

Project Sketch: Heterogeneous Treatment Effects (HTE)

Master's project idea - Department of Mathematics, University of Oslo

What is this about?

Heterogeneous Treatment Effects (HTE) study how the causal effect of an intervention varies across individuals or subgroups. Instead of a single average effect, we target the conditional average treatment effect (CATE), which reveals who benefits most (or least) and supports personalized decisions.

$$\tau(x) = \mathbb{E}[Y(1) - Y(0) \mid X = x]$$

Why this matters in practice

HTE is central to precision medicine, education, public policy, and online platforms: it helps allocate scarce resources to those who benefit most, reduce harm in subgroups with adverse effects, and design fair, effective, and transparent treatment regimes. Organizations increasingly need data-driven policies that adapt to context, and HTE delivers the statistical foundation for such individualized decisions.

Key challenges (statistical and applied)

- Confounding & overlap: making effects identifiable and stable when covariates are high-dimensional and overlap is limited.
- Flexibility vs. overfitting: learning rich response surfaces while controlling complexity and ensuring smoothness.
- Uncertainty quantification: valid intervals/bands for CATE and downstream policy value, robust to model misspecification.
- Deployment: translating estimates into implementable treatment rules that are interpretable, fair, and cost-aware.

Project focus

This project explores regularization for HTE estimation - either Bayesian (e.g., Gaussian processes, hierarchical shrinkage priors, or Bayesian additive models) or machine-learning-based (e.g., penalized trees/forests, boosting, spline/RKHS penalties). The goal is to build estimators that are flexible yet smooth, with proper uncertainty quantification (posterior intervals, bootstrap, or conformal methods), and to use them to construct treatment regimes such as $d(x) = 1\{\tau(x) > c\}$, evaluate policy value, and study fairness/trade-offs.

Simple maths of the problem

Identification under standard assumptions (ignorability and overlap):

$$\tau(x) = \mathbb{E}[Y \mid A = 1, X = x] - \mathbb{E}[Y \mid A = 0, X = x]$$

Learning objective with regularization and a flexible model f :


$$\min_f \sum_i \ell(Y_i, f(X_i, A_i)) + \lambda \Omega(f), \quad \tau(x) = f(x, 1) - f(x, 0)$$

Policy and value evaluation:

$$d(x) = \mathbf{1}\{\tau(x) > c\}, \quad V(d) = \mathbb{E}[Y^{d(x)}]$$

Didactic plots

Figure 1 shows a CATE curve with an uncertainty band: heterogeneity suggests that simple averages can miss who actually benefits, and the band highlights the need for regularization and valid uncertainty quantification. Figure 2 displays an uplift curve: sorting individuals by estimated benefit illustrates how targeting can increase average outcomes and motivates policy learning and evaluation. Figure 3 depicts a simple decision rule $d(x)=\mathbf{1}\{\hat{\tau}(x)>0\}$, connecting estimation to actionable treatment regimes and opening questions about fairness, risk-sensitive thresholds, and robustness under limited overlap.



Professor & Postdoc contacts

For supervision and questions, please contact:

- Professor Johan Pensar <johanpen@uio.no>
- Postdoc Johan de Aguas <johanmd@math.uio.no>

Deliverables: an HTE estimator with regularization, empirical evaluation on a public dataset, and a short report with policy analysis.