

A tutorial on individualized treatment effect prediction from randomized trials with a binary endpoint

Jeroen Hoogland¹ | Joanna IntHout² | Michail Belias² |
Maroeska M. Rovers² | Richard D. Riley³ | Frank E. Harrell Jr⁴ |
Karel G. M. Moons^{1,5} | Thomas P. A. Debray^{1,5} | Johannes B. Reitsma^{1,5}

¹Julius Center for Health Sciences and Primary Care, University Medical Center Utrecht, Utrecht University, Utrecht, the Netherlands

²Radboud Institute for Health Sciences (RIHS), Radboud University Medical Center, Nijmegen, the Netherlands

³School of Medicine, Keele University, Keele, UK

⁴Department of Biostatistics, Vanderbilt University School of Medicine, Nashville, Tennessee, USA

⁵Cochrane Netherlands, University Medical Center Utrecht, Utrecht University, Utrecht, the Netherlands

Correspondence

Jeroen Hoogland, Julius Center for Health Sciences and Primary Care, University Medical Center Utrecht, Utrecht University, Utrecht, the Netherlands.
Email: j.hoogland-2@umcutrecht.nl

Randomized trials typically estimate average relative treatment effects, but decisions on the benefit of a treatment are possibly better informed by more individualized predictions of the absolute treatment effect. In case of a binary outcome, these predictions of absolute individualized treatment effect require knowledge of the individual's risk without treatment and incorporation of a possibly differential treatment effect (ie, varying with patient characteristics). In this article, we lay out the causal structure of individualized treatment effect in terms of potential outcomes and describe the required assumptions that underlie a causal interpretation of its prediction. Subsequently, we describe regression models and model estimation techniques that can be used to move from average to more individualized treatment effect predictions. We focus mainly on logistic regression-based methods that are both well-known and naturally provide the required probabilistic estimates. We incorporate key components from both causal inference and prediction research to arrive at individualized treatment effect predictions. While the separate components are well known, their successful amalgamation is very much an ongoing field of research. We cut the problem down to its essentials in the setting of a randomized trial, discuss the importance of a clear definition of the estimand of interest, provide insight into the required assumptions, and give guidance with respect to modeling and estimation options. Simulated data illustrate the potential of different modeling options across scenarios that vary both average treatment effect and treatment effect heterogeneity. Two applied examples illustrate individualized treatment effect prediction in randomized trial data.

KEYWORDS

causal inference, personalized medicine, prediction, regression, treatment effect

1 | INTRODUCTION

Prediction of risk and prediction of treatment effect are two key components in modern medicine and personalized healthcare. On the one hand, risk predictions are classically functions of multiple patient characteristics. They include predictions of the risk of having a specific health outcome or condition (diagnosis) or of developing a future health outcome (prognosis). Also, risk predictions vary naturally across patients, are descriptive, and can be uniformly expressed as probabilities.¹ Importantly, risk prediction models are generally descriptive and are *not* intended to reflect the causal mechanism; in particular, included predictor effects in the model are not intended to reflect the extent to which their removal or modification would change an individual's prediction. On the other hand, predictions of treatment effect *do* express an expected difference due to modification of the treatment condition. They have classically been studied on a group level (eg, treated group vs control group), often assume a constant effect across individuals, have a causal interpretation, and are traditionally expressed using relative effect measures (eg, odds ratio, relative risk, or hazard ratio).²

Risk predictions and treatment effect estimation are two important areas of research but have largely developed in separation, leading to an apparent contradiction between methods for prediction and methods for causal inference. However, answers to many important questions need to bridge the divide. For instance, “How will a possible treatment change predicted outcome risk?” or “Is there variability in the effect of this treatment across patients (ie, differential treatment effect)?” These questions involve both the causal effect of treatment on a targeted health outcome and the adequate incorporation of associations with individual patient characteristics.

It is exactly these types of questions that need to be answered to provide more tailored, stratified, personalized, or precision medicine.³⁻⁵ The limitations of average relative treatment effects have long been recognized and the promise of a more individualized yet evidence-based approach has been enticing. Such a strategy requires focus on heterogeneity between patients and its relation to risk of the outcome of interest and variability in treatment effect. Also, moving toward more individualized estimates inherently means moving to more absolute expressions of variability in risk and treatment effect that are interpretable on the individual level.^{6,7} For example, predictions of the absolute risk of a future event under different treatment conditions provide a natural basis for shared decision-making.

In this tutorial, we aim to give a platform for statisticians and other researchers embarking on the prediction of individualized treatment effect. We focus on the risk of developing a binary outcome or endpoint, and aim to combine the highly conditional nature of typical risk prediction modeling with causal inference about treatment effectiveness. While the problem is well-known, it is very complex and therefore, we will cut it down to its essentials in the setting of a randomized trial, and mainly limit our scope to regression-based methods. We discuss the importance of a clear definition of the estimand of interest, provide insight into the required assumptions, and give guidance with respect to the modeling and estimation options. Key considerations with respect to the choice of modeling and estimation methods are further illustrated in a simulation study and two applied examples.

2 | DEFINING INDIVIDUALIZED TREATMENT EFFECT

The main idea underlying our endeavor is that the effect of treatment may be different for each individual, and that it may be beneficial to personalize or individualize its estimate. In the context of risk prediction, this implies that treatment causes a change in predicted outcome risk that may vary across individuals conditional on their characteristics. In other words, a personalized or individualized treatment effect describes the effect of modifying a treatment condition (ie, setting its value) while controlling for (ie, conditioning on) that individual's characteristics. We restrict our description to settings in which variables besides treatment do not have a causal interpretation, since this nicely aligns with the typical design of a randomized trial. This lack of causal interpretation for the set of variables conditioned on, is typical for classical prediction modeling. While the inner workings of a model that simultaneously describes both causal and mere associative relations may not need to discern between these different roles, they are of importance when interpreting the model.

To that effect, distinguishing between variables that do and do not have a causal interpretation is helpful for a precise definition of the individualized treatment effect of interest. Two common approaches to make this distinction are the $do(\cdot)$ operator introduced by Pearl et al⁸ and the potential outcomes framework popularized by Rubin.⁹ The $do(\cdot)$ operator is an operator that describes the effect of setting or modifying a variable to take a certain value (eg, $P(Y = y|do(X = x))$) and clearly separates this case from classical conditional notation (eg, $P(Y = y|X = x)$). The potential outcomes framework, as popularized by Rubin,^{9,10} allows for a formal distinction at the level of the outcomes that arises when the causal variable takes on different values. For instance, if interest is in the causal effect of a treatment, and treatment takes values $a \in \mathcal{A}$,

$Y^{A=a}$ denotes the potential outcome for treatment a . In case of a treatment variable that can be set to 0 (control) or 1 (treated), the two potential outcomes are $Y^{a=0}$ and $Y^{a=1}$. The notation easily allows for conditioning, such that the effect δ of treatment on the risk of an event, conditional on covariates \mathbf{X} , can be written as

$$\delta(\mathbf{x}) = P(Y^{a=1} = 1 | \mathbf{X} = \mathbf{x}) - P(Y^{a=0} = 1 | \mathbf{X} = \mathbf{x}), \quad (1)$$

where bold face indicates vectors. The same quantity could be written in $do(\cdot)$ notation as

$$\delta(\mathbf{x}) = P(Y = 1 | do(A = 1), \mathbf{X} = \mathbf{x}) - P(Y = 1 | do(A = 0), \mathbf{X} = \mathbf{x}). \quad (2)$$

For our purposes, the differences between these frameworks are not of interest and we adopt the potential outcomes framework throughout the remainder of the article for reasons of familiarity in statistical research.

A final remark on the nature of “individualized” or “personalized” is in place: the estimand of interest is not truly individual, but relates to groups of individuals sharing a covariate pattern. A truly individual treatment effect can never be observed since only one potential outcome can be observed at any time (or equivalently, only one treatment can be assigned). This problem has been referred to as the fundamental problem of causal inference.¹¹ Acknowledging this, for a dichotomous outcome Y , we define the individualized treatment effect ($\delta(\mathbf{x}_i)$) for individual i with covariate vector \mathbf{x}_i as

$$\delta(\mathbf{x}_i) = P(Y_i^{a=1} = 1 | \mathbf{X} = \mathbf{x}_i) - P(Y_i^{a=0} = 1 | \mathbf{X} = \mathbf{x}_i). \quad (3)$$

The individualized treatment effect $\delta(\mathbf{x}_i)$ can be interpreted as the expected difference in outcome risk for an individual with covariate values \mathbf{x}_i under two different treatment conditions. This definition is easily extended to ≥ 2 treatment conditions, but we will focus on a setting with 2 treatment conditions. While we will focus on $\delta(\mathbf{x}_i)$ as our estimand of interest, note that important information is lost when only looking at the difference between potential outcomes. Therefore, the main task is to predict the conditional risk of both potential outcomes, which provides a more complete picture and directly leads to an estimate of $\delta(\mathbf{x}_i)$. Section 3 first discusses the assumptions that allow estimation of $\delta(\mathbf{x}_i)$ from randomized trial data. Subsequently, Sections 4 and 5 describe specification and estimation of regression models for the purpose of predicting $\delta(\mathbf{x}_i)$.

3 | IDENTIFIABILITY ASSUMPTIONS

The individualized treatment effect specified in Equation (3) is written as a difference between two potential outcomes. However, in practice only a single potential outcome will be observed for each individual. Identification of $\delta(\mathbf{x}_i)$ based on the observed data requires assumptions to supplement the data. The necessary identifiability assumptions are consistency, exchangeability, and positivity. We here shortly introduce the fundamentals as relevant to our setting; excellent introductory¹² and comprehensive texts on causal inference are available elsewhere.¹³

Consistency refers to equality between the observed outcome Y and the potential outcome for the actually assigned treatment Y^a . When A takes on value 0 (control) or 1 (treated), this can be expressed as $Y = AY^{a=1} + (1 - A)Y^{a=0}$. This assumption holds when the data reflect well-defined treatments. As a counter example, consider a situation in which there is much variability in the active treatment (eg, different starting times, intensity or dosage, duration) but these are all just labeled $a = 1$: the causal contrast $Y^{a=1} - Y^{a=0}$ is no longer clearly defined. As is clear from Equation (3), the contrast of interest actually requires consistency conditional on covariates \mathbf{X} . This extension is trivial if marginal consistency holds.

Exchangeability requires that the potential outcomes are independent of treatment assignment ($Y^a \perp\!\!\!\perp A$ for all a). In other words, the actually assigned treatment does not predict the potential outcome.¹² As an example, consider a two-arm study of a new treatment: in terms of potential outcomes, exchangeability with respect to treatment here implies that it does not matter which arm received the new treatment. Since our interest is in a conditional treatment effect, exchangeability should hold conditionally ($Y^a \perp\!\!\!\perp A | \mathbf{X}$ for all a). While this is a challenging assumption to satisfy in general, it holds automatically in the context of a randomized trial. After either marginal randomization (ie, a common probability of treatment for all) or conditional randomization (ie, with the probability of treatment depending on covariates, also known as stratified randomization), conditional exchangeability holds when conditioning on (at least) the variables used during randomization. It is important to realize that randomization only provides exchangeability at baseline, and the causal contrast $\delta(\mathbf{x}_i)$ at that time therefore reflects an intention-to-treat effect. Any conditioning on post-randomization

information is no longer protected by randomization and the exchangeability assumption will no longer be guaranteed to hold.^{14,15} For instance, estimation of a per protocol individualized treatment effect would require further assumptions such as absence of any unmeasured confounders and correct specification of all confounders.¹⁶ We here limit our overview to intention-to-treat effects.

Positivity reflects the assumption that each patient should have a nonzero probability of either treatment assignment, which is clearly fulfilled in case of a randomized study.

A final assumption that is often made is the assumption of *no interference*, stating that the potential outcomes for one individual do not depend on treatment assignment of other individuals. While not strictly necessary, the situation quickly grows in complexity without this assumption since the potential outcome definitions would then have to incorporate the dependence on other units.^{9,16} The combination of consistency and no interference is also often referred to as the *stable unit treatment value assumption* (SUTVA).⁹

The definition of $\delta(\mathbf{x}_i)$ in Equation (3) assumes no interference, which follows from the fact that the potential outcome for individual i only depends on the individual's own covariate status and treatment assignment. Further assuming positivity provides a causal interpretation of the treatment effect conditional on covariates \mathbf{x}_i , with i from 1, ... , n . Finally, consistency and exchangeability are necessary to rewrite the estimand in terms of observed variables only:

$$\begin{aligned}\delta(\mathbf{x}_i) &= P(Y_i^{a=1} = 1 | \mathbf{X} = \mathbf{x}_i) - P(Y_i^{a=0} = 1 | \mathbf{X} = \mathbf{x}_i) \\ &= P(Y_i^{a=1} = 1 | A = 1, \mathbf{X} = \mathbf{x}_i) - P(Y_i^{a=0} = 1 | A = 0, \mathbf{X} = \mathbf{x}_i) \quad (\text{by exchangeability}) \\ &= P(Y_i = 1 | A = 1, \mathbf{X} = \mathbf{x}_i) - P(Y_i = 1 | A = 0, \mathbf{X} = \mathbf{x}_i) \quad (\text{by consistency}).\end{aligned}\tag{4}$$

In summary, the identifiability assumptions allow $\delta(\mathbf{x}_i)$ to be estimated from the observed data. While Equation (4) essentially allows for fully nonparametric estimation of $\delta(\mathbf{x}_i)$, there is usually insufficient data to do so when interest is in a highly conditional treatment effect (as is the case for an individualized treatment effect). That is, there will not be sufficient cases with $\mathbf{X} = \mathbf{x}_i$ under both treatments to reliably estimate $\delta(\mathbf{x}_i)$. This brings us to the need of a model for $P(Y_i = 1 | A = a_i, \mathbf{X} = \mathbf{x}_i)$ to smooth over the gaps in the observed set of all \mathbf{x}_i across both treatments.

4 | MODELS FOR THE PREDICTION OF INDIVIDUALIZED TREATMENT EFFECT

With the identifiability conditions in place for a causal interpretation of $\delta(\mathbf{x}_i)$ as estimated based on the observed trial data only, the remaining problem can be recognized as a typical prediction modeling problem (Equation 4). Therefore, well-established modeling techniques can be used to model the required conditional risks. While a vast array of possible prediction modeling techniques is available, we will focus on modeling techniques that have a basis in generalized linear modeling. More specifically, due to the binary outcome, we will focus on methods that have a basis in logistic regression, which has been the mainstay method for clinical prediction models in settings with a binary endpoint.¹ Key features of logistic regression include that it directly provides the probabilistic estimates of interest¹⁷ and has well-known properties. Also, it is a possibly parsimonious model family for the task at hand as explained in the next section.

4.1 | Homogeneous treatment effect

The aim is to model the observed outcomes \mathbf{Y} as a function of treatment assignment A and covariates \mathbf{X} to provide estimates for the right-hand side of Equation (4) and hence $\delta(\mathbf{x}_i)$. While the treatment effect of interest is expressed in terms of a difference in probabilities when the outcome is binary, this may not be the most appropriate scale to model treatment effect. The reason that the logistic model might provide a parsimonious model to predict the required conditional probabilities, is that a constant effect on the log odds scale has a valid interpretation across the entire range of predicted probabilities. For instance, the effect of treatment on outcome risk could really be constant on the log odds scale regardless of outcome risk in absence of treatment. This property does not hold for linear probability models or relative risk models, but very similar results may be obtained for probit models. As a quick reminder, Figure 1 shows the logistic link that transforms the log odds or linear predictor scale to the probability scale. A constant treatment effect on the log odds scale has a large effect on the probability scale when the linear predictor equals zero and approaches zero for very low and high linear predictor values. This nicely reflects the difference in the amount of wiggle room when the control outcome

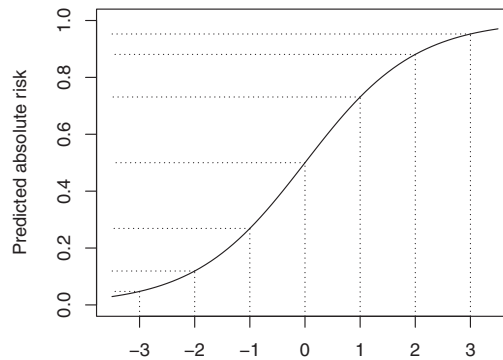


FIGURE 1 The translations of effects from the linear predictor (LP) scale to the absolute scale vary depending on location as shown for the logistic link function ($\frac{1}{1+e^{-LP}}$). Therefore, a constant or homogeneous relative effect, here shown as a 1-point difference on the linear predictor (ie, log odds) scale for ease of exposition, has different implications on the absolute risk scale. For example, for a patient with a control outcome risk of 50% (LP = 0), a treatment effect of -1 on the log odds scale reduces predicted absolute risk to 27% (LP = -1, resulting in an absolute risk reduction of 23%). For a patient with a control outcome risk of 27% (LP = -1), the same treatment effect on the log odds scale leads to a predicted risk of 12% (LP = -2, resulting in an absolute risk reduction of 15%)

risk reflects unpredictability vs near certainty, respectively. The reason to emphasize these well-known properties is to highlight that a very simple model on the log odds scale may very well lead to potentially relevant differences between individuals on the level of $\delta(\mathbf{x}_i)$ (ie, differences in absolute risk).

The simplest way to create such a model is to assume absence of any interaction between treatment and the other covariates. Thus, the (conditional) treatment effect is assumed to be *homogeneous* or constant on the log odds scale. In terms of notation, let Y_i denote the independent dichotomous *observed* outcomes, assumed to have patient specific mean μ_i (ie, $Y_i \stackrel{i.i.d.}{\sim} \text{Bernoulli}(\mu_i)$). Also, with vector \mathbf{x}_i denoting the p individual characteristics and a_i being the treatment indicator as before, this leads to the following simple logistic regression model

$$\text{logit}(P(Y_i = 1|A = a_i, \mathbf{X} = \mathbf{x}_i)) = \beta_0 + \beta_1 a_i + \boldsymbol{\beta}^\top \mathbf{x}_i \tag{5}$$

with regression parameters β_0 , β_1 , and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$. The key assumptions for this logistic model including a conditional *homogeneous* treatment effect β_1 are (i) appropriateness of the logistic link function, (ii) linearity in the parameters, and (iii) additivity of at least the treatment effect on the log odds scale (ie, there are no treatment-covariate interactions on the log odds scale). The linearity assumption on the covariate contributions could easily be relaxed, allowing for global or local transformations of \mathbf{x}_i such as polynomials and splines.

The predicted individualized treatment effect follows directly after estimation of the model parameters*:

$$\hat{\delta}(\mathbf{x}_i) = \frac{1}{1 + e^{-(\hat{\eta}_i + \hat{\beta}_1)}} - \frac{1}{1 + e^{-\hat{\eta}_i}}, \tag{6}$$

where $\hat{\eta}_i = \hat{\beta}_0 + \hat{\boldsymbol{\beta}}^\top \mathbf{x}_i$. Strict additivity of the treatment effect on the log odds scale may provide a parsimonious model for the analysis and is easily translated to the more interpretable scale of $\delta(\mathbf{x}_i)$ where this additivity does not hold. Note that in case of such a homogeneous treatment effect on the log odds scale, variability in \mathbf{x}_i is the driving force behind any differentiation on the level of $\hat{\delta}(\mathbf{x}_i)$. This variability in \mathbf{x}_i corresponds to variability in prognosis across individuals under control treatment (ie, variability in $Y_i^{a=0}$).

4.2 | Heterogeneous or differential treatment effect

As an extension of homogeneous treatment effect models, the relative treatment effect can also be allowed to depend on the other covariates. We will refer to such nonadditivity of the treatment effect on the relative scale as heterogeneity

*For our current goal and scope, inferring individualized treatment effect from predictions under the relevant potential outcomes is the end-point. However, it is interesting to note that this type of conditional potential outcome predictions can also serve as input for substitution estimators aiming for more marginal estimates, such as the parametric g-formula¹³ and estimators of marginal risk difference and marginal risk ratio based on logistic models.¹⁸

of treatment effect (HTE) or differential treatment effect. We note that there is no single accepted definition of the term HTE in the literature, and that it is sometimes used in a broader sense to also include the variability in $\delta(\mathbf{x}_i)$ that may result from a homogeneous treatment effect model.⁴ We use HTE in its narrow sense (ie, restricted to nonadditive relative effects), since this allows one to distinguish between possible variability in the way treatment affects different individuals (homogeneous vs heterogeneous) and variability among individuals that does not relate to treatment effect (ie, variability in expected prognosis under control treatment $P(Y_i = 1|A = 0, \mathbf{X} = \mathbf{x}_i)$).

The homogeneous treatment effect model in Equation (5) can easily be extended to allow for HTE by inclusion of treatment-covariate interactions. The model then becomes

$$\text{logit}(P(Y_i = 1|A = a_i, \mathbf{X} = \mathbf{x}_i)) = \beta_0 + \beta_1 a_i + \boldsymbol{\beta}_m^\top \mathbf{x}_i + \boldsymbol{\beta}_z^\top \mathbf{z}_i a_i, \quad (7)$$

where \mathbf{z}_i is a subset of \mathbf{x}_i , $\boldsymbol{\beta}_m$ includes the coefficients for the main effects of \mathbf{x}_i , and $\boldsymbol{\beta}_z$ includes the coefficients for treatment-covariate interactions. As before, the space of the measured \mathbf{x}_i can be expanded using global or local transformations. These transformations need not be the same in \mathbf{z}_i : the functional form of the effect of a covariate may depend on treatment status. Also, note that when \mathbf{z}_i equals \mathbf{x}_i (ie, all covariates are involved in treatment-covariate interactions), an exactly equivalent parameterization can be obtained by specification of separate models for the treated group and the control group with just an intercept and main covariate effects, which separates the models for both potential outcomes. Additional details are provided in the online supporting material (Part A).

Based on the model in Equation (7), the predicted individualized treatment effect again follows easily from the parameter estimates. In analogy to Equation (6), the prediction of $\delta(\mathbf{x}_i)$ can be derived according to

$$\hat{\delta}(\mathbf{x}_i) = \frac{1}{1 + e^{-(\hat{\eta}_i + \hat{\beta}_1 + \hat{\boldsymbol{\beta}}_m^\top \mathbf{x}_i)}} - \frac{1}{1 + e^{-\hat{\eta}_i}}. \quad (8)$$

While this prediction comes nice and easy in theory, the challenge lies in precise model specification and the estimation of the model parameters of these relatively complex models.

5 | MODEL ESTIMATION

The preferred method for estimation of the prediction models of interest depends on the relation between model complexity and the amount of signal in the data. In medical statistics, the amount of variability in the outcome of interest that can be explained is often low to moderate, which is the main reason for the need for large sample sizes. In case of insufficient sample size, models are prone to overfitting, which describes the situation where the model captures part of the noise in the data. Overfitting is an important concern since it limits generalizability of the model. In this section, we discuss the need for methods that mitigate the susceptibility to overfitting.

5.1 | Maximum likelihood

Estimates for the proposed logistic regression models can be derived with standard maximum likelihood estimation for generalized linear models.^{19,20} However, even after taking all content knowledge into account, models are often still overly complex with respect to the available sample size. This may already hold for a logistic model of homogeneous treatment effect when the number of covariates is large, or their functional form allowed to be complex and the sample size is relatively small.²¹ It has long been recognized that standard maximum likelihood estimation of logistic models is problematic in these settings due to finite sample bias, perfect separation, collinearity, and overfitting.^{17,22} Models including many nonadditive effects such as treatment-covariates interactions are even more prone and necessitate either strong prior assumptions or restrictions during model estimation.

5.2 | Penalized maximum likelihood

Regression based on penalized maximum likelihood (pML) estimation has emerged as a method that can, at least to some degree, cope with relatively complex models.²³⁻²⁵ These methods penalize the log-likelihood for the magnitude of

regression coefficients other than the intercept. This penalty introduces bias toward zero on the estimated (non-intercept) coefficients, and thus toward the overall outcome incidence for predictions. In other words, it introduces a bias that reduces variability at the level of the predictions. This balance is also known as the bias-variance trade-off. Well-known pML methods include ridge regression, lasso regression, and the elastic net which includes the first two as special cases.²⁶ The ridge penalty is a smooth penalty on squared size of the regression parameters and leads to shrinkage of the estimated coefficients. It was originally developed to deal with collinearity and tends to distribute the weight among collinear variables.²³ The lasso penalty is on the absolute value of the coefficients and leads to both shrinkage and selection.²⁴ It has a tendency to select among collinear variables. The elastic net penalty is a weighted balance of the ridge and lasso penalty. The required degree of penalization is a tuning parameter that needs to be estimated from the data, which is most commonly done by means of cross-validation. Importantly, this estimation involves uncertainty that is most problematic when accurate penalization is needed the most (ie, small data sets and/or low signal relative to noise).^{27,28} With respect to the equivalence of a model including all treatment-covariate interactions on the one hand and separate modeling in the treated and control group on the other hand (Section 4.2), note that this equivalence no longer holds in case of pML. Details are provided in the online supporting material (Part A).

5.2.1 | Shrinkage and/or selection

Penalization in clinical prediction modeling

In some settings, the underlying process to be modeled is fairly well known, and therefore, the same holds for the elements that should be included in a model. While such a setting does not require selection of parameters, shrinkage may still be beneficial in terms of prediction accuracy when sample size is limited with respect to model complexity. In contrast, the available data may be very rich while the underlying process not well understood, but thought to be sparse. Penalization approaches that provide selection (ie, sparse solutions) have been successfully applied to prediction problems with many covariates (possibly more than cases, ie, $p \gg n$) where selection of variables is key and there is insufficient content knowledge to do so, such as the selection of possibly important signals from microarray data.²⁹

In the typical context of clinical prediction modeling, the properties of the problem are somewhere in-between these two extremes. In the best-case scenario, clear prespecification of a model might be possible. Often however, even though the data are typically low-dimensional, some further variable selection may still be required. The choice of penalty, and hence the need for selection, therefore heavily depends on the state of content knowledge and the amount of data available. An issue that may guide the selection of a penalty function is the need for an honest representation of the relative weights of model parameters. Ridge regression tends to share the regression weights between parameters that are correlated with the outcome *and* with each other,³⁰ while lasso tends to select among such variables.^{29(chap4)} As an example, if two highly correlated covariates are equally predictive of the outcome, ridge will keep both in the model with approximately equal weight. Lasso will remove the one variable that happens to have a slightly weaker association with the outcome in the current sample. Both representations can be useful, but they serve a different purpose.

Penalization and modeling of heterogeneous treatment effect

Models of heterogeneous treatment effect include treatment-covariate interactions. Such interactions are always harder to estimate than the overall (main) treatment effect. In terms of selection, heterogeneous treatment effect models encounter the variable selection issue twice: for the main effects and for the treatment-covariate interactions. Usually, lower order (main) effects are kept in the model for each component of an interaction. However, both lasso regression and elastic net regression do not respect this hierarchical nature. To that effect, a hierarchical group lasso (HGL) algorithm has been developed that does respect the hierarchy between main effects and interaction effects.³¹ In short, variable selection is achieved on a group level that is allowed to be hierarchical, such that main effect groups will be in the model when they are part of any interaction. For problems with non-overlapping groups, regular group lasso algorithms are also widely available.³²

6 | MODEL COMPLEXITY

The current state of subject matter or content knowledge, the type of process under study, the strength of the associations in the data, and the final purpose of the model, all weigh into the decision on the best balance between prior model

specification and more data-driven modeling methods. We have described both specification and estimation of logistic models that can be used to predict treatment effect. In this section, we will shortly touch upon the complementary roles of content knowledge and penalization, the concern of overfitting when aiming for out-of-sample prediction, and the degree to which model complexity can be left to the data-drive methods.

6.1 | Content knowledge and penalization

Content knowledge, general statistical knowledge, and data-driven methods such as penalization should work synergistically to arrive at the most parsimonious model that reflects current content knowledge as updated by the data. Ideally, content knowledge includes knowledge on nonlinear covariate contributions, interactions among covariates, and especially knowledge on covariates that might interact with treatment. Unfortunately, such knowledge is often very limited. Consequently, the number of parameters that one wants to estimate may still be relatively large even after all content knowledge is exhausted. While data driven approaches can help to nudge such models in the right direction, they provide no panacea or substitution of content knowledge.

6.2 | Overfitting and prediction accuracy

For our aim to predict individualized treatment effect, the conditional model of treatment effect is intended to generalize and therefore overfitting is a key concern. The primary challenge is to balance model complexity with respect to the amount of available data in a way that generalizes beyond the original sample. The preferred way to do so is to limit model complexity based on content knowledge, supplemented by use of penalization. The relation between prediction accuracy, model complexity, effective sample size, and the strength of the associations in the data is well known and has been described in the context of risk prediction (eg, Riley et al²¹ and van Smeden et al²²). However, estimation of $\delta(\mathbf{x}_i)$ requires the difference between two outcome risk predictions (under the two to be compared treatments) to be accurate. While these two predictions will be highly correlated since they arise from the same individual, the expected error will invariably be larger than for a single prediction. To our knowledge, there is no guidance available on the necessary conditions with respect to effective sample size and expected explained variation for accurate prediction of risk differences within individuals. Our simulation study below provides some first insights.

6.3 | “Risk modeling” vs “effect modeling”

The use of treatment-covariate interactions to model heterogeneous treatment effect (eg, as in Equation 7) has also been referred to as “effect modeling” and has been distinguished from “risk modeling”.^{4,5,33-35} In risk modeling, treatment effect variability is evaluated as a function of outcome risk, where outcome risk is a function of the covariates. Therefore, it can be seen as a data reduction method. For example, in risk modeling the linear predictor scores $\hat{\eta}_i$ resulting from a model for $P(Y_i = 1|A = 0, \mathbf{X} = \mathbf{x}_i)$ can be interacted with treatment instead of the full multi-dimensional representation of \mathbf{x}_i . That is,

$$\text{logit}(P(Y_i = 1|A = a_i, \hat{\eta}_i)) = \beta_t a_i + \hat{\eta}_i + f(\hat{\eta}_i) a_i, \quad (9)$$

where $f(\cdot)$ denotes a possibly flexible function and $\hat{\eta}_i$ is used as an offset. The general idea is that $f(\hat{\eta}_i)$ is a much simpler structure to estimate than many individual treatment-covariate interactions. Therefore, it is supposed to fill a gap in situations where content knowledge is insufficient to limit model complexity to something that can be reliably estimated in the data.

From a statistical point of view, “risk modeling” does of course reduce the risk of overfitting, since it restricts the modeled treatment effect heterogeneity to be a function of a scalar. However, the price to pay is that HTE is thereby forced to be proportional to the main effects of \mathbf{x}_i . This implies (i) that all covariates that have a main effect also modify the relative treatment effect, and (ii) that the effect of each element of \mathbf{x}_i on HTE has the same direction as its effect on outcome risk. These are strong assumptions that have no clear biological substrate and no clear statistical preference over

other data reduction methods such as principal component analysis.[†] Nonetheless, the idea behind risk modeling does reflect recognition of the danger of overfitting when modeling many treatment-covariate interactions, which remains an important issue when modeling heterogeneous treatment effect.

6.4 | Tree-based methods

Many different modeling techniques are available under the machine learning umbrella. While we have limited our scope to regression-based methods, there are other methods that could be used instead. In particular, several tree-based methods have been developed for the specific purpose of individualized treatment effect prediction. While an in-depth discussion is beyond our scope, we here provide several key references. Wager and Athey³⁶ extend the well-known random forest algorithm by Breiman³⁷ to enable individualized treatment effect prediction and provide a very thorough overview of the required conditions for causal inference, including asymptotic theory. Also, they provide an overview of the literature on forest-based algorithms for estimating heterogeneous treatment effect. Lu et al³⁸ provide a clear exposition of individualized treatment effect prediction and the potential outcome framework, and provide empirical and simulation results on a wide variety of random forest methods for causal inference, including virtual twins, counterfactual random forests, the aforementioned causal forests and Bayesian adaptive regression trees. Generalized random forests³⁹ constitute a more recent addition to the literature and form a much broader method that can also be used for causal individualized treatment effect estimation (and effectively encompass the work by Wager and Athey³⁶). Lastly, model-based recursive partitioning is a somewhat different tree-based approach that incorporates parametric models into the tree.⁴⁰ Such a parametric model can for instance describe control outcome risk and relative treatment effect. Node-splitting then occurs on the variable that generates most instability in the parameters for this model. Seibold et al⁴¹ have developed a model-based recursive partitioning random forest to identify treatment effect heterogeneity. Model-based recursive partitioning is most closely related to the methods discussed in this tutorial since it can differentiate between heterogeneity of treatment effect on the relative scale and differential outcome risk that may be related to a homogeneous treatment effect (ie, constant log odds of treatment).

Beyond the predictions of individualized treatment effect, there have been efforts to find subgroups that might need different treatment based on a fitted causal forest,⁴² and efforts to further explain random forest-based treatment effect predictions based on covariate data.³⁸ Both papers go through considerable lengths to disentangle and interpret predicted individualized treatment effect differences. While this may lead to interesting hypotheses, this should be a careful undertaking. An illustration of the things that could go wrong is available in Rigdon et al,⁴³ showing high false discovery rates if such care is not taken.

To the best of our knowledge, direct comparisons of regression-based and other methods for the prediction of individualized treatment effect have not been performed yet. In general, it can be expected that more flexible models are more prone to overfitting and require more data.⁴⁴ A challenge in the comparison of different methods in simulation studies is that specific data generating mechanisms favor specific methods. For instance, a data generating mechanism with linear and additive effects will favor regression methods; a mechanism generating subgroups based on cut-offs will favor tree-based algorithms. An interesting method to acknowledge these effects and check robustness is provided by Austin et al,⁴⁵ and could also be applied in the context of individualized treatment effect prediction in future work comparing a wider set of methods.

7 | LEARNING FROM SIMULATIONS

We conducted a simulation study to illustrate the consequences of the choice of model and estimation method when predicting individualized treatment effect. Such simulations are especially helpful when predicting potential outcomes,

[†]The preference to model HTE as a function of outcome risk was originally motivated as “outcome risk is a mathematical determinant of treatment effect”,⁴ along with a reference to the fact that the [marginal] odds ratio [of exposure to treatment] equals $EER/(1 - EER) \div CER/(1 - CER)$ [where EER is the experimental outcome prevalence and CER is the control outcome prevalence]. However, the same equation can be used to show that independence is a possibility: (i) only the $CER/(1 - CER)$ part depends on control outcome risk and is positive and finite for any control prevalence in the (0, 1) interval, and (ii) combined with any valid odds ratio, $EER/(1 - EER)$ is also positive and finite and therefore maps back to a prevalence among the treated in the (0, 1) interval. This also holds for the conditional treatment effect and was in fact one of the reasons to prefer logistic regression as explained in Section 4.1.

since they can never be observed directly in practice. Also, the ability to manipulate the data generating mechanism into several interesting settings has a clear illustrative advantage. In the design of our simulation study, we adhered to the general guidelines proposed by Burton et al.⁴⁶ and Morris et al.⁴⁷

7.1 | Data generating mechanisms

The data generating mechanism was parametric and was based on a logistic model (Equation 7). Settings varied across the full factorial combination of varying total sample size (400, 1200, 3600), presence/absence of a main treatment effect ($\beta_t = \ln(0.6)$ or $\beta_t = \ln(1)$), and presence/absence of heterogeneity of treatment effect. The treatment indicators a_i were independent samples from a Bernoulli distribution with probability 0.5. Twelve covariates were drawn from a multivariate standard normal distribution with a compound symmetric covariance matrix ($\rho = 0.1$). Main effect coefficients β_m were of exponentially decreasing size ($\beta_{m,1}, \dots, \beta_{m,12} = 2^{-\frac{0}{2}}, 2^{-\frac{1}{2}}, 2^{-\frac{2}{2}}, \dots, 2^{-\frac{11}{2}}$) to reflect (i) decreasing added value of consecutive variables, and (ii) that it is unlikely for variables that are included in a risk model to have truly zero coefficients. In settings with a homogeneous treatment effect, there were no treatment-covariate interactions (ie, $\beta_z = \mathbf{0}$). In settings with a heterogeneous treatment effect, β_z was equal to $-\frac{1}{2}, -\frac{1}{4},$ and $-\frac{1}{8}$ for $\beta_{z,10}, \beta_{z,11},$ and $\beta_{z,12},$ respectively, and included a small random perturbation that was generated once for all simulations ($\leq |0.05|$) for $\beta_{z,1}, \dots, \beta_{z,9}$. In each simulation setting, the intercept was chosen such that the true underlying outcome prevalence in the control arm was 25% (details provided in the online supporting material, Part B). Nagelkerke R^2 , as measured between the true conditional probabilities for the assigned treatment (ie, $P(Y_i^{A=a_i} = 1 | \mathbf{X} = \mathbf{x}_i)$) and the observed events Y_i in a large sample, was approximately 0.4 for all settings.

To get more insight into the different simulation settings, the distribution of $\delta(\mathbf{x}_i)$ for each of the data generating mechanisms is shown in Figure 2. Note that (i) there is no variability at all in the upper right figure (due to absence of any treatment effect); (ii) variability in the upper left figure is due to variability in control risk across patients, and (iii) variability in the lower two figures is due to heterogeneity in both control risk and treatment effect.

7.2 | Model development

Within each simulation run, both a development and a validation data set were simulated for each of the simulation settings. That is, both data sets were always generated according to the same data generating mechanism. The size of the model development sets matched the simulation settings (ie, 400, 1200, or 3600), and the size of the validation sets was always equal to 10 000 observations.

Table 1 provides an overview of the evaluated methods. The overall absolute treatment effect is just the marginal version of $\delta(\mathbf{x}_i)$ as marginalized over all covariates. Its estimate of $\delta(\mathbf{x}_i)$ is just the $\hat{\delta}$, the difference in mean outcome incidence between treatment arms. For the homogeneous treatment effect models, all covariates entered the model only as main effects. In case of heterogeneous treatment effect models, all covariates entered the model as both main effects and interactions with treatment. The selected penalty parameter for ridge, lasso, and HGL was the one with the smallest deviance in 10-fold cross-validation. Note that both lasso and HGL may set coefficients to exactly zero (ie, perform selection). Also, while HGL can search the entire interaction space, the current implementation was limited to treatment-covariate interactions in parallel to the other methods. A final variation of the HTE methods was a “content knowledge” (CK) setting, where we assumed that content knowledge suggests that only the first 8 variables have important main effects, and only covariates 9 to 12 are likely treatment interaction candidates. Ridge regression was used to estimate such a CK-based model. The “risk modeling” implementation included a risk model estimated in the control group based on main effects for all covariates and a linear treatment with risk-score interaction. Both standard maximum likelihood and ridge regression were performed for the risk model. A significance-based approach was implemented as a final comparison. Starting from a homogeneous treatment effect model including all covariates, a likelihood ratio test was performed for the treatment effect coefficient. When nonsignificant, the treatment coefficient was removed from the model (leaving only main covariate effects). When significant, all treatment-covariate interactions were added to the model and a second likelihood ratio test was performed to evaluate their joint significance. Treatment-covariate interactions were kept in the model when this joint test was significant and removed otherwise. All tests used an α level of 0.05. In addition to the methods in Table 1, HTE was modeled using separate prediction models per treatment arm, and thus per potential outcome, as described in the online supporting material (Part A).

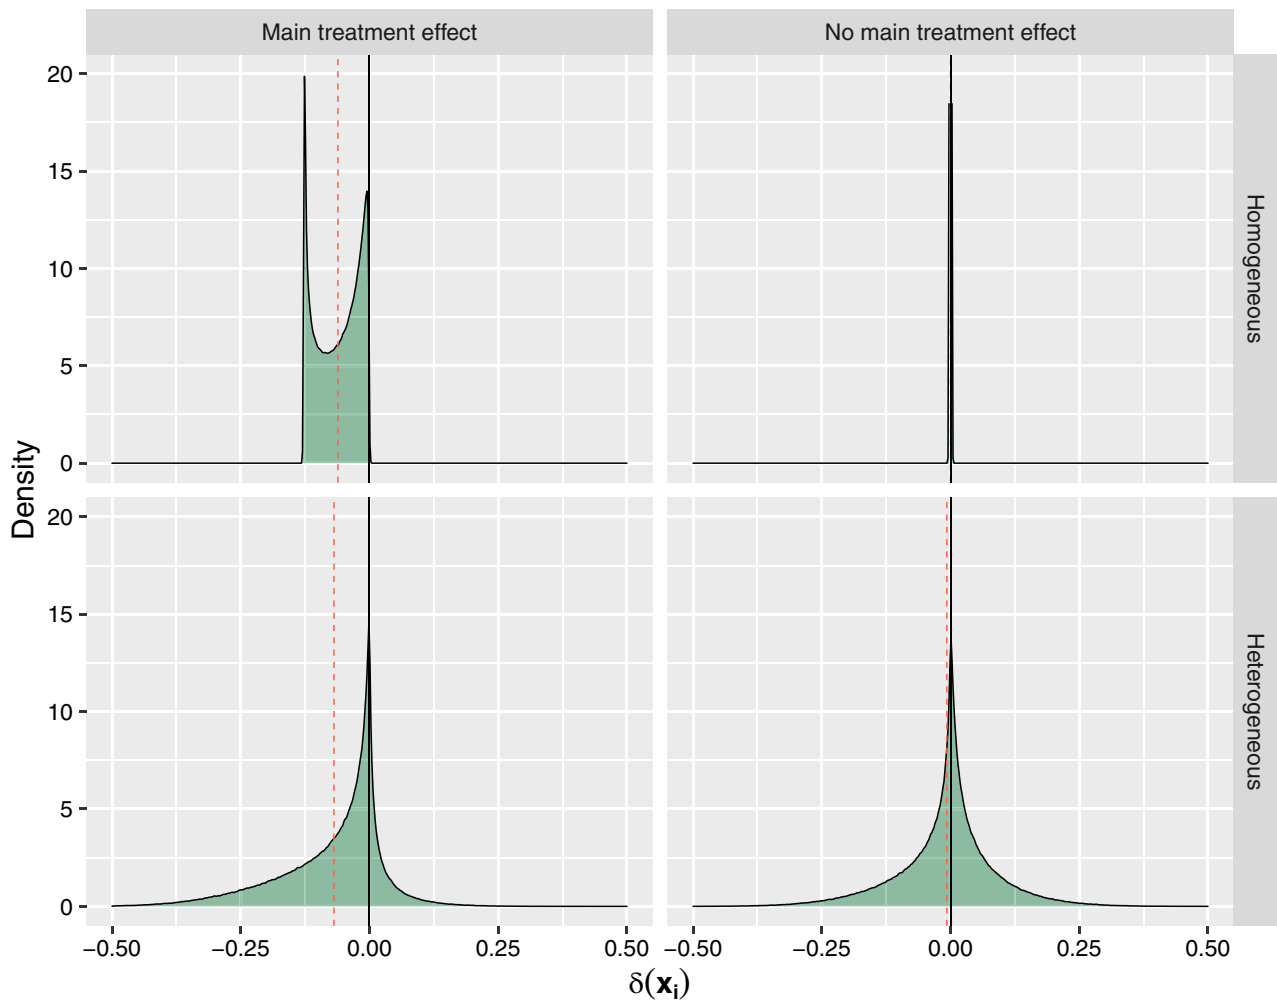


FIGURE 2 Distribution of $\delta(\mathbf{x}_i)$ according to the data generating mechanism for each of the simulation settings. The quadrants correspond to settings with a homogeneous main treatment effect (upper left), absence of any treatment effect (upper right), heterogeneous treatment effect in presence of a main treatment effect (lower left), and heterogeneous treatment effect in absence of a main treatment effect (lower right). The red dotted lines provide the mean of $\delta(\mathbf{x}_i)$ per setting. Note that all of the mass in the upper right figure is on a spike at $\delta(\mathbf{x}_i) = 0$ [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 1 Implemented methods toward the prediction of individualized treatment effect ($\hat{\delta}(\mathbf{x}_i)$)

Regression model	Equations	Estimation method
Overall absolute treatment effect (overall)	–	ML
Homogeneous treatment effect (HOM)	(5) and (6)	ML, ridge
Heterogeneous treatment effect (HTE, HTE-CK) ^a	(7) and (8)	ML, ridge ^b , lasso, HGL
“Risk modeling” (RM)	(9)	ML, ridge
Significance-based (SB)	–	ML

Abbreviations: HGL, hierarchical group lasso; ML, maximum likelihood.

^aModel specification differs between the default case (HTE) modeling all main effects and treatment-covariate interactions, and the content knowledge case (HTE-CK) modeling a selection of main effects and an treatment-covariate effects (details are described in Section 7.2).

^bOnly ridge was used for the HTE-CK model.

7.3 | Model evaluation

Each of the methods provides a prediction vector $\hat{\delta}$ with elements $\hat{\delta}_i$ (short for $\hat{\delta}(x_i)$) as derived in the validation sample. These predictions can be compared to the known δ based on the data generating mechanism. The root mean squared prediction error (rMSPE) between the elements of these two vectors was used to quantify the prediction errors according to

$$\text{rMSPE} = \sqrt{n^{-1}(\hat{\delta} - \delta)^\top (\hat{\delta} - \delta)}. \quad (10)$$

The root was taken to arrive at an expression of the error on the risk difference scale. In addition, the 0.9-quantile of absolute prediction errors was derived for both $\hat{\delta}$ and predicted risk. Supplementing these single figure summaries, calibration plots were derived for both $\hat{\delta}$ and predicted risk, where predictions were cut into 20 equal-size quantiles groups and compared to the true values based on the data-generating mechanism.

7.4 | Statistical software

All analyses were performed in R statistical software version 3.5.0.⁴⁸ The R script for exact replication of the simulation study is available as online supplementary material. Logistic regression models based on maximum likelihood estimation were fitted using `glm()`. Ridge and lasso implementations were based on the `glmnet` package.²⁶ HGL was implemented using the `glinetnet` package.³¹

7.5 | Simulation study results

We here synthesize the results of 250 simulation runs in terms of rMSPE and calibration of the predicted individualized treatment effect.

7.5.1 | Average root mean squared prediction error

The main simulation results with respect to rMSPE are shown in Figure 3, which provides a summary of the error that can be expected in the long run across settings and methods. Five key observations can be made across all settings. First, conditioning on main effects of the available covariates as in a homogeneous treatment effect model was always beneficial when compared to the fully marginal $\hat{\delta}$. This confirms the idea that a conditional estimand ($\delta(x_i)$) requires a conditional estimator. Second, HTE model accuracy was especially sensitive to sample size, which is in line with expectation due to the amount of parameters that needs to be estimated. Third, penalization is key and improved estimation of both homogeneous and heterogeneous treatment effect models up to large sample sizes. Fourth, penalization does *not* remove the risk of overfitting for complex models in small sample size settings. Heterogeneous treatment effect models could not be reliably fitted in small samples regardless of the estimation method. Fifth, utilization of content knowledge was the most effective way to reduce HTE model complexity. Lasso and HGL did not catch up, even though the content knowledge simulated here was not entirely correct and the lasso models did start from the correct set of variables given the data generating mechanism. The reduction in the potential set of treatment-covariate interaction variables in the content knowledge-based model was very effective, even though this missed out on small interaction effects. Nonetheless, apparent content knowledge could of course not help out when it was wrong, such as in the complete null setting (no main treatment effect, no heterogeneity). Lastly, risk modeling was not the preferred method in any of the simulation settings. This of course depends on the data generating mechanism where HTE was not fully explained by a risk model, but this would also not be expected in practice and serves the purpose of showing that risk modeling may miss important treatment effect heterogeneity. Additional rMSPE results for comparing treatment-covariate interaction modeling with separate prediction models per treatment arm, and thus per potential outcome, are described in the online supporting material (Part A). In short, treatment-covariate interaction modeling by means of lasso regression was best across all settings when compared to per arm modeling. The differences between treatment-covariate interaction modeling and per arm modeling were more nuanced for ridge regression.

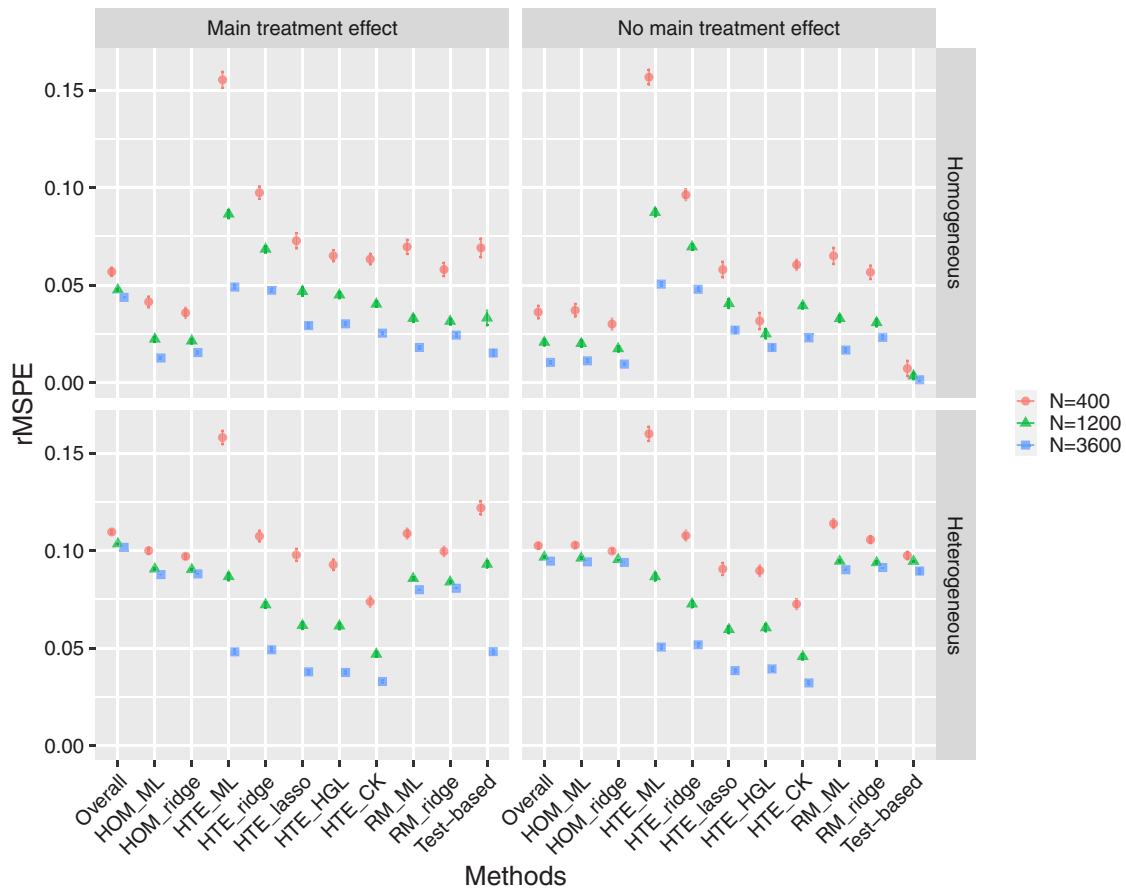


FIGURE 3 Simulation study results: average root mean squared prediction error (rMSPE) of the predicted treatment effects (over 250 simulations) with ± 2 SE error bars for all simulation settings. The quadrants correspond to settings with a homogeneous main treatment effect (upper left), absence of any treatment effect (upper right), heterogeneous treatment effect in presence of a main treatment effect (lower left), and heterogeneous treatment effect in absence of a main treatment effect (lower right). Note that the standard errors are often so small that they are obscured by the mean estimates. Abbreviations: homogeneous treatment effect models (HOM), heterogeneous treatment effect models with treatment-covariate interactions (HTE), and heterogeneous treatment effect models based on risk modeling (RM), as estimated by means of maximum likelihood (ML), ridge or lasso regression, or hierarchical group lasso (HGL), with CK standing for content knowledge [Colour figure can be viewed at wileyonlinelibrary.com]

7.5.2 | Calibration

The online supplementary material (Part C) provides calibration plots of $\hat{\delta}(x_i)$ across methods and settings. It provides further insight into the distribution of the errors and the degree of variability across replications. Several important observations can be made that supplement the conclusions based on rMSPE. First, calibration curves at least pass the (0, 0) point, and the size of the errors increases as predictions move away from zero. Therefore, prediction errors were much smaller around the harm-benefit boundary (ie, $\delta(x_i) = 0$). Second, calibration in individual small samples could be far off even if average performance across simulations was good. For instance, fitting a ridge regression homogeneous treatment effect model in accordance with the data generating mechanism could still lead to substantial overfitting or underfitting in any individual data set. These findings on the risk difference scale are in line with earlier results on direct prediction.^{27,28} Third, in small and even medium sample sizes, even penalized heterogeneous treatment effect models overfitted to such an extent that they falsely predicted harm for a subpopulation when in fact there was none (as in the homogeneous treatment effect settings). Fourth, when the data generating mechanism was heterogeneous, calibration of predicted treatment effect was quite reasonable for penalized HTE models in medium and large sample sizes. Note that, while useful for illustrative purposes in this simulation setting, these calibration plots are not available in practice since they compare predictions against the true individual treatment benefits (which do not even have an observed individual level equivalent).

8 | APPLIED EXAMPLES

8.1 | Acute otitis media

For illustrative purposes, we analyzed data from a randomized, double-blind, placebo-controlled trial of amoxicillin for clinically diagnosed acute otitis media (AOM) in children 6 months to 5 years of age.⁴⁹ This trial included 512 children and collected baseline data on antibiotic treatment received, sex, presence of recurrent AOM, fever, bilateral occurrence, ear pain, presence of a runny nose, cough, tympanic membrane abnormality, and age. All variables but the latter were dichotomous. The endpoint analyzed here is the same as reported by Rovers et al.⁵⁰ positive when either fever or ear pain was present after 3 days of follow-up. While not truly binary, composite endpoints occur frequently in practice. Thus, data were available on a total of 9 patient characteristics, treatment, and on a composite dichotomous endpoint. All in all, there were 147 events.

We first fitted a logistic regression model on the full data set with main effects for treatment and the 9 patient characteristics. The estimated log odds of treatment was statistically insignificant, but in the expected direction ($\hat{\beta}_t = -0.34$, SE = 0.20). The apparent Nagelkerke R^2 for this model was only 8.8%, and a larger sample size would be generally be required for prediction model development in such low signal settings.⁵¹ Considering the simulation study results, the sample size, and the low amount of signal with respect to predicted risk, the starting point for any prediction of individualized treatment effect in these data is very weak. Nonetheless, it is interesting to see whether the proposed methods indeed show a lack of predictive ability. To that effect, we internally validated all of the methods evaluated in the simulation study (except for the use of content knowledge) in 100 bootstrap samples. The lowest out-of-sample Brier score (0.202) was obtained with the homogeneous treatment effect model fitted by means of ridge regression. However, the accompanying out-of-sample Nagelkerke R^2 was near-zero and even negative for more flexible models. Together, these findings show a lack of strong support for a nonzero average treatment effect and for any ability to personalize treatment effect.

8.2 | International Stroke Trial

The International Stroke Trial (IST) was a large randomized open trial comparing no antithrombotic treatment, aspirin treatment, and subcutaneous heparin treatment in a total of 19 435 patients with acute ischemic stroke.⁵² The individual patient data from the IST trial are available for public use and can be downloaded from the web.⁵³ For the current applied example, we used data from patients randomized between no treatment ($n = 4860$) and aspirin treatment ($n = 4858$). Interestingly, the effect of aspirin treatment on the combined endpoint of death or dependency at 6 months was evaluated conditional on a prognostic score in the original article (and found to be effective on average). The prognostic score was based on age, sex, state of consciousness, and eight other neurological symptoms evaluated at baseline. All variables except for age (continuous) and sex (dichotomous) are categorical variables with three levels. The total number of events was 6043. Ninety-nine patients with a missing outcome were omitted; covariates were complete.

We first fitted a logistic regression model on the full data set with a main effect for treatment and main effects for the sex, state of consciousness and the eight neurological signs. Categorical variables were dummy coded. The estimated log odds of treatment in this model was -0.11 (SE 0.048), corresponding to an odds ratio of 0.90, and the apparent C -statistic and Nagelkerke R^2 were 0.79 and 0.31, respectively. Even though the relative treatment effect is small, the presence of an average effect, the ability to explain a substantial part of the risk of an event, and the amount of available data provide a good starting position when aiming to individualize treatment effect prediction. In terms of the simulation study results, the expectancy is that a HTE model would pick up heterogeneity if it is present and that a homogeneous model would be preferred otherwise.

For illustrative purposes, we therefore examine the predictions from a ridge homogeneous treatment effect model (HOM-ridge) and a hierarchical group lasso HTE (HGL-HTE) model as fitted in the entire sample. In both cases, the penalty parameter with the lowest 10-fold cross-validation deviance was selected. Figure 4 shows the very high correlation (0.996) between risks predicted from either model (upper left panel). Nonetheless, the distribution of $\hat{\delta}(\mathbf{x}_i)$ is quite different for both models (upper right and lower left panel). The HOM-ridge model hardly discriminates w.r.t. treatment effect, whereas the HTE-HGL model predicts harm for a substantial part of the population. The lower right panel shows the relation between $\hat{\delta}(\mathbf{x}_i)$ as predicted from both models.

The main question is whether any, and if so, which of the predictions of $\delta(\mathbf{x}_i)$ can be expected to generalize. A key limitation is that $\hat{\delta}(\mathbf{x}_i)$ cannot be validated directly in real data. A possible approximation is to check whether groups based

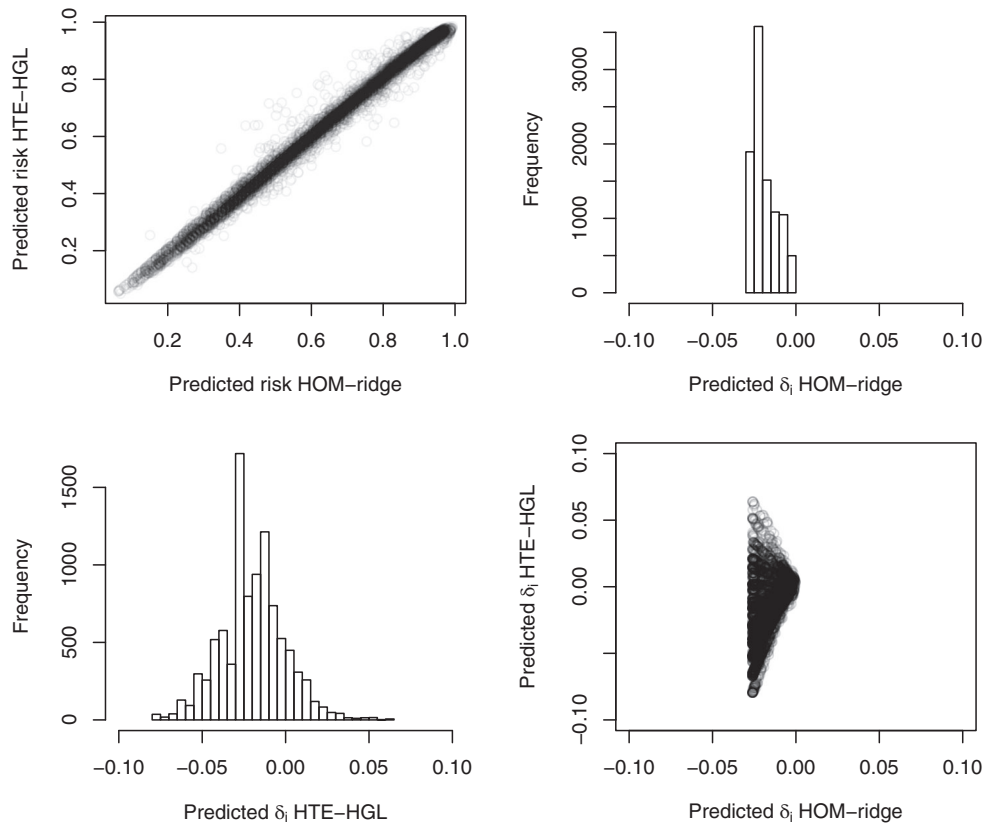


FIGURE 4 International Stroke Trial (IST) data. The upper left panel shows the relation between predicted risks based on the ridge homogeneous treatment effect model (HOM-ridge) and the hierarchical group lasso HTE (HTE-HGL). The histograms show the distribution of predicted individualized treatment effect for the same models. The lower right panel shows their mutual relation

on quantiles of $\hat{\delta}(\mathbf{x}_i)$ relate to observed treatment effect within these same groups. This is essentially a group level effort to approximate calibration at the level of $\hat{\delta}(\mathbf{x}_i)$. The idea is to make use of the fact that the potential outcomes, and thus $\delta(\mathbf{x}_i)$, are independent of treatment assignment (exchangeability). Figure 5 shows apparent and bootstrap results. As can be seen, the apparent calibration of predicted $\hat{\delta}(\mathbf{x}_i)$ is not good for HOM-ridge and reasonable for HTE-HGL. However, in both cases, bootstrap results show such a high degree of variation and lack of trend that the prediction of $\hat{\delta}(\mathbf{x}_i)$ cannot be trusted. In case of the HOM-ridge model, this may relate to the fact that such small variations in $\hat{\delta}(\mathbf{x}_i)$ cannot just not be retrieved from limited out-of-sample cases. In case of the HTE-HGL model, it implies that even in such a large sample, penalized models may still overfit.

To further illustrate this, we again applied all methods evaluated in the main simulation study (except for the application of content knowledge). We evaluated model fit by means of out-of-sample Brier score and Nagelkerke R^2 in 100 bootstrap replications. The best Brier score was achieved for the homogeneous treatment effect model fitted with standard maximum likelihood, closely followed by the equivalent ridge version and the HGL and ridge HTE model (0.17910, 0.17915, 0.17921, and 0.17933, respectively). The same group of methods had the highest Nagelkerke R^2 values (0.3055, 0.3059, 0.3062, and 0.3066, respectively), with little difference between them. All in all, the differences in bootstrap-corrected estimates of overall fit were very small even though the number of model parameters differed substantially across models. Such results, and the extremely high correlation between risk predictions from different models (as in the left upper panel of Figure 4), may already be interpreted as red flags with respect to overfitting of the HTE models.

9 | PRACTICAL CONSIDERATIONS

In general, as with regular prediction modeling, overfitting is an important concern and sample size and penalization are key. We refer to recent guidance papers on sample size^{21,22,51} and the use of penalization in clinical prediction

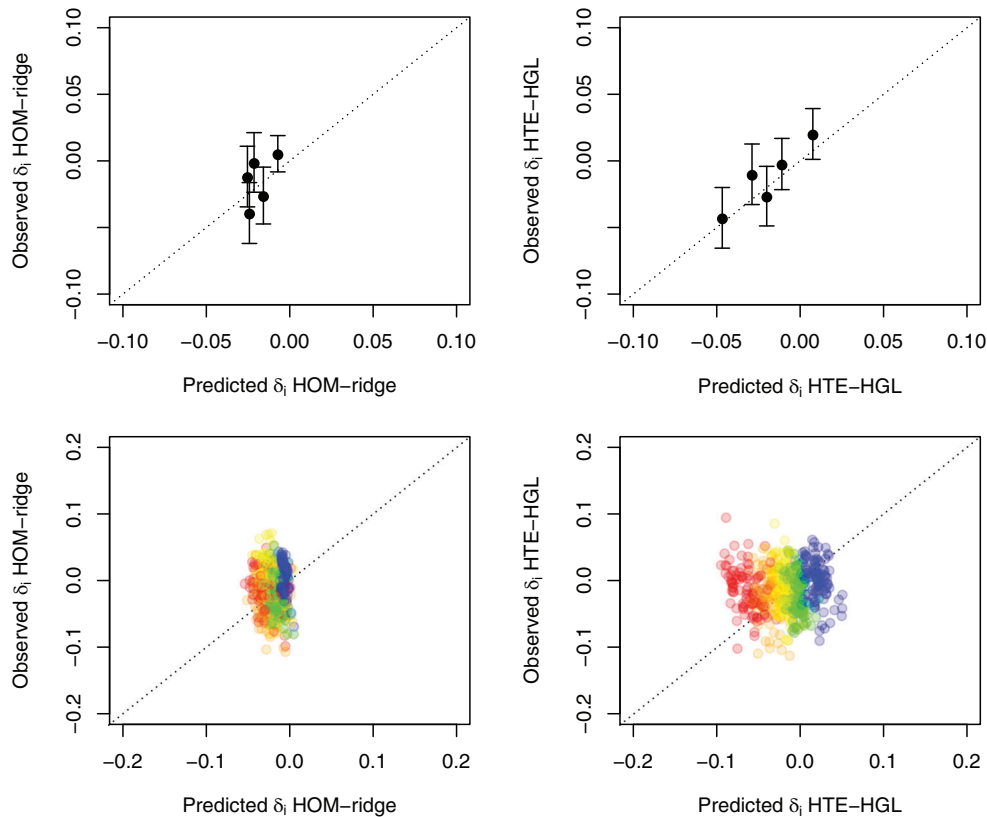


FIGURE 5 International Stroke Trial (IST) data. The upper panels show apparent calibration of treatment effect for the ridge homogeneous treatment effect model (HOM-ridge) and the hierarchical group lasso HTE (HTE-HGL). Dots describe mean $\hat{\delta}(x_i)$ within quintiles of $\hat{\delta}(x_i)$ (x -axis) vs the marginal treatment effect within those same groups (y -axis). Bars provide ± 1 SE. The lower panels show the out-of-sample estimates obtained from 100 bootstrap replications with red representing the lowest quintile $\hat{\delta}(x_i)$, followed by orange, yellow, green, and blue for the highest quintile of $\hat{\delta}(x_i)$ [Colour figure can be viewed at wileyonlinelibrary.com]

modeling.^{27,28} These guidelines provide a lower bound for the sample size when developing models to predict individualized treatment effects. The required sample size will further increase when including treatment-covariate interactions.

With respect to the choice of modeling approach, the HGL performed well across settings. HGL best captured treatment effect heterogeneity when present and had acceptable overfitting otherwise. It clearly outperformed HTE models based on ridge, lasso, or unpenalized maximum likelihood. In practice, a penalized homogeneous treatment effect model provides an important alternative model that is less prone to overfitting, but also less capable of capturing HTE. This modeling approach therefore appears particularly promising when RCTs are relatively small, or when there is (almost) no treatment-effect modification. Cross-validation or bootstrapping may be used to choose between HGL and homogeneous treatment effect models. The IST applied example provided an example of such a bootstrapping approach, showing the evaluation of well-known measures of overall fit (Brier score and Nagelkerke R^2),^{1,54} and a proposal to visually evaluate performance on the level of aggregated individualized treatment effect.

Nonetheless, samples that are very small with respect to the number of model parameters of interest should raise caution.^{21,22,51} While penalization is clearly beneficial on average, accurate estimation of the penalization parameter itself cannot be expected in small sample size settings and there is no way to know whether this affects any particular model in practice.^{27,28} Also, the empirical evaluation of individualized treatment effect models is still in its infancy. Both group level evaluation of predicted individualized treatment effects and individual level evaluation of outcome risk predictions are only indirect approximations of the performance of individualized treatment effect prediction. Due to the insensitivity of the available measures, we expect model comparisons to be conservative in the sense that they will only prefer a heterogeneous model if the evidence is quite strong. For instance, HTE needs to be large enough to substantially affect overall outcome predictions, or sample size needs to be large enough to reliably assess aggregate predicted vs observed treatment effect. Also, complex models suffer more from the decreased variability in bootstrap samples.

Importantly, any measure based on observed data can only be an approximation of the performance of a potential outcome prediction model, with the remainder resting on assumptions.⁵⁵ The plausibility of the identifiability assumptions is an important part of model evaluation, including thoughts about their transportability, and requires thinking instead of measuring. To the knowledge of the authors, there is no specific guidance on the validation of the potential outcome type of individualized treatment effect prediction models beyond the guidance provided in this tutorial. This remains an important open area of research that needs further study.

Lastly, a note on the interpretation of any identified treatment effect heterogeneity. Randomization of treatment (or unconfoundedness after appropriate modeling of confounders in observational settings) supports causal inference with respect to treatment in a subgroup (some covariate status). This is subtly, but importantly, different from the assumption that subgroup membership causes or explains a change in treatment effect. For instance, subgroup membership might just be correlated with an unobserved underlying cause of treatment effect heterogeneity. Such implicit inverting of the causal relation underlies the exploration of subgroups as identified by individualized treatment effect prediction to (i) inform treatment decisions, or (ii) “explain” treatment effect heterogeneity. However, such conclusions require randomization or unconfoundedness of subgroup membership. Care should be taken to emphasize that the estimated treatment effect describes the causal effect of treatment within a given subgroup, and not necessarily the other way around (ie, it does not warrant the interpretation that the characteristics defining a subgroup cause differential treatment effect). Therefore, while predicted treatment effect heterogeneity may provide interesting hypotheses about its causal structure, it does not provide answers without further thinking about, and analysis of, the causal pathways involved.

10 | DISCUSSION

We have provided an overview of the process of individualized treatment effect prediction in the context of a randomized trial with a binary endpoint. To that effect, we have described the integration of key elements from the fields of causal inference and clinical prediction research. These methods can be used to expand on the mainstay analysis of randomized trials, and may help to uncover between-subject heterogeneity in terms of predicted outcome risk and treatment effect.

With respect to causal inference, we focused on the causal nature of the question of interest and a clear definition of individualized treatment effect based on the potential outcomes framework. From there, we explained the necessary assumptions to identify the individualized treatment effect based on the observed data. While such effects can in principle be estimated nonparametrically, further modeling is beneficial and allows straightforward comparison of treatment effect conditional on many covariates. Even though the prediction problem itself could be solved without any reference to causal inference, going through the motions increases clarity of the research question and gains understanding of the requirements for a causal interpretation of the final model.

With respect to prediction modeling, we focused on the need for a parsimonious model with validity across the risk scale (log odds), while maintaining an interpretable scale for the final result (the risk difference scale). Specification of models for a homogeneous treatment effect (constant relative effect) and differential or heterogeneous treatment effect were described in detail. Subsequently, the relation between prior knowledge and data-driven methodology was examined, revealing the need for both. In line with general sample size guidance when developing a multivariable prediction model,⁵¹ sufficient sample size was important for accurate individualized treatment effects predictions and model stability.

While all of the required ingredients for individualized treatment effect prediction are well-known, their successful combination constitutes a challenging problem that is on the boundary of what can be observed in empirical data. Our simulation study, with a known data generating mechanism, provided clear insight into methods that are able to pick up heterogeneity in sufficiently large samples, while limiting the amount of overfitting in absence of heterogeneity. While very informative, actual analysis of observed data will have to rely on dichotomous outcomes from subjects observed under a single treatment condition that are only incompletely matched across treatment groups. The best way to evaluate the performance of individualized treatment effect prediction models is an open question. We described a bootstrap-based internal validation approach that decreases the risk of overfitting. A very recent contribution to the literature on potential outcome prediction and individualized treatment effect describes a very similar split sample approach.⁵⁶ Also, a novel type of c-index has been suggested to measure discriminative performance of individualized treatment effect predictions.⁵⁷ Nguyen et al⁵⁶ provide some cautionary notes on its interpretation and estimated standard error in their appendix.

BOX 1: Key points and recommendations

- It is important to clearly define the individualized treatment effect of interest and to be aware of the identifiability assumptions underlying its causal interpretation.
- Analogous to causal inference with respect to average treatment effect, randomization of treatment greatly facilitates causal inference with respect to predicted individualized treatment effects.
- Logistic regression provides a parsimonious model to predict absolute individualized treatment effect (ie, treatment effect on the risk difference scale) in new patients. Even in absence of treatment-covariate interaction (ie, homogeneous patient-level treatment effect on odds ratio scale), a logistic model accounting for individual patient characteristics (prognostic factors) can lead to meaningful differentiation in terms of absolute treatment effect.
- The inclusion of treatment-covariate interactions to model heterogeneous treatment effect (ie, patient-level heterogeneity in the treatment effect on the odds ratio scale) should preferably be based on biological and clinical rationale and be informed by statistical evidence from prior studies.
- Sample size and penalized estimation are of key importance for accurate individualized treatment effect prediction; penalization alone does not guarantee accurate predictions in new individuals when sample size is insufficient with respect to model complexity. Existing guidelines on clinical prediction modeling provide a lower bound for the sample size needed in case of individualized treatment effect prediction.^{21,22,51}
- In practice, bootstrap internal validation of likelihood-based measures of overall fit (eg, R^2 , AIC), mean squared prediction error (eg, Brier score), and aggregate observed vs expected measures of treatment effect variability (as in Section 8.2) help to choose among competing models. It is recommended to include a homogeneous treatment effect model as a starting point (reference model) and to consider more complex models based on biological, clinical, and statistical evidence.
- Future work is needed to further delineate best practices for the evaluation of individualized treatment effect predictions and models, hence also improving model comparison and validation procedures.
- Purely explorative indications of heterogeneous treatment effect provide an interesting starting point for further research (eg, into the causal structure of the heterogeneity) and require external validation.
- This tutorial handles causal prediction of treatment effect on a binary outcome, conditional on individual level patient characteristics. This should not be confused with a causal interpretation of the effect of individual level patient characteristics on the effect of treatment.

The prediction of individualized or personalized treatment effect is an active field of research. Recent broad overviews on predictive approaches toward heterogeneity of treatment effect are available elsewhere and include a comprehensive overview of applied papers.^{5,58} Related work approaches the problem from the missing data perspective.^{59,60} Also, work has been done on individualized treatment effect prediction for optimal treatment selection^{61,62} and selection of patients for future studies.^{63,64} All in all, the literature on personalized medicine approaches that use prediction modeling is vast and too extensive to cite here. What we add is a clear, principled and from the ground-up overview that integrates prediction modeling with causal inference and accentuates the importance of study design features.

To limit the scope, we did not venture into the incorporation of post-randomization measurements, dropout and selection bias, and observational data. Such topics require careful attention to the exchangeability assumption, which is no longer fulfilled by the study design and needs further assumptions and careful modeling with respect to all possible confounders. A recent scoping review provides an overview of the literature with respect to methods for causal prediction that extend to observational data.⁶⁵ Also, where we have focused on intention to treat estimates of point exposure treatment, different settings and questions require further thought on the relevant definition of the estimand.⁶⁶ As a second limitation, we provided a small simulation study covering a limited number of settings that was designed for illustrative purposes. The setup was such that development and test sets were generated from the same data generating mechanism. In practice, there may be differences between these two settings that are not captured by the models, and the uncertainty that accompanies these unknowns may overshadow relatively small gains realized by more complex models.⁶⁷

More general limitations pertain to the typical randomized trial design that provides the data to be used for individualized treatment effect prediction. Other designs, such as $N \times 1$ trials and cross-over designs may provide more direct within-person comparability, and thereby also provide information on the stability of treatment response for an individual.⁶⁸ However, these designs are infeasible for many conditions and have their own set of challenges.⁶⁹ More

importantly, randomized trials are typically designed to be of sufficient sample size to reveal an anticipated average relative treatment effect. Therefore, randomized trials are not designed for complex prediction modeling. Hence, if we want to walk down the avenue of individualized treatment effect modeling, we will either have to design trials with this purpose in mind, or have to find more creative ways to amplify our data. This could include the analysis of individual patient data from multiple randomized trials, or even the use of nonrandomized studies for the estimation of outcome risk under a control condition.⁶⁵ Besides clear opportunities, such approaches also bring about many new challenges. For instance, typical challenges that occur in clustered data settings (eg, between-study heterogeneity) have been comprehensively illustrated in a recent tutorial on the examination of heterogeneous (relative) treatment effect in patient-level data from multiple randomized trials.⁷⁰ The implications of such challenges in the context of causal prediction research require further study.

11 | CONCLUDING REMARKS

We hope that our overview on the basics underlying individualized treatment effect prediction in binary endpoint settings is a useful guide and starting point for statisticians interested in this area. The successful implementation of individualized treatment effect prediction requires careful thought on the exact nature of the question and estimand(s) of interest, the causal and modeling assumptions relied on, and the ever-present bias-variance trade-off that requires even greater care than usual when working with potential outcomes. Sample size considerations are important in all areas of research and there is increasing awareness on the need for larger sample sizes when developing prediction models and examining treatment-covariates interaction. Future work is needed on the validation of models for predicted individualized treatment effect, their role in uncovering sources of heterogeneity, and ways to account for the clustered nature of many data sets. Also, beyond the frequentist framework, the basis for a fully Bayesian approach has long been recognized⁷¹ and could combine the advantage of penalization with a more thorough view on the posterior distribution of the model parameters. A summary of key recommendations and findings is provided in Box 1.

ACKNOWLEDGEMENTS

We thank the reviewers and the editorial team for their careful reading, thoughtful comments, and constructive remarks that were greatly beneficial to the manuscript. We also like to thank the researchers involved in the original otitis media⁴⁹ and stroke trials⁵² for use of their data.

Jeroen Hoogland, Michail Belias, Joanna IntHout, Maroeska M. Rovers, Thomas P. A. Debray and Johannes B. Reitsma acknowledge financial support from the Netherlands Organisation for Health Research and Development (grant 91215058). Thomas P. A. Debray also acknowledges financial support from the Netherlands Organisation for Health Research and Development (grant 91617050)

DATA AVAILABILITY STATEMENT


Data for the International Stroke Trial applied example are publicly available.⁵³ Data for the otitis media applied example not available for sharing since they contain privacy sensitive data according to the General Data Protection Regulation. R scripts to perform the simulation study, including data generation and analysis, are available for sharing.

ORCID


Jeroen Hoogland  <https://orcid.org/0000-0002-2397-6052>

Joanna IntHout  <https://orcid.org/0000-0002-6127-0747>

Michail Belias  <https://orcid.org/0000-0003-2024-0231>

Maroeska M. Rovers  <https://orcid.org/0000-0002-3095-170X>

Richard D. Riley  <https://orcid.org/0000-0001-8699-0735>

Frank E. Harrell Jr  <https://orcid.org/0000-0002-8271-5493>

Karel G. M. Moons  <https://orcid.org/0000-0003-2118-004X>

Thomas P. A. Debray  <https://orcid.org/0000-0002-1790-2719>

REFERENCES

1. Steyerberg E. *Clinical Prediction Models*. Statistics for Biology and Health. New York, NY: Springer; 2009.
2. Senn S. *Statistical Issues in Drug Development*. 2nd ed. Chichester, England; Hoboken, NJ: John Wiley & Sons; 2007 OCLC: ocn180907943.

3. Glasziou PP, Irwig LM. An evidence based approach to individualising treatment. *BMJ*. 1995;311:1356-1359.
4. Kent DM, Steyerberg E, van Klaveren D. Personalized evidence based medicine: predictive approaches to heterogeneous treatment effects. *BMJ*. 2018;363:k4245.
5. Kent DM, van Klaveren D, Paulus JK, et al. The predictive approaches to treatment effect heterogeneity (PATH) statement: explanation and elaboration. *Ann Internal Med*. 2020;172:W1-W25.
6. Dahabreh IJ, Hayward R, Kent DM. Using group data to treat individuals: understanding heterogeneous treatment effects in the age of precision medicine and patient-centred evidence. *Int J Epidemiol*. 2016;45(6):2184-2193.
7. Murray EJ, Caniglia EC, Swanson SA, Hernández-Díaz S, Hernán MA. Patients and investigators prefer measures of absolute risk in subgroups for pragmatic randomized trials. *J Clin Epidemiol*. 2018;103:10-21.
8. Pearl J, Glymour M, Jewell NP. *Causal Inference in Statistics: A Primer*. Chichester, West Sussex: Wiley; 2016.
9. Rubin DB. Causal inference using potential outcomes: design, modeling, decisions. *J Am Stat Assoc*. 2005;100:322-331.
10. Rubin DB. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educ Psychol*. 1974;66(5):688-701.
11. Holland P. Statistics and causal inference. *J Am Stat Assoc*. 1986;81(396):945-960.
12. Hernan MA. A definition of causal effect for epidemiological research. *J Epidemiol Community Health*. 2004;58:265-271.
13. Hernán MA, Robins JM. *Causal Inference: What If*. Boca Raton, FL: Chapman & Hall/CRC Press; 2020.
14. Hernán MA, Hernández-Díaz S, Robins JM. Randomized trials analyzed as observational studies. *Ann Internal Med*. 2013;159(8):560-562.
15. Hernán MA, Robins JM. Per-protocol analyses of pragmatic trials. *N Engl J Med*. 2017;377:1391-1398.
16. Goetghebeur E, le Cessie S, De Stavola B, Moodie EE, Waernbaum I, The Topic Group Causal Inference (TG7) of the STRATOS Initiative. Formulating causal questions and principled statistical answers. *Stat Med*. 2020;39(30):4922-4948.
17. Spiegelhalter DJ. Probabilistic prediction in patient management and clinical trials. *Stat Med*. 1986;5:421-433.
18. Austin PC. Absolute risk reductions, relative risks, relative risk reductions, and numbers needed to treat can be obtained from a logistic regression model. *J Clin Epidemiol*. 2010;63:2-6.
19. Nelder JA, Wedderburn RWM. Generalized linear models. *J Royal Stat Soc Ser A (General)*. 1972;135(3):370-384.
20. Agresti A. *Foundations of Linear and Generalized Linear Models*. Wiley Series in Probability and Statistics. Hoboken, NJ: John Wiley & Sons Inc; 2015.
21. Riley RD, Ensor J, Snell KIE, et al. Calculating the sample size required for developing a clinical prediction model. *BMJ*. 2020;368:m441.
22. van Smeden M, Moons KG, de Groot JA, et al. Sample size for binary logistic prediction models: beyond events per variable criteria. *Stat Methods Med Res*. 2019;28(8):2455-2474.
23. Hoerl AE, Kennard RW. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics*. 1970;12:55-67.
24. Tibshirani R. Regression shrinkage and selection via the lasso. *J Royal Stat Soc Ser B (Methodol)*. 1996;58(1):267-288.
25. Efron B, Hastie T, Johnstone I, Tibshirani R. Least angle regression. *Ann Stat*. 2004;32(2):407-499.
26. Friedman J, Hastie T, Tibshirani R. Regularization paths for generalized linear models via coordinate descent. *J Stat Softw*. 2010;33(1):1-22.
27. Van Calster B, van Smeden M, De Cock B, Steyerberg EW. Regression shrinkage methods for clinical prediction models do not guarantee improved performance: simulation study. *Stat Methods Med Res*. 2020;29:3166-3178.
28. Riley RD, Snell KI, Martin GP, et al. Penalisation and shrinkage methods produced unreliable clinical prediction models especially when sample size was small. *J Clin Epidemiol*. 2020;132:88-96.
29. Hastie T, Tibshirani R, Wainwright M. No. 143 in Monographs on Statistics and Applied Probability. *Statistical Learning with Sparsity: The Lasso and Generalizations*. Boca Raton, FL: CRC Press, Taylor & Francis Group; 2015.
30. Hastie T. Ridge regularization: an essential concept in data science. *Technometrics*. 2020;62:426-433.
31. Lim M, Hastie T. Learning interactions via hierarchical group-lasso regularization. *J Comput Graph Stat*. 2015;24:627-654.
32. Yang Y, Zou H. A fast unified algorithm for solving group-lasso penalized learning problems. *Stat Comput*. 2015;25:1129-1141.
33. Sussman JB, Kent DM, Nelson JP, Hayward RA. Improving diabetes prevention with benefit based tailored treatment: risk based reanalysis of diabetes prevention program. *BMJ*. 2015;350:h454.
34. Kent DM, Rothwell PM, Ioannidis JP, Altman DG, Hayward RA. Assessing and reporting heterogeneity in treatment effects in clinical trials: a proposal. *Trials*. 2010;11(85):1-11.
35. Burke JF, Hayward RA, Nelson JP, Kent DM. Using internally developed risk models to assess heterogeneity in treatment effects in clinical trials. *Circ Cardiovasc Qual Outcomes*. 2014;7(1):163-169.
36. Wager S, Athey S. Estimation and inference of heterogeneous treatment effects using random forests. *J Am Stat Assoc*. 2018;113:1228-1242.
37. Breiman L. Random forests. *Mach Learn*. 2001;45(1):5-32.
38. Lu M, Sadiq S, Feaster DJ, Ishwaran H. Estimating individual treatment effect in observational data using random forest methods. *J Comput Graph Stat*. 2018;27:209-219.
39. Athey S, Tibshirani J, Wager S. Generalized random forests. *Ann Stat*. 2019;47(2):1148-1178.
40. Zeileis A, Hothorn T, Hornik K. Model-based recursive partitioning. *J Comput Graph Stat*. 2008;17:492-514.
41. Seibold H, Zeileis A, Hothorn T. Model-based recursive partitioning for subgroup analyses. *Int J Biostat*. 2016;12(1):45-63.
42. Scarpa J, Bruzeilius E, Doupe P, Le M, Faghmous J, Baum A. Assessment of risk of harm associated with intensive blood pressure management among patients with hypertension who smoke: a secondary analysis of the systolic blood pressure intervention trial. *JAMA Netw Open*. 2019;2:e190005.
43. Rigdon J, Baiocchi M, Basu S. Preventing false discovery of heterogeneous treatment effect subgroups in randomized trials. *Trials*. 2018;19:382.

44. van der Ploeg T, Austin PC, Steyerberg EW. Modern modelling techniques are data hungry: a simulation study for predicting dichotomous endpoints. *BMC Med Res Methodol.* 2014;14:137.
45. Austin PC, Harrell FE, Steyerberg EW. Predictive performance of machine and statistical learning methods: impact of data-generating processes on external validity in the “large N, small p” setting. *Stat Methods Med Res.* 2021;30(6):1465-1483.
46. Burton A, Altman DG, Royston P, Holder RL. The design of simulation studies in medical statistics. *Stat Med.* 2006;25:4279-4292.
47. Morris TP, White IR, Crowther MJ. Using simulation studies to evaluate statistical methods: using simulation studies to evaluate statistical methods. *Stat Med.* 2019;38(11):2074-2102.
48. R Core Team. *R: A Language and Environment for Statistical Computing.* Vienna, Austria: R Foundation for Statistical Computing; 2017.
49. Le Saux N. A randomized, double-blind, placebo-controlled noninferiority trial of amoxicillin for clinically diagnosed acute otitis media in children 6 months to 5 years of age. *Can Med Assoc J.* 2005;172:335-341.
50. Rovers MM, Glasziou P, Appelman CL, et al. Antibiotics for acute otitis media: a meta-analysis with individual patient data. *Lancet.* 2006;368:1429-1435.
51. Riley RD, Snell KI, Ensor J, et al. Minimum sample size for developing a multivariable prediction model: Part II - binary and time-to-event outcomes. *Stat Med.* 2019;38:1276-1296.
52. International Stroke Trial Collaborative Group. The International Stroke Trial (IST): a randomised trial of aspirin, subcutaneous heparin, both, or neither among 19 435 patients with acute ischaemic stroke. *Lancet.* 1997;349:1569-1581.
53. Sandercock PA, Niewada M, Czlonkowska A, International Stroke Trial Collaborative Group. The International Stroke Trial database. *Trials.* 2011;12:101.
54. Harrell FE. *Regression Modeling Strategies: With Applications to Linear Models, Logistic and Ordinal Regression, and Survival Analysis.* Springer Series in Statistics. 2nd ed. Cham, Switzerland; Heidelberg, Germany; New York, NY: Springer; 2015 OCLC 922304565.
55. Dickerman BA, Hernán MA. Counterfactual prediction is not only for causal inference. *Eur J Epidemiol.* 2020;35:615-617.
56. Nguyen T-L, Collins GS, Landais P, Le Manach Y. Counterfactual clinical prediction models could help to infer individualized treatment effects in randomized controlled trials—an illustration with the International Stroke Trial. *J Clin Epidemiol.* 2020;125:47-56.
57. van Klaveren D, Steyerberg EW, Serruys PW, Kent DM. The proposed ‘concordance-statistic for benefit’ provided a useful metric when modeling heterogeneous treatment effects. *J Clin Epidemiol.* 2018;94:59-68.
58. Rekkas A, Paulus JK, Raman G, et al. Predictive approaches to heterogeneous treatment effects: a scoping review. *BMC Med Res Methodol.* 2020;20:264.
59. Lamont A, Lyons MD, Jaki T, et al. Identification of predicted individual treatment effects in randomized clinical trials. *Stat Methods Med Res.* 2018;27(1):142-157.
60. Buuren S v. *Flexible Imputation of Missing Data.* Chapman & Hall/CRC Interdisciplinary Statistics Series. 2nd ed. Boca Raton, FL: CRC Press, Taylor & Francis Group; 2018.
61. Cai T, Tian L, Wong PH, Wei LJ. Analysis of randomized comparative clinical trial data for personalized treatment selections. *Biostatistics.* 2011;12:270-282.
62. Kang C, Janes H, Huang Y. Combining biomarkers to optimize patient treatment recommendations: combine markers for treatment selection. *Biometrics.* 2014;70:695-707.
63. Zhao L, Tian L, Cai T, Claggett B, Wei LJ. Effectively selecting a target population for a future comparative study. *J Am Stat Assoc.* 2013;108:527-539.
64. Li J, Zhao L, Tian L, et al. A predictive enrichment procedure to identify potential responders to a new therapy for randomized, comparative controlled clinical studies: predictive enrichment. *Biometrics.* 2016;72:877-887.
65. Lin L, Sperrin M, Jenkins DA, Martin GP, Peek N. A scoping review of causal methods enabling predictions under hypothetical interventions. *Diagn Prognost Res.* 2021;5:3.
66. van Geloven N, Swanson S, Ramspek C, et al. Prediction meets causal inference: the role of treatment in clinical prediction models. *Eur J Epidemiol.* 2020;35:619-630. arXiv: 2004.06998.
67. Hand DJ. Classifier technology and the illusion of progress. *Stat Sci.* 2006;21:1-14.
68. Senn S. Statistical pitfalls of personalized medicine. *Nature.* 2018;563:619-621.
69. Senn S. *Cross-Over Trials in Clinical Research.* Statistics in Practice. 2nd ed. Chichester, England; New York, NY: J. Wiley; 2002.
70. Riley RD, Debray TP, Fisher D, et al. Individual participant data meta-analysis to examine interactions between treatment effect and participant-level covariates: statistical recommendations for conduct and planning. *Stat Med.* 2020;39(15):2115-2137.
71. Simon R. Bayesian subset analysis: application to studying treatment-by-gender interactions. *Stat Med.* 2002;21:2909-2916.

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of this article.

How to cite this article: Hoogland J, IntHout J, Belias M, et al. A tutorial on individualized treatment effect prediction from randomized trials with a binary endpoint. *Statistics in Medicine.* 2021;40(26):5961-5981. doi: 10.1002/sim.9154