FEDERATED CAUSAL INFERENCE: MULTI-SOURCE ATE ESTIMATION BEYOND META-ANALYSIS

Aurélien Bellet (Inria, Montpellier, France)

Joint work with **Rémi Khellaf** and Julie Josse

2024 IMS International Conference on Statistics and Data Science (ICSDS)
Session “Federated Causal Inference Meets Meta-Analysis”
December 17, 2024

Goal of causal inference: measure the effect of a treatment on an outcome

Randomized Controlled Trials (RCTs) are the gold standard but limited scope (stringent eligibility criteria, limited sample size...)

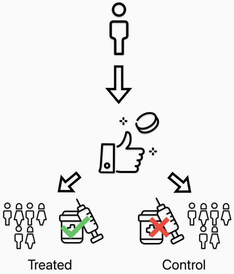

Meta-analysis (aggregating estimated effects from multiple studies) is at the top of the pyramid of evidence

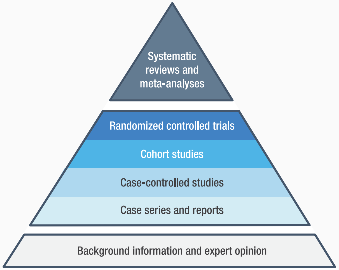

- Meta-analyses still face significant challenges:
  - Data heterogeneity across studies (sample sizes, populations, center effects...)
  - Difficulty to share individual-level data due to data silos and personal data regulations
- **Our work** bridges causal inference and federated learning [Kairouz et al., 2021] to better estimate average treatment effects from decentralized data sources
  1. We consider several estimators with varying communication costs
  2. We study their statistical performance under various types of data heterogeneity
  3. We validate on numerical simulations and provide guidelines for practitioners

RELATED WORK IN FEDERATED CAUSAL INFERENCE

- [Xiong et al., 2023] propose to use one-shot aggregation to federate the outcome or propensity score model, but it is unclear when these estimators should be preferred to other methods (e.g., traditional meta-analysis)
- [Vo et al., 2022b] employ a Bayesian framework using Gaussian processes but is restricted to uniform data distributions across sources
- [Vo et al., 2022a, Han et al., 2021, Han et al., 2023, Makhija et al., 2024, Guo et al., 2024] focus on transferring causal estimates from one source to another, while our work aims to estimate causal effects across the joint population

Problem Setting

- Estimate effect of treatment W on outcome Y given covariates X, with $W_i \sim \mathcal{B}(p)$
- Average Treatment Effect (ATE) measured as a risk difference $\tau = \mathbb{E}[Y_i(1) - Y_i(0)]$

Obs.	Covariates			Treatment	Outcome	Potential Outcomes	
i	X_1	X_2	X_3	W	Y	$Y^{(1)}$	$Y^{(0)}$
1	2.3	1.5	M	1	3.2	3.2	??
2	2.2	3.1	F	0	2.8	??	2.8
3	3.5	2.0	F	1	2.1	2.1	??
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
$n-1$	3.7	2.0	F	0	2.8	??	2.8
n	2.5	1.7	M	1	3.2	3.2	??

- 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			Treatment	Outcomes
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

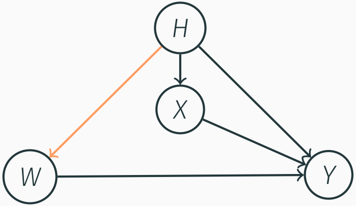

How to estimate τ without pooling together individual-level data?

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			Treatment	Outcomes
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**

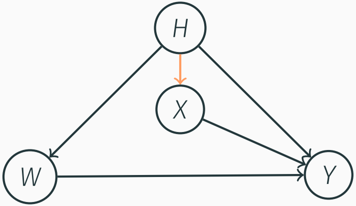

How to estimate τ without pooling together individual-level data?

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			Treatment	Outcomes
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**

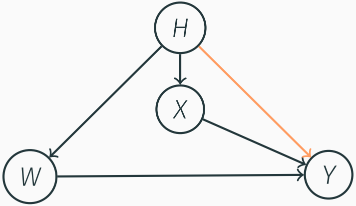

How to estimate τ without pooling together individual-level data?

- For now, same linear outcome model for all studies:

$$\forall k : \quad Y_{k,i}^{(w)} = c^{(w)} + X_{k,i}\beta^{(w)} + \varepsilon_{k,i}^{(w)}, \quad \text{with } \mathbb{E}\left[X_k^\top \varepsilon_{k,i}^{(w)}\right] = 0, \mathbb{V}\left(\varepsilon_{k,i}^{(w)} \mid X_k\right) = \sigma^2$$

- Standard assumptions (consistency, positivity, unconfoundedness)
- Ideal baseline: estimator $\hat{\tau}_{\mathrm{pool}} = \frac{1}{n}\sum_{i=1}^{n} X_i'(\hat{\theta}_{\mathrm{pool}}^{(1)} - \hat{\theta}_{\mathrm{pool}}^{(0)})$ on pooled data, where

  $\hat{\theta}_{\mathrm{pool}}^{(w)} = (\hat{c}_{\mathrm{pool}}^{(w)}, \hat{\beta}_{\mathrm{pool}}^{(w)}) = \left({X'^{(w)}}^\top X'^{(w)}\right)^{-1} {X'^{(w)}}^\top Y^{(w)}$ is the OLS estimator and $X'^{(w)} = [1, X^{(w)}]$
- $\hat{\tau}_{\mathrm{pool}}$ always has lower variance than the simple difference-in-means estimator [Benkeser et al., 2021, Lei and Ding, 2021]

Federated Estimators

THREE TYPES OF FEDERATED ESTIMATORS

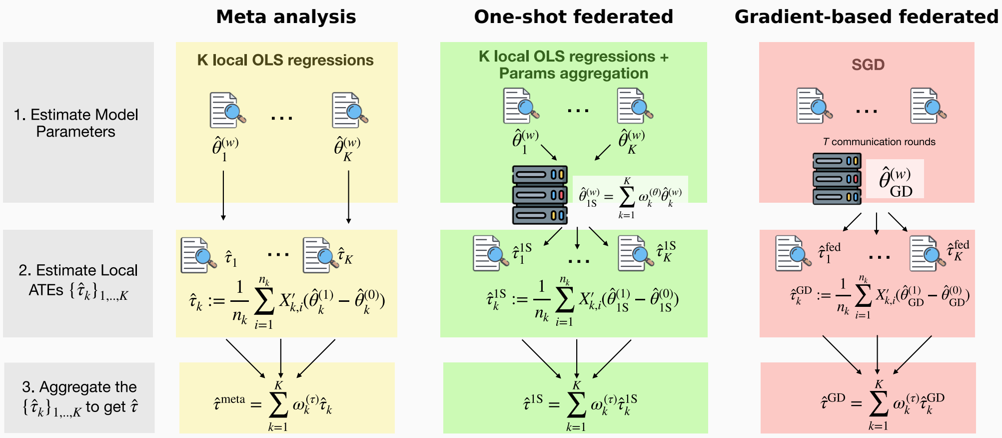

- Meta and one-shot require local sample size $n_k^{(w)} \geq d$ for k, w
- Aggregation: sample size weights (SW) or inverse variance weights (IVW)

Comparison of the Estimators

Homogeneous setting

HOMOGENEOUS SETTING

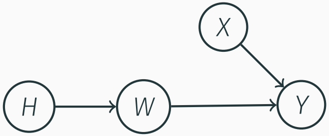

- The source membership variable H only affects the treatment allocation scheme
- Let $W_{k,i} \sim \mathcal{B}(p_k)$

Estimators are unbiased but differ by their asymptotic variance and communication costs:

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

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

NUMERICAL ILLUSTRATION ($K = 5$ AND $d = 10$)

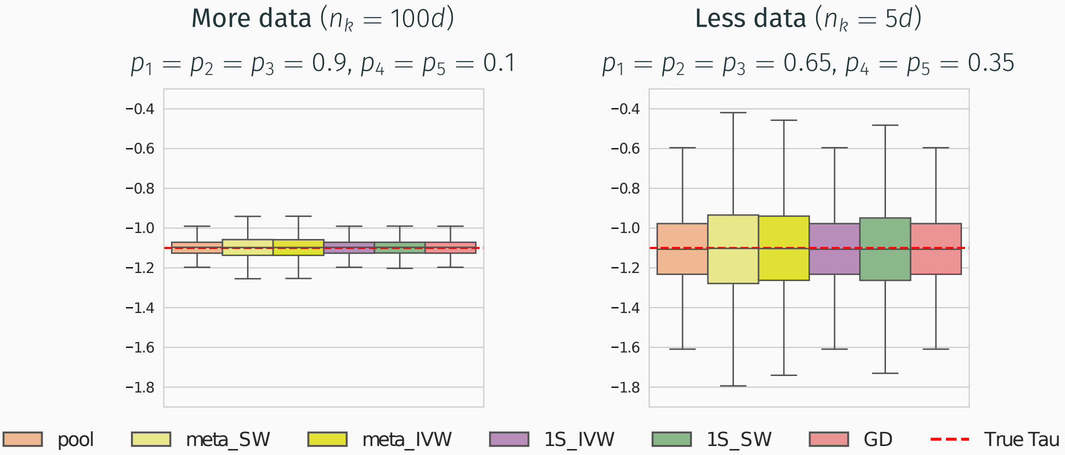

Comparison of the Estimators

Heterogeneous Distributions

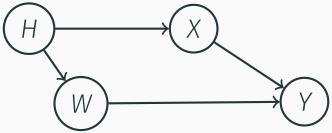

- Distributional shift across sources: $H \not\perp\!\!\!\perp X \implies \tau_k \neq \tau_{k'}$
- Global ATE is given by $\tau = \sum_{k=1}^{K} \rho_k \tau_k$ with $\rho_k = \mathbb{P}(H = k) = \mathbb{E}\left[\frac{n_k}{n}\right]$

- $\hat{\tau}_{\mathrm{meta-IVW}}$ is biased because inverse variance weights give biased estimates of the ρ_k
- $\mathbb{V}^{\infty}(\hat{\tau}_{\mathrm{pool}}) = \mathbb{V}^{\infty}(\hat{\tau}_{\mathrm{GD}}) = \mathbb{V}^{\infty}(\hat{\tau}_{\mathrm{1S-IVW}}) \leq \mathbb{V}^{\infty}(\hat{\tau}_{\mathrm{meta-SW}})$
- $\hat{\tau}_{\mathrm{1S-SW}}$ is robust to heterogeneous covariances $\{\Sigma_k\}_k$ but has larger variance for different means $\{\mu_k\}_k$

$$X_k \sim \mathcal{N}(\mu_k, \Sigma_k)$$

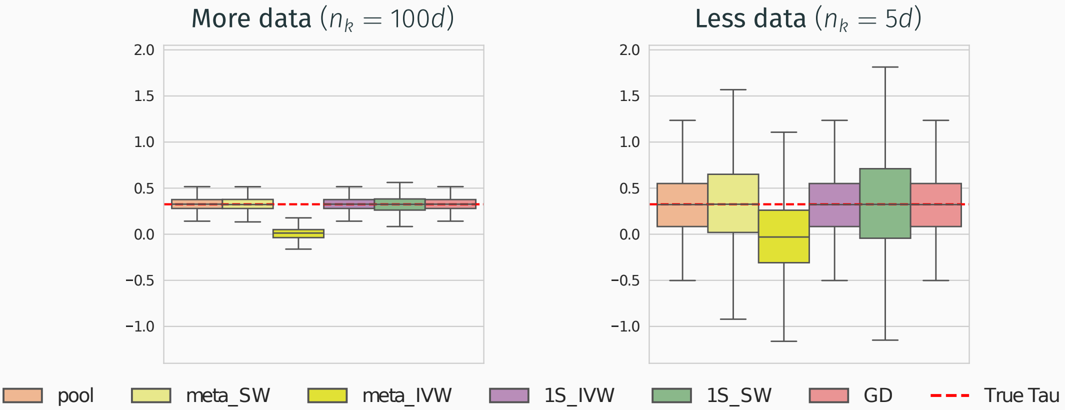

Comparison of the Estimators

Presence of Center Effects

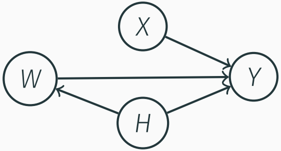

- Studies may have different baselines in individual outcomes due to varying practices or organizational contexts (e.g. hospital specialized in oncology)
- We model this by a fixed effect of the source H onto the outcome Y:

$$Y_{k,i}^{(w)} = c^{(w)} + h_k + X_{k,i}\beta^{(w)} + \varepsilon_i(w)$$

(Note: the CATEs $\mathbb{E}[Y(1) - Y(0)|X, H]$ remain the same across sources)

- Caution: H is now a confounder!

- $\hat{\tau}_{\mathrm{meta-SW}}$ and $\hat{\tau}_{\mathrm{meta-IVW}}$ naturally account for the center effects
- Other federated estimators are biased and need to be adjusted
    - One-shot estimators: share and aggregate only the covariates coefficients $\hat{\beta}_k$, while keeping intercepts local
    - GD estimators: add H as an additional covariate

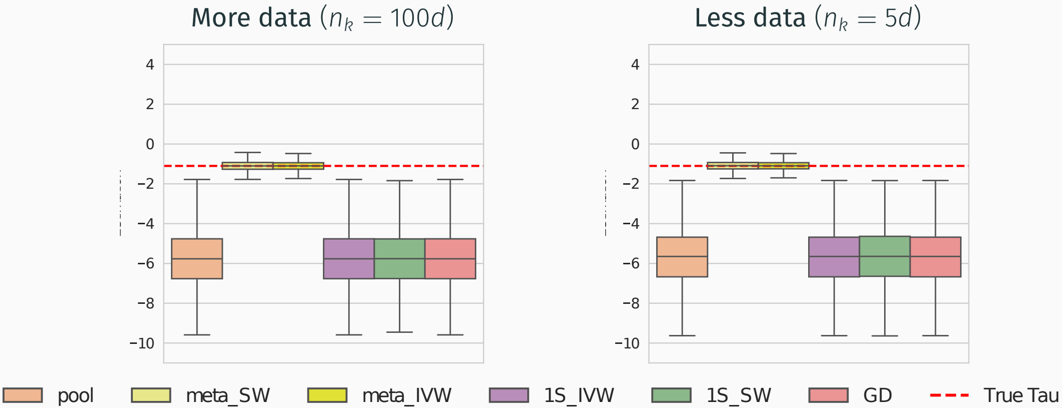
More data ($n_k = 100d$)
Less data ($n_k = 5d$)
4
2
0
−2
−4
−6
−8
−10
pool
meta_SW
meta_IVW
1S_IVW
1S_SW
GD
True Tau

NUMERICAL ILLUSTRATION

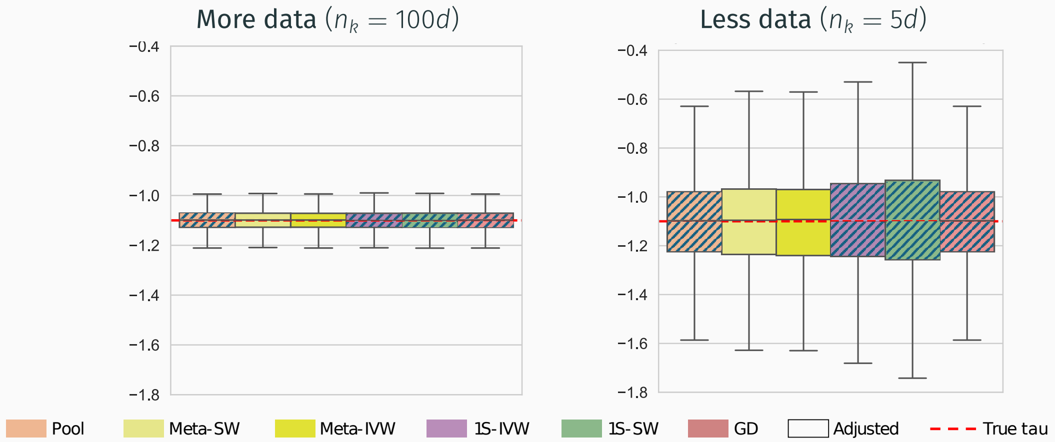

Conclusion & Perspectives

SUMMARY: DECISION DIAGRAM FOR PRACTITIONERS

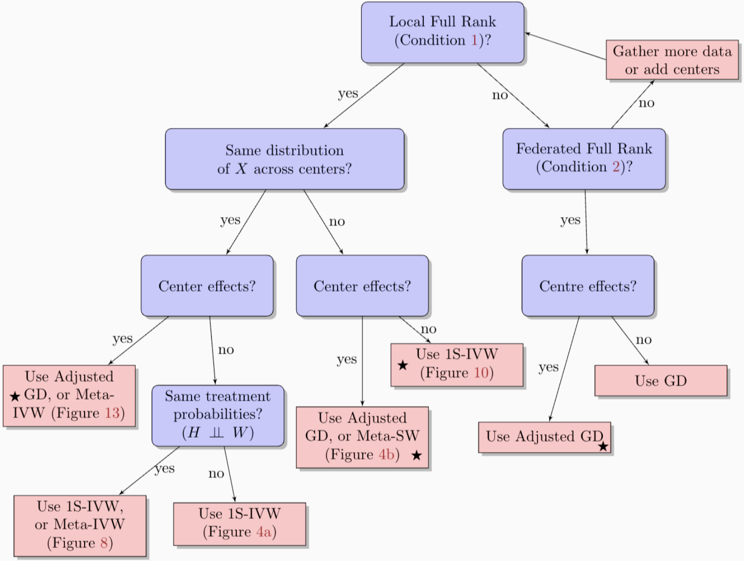

Figure 6: Decision Diagram for Practitionners. The sign ★ denotes scenarios where the DM estimator is biased.

SOME PERSPECTIVES

- Extend to observational studies (e.g., federated IPW and AIPW) and nonlinear models
- Handle covariate mismatch across sources
- Consider non-collapsible causal measures (e.g., odds ratio)
- Provide robust privacy guarantees (differential privacy)

THANK YOU FOR YOUR ATTENTION!

QUESTIONS?

[Benkeser et al., 2021] Benkeser, D., Díaz, I., Luedtke, A., Segal, J., Scharfstein, D., and Rosenblum, M. (2021).
Improving precision and power in randomized trials for covid-19 treatments using covariate adjustment, for binary, ordinal, and time-to-event outcomes.
Biometrics, 77(4):1467–1481.

[Guo et al., 2024] Guo, T., Karimireddy, S. P., and Jordan, M. I. (2024).
Collaborative heterogeneous causal inference beyond meta-analysis.
arXiv preprint arXiv:2404.15746.

[Han et al., 2021] Han, L., Hou, J., Cho, K., Duan, R., and Cai, T. (2021).
Federated adaptive causal estimation (face) of target treatment effects.
arXiv preprint arXiv:2112.09313.

[Han et al., 2023] Han, L., Shen, Z., and Zubizarreta, J. R. (2023).
Multiply robust federated estimation of targeted average treatment effects.
In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S., editors, *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.

[Kairouz et al., 2021] Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., Bonawitz, K., Charles, Z., Cormode, G., Cummings, R., et al. (2021).
Advances and open problems in federated learning.
Foundations and trends® in machine learning, 14(1–2):1–210.

[Lei and Ding, 2021] Lei, L. and Ding, P. (2021).
Regression adjustment in completely randomized experiments with a diverging number of covariates.
Biometrika, 108(4):815–828.

[Makhija et al., 2024] Makhija, D., Ghosh, J., and Kim, Y. (2024).
Federated learning for estimating heterogeneous treatment effects.
CoRR, abs/2402.17705.

[Vo et al., 2022a] Vo, T. V., Bhattacharyya, A., Lee, Y., and Leong, T.-Y. (2022a).
An adaptive kernel approach to federated learning of heterogeneous causal effects.
Advances in Neural Information Processing Systems, 35:24459–24473.

[Vo et al., 2022b] Vo, T. V., Lee, Y., Hoang, T. N., and Leong, T.-Y. (2022b).
Bayesian federated estimation of causal effects from observational data.
In *UAI*.

[Xiong et al., 2023] Xiong, R., Koenecke, A., Powell, M., Shen, Z., Vogelstein, J. T., and Athey, S. (2023).
Federated causal inference in heterogeneous observational data.
Statistics in Medicine, 42(24):4418–4439.

FULL HETEROGENEITY – GRAPHICAL MODEL

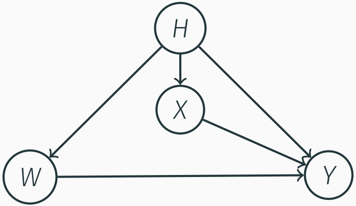

Different $h_k, p_k, \mu_k, \Sigma_k$