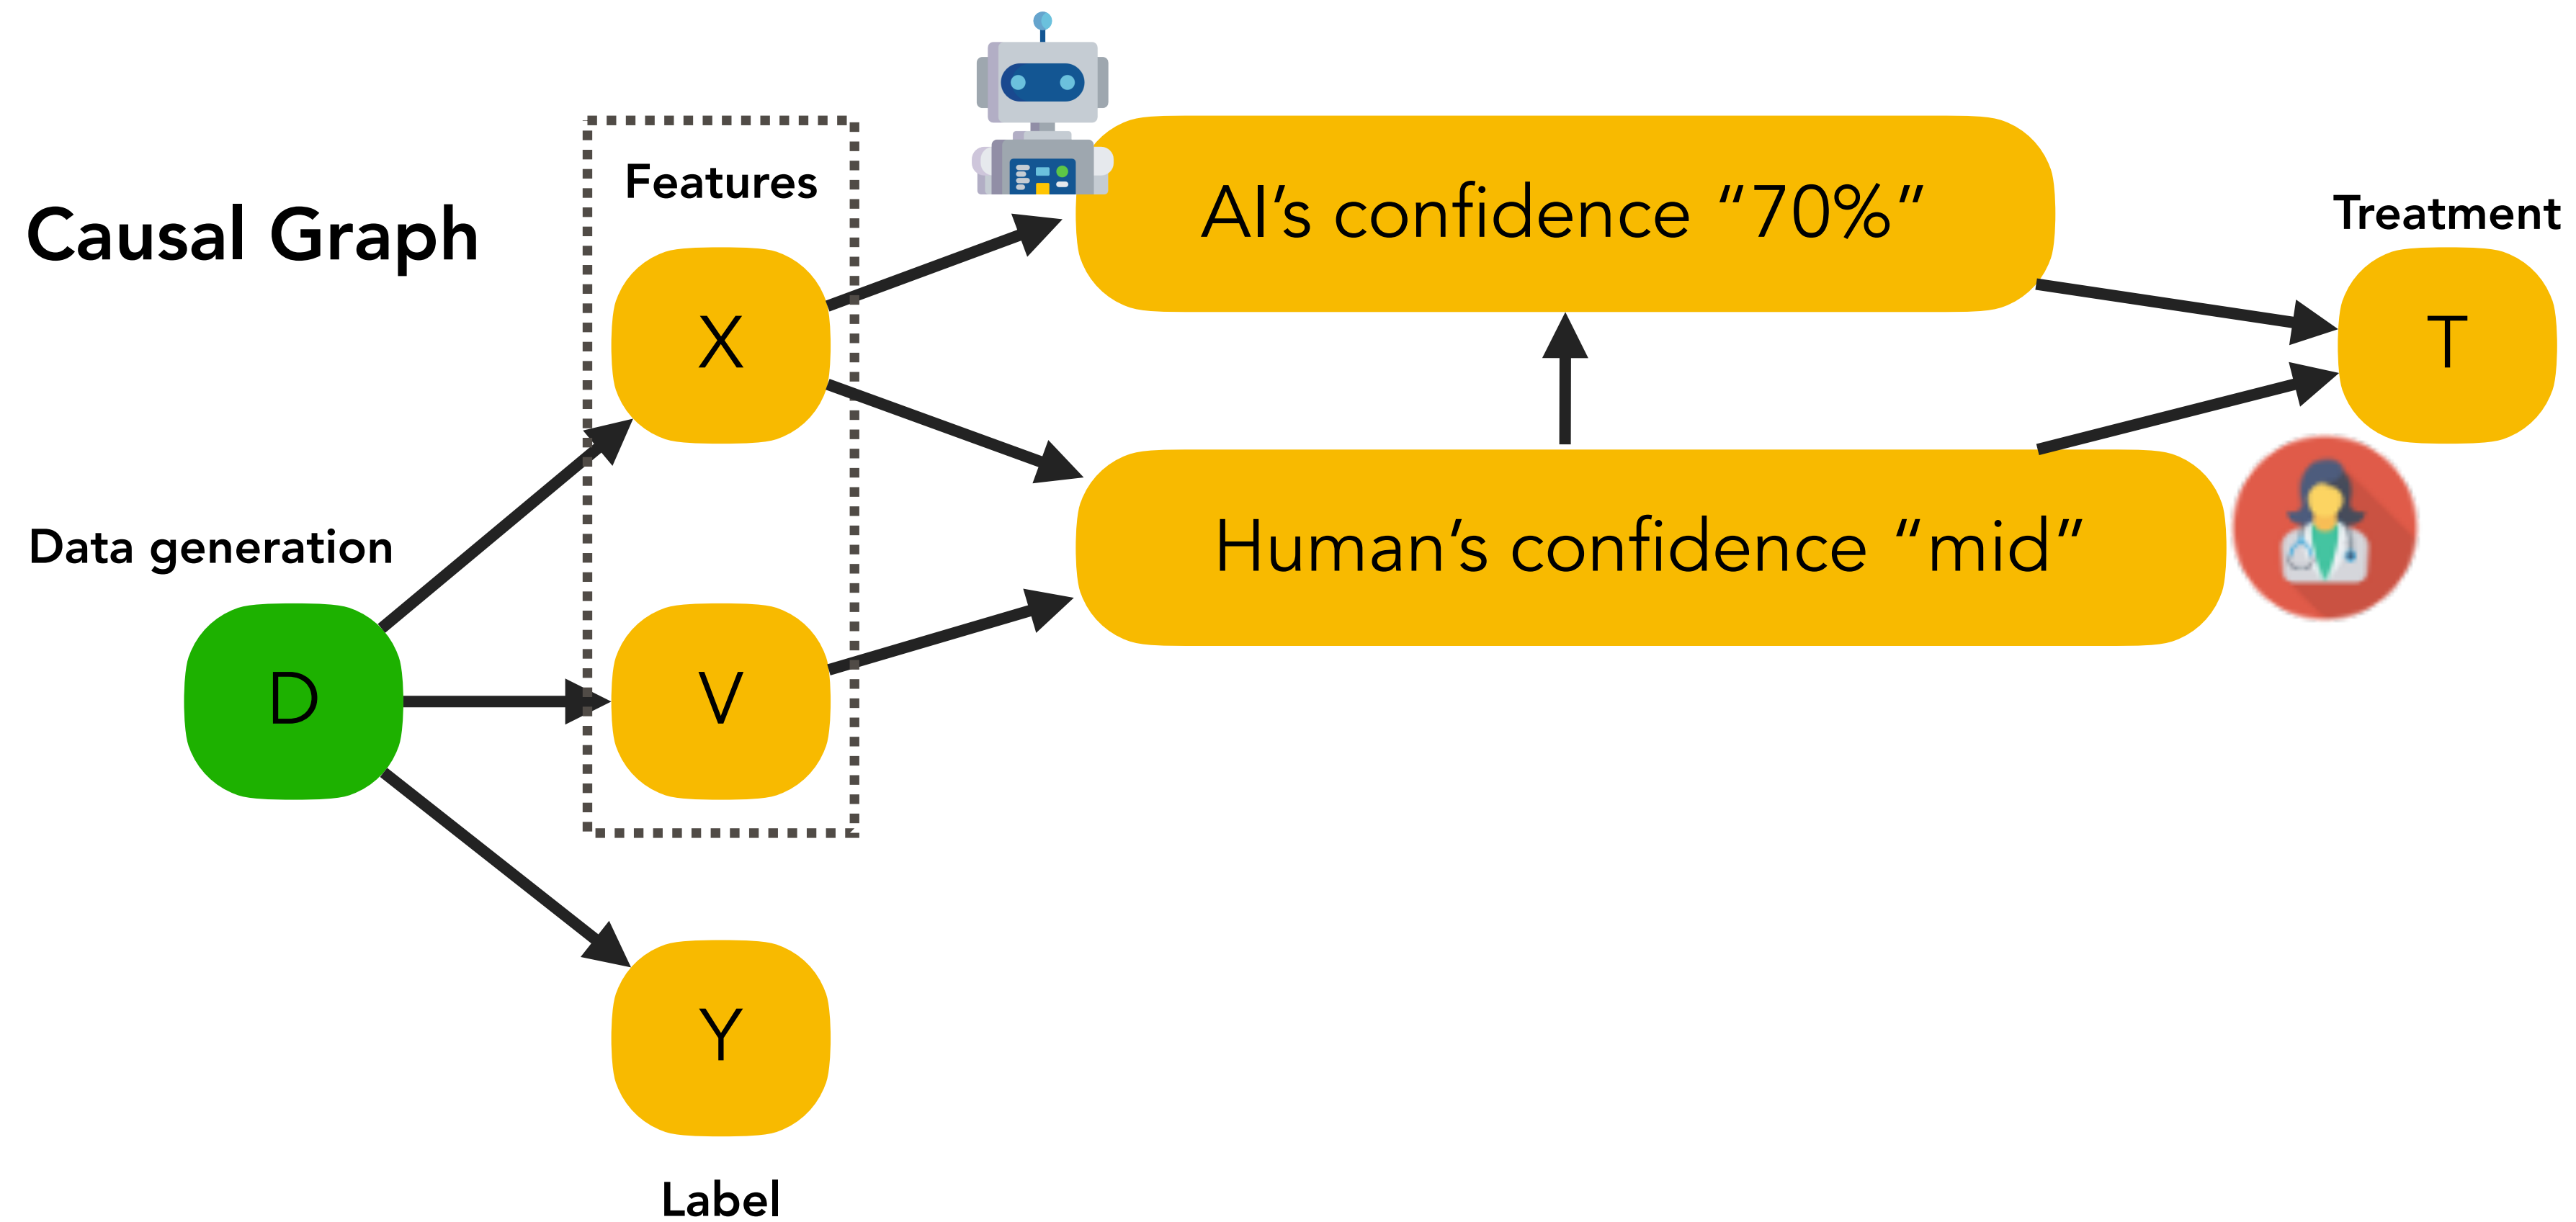
# Calibration



# Calibration

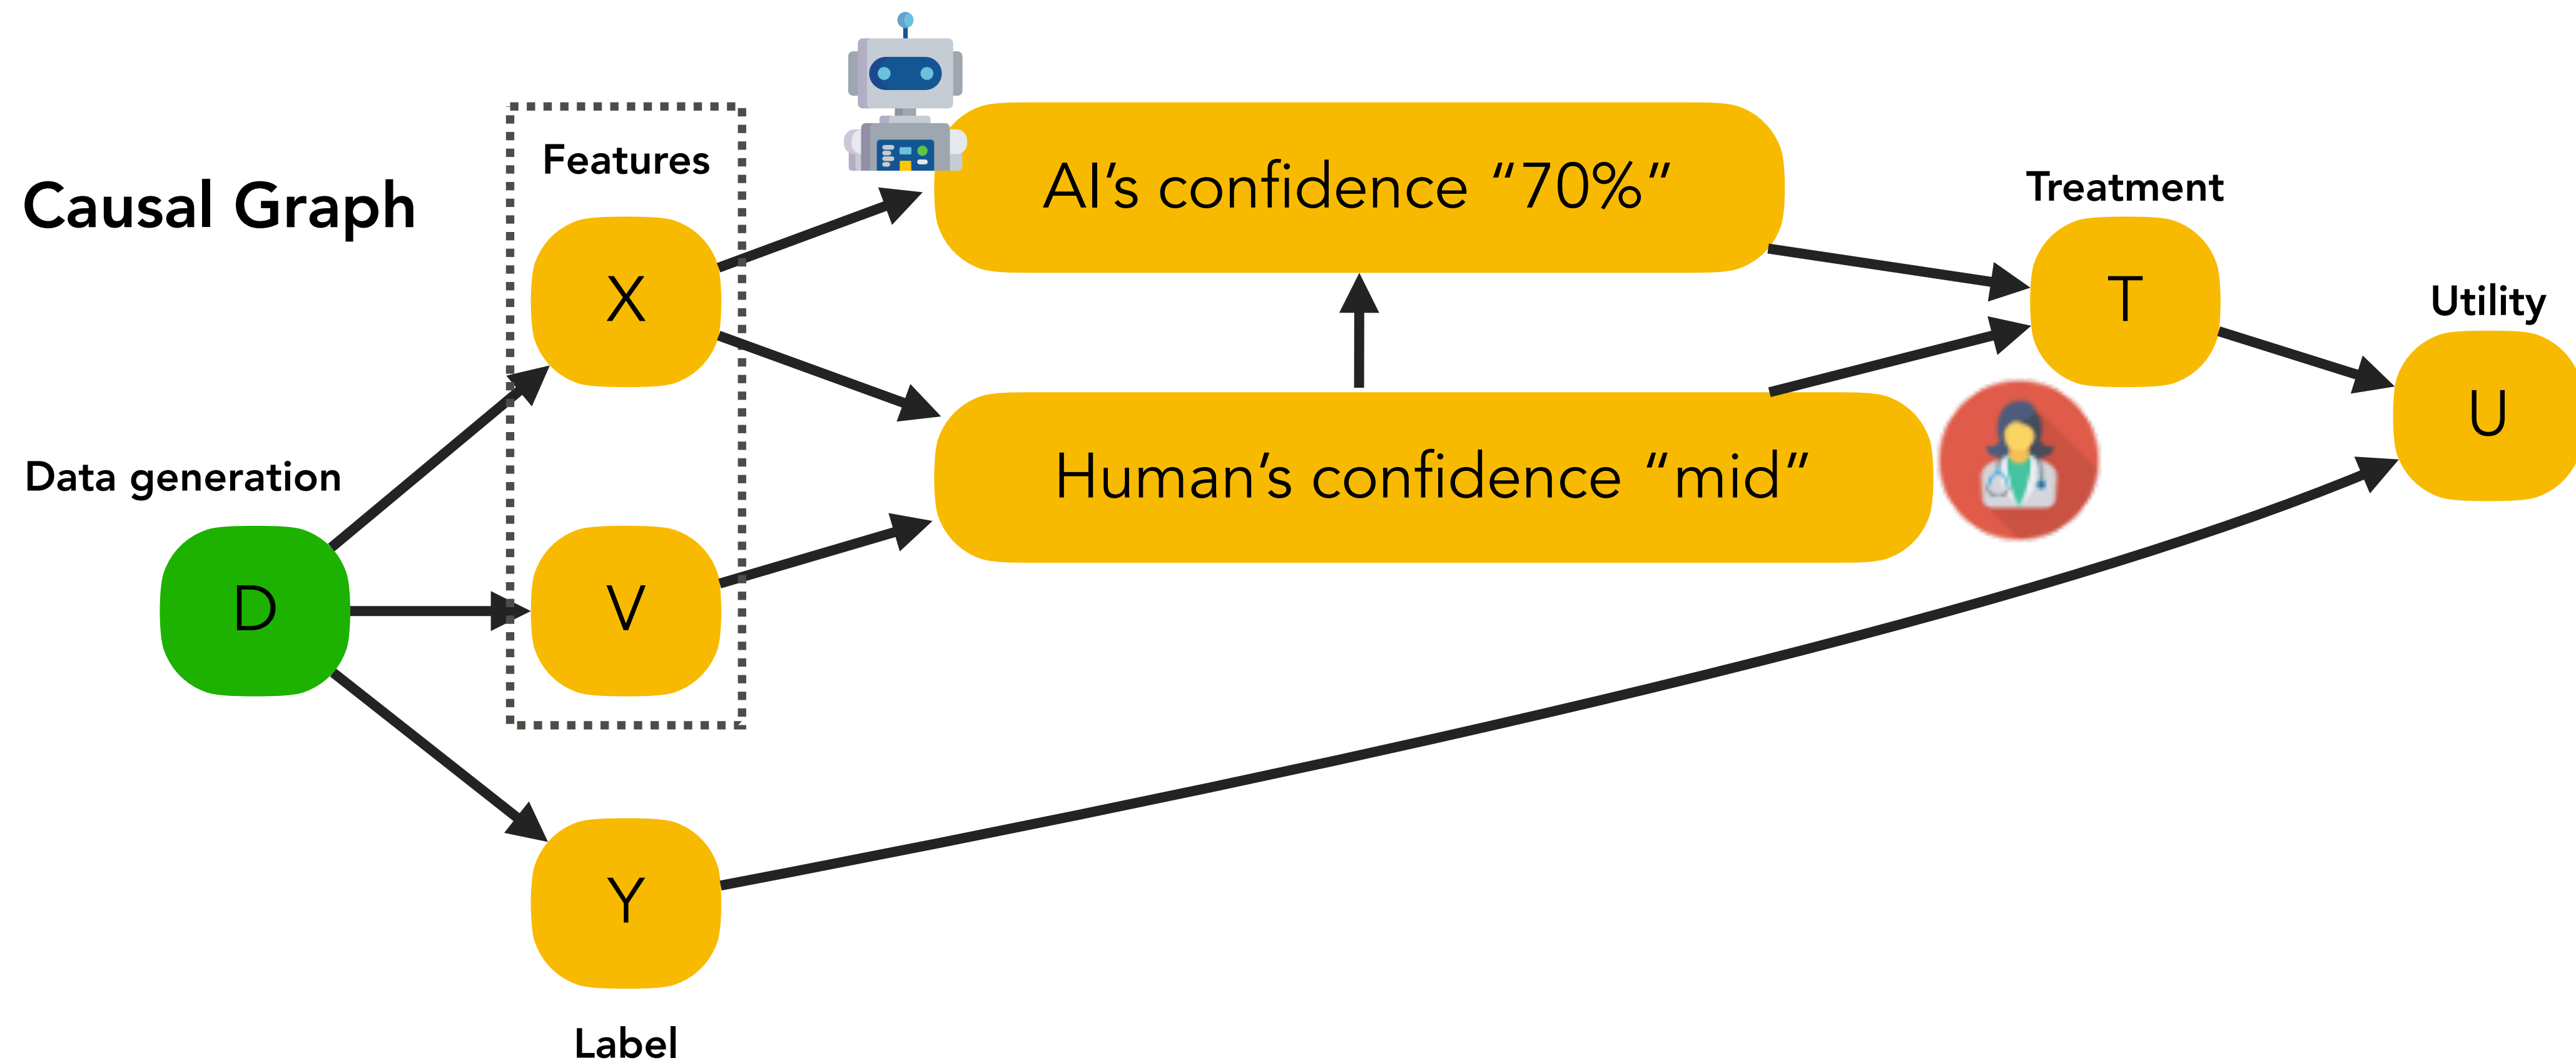


# Calibration



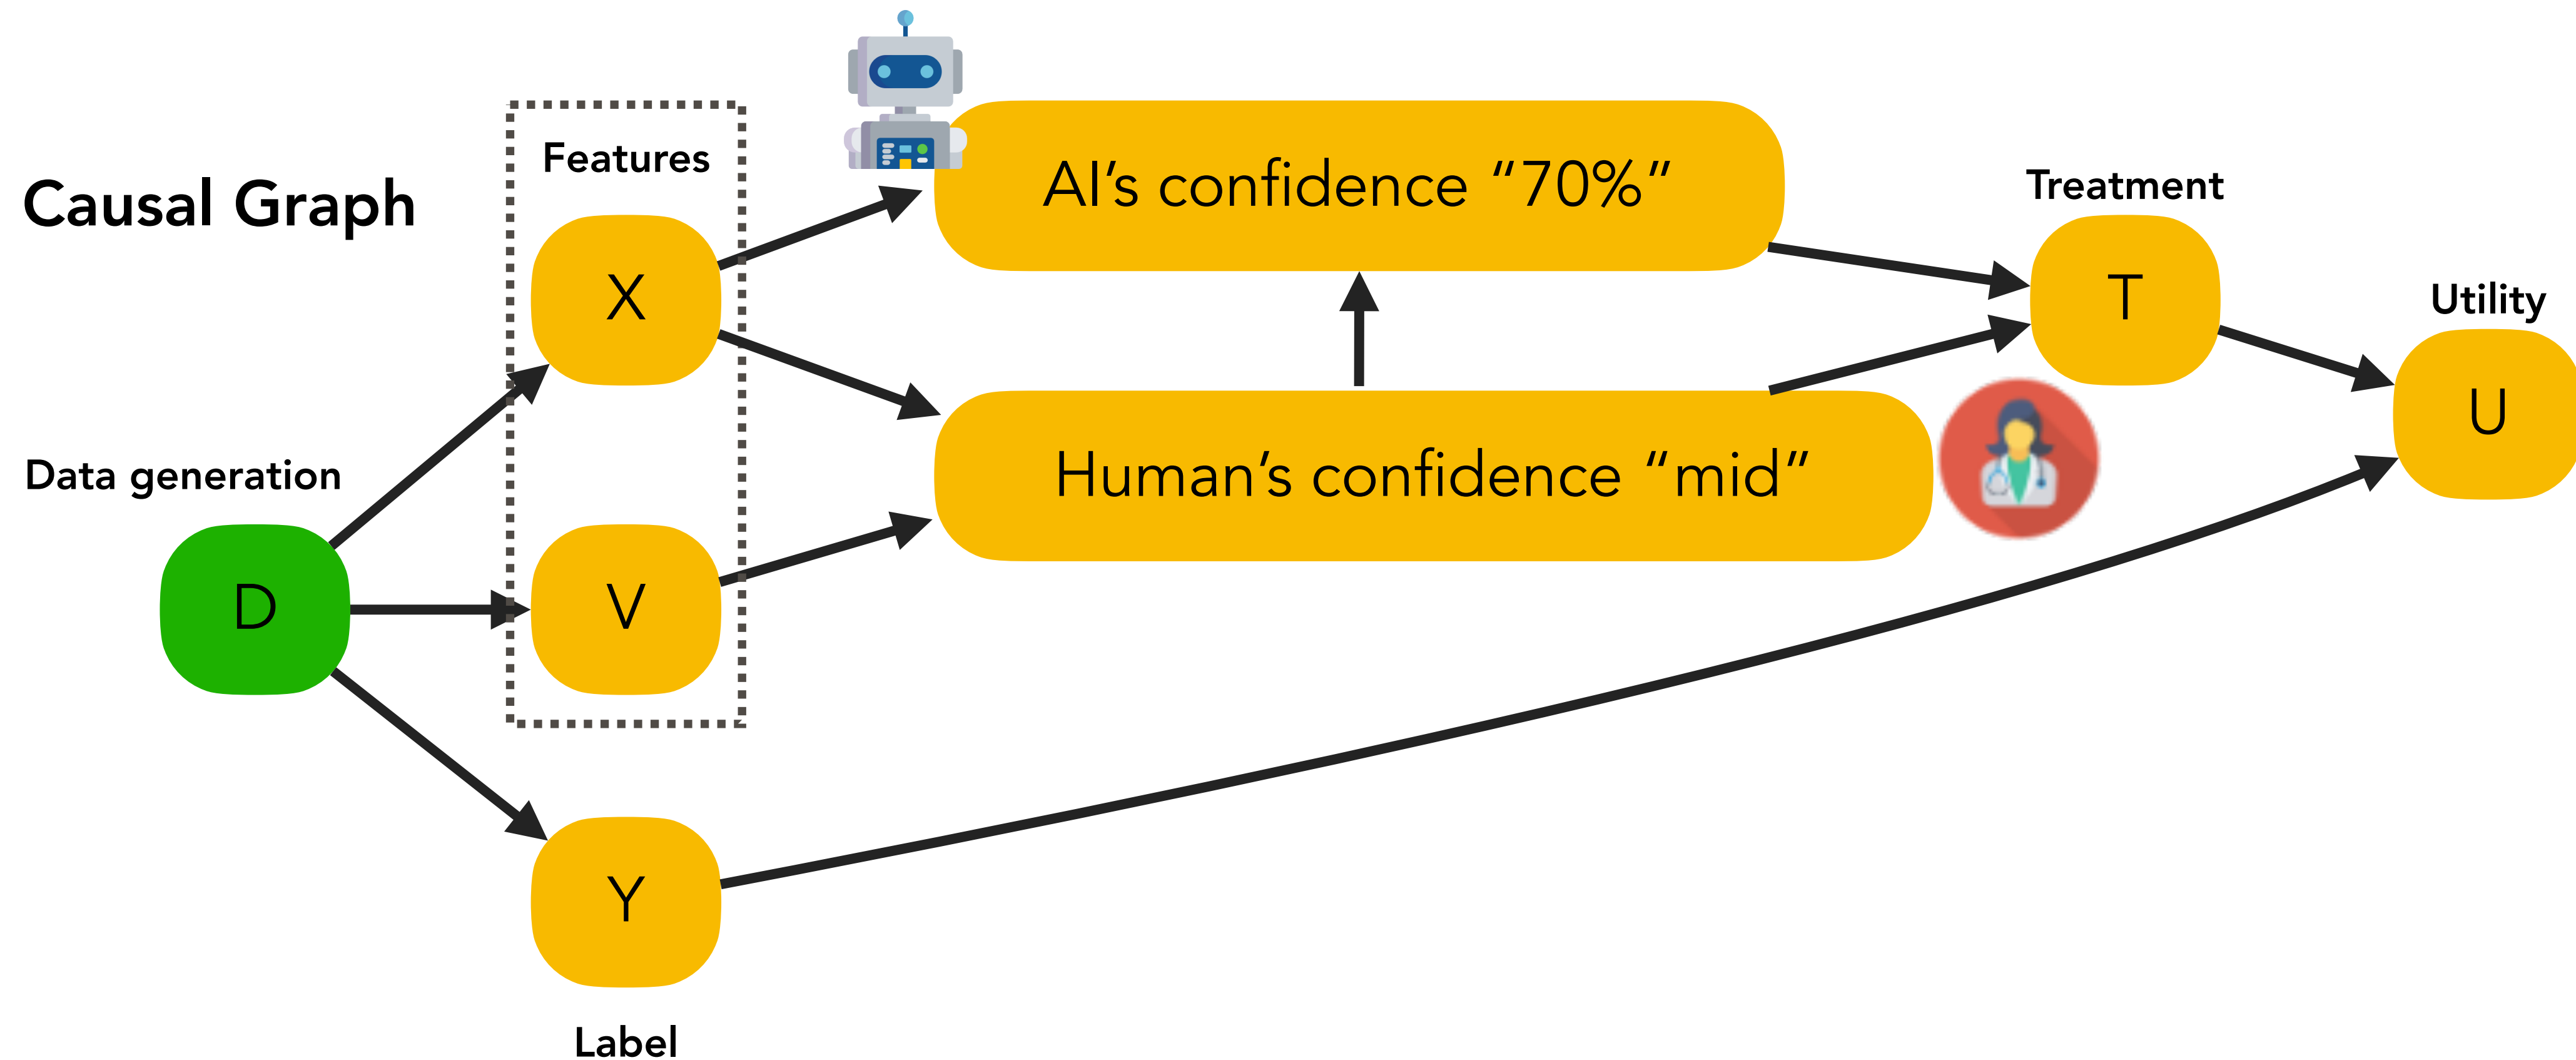


# Calibration





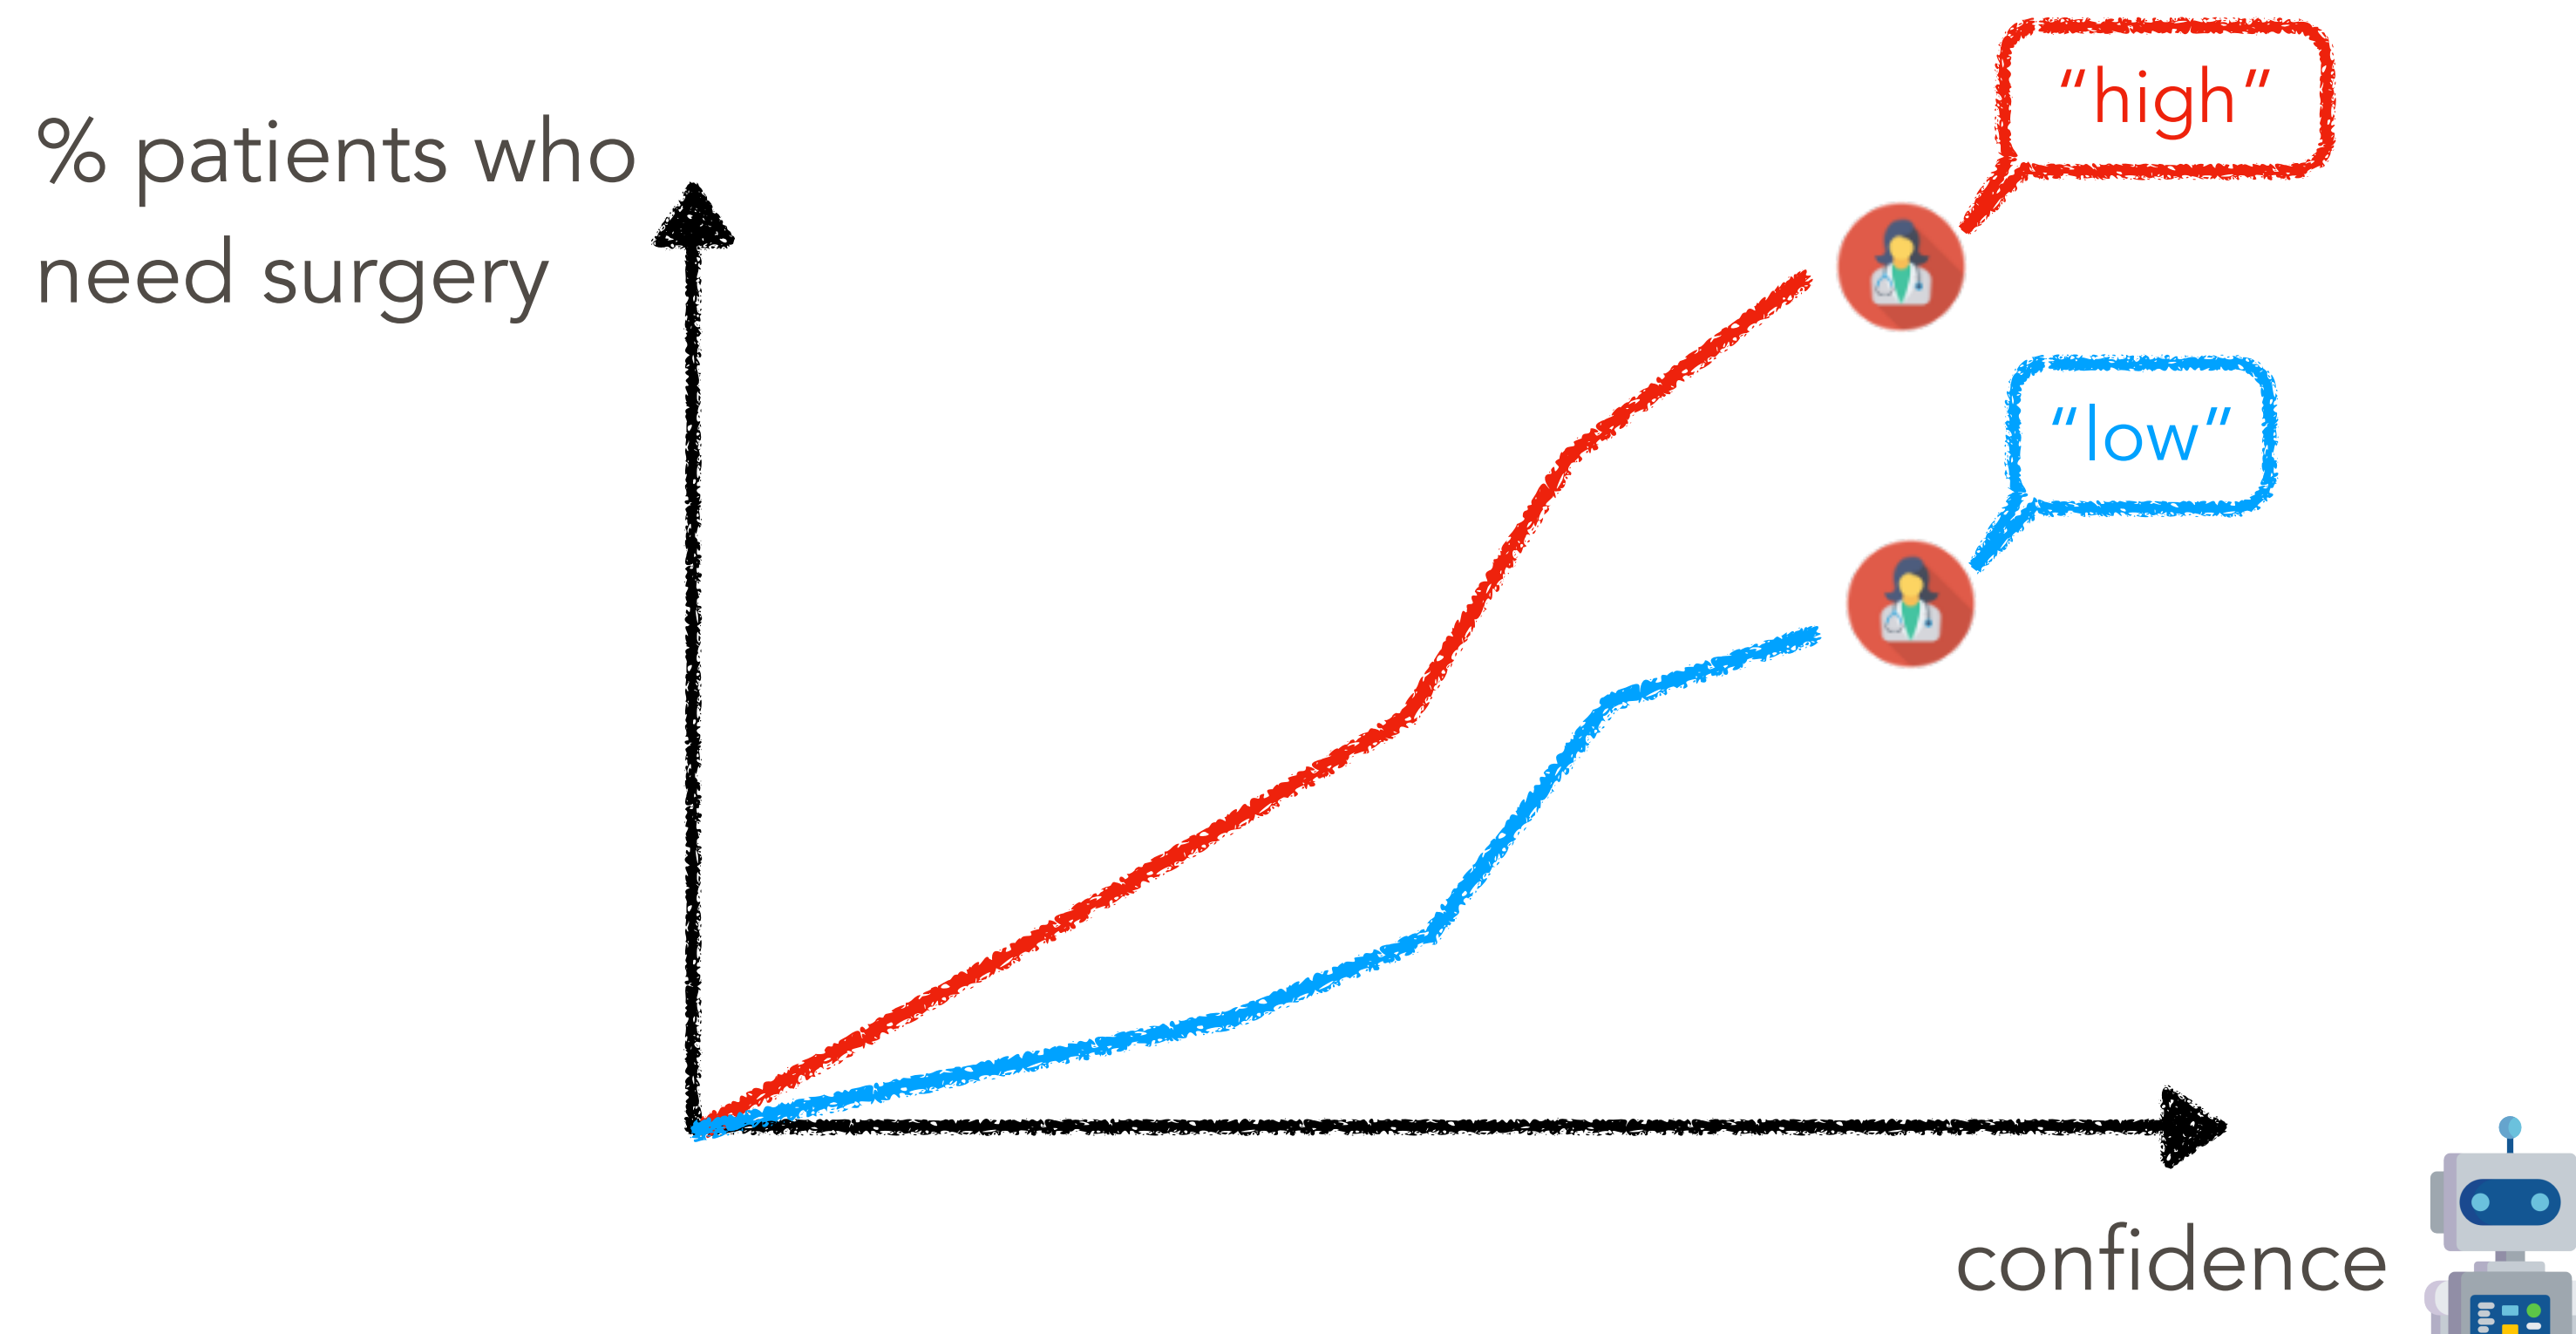
# Calibration



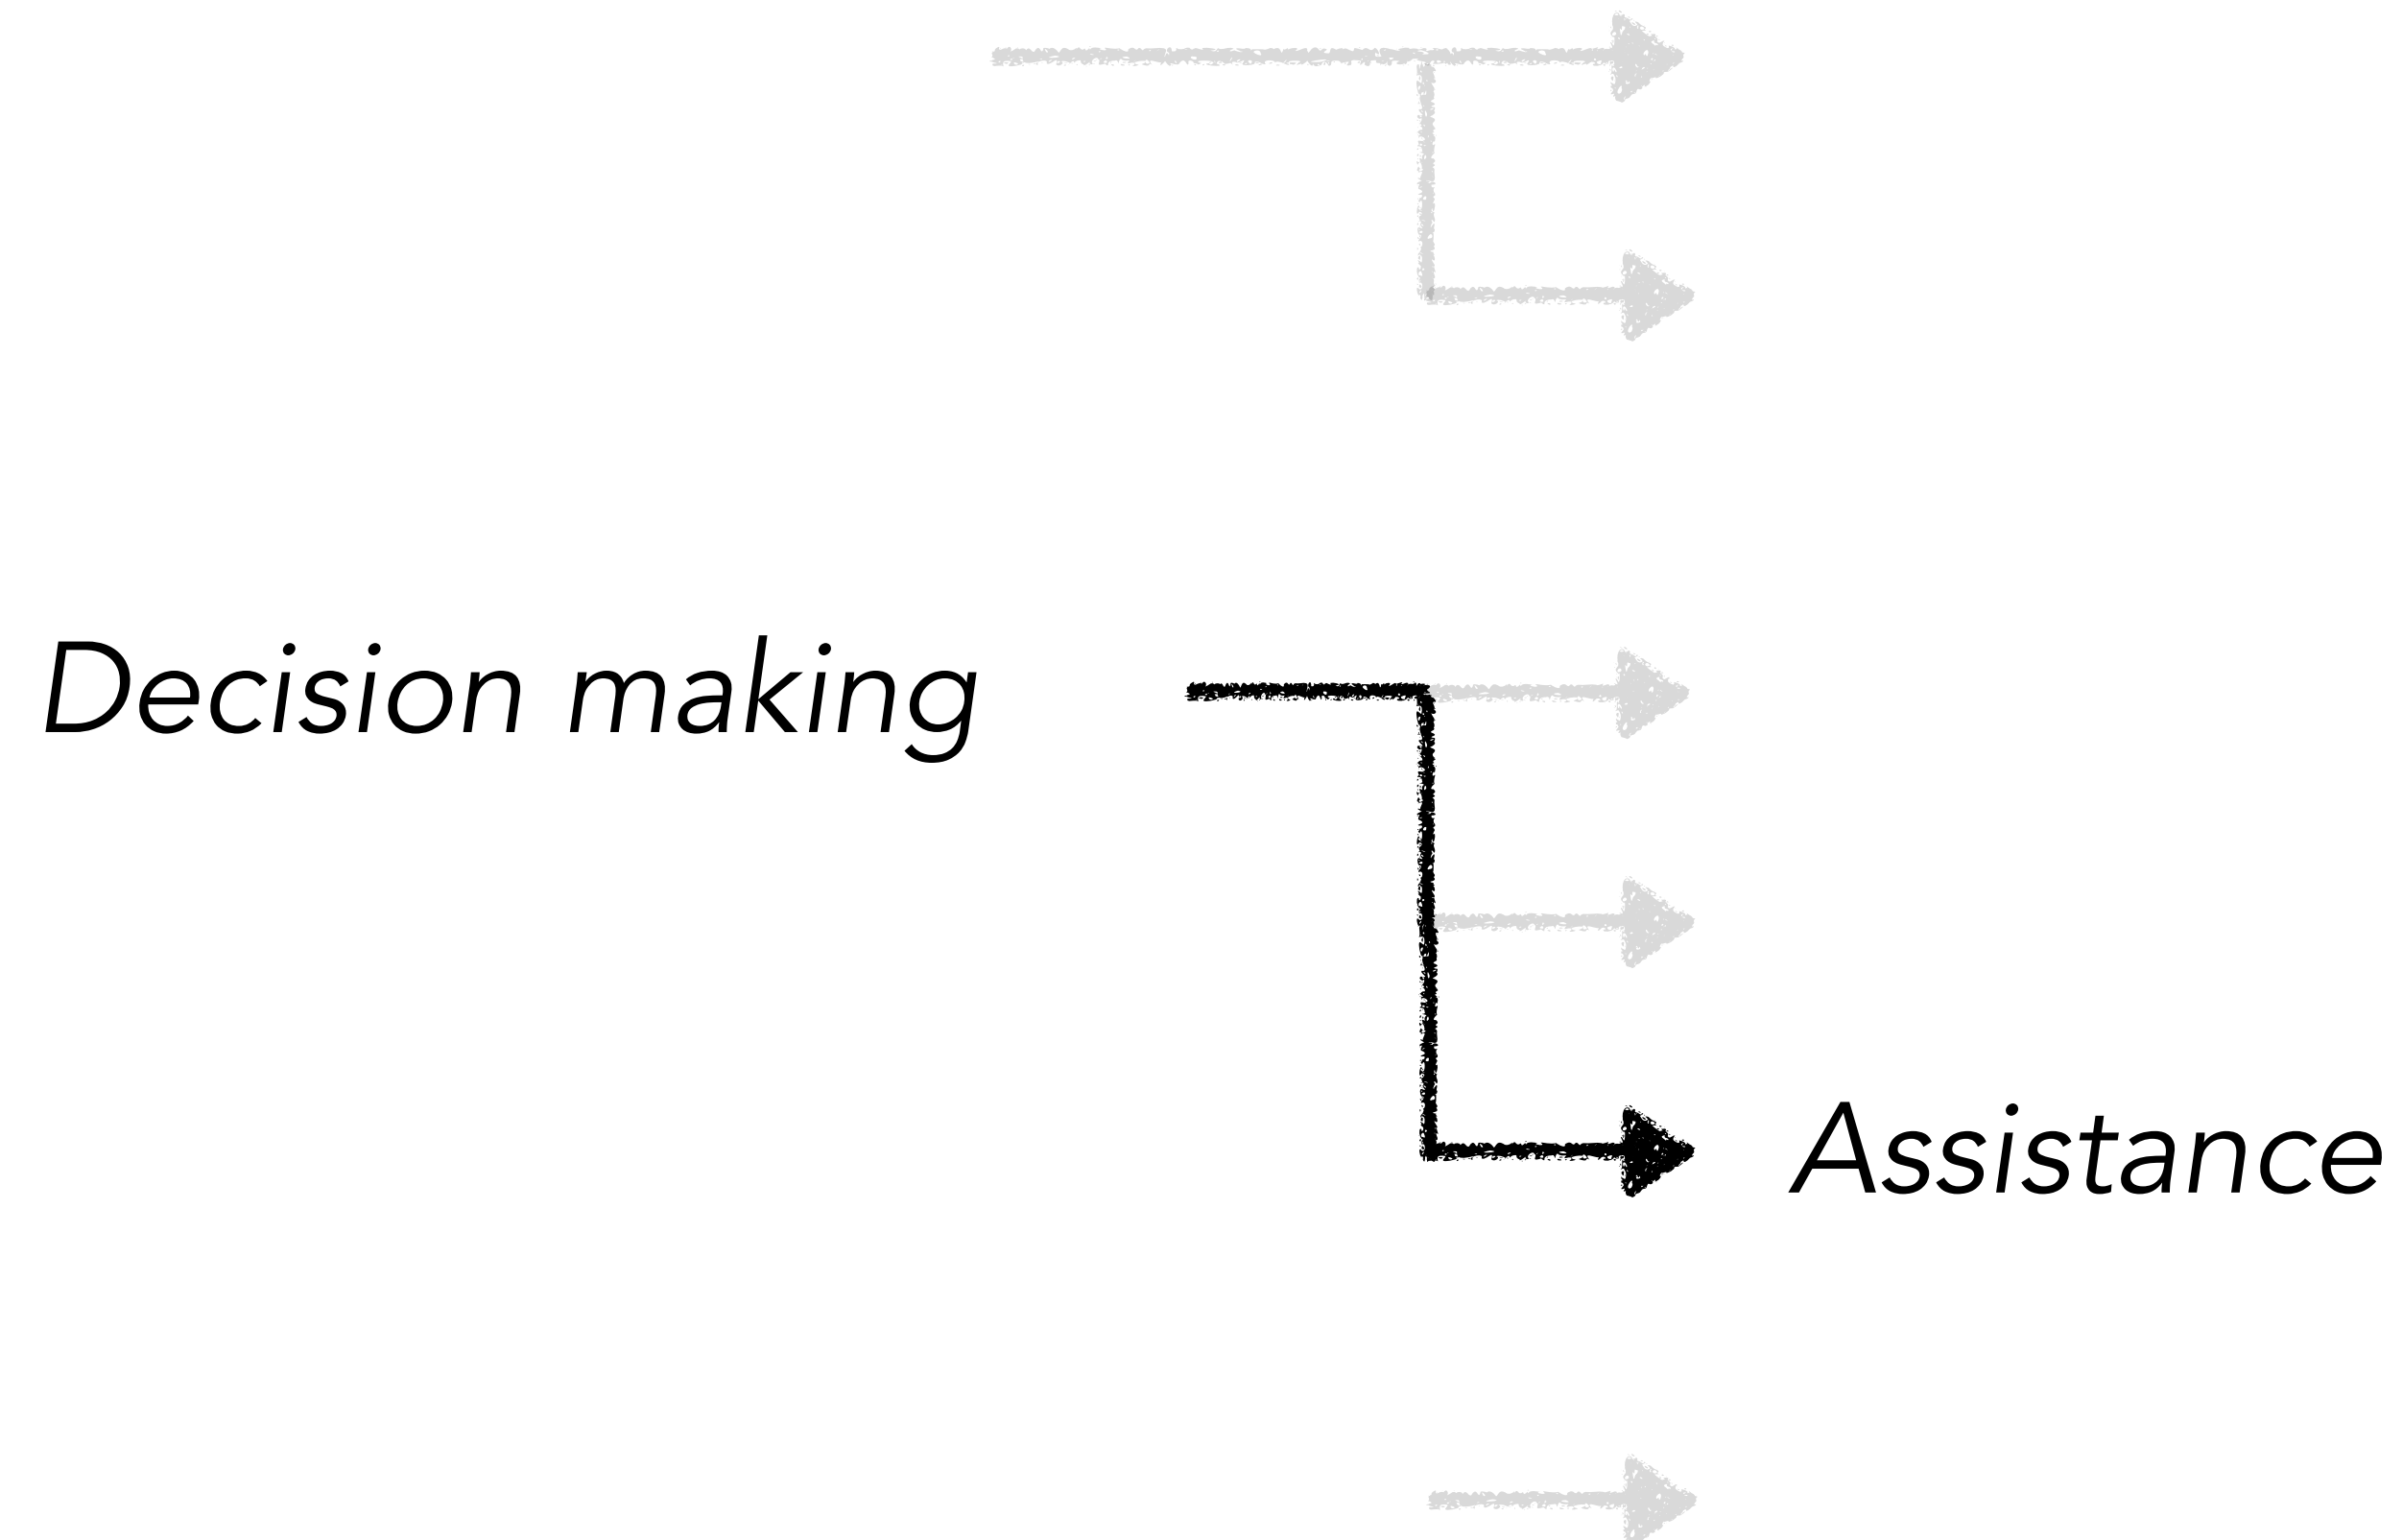
There exist instances of this decision making process in which any monotonic decision policy based on calibrated AI predictions is suboptimal.

# Calibration

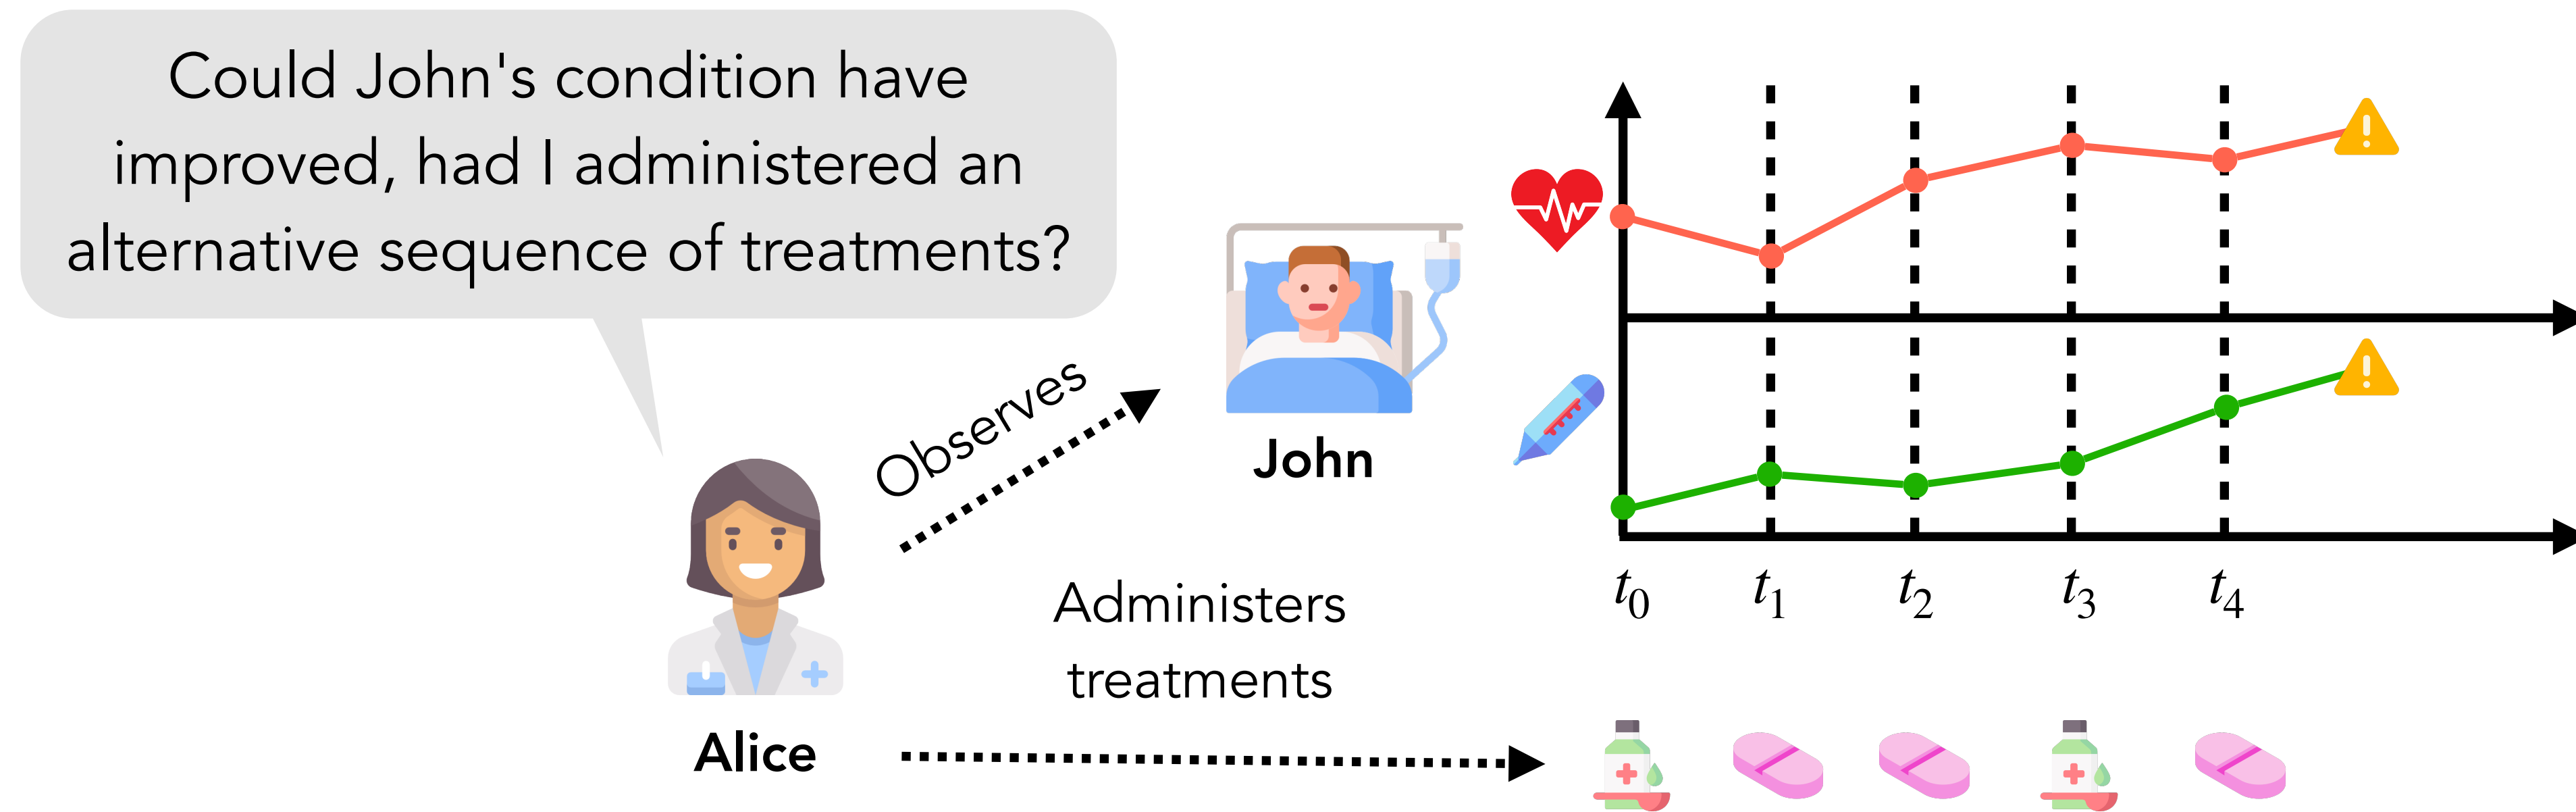
To make sure the level of trust the optimal decision maker needs to place on predictions is (always) monotone on the confidence values, one can use **multiplicative calibration**.



# Use cases of counterfactuals in machine learning



# AI-assisted counterfactuals in sequential decision making



# Alternative sequence of treatments as counterfactuals

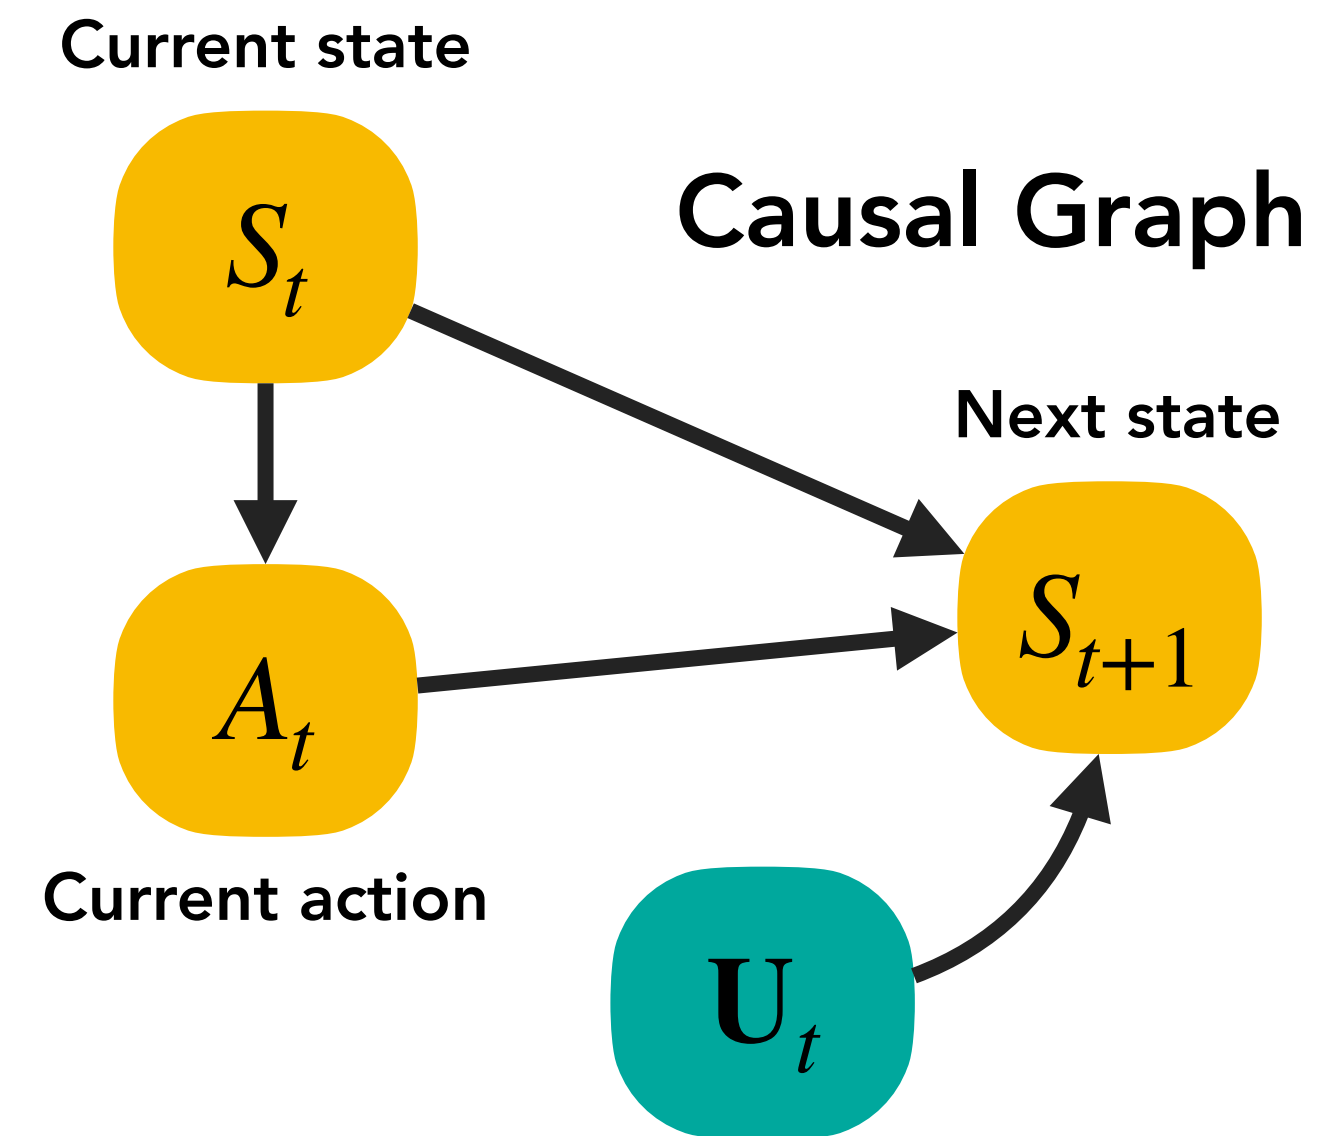
## Structural Causal Model $\mathcal{M}$

$$S_{t+1} := g_S(S_t, A_t, \mathbf{U}_t)$$

$$A_t := g_A(S_t, \mathbf{V}_t)$$

$$\mathbf{U}_t \sim P(\mathbf{U})$$

$$\mathbf{V}_t \sim P(\mathbf{V})$$



...

# Alternative sequence of treatments as counterfactuals

## Structural Causal Model $\mathcal{M}$

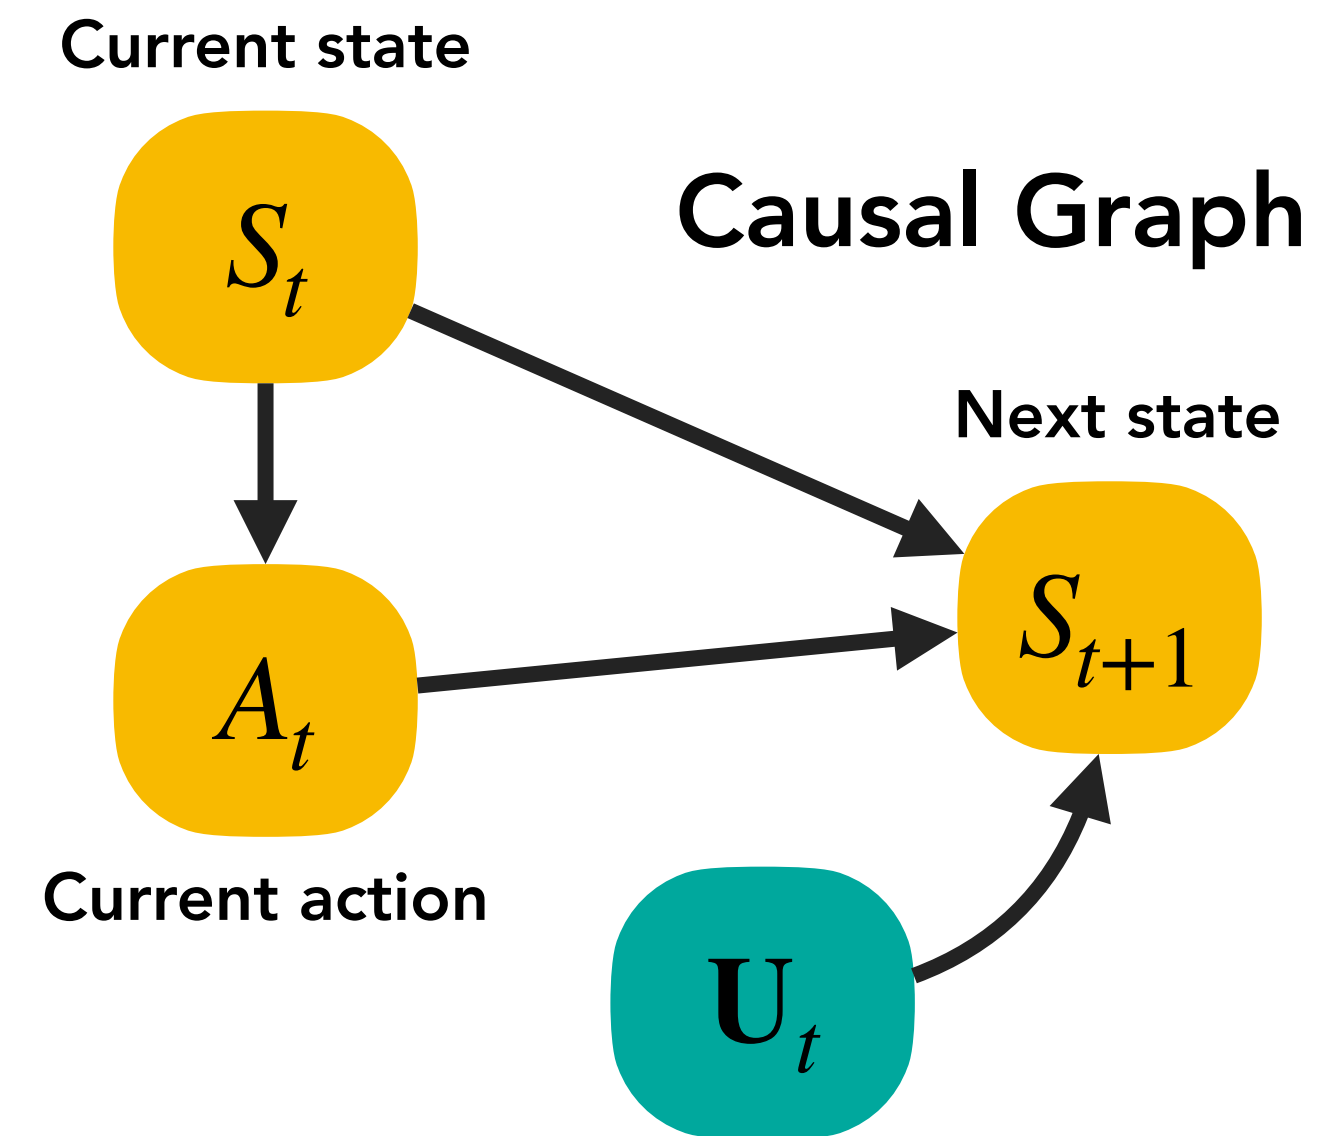
$$S_{t+1} := g_S(S_t, A_t, \mathbf{U}_t)$$

$$A_t := g_A(S_t, \mathbf{V}_t)$$

$$\mathbf{U}_t \sim P(\mathbf{U})$$

$$\mathbf{V}_t \sim P(\mathbf{V})$$

At state  $S_t = s_t$ , the doctor took action  $A_t = a_t$ , what would have happened had the doctor taken action  $a' \neq a_t$ ?



...

# Alternative sequence of treatments as counterfactuals

**Modified Structural Causal Model**  $\mathcal{M}_{\{S_t=s_t, A_t=a_t\}}$

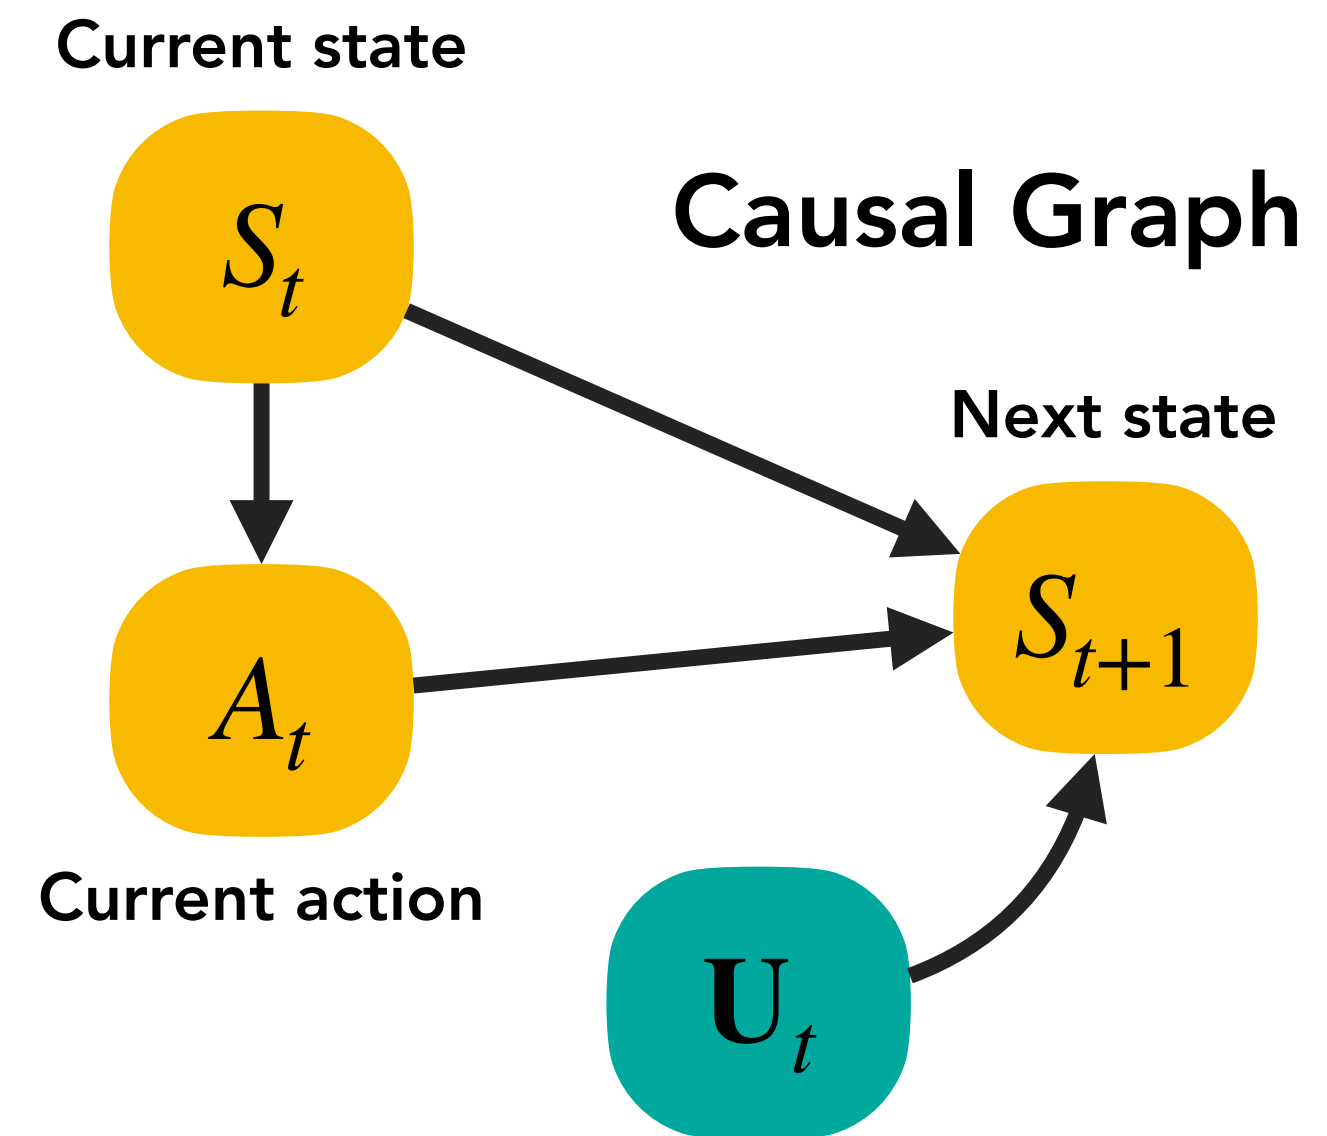
$$S_{t+1} := g_S(S_t, A_t, \mathbf{U}_t)$$

$$A_t := g_A(S_t, \mathbf{V}_t)$$

$$\mathbf{U}_t \sim P(\mathbf{U} \mid S_t = s_t, A_t = a_t)$$

$$\mathbf{V}_t \sim P(\mathbf{V} \mid S_t = s_t)$$

Posterior distribution  
of the noises



**At state  $S_t = s_t$ , the doctor took action  $A_t = a_t$ , what would have happened had the doctor taken action  $a' \neq a_t$ ?**

...

# Alternative sequence of treatments as counterfactuals

**Modified Structural Causal Model**  $\mathcal{M}_{\{S_t=s_t, A_t=a_t\}}$

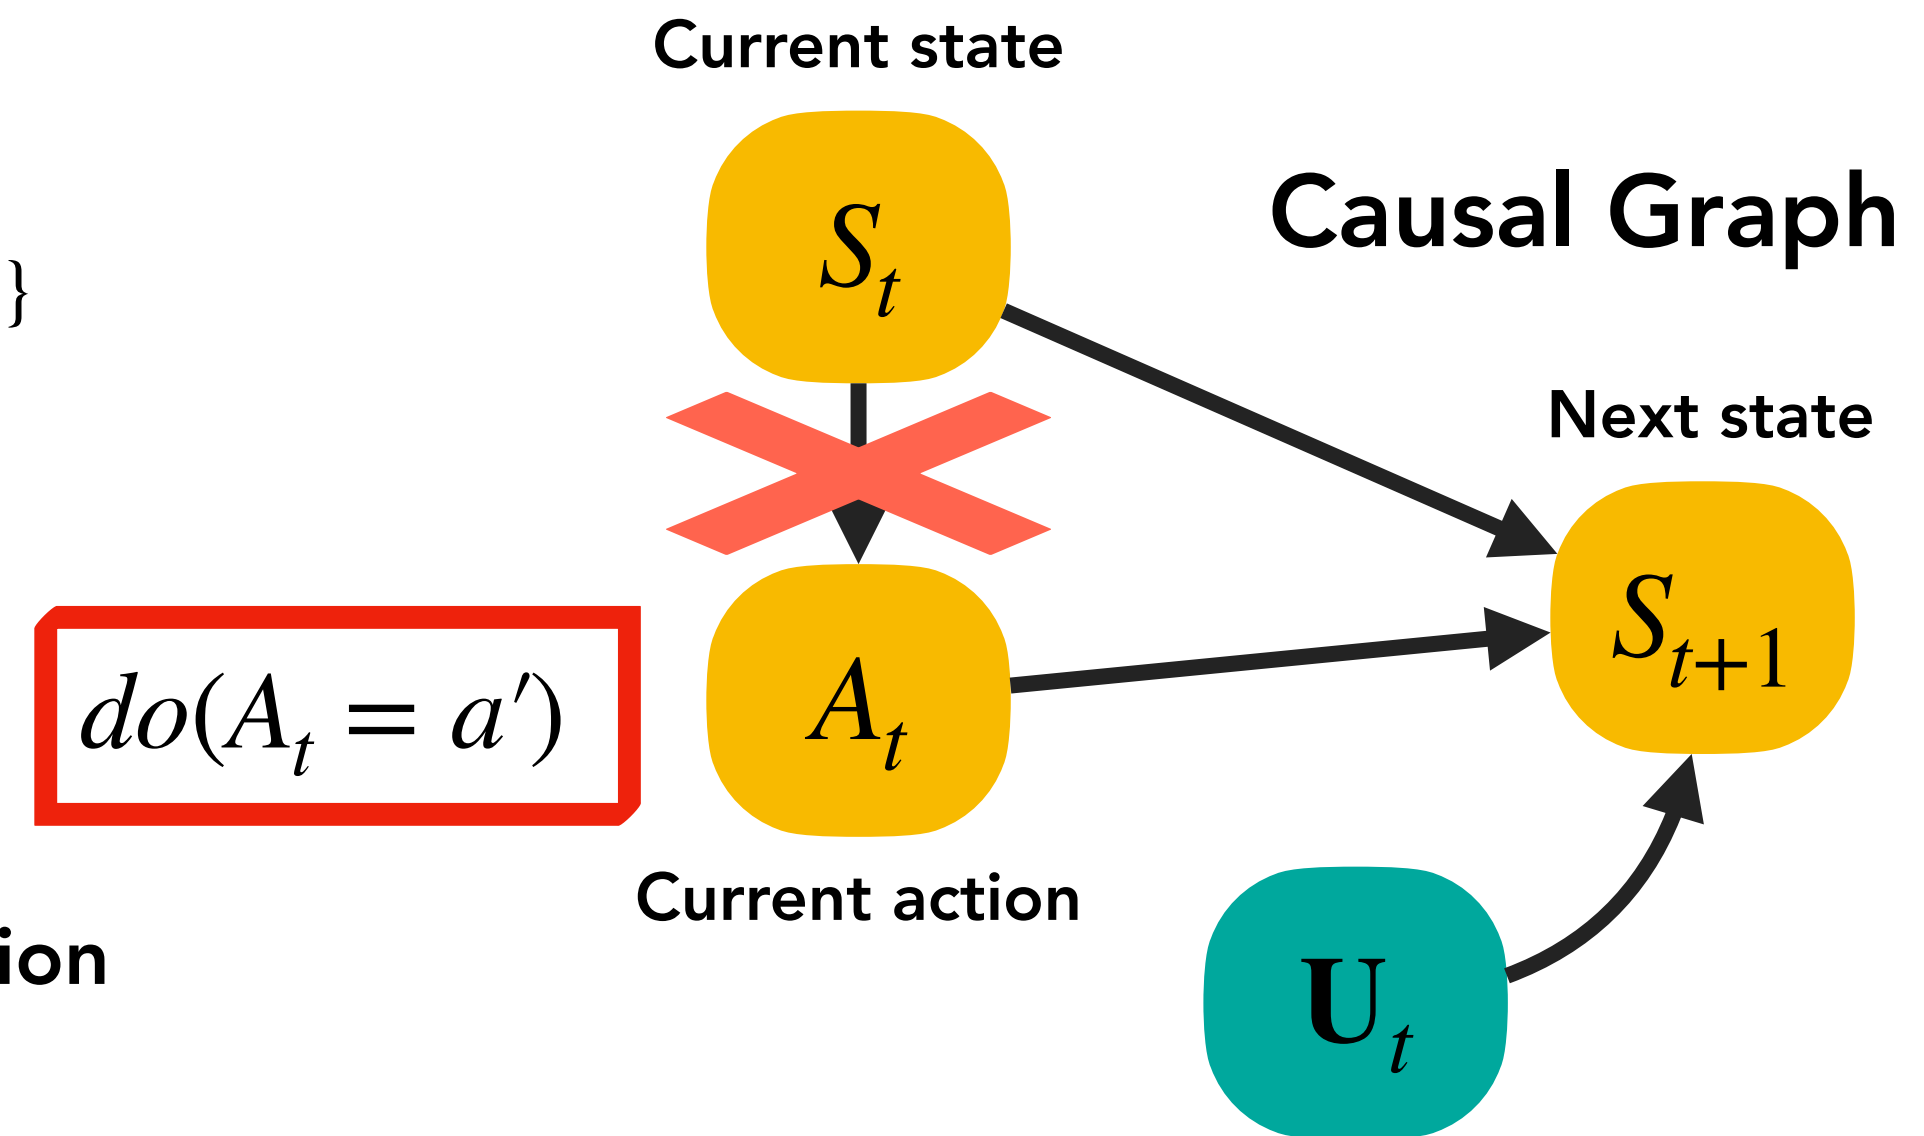
$$S_{t+1} := g_S(S_t, A_t, \mathbf{U}_t)$$

$$A_t := \cancel{g_A(S_t, \mathbf{V}_t)} \quad A_t := a'$$

$$\mathbf{U}_t \sim P(\mathbf{U} \mid S_t = s_t, A_t = a_t)$$

$$\mathbf{V}_t \sim P(\mathbf{V} \mid S_t = s_t)$$

Posterior distribution  
of the noises



At state  $S_t = s_t$ , the doctor took action  $A_t = a_t$ , what would have happened had the doctor taken action  $a' \neq a_t$ ?

...



# Alternative sequence of treatments as counterfactuals

**Modified Structural Causal Model**  $\mathcal{M}_{\{S_t=s_t, A_t=a_t\}}$

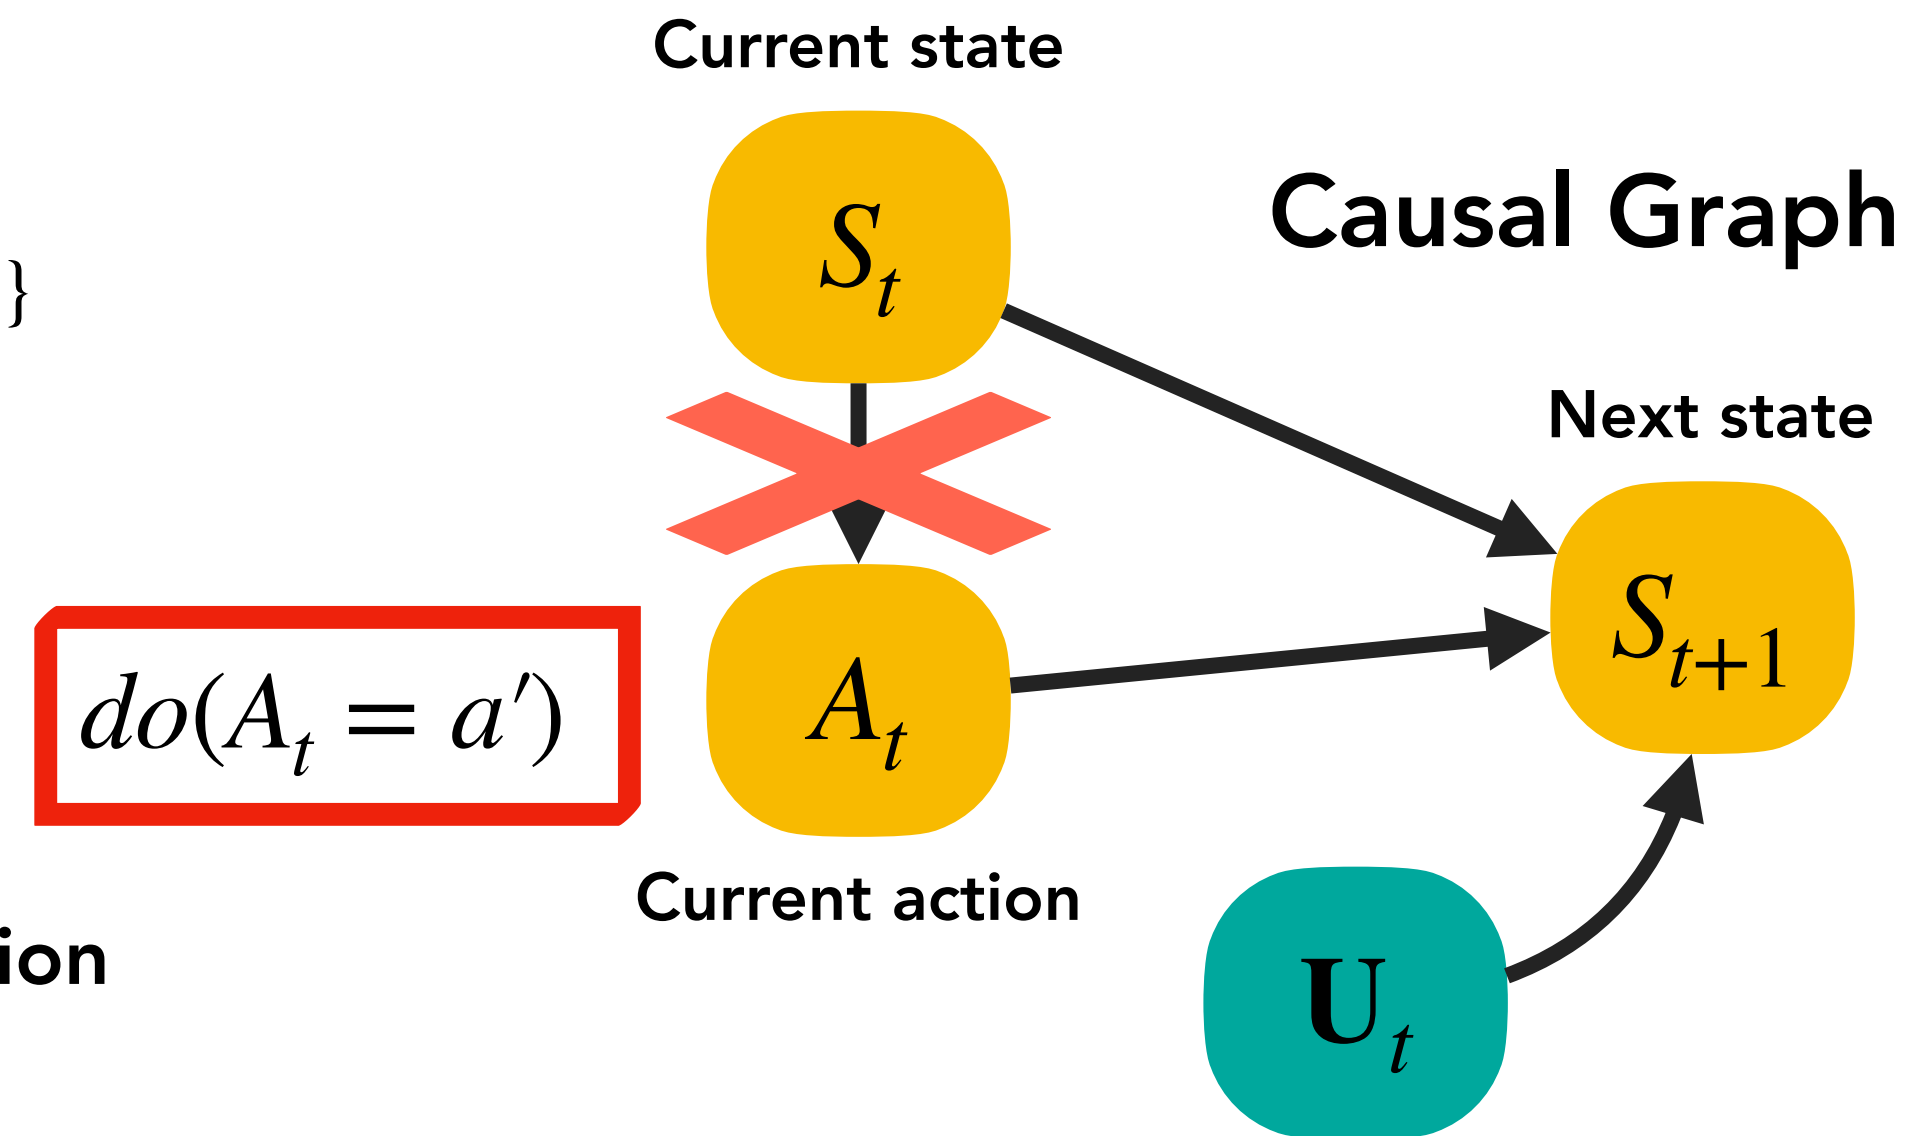
$$S_{t+1} := g_S(S_t, A_t, \mathbf{U}_t)$$

$$A_t := \cancel{g_A(S_t, \mathbf{V}_t)} \quad A_t := a'$$

$$\mathbf{U}_t \sim P(\mathbf{U} \mid S_t = s_t, A_t = a_t)$$

$$\mathbf{V}_t \sim P(\mathbf{V} \mid S_t = s_t)$$

Posterior distribution  
of the noises



At state  $S_t = s_t$ , the doctor took action  $A_t = a_t$ , what would have happened had the doctor taken action  $a' \neq a_t$ ?

$$S_{t+1} \sim P^{\mathcal{M}} \mid S_t=s_t, A_t=a_t; do(A_t=a') (S_{t+1})$$

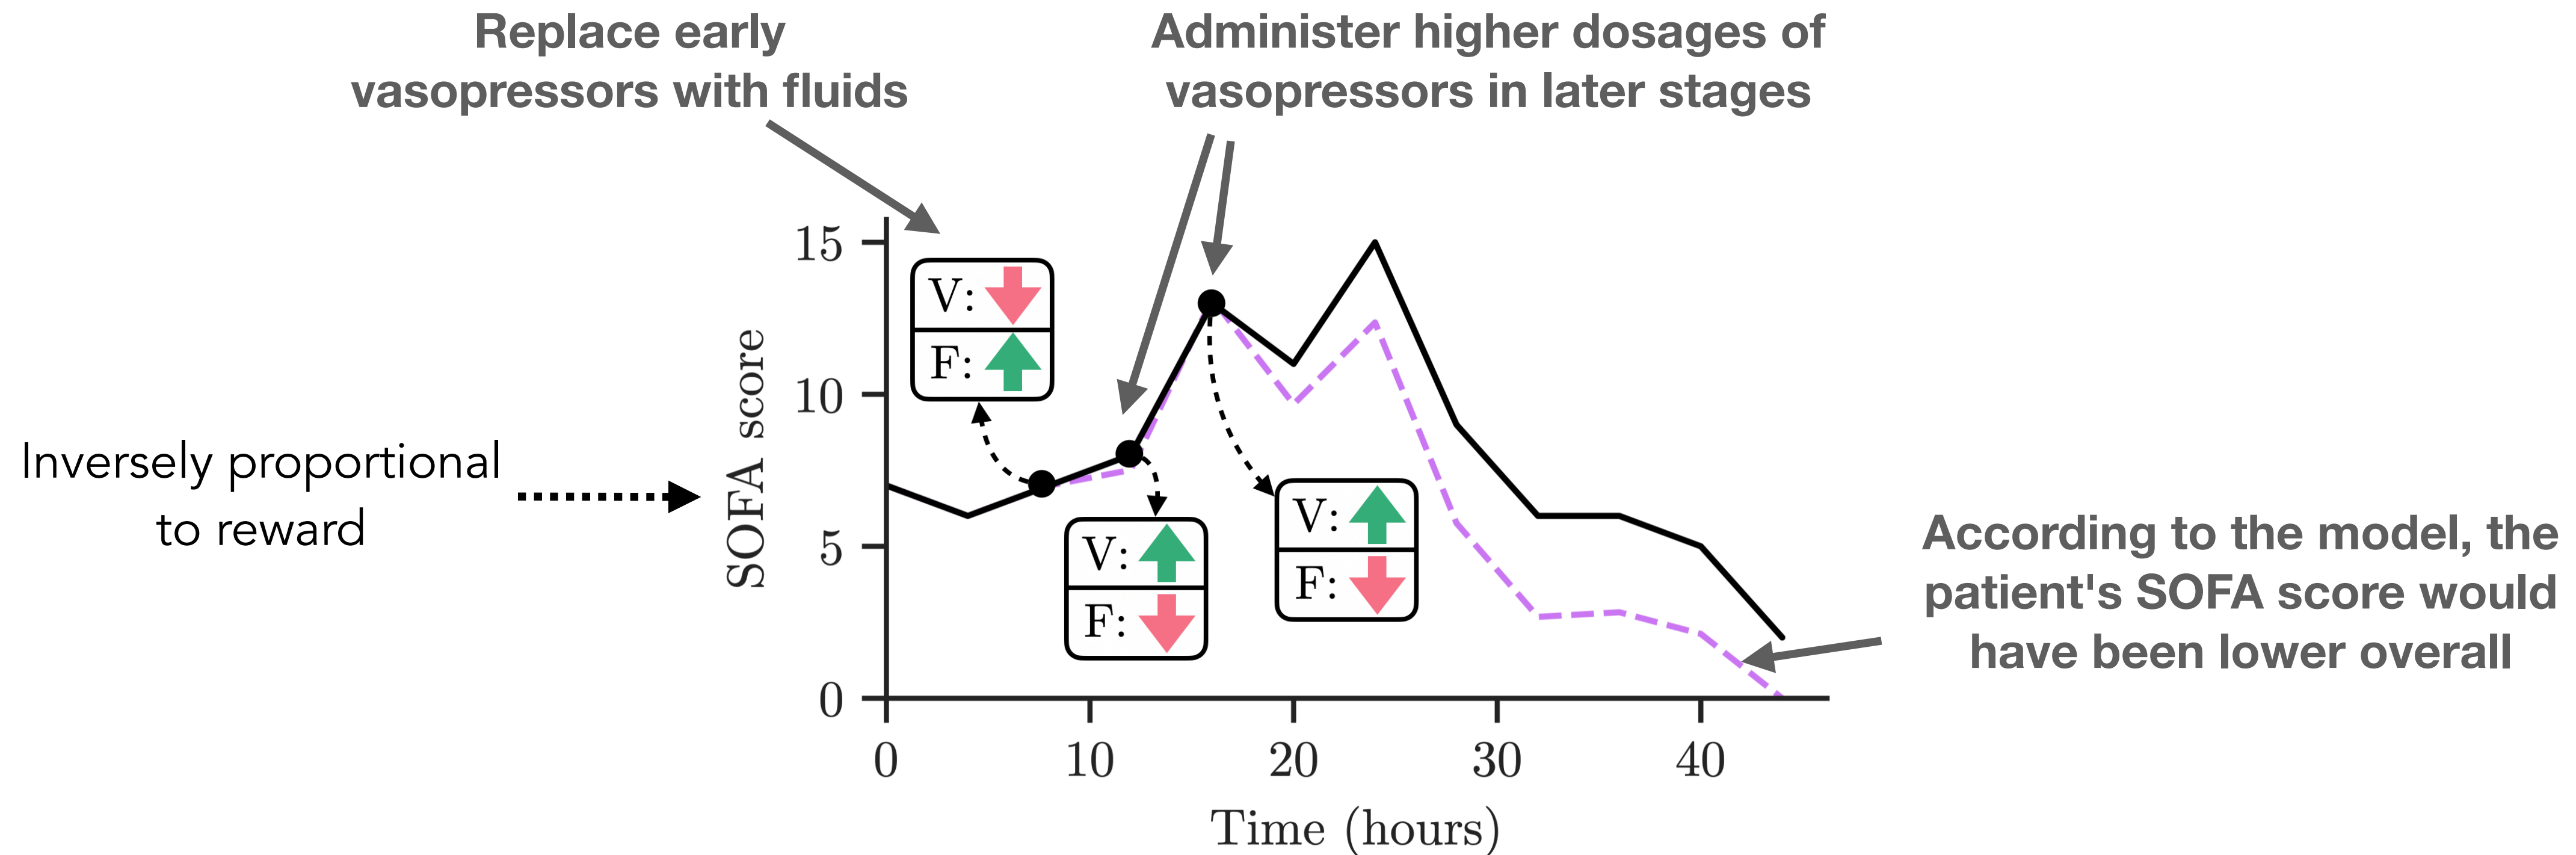
...

# Counterfactually optimal action sequences

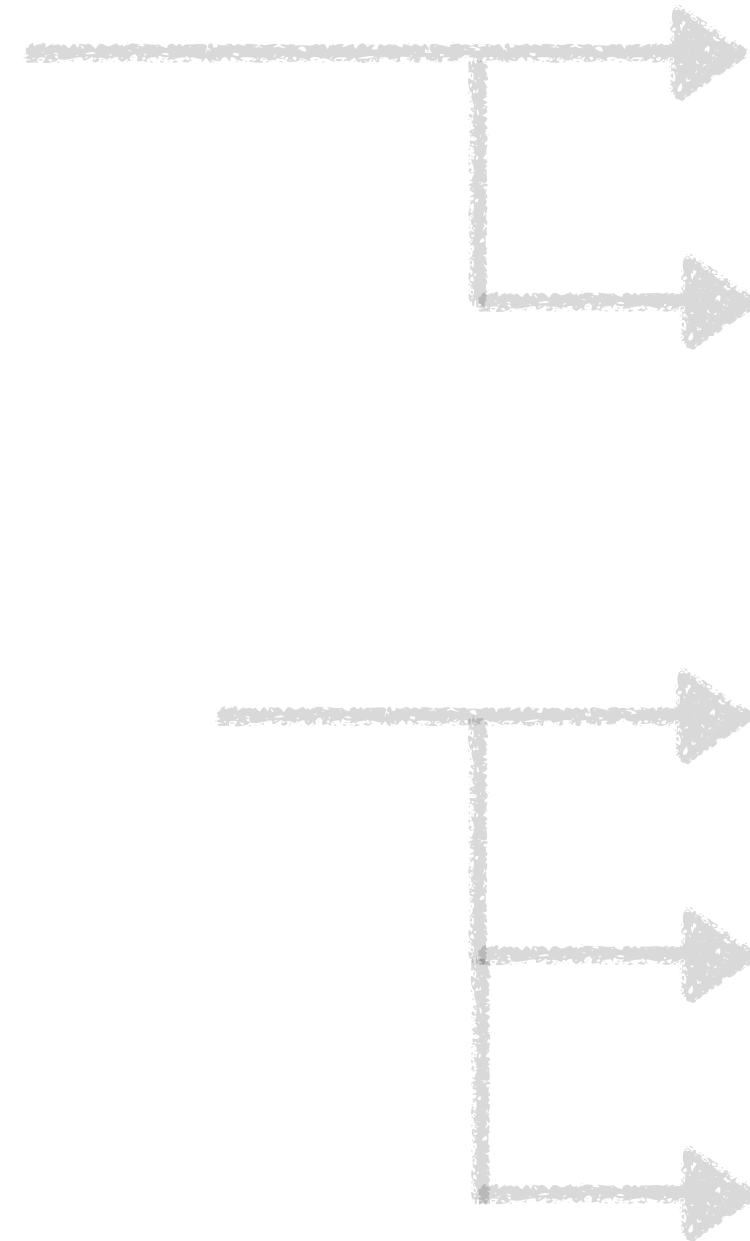
Given the counterfactual transition probabilities  $S_{t+1} \sim P^{\mathcal{M}} | S_t=s_t, A_t=a_t; do(A_t=a')$  ( $S_{t+1}$ ) and a reward function  $r(s, a)$ , one may find alternative sequence of actions  $a'_1, \dots, a'_{T-1}$  close to the observed actions  $a_1, \dots, a_{T-1}$  that maximizes the average counterfactual reward.

# Counterfactually optimal action sequences

Given the counterfactual transition probabilities  $S_{t+1} \sim P^{\mathcal{M}} | S_t=s_t, A_t=a_t; do(A_t=a') (S_{t+1})$  and a reward function  $r(s, a)$ , one may find alternative sequence of actions  $a'_1, \dots, a'_{T-1}$  close to the observed actions  $a_1, \dots, a_{T-1}$  that maximizes the average counterfactual reward.



# Use cases of counterfactuals in machine learning



*Reinforcement learning* → *Training*

# Counterfactually-guided training in reinforcement learning

In reinforcement learning, given a transition probability  $P(s' | s, a)$  and a reward function  $r(s, a)$ , the goal is to design an action policy  $a := \pi(s)$  with the highest average reward, i.e.

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}_{\tau \sim \pi, P} [R(\tau)] \quad \text{where} \quad R(\tau) = \sum_{t=1}^T R(s_t, a_t)$$

# Counterfactually-guided training in reinforcement learning

In reinforcement learning, given a transition probability  $P(s' | s, a)$  and a reward function  $r(s, a)$ , the goal is to design an action policy  $a := \pi(s)$  with the highest average reward, i.e.

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}_{\tau \sim \pi, P} [R(\tau)] \quad \text{where} \quad R(\tau) = \sum_{t=1}^T R(s_t, a_t)$$

Counterfactually-guided training refers to the evaluation of the above expectation using data gathered via an action policy  $\pi' \neq \pi$  and counterfactual reasoning

# Counterfactually-guided training in reinforcement learning

In reinforcement learning, given a transition probability  $P(s' | s, a)$  and a reward function  $r(s, a)$ , the goal is to design an action policy  $a := \pi(s)$  with the highest average reward, i.e.

$$\pi^* = \operatorname{argmax}_{\pi} \mathbb{E}_{\tau \sim \pi, P} [R(\tau)] \quad \text{where} \quad R(\tau) = \sum_{t=1}^T R(s_t, a_t)$$

Counterfactually-guided training refers to the evaluation of the above expectation using data gathered via an action policy  $\pi' \neq \pi$  and counterfactual reasoning

## Structural Causal Model $\mathcal{M}$

$$S_{t+1} := g_S(S_t, A_t, \mathbf{U}_t)$$

$$A_t := \pi'(S_t)$$

$$\mathbf{U}_t \sim P(\mathbf{U})$$

Key idea:

$$E_{S_t, a_t \sim P^{\mathcal{M}}} \left[ \underbrace{P^{\mathcal{M}} | S_t=s_t, A_t=a_t; do(A_t=\pi(S_t))}_{\text{Counterfactual probability}} \right] = \underbrace{P^{\mathcal{M}}; do(A_t=\pi(S_t))}_{\text{Interventional probability}}$$

Observational  
probability

Counterfactual  
probability

Interventional  
probability



# Use cases of counterfactuals in machine learning

