Published in final edited form as:

Biometrics. 2017 December ; 73(4): 1355–1365. doi:10.1111/biom.12677.

Bayesian Population Finding with Biomarkers in a Randomized Clinical Trial

Satoshi Morita^{1,*} and Peter Müller²

¹Department of Biomedical Statistics and Bioinformatics, Kyoto University Graduate School of Medicine, Kyoto, Japan

²Department of Mathematics, University of Texas, Austin, TX, U.S.A

Summary

The identification of good predictive biomarkers allows investigators to optimize the target population for a new treatment. We propose a novel utility-based Bayesian population finding (BaPoFi) method to analyze data from a randomized clinical trial with the aim of finding a sensitive patient population. Our approach is based on casting the population finding process as a formal decision problem together with a flexible probability model, Bayesian additive regression trees (BART), to summarize observed data. The proposed method evaluates enhanced treatment effects in patient subpopulations based on counter-factual modeling of responses to new treatment and control for each patient. In extensive simulation studies, we examine the operating characteristics of the proposed method. We compare with a Bayesian regression-based method that implements shrinkage estimates of subgroup-specific treatment effects. For illustration, we apply the proposed method to data from a randomized clinical trial.

Keywords

Bayesian additive regression trees; Bayesian decision problem; Biomarkers; Population finding; Randomized clinical trial

1. Introduction

We consider the problem of setting up a clinical trial to find a sensitive subpopulation for a new treatment. We approach the problem as a Bayesian decision problem. The main contributions are the separation of the probability model and the criterion for the choice of the sensitive subpopulation; related to this, the use of flexible nonparametric Bayesian probability models for the outcome; the practical feasibility of the proposed approach in moderate size phase II as well as larger phase III studies; and the comparison with standard methods (Dixon and Simon, 1991). We do not require a particular inference model. The approach is valid with any sufficiently flexible outcome model.

*Address for correspondence: Satoshi Morita, PhD, Department of Biomedical Statistics and Bioinformatics, Kyoto University Graduate School of Medicine, Address: 54 Kawahara-cho, Shogoin, Sakyo-ku, Kyoto 606-8507, Japan, smorita@kuhp.kyoto-u.ac.jp
Phone: +81-75-751-4717, Fax: +81-75-751-4767.

Conflict of interest statement: There are no conflicts of interest in this study.

The importance of examining predictive biomarkers for efficacy of a new treatment in a clinical development has been extensively discussed, among others, by McShane *et al.* (2009), Seymour *et al.* (2010), and Renfro *et al.* (2016). New drug development should aim to optimize the target population of patients for treatment by appropriately focusing on patients who are most likely to benefit from a new treatment. Identifying such sensitive subpopulations of patients is increasingly more important in clinical development as more targeted therapies are considered.

We propose a novel utility-based Bayesian population finding (BaPoFi) method to find a sensitive population in the context of a randomized clinical trial (RCT). More specifically, we develop a statistical tool to inform a “Go / No-go” decision making at the end of an RCT, that is, on whether or not to further examine a new treatment in a particular patient population in a subsequent trial. In general, the aim may be to find a patient population for which a new treatment is particularly effective, or for which it is particularly harmful, such that the identified population can be targeted for recruitment, respectively exclusion in a subsequent clinical trial. In this paper, we focus on the earlier only. The desired subpopulation is characterized by patient baseline covariates. Hereafter, we use “biomarkers” to refer to such predictive baseline covariates.

Let H_a denote a hypothesis of an enhanced treatment effect for a selected patient subpopulation. Let H_0 and H_1 denote the overall null hypothesis of no treatment effect and the overall alternative hypothesis of a common treatment effect for the entire population of patients, respectively. In a randomized phase II (RPII) trial, one may make a decision between i) H_0 : no-go to phase III and terminate clinical development, ii) H_a : go to phase III in a selected subpopulation (focused phase III trial), iii) H_1 : go to phase III for the entire population, and possibly iv) H_a and H_1 : launch two separate phase III trials in a selected subpopulation and the entire population, respectively. As for a randomized phase III (RPIII) trial, one may i) H_0 : conclude that a new treatment is not effective and consider no further development, ii) H_a : conduct a new (additional) RPIII trial to confirm an enhanced treatment effect of a new therapy in a selected subpopulation, or iii) H_1 : confirm the treatment effect of a new therapy in the overall patient population (without further study).

We focus on a two-arm clinical trial and do not consider adaptive randomization or interim analysis during the trial. Using all data observed in an RCT, we aim to accurately identify the predictors of efficacy for a new therapy if there are and precisely finding the sensitive subpopulation defined by those predictor(s). For application, we use continuous covariates such as genetic biomarkers (e.g., gene expression level) to define patient sub-populations. Our approach is based on casting the population finding process as a formal decision problem which is defined on top of any flexible inference model. We only focus on the action of reporting a subgroup of patients who might most benefit from a new treatment. Flexible probability models that could be used with this approach include multivariate adaptive regression splines (MARS) (Friedman, 1991), classification and regression trees (CART) (Chipman *et al.*, 1998; Denison *et al.*, 1998), Bayesian additive regression trees (BART) (Chipman *et al.*, 2010), Gaussian processes (Rasmussen and Williams, 2006), etc. Methodologically, we separate the construction of the assumed sampling model and an

action space to address the problem of population-finding. In particular, the reported subgroup is not linked with the parameters of the probability model.

There is substantial literature on analytical methods to find patient subpopulations with enhanced treatment effects. Dixon and Simon (1991) and Simon (2002) proposed Bayesian shrinkage estimators of subgroup-specific treatment effects, evaluated using a regression model with interaction effects between treatment and covariates. Sivaganesan *et al.* (2010) developed a Bayesian model selection approach. Ruberg *et al.* (2010) and Lipkovich *et al.* (2011) used a tree-based method, in particular, a recursive partitioning method to identify patient subgroups with significantly different response to treatment. Foster *et al.* (2011) developed tree-based algorithms to evaluate enhanced treatment effects in patient subgroups. Their method, called “Virtual Twins”, was based on the idea of counter-factual modeling to predict responses to new treatment and control for each subject. Sivaganesan *et al.* (2016), in parallel work, propose a similar Bayesian decision theoretic approach for subgroup analysis. A good recent review appears in Jones *et al.* (2011). For decision making related to the study design of a subsequent RPIII trial (e.g., standard, enrichment, or biomarker-stratified designs), it may be useful to develop an RPII study design based on pre-defined biomarker-positive and biomarker-negative subgroups, as in Fredlin *et al.* (2012). However, this is a different framework from our proposed method which includes the construction of candidate subgroups as part of the inference. Throughout the discussion, we use examples and language specific to a single treatment choice and a single outcome. However, the approach remains valid in the context of a comparison of treatment regimens with multi-step treatment choices, such as, for example, front-line and salvage therapy in a chemotherapy trial, as in Xu *et al.* (2016).

This paper is organized as follows. In Section 2, we provide motivating examples. We explain a probability model to summarize data for inference of heterogeneous treatment effects and subgroups in Section 3. Section 4 presents our proposed approach to finding a sensitive subpopulation as a decision problem. We conduct extensive simulation studies to examine the operating characteristics of the proposed method in Section 5. Section 6 shows an application of the proposed method. We close with a brief discussion in Section 7.

2. Motivating examples

We consider the development of a new therapy in an oncology clinical trials. We assume a situation where the number of candidate biomarkers could be narrowed, say to 5 ~ 20 markers, on the basis of preclinical and clinical data. Many RPII and RPIII clinical trials now routinely include extensive collection of predictive biomarkers that could be used to identify sensitive patient-populations to a new treatment. For example, Tsuburaya *et al.* (2012) describe an RPII trial for advanced gastric cancer that was designed to compare two types of oral fluoropyrimidine. The two treatments show several different characteristics in the mechanisms of their antitumor effect. On the basis of preclinical and clinical data, this RPII trial prespecified five mRNA expression levels as candidate biomarkers to determine which agent should be used for each patient. For another example, Okamoto *et al.* (2010) describe an RPIII trial that compares two combination chemotherapies in patients with advanced non-small cell lung cancer. Later, in Section 6, we will show an application of the

proposed method to the data from this RPIII trial. The protocol includes recording of 17 continuous biomarkers (16 mRNA and 1 protein expression levels). As efficacy outcome we use maximal tumor size shrinkage percentage (maxTS%). This and other continuous scale measurements of change in tumor size have been increasingly used to evaluate the efficacy of new treatments, in addition to more conventional endpoints such as tumor response, progression-free survival, and overall survival (Karrison *et al.*, 2007).

3. Heterogeneous treatment effects and subgroups

Our approach is based on casting inference for population finding as a formal decision problem. The basic components of a decision-theoretic setup include an action space $\mathcal{A}$ of possible decisions $a \in \mathcal{A}$, a probability model $p(\mathbf{y}, \boldsymbol{\mu})$ for data $\mathbf{y}$ and parameters $\boldsymbol{\mu}$, and a utility function $u(a, \boldsymbol{\mu}, \mathbf{y})$. Data may include observed data and future data. We will use $\tilde{\mathbf{y}}$ to separately indicate future data when needed. The parameters $\boldsymbol{\mu}$ could, for example, be regression parameters representing main and interaction effects. See Sections 3.1 and 4.3 for specific examples. Also, we use $p(\cdot)$ to generically denote probability models, and introduce separate notation for different models only when context and arguments do not clarify the use. A utility function $u(\cdot)$ quantifies relative preferences for hypothetical future outcomes ($\tilde{\mathbf{y}}$) and assumed parameter values ($\boldsymbol{\mu}$) under alternative decisions a , given observed data ($\mathbf{y}$). It can be argued that a rational decision maker should choose an action in $\mathcal{A}$ to maximize the expectation of u (Robert, 2007). The expectation is with respect to p , conditioning on all observed data, and marginalizing over all parameters and all future data. We will use $U(a)$ to denote expected utility.

3.1 Notation

We consider a RCT with a total sample size of n patients. For patient i , let y_i denote the observed outcome variable and let z_i denote a treatment indicator with $z_i = 0$ for control (C) and $z_i = 1$ for experimental therapy (N). Let $\mathbf{y} = (y_1, \dots, y_n)$ and $\mathbf{z} = (z_1, \dots, z_n)$ denote vectors of outcomes and treatment indicators for all n patients, respectively. Let $\mathbf{x}_i = (x_{i,1}, \dots, x_{i,p})$ denote the patient baseline covariates, and let $\mathbf{X} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ denote the covariates for all n patients. In the upcoming discussion, we consider an RCT with $n = 100$ to 300 and $p = 5$ to 20. For the moment, we assume a continuous outcome variable y_i . But extension to binary outcomes is straightforward. In summary, $(\mathbf{y}, \mathbf{z}, \mathbf{X})$ is all data observed in n patients in the clinical trial. Let $\boldsymbol{\mu}$ denote parameters that index a sampling model $p(\mathbf{y} | \boldsymbol{\mu}, \mathbf{z}, \mathbf{X})$. We take a Bayesian approach to inference and complete the model with a prior $p(\boldsymbol{\mu})$. Together, the sampling model and prior define the joint probability model $p(\mathbf{y}, \boldsymbol{\mu})$ that we mentioned in the earlier general outline as $p(\mathbf{y}, \boldsymbol{\mu}) = p(\mathbf{y} | \boldsymbol{\mu}, \mathbf{z}, \mathbf{X}) \cdot p(\boldsymbol{\mu})$ (the additional $\mathbf{z}, \mathbf{X}$ in the conditioning set are a slight abuse of notation to indicate the covariates in a regression model). This implies a posterior probability model $p(\boldsymbol{\mu} | \mathbf{z}, \mathbf{X}, \mathbf{y})$ and a posterior predictive distribution for a future observation $i = n + 1$

$$p(\tilde{y}_{n+1} | z_{n+1}, \mathbf{x}_{n+1}, \mathbf{z}, \mathbf{X}, \mathbf{y}) = \int p(\tilde{y}_{n+1} | \boldsymbol{\mu}, z_{n+1}, \mathbf{x}_{n+1}) p(\boldsymbol{\mu} | \mathbf{z}, \mathbf{X}, \mathbf{y}) d\boldsymbol{\mu}. \quad (1)$$

The motivating study considers a decision for a go versus no-go decision about further clinical development of a new treatment. We therefore use a utility function focused on the evaluation of enhanced treatment effects for a future patient population, evaluating expected utility by integration with respect to the posterior predictive distribution (1). We will introduce details about the probability model later. For the moment, we only require that there is such a model, but do not need to refer to details.

Our approach includes a characterization of subpopulations on the basis of discretized covariates. Taking typical sample sizes of clinical trials into account, we consider trichotomizing each covariate x_j , for $j = 1, \dots, p$. Let L, M, H generically denote trichotomized covariate values (for “low”, “medium” and “high”). We use subsets $w_j \subset \{L, M, H\}$ to describe subgroups of patients. For example, a recommended sensitive subpopulation may consist of patients who have moderate or higher value of x_j . We describe this subpopulation as $w_j = \{M, H\}$. Some minor extensions of this notation are introduced below. Our method can also handle ordered-categorical covariates, e.g., immunohistochemistry (IHC) data, without large efforts. For binary (or trinary) covariates, one may use them as they are.

3.2 Subgroups

We report a subgroup of patients who might most benefit from a new treatment, N , as

$$a = (J, \mathbf{w}), \quad (2)$$

where $J \subseteq \{1, \dots, p\}$ and $\mathbf{w} = \{w_j, j \in J\}$. We call the covariates defining a sensitive subgroup the predictive covariates. We write $\{i : \mathbf{x}_i \in a\}$ as short for $\{i : x_{i,j} \in w_j \text{ for } j \in J\}$. Here and elsewhere we use the notation $x_{i,j} \in w_j$ to indicate that the trichotomized value of $x_{i,j}$ is in $w_j \subset \{L, M, H\}$. Then $S(a) = \{i : \mathbf{x}_i \in a\}$ are patients within the subgroup selected by a .

In the present study, we consider only sensitive subgroups that are defined by one or two predictive covariates. We introduce this restriction because, three or more predictors may not be interpretable or practically useful, and also because we wish to keep the computational effort reasonable. Thus, we restrict to $J = \{j\}$ or $J = \{j, k\}$. For example, we may report that x_j is the predictive covariate and the patient subgroup $a = (j, w_j = \{M, H\})$ is sensitive to treatment N . We introduce two more special actions, including $a = H_0$ for reporting that treatment N shows no efficacy in any populations, and $a = H_1$ for reporting that treatment N is equally effective for the entire patient population. Note that we are using H_0 and H_1 here to refer to decisions. The double use of notation does not cause any ambiguities, as it is always clear from the context whether H_0 or H_1 refer to the decision or to the hypothesis.

3.3 Inference of heterogeneous treatment effects

To examine heterogeneous treatment effects depending on patient covariates, we use the notion of counter-factual models, in which two potential outcomes for each patient, one under each treatment, are introduced.

That is, we replace the single outcome y_{n+1} in (1) by a pair $(y_{n+1}(C), y_{n+1}(N))$ with $p(y(C), y(N) | \boldsymbol{\mu}, \mathbf{x}_{n+1}) = p(y_{n+1} = y(C) | z_{n+1} = 0, \boldsymbol{\mu}, \mathbf{x}_{n+1}) \cdot p(y_{n+1} = y(N) | z_{n+1} = 1, \boldsymbol{\mu}, \mathbf{x}_{n+1})$. In actual data, of course only one element of the pair is observed (Rubin, 1978; Hill, 2011).

We then define the predictive conditional treatment effect (PCTE) for patient $i = n + 1$ as

$$\text{PCTE}(\mathbf{x}_{n+1}) = E(\tilde{y}_{n+1}(N) - \tilde{y}_{n+1}(C) | \mathbf{x}_{n+1}), \quad (3)$$

where the expectation is with respect to the posterior predictive distribution (1). Positive values of PCTE indicate higher treatment effect of N for patient $n+1$ than C . For later reference, we note that this expectation may be evaluated using a simulation-based numerical approximation. Next, we define the predictive average treatment effect (PATE) in a selected subgroup by averaging the PCTE($\mathbf{x}_i$) over $S(a)$ as

$$\text{PATE}_{SS}(a) = \frac{1}{|S(a)|} \cdot \sum_{i \in S(a)} \text{PCTE}(\mathbf{x}_i). \quad (4)$$

The sum is over all observed patients in $S(a)$, implying the use of the empirical distribution to average over patients with covariates in subgroup a . We will give the details of these computations in Section 4.3. Although, in this study, our interest is on one-sided treatment effects of N as in (3), it is easy to extend the above framework to two-sided effects if desired.

4. Selecting a sensitive subgroup as a decision problem

4.1 Selection criterion

We define the following utility function to formalize preferences across possible predictive outcomes ($\tilde{y}$) and actions of subgroup reporting (a). For $a \in H_0$ we use

$$u(a, \tilde{y}) = m(\tilde{y}(N) - \tilde{y}(C)) \cdot g(|S(a)|) \cdot h(|J|), \quad (5)$$

where $m(\cdot)$ formalizes the clinical preference for larger effect sizes, and $g(\cdot)$ and $h(\cdot)$ represent preferences for larger subpopulation size and for a more parsimonious description of the subpopulation. In other words, we favor a subpopulation with a meaningful difference in the expectation of outcome between treatment arms (relative to a clinically-decided threshold, $\delta > 0$), a large population size, and a parsimonious description with few covariates. The criterion (5) does not involve $\boldsymbol{\mu}$ (but could be changed to include $\boldsymbol{\mu}$ without any substantial change to the following discussion).

Marginalizing (5) with respect to $\tilde{y}$, that is, integrating with respect to (1), defines expected utility $U(a) = \int \int u(a, \tilde{y}) p(\tilde{y} | \boldsymbol{\mu}, \mathbf{z}, \mathbf{x}) p(\boldsymbol{\mu} | \mathbf{z}, \mathbf{X}, \mathbf{y}) d\tilde{y} d\boldsymbol{\mu}$. The details of evaluating the expectation with respect to $\tilde{y}$ will be given later in Section 4.3. Also, in the statement of

$U(a)$ we now add the special case $a = H_0$. In summary, we propose the following utility function. We substitute the identity for the function $m(\cdot)$ in (5), recognize the posterior predictive integral as the PATE, use specific choices for $g(\cdot)$ and $h(\cdot)$, and add the special case for H_0 to define:

$$U(a) = \begin{cases} \{\text{PATE}_{SS}(a) - \delta\} \cdot \frac{|S(a)+1|^\phi}{(|J|+1)^\zeta} & \text{if } a \neq H_0 \\ u_0 & \text{if } a = H_0, \end{cases} \quad (6)$$

where $\text{PATE}_{SS}(a)$ is computed by (4) and the positive constants (ϕ, ζ) are tuning parameters. Later, in Section 5.2, we will show how (ϕ, ζ) can be used to achieve desired frequentist operating characteristics. To specify δ , one may solicit a minimum clinically meaningful difference in the expected outcome from clinicians. For the action $a = H_1$, we compute $\text{PATE}_{SS}(a)$ as in (4) over all n patients and evaluate $U(a)$ with $|S(a)| = n$ and $|J| = 0$. Finally, u_0 specifies the utility for the action $a = H_0$. It is a convenient tuning parameter to achieve a desired type I error rate. We will show an example in Section 5.2. For example, when $\text{PATE}_{SS}(a) < \delta$ for all subgroups, indicating that new treatment is not beneficial for any patient populations, then u_0 should be specified to a value $u_0 > U(a)$ for all $a \neq H_0$. Practically, this is achieved by selecting u_0 sufficiently large to maintain a desired type-I error rate. We discuss this process in more detail later, in Section 5.2.

Importantly, the criterion $U(a)$ in (6) is not used for model selection (although it resembles some commonly used criteria). As we shall discuss below, we use a flexible regression model for fitting the data with software that implicitly uses model averaging for optimal fit. There is no notion of model selection. Our approach separates the subgroup report a and the probability model that is used to fit the data. For example, the model might include a statistically highly significant interaction of some covariates, but we might not choose to report any subgroups characterized by those covariates if the corresponding effects are unattractive clinically (or for marketing).

4.2 Optimal selection

The optimal decision is formally described as

$$a^* = \arg \max_a U(a). \quad (7)$$

The rule a^* is known as the Bayes rule. In addition to a^* , we also recommend to report slightly suboptimal choices. Say, the top 5 instead of only the formal Bayes rule a^* in (7). This mitigates undesirable sensitivity with respect to technical details in the utility function (5) and the assumed probability model (related to technical convenience rather than expert judgment).

As described above, we aim to find one sensitive subgroup using one covariate or a set of two covariates. To achieve this goal, we trichotomize each covariate x_j , $j = 1, \dots, p$. For continuous covariates, we use the 33% and 67% quantiles (Q_j^{33}, Q_j^{67}) of x_j to define three

categories with respect to each x_j using the empirical distribution. The three categories define L, M, H . Consider first the case $|J| = 1$. We compute $\text{PATE}_{SS}(a)$ for each of the five subsets (illustrated in Figure 1): $w_j = \{L\}, \{M\}, \{H\}, \{L, M\}$, and $\{M, H\}$ with respect to every covariate $x_j, j = 1, \dots, p$. We omit the counter-intuitive choice $w_j = \{L, H\}$. It is unlikely that a sensitive subgroup consists of such distant subsets. For the two-covariate cases, $J = \{j, k\}$, we construct rectangular subgroups and L -shaped subgroups as illustrated in Figure 1. A rectangular subgroup is constructed as $\{x : x_{ij} \in w_j, x_{ik} \in w_k\}$. L -shaped subgroups are an ad-hoc way of slightly generalizing the rectangular subgroups to $\{x : x_{ij} \in w_j \text{ or } x_{ik} \in w_k\}$ (by allowing the union of $x_{ij} \in w_j$ or $x_{ik} \in w_k$ instead of the intersection). In both cases, we exclude again implausible distant subgroups with $w_j = \{L, H\}$ or $w_k = \{L, H\}$.

Some choices of the described approach are reasonable ad-hoc choices and could be replaced by alternatives. Instead of reporting the top 5 subgroups, one could report all actions with expected utility within some range of $U(a^*)$. For example, let dU denote the difference between maximum and minimum $U(a)$, over all possible actions. One could replace the top 5 by the set of all actions a with $U(a) - U(a^*) - 0.1 dU$. Similarly, the restriction to at most two covariates is an arbitrary constraint to favor parsimony and avoid excessive computational complexity. If desired, one could allow for more covariates to characterize the subgroups. In the case of three or more covariates, we would restrict to rectangular subgroups. Also, to keep computation manageable, we would restrict the search to subsets of the Bayes rule a^* with $|J| \leq 2$. Finally, the use of terciles Q^{33} and Q^{66} as thresholds could be modified. However, the choice of thresholds should be restricted to imply subsets with some minimum subsample size in the available data. Possible alternatives are medians or even quartiles. We do not further pursue such variations to avoid an unnecessarily complicated exposition.

4.3 Probability model

The proposed approach works with any probability model to summarize observed data, including any machine learning method. We use a Bayesian nonparametric method, BART (Chipman *et al.*, 2010; Hill, 2011), for $p(y | z, \mathbf{X}, \boldsymbol{\mu})$ and $p(\boldsymbol{\mu})$. BART consists of two parts, a sum-of-trees model (random forest) and a regularization prior on the parameters of the model. BART uses a random forest to define a regression mean function in $y_i = \eta_{BT}(z_i, \mathbf{x}_i, \boldsymbol{\mu}) + \varepsilon_i$, where ε_i is a normal residual $\sim i.i.d. N(0, \sigma^2)$. Here, $\boldsymbol{\mu}$ denotes generically all BART parameters including in particular terminal node parameters that represent main and interaction effects. The additive BART model with multivariate components allows to naturally incorporate not only main effects but also interaction effects, thereby giving BART excellent predictive capabilities. If needed, BART can easily be extended to deal with a binary outcome $y_i = (0 \text{ or } 1)$ using a probit model for classification, $p(y_i = 1 | z_i, \mathbf{x}_i, \boldsymbol{\mu}) = \Phi\{\eta_{BT}(z_i, \mathbf{x}_i, \boldsymbol{\mu})\}$, where $\Phi(\cdot)$ is the standard normal cumulative distribution function.

To evaluate $\text{PCTE}(\mathbf{x}_j)$ in (4) with respect to $\tilde{y}$, we use

$$E(\tilde{y}_{n+1} | z_{n+1}, \mathbf{x}_{n+1}, \mathbf{z}, \mathbf{X}, \mathbf{y}) = E\{\eta_{BT}(z_{n+1}, \mathbf{x}_{n+1}, \boldsymbol{\mu}) | \mathbf{z}, \mathbf{X}, \mathbf{y}\}, \quad (8)$$

where the expectation is with respect to the posterior distribution $p(\mu | z, X, y)$. The posterior predictive means (8) are available as part of the BART output, as implemented in the R Package `BayesTree` (Core Team: R, 2014). The details of implementing BART using `BayesTree` are given in the Appendix A. In brief, `BayesTree` requires two datasets, one with the actually observed treatment assignments z and another data set with reversed treatment assignment $(1 - z)$ to obtain posterior predictive counter-factual outcomes, as needed to evaluate $PCTE(x_i)$.

The BART model includes hyperparameters for the prior probability of a tree which is specified as $\nu(1 + d)^{-\xi}$. Here, $\nu \in (0, 1)$ and $\xi \in [0, \infty)$ denote the base and power hyperparameters, respectively, and d denotes the tree depth (using the notation in Chipman *et al.*, 2010). We use the default settings in `BayesTree` with $\nu = 0.95$ and $\xi = 2.0$. Also, we use the default number of trees = 200. We assume a conjugate normal prior for the mean parameters of the terminal nodes and assume a conjugate inverse- χ^2 prior for σ^2 (Chipman *et al.*, 2010). The choice of these conditionally conjugate priors has substantial computational benefits and is the default setting of `BayesTree`.

This completes the description of the proposed BaPoFi method. For reference we briefly summarize the method as a step by step algorithm.

Step 1. Prepare a dataset consisting of the outcome variable y , the treatment indicator z , the reversed treatment assignment $(1 - z)$, and patient baseline covariates X .

Step 2. Discretize continuous covariates.

Step 3. Specify a set of possible candidate subgroups, restricting the shape of the subgroups and the number of covariates to keep the computational effort reasonable.

Step 4. Complete the specification of the utility function (5) with specific $m(\cdot)$, $g(\cdot)$, and $h(\cdot)$ to reflect clinical and marketing preferences.

Step 5. Evaluate $PATE_{SS}(a)$ in (4) and expected utility $U(a)$ for all the candidate subgroups.

Step 6. Report the top 5 subgroups in terms of $U(a)$.

The computational time of BaPoFi to analyze data with, for example, $n = 300$ patients and $p = 10$ covariates is 1.4 seconds on an Intel Core i7-5930K CPU running at 3.50 GHz with 16 GB internal RAM.

5. Simulation study design

We validate the proposed approach using a large scale simulation study. We use a balanced design with equal (1:1) allocation to the two treatment arms (new treatment N , versus control C) in an RCT. We compare the proposed method (BaPoFi) with a Bayesian subgroup analysis based on Dixon and Simon (1991) (DS). The DS method considers the posterior distribution for subgroup-specific treatment effects under a regression model with interaction effects (details are given in the Appendix B). We simulate data according to a normal linear regression model as explained below in Section 5.1. The simulation truth highly favors the DS method over BaPoFi, since the data generating model almost matches the analysis model

of DS. We performed the simulation study with ten covariates ($p = 10$) with a maximum number of $n = 300$ patients. Finally, we carried out some additional simulation with $p = 5$ and smaller sample sizes $n = 100, 150, 200, 250$, and 300 .

5.1 Data generation

To evaluate and compare BaPoFi versus DS, we carried out 1,000 repeat simulations of hypothetical trials under eleven assumed true scenarios. We evaluated several summaries under this repeat simulation, that is, frequentist probabilities and expectations. Some of the summaries, like true negative rate, etc., require the notion of a “true subpopulation” under the simulation truth. However, the simulation truth only defines a true sampling model. Let $p^0(y | z, \mathbf{x}, \boldsymbol{\theta}^0)$ denote this simulation truth, which can be entirely different from the analysis model in (1) (see below for details). Importantly, the specification of $p^0(\cdot)$ does not include any notion of true subgroups. Instead we define a true subgroup as follows. For the BaPoFi method, let $U^0(a)$ denote the utility (6) under the simulation truth, that is, we replace the expectation in (3) by an expectation under $p^0(y | z, \mathbf{x}, \boldsymbol{\theta}^0)$. Under each simulation scenario, we computed $U^0(a)$ for all subgroups constructed by trichotomizing one or two covariates as explained in Section 4.2. The subgroup with the highest $U^0(a)$ value is the “simulation truth” for the subgroup report.

We simulate outcomes y_i under an assumed simulation truth $p^0(y_i | z_i, \mathbf{x}_i, \boldsymbol{\theta}^0)$, using eleven different scenarios for p^0 . Details are given in the supplementary materials. Briefly summarized, Scenario 1 is an overall null case (H_0) and Scenario 2 is an overall alternative case (H_1) with an overall treatment effect of arm N . As shown in Figure 2, Scenarios 3 (1-L), 4 (1-M), and 5 (1-S) are scenarios with one predictive covariate (x_1) and large (67%, scenario 3), moderate (50%, scenario 4), and small (25%, scenario 5) subpopulations. Here percent are the subgroup size as % of the entire population size. Scenarios 6 (2-M), 7 (2-S) and 8 (2-VS) are scenarios with two predictive covariates (x_1, x_2) and moderate (45%, scenario 6), small (25%, scenario 7), and very small (6%, scenario 8) subpopulations. Scenario 9 (2-M^{Ls}) is a case with two predictive covariates (x_1, x_2) and an L -shaped sensitive subpopulation of moderate size (55%).

We assume two more scenarios with true sensitive subpopulations that are neither rectangular nor L -shaped: Scenarios 10 (2-S^{nonRc}) and 11 (3-S^{nonRc}) in which those subpopulations are small-sized (25%) and defined by two (x_1, x_2) and three (x_1, x_2, x_3) predictive covariates, respectively.

The simulation truth was chosen to strongly favor the DS method, particularly for Scenarios 2 to 9, by assuming a linear model that almost matches the DS analysis model. We use the following linear model with a linear term $g(\cdot)$ to generate data,

$$y_i = g(z_i, \mathbf{x}_i, \boldsymbol{\theta}^0) + e_i, \quad e_i \sim N(0, s^2), \quad (9)$$

where $\boldsymbol{\theta}^0 = (\alpha^0, \boldsymbol{\beta}^0, \gamma^0)$ with α^0 for the overall treatment effect of arm N , $\boldsymbol{\beta}^0$ for the covariates, and γ^0 for the interaction effect. We assumed that all covariates were continuous,

and generated $x_{i,j}$ independently from a log-normal $LN(0,1)$. We set the residual variance $s^2=2^2$.

For Scenario 2 (H_1), we simulated data according to (9) with the linear term $g(z_i, \mathbf{x}_i, \alpha^0, \beta^0) = \alpha^0 z_i + h(\mathbf{x}_i, \beta^0)$ and $h(\mathbf{x}_i, \beta^0) = 0.1x_{1,i} + 0.2x_{2,i} - 0.1x_{3,i} - 0.1x_{4,i} + 0.2x_{5,i} + 0.2x_{1,i}x_{3,i}$.

For Scenarios 3 to 5, we used the linear term

$g(z_i, \mathbf{x}_i, \alpha^0, \beta^0, \gamma_1^0) = \alpha^0 z_i + h(\mathbf{x}_i, \beta^0) + \gamma_1^0 z_i I(x_{1,i} \geq Q_1)$. Here $I(\cdot)$ denotes an indicator function, and Q_j is the cutoff value for x_j under the simulation truth such that the subpopulation size is as indicated earlier. For Scenarios 6 to 8, we used the linear term

$g(z_i, \mathbf{x}_i, \alpha^0, \beta^0, \gamma_{1.2}^0) = \alpha^0 z_i + h(\mathbf{x}_i, \beta^0) + \gamma_{1.2}^0 z_i I(x_{1,i} \geq Q_1 \cap x_{2,i} \geq Q_2)$. Under Scenario 9, we replaced the interaction term with $\gamma_{1.2}^0 z_i I(x_{1,i} \geq Q_1 \cup x_{2,i} \geq Q_2)$.

For Scenario 10, we used the linear term

$g(z_i, \mathbf{x}_i, \alpha^0, \beta^0, \gamma_{1.2}^0) = \alpha^0 z_i + h(\mathbf{x}_i, \beta^0) + \gamma_{1.2}^0 z_i I(x_{1,i} + x_{2,i} \geq Q_{1.2})$. Under Scenario 11, we replaced $I(x_{1,i} + x_{2,i} \geq Q_{1.2})$ with $I(x_{1,i} + x_{2,i} + x_{3,i} \geq Q_{1.2.3})$. We specified the bounds $Q_{1.2}$ and $Q_{1.2.3}$ such that the proportion of patients with $x_{1,i} + x_{2,i} \geq Q_{1.2}$ and the proportion of patients with $x_{1,i} + x_{2,i} + x_{3,i} \geq Q_{1.2.3}$ were equal to 25%.

In addition, each scenario is also characterized by the standardized mean difference (SMD) between arms N and C in the sensitive and non-sensitive subpopulations. The SMD is defined as $SMD^S = (\alpha^0 + \gamma^0)/s$ for the sensitive subpopulation and $SMD^{NS} = \alpha^0/s$ for the non-sensitive subpopulation, respectively. The SMD^S values used in the simulation study are shown in Table 1.

In some cases, the simulation truth $p^0(\cdot)$ includes features that could naturally be called 'true subgroups'. For example, a term that involves an indicator $I(x_1 > Q_1)$ might suggest a corresponding subgroup. However, depending on the PATE and the subsample size, such a subset might or might not be of clinical and marketing interest. This highlights the separation of the subgroup report on one hand, and inference for the probability model on the other hand.

5.2 Frequentist operating characteristics

Recall that the BaPoFi method reports the top 5 subgroups with the five highest utilities $U(a)$. Letting the superscript c denote the absence of the specific report in the top 5 subgroups, we evaluated several errors, as frequentist rates. We use the following six error rates $\Pr(a | b)$, where a (a^c) indicates that a decision is (is not) among the top 5 decisions and b behind the conditioning bar refers to a simulation truth (with a slight abuse of the conditioning bar). Here $\Pr(\cdot)$ denotes frequentist probabilities under repeat simulation. Thus, for example, $\Pr(H_0 | H_0)$ is the proportion of trials under the simulation truth H_0 in which the top 5 reports contain $a = H_0$.

True decision rate (TDR)

- 1) True negative rate (TNR) = $\Pr(H_0 | H_0)$
- 2) True positive rate (TPR) = $\Pr(H_1 | H_1)$,

- 3) True subgroup rate (TSR) = $\Pr(H_a | H_a)$,

False decision rate (FDR)

- 4) False negative rate (FNR) = $\Pr(H_0 \text{ and } H_a^c | H_a)$ under H_a or $\Pr(H_0 \text{ and } H_1^c | H_1)$ under H_1 ,
- 5) False positive rate (FPR) = $\Pr(H_1 \text{ and } H_a^c | H_a)$.
- 6) False subgroup rate (FSR) = $1 - \Pr(H_a \text{ or } H_0 \text{ or } H_1 | H_a)$ or $1 - \Pr(H_0 \text{ or } H_1 | H_1)$

The 6 rates can be conveniently summarized by true decision rate (TDR) and false decision rate (FDR), as indicated. TDR unifies TNR, TPR, and TSR, with the simulation truth determining which one applies. And FDR is the sum of FNR, FSR and, in the case of a simulation truth H_a , FPR. In other words, under H_a and H_1 , FSR is the proportion of trials which report a false subgroup. Note that type I error (TIE) is included as $1 - \text{TNR} = \Pr(H_0^c | H_0)$.

We use the tuning parameters ϕ , ζ , u_0 for BaPoFi and π^* of the DS method to achieve desired levels of the described operating characteristics. The use of frequentist operating characteristics to set tuning parameters is common practice in Bayesian clinical trial design, simply because review committees almost always require summaries under repeat simulations. For example, u_0 is adjusted to achieve desired type-I error rate under a null simulation truth. In our implementation, we varied ϕ and ζ on a grid between 0.10 and 0.50 (in increments of 0.05).

After determining $(\phi, \zeta) = (0.25, 0.15)$, we set $u_0 = -3.17$. Repeat simulations were carried out with $p = 10$ and $n = 300$.

When using the approach for a real clinical study, one may specify the tuning parameters for BaPoFi, based on results of the repeat simulation using the actual numbers of patients and covariates, and the empirical sample variance of an outcome variable.

5.3 Simulation results

Table 1 summarizes simulation results with four operating characteristics for $n = 300$ and $p = 10$ under the 11 clinical scenarios. Under [Scenario 2](#), we find low TDR for DS. DS frequently fails to include (the true) H_1 among the top 5 subgroup reports. While BaPoFi and DS are almost comparable in terms of TDR under [Scenario 3 and 4](#), BaPoFi shows a more favorable TDR as SMD^S becomes larger under [Scenario 5](#). Recall that under Scenario 5 the size of the sensitive patient subpopulation was small, at 25% of the population. Under [Scenarios 6 and 7](#), DS reports desirable error rates, while BaPoFi maintains acceptable rates. Recall that Scenario 7 perfectly matches the assumed analysis model under the DS method. Similarly to Scenario 5, performance of BaPoFi improves as SMD^S increased under [Scenario 8](#), again with a very small subpopulation size (6% of the total population size). This good performance of BaPoFi to find a small-sized sensitive subpopulation may be owing to the finer partitioning with trichotomized covariates. Under [Scenarios 9 and 10](#) with the challenging L -shaped and non-rectangular simulation truths, BaPoFi continues to work reasonably well, suggesting higher flexibility of BaPoFi in subpopulation finding. Finally,

BaPoFi shows favorable TDR under Scenario 11 with the non-rectangular subgroup defined by three covariates.

Although BaPoFi reports $\text{FNR} = 0$ for Scenarios 2 to 11, this need not always be the case. We carried out additional simulations under Scenario 3, but now with $\text{SMD}^S = 0.3$ and find $\text{FNR} = 0.22$.

Figure 3 shows TDR (=TSR) values under BaPoFi and DS, respectively, with $p = 5$ and sample sizes $n = 100, 150, 200, 250$ and 300 , respectively under Scenarios 4, 5, 7, and 8. We used $p = 5$ because particularly the DS method was expected not to work well for higher dimensions, $p = 10$, and small total sample size, e.g., $n = 100$ or 150 . We evaluated TDR of the two methods with $\text{SMD}^S = 1.5$ under Scenarios 4 and 5 and with $\text{SMD}^S = 2.0$ under Scenarios 7 and 8, respectively. Under Scenario 4, BaPoFi shows higher TDR than DS for larger sample sizes. Under Scenario 7, DS outperforms BaPoFi, as expected, regardless of sample size. This exploratory investigation suggests that when the number of covariates of interests is small, say $p = 5$, a sample size greater than or equal to 200 may be required to provide at least 60% TDR to find a moderate size (say 25 to 50% of n) subpopulation. Under Scenarios 5 and 8, BaPoFi reports higher TDR, although still short of desirable levels. This simulation suggests that BaPoFi can be used to find a small-sized sensitive subpopulation under sample sizes that are commonly used in a phase II clinical trials if the enhanced treatment effect is sufficiently high. In fact, under Scenario 8 with $\text{SMD}^S = 3.0$, the BaPoFi method provided 0.62, 0.70, and 0.76 of TDR under $n = 200, 250$, and 300 , respectively.

6. Application to a phase III trial

We applied the proposed BaPoFi method to data from a randomized phase III trial to compare two combination chemotherapies, S-1 plus carboplatin (N) versus paclitaxel plus carboplatin (C), in patients with advanced non-small cell lung cancer (Okamoto *et al.*, 2010). For all patients enrolled in this phase III trial, gene expression of pharmacologically relevant genes of S-1 or paclitaxel were recorded. Specifically, 16 mRNA ($\text{mR}_1 - \text{mR}_{16}$) and 1 protein (Pn_1) expression levels were recorded, that is, a total of $p = 17$ continuous biomarkers, for $n = 267$ patients. Of them 130 were randomized to treatment C and 137 were randomized to receive treatment N. As discussed before, we used $\text{maxTS\%}$, the maximal tumor size shrinkage percentage from baseline, as the continuous outcome variable.

To implement BaPoFi, we first fix the design parameters, δ , ϕ , ζ , and u_0 . Working with a preliminary estimate of the standard deviation of $\text{maxTS\%}$ at 0.276 (based on the empirical variance in the data), we set $\delta = 0.07$ as nearly 25% of the standard deviation. We determined to use $(\phi, \zeta) = (0.25, 0.15)$ based on the earlier reported simulation study. Given $n = 267$, $p = 17$, and the other fixed design parameters, we set $u_0 = -0.039$ to control T1E at 0.05. As an ad-hoc sensitivity analysis, we also include alternative analysis using $(\phi, \zeta) = (0.35, 0.25)$ and $(0.15, 0.35)$.

Figure 4 shows the top 5 reported subgroups, in descending order from the left, using the same format as in Figure 1. Based on the frequency and order of the biomarkers that appear in the top 5, mR_6 may be the most probable sensitive biomarker for treatment N. Three

additional biomarkers, mR_2 , mR_{11} , and Pn_1 are also promising candidates to identify sensitive subpopulations for treatment N . As shown in the second row, the analysis with larger $\phi = 0.35$, reported larger subpopulations than the first row. Recalling that ϕ represents preference for a larger population size in the utility function, this is to be expected. More importantly, note that the almost same biomarkers were selected. In contrast, as shown in the bottom row, the analysis with larger $\zeta = 0.35$ (representing a stronger preference for parsimony) more frequently reports subgroups defined by a single biomarker.

In addition, it may be worth further investigating subgroups constructed by each of or a combination of these four biomarkers using other, alternative efficacy outcome data. For example, progression-free survival and overall survival, which are usually examined as the primary and secondary endpoints in cancer clinical trials, may be useful to strengthen a clinical interpretations of the analysis results.

BaPoFi subgroup reporting and the sensitivity analyses provides clinical researchers with useful information as presented above. However, results should not be blindly over-interpreted. We propose the following guidelines about the use of results in a formal Go/No-go decision. First, before starting BaPoFi analyses, summarize useful information (prior knowledge) from basic research and previous clinical trials to draw candidate biomarkers. If BaPoFi analyses consistently select sensitive subpopulations on the basis of one or several promising biomarker(s), and those are consistent with prior knowledge, then consider designing a new clinical trial in the selected subpopulation. If, on the other hand, the BaPoFi analysis suggests the entire population as optimal subgroup ($a = H_1$), then consider going to phase III for the overall population. If there is substantial variation in selected biomarkers and subgroups in additional sensitivity analyses, but some consistency with prior information is observed, it may not be possible to simply derive a final conclusion (this is the most difficult case). One may take into account results from additional analyses of alternative efficacy outcomes in seemingly promising subgroups selected by BaPoFi analyses. Finally, if no consistent result is obtained within BaPoFi analyses and there is a conflict with prior information, consider no further development.

7. Discussion

The proposed BaPoFi method addresses the population-finding problem by separating the specification of an action space versus the construction of a flexible probability model for the outcomes. We used BART as a flexible probability model for the outcomes. However, any other flexible probability model could be used. Similarly, instead of the rectangular subgroups that we (mainly) used, one could allow for patient subgroups with any other shapes, such as the L -shaped subgroup defined in the simulation scenario 9. We reported the top 5 subgroups rather than the strictly optimal subgroup to mitigate undesirable sensitivity with respect to technical details in the utility function. However, clinical investigators may need more focused recommendations. In order to derive a final conclusion on the subpopulation-finding, it may be vital to take prior medical knowledge into consideration.

Some more limitations remain. A decision theoretic approach often depends on the arbitrary choice of design parameters in a utility function. Controlling the type I error rate and other

frequentist rates mitigates the problem. In the present study, we restrict the search to at most two covariates and rectangular-type subgroups to keep computation manageable. However, interesting subpopulations might be more complicated in terms of dimension and/or shape, including, for example, non-rectangular subgroups defined by more than two covariates.

Although we focus on an RCT comparing two treatment arms, extensions to accommodate observational studies or clinical trials with more than two arms are possible. In observational studies, the potential outcomes $\tilde{y}_{n+1}(N)$ and $\tilde{y}_{n+1}(C)$ in (3) may be dependent with treatment assignment z due to baseline confounding covariates (Hill, 2011). Letting $\mathbf{v}$ denote a vector of confounding covariates, additional conditioning on $\mathbf{v}$ in (3) could be used in non-randomized study settings. Another generalization is to studies with more than two treatment arms, for example, a three-arm clinical trial comparing two new treatments, N_1 and N_2 , with control. Defining potential outcomes $\tilde{y}_{n+1}(C)$, $\tilde{y}_{n+1}(N_1)$, and $\tilde{y}_{n+1}(N_2)$ and extending (3) to evaluate the two comparisons $E(\tilde{y}_{n+1}(N_1) - \tilde{y}_{n+1}(C) \mid \mathbf{x}_{n+1})$ and $E(\tilde{y}_{n+1}(N_2) - \tilde{y}_{n+1}(C) \mid \mathbf{x}_{n+1})$ generalizes BaPoFi to three-arm trials.

Another limitation is that the described approach does not deal with missing (not available) biomarker data as it may occur in real data analyses and it does not consider adaptive randomization and/or interim analysis during a trial. But the current method could be easily extended to deal with such modifications if desired. Finally, the performance of BaPoFi depends on the cutoff points to discretize (trichotomize) covariates. Determining optimal cutoff points needs larger sample sizes and is beyond the scope of a phase II clinical trial.

8. Supplementary Materials

Supplementary materials referenced in Section 5 is available with this article at the *Biometrics* website on Wiley Online Library. The R-code for BaPoFi and sample data are also available at the Biometrics website on Wiley Online Library.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Drs. Isamu Okamoto and Kazuto Nishio for providing us with the precious data. Satoshi Morita's research was partially supported by a Grant-in-Aid for Scientific Research C-24500345 from the Ministry of Health, Labour and Welfare of Japan. Peter Müller's research was supported by NIH/NCI 1 R01 CA207195-01. We thank the Editor, the Associate Editor, and the three referees for their thoughtful and constructive comments and suggestions.

References

- Chipman HA, George EI, McCulloch RE. Bayesian CART Model Search. *The Journal of the American Statistical Association*. 1998; 93:935–948.
- Chipman HA, George EI, McCulloch RE. BART: Bayesian additive regression trees. *The Annals of Applied Statistics*. 2010; 4:266–298.
- Core Team: R. A language and environment for statistical computing. R Foundation for Statistical Computing; Vienna, Austria: 2014. <http://www.R-project.org/>
- Denison DGT, Mallick BK, Smith AFM. A Bayesian CART algorithm. *Biometrika*. 1998; 85:363–377.

- Dixon DO, Simon R. Bayesian Subset Analysis. *Biometrics*. 1991; 47:871–881. [PubMed: 1742443]
- Foster JC, Taylor JM, Ruberg SJ. Subgroup identification from randomized clinical trial data. *Statistics in Medicine*. 2011; 30:2867–2880. [PubMed: 21815180]
- Friedman JH. Multivariate adaptive regression splines. *The Annals of Statistics*. 1991; 19:1–67.
- Gilks, W., Richardson, S., Spiegelhalter, D. *Markov Chain Monte Carlo in Practice*. Chapman & Hall; London: 1996.
- Hill JL. Bayesian Nonparametric Modeling for Causal Inference. *Journal of Computational and Graphical Statistics*. 2011; 20:217–240.
- Jain RK, Lee JJ, Ng C, Hong D, Gong J, Naing A, et al. Change in tumor size by RECIST correlates linearly with overall survival in phase I oncology studies. *Journal of Clinical Oncology*. 2012; 30:2684–2690. [PubMed: 22689801]
- Jones HE, Ohlssen DI, Neuenschwander B, Racine A, Branson M. Bayesian models for subgroup analysis in clinical trials. *Clinical Trials*. 2011; 8:129–143. [PubMed: 21282293]
- Karrison TG1, Maitland ML, Stadler WM, Ratain MJ. Design of phase II cancer trials using a continuous endpoint of change in tumor size: application to a study of sorafenib and erlotinib in non small-cell lung cancer. *The Journal of the National Cancer Institute*. 2007; 99:1455–1461. [PubMed: 17895472]
- Lipkovich I, Dmitrienko A, Denne J, Enas G. Subgroup identification based on differential effect search—a recursive partitioning method for establishing response to treatment in patient subpopulations. *Statistics in Medicine*. 2011; 30:2601–2621. [PubMed: 21786278]
- McShane LM, Hunsberger S, Adjei AA. Effective incorporation of biomarkers into phase II trials. *Clinical Cancer Research*. 2009; 15:1898–1905. [PubMed: 19276274]
- Okamoto I, Yoshioka H, Morita S, Ando M, Takeda K, Seto T, et al. Phase III trial comparing oral S-1 plus carboplatin with paclitaxel plus carboplatin in chemotherapy-naïve patients with advanced non-small-cell lung cancer: results of a west Japan oncology group study. *Journal of Clinical Oncology*. 2010; 28:5240–5246. [PubMed: 21079147]
- Rasmussen, CE., Williams, CKI. *Gaussian Processes for Machine Learning*. The MIT Press; Cambridge: 2006.
- Renfro LA, Mallick H, An MW, Sargent DJ, Mandrekar SJ. Clinical trial designs incorporating predictive biomarkers. *Cancer Treatment Reviews*. 2016; 43:74–82. [PubMed: 26827695]
- Robert, CP. *The Bayesian Choice*. Springer-Verlag; New York: 2007.
- Ruberg SJ, Chen L, Wang Y. The mean does not mean as much anymore: finding sub-groups for tailored therapeutics. *Clinical Trials*. 2010; 7:574–583. [PubMed: 20667935]
- Rubin DB. Bayesian Inference for Causal Effects: The Role of Randomization. *The Annals of Statistics*. 1978; 6:34–58.
- Seymour L, Ivy SP, Sargent D, Spriggs D, Baker L, Rubinstein L, et al. The design of phase II clinical trials testing cancer therapeutics: consensus recommendations from the clinical trial design task force of the national cancer institute investigational drug steering committee. *Clinical Cancer Research*. 2010; 16:1764–1769. [PubMed: 20215557]
- Simon R. Bayesian subset analysis: application to studying treatment-by-gender interactions. *Statistics in Medicine*. 2002; 21:2909–2916. [PubMed: 12325107]
- Sivaganesan S, Laud PW, Müller P. A Bayesian subgroup analysis with a zero-enriched Polya Urn scheme. *Statistics in Medicine*. 2011; 30:312–323. [PubMed: 21225894]
- Sivaganesan S, Müller P, Huang B. Subgroup finding via Bayesian additive regression trees. 2016 (submitted).
- Tsuburaya A, Morita S, Kodera Y, Kobayashi M, Shitara K, Yamaguchi K, et al. A randomized phase II trial to elucidate the efficacy of XP and SP as the first-line treatment for advanced gastric cancer: XP ascertainment vs. SP randomized PII (XParTS II) trial. *BMC Cancer*. 2012; 12:307. [PubMed: 22824079]
- Xu Y, Müller P, Wahed AS, Thall PF. Bayesian nonparametric estimation for dynamic treatment regimes with sequential transition times. *The Journal of the American Statistical Association*. 2016; 111:921–950. [PubMed: 28018015]

Appendix A

Implementation of BART using BayesTree

Using the R Package, BayesTree (Core Team: R, 2014), run the function `bart(x.train = xt, y.train=y, x.test=xp)` with `xt` containing z and all X , that is $[z, X]$, and `xp` containing reversed treatment assignment data $[1-z, X]$. In this way, using `x.test = xp`, obtain the predicted values for the potential outcomes that would be observed under the reversed treatment assignment. Then, compare the predicted outcome values using `x.train = xt` and those using `x.test = xp` to obtain $E\{\eta_{BT}(z_i, \mathbf{x}_i, \boldsymbol{\mu}) | z, X, y\}$ and $E\{\eta_{BT}(1 - z_i, \mathbf{x}_i, \boldsymbol{\mu}) | z, X, y\}$ in (8), respectively, for patient $i = 1, \dots, n$, where the expectation is with respect to the posterior distribution $p(\boldsymbol{\mu} | z, X, y)$. Finally, compute the average of the after burn-in MCMC samples (Gilks *et al.*, 1996), $\eta_{BT}^{[1]}(z=1, \mathbf{x}_i, \boldsymbol{\mu}), \dots, \eta_{BT}^{[L]}(z=1, \mathbf{x}_i, \boldsymbol{\mu})$ and that of $\eta_{BT}^{[1]}(z=0, \mathbf{x}_i, \boldsymbol{\mu}), \dots, \eta_{BT}^{[L]}(z=0, \mathbf{x}_i, \boldsymbol{\mu})$ for sufficiently large L , to approximate $E\{\tilde{y}_i(N) | \mathbf{x}_i, z, X, y\}$ and $E\{\tilde{y}_i(C) | \mathbf{x}_i, z, X, y\}$, respectively. These average values are obtained as the BART outputs: `yhat.tain.mean` and `yhat.test.mean`, respectively.

Appendix B

Dixon and Simon's method

Suppressing the patient index i , let x_j^d denote the dichotomized covariate, dichotomized at the median value, coding 0 for lower or 1 for higher than the median value. Referring to Eq. (1.2) on page 872 in Dixon and Simon (1991) (DS), we used the following linear model with the vector of model parameters $\boldsymbol{\mu}_{DS} = (\nu, \tau, \xi_1, \dots, \xi_p, v_{1,1}, \dots, v_{1,p}, v_{2,1,2}, \dots, v_{2,(p-1),p})$, and treatment indicator z to analyze data,

$$E(y | z, \mathbf{x}^d, \boldsymbol{\mu}_{DS}) = \nu + z\tau + \sum_j \left(x_j^d \xi_j + z x_j^d v_{1,j} \right) + \sum_{j,k \text{ for } j < k} z x_j^d x_k^d v_{2,j,k}, \quad (\text{A1})$$

where $v_{1,j}$ and $v_{2,j,k}$ are the model parameters for the interaction between the treatment and covariate $x_j^d, j = 1, \dots, p$, and the interaction between the treatment and a combination of two covariates $(x_j^d, x_k^d), 1 \leq j < k \leq p$, respectively. Let $E(y | \boldsymbol{\mu}_{DS}, z, a)$ denote $E(y | z, \mathbf{x}^d, \boldsymbol{\mu}_{DS})$ averaged with respect to the distribution of the other covariate levels in the subgroup described by a , as recommended in Dixon and Simon (1991) (page 878). The weight reflect the relative frequency of patients with a particular character determined by the one or two covariates. Using (A1), we define the treatment effect as

$$\psi(\boldsymbol{\mu}_{DS}, a) = E(y | \boldsymbol{\mu}_{DS}, z=1, a) - E(y | \boldsymbol{\mu}_{DS}, z=0, a) \quad (\text{A2})$$

with respect to a for all the subgroups constructed by one or two dichotomized covariates ($a=H_a$) and for the entire population of patients ($a=H_1$). Referring to Dixon and Simon

(1991), we evaluated the posterior probability, $PP_{DS}(a) = p(\psi(\mu_{DS}, a) > \delta \mid z, y, \mathbf{X})$ to summarize the analysis results. Recall that δ denotes the efficacy threshold. We report a as a sensitive subgroup if $PP_{DS} > \pi^*$.

For the simulation truth, we use a different definition for DS from that for BaPoFi. This is because DS uses subgroups indexed by dichotomized covariates only. Let $\psi^0(a)$ denote the expected treatment effect (A2) with respect to $p^0(y \mid z, \mathbf{x}, \boldsymbol{\theta}^0)$. Under each scenario, the subgroup with the highest $\psi^0(a)$ value, among all possible subgroup reportings based on the dichotomized covariates, is the “simulation truth” for the subgroup report under DS. If multiple subgroups had the same highest value, we selected the subgroup with the largest patient-population size. The use of slightly different “true subgroup” constructions for the two methods allows for a fair comparison, as DS is constructed to report subgroups based on dichotomized variables, and uses a different criterion than BaPoFi. Similarly to the BaPoFi method, we used the “top 5 reports” also in the DS method, reporting the ones with the five highest PP_{DS} among all subgroups with $PP_{DS} > \pi^*$.

Figure 1.

Illustrations of trichotomizing one (x_1) or two (x_1, x_2) covariates to construct patient subgroups with their 33% and 67% point cutoff values denoted by Q_j^{33} and Q_j^{67} , $j = 1, 2$, respectively. Denoting $w(x_j) \subset \{L, M, H\}$, where the levels L, M , and H indicate the subgroups of patients with low, moderate, and high values of x_j , respectively, areas in shadow show the constructed subgroups.

Figure 2.

Clinical scenarios for the simulation study. The new treatment provides a higher efficacy in the area in shadow than the rest, that is, the area shows the interaction effect between treatment and covariates (x_1 , x_2 , x_3), under each scenario. Scenario 1 is the overall null case (H_0), while Scenario 2 is overall alternative case (H_1).

Figure 3.

True decision rates (= true subgroup rates) for the BaPoFi (solid line) and DS (dashed line) methods with the number of covariates $p = 5$ for the total sample size values $n = 100, 150, 200, 250, 300$, under Scenarios 4, 5, 7, and 8. $\text{SMD}^S = 1.5$ under Scenarios 4 and 5, and $\text{SMD}^S = 2.0$ under Scenarios 7 and 8.

Figure 4.

Application of the BaPoFi method to a phase III clinical trial, reporting the top 5 subgroups with $(\phi, \zeta) = (0.25, 0.15)$, $(0.35, 0.25)$, and $(0.15, 0.35)$.

Table 1

Simulation results of comparison between the BaPoFi and DS methods with ($n=$) 300 patients and ($p=$) 10 covariates, summarized in terms of true decision rate (TDR) and false decision rates (FDRs). TDR unifies true negative rate (TNR) under Scenario 1, true positive rate (TPR) under Scenario 2, and true subgroup rate (TSR) under the others. FDRs represent false negative rate (FNR), false positive rate (FPR), and false subgroup rate (FSR). The simulation scenarios are characterized by the number of predictive covariates (0, 1, 2, 3), the size (L: large, M: moderate, S: small, VS: very small), and shape (rectangular, L -shaped (L_s), *non-rectangular* (*nonRect*)) of sensitive subpopulation, and the standardized mean differences between new and control treatments in the sensitive subpopulation, SMD^{δ} . SMD in the non-sensitive subpopulation is set at 0.5 for Scenarios 2 to 11. SMD taking a higher value than 1 indicates that the treatment effect of N beyond the efficacy threshold.

Scenario	SMD^{δ}	Method	FDRs		
			TDR	FNR	FPR
1 H_0	0	BaPoFi	0.95	-	-
2 H_1	1.5	DS	0.95	-	-
		BaPoFi	0.99	0	0.01
3 1-L	1.5	DS	0.65	0.10	-
		BaPoFi	0.66	0	0.07
4 1-M	1.5	DS	0.67	0.19	0
		BaPoFi	0.68	0	0.01
5 1-H	1.5	DS	0.67	0.18	0
		BaPoFi	0.13	0	0.87
6 2-M	2.0	DS	0.02	0.77	0
		BaPoFi	0.63	0	0.37
7 2-S	2.0	DS	0.22	0.46	0
		BaPoFi	0.54	0	0.12
8 2-VS	2.0	DS	0.78	0.08	0
		BaPoFi	0.60	0	0.01
9 2-M ^{L_s}	2.0	DS	0.91	0.09	0
		BaPoFi	0.22	0	0.78
	3.0	DS	0.03	0.94	0
		BaPoFi	0.48	0	0.52
		DS	0.09	0.85	0
		BaPoFi	0.29	0	0.14

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Scenario	SMD ^s	Method	TDR	FDRs		
				FNR	FPR	FSR
10	2.0	DS	0	0.06	0.02	0.92
		BaPoFi	0.62	0	0.01	0.37
		DS	0.21	0.56	0	0.23
11	2.0	BaPoFi	0.73	0	0.01	0.26
		DS	0.32	0.58	0	0.10