Personalized Care Through Causal Learning: From Data to Decisions

Julie Josse. Senior Researcher Inria 2020-; Prof. CMAP X 2016-2020

Lead Inria-Inserm PreMeDICaL team: personalized medicine by data integration & causal learning

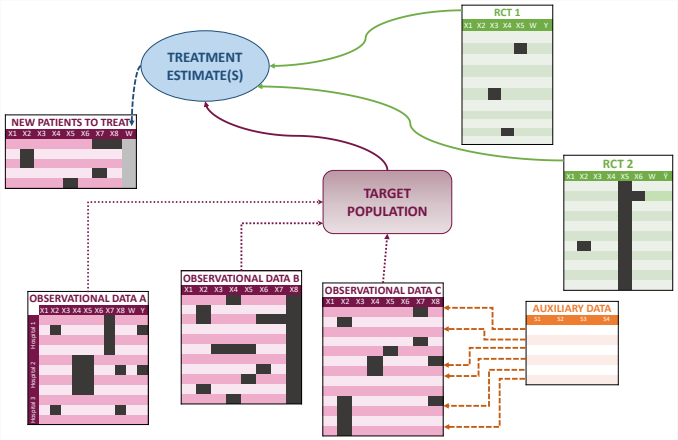

(Online) Decision support tool with quantified uncertainty

Ex: Traumatrix project[1]: Reducing under and over triage for improved resource allocation in trauma care

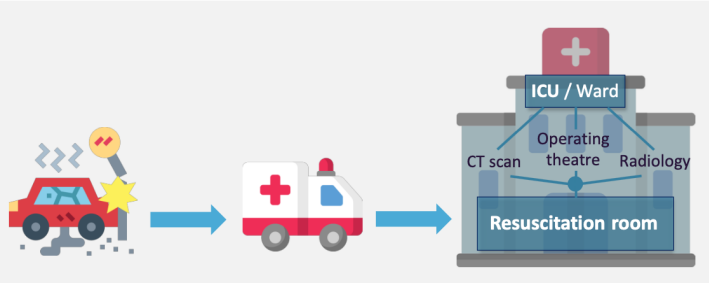

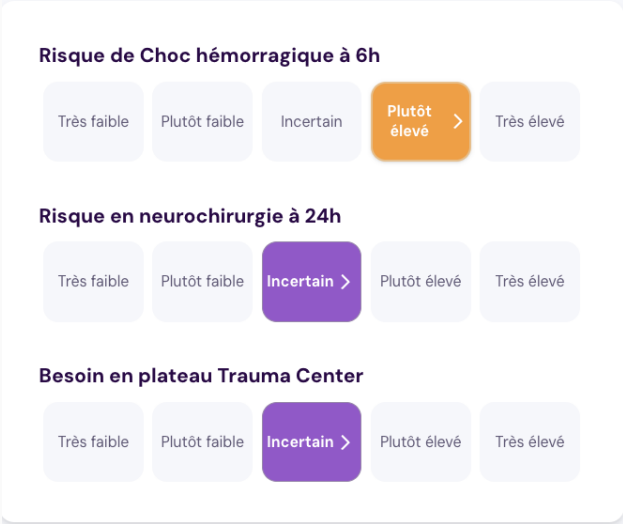

Major trauma: brain injuries or hemorrhagic shock from car accidents, falls, stab wounds, etc. $\Rightarrow$ requires specialized care/resources in "trauma centers"

Many patients are misdirected: human/ economical costs

Clinical trial launched in 2025: real-time implementation of Machine Learning models in ambulance dispatch via a mobile data collection application

[1] www.traumabase.eu - https://www.traumatrix.fr/

Personalization of treatment recommendation

Ex: Estimating treatment effect from the Traumabase data

- ▷ 40000 trauma patients
- ▷ 300 heterogeneous features from pre-hospital and in-hospital settings
- ▷ 40 trauma centers, 4000 new patients per year

Center	Accident	Age	Sex	Weight	Lactacte	Blood Press.	TXA.	Y
Beaujon	fall	54	m	85	NA	180	treated	0
Pitie	gun	26	m	NA	NA	131	untreated	1
Beaujon	moto	63	m	80	3.9	145	treated	1
Pitie	moto	30	w	NA	NA	107	untreated	0
HEGP	knife	16	m	98	2.5	118	treated	1
⋮								⋱

⇒ **Estimate causal effect** (with missing values[2]): Administration of the **treatment** *tranexamic acid (TXA)*, given within 3 hours of the accident, on the **outcome** (Y) *28 days in-hospital mortality* for trauma brain patients

[2] Mayer, I., Wager, S. & **J.J**. (2020). Doubly robust treatment effect estimation with incomplete confounders. *Annals Of Applied Statistics. (implemented in package* `grf`*).*

Causal inference: "what would happen if?"

Causal inference (simplest) question

Assume a policy/intervention/**treatment** W causes an **outcome** Y

Aim: **estimate the effect** as accurately as possible (bias & variance)

Causal inference: "what would happen if?"

Potential Outcome framework (Neyman, 1923; Rubin, 1974)

- $(\underbrace{X}_{\text{covariates}} , \overbrace{W}^{\text{treatment}} , \underbrace{Y(1), Y(0)}_{\textit{potential outcomes}}) \in \mathbb{R}^d \times \{0, 1\} \times \mathbb{R} \times \mathbb{R}$
- Individual **causal effect** of the binary treatment: $\Delta_i = Y_i(1) - Y_i(0)$

Problem: Δ_i never observed (only one outcome is observed per indiv.)

Covariates			Treatment	Outcome(s)	
X_1	X_2	X_3	W	Y(0)	Y(1)
1.1	20	F	1	**?**	200
-6	45	F	0	10	**?**
0	15	M	1	**?**	150
	...		...	...	...
-2	52	M	0	100	**?**

Causal inference: "what would happen if?"

Potential Outcome framework (Neyman, 1923; Rubin, 1974)

- $(\underbrace{X}_{\text{covariates}} , \overbrace{W}^{\text{treatment}} , \underbrace{Y(1), Y(0)}_{\textit{potential outcomes}}) \in \mathbb{R}^d \times \{0, 1\} \times \mathbb{R} \times \mathbb{R}$
- Individual **causal effect** of the binary treatment: $\Delta_i = Y_i(1) - Y_i(0)$

Problem: Δ_i never observed (only one outcome is observed per indiv.)

Covariates			Treatment	Outcome(s)	
X_1	X_2	X_3	W	Y(0)	Y(1)
1.1	20	F	1	**?**	200
-6	45	F	0	10	**?**
0	15	M	1	**?**	150
	...		...	...	...
-2	52	M	0	100	**?**

Average Treatment Effect (ATE): $\tau = \mathbb{E}[\Delta_i] = \mathbb{E}[Y_i(1) - Y_i(0)]$

ATE with **Risk Difference**: difference of the average outcome had everyone gotten treated and the average outcome had nobody gotten treatment

Data sources & evidences to estimate the treatment effect

Randomized Controlled Trial (RCT)

- **gold standard** (allocation )
- same covariate distributions in treated and control groups
  $\Rightarrow$ High **internal** validity

Observational data

Data sources & evidences to estimate the treatment effect

Randomized Controlled Trial (RCT)

- **gold standard** (allocation )
- same covariate distributions in treated and control groups
  $\Rightarrow$ High **internal** validity

- expensive, long, ethical limitations
- small sample size: restrictive inclusion criteria
  $\Rightarrow$ No personalized medicine
- **trial sample different from the population eligible for treatment**
  $\Rightarrow$ Low **external** validity

Observational data

Data sources & evidences to estimate the treatment effect

Randomized Controlled Trial (RCT)

- **gold standard** (allocation )
- same covariate distributions in treated and control groups
  $\Rightarrow$ High **internal** validity

- expensive, long, ethical limitations
- small sample size: restrictive inclusion criteria
  $\Rightarrow$ No personalized medicine
- **trial sample different from the population eligible for treatment**
  $\Rightarrow$ Low **external** validity

Observational data

- low cost
- large amounts of data (registries, biobanks, EHR, claims)
  $\Rightarrow$ patient's heterogeneity
- **representative of the target populations**
  $\Rightarrow$ High **external** validity

Data sources & evidences to estimate the treatment effect

Randomized Controlled Trial (RCT)

- **gold standard** (allocation )
- same covariate distributions in treated and control groups
  $\Rightarrow$ High **internal** validity

- expensive, long, ethical limitations
- small sample size: restrictive inclusion criteria
  $\Rightarrow$ No personalized medicine
- **trial sample different from the population eligible for treatment**
  $\Rightarrow$ Low **external** validity

Observational data

- "big data": low quality
- lack of a controlled design opens the door to **confounding bias**
  $\Rightarrow$ Low **internal** validity

- low cost
- large amounts of data (registries, biobanks, EHR, claims)
  $\Rightarrow$ patient's heterogeneity
- **representative of the target populations**
  $\Rightarrow$ High **external** validity

Leverage both RCT and observational data

RCT

\+ No confounding

– Trial sample different from the population eligible for treatment

(big) Observational data

– Confounding

\+ Representative of the target population

We can use both to [3] ...

- ▷ ...validate observational methods, correct for confounding bias
- ▷ ...improve estimation of heterogeneous treatment effects
- ▷ ...**generalize the treatment effect to a target population** (data fusion, transportability, recovery from selection bias)[4],[5]

[3] Colnet, et al. **J.J.** (2022). Causal inf. for combining RCT & obs. studies. *Statistical Science*.
[4] Elias Bareinboim & Judea Pearl. (2016). Causal inference & the data-fusion problem. *PNAS*.
[5] Dahabreh, Haneuse, Robins, Robertson, Buchanan, Stuart, Hernan. (2021). Study Designs for Extending Causal Inferences From a RCT to a Target Population *American J. of Epidemiology*.

Leverage both RCT and observational data

RCT

\+ No confounding

– Trial sample different from the population eligible for treatment

(big) Observational data

– Confounding

\+ Representative of the target population

We can use both to [3] ...

▷ ...validate observational methods, correct for confounding bias

▷ ...improve estimation of heterogeneous treatment effects

▷ ...**generalize the treatment effect to a target population** (data fusion, transportability, recovery from selection bias)[4],[5]

> The FDA has greenlighted the usage of the drug *Ibrance* to men with breast cancer, though clinical trials were performed only on women.

→ Reduce drug approval times and costs

---

[3] Colnet, et al. **J.J.** (2022). Causal inf. for combining RCT & obs. studies. *Statistical Science*.

[4] Elias Bareinboim & Judea Pearl. (2016). Causal inference & the data-fusion problem. *PNAS*.

[5] Dahabreh, Haneuse, Robins, Robertson, Buchanan, Stuart, Hernan. (2021). Study Designs for Extending Causal Inferences From a RCT to a Target Population *American J. of Epidemiology*.

Predicting treatment effects from 1 trial to another population

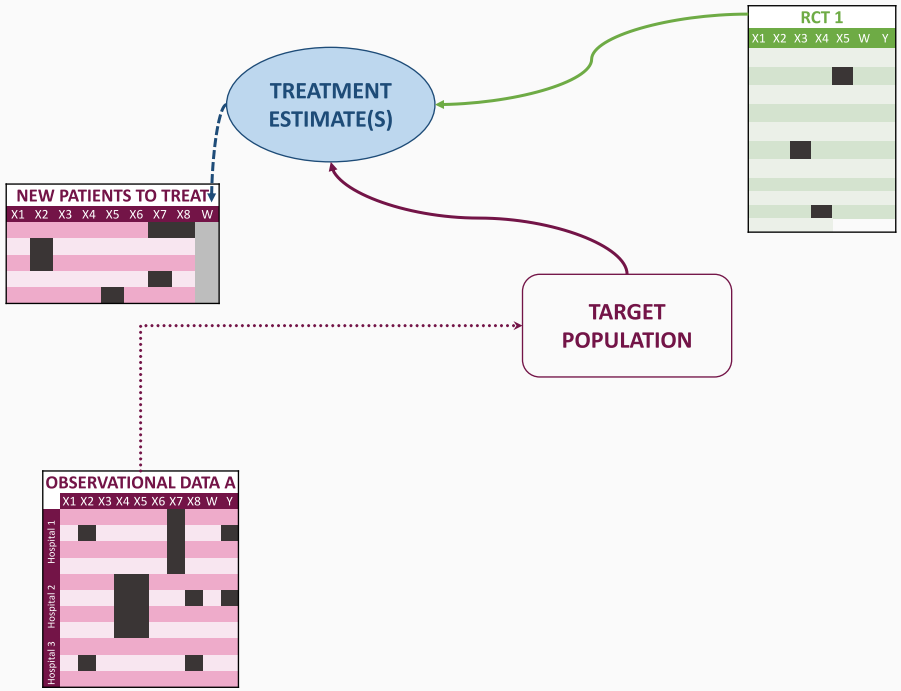




Bénédicte Colnet (Corps des Mines, French social security's direction), Imke Mayer (Owkin)
Erwan Scornet (X - Sorbonne Université), Gaël Varoquaux (Inria)

Generalization task from one RCT to a target population

Two data sources:

- A trial of size n with $p_R(x)$ the probability of observing individual with $X = x$,
- A sample of the target population of interest – for e.g. a national cohort (resp. m and $p_T(x)$).

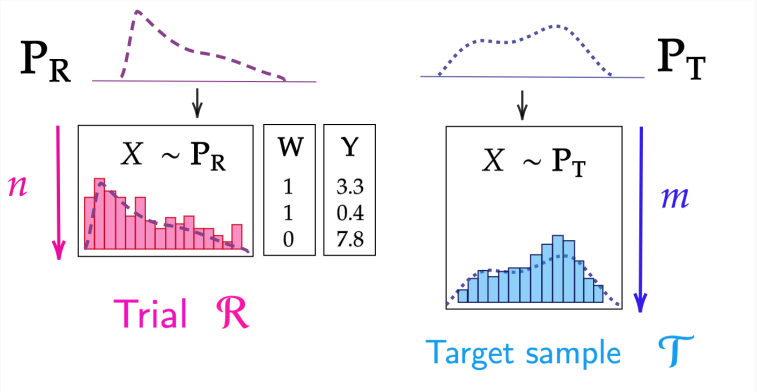

Generalization task from one RCT to a target population

Two data sources:

- A trial of size n with $p_R(x)$ the probability of observing individual with $X = x$,
- A sample of the target population of interest – for e.g. a national cohort (resp. m and $p_T(x)$).

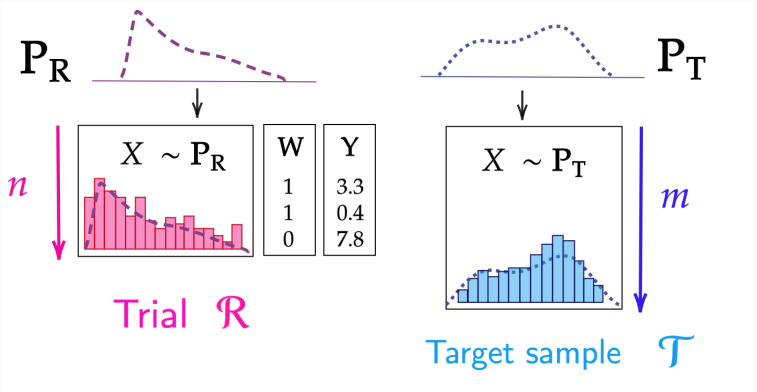

Covariates distribution not the same in the RCT & target pop:

$$p_R(x) \neq p_T(x) \Rightarrow \underbrace{\tau_R := \mathbb{E}_R[Y(1) - Y(0)]}_{\text{ATE in the RCT}} \neq \underbrace{\mathbb{E}_T[Y(1) - Y(0)] := \tau_T}_{\text{Target ATE}}$$

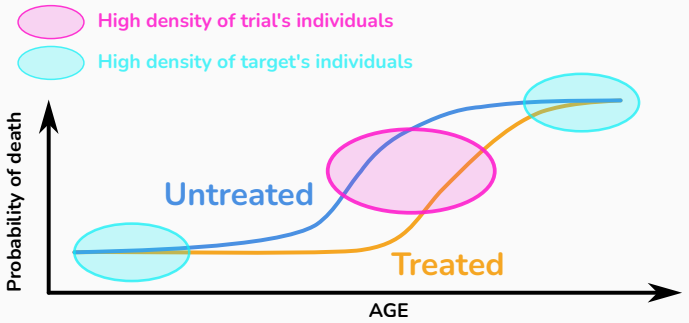

Assumptions for ATE identifiability in generalization

Overlap assumption[6]

$$\forall x \in \mathbb{X}, p_R(x) > 0 \text{ and } \text{supp}(P_T(X)) \subset \text{supp}(P_R(X))$$

The observational covariate support is included in the RCT's support. Every individual in the target population could have been selected into the trial

Transportability of the conditional average treatment effect (CATE)[7]

$$\underbrace{\mathbb{E}_R[Y(1) - Y(0) \mid X]}_{\tau_R(X)} = \underbrace{\mathbb{E}_T[Y(1) - Y(0) \mid X]}_{\tau_T(X)}$$

Need to know which variables are shifted treatment effect modifiers

The treatment effect depends on covariates in the same way in the source (RCT) and target population

[6] If this is too strong, we could generalize on a different target population: the target population for which eligibility criteria of the trial are ensured

[7] Equivalent formulation with sampling mechanism S: $(Y(1) - Y(0)) \perp\!\!\!\perp S \mid X$

Identifiability and estimation for generalization: weighting

Generalization of local effects (i.e. conditional effects/strata)

$$\begin{aligned}\tau_{\mathrm{T}} &= \mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)] = \mathbb{E}_{\mathrm{T}}[\mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)|X]] \\ &= \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{T}}(X)\right] = \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{R}}(X)\right] \text{ Transportability CATE} \\ &= \mathbb{E}_{\mathrm{R}}\left[\frac{p_{\mathrm{T}}(X)}{p_{\mathrm{R}}(X)}\tau_{\mathrm{R}}(X)\right]\end{aligned}$$

IPSW: inverse propensity sampling weighting

$$\hat{\tau}_{\pi,n,m} = \frac{1}{n}\sum_{i\in\mathcal{R}} \frac{\hat{p}_{\mathrm{T}}(X_i)}{\hat{p}_{\mathrm{R}}(X_i)} Y_i\left(\frac{W_i}{\pi} - \frac{1-W_i}{1-\pi}\right),$$

$\hat{p}_{\mathrm{R},n}(x) := \frac{1}{n}\sum_{i\in\mathcal{R}} \mathbb{1}_{X_i=x}$, π proba. of treatment assignment in trial

Re-weight, so that the trial follows the target sample's distribution

Idea of relying on an external representative sample to reweight is recent[8]

[8] Cole & Stuart. (2010). Generalizing from RCT to target pop. *American J. of Epidemiology.*

Identifiability and estimation for generalization: weighting

Generalization of local effects (i.e. conditional effects/strata)

$$\begin{aligned}\tau_{\mathrm{T}} &= \mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)] = \mathbb{E}_{\mathrm{T}}[\mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)|X]] \\ &= \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{T}}(X)\right] = \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{R}}(X)\right] \text{ Transportability CATE} \\ &= \mathbb{E}_{\mathrm{R}}\left[\frac{p_{\mathrm{T}}(X)}{p_{\mathrm{R}}(X)}\tau_{\mathrm{R}}(X)\right]\end{aligned}$$

IPSW: inverse propensity sampling weighting

$$\hat{\tau}_{\pi,n,m} = \frac{1}{n}\sum_{i\in\mathcal{R}} \frac{\hat{p}_{\mathrm{T}}(X_i)}{\hat{p}_{\mathrm{R}}(X_i)} Y_i\left(\frac{W_i}{\pi} - \frac{1-W_i}{1-\pi}\right),$$

$\hat{p}_{\mathrm{R},n}(x) := \frac{1}{n}\sum_{i\in\mathcal{R}} \mathbb{1}_{X_i=x}$, π proba. of treatment assignment in trial

Re-weight, so that the trial follows the target sample's distribution

Idea of relying on an external representative sample to reweight is recent[8]

[8] Cole & Stuart. (2010). Generalizing from RCT to target pop. *American J. of Epidemiology.*

Identifiability and estimation for generalization: weighting

Generalization of local effects (i.e. conditional effects/strata)

$$\begin{aligned}\tau_{\mathrm{T}} &= \mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)] = \mathbb{E}_{\mathrm{T}}[\mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)|X]] \\ &= \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{T}}(X)\right] = \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{R}}(X)\right] \text{ Transportability CATE} \\ &= \mathbb{E}_{\mathrm{R}}\left[\frac{p_{\mathrm{T}}(X)}{p_{\mathrm{R}}(X)}\tau_{\mathrm{R}}(X)\right]\end{aligned}$$

IPSW: inverse propensity sampling weighting

$$\hat{\tau}_{\pi,n,m} = \frac{1}{n}\sum_{i\in\mathcal{R}} \frac{\hat{p}_{\mathrm{T}}(X_i)}{\hat{p}_{\mathrm{R}}(X_i)} Y_i\left(\frac{W_i}{\pi} - \frac{1-W_i}{1-\pi}\right),$$

$\hat{p}_{\mathrm{R},n}(x) := \frac{1}{n}\sum_{i\in\mathcal{R}} \mathbb{1}_{X_i=x}$, π proba. of treatment assignment in trial

Re-weight, so that the trial follows the target sample's distribution

Idea of relying on an external representative sample to reweight is recent[8]

[8] Cole & Stuart. (2010). Generalizing from RCT to target pop. *American J. of Epidemiology.*

Identifiability and estimation for generalization: weighting

Generalization of local effects (i.e. conditional effects/strata)

$$\begin{aligned}\tau_{\mathrm{T}} &= \mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)] = \mathbb{E}_{\mathrm{T}}[\mathbb{E}_{\mathrm{T}}[Y_i(1) - Y_i(0)|X]] \\ &= \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{T}}(X)\right] = \mathbb{E}_{\mathrm{T}}\left[\tau_{\mathrm{R}}(X)\right] \text{ Transportability CATE} \\ &= \mathbb{E}_{\mathrm{R}}\left[\frac{p_{\mathrm{T}}(X)}{p_{\mathrm{R}}(X)}\tau_{\mathrm{R}}(X)\right]\end{aligned}$$

IPSW: inverse propensity sampling weighting

$$\hat{\tau}_{\pi,n,m} = \frac{1}{n}\sum_{i\in\mathcal{R}} \frac{\hat{p}_{\mathrm{T}}(X_i)}{\hat{p}_{\mathrm{R}}(X_i)} Y_i\left(\frac{W_i}{\pi} - \frac{1-W_i}{1-\pi}\right),$$

$\hat{p}_{\mathrm{R},n}(x) := \frac{1}{n}\sum_{i\in\mathcal{R}} \mathbb{1}_{X_i=x}$, π proba. of treatment assignment in trial

Re-weight, so that the trial follows the target sample's distribution

Idea of relying on an external representative sample to reweight is recent[8]

Open questions remain: Impact of the two data sources' sizes n & m?

[8] Cole & Stuart. (2010). Generalizing from RCT to target pop. *American J. of Epidemiology.*

Reweighting the RCT: finite sample & asymptotic analysis[9]

$$\hat{\tau}_{\pi,n,m} = \frac{1}{n} \sum_{i \in \mathcal{R}} \frac{\overbrace{\hat{p}_{\mathrm{T}}(X_i)}^{\text{Estimated with } \mathcal{T}}}{\underbrace{\hat{p}_{\mathrm{R}}(X_i)}_{\text{Estimated with } \mathcal{R}}} Y_i \left(\frac{W_i}{\pi} - \frac{1 - W_i}{1 - \pi} \right),$$

Asymptotic properties

Letting $\lim_{n,m \to \infty} m/n = \lambda \in [0, \infty]$

$$\lim_{n,m \to \infty} \min(n, m) \operatorname{Var}\left[\hat{\tau}_{\pi,n,m}\right] = \min(1, \lambda) \left(\frac{\operatorname{Var}\left[\tau(X)\right]}{\lambda} + V_{so} \right)$$

Variance depends on the size of the two data sets, n and m

- If target >> trial, $m/n \to \infty$, (i.e., $\lambda = \infty$): asymptotic variance = Semi-Oracle's one and depends on the ratio of probabilities.
- If trial >> target $m/n \to 0$, (i.e., $\lambda = 0$): asymptotic variance depends on var. of Conditional Average Treatment Effect $\operatorname{Var}\left[\tau(X)\right]$.

[9] Colnet, **J.J** et al. 2022. Reweighting the RCT for generalization: finite sample analysis and variable selection. *JRSSA*.

Reweighting the RCT: finite sample & asymptotic analysis[9]

$$\hat{\tau}_{\pi,n,m} = \frac{1}{n} \sum_{i \in \mathcal{R}} \frac{\overbrace{\hat{p}_{\mathrm{T}}(X_i)}^{\text{Estimated with } \mathcal{T}}}{\underbrace{\hat{p}_{\mathrm{R}}(X_i)}_{\text{Estimated with } \mathcal{R}}} Y_i \left(\frac{W_i}{\pi} - \frac{1 - W_i}{1 - \pi} \right),$$

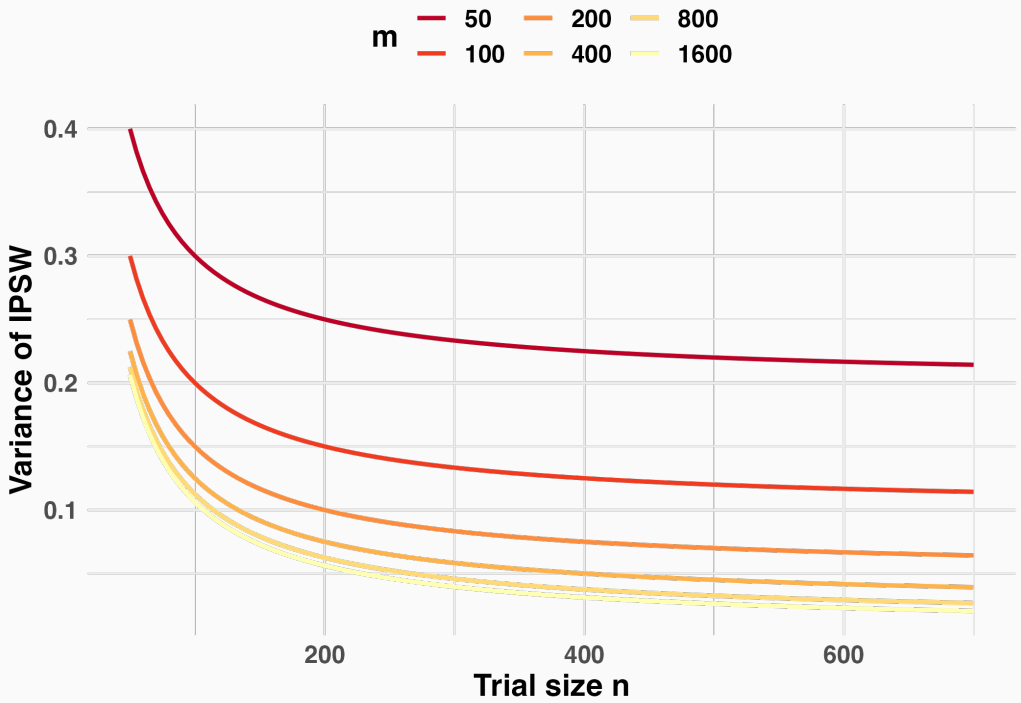

Impact for data collection

[9] Colnet, **J.J** et al. 2022. Reweighting the RCT for generalization: finite sample analysis and variable selection. *JRSSA*.

Generalization from Crash 3 trial[10] to the Traumabase

CRASH3

- Multi-centric RCT - 29 countries
- 9000 individuals - develp. countries
- Positive effect for moderately injured patients

Traumabase

- Observational sample
- 8200 patients with brain trauma
- Deleterious/No evidence for an effect of TXA

Comparison of trials, observational data, and generalization estimates

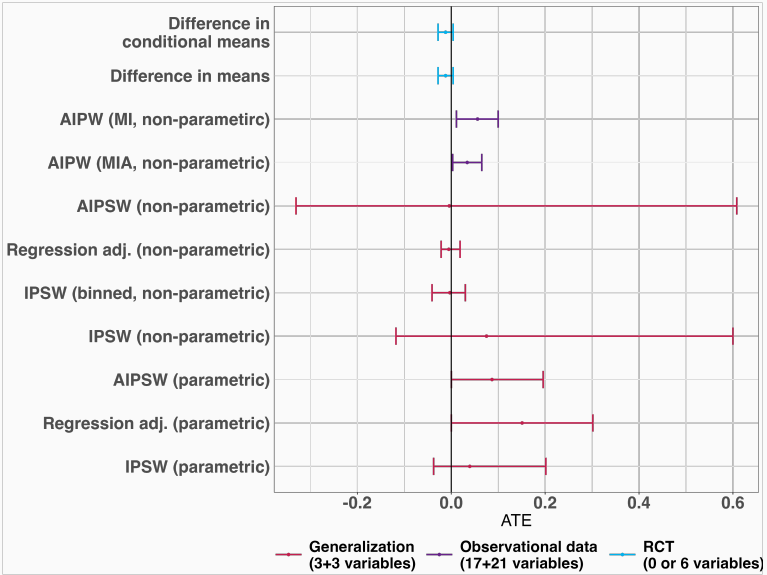

x-axis: Estimation of the Average Treatment Effect, Confidence intervals with bootstrap
y-axis: Estimation methods (estimation of nuisances: parametric: logistic regression - non parametric: forests)

[10](2019). Effects of tranexamic acid on death in patients with acute trauma. brain injury. *Lancet.*

Many medical and statistical challenges

• 1) Shifted effect modifiers not available in Traumabase[11]. Missing covariates in one/both sets: **sensitivity analysis**

			Covariates			Treat	Outcomes
	Set	S	X_1	X_2	X_3	W	Y
1	$\mathcal{R}$	1	1.1	20	NA	1	24.1
	$\mathcal{R}$	1	-6	45	NA	0	26.3
n	$\mathcal{R}$	1	0	15	NA	1	23.5
$n+1$	$\mathcal{O}$	?	-1	35	7.1		
$n+2$	$\mathcal{O}$	?	-2	52	2.4		
	$\mathcal{O}$	?		...			
$n+m$	$\mathcal{O}$	?	-2	22	3.4		

[11] Colnet, **J.J**, et al. 2022. Generalizing a causal effect: sensitivity analysis. *J. of Causal Inference*.

[12] Mayer, **J.J.** 2021. Generalizing effects with incomplete covariates *Biometrical Journal.*

[13] Colnet, **J.J** et al. 2023. Reweighting the RCT for generalization: finite sample analysis and variable selection. *JRSSC.*

[14] Colnet, **J.J** et al. 2024. Risk-Ratio, Odds-ratio, wich causal measure is easier to generalize?

Many medical and statistical challenges

• 1) Shifted effect modifiers not available in Traumabase[11]. Missing covariates in one/both sets: **sensitivity analysis**

			Covariates			Treat	Outcomes
	Set	S	X_1	X_2	X_3	W	Y
1	$\mathcal{R}$	1	1.1	20	NA	1	24.1
	$\mathcal{R}$	1	-6	45	NA	0	26.3
n	$\mathcal{R}$	1	0	15	NA	1	23.5
$n+1$	$\mathcal{O}$	?	-1	35	7.1		
$n+2$	$\mathcal{O}$	?	-2	52	2.4		
	$\mathcal{O}$	?		...			
$n+m$	$\mathcal{O}$	?	-2	22	3.4		

• 2) **Missing values:** Missing values (NA) in both RCT and Obs data[12]

[11] Colnet, **J.J**, et al. 2022. Generalizing a causal effect: sensitivity analysis. *J. of Causal Inference*.
[12] Mayer, **J.J.** 2021. Generalizing effects with incomplete covariates *Biometrical Journal.*
[13] Colnet, **J.J** et al. 2023. Reweighting the RCT for generalization: finite sample analysis and variable selection. *JRSSC.*
[14] Colnet, **J.J** et al. 2024. Risk-Ratio, Odds-ratio, wich causal measure is easier to generalize?

Many medical and statistical challenges

• 1) Shifted effect modifiers not available in Traumabase[11]. Missing covariates in one/both sets: **sensitivity analysis**

			Covariates			Treat	Outcomes
	Set	S	X_1	X_2	X_3	W	Y
1	$\mathcal{R}$	1	1.1	20	NA	1	24.1
	$\mathcal{R}$	1	-6	45	NA	0	26.3
n	$\mathcal{R}$	1	0	15	NA	1	23.5
$n+1$	$\mathcal{O}$	?	-1	35	7.1		
$n+2$	$\mathcal{O}$	?	-2	52	2.4		
	$\mathcal{O}$	?		...			
$n+m$	$\mathcal{O}$	?	-2	22	3.4		

• 2) **Missing values:** Missing values (NA) in both RCT and Obs data[12]

• 3) Which covariates should be include? Would adding prognostic variables reduce the variance as in the classical case?[13]

[11] Colnet, **J.J**, et al. 2022. Generalizing a causal effect: sensitivity analysis. *J. of Causal Inference*.
[12] Mayer, **J.J.** 2021. Generalizing effects with incomplete covariates *Biometrical Journal.*
[13] Colnet, **J.J** et al. 2023. Reweighting the RCT for generalization: finite sample analysis and variable selection. *JRSSC.*
[14] Colnet, **J.J** et al. 2024. Risk-Ratio, Odds-ratio, wich causal measure is easier to generalize?

Many medical and statistical challenges

• 1) Shifted effect modifiers not available in Traumabase[11]. Missing covariates in one/both sets: **sensitivity analysis**

			Covariates			Treat	Outcomes
	Set	S	X_1	X_2	X_3	W	Y
1	$\mathcal{R}$	1	1.1	20	NA	1	24.1
	$\mathcal{R}$	1	-6	45	NA	0	26.3
n	$\mathcal{R}$	1	0	15	NA	1	23.5
$n+1$	$\mathcal{O}$	?	-1	35	7.1		
$n+2$	$\mathcal{O}$	?	-2	52	2.4		
	$\mathcal{O}$	?		...			
$n+m$	$\mathcal{O}$	?	-2	22	3.4		

• 2) **Missing values:** Missing values (NA) in both RCT and Obs data[12]

• 3) Which covariates should be include? Would adding prognostic variables reduce the variance as in the classical case?[13]

• 4) Clinicians are more interested in the **risk ratio** than the risk difference[14]

[11] Colnet, **J.J**, et al. 2022. Generalizing a causal effect: sensitivity analysis. *J. of Causal Inference*.
[12] Mayer, **J.J.** 2021. Generalizing effects with incomplete covariates *Biometrical Journal.*
[13] Colnet, **J.J** et al. 2023. Reweighting the RCT for generalization: finite sample analysis and variable selection. *JRSSC.*
[14] Colnet, **J.J** et al. 2024. Risk-Ratio, Odds-ratio, wich causal measure is easier to generalize?

Comparing two average situations

Binary outcome: $\mathbb{P}[Y(w) = 1] = \mathbb{E}[Y(w)]$ and $\mathbb{P}[Y(w) = 0] = 1 - \mathbb{E}[Y(w)]$.

Absolute measures

$$\tau^{\text{RD}} := \mathbb{E}[Y(1)] - \mathbb{E}[Y(0)], \qquad \tau^{\text{NNT}} := (\tau^{\text{RD}})^{-1}.$$

• Number Needed to Treat (NNT): how many individuals should be treated to observe one individual answering positively to treatment.

Relative measures

$$\tau^{\text{RR}} := \frac{\mathbb{E}[Y(1)]}{\mathbb{E}[Y(0)]}, \quad \tau^{\text{SR}} := \frac{\mathbb{P}[Y(1) = 0]}{\mathbb{P}[Y(0) = 0]} = \frac{1 - \mathbb{E}[Y(1)]}{1 - \mathbb{E}[Y(0)]},$$

$$\tau^{\text{OR}} := \frac{\mathbb{P}[Y(1) = 1]}{\mathbb{P}[Y(1) = 0]} \left(\frac{\mathbb{P}[Y(0) = 1]}{\mathbb{P}[Y(0) = 0]}\right)^{-1}$$

• A null effect now corresponds to a Risk Ratio of 1
• Survival Ratio (SR) corresponds to the RR with swapped labels Y
• RR is not symmetric to the choice of outcome 0 and 1 –e.g. counting the living or the dead while Odds Ratio (OR) is

Different treatment measures give different impressions

An example: Randomized Control Trial (RCT) from Cook and Sackett (1995)

- $Y = 1$ stroke in 5 years and $Y = 0$ no stroke
- W antihyperintensive therapy
- Feature X (blood pressure), $X = 1$ low baseline risk (15/1000 versus 2/10)

$$\mathbb{P}[Y(0) = 1 \mid X = 0] \geq \mathbb{P}[Y(0) = 1 \mid X = 1]$$

	τ_{RD}	τ_{RR}	τ_{SR}	τ_{NNT}	τ_{OR}
All (P_R)	-0.0452	0.6	1.05	22	0.57
X = 1	-0.006	0.6	1.01	167	0.6
X = 0	-0.08	0.6	1.1	13	0.545

- RD: treatment reduces by 0.045 the probability to suffer from a stroke
- RR: the treated has 0.6 $\times$ the risk of having a stroke comp. with the control
- SR: increased chance of not having a stroke when treated (factor 1.05).
- NNT: one has to treat 22 people to prevent one additional stroke
- OR $\approx$ RR in a stratum where prevalence of the outcome is low

Different treatment measures give different impressions

An example: Randomized Control Trial (RCT) from Cook and Sackett (1995)

- $Y = 1$ stroke in 5 years and $Y = 0$ no stroke
- W antihyperintensive therapy
- Feature X (blood pressure), $X = 1$ low baseline risk (15/1000 versus 2/10)

$$\mathbb{P}[Y(0) = 1 \mid X = 0] \geq \mathbb{P}[Y(0) = 1 \mid X = 1]$$

	τ_{RD}	τ_{RR}	τ_{SR}	τ_{NNT}	τ_{OR}
All (P_R)	-0.0452	0.6	1.05	22	0.57
$X = 1$	-0.006	0.6	1.01	167	0.6
$X = 0$	-0.08	0.6	1.1	13	0.545

- RD is heterogeneous with X while RR is homogeneous with X
- Heterogeneity's property defined w.r.t. *(i)* covariates & *(ii)* a measure
- Impact of the baseline risk: with 3% baseline mortality reduced to 1% by treatment, RD shows a 0.02 drop, while RR shows controls have three times the risk: RD suggests a small effect; RR highlights a larger one

Formalization of causal measures's properties: toward guidance

The age-old question of how to report effects

Source: Wikipedia

`` We wish to decide whether we shall count the failures or the successes and whether we shall make relative or absolute comparisons"

— Mindel C. Sheps, New England Journal of Medicine, in 1958

The choice of the measure is still actively discussed

e.g. Spiegelman and VanderWeele, 2017; Baker and Jackson, 2018; Feng et al., 2019; Doi et al., 2022; Xiao et al., 2021, 2022; Huitfeldt et al., 2021; Lapointe-Shaw et al., 2022; Liu et al., 2022 …

— CONSORT guidelines recommend to report all of them

Risk ratio, odds ratio, risk difference

Which causal measure is easier to generalize?

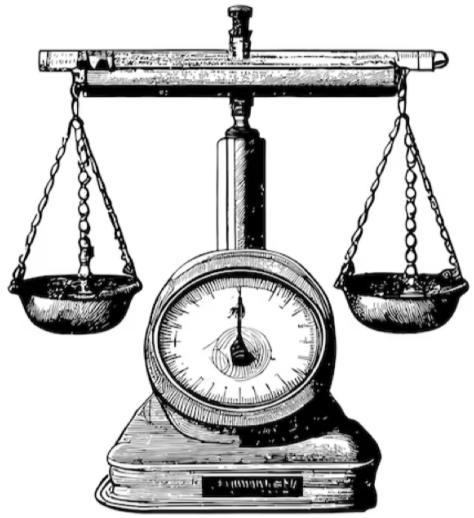

A desirable property: collapsibility

Collapsibility: Population's effect is equal to a weighted sum of local effects (conditional effects)

Direct collapsibility - weights are equal to population's proportions

$$\tau = \mathbb{E}\left[\tau(X)\right]$$

- Risk Difference is directly collapsible

	τ_{RD}	τ_{RR}	τ_{SR}	τ_{NNT}	τ_{OR}
All (P_R)	-0.0452	0.6	1.05	22	0.57
X = 1	-0.006	0.6	1.01	167	0.6
X = 0	-0.08	0.6	1.1	13	0.545

$$\tau_R^{RD} = p_R(X=1) \times \tau_R^{RD}(X=1) + p_R(X=0) \times \tau_R^{RD}(X=0)$$

$$-0.0452 = -0.47 \times 0.006 - 0.53 \times 0.08.$$

Useful for generalization! (replacing p_R by p_T)

Summary of causal measure properties

Direct collapsibility

$$\mathbb{E}[\tau(X)] = \tau$$

Collapsibility: weights depend on the baseline distribution $Y(0)$

$$\mathbb{E}[w(X, P(X, Y(0)))\,\tau(X)] = \tau \quad \text{with } w \geq 0,\ \mathbb{E}[w(X, P(X, Y(0)))] = 1$$

Logic respecting (Simpson paradox)

$$\tau \in \left[\min_x(\tau(x)), \max_x(\tau(x))\right].$$

Ex. OR: Overall population, $\tau_{\text{OR}} \approx 0.26$ $\tau_{\text{OR}|F=1} \approx 0.167$ and $\tau_{\text{OR}|F=0} \approx 0.166$

Measure	Dir. collapsible	Collapsible	Logic-respecting
Risk Difference	Yes	Yes	Yes
Number Needed to Treat	No	No	Yes
Risk Ratio	No	Yes	Yes
Survival Ratio	No	Yes	Yes
Odds Ratio	No	No	No

Back to generalizability from one RCT to a Target pop.

Generalizing	Conditional Outcome	Local effects/CATE
Assumption	$\mathbb{E}_R[Y(w) \mid X] = \mathbb{E}_T[Y(w) \mid X]$	$\tau_R(X) = \tau_T(X)$
Variables	All shifted prognostic covariates	All shifted effect modifiers
Identification	$\mathbb{E}_T[Y(w)] = \mathbb{E}_T[\mathbb{E}_R[Y(w) \mid X]]$	$\mathbb{E}_R\left[\frac{p_T(X)}{p_R(X)} w_T(Y(0), X)\, \tau_R(X)\right]$
Estimation	Regression (G-formula)	Weighting

- Generalize local effects only for collapsible measures, need info. on $Y^{(0)}$
- Generalizing conditional outcome require stronger assumptions
- Estimation challenges: using with parametric models (logistic reg.) or non-parametric ones (i.e. random forests). Ex for RR[15]: doubly robust estim. using influence function theory, convergence rate, confidence intervals, etc.

[15] Boughdiri, **J.J.**, Scornet. (2024). Estimating Risk Ratios in Causal Inference.

From one to multiple Randomized Control Trials (RCTs)

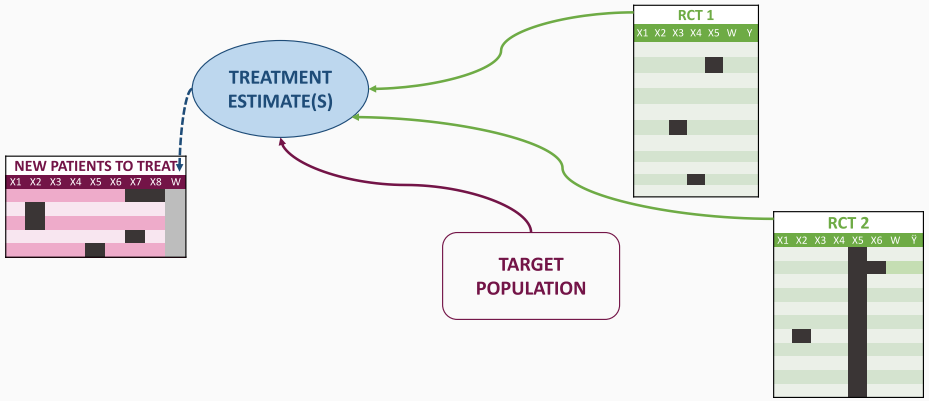

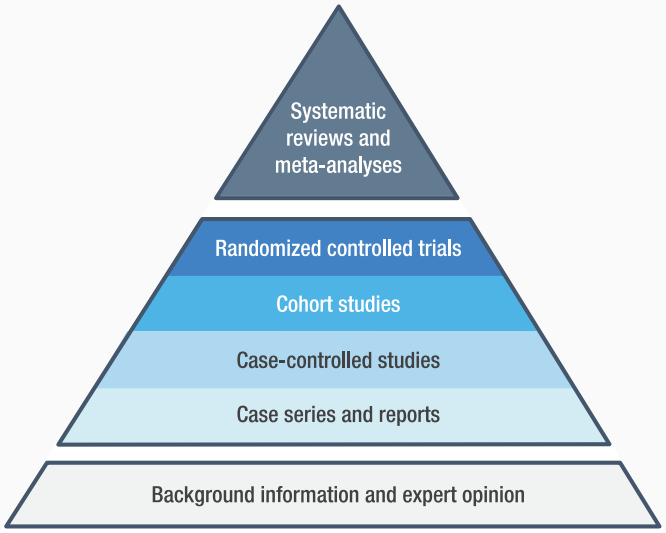

Meta-analysis (aggregating estimated effects from multiple studies) is at the top of the pyramid of evidence based medicine.

Meta-analysis still faces significant challenges:

- Be careful with aggregation of causal measures
- **Heterogeneity** across studies: sample size, population, center effects
- **Difficulty to share individual-level data**: **data silos** & **regulations**

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

initialize model

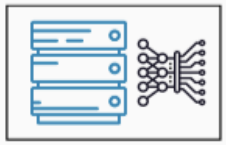

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

each party makes an update using its local dataset

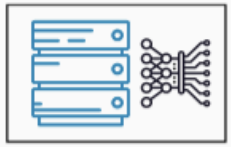

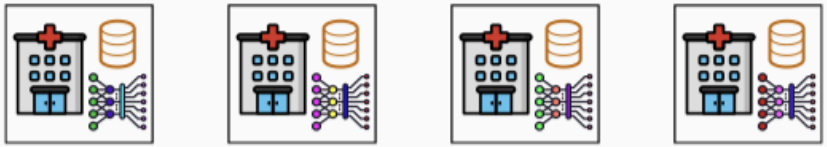

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

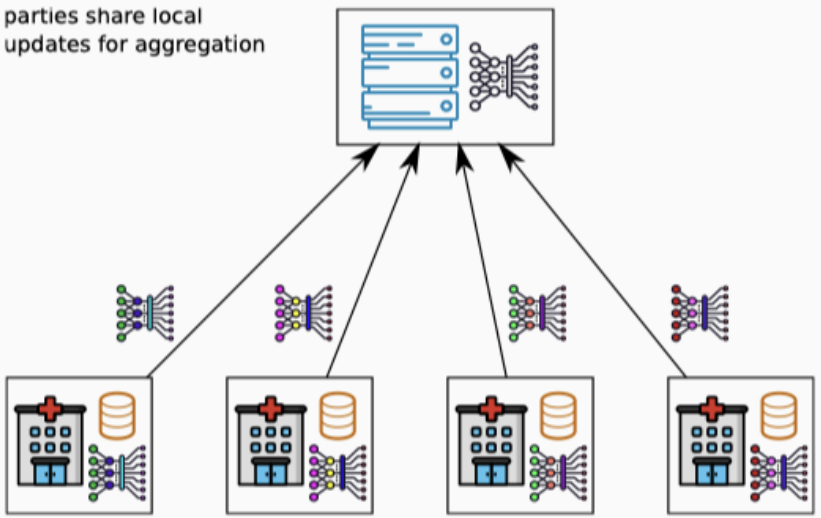

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

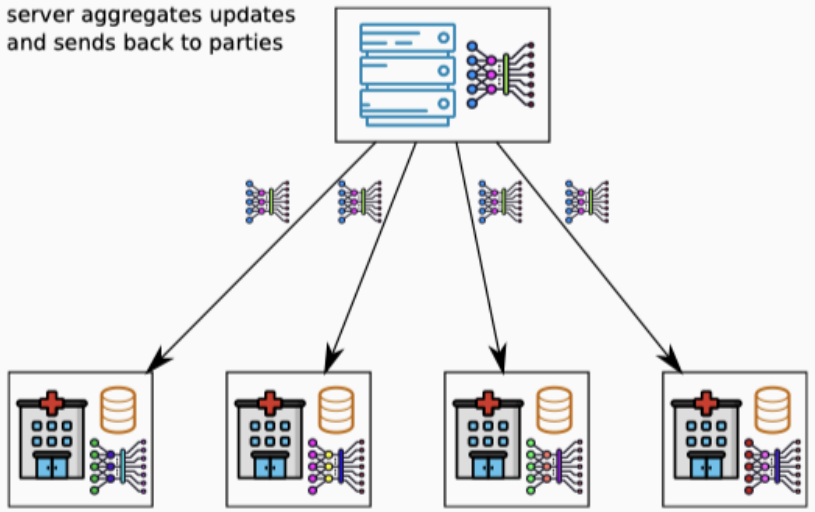

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

parties update their copy of the model and iterate

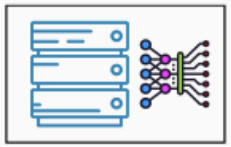

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Going beyond meta-analysis with federated causal inference[16]

A BASELINE FL ALGORITHM: FEDAVG [McMahan et al., 2017]

parties update their copy of the model and iterate

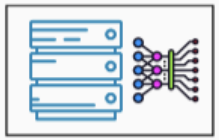

Algorithm FedAvg (server-side)

initialize θ
for each round $t = 0, 1, \ldots$ **do**
  for each party k in parallel **do**
    $\theta_k \leftarrow \text{ClientUpdate}(k, \theta)$
  $\theta \leftarrow \frac{1}{K} \sum_{k=1}^{K} \theta_k$

Algorithm ClientUpdate(k, θ)

Parameters: # steps L, step size η
for $1, \ldots, L$ **do**
  $\theta \leftarrow \theta - \eta \nabla F(\theta; \mathcal{D}_k)$
send θ to server

- Numerous extensions / improvements: fully decentralized (no server), dealing with highly heterogeneous data, privacy, fairness, compression... [Kairouz et al., 2021]

Bridging causal inference and federated learning to improve treatment effect estimation from **decentralized data sources**

[16] Rémi Khellaf, Aurelien Bellet, **J.J**. (2025). Multi-centric ATE estimation *AISTAT*.

Our setting: decentralized heterogeneous RCTs

We consider K **decentralized and potentially heterogeneous** RCTs (studies) from different sources and want to estimate the ATE given by $\tau = \mathbb{E}\left(\mathbb{E}(Y^{(1)} - Y^{(0)} \mid H)\right)$

Source	Obs.	Covariates			Treat.	Outcome
H	i	X_1	X_2	X_3	W	Y
1	1	2.3	1.5	M	1	3.2
1	2	2.2	3.1	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	1	4.5	5.0	F	1	4.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	1	3.7	2.0	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	n_K	2.5	1.7	M	0	3.2

Our setting: decentralized heterogeneous RCTs

We consider K **decentralized and potentially heterogeneous** RCTs (studies) from different sources and want to estimate the ATE given by $\tau = \mathbb{E}\left(\mathbb{E}(Y^{(1)} - Y^{(0)} \mid H)\right)$

Source	Obs.	Covariates			Treat.	Outcome
H	i	X_1	X_2	X_3	W	Y
1	1	2.3	1.5	M	1	3.2
1	2	2.2	3.1	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	1	4.5	5.0	F	1	4.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	1	3.7	2.0	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	n_K	2.5	1.7	M	0	3.2

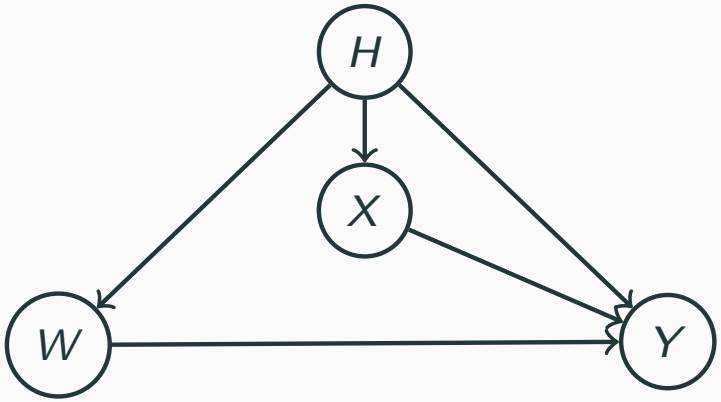

Our setting: decentralized heterogeneous RCTs

We consider K **decentralized and potentially heterogeneous** RCTs (studies) from different sources and want to estimate the ATE given by $\tau = \mathbb{E}\left(\mathbb{E}(Y^{(1)} - Y^{(0)} \mid H)\right)$

Source	Obs.	Covariates			Treat.	Outcome
H	i	X_1	X_2	X_3	W	Y
1	1	2.3	1.5	M	1	3.2
1	2	2.2	3.1	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	1	4.5	5.0	F	1	4.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	1	3.7	2.0	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	n_K	2.5	1.7	M	0	3.2

Heterogeneity in
treatment allocation

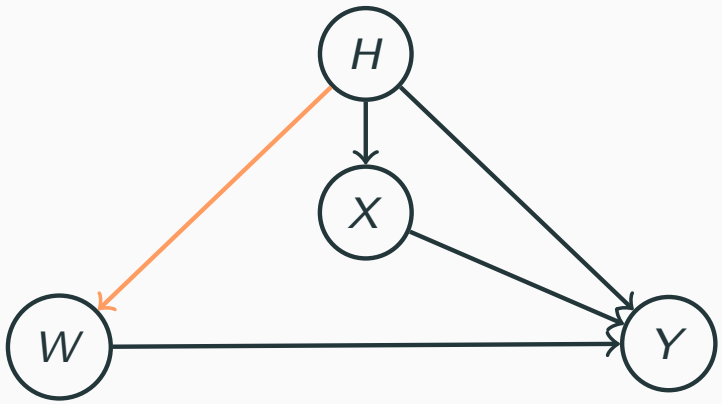

Our setting: decentralized heterogeneous RCTs

We consider K **decentralized and potentially heterogeneous** RCTs (studies) from different sources and want to estimate the ATE given by $\tau = \mathbb{E}\left(\mathbb{E}(Y^{(1)} - Y^{(0)} \mid H)\right)$

Source	Obs.	Covariates			Treat.	Outcome
H	i	X_1	X_2	X_3	W	Y
1	1	2.3	1.5	M	1	3.2
1	2	2.2	3.1	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	1	4.5	5.0	F	1	4.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	1	3.7	2.0	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	n_K	2.5	1.7	M	0	3.2

Heterogeneity in
covariates distribution

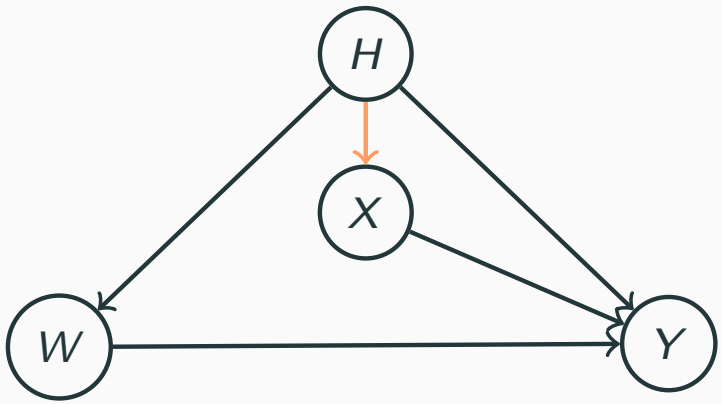

Our setting: decentralized heterogeneous RCTs

We consider K **decentralized and potentially heterogeneous** RCTs (studies) from different sources and want to estimate the ATE given by $\tau = \mathbb{E}\left(\mathbb{E}(Y^{(1)} - Y^{(0)} \mid H)\right)$

Source	Obs.	Covariates			Treat.	Outcome
H	i	X_1	X_2	X_3	W	Y
1	1	2.3	1.5	M	1	3.2
1	2	2.2	3.1	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	1	4.5	5.0	F	1	4.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	1	3.7	2.0	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	n_K	2.5	1.7	M	0	3.2

Presence of
center effects

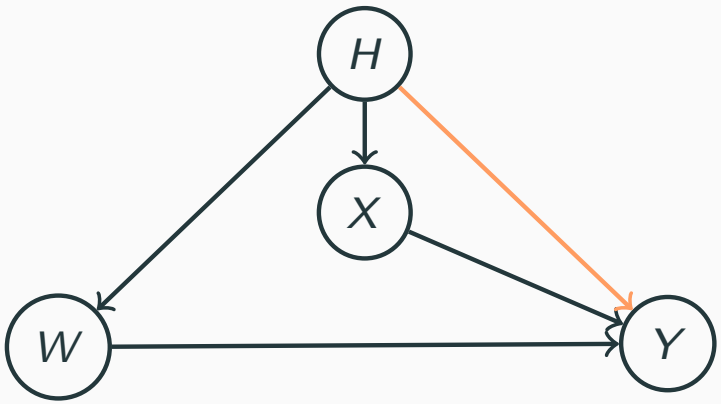

Our setting: decentralized heterogeneous RCTs

We consider K **decentralized and potentially heterogeneous** RCTs (studies) from different sources and want to estimate the ATE given by $\tau = \mathbb{E}\left(\mathbb{E}(Y^{(1)} - Y^{(0)} \mid H)\right)$

Source	Obs.	Covariates			Treat.	Outcome
H	i	X_1	X_2	X_3	W	Y
1	1	2.3	1.5	M	1	3.2
1	2	2.2	3.1	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
2	1	4.5	5.0	F	1	4.1
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	1	3.7	2.0	F	0	2.8
⋮	⋮	⋮	⋮	⋮	⋮	⋮
K	n_K	2.5	1.7	M	0	3.2

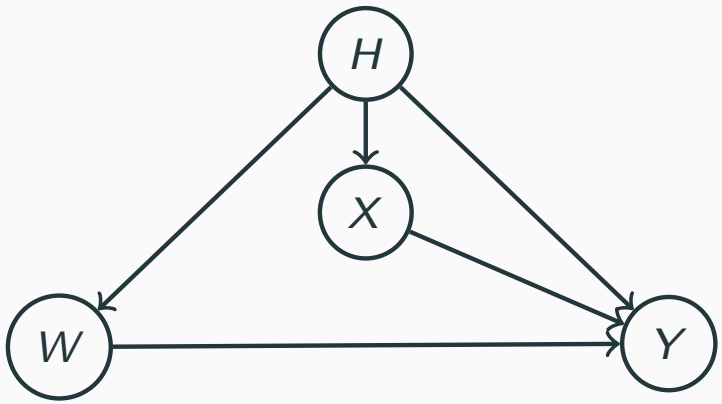

How to estimate τ without pooling together individual-level data?

Three types of federated estimators

Ex: linear outcome model for all studies $\forall k$: $Y_{k,i}^{(w)} = c^{(w)} + X_{k,i}\beta^{(w)} + \varepsilon_{k,i}^{(w)}$

Baseline: estimator $\hat{\tau}_{\text{pool}} = \frac{1}{n}\sum_{i=1}^{n} X_i'(\hat{\theta}_{\text{pool}}^{(1)} - \hat{\theta}_{\text{pool}}^{(0)})$ on pooled data

$\hat{\theta}_{\text{pool}}^{(w)} = (\hat{c}_{\text{pool}}^{(w)}, \hat{\beta}_{\text{pool}}^{(w)}) = \left(X'^{(w)^\top} X'^{(w)}\right)^{-1} X'^{(w)^\top} Y^{(w)}$ with $X'^{(w)} = [1, X^{(w)}]$

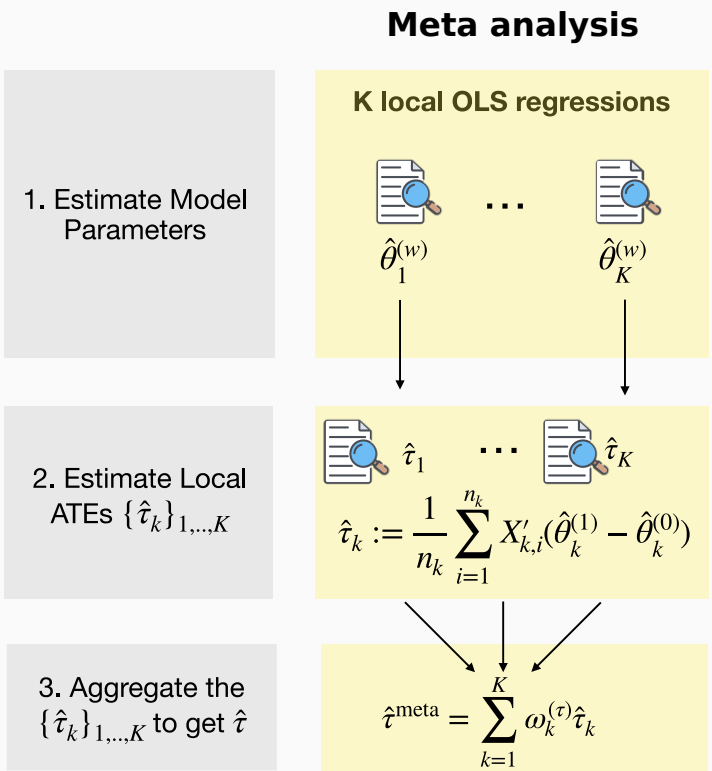

Aggregation w_k: sample size weights (SW) or inverse variance weights (IVW)

Three types of federated estimators

Ex: linear outcome model for all studies $\forall k$: $Y_{k,i}^{(w)} = c^{(w)} + X_{k,i}\beta^{(w)} + \varepsilon_{k,i}^{(w)}$

Baseline: estimator $\hat{\tau}_{\text{pool}} = \frac{1}{n}\sum_{i=1}^{n} X_i'(\hat{\theta}_{\text{pool}}^{(1)} - \hat{\theta}_{\text{pool}}^{(0)})$ on pooled data

$\hat{\theta}_{\text{pool}}^{(w)} = (\hat{c}_{\text{pool}}^{(w)}, \hat{\beta}_{\text{pool}}^{(w)}) = \left(X'^{(w)\top} X'^{(w)}\right)^{-1} X'^{(w)\top} Y^{(w)}$ with $X'^{(w)} = [1, X^{(w)}]$

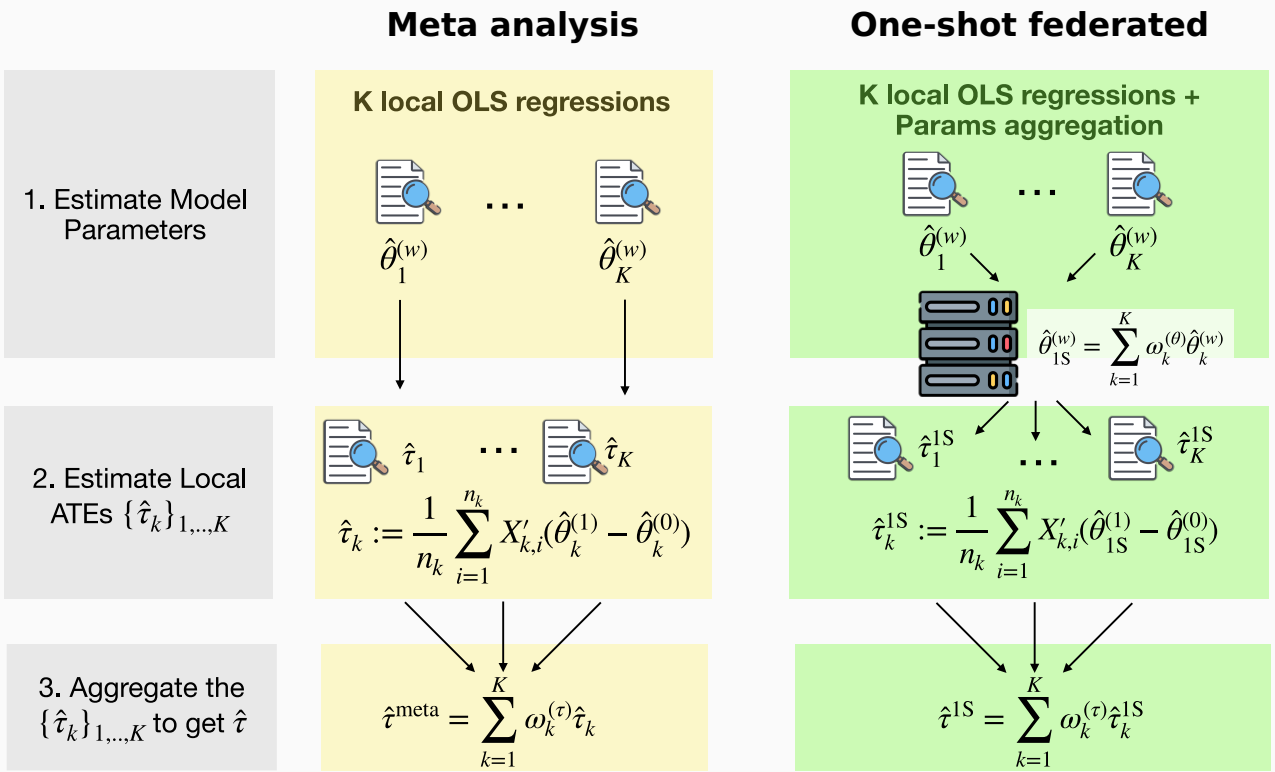

Aggregation w_k: sample size weights (SW) or inverse variance weights (IVW)

Three types of federated estimators

Ex: linear outcome model for all studies $\forall k$: $Y_{k,i}^{(w)} = c^{(w)} + X_{k,i}\beta^{(w)} + \varepsilon_{k,i}^{(w)}$

Baseline: estimator $\hat{\tau}_{\text{pool}} = \frac{1}{n}\sum_{i=1}^{n} X_i'(\hat{\theta}_{\text{pool}}^{(1)} - \hat{\theta}_{\text{pool}}^{(0)})$ on pooled data

$\hat{\theta}_{\text{pool}}^{(w)} = (\hat{c}_{\text{pool}}^{(w)}, \hat{\beta}_{\text{pool}}^{(w)}) = \left(X'^{(w)\top} X'^{(w)}\right)^{-1} X'^{(w)\top} Y^{(w)}$ with $X'^{(w)} = [1, X^{(w)}]$

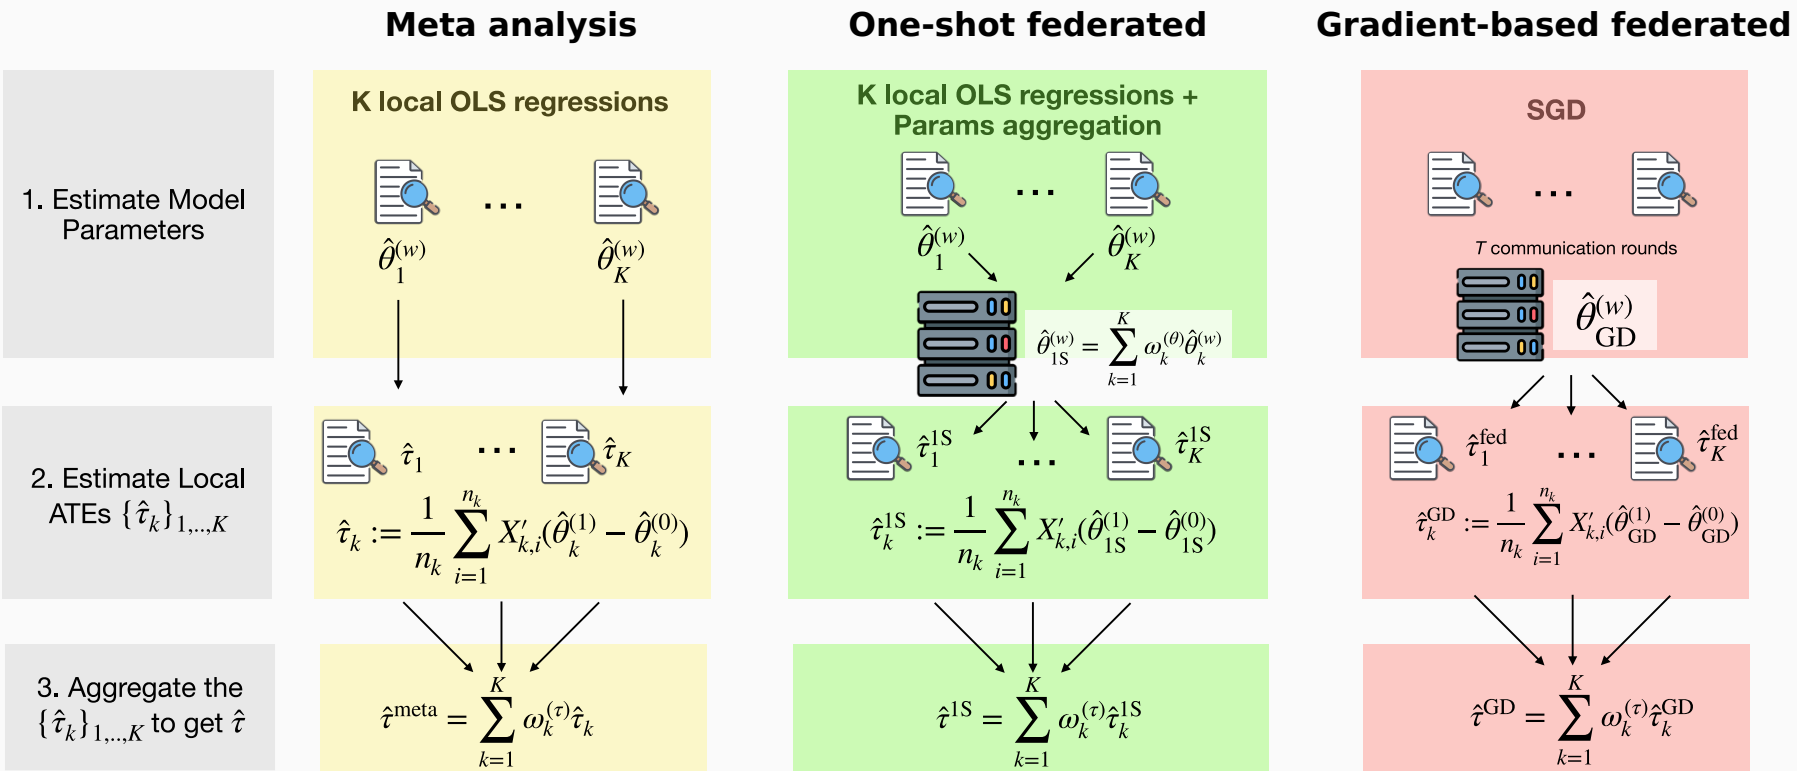

Aggregation w_k: sample size weights (SW) or inverse variance weights (IVW)

Statistical perf. & communication costs

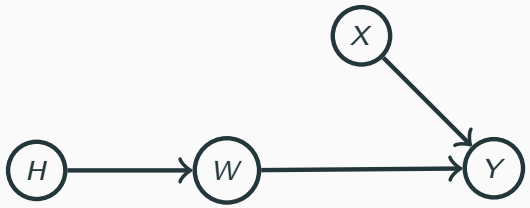

Heterogeneity: Source membership H only affects treatment allocation: $W_{k,i} \sim \mathcal{B}(p_k)$

Unbiased estimators but different asymptotic variance & communication costs:

Estimator	$\mathbb{V}^\infty$	Com. rounds	Com. cost
$\hat{\tau}_{\text{Meta-SW}}$	$\frac{\sigma^2}{n}\sum_{k=1}^{K}\frac{\rho_k}{p_k(1-p_k)}+\frac{1}{n}\|\beta^{(1)}-\beta^{(0)}\|_\Sigma^2$	1	$O(1)$
$\hat{\tau}_{\text{Meta-IVW}}$	$\Big(\sum_{k=1}^{K}(\sigma^2\frac{n\rho_k}{p_k(1-p_k)}+\frac{1}{n_k}\|\beta^{(1)}-\beta^{(0)}\|_\Sigma^2)^{-1}\Big)^{-1}$	1	$O(1)$
$\hat{\tau}_{\text{1S-SW}}$	V_{pool}	2	$O(d)$
$\hat{\tau}_{\text{1S-IVW}}$	V_{pool}	2	$O(d^2)$
$\hat{\tau}_{\text{GD}}$	V_{pool}	$T+1$	$O(Td)$
$\hat{\tau}_{\text{pool}}$	$V_{\text{pool}}=\frac{\sigma^2}{n}\frac{1}{p(1-p)}+\frac{1}{n}\|\beta^{(1)}-\beta^{(0)}\|_\Sigma^2$	—	—

with $\rho_k = \mathbb{P}(H=k) = \mathbb{E}\left[\frac{n_k}{n}\right]$ and $p = \sum_{k=1}^{K}\frac{n_k}{n}p_k$

Federated Causal Inference/Generalization

Federated RCTs: Guidelines Meta & GD predilection regimes

- Small **sample size**: Gradient Descent: other need $n_k^{(w)} \geq d$ for k, w
- **Heterogeneity**: Shift across sources ($\hat{\tau}_{\text{meta}-\text{IVW}}$ biased); different baseline outcomes ($\hat{\tau}_{\text{meta}}$ handles center effects, $\hat{\tau}_{\text{GD}}$ needs adjustment/prior knowledge on the model)

Multiple Randomized Control Trials, Multiple Observational data, Multiple Causal Measures

Real world data strenghten clinical evidences

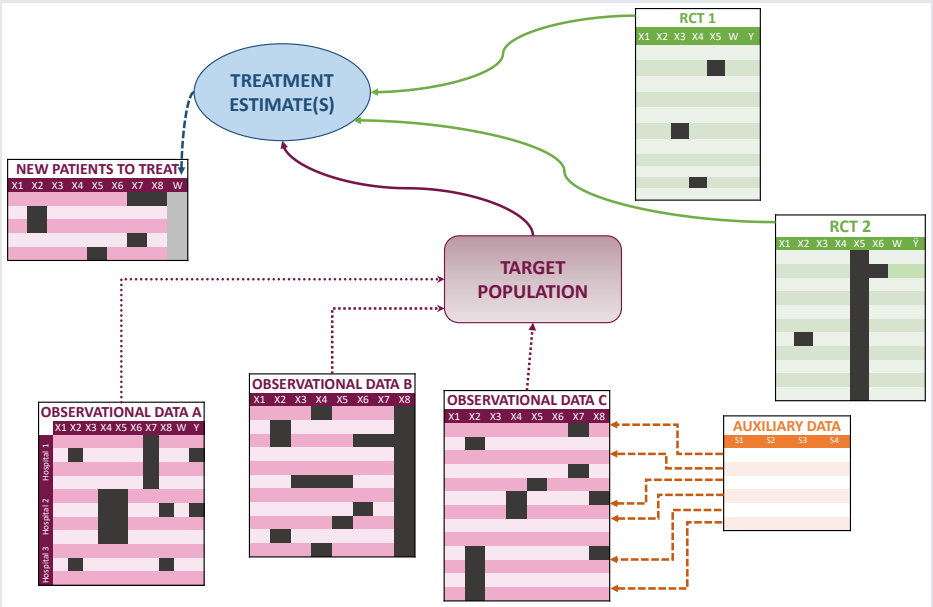

Challenges for personalized treatment effects recommendation

Formalization needed due to causal measure subtleties & aggregation

On going/future work:

- Provide robust privacy guarantees (differential privacy)
- Complex outcome[17]/treatment/features: distributions, survival, time
- Policy learning: which treatment to give to each patient at what time?
- Uncertainty quantification in treatment recommendation

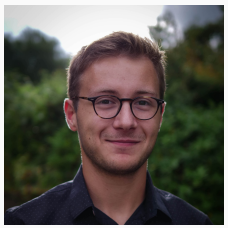





Clément Berenfeld, Ahmed Boudghiri, Mathieu Even, Agathe Chabassier, Laura Fuentes, Rémi Khellaf, Charlotte Voinot - Some are funded by PEPR Santé numérique SMATCH.


Aurelien Bellet (Inria), Erwan Scornet (Sorbone Univ.)

[17] Even, **J.J.** (2025). Rethinking the win ratio: causal framework for hierarchical outcome Analysis