

Recovering Sparse and Interpretable Subgroups with Heterogeneous Treatment Effects with Censored Time-to-Event Outcomes

Chirag Nagpal

Vedant Sanil

Artur Dubrawski

Auton Lab, Carnegie Mellon

CHIRAGN@CS.CMU.EDU

VSANIL@ANDREW.CMU.EDU

AWD@CS.CMU.EDU

Abstract

Studies involving both randomized experiments as well as observational data typically involve time-to-event outcomes such as time-to-failure, death or onset of an adverse condition. Such outcomes are typically subject to censoring due to loss of follow-up and established statistical practice involves comparing treatment efficacy in terms of hazard ratios between the treated and control groups. In this paper we propose a statistical approach to recovering sparse phenogroups (or subtypes) that demonstrate differential treatment effects as compared to the study population. Our approach involves modelling the data as a mixture while enforcing parameter shrinkage through structured sparsity regularization. We propose a novel inference procedure for the proposed model and demonstrate its efficacy in recovering sparse phenotypes across large landmark real world clinical studies in cardiovascular health.

Keywords: Time-to-Event, Survival Analysis, Heterogeneous Treatment Effects, Hazard Ratio

1. Introduction

Data driven decision making across multiple disciplines including healthcare, epidemiology, econometrics and prognostics often involves establishing efficacy of an intervention when outcomes are measured in terms of the time to an adverse event, such as death, failure or onset of a critical condition. Typically the analysis of such studies involves assigning a patient population to two or more different treatment arms often called the ‘treated’ (or ‘exposed’) group and the ‘control’ (or ‘placebo’) group and observing whether the populations experience an adverse event (for instance death or onset of a disease) over the study period at a rate that is higher (or lower) than for the control group. Efficacy of a treatment is thus established by comparing the relative difference in the rate of event incidence between the two arms called the hazard ratio. However, not all individuals benefit equally from an intervention. Thus, very often potentially beneficial interventions are discarded even though there might exist individuals who benefit, as the population level estimates of treatment efficacy are inconclusive.

In this paper we assume that patient responses to an intervention are typically heterogeneous and there exists patient subgroups that are unaffected by (or worse, **harmed**) by the intervention being assessed. The ability to discover or phenotype these patients is thus clinically useful as it would allow for more precise clinical decision making by identifying individuals that actually benefit from the intervention being assessed.

Towards this end, we propose **Sparse Cox Subgrouping**, (SCS) a latent variable approach to model patient subgroups that demonstrate heterogeneous effects to an intervention. As opposed to existing literature in modelling heterogeneous treatment effects with censored time-to-event outcomes our approach involves structured regularization of the covariates that assign individuals to

subgroups leading to parsimonious models resulting in phenogroups that are interpretable. We release a python implementation of the proposed SCS approach as part of the auton-survival package (Nagpal et al., 2022b) for survival analysis:

<https://autonlab.github.io/auton-survival/>

2. Related Work

Large studies especially in clinical medicine and epidemiology involve outcomes that are time-to-events such as death, or an adverse clinical condition like stroke or cancer. Treatment efficacy is typically estimated by comparing event rates between the treated and control arms using the Proportional Hazards (Cox, 1972) model and its extensions.

Identification of subgroups in such scenarios has been the subject of a large amount of traditional statistical literature. Large number of such approaches involve estimation of the factual and counterfactual outcomes using separate regression models (T-learner) followed by regressing the difference between these estimated potential outcomes. Within this category of approaches, Lipkovich et al. (2011) propose the subgroup identification based on differential effect search (SIDES) algorithm, Su et al. (2009) propose a recursive partitioning method for subgroup discovery, Dusseldorp and Mechelen (2014) propose the qualitative interaction trees (QUINT) algorithm, and Foster et al. (2011) propose the virtual twins (VT) method for subgroup discovery involving decision tree ensembles. We include a parametric version of such an approach as a competing baseline.

Identification of heterogeneous treatment effects (HTE) is also of growing interest to the machine learning community with multiple approaches involving deep neural networks with balanced representations (Sontag et al., 2017; Johansson et al., 2020), generative models Louizos et al. (2017) as well as Non-Parametric methods involving random-forests (Wager and Athey, 2018) and Gaussian Processes (Alaa and Van Der Schaar, 2017). There is a growing interest in estimating HTEs from an interpretable and trustworthy standpoint (Lee et al., 2020; Nagpal et al., 2020; Morucci et al., 2020; Wu et al., 2022; Crabbé et al., 2022). Wang and Rudin (2022) propose a sampling based approach to discovering interpretable rule sets demonstrating HTEs.

However large part of this work has focused extensively on outcomes that are binary or continuous. The estimation of HTEs in the presence of censored time-to-events has been limited. Xu et al. (2022) explore the problem and describe standard approaches to estimate treatment effect heterogeneity with survival outcomes. They also describe challenges associated with existing risk models when assessing treatment effect heterogeneity in the case of cardiovascular health.

There has been some initial attempt to use neural network for causal inference with censored time-to-event outcomes. Curth et al. (2021) propose a discrete time method along with regularization to match the treated and control representations. Chapfuwa et al. (2021)'s approach is related and involves the use of normalizing flows to estimate the potential time-to-event distributions under treatment and control. While our contributions are similar to Nagpal et al. (2022a), in that we assume treatment effect heterogeneity through a latent variable model, our contribution differs in that 1) Our approach is free of the expensive Monte-Carlo sampling procedure and 2) Our generalized EM inference procedure allows us to naturally incorporate structured sparsity regularization, which helps recovers phenogroups that are parsimonious in the features they recover that define subgroups.

Survival and time-to-event outcomes occur pre-eminently in areas of cardiovascular health. One such area is reducing combined risk of adverse outcomes from atherosclerotic disease¹ (Herrington et al., 2016; Furberg et al., 2002; Group, 2009; Buse et al., 2007) The ability of recovering groups with differential benefits to interventions can thus lead to improved patient care through framing of optimal clinical guidelines.

3. Proposed Model: Sparse Cox Subgrouping




Figure 1: Potential outcome distributions under the assumptions of treatment effect heterogeneity. **Case 1:** Amongst the treated population, conditioned on the latent Z , there are two subgroups that **benefit** and are **unaffected** by the intervention. **Case 2:** There is an additional latent subgroup conditioned on which, the treated population is **harmed** with a worse survival rate.

Notation As is standard in survival analysis, we assume that we either observe the true time-to-event or the time of censoring $U = \min\{T, C\}$ indicated by the censoring indicator defined as $\Delta = \mathbf{1}\{T < C\}$. We thus work with a dataset of right censored observations in the form of 4-tuples, $\mathcal{D} = \{(\mathbf{x}_i, \delta_i, \mathbf{u}_i, \mathbf{a}_i)\}_{i=1}^n$, where $\mathbf{u}_i \in \mathbb{R}^+$ is the time-to-event or censoring as indicated by $\delta_i \in \{0, 1\}$, $\mathbf{a}_i \in \{0, 1\}$ is the indicator of treatment assignment, and \mathbf{x}_i are individual covariates that confound the treatment assignment and the outcome.

ASSUMPTION 1 (Independent Censoring). *The time-to-event T and the censoring distribution C are independent conditional on the covariates X and the intervention A .*

Model Consider a maximum likelihood approach to model the data \mathcal{D} the set of parameters Ω . Under Assumption 1 the likelihood of the data \mathcal{D} can be given as,

$$\mathcal{L}(\Omega; \mathcal{D}) \propto \prod_{i=1}^{|\mathcal{D}|} \lambda(u_i | X = \mathbf{x}_i, A = \mathbf{a}_i)^{\delta_i} S(u_i | X = \mathbf{x}_i, A = \mathbf{a}_i), \quad (1)$$

1. A class of related clinical conditions from increasing deposits of plaque in the arteries, leading to Stroke, Myocardial Infarction and other Coronary Heart Diseases.

here $\lambda(t) = \lim_{\Delta t \rightarrow 0} \frac{\mathbb{P}(t \leq T < t + \Delta t | T \geq t)}{\Delta t}$ is the hazard and $S(t) = \mathbb{P}(T > t)$ is the survival rate.

ASSUMPTION 2 (PH). *The distribution of the time-to-event T conditional on the covariates and the treatment assignment obeys proportional hazards.*

From Assumption 2 (Proportional Hazards), an individual with covariates ($X = \mathbf{x}$) under intervention ($A = \mathbf{a}$) under a Cox model with parameters β and treatment effect ω is given as

$$\lambda(t|A = \mathbf{a}, X = \mathbf{x}) = \lambda_0(t) \exp(\beta^\top \mathbf{x} + \omega \cdot \mathbf{a}), \quad (2)$$

Here, $\lambda_0(\cdot)$ is an infinite dimensional parameter known as the base survival rate. In practice in the Cox's model the base survival rate is a nuisance parameter and is estimated non-parametrically. In order to model the heterogeneity of treatment response. We will now introduce a latent variable $Z \in \{0, 1, -1\}$ that mediates treatment response to the model,

$$\begin{aligned} \lambda(t|A = \mathbf{a}, X = \mathbf{x}, Z = \mathbf{k}) &= \lambda_0(t) \exp(\beta^\top \mathbf{x}) \exp(\omega)^{k\mathbf{a}}, \\ \text{and, } \mathbb{P}(Z = \mathbf{k}|X = \mathbf{x}) &= \frac{\exp(\boldsymbol{\theta}_k^\top \mathbf{x})}{\sum_j \exp(\boldsymbol{\theta}_j^\top \mathbf{x})}. \end{aligned} \quad (3)$$

Here, $\omega \in \mathbb{R}$ is the treatment effect, and $\boldsymbol{\theta} \in \mathbb{R}^{k \times d}$ is the set of parameters that mediate assignment to the latent group Z conditioned on the confounding features \mathbf{x} . Note that the above choice of parametrization naturally enforces the requirements from the model as in Figure 1. Consider the following scenarios,

Case 1: The study population consists of two sub-strata ie. $Z \in \{0, +1\}$, that are benefit and are unaffected by treatment respectively.

Case 2: The study population consists of three sub-strata ie. $Z \in \{0, +1, -1\}$, that benefit, are harmed or unaffected by treatment respectively.

Following from Equations 1 & 3, the complete likelihood of the data \mathcal{D} under this model is,

$$\begin{aligned} \mathcal{L}(\boldsymbol{\Omega}; \mathcal{D}) &= \prod_{i=1}^{|\mathcal{D}|} \sum_{k \in Z} \left(\lambda_0(u_i) h(\mathbf{x}, \mathbf{a}, \mathbf{k}) \right)^{\delta_i} S_0(u_i)^{h(\mathbf{x}, \mathbf{a}, \mathbf{k})} \mathbb{P}(Z = k | X = \mathbf{x}_i) \\ \text{where, } \ln h(\mathbf{x}, \mathbf{a}) &= \beta^\top \mathbf{x} + \mathbf{k} \cdot \mathbf{a} \cdot \omega \text{ and } \ln S_0(\cdot) = -\mathbf{\Lambda}_0(\cdot), \end{aligned} \quad (4)$$

Note that $\mathbf{\Lambda}_0(\cdot) = \int_0^t \lambda_0(\cdot)$ is the infinite dimensional cumulative hazard and is inferred when learning the model. We will notate the set of all learnable parameters as $\boldsymbol{\Omega} = \{\boldsymbol{\theta}, \beta, \omega, \mathbf{\Lambda}_0\}$.

Shrinkage In retrospective analysis to recover treatment effect heterogeneity a natural requirement is parsimony of the recovered subgroups in terms of the covariates to promote model interpretability. Such parsimony can be naturally enforced through appropriate shrinkage on the coefficients that promote sparsity. We want to recover phenogroups that are 'sparse' in $\boldsymbol{\theta}$. We enforce sparsity in the parameters of the latent Z gating function via a group ℓ_1 (Lasso) penalty. The final loss function to be optimized including the group sparsity regularization term is,

$$\begin{aligned} \mathcal{L}(\boldsymbol{\Omega}; \mathcal{D}) + \epsilon \cdot \mathcal{R}(\boldsymbol{\theta}) \text{ where, } \mathcal{R}(\boldsymbol{\theta}) &= \sum_d \sqrt{\sum_{k \in Z} (\boldsymbol{\theta}_d^k)^2} \\ \text{and } \epsilon > 0 &\text{ is the strength of the shrinkage parameter.} \end{aligned} \quad (5)$$

Identifiability Further, to ensure identifiability we restrict the gating parameters for the ($Z = 0$) to be 0. Thus $\theta_1 = 0$.

Inference We will present a variation of the **Expectation Maximization** algorithm to infer the parameters in Equation 3. Our approach differs from Nagpal et al. (2022a, 2021) in that it does not require stochastic Monte-Carlo sampling. Further, our generalized EM inference allows for incorporation of the structured sparsity in the **M-Step**.

A Semi-Parametric $Q(\cdot)$ Note that the likelihood in Equation 3 is semi-parametric and consists of parametric components and the infinite dimensional base hazard $\Lambda(\cdot)$. We define the $Q(\cdot)$ as:

$$Q(\Omega; \mathcal{D}) = \sum_{i=1}^n \sum_{k \in \mathcal{Z}} \gamma_i^k \left(\ln \mathbf{p}_\theta(Z = k | X = \mathbf{x}_i) + \ln \mathbf{p}_{\omega, \beta, \Lambda}(T | Z = k, X = \mathbf{x}_i) \right) + \mathcal{R}(\theta)$$

The E-Step Requires computation of the posteriors counts $\gamma := \mathbf{p}(Z = k | T, X = \mathbf{x}, A = \mathbf{a})$.

RESULT 1 (Posterior Counts). *The posterior counts γ for the latent Z are estimated as,*

$$\begin{aligned} \gamma^k &= \widehat{\mathbb{P}}(Z = k | X = \mathbf{x}, A = \mathbf{a}, \mathbf{u}) \\ &= \frac{\mathbb{P}(\mathbf{u} | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) \mathbb{P}(Z = \mathbf{k} | X = \mathbf{x})}{\sum_k \mathbb{P}(\mathbf{u} | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) \mathbb{P}(Z = \mathbf{k} | X = \mathbf{x})} \\ &= \frac{\mathbf{h}(\mathbf{x}, \mathbf{a}, \mathbf{k})^{\delta_i} \widehat{\mathbf{S}}_0(\mathbf{u})^{h(\mathbf{x}, \mathbf{a}, \mathbf{k})} \exp(\boldsymbol{\theta}_k^\top \mathbf{x})}{\sum_{j \in \mathcal{Z}} \mathbf{h}(\mathbf{x}, \mathbf{a}, \mathbf{j})^{\delta_i} \widehat{\mathbf{S}}_0(\mathbf{u})^{h(\mathbf{x}, \mathbf{a}, \mathbf{j})} \exp(\boldsymbol{\theta}_j^\top \mathbf{x})}. \end{aligned} \quad (6)$$

For a full discussion on derivation of the $Q(\cdot)$ and the posterior counts please refer to Appendix B

The M-Step Involves maximizing the $Q(\cdot)$ function. Rewriting the $Q(\cdot)$ as a sum of two terms,

$$Q(\Omega) = \underbrace{\sum_{i=1}^n \sum_{k \in \mathcal{Z}} \gamma_i^k \ln \mathbf{p}_{\omega, \beta, \Lambda_0}(T | Z = k, X = \mathbf{x}_i, A = \mathbf{a}_i)}_{\mathbf{A}(\omega, \beta, \Lambda_0)} + \underbrace{\sum_{i=1}^n \sum_{k \in \mathcal{Z}} \gamma_i^k \ln \mathbf{p}_\theta(Z = k | X = \mathbf{x}_i)}_{\mathbf{B}(\theta)} + \mathcal{R}(\theta)$$

RESULT 2 (Weighted Cox model). *The term \mathbf{A} can be rewritten as a weighted Cox model and thus optimized using the corresponding ‘partial likelihood’,*

Updates for $\{\beta, \omega\}$: The partial-likelihood, $\mathcal{P}\mathcal{L}(\cdot)$ under sampling weights (Binder, 1992) is

$$\mathcal{P}\mathcal{L}(\Omega; \mathcal{D}) = \sum_{i=1, \delta_i=1}^n \sum_{k \in \mathcal{Z}} \gamma_i^k \left(\ln \mathbf{h}_k(\mathbf{x}_i, \mathbf{a}_i, \mathbf{k}) - \ln \sum_{j \in \text{RiskSet}(u_i)} \sum_{k \in \mathcal{Z}} \gamma_j^k \mathbf{h}_k(\mathbf{x}_j, \mathbf{a}_j, \mathbf{k}) \right) \quad (7)$$

Here RiskSet(\cdot) is the ‘risk set’ or the set of all individuals who haven’t experienced the event till the corresponding time, i.e. RiskSet(t) := $\{i : u_i > t\}$. $\mathcal{P}\mathcal{L}(\cdot)$ is convex in $\{\beta, \omega\}$ and we update these with a gradient step.

Updates for Λ_0 : The base hazard Λ_0 are updated using a weighted Breslow's estimate (Breslow, 1972; Lin, 2007) assuming the posterior counts γ to be sampling weights (Chen, 2009),

$$\widehat{\Lambda}_0(t)^+ = \sum_{i=1}^n \sum_{k \in \mathcal{Z}} \mathbf{1}\{u_i < t\} \frac{\gamma_i^k \cdot \delta_i}{\sum_{j \in \text{RiskSet}(u_i)} \sum_{k \in \mathcal{Z}} \gamma_j^k \mathbf{h}_k(\mathbf{x}_j, \mathbf{a}_j, \mathbf{k})} \quad (8)$$

Term \mathbf{B} is a function of the gating parameters θ that determine the latent assignment Z along with sparsity regularization. We update \mathbf{B} using a Proximal Gradient update as is the case with Iterative Soft Thresholding (ISTA) for group sparse ℓ_1 regression.

Updates for θ : The proximal update for θ including the group regularization (Friedman et al., 2010) term $\mathcal{R}(\cdot)$ is,

$$\widehat{\theta}^+ = \text{prox}_{\eta\epsilon} \left(\theta - \frac{d}{d\theta} \mathbf{B}(\theta) \right), \text{ where } \text{prox}_{\eta\epsilon}(\mathbf{y}) = \frac{\mathbf{y}}{\|\mathbf{y}\|_2} \max\{0, \|\mathbf{y}\|_2 - \eta\epsilon\}. \quad (9)$$

All together the inference procedure is described in Algorithm 1.

Algorithm 1: Parameter Learning for SCS with a Generalized EM

Input : Training set, $\mathcal{D} = \{(\mathbf{x}_i, u_i, \mathbf{a}_i, \delta_i)_{i=1}^n\}$; maximum EM iterations, B , step size η

```

while <not converged> do
  for  $b \in \{1, 2, \dots, B\}$  do
    E-STEP .....
     $\gamma_i^k = \frac{\mathbf{h}(\mathbf{x}, \mathbf{a}, \mathbf{k})^{\delta_i} \widehat{S}_0(\mathbf{u})^{\mathbf{h}(\mathbf{x}, \mathbf{a}, \mathbf{k})} \exp(\theta_k^\top \mathbf{x})}{\sum_{j \in \mathcal{Z}} \mathbf{h}(\mathbf{x}, \mathbf{a}, \mathbf{j})^{\delta_i} \widehat{S}_0(\mathbf{u})^{\mathbf{h}(\mathbf{x}, \mathbf{a}, \mathbf{j})} \exp(\theta_j^\top \mathbf{x})}$   $\triangleright$  Compute posterior counts (Equation 6).
    M-STEP .....
     $\widehat{\beta}^+ \leftarrow \widehat{\beta} - \eta \cdot \nabla_{\beta} \mathcal{P}\mathcal{L}(\beta, \mathbf{w}; \mathcal{D})$ 
     $\widehat{\mathbf{w}}^+ \leftarrow \widehat{\mathbf{w}} - \eta \cdot \nabla_{\mathbf{w}} \mathcal{P}\mathcal{L}(\beta, \mathbf{w}; \mathcal{D})$   $\triangleright$  Gradient descent update.
     $\widehat{\Lambda}_0(t)^+ \leftarrow \sum_{i=1}^n \sum_{k \in \mathcal{Z}} \mathbf{1}\{u_i < t\} \frac{\gamma_i^k \cdot \delta_i}{\sum_{j \in \text{RiskSet}(u_i)} \sum_{k \in \mathcal{Z}} \gamma_j^k \mathbf{h}_k(\mathbf{x}_j, \mathbf{a}_j, \mathbf{k})}$   $\triangleright$  Breslow (1972)'s estimator.
     $\widehat{\theta}^+ \leftarrow \widehat{\theta} - \eta \cdot \nabla_{\theta} \mathbf{B}(\theta)$   $\triangleright$  Update  $\theta$  with gradient of  $\widehat{Q}$ .
     $\widehat{\theta}^+ \leftarrow \text{prox}_{\epsilon\eta}(\widehat{\theta})$   $\triangleright$  Proximal update.
  end
end

```

Return: learnt parameters Ω ;

4. Experiments

In this section we describe the experiments conducted to benchmark the performance of SCS against alternative models for heterogeneous treatment effect estimation on multiple studies including a synthetic dataset and multiple large landmark clinical trials for cardiovascular diseases.

4.1. Simulation




Figure 2: a) Population level Kaplan-Meier Estimates of the Survival Distribution stratified by the treatment assignment. b) Distribution of the Latent Z in X and the recovered decision boundary by SCS. c) Receiver Operator Characteristics of SCS in recovering the true phenotype.




Figure 3: The phenotypes recovered with Sparse Cox Subgrouping on the Synthetic Data. As expected, the recovered phenotypes conform to the modelling assumptions as in Figure 4.

In this section we first describe the performance of the proposed Sparse Cox Subgrouping approach on a synthetic dataset designed to demonstrate heterogeneous treatment effects. We randomly assign individuals to the treated or control group. The latent variable Z is drawn from a uniform categorical distribution that determines the subgroup,

$$A \sim \text{Bernoulli}(1/2), \quad Z \sim \text{Categorical}(1/3)$$

Conditioned on Z we sample $X_{1:2} \sim \text{Normal}(\mu_z, \sigma_z)$ as in Figure 2 that determine the conditional Hazard Ratios $\text{HR}(k)$, and randomly sample noisy covariates $X_{3:6} \sim \text{Uniform}(-1, 1)$. The true

time-to-event T and censoring times C are then sampled as,

$$T|(X = \mathbf{x}, A = \mathbf{a}, Z = \mathbf{k}) \sim \text{Gompertz}(\beta = 1, \eta = 0.25 \cdot \text{HR}(k)^{\alpha}), \quad C|T \sim \text{Uniform}(0, T)$$

Finally we sample the censoring indicator $\Delta \sim \text{Bernoulli}(0.8)$ and set the observed time-to-event,

$$U = T \text{ if } \Delta = 1, \text{ else we set } U = C.$$

Figure 2 presents the ROC curves for SCS’s ability to identify the groups with enhanced and diminished treatment effects respectively. In Figure 3 we present Kaplan-Meier estimators of the Time-to-Event distributions conditioned on the predicted Z by SCS. Clearly, SCS is able to identify the phenogroups corresponding to differential benefits.

4.2. Recovering subgroups demonstrating Heterogeneous Treatment Effects from Landmark studies of Cardiovascular Health

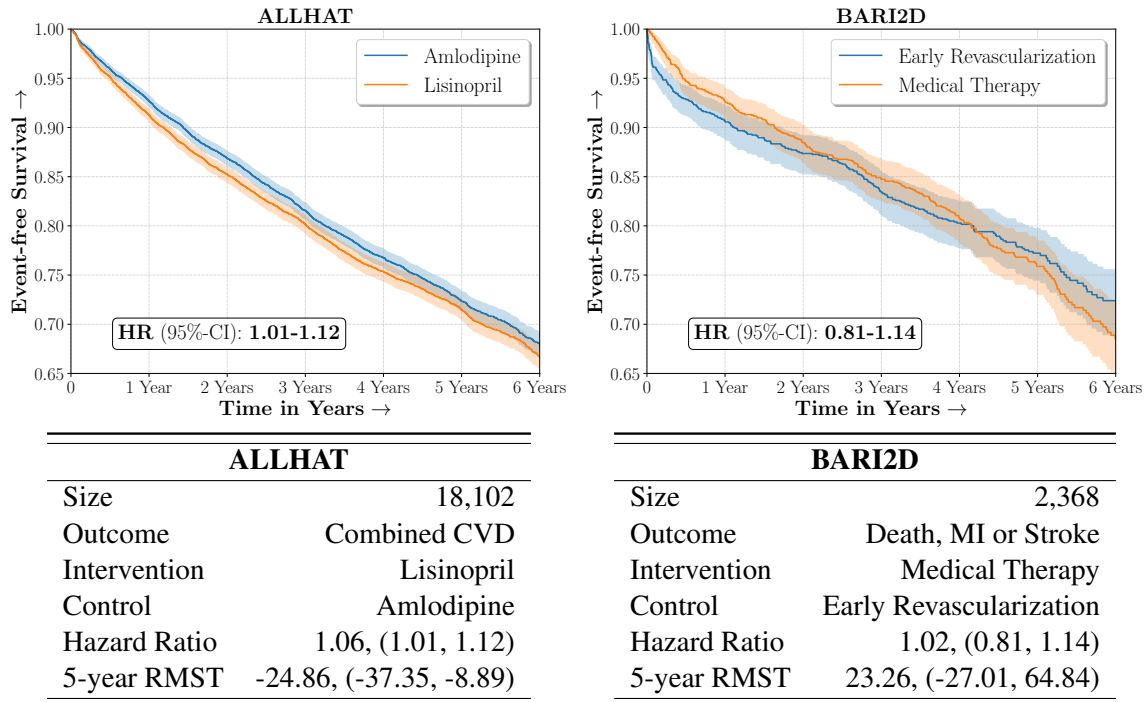


Figure 4: Event-free Kaplan-Meier survival curves stratified by the treatment assignment and summary statistics for the **ALLHAT** and **BARI2D** studies. (Combined CVD: Coronary Heart Disease, Stroke, other treated angina, fatal or non-fatal Heart Failure, and Peripheral Arterial Disease.)

Antihypertensive and Lipid-Lowering Treatment to Prevent Heart Attack (Furberg et al., 2002)

The ALLHAT study was a large randomized experiment conducted to assess the efficacy of multiple classes of blood pressure lowering medicines for patients with hypertension in reducing risk of adverse cardiovascular conditions. We considered a subset of patients from the original **ALLHAT** study who were randomized to receive either Amlodipine (a calcium channel blocker) or Lisinopril (an Angiotensin-converting enzyme inhibitor). Overall, Amlodipine was found to be more efficacious than Lisinopril in reducing combined risk of cardio-vascular disease.

Bypass Angioplasty Revascularization Investigation in Type II Diabetes (Group, 2009)

Diabetic patients have been traditionally known to be at higher risk of cardiovascular disease however appropriate intervention for diabetics with ischemic heart disease between surgical coronary revascularization or management with medical therapy is widely debated. The **BARI2D** was a large landmark experiment conducted to assess efficacy between these two possible medical interventions. Overall **BARI2D** was inconclusive in establishing the appropriate therapy between Coronary Revascularization or medical management for patients with Type-II Diabetes.

Figure 4 presents the event-free survival rates as well as the summary statistics for the studies. In our experiments, we included a large set of confounders collected at baseline visit of the patients which we utilize to train the proposed model. A full list of these features are in Appendix A.

4.3. Baselines

Cox PH with ℓ_1 Regularized Treatment Interaction (COX-INT)

We include treatment effect heterogeneity via interaction terms that model the time-to-event distribution using a proportional hazards model as in Kehl and Ulm (2006). Thus,

$$\lambda(t|X = \mathbf{x}, A = \mathbf{a}) = \lambda_0(t) \exp(\beta^\top \mathbf{x} + \mathbf{a} \cdot \theta^\top \mathbf{x}) \tag{10}$$

The interaction effects θ are regularized with a lasso penalty in order to recover a sparse phenotyping rule defined as $G(\mathbf{x}) = \theta^\top \mathbf{x}$.

Binary Classifier with ℓ_1 Regularized Treatment Interaction (BIN-INT)

Instead of modelling the time-to-event distribution we directly model the thresholded survival outcomes $Y = \mathbf{1}\{T < t\}$ at a five year time horizon using a log-linear parameterization with a logit link function. As compared to COX-INT, this model ignores the data-points that were right-censored prior to the thresholded time-to-event, however it is not sensitive to the strong assumption of Proportional Hazards.

$$\mathbb{E}[T > t|X = \mathbf{x}, A = \mathbf{a}] = \sigma(\beta^\top \mathbf{x} + \beta_0 + \mathbf{a} \cdot \theta^\top \mathbf{x}),$$

and, $\sigma(\cdot)$ is the logistic link function. (11)

Cox PH T-Learner with ℓ_1 Regularized Logistic Regression (COX-TLR)

We train two separate Cox Regression models on the treated and control arms (T-Learner) to estimate the potential outcomes under treatment ($A = 1$) and control ($A = 0$). Motivated from the ‘Virtual Twins’ approach as in Foster et al. (2011), a logistic regression with an ℓ_1 penalty is trained to estimate if the risk of the potential outcome under treatment is higher than under control. This logistic regression is then employed as the phenotyping function $G(\cdot)$ and is given as,

$$G(\mathbf{x}) = \mathbb{E}[\mathbf{1}\{f_1(\mathbf{x}, t) > f_0(\mathbf{x}, t)\}|X = \mathbf{x}]$$

where, $f_a(\mathbf{x}, t) = \mathbb{P}(T > t|\text{do}(A = \mathbf{a}), X = \mathbf{x})$. (12)

The above models involving sparse ℓ_1 regularization were trained with the glmnet (Friedman et al., 2009) package in R.

The ACC/AHA Long term Atheroscleoratic Cardiovascular Risk Estimate ²

The American College of Cardiology and the American Heart Association model for estimation of risk of Atherosclerotic disease risk (Goff Jr et al., 2014) involves pooling data from multiple observational cohorts of patients followed by modelling the 10-year risk of an adverse cardiovascular condition including death from coronary heart disease, Non-Fatal Myocardial Infarction or Non-fatal Stroke. While the risk model was originally developed to assess factual risk in the observational sense, in practice it is also employed to assess risk when making counterfactual decisions.

4.4. Results and Discussion

Protocol We compare the performance of SCS and the corresponding competing methods in recovery of subgroups with enhanced (or diminished treatment effects). For each of these studies we stratify the study population into equal sized sets for training and validation while persevering the proportion of individuals that were assigned to treatment and experienced the outcome in the follow up period. The models were trained on the training set and validated on the held-out test set. For each of the methods we experiment with models that do not enforce any sparsity ($\epsilon = 0$) as well as tune the level of sparsity to recover phenotyping functions that involve 5 and 10 features. The subgroup size are varied by controlling the threshold at which the individual is assigned to a group. Finally, the treatment effect is compared in terms of Hazard Ratios, Risk Differences as well as Restricted Mean Survival Time over a 5 Year event period.

Results We present the results of SCS versus the baselines in terms of Hazard Ratios on the **ALLHAT** and **BARI2D** datasets in Figures 5 and 6. In the case of **ALLHAT**, SCS consistently recovered phenogroups with more pronounced (or diminished) treatment effects. On external validation on the heldout dataset, we found a subgroup of patients that had similar outcomes whether assigned to Lisinopril or Amlodipine, whereas the other subgroup clearly identified patients that were harmed with Lisinopril. The group harmed with Lisinopril had higher Diastolic BP. On the other hand, patients with Lower kidney function did not seem to benefit from Amlodipine.

In the case of **BARI2D**, SCS recovered phenogroups that were both harmed as well as benefitted from just medical therapy without revascularization. The patients who were harmed from Medical therapy were typically older, on the other hand the patients who benefitted primarily included patients who were otherwise assigned to receive PCI instead of CABG revascularization, suggesting PCI to be harmful for diabetic patients.

Tables 3 and 4 present the features that were selected by the proposed model for the studies. Additionally, we also report tabulated results involving metrics like risk difference and restricted mean survival time in the Appendix C.

5. Concluding Remarks

We presented Sparse Cox Subgrouping (SCS) a latent variable approach to recover subgroups of patients that respond differentially to an intervention in the presence of censored time-to-event outcomes. As compared to alternative approaches to learning parsimonious hypotheses in such settings,

². <https://tools.acc.org/ascvd-risk-estimator-plus/>

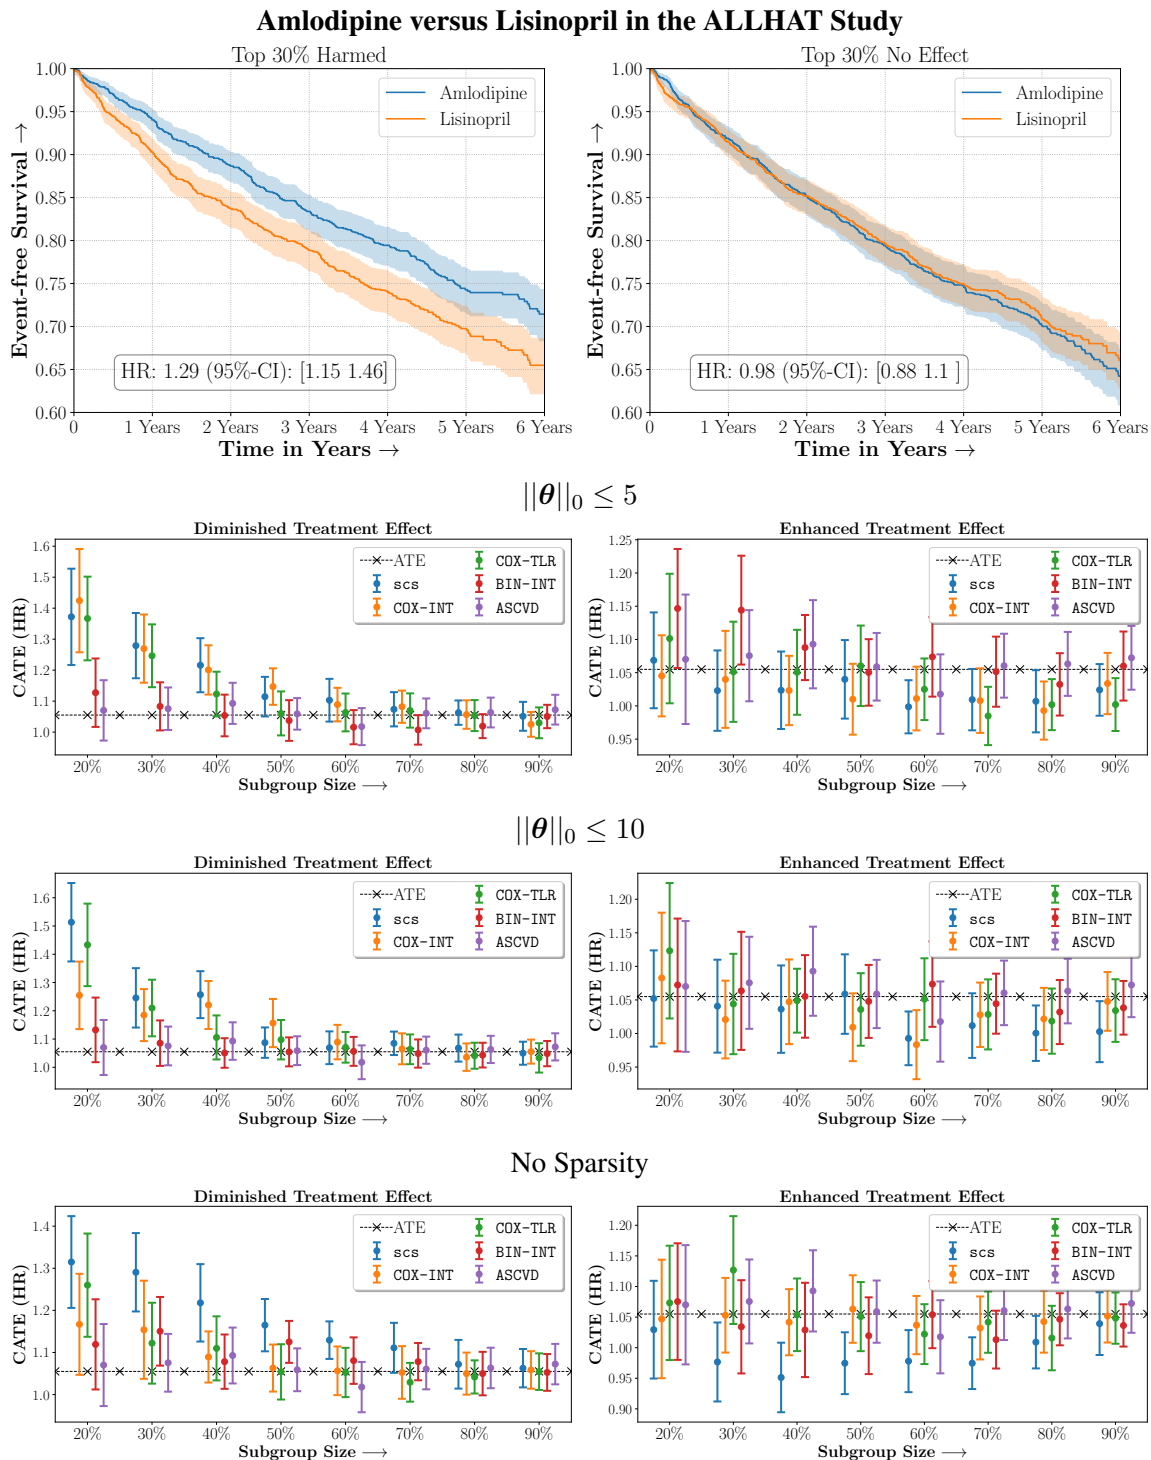
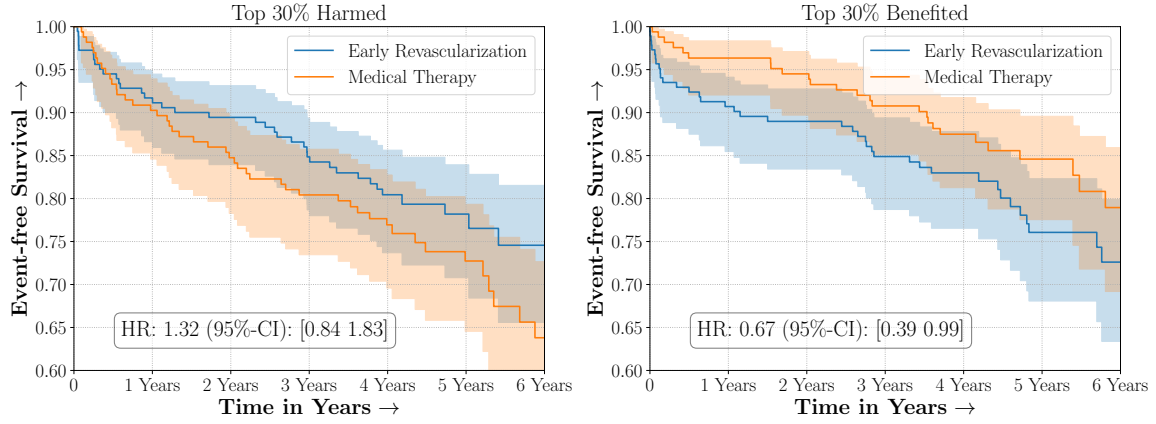



Figure 5: Conditional Average Treatment Effect in Hazard Ratio versus subgroup size for the latent phenogroups extracted from the ALLHAT study.


Early Revascularization versus Medical Therapy in the BARI2D Study



$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity




Figure 6: Conditional Average Treatment Effect in Hazard Ratio versus subgroup size for the latent phenogroups extracted from the **BARI2D** study.

our proposed model recovered hypotheses with more pronounced treatment effects which we validated on multiple studies for cardiovascular health.

While powerful in its ability to recover parsimonious subgroups there exists limitations in SCS in its current form. The model is sensitive to proportional hazards and may be ill-specified when the proportional hazards assumptions are violated as is evident in many real world clinical studies (Maron et al., 2018; Bretthauer et al., 2022). Another limitation is that SCS in its current form looks at only a single endpoint (typically death, or a composite of multiple adverse outcome). In practice however real world studies typically involve multiple end-points. We envision that extensions of SCS would allow patient subgrouping across multiple endpoints, leading to discovery of actionable sub-populations that similarly benefit from the intervention under assessment.

References

- Ahmed M Alaa and Mihaela Van Der Schaar. Bayesian inference of individualized treatment effects using multi-task gaussian processes. *Advances in neural information processing systems*, 30, 2017.
- David A Binder. Fitting cox’s proportional hazards models from survey data. *Biometrika*, 79(1): 139–147, 1992.
- Norman E Breslow. Contribution to discussion of paper by dr cox. *J. Roy. Statist. Soc., Ser. B*, 34: 216–217, 1972.
- Michael Bretthauer, Magnus Løberg, Paulina Wieszczy, Mette Kalager, Louise Emilsson, Kjetil Garborg, Maciej Rupinski, Evelien Dekker, Manon Spaander, Marek Bugajski, et al. Effect of colonoscopy screening on risks of colorectal cancer and related death. *New England Journal of Medicine*, 2022.
- John B Buse, ACCORD Study Group, et al. Action to control cardiovascular risk in diabetes (accord) trial: design and methods. *The American journal of cardiology*, 99(12):S21–S33, 2007.
- Paidamoyo Chapfuwa, Serge Assaad, Shuxi Zeng, Michael J Pencina, Lawrence Carin, and Ricardo Henao. Enabling counterfactual survival analysis with balanced representations. In *Proceedings of the Conference on Health, Inference, and Learning*, pages 133–145, 2021.
- Yi-Hau Chen. Weighted breslow-type and maximum likelihood estimation in semiparametric transformation models. *Biometrika*, 96(3):591–600, 2009.
- David R Cox. Regression models and life-tables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 34(2):187–202, 1972.
- Jonathan Crabbé, Alicia Curth, Ioana Bica, and Mihaela van der Schaar. Benchmarking heterogeneous treatment effect models through the lens of interpretability. *arXiv preprint arXiv:2206.08363*, 2022.
- Alicia Curth, Changhee Lee, and Mihaela van der Schaar. Survite: Learning heterogeneous treatment effects from time-to-event data. *Advances in Neural Information Processing Systems*, 34: 26740–26753, 2021.

- Elise Dusseldorp and Iven Mechelen. Qualitative interaction trees: A tool to identify qualitative treatment-subgroup interactions. *Statistics in medicine*, 33, 01 2014. doi: 10.1002/sim.5933.
- Jared C Foster, Jeremy MG Taylor, and Stephen J Ruberg. Subgroup identification from randomized clinical trial data. *Statistics in medicine*, 30(24):2867–2880, 2011.
- Jerome Friedman, Trevor Hastie, Rob Tibshirani, et al. glmnet: Lasso and elastic-net regularized generalized linear models. *R package version*, 1(4):1–24, 2009.
- Jerome Friedman, Trevor Hastie, and Robert Tibshirani. A note on the group lasso and a sparse group lasso. *arXiv preprint arXiv:1001.0736*, 2010.
- Curt D Furberg et al. Major outcomes in high-risk hypertensive patients randomized to angiotensin-converting enzyme inhibitor or calcium channel blocker vs diuretic: the antihypertensive and lipid-lowering treatment to prevent heart attack trial (allhat). *Journal of the American Medical Association*, 2002.
- David C Goff Jr, Donald M Lloyd-Jones, Glen Bennett, Sean Coady, Ralph B D’agostino, Raymond Gibbons, Philip Greenland, Daniel T Lackland, Daniel Levy, Christopher J O’donnell, et al. 2013 acc/aha guideline on the assessment of cardiovascular risk: a report of the american college of cardiology/american heart association task force on practice guidelines. *Circulation*, 129(25_suppl_2):S49–S73, 2014.
- BARI 2D Study Group. A randomized trial of therapies for type 2 diabetes and coronary artery disease. *New England Journal of Medicine*, 360(24):2503–2515, 2009.
- William Herrington, Ben Lacey, Paul Sherliker, Jane Armitage, and Sarah Lewington. Epidemiology of atherosclerosis and the potential to reduce the global burden of atherothrombotic disease. *Circulation research*, 118(4):535–546, 2016.
- Fredrik D Johansson, Uri Shalit, Nathan Kallus, and David Sontag. Generalization bounds and representation learning for estimation of potential outcomes and causal effects. *arXiv preprint arXiv:2001.07426*, 2020.
- Victoria Kehl and Kurt Ulm. Responder identification in clinical trials with censored data. *Computational Statistics & Data Analysis*, 50(5):1338–1355, 2006.
- Kwonsang Lee, Falco J Bargagli-Stoffi, and Francesca Dominici. Causal rule ensemble: Interpretable inference of heterogeneous treatment effects. *arXiv preprint arXiv:2009.09036*, 2020.
- DY Lin. On the breslow estimator. *Lifetime data analysis*, 13(4):471–480, 2007.
- Ilya Lipkovich, Alex Dmitrienko, Jonathan Denne, and Gregory Enas. Subgroup identification based on differential effect search (sides) – a recursive partitioning method for establishing response to treatment in patient subpopulations. *Statistics in medicine*, 30:2601–21, 07 2011. doi: 10.1002/sim.4289.
- Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. Causal effect inference with deep latent-variable models. *Advances in neural information processing systems*, 30, 2017.

- David J Maron, Judith S Hochman, Sean M O’Brien, Harmony R Reynolds, William E Boden, Gregg W Stone, Sripal Bangalore, John A Spertus, Daniel B Mark, Karen P Alexander, et al. International study of comparative health effectiveness with medical and invasive approaches (ischemia) trial: rationale and design. *American heart journal*, 201:124–135, 2018.
- Marco Morucci, Vittorio Orlandi, Sudeepa Roy, Cynthia Rudin, and Alexander Volfovsky. Adaptive hyper-box matching for interpretable individualized treatment effect estimation. In *Conference on Uncertainty in Artificial Intelligence*, pages 1089–1098. PMLR, 2020.
- Chirag Nagpal, Dennis Wei, Bhanukiran Vinzamuri, Monica Shekhar, Sara E Berger, Subhro Das, and Kush R Varshney. Interpretable subgroup discovery in treatment effect estimation with application to opioid prescribing guidelines. In *Proceedings of the ACM Conference on Health, Inference, and Learning*, pages 19–29, 2020.
- Chirag Nagpal, Steve Yadlowsky, Negar Rostamzadeh, and Katherine Heller. Deep cox mixtures for survival regression. In *Machine Learning for Healthcare Conference*, pages 674–708. PMLR, 2021.
- Chirag Nagpal, Mononito Goswami, Keith Dufendach, and Artur Dubrawski. Counterfactual phenotyping with censored time-to-events. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD ’22, page 3634–3644, New York, NY, USA, 2022a. Association for Computing Machinery. ISBN 9781450393850. doi: 10.1145/3534678.3539110. URL <https://doi.org/10.1145/3534678.3539110>.
- Chirag Nagpal, Willa Potosnak, and Artur Dubrawski. auton-survival: an open-source package for regression, counterfactual estimation, evaluation and phenotyping with censored time-to-event data. *arXiv preprint arXiv:2204.07276*, 2022b.
- David Sontag, Uri Shalit, and Fredrik Johansson. Estimating individual treatment effect: Generalization bounds and algorithms. 01 2017.
- Xiaogang Su, Chih-Ling Tsai, Hansheng Wang, David Nickerson, and Bogong Li. Subgroup analysis via recursive partitioning. *Journal of Machine Learning Research*, 10:141–158, 02 2009. doi: 10.2139/ssrn.1341380.
- Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.
- Tong Wang and Cynthia Rudin. Causal rule sets for identifying subgroups with enhanced treatment effects. *INFORMS Journal on Computing*, 2022.
- Han Wu, Sarah Tan, Weiwei Li, Mia Garrard, Adam Obeng, Drew Dimmery, Shaun Singh, Hanson Wang, Daniel Jiang, and Eytan Bakshy. Interpretable personalized experimentation. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 4173–4183, 2022.
- Yizhe Xu, Nikolaos Ignatiadis, Erik Sverdrup, Scott Fleming, Stefan Wager, and Nigam Shah. Treatment heterogeneity with survival outcomes. *arXiv preprint arXiv:2207.07758*, 2022.

Appendix A. Additional Details on the ALLHAT and BARI 2D Case Studies

Tables 1 and 2 represent additional confounding variables found in the **ALLHAT** and **BARI2D** trials respectively.

Name	Description
ETHNIC	Ethnicity
SEX	Sex of Participant
ESTROGEN	Estrogen supplementation
BLMEDS	Antihypertensive treatment
MISTROKE	History of Stroke
HXCABG	History of coronary artery bypass
STDEPR	Prior ST depression/T-wave inversion
OASCVD	Other atherosclerotic cardiovascular disease
DIABETES	Prior history of Diabetes
HDLLT35	HDL cholesterol <35mg/dl; 2x in past 5 years
LVHECG	LVH by ECG in past 2 years
WALL25	LVH by ECG in past 2 years
LCHD	History of CHD at baseline
CURSMOKE	Current smoking status.
ASPIRIN	Aspirin use
LLT	Lipid-lowering trial
AGE	Age upon entry
BLWGT	Weight upon entry
BLHGT	Height upon entry
BLBMI	Body Mass Index upon entry
BV2SBP	Baseline SBP
BV2DBP	Baseline DBP
APOTAS	Baseline serum potassium
BLGFR	Baseline est glomerular filtration rate
ACHOL	Total Cholesterol
AHDL	Baseline HDL Cholesterol
AFGLUC	Baseline fasting serum glucose

Table 1: List of confounding variables used for experiments involving the **ALLHAT** dataset.

Name	Description
hxmi	History of MI
age	Age upon entry
dbp_stand	Standing diastolic BP
sbp_stand	Standing systolic BP
sex	Sex
asp	Aspirin use
smkcat	Cigarette smoking category
betab	Beta blocker use
ccb	Calcium blocker use
hxhtn	History of hypertension requiring tx
insulin	Insulin use
weight	Weight (kg) upon entry
bmi	BMI upon entry
qabn	Abnormal Q-Wave
trig	Triglycerides (mg/dl) upon entry
dmdur	Duration of diabetes mellitus
ablvef	Left ventricular ejection fraction <50%
race	Race
priorrev	Prior revascularization
hxcva	Cerebrovascular accident
screat	Serum creatinine (mg/dl)
hmg	Statin
hxhypo	History of hypoglycemic episode
hba1c	Hemoglobin A1c(%)
priorstent	Prior stent
spotass	Serum Potassium(mEq/L)
hispanic	Hispanic ethnicity
tchol	Total Cholesterol
hdl	HDL Cholesterol
insul_circ	Circulating insulin (IU/ml)
tzd	Thiazolidinedione
ldl	LDL Cholesterol
tabn	Abnormal T-waves
nsgn	Nonsublingual nitrate
sulf	Sulfonylurea
hxchf	Histoty of congestive heart failure req tx
arb	Angiotensin receptor blocker
acr	Urine albumin/creatinine ratio mg/g
diur	Diuretic
apa	Anti-platelet
hxchl	Hypercholesterolemia req tx
acei	ACE inhibitor
abilow	Low ABI (<= 0.9)
biguanide	Biguanide
stavn	Abnormal ST depression

Table 2: List of confounding variables used for experiments involving the **BARI2D** dataset.

ALLHAT		BARI2D	
Name	Description	Name	Description
BV2SBP	Baseline Seated Diastolic Pressure	age	Age upon entry
BLGFR	Baseline est Glomerular Filtration Rate	asp	Aspirin use
BLMEDS	Antihypertensive Treatment	hxhtn	History of hypertension requiring tx
CURSMOKE	Current Smoking Status	hxchl	Hypercholesterolemia req tx
SEX	Sex of Participant	priorstent	Prior stent

 Table 3: List of selected features with sparsity level: $\|\theta\|_0 \leq 5$

ALLHAT		BARI2D	
Name	Description	Name	Description
BV2SBP	Baseline Seated Diastolic Pressure	age	Age upon entry
BLGFR	Baseline est Glomerular Filtration Rate	asp	Aspirin use
BLMEDS	Antihypertensive Treatment	hxhtn	History of hypertension requiring tx
CURSMOKE	Current Smoking Status	hxchl	Hypercholesterolemia req tx
SEX	Sex of Participant	priorstent	Prior stent
ASPIRIN	Aspirin Use	acei	ACE Inhibitor
ACHOL	Total Cholesterol	acr	Urine albumin/creatinine ratio mg/g
BLWGT	Weight upon entry	insul_circ	Circulating insulin
BMI	Body mass index upon entry	screat	Serum creatinine (mg/dl)
OASCVD	Other atherosclerotic cardiovascular disease	tchol	Total Cholesterol

 Table 4: List of selected features with sparsity level: $\|\theta\|_0 \leq 10$

Appendix B. Derivation of the Inference Algorithm

Censored Instances: Note that in the case of the censored instances we will condition on the thresholded survival ($T > u$). The the posterior counts thus reduce to:

$$\begin{aligned} \gamma^k &= \mathbb{P}(Z = k | X = \mathbf{x}, A = \mathbf{a}, T > u) \\ &= \frac{\mathbb{P}(T > t | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) p(Z = \mathbf{k} | X = \mathbf{x})}{\sum_k \mathbb{P}(T > t | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) \mathbb{P}(Z = \mathbf{k} | X = \mathbf{x})} \end{aligned} \quad (13)$$

Here, $\mathbb{P}(T > t | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) = \exp(-\Lambda(t))^{h(\mathbf{x}, \mathbf{a}, \mathbf{k})}$

Uncensored Instances The posteriors are $\gamma^k = p_{\theta}(Z = k | X = \mathbf{x}, T = u)$,

Posteriors for the uncensored data are more involved and involve the base hazard $\lambda_0(\cdot)$. Posteriors for uncensored data are independent of the base hazard function, $\lambda_0(\cdot)$ as,

$$\gamma^k = \frac{\lambda_0(u) h_k(\mathbf{x}, \mathbf{a}) S_0(u_i)^{h_k(\mathbf{x}, \mathbf{a})}}{\sum_k \lambda_0(u) h_k(\mathbf{x}, \mathbf{a}) S_0(u)^{h_k(\mathbf{x}, \mathbf{a})}} = \frac{h_k(\mathbf{x}, \mathbf{a}) S_0(u_i)^{h_k(\mathbf{x}, \mathbf{a})}}{\sum_k h_k(\mathbf{x}, \mathbf{a}) S_0(u_i)^{h_k(\mathbf{x}, \mathbf{a})}}$$

Combining Equations 13 and 14 we arrive at the following estimate for the posterior counts

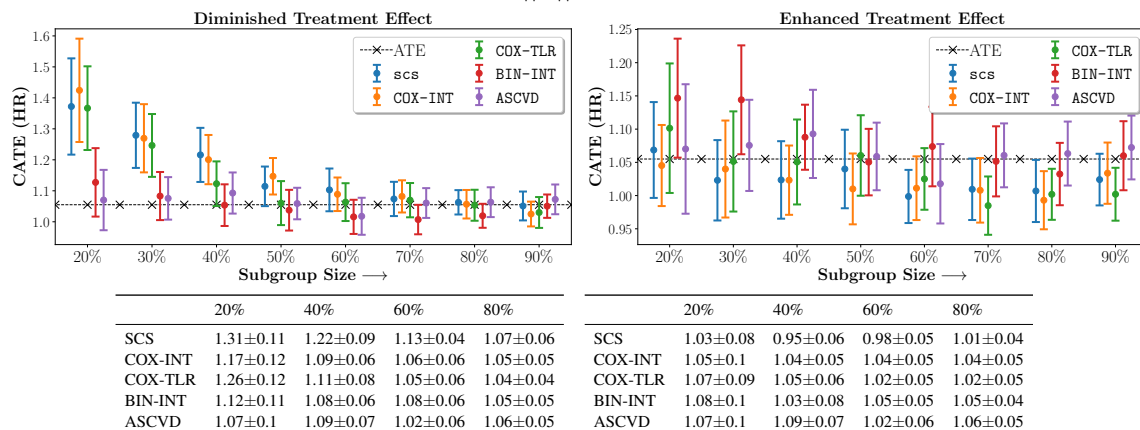
$$\begin{aligned} \gamma^k &= \hat{\mathbb{P}}(Z = k | X = \mathbf{x}, A = \mathbf{a}, u) \\ &= \frac{\mathbb{P}(u | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) \mathbb{P}(Z = \mathbf{k} | X = \mathbf{x})}{\sum_k \mathbb{P}(u | Z = \mathbf{k}, X = \mathbf{x}, A = \mathbf{a}) \mathbb{P}(Z = \mathbf{k} | X = \mathbf{x})} \\ &= \frac{h(\mathbf{x}, \mathbf{a}, \mathbf{k})^{\delta_i} \hat{S}_0(u)^{h(\mathbf{x}, \mathbf{a}, \mathbf{k})} \exp(\theta_{\mathbf{k}}^\top \mathbf{x})}{\sum_{j \in \mathcal{Z}} h(\mathbf{x}, \mathbf{a}, \mathbf{j})^{\delta_i} \hat{S}_0(u)^{h(\mathbf{x}, \mathbf{a}, \mathbf{j})} \exp(\theta_{\mathbf{j}}^\top \mathbf{x})}. \end{aligned} \quad (14)$$

Appendix C. Additional Results

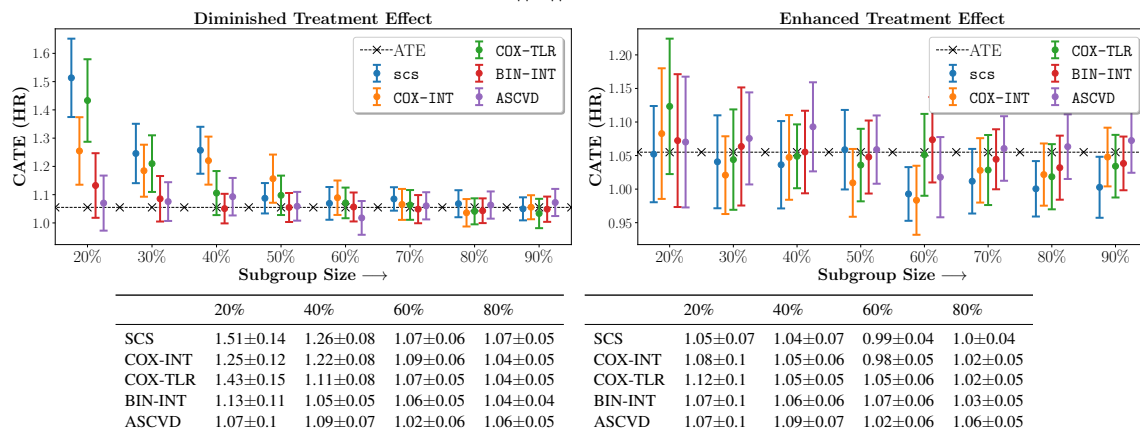
Figures 7, 8, 9 present tabulated metrics on **ALLHAT** with Hazard Ratio, Risk Difference and Restricted Mean Survival Time respectively. Figures 10, 11, 12 present tabulated metrics **BARI2D** with Hazard Ratio, Risk Difference and Restricted Mean Survival Time metrics respectively.

SPARSE COX SUBGROUPS

$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity

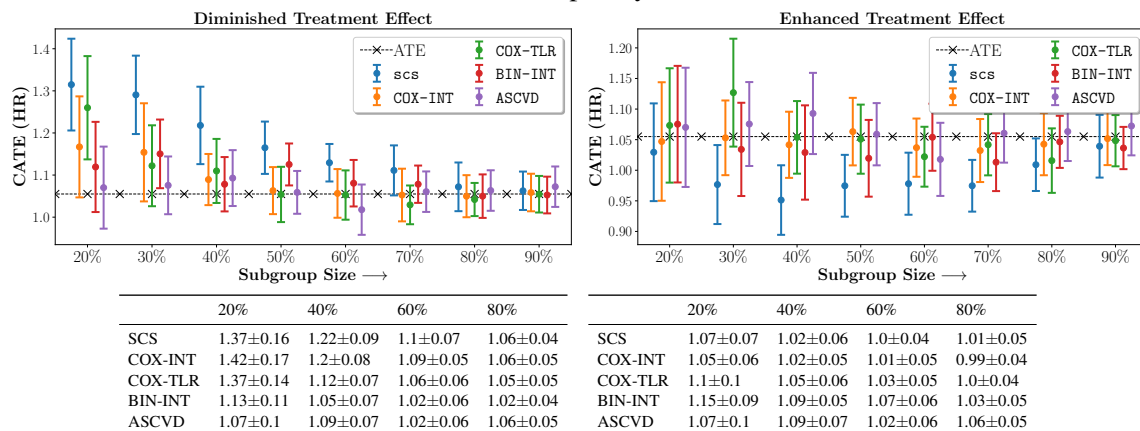
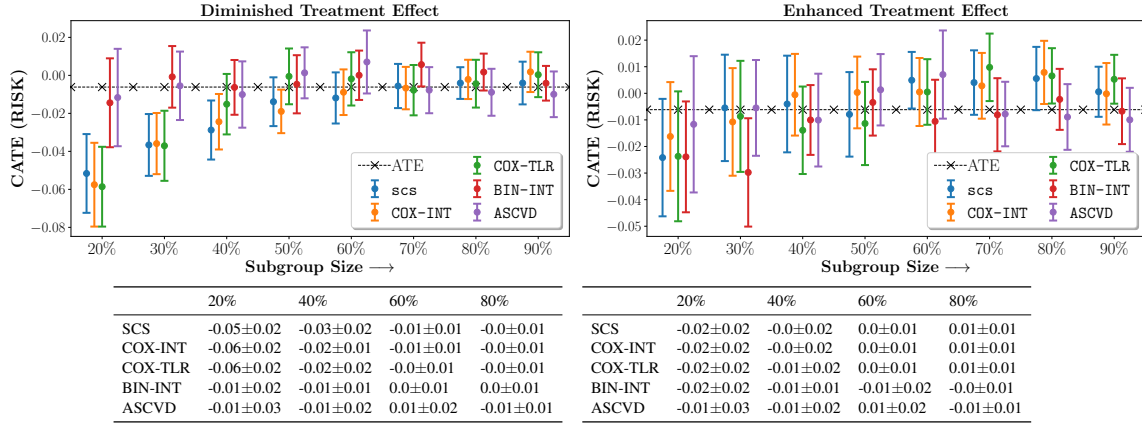
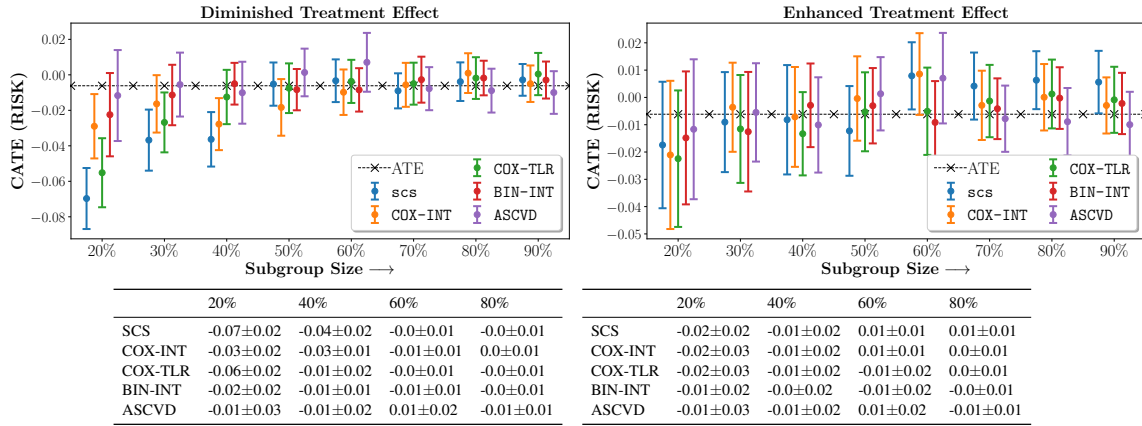


Figure 7: Conditional Average Treatment Effect in Hazard Ratio versus subgroup size for the latent phenogroups extracted from the ALLHAT study.

$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity

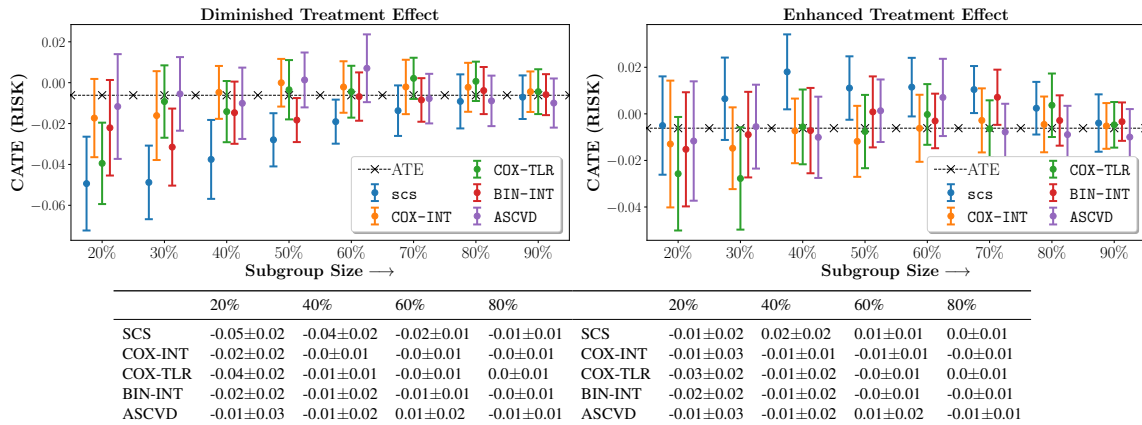



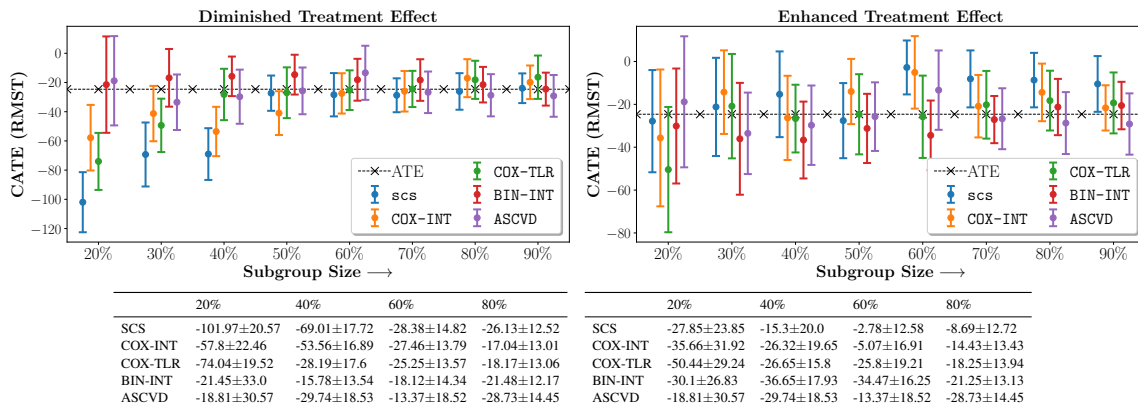
Figure 8: Conditional Average Treatment Effect in Risk versus subgroup size for the latent phenogroups extracted from the ALLHAT study.

SPARSE COX SUBGROUPS

$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity


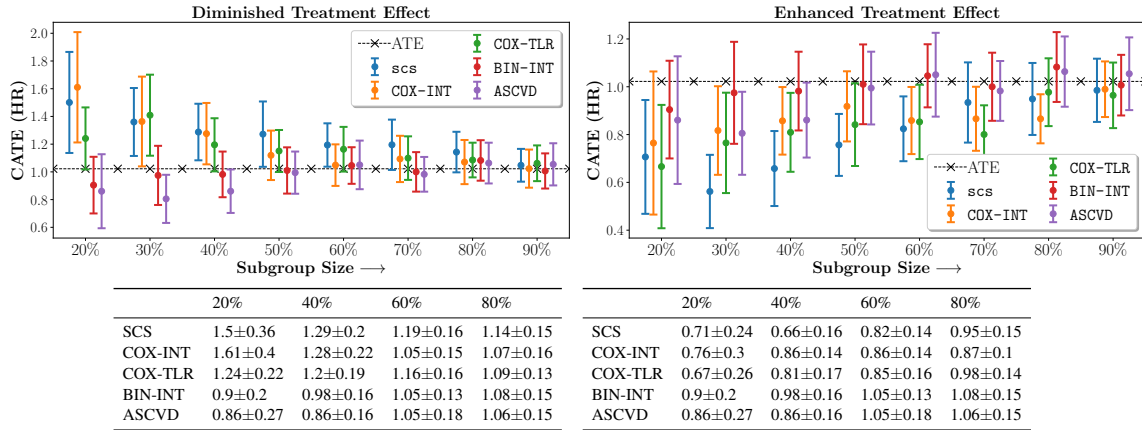
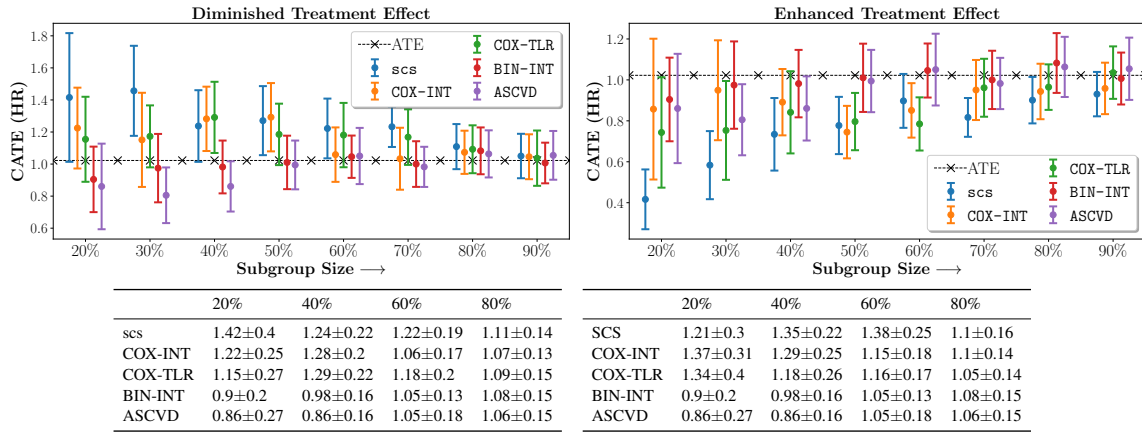


Figure 9: Conditional Average Treatment Effect in Restricted Mean Survival Time versus subgroup size for the latent phenogroups extracted from the ALLHAT study.

$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity

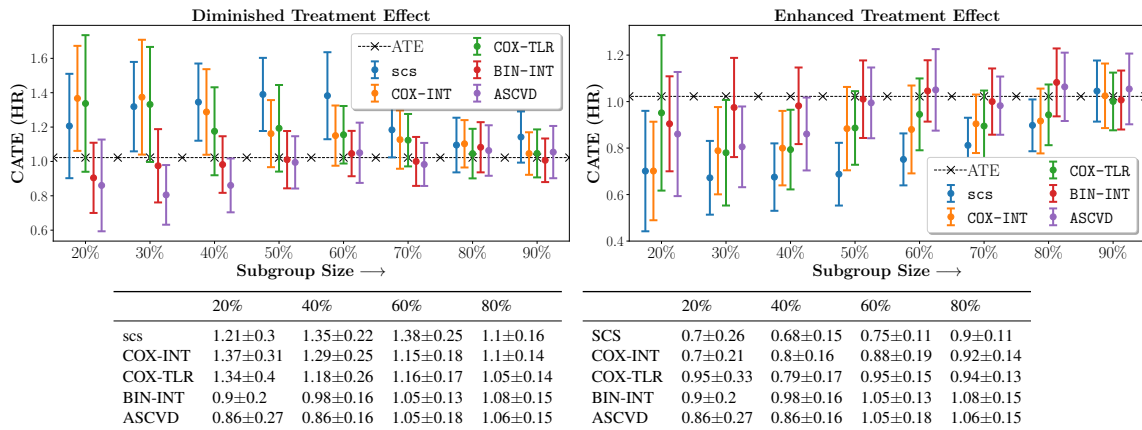
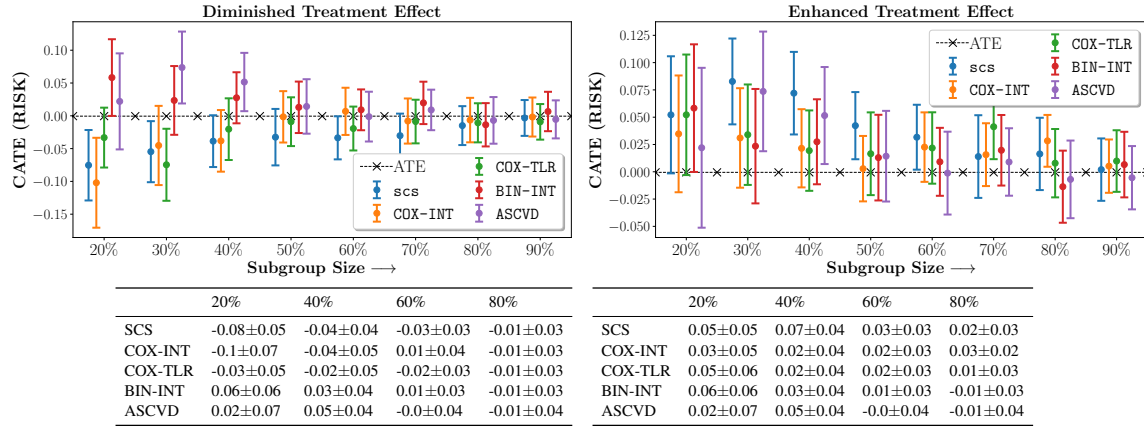


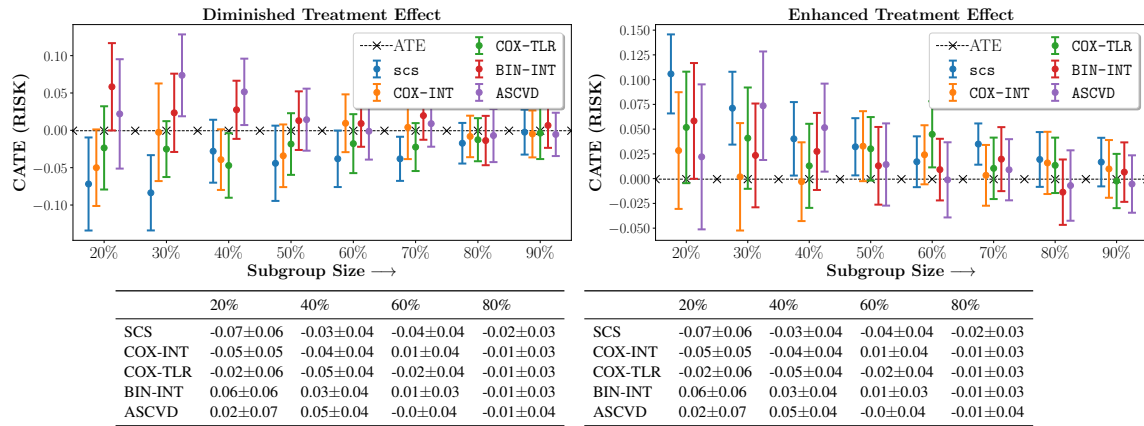
Figure 10: Conditional Average Treatment Effect in Hazard Ratio versus subgroup size for the latent phenogroups extracted from the **BARI 2D** study.

SPARSE COX SUBGROUPS

$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity

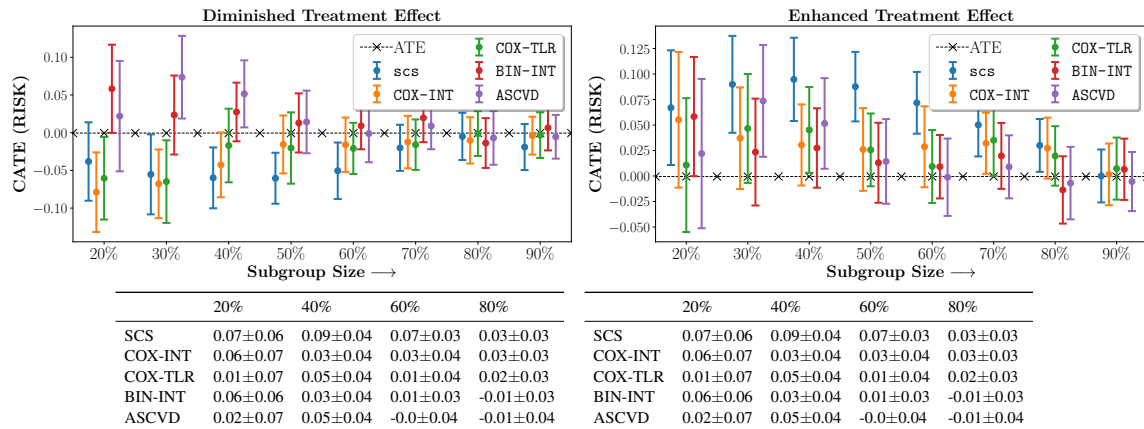




Figure 11: Conditional Average Treatment Effect in Risk versus subgroup size for the latent phenogroups extracted from the **BARI 2D** study.

$$\|\theta\|_0 \leq 5$$



$$\|\theta\|_0 \leq 10$$



No Sparsity




Figure 12: Conditional Average Treatment Effect in Restricted Mean Survival Time versus subgroup size for the latent phenogroups extracted from the **BARI 2D** study.