

This is a self-archived version of an original article. This version may differ from the original in pagination and typographic details.

Author(s): Reinikainen, Jaakko; Karvanen, Juha

Title: Bayesian subcohort selection for longitudinal covariate measurements in follow-up studies

Year: 2022

Version: Published version

Copyright: © 2022 The Authors. Statistica Neerlandica published by John Wiley & Sons Ltd on

Rights: CC BY 4.0

Rights url: <https://creativecommons.org/licenses/by/4.0/>

Please cite the original version:

Reinikainen, J., & Karvanen, J. (2022). Bayesian subcohort selection for longitudinal covariate measurements in follow-up studies. *Statistica Neerlandica*, 76(4), 372-390.
<https://doi.org/10.1111/stan.12264>

Bayesian subcohort selection for longitudinal covariate measurements in follow-up studies

Jaakko Reinikainen^{1,2}  | Juha Karvanen²

¹Finnish Institute for Health and Welfare, Helsinki, Finland

²Department of Mathematics and Statistics, University of Jyväskylä, Jyväskylä, Finland

Correspondence

Jaakko Reinikainen, Finnish Institution for Health and Welfare, Helsinki, Finland.
Email: jaakko.reinikainen@thl.fi

Funding information

Emil Aaltosen Säätiö

Abstract

We propose an approach for the planning of longitudinal covariate measurements in follow-up studies where covariates are time-varying. We assume that the entire cohort cannot be selected for longitudinal measurements due to financial limitations, and study how a subset of the cohort should be selected optimally, in order to obtain precise estimates of covariate effects in a survival model. In our approach, the study will be designed sequentially utilizing the data collected in previous measurements of the individuals as prior information. We propose using a Bayesian optimality criterion in the subcohort selections, which is compared with simple random sampling using simulated and real follow-up data. Our work improves the computational approach compared to the previous research on the topic so that designs with several covariates and measurement points can be implemented. As an example we derive the optimal design for studying the effect of body mass index and smoking on all-cause mortality in a Finnish longitudinal study. Our results support the conclusion that the precision of the estimates can be clearly improved by optimal design.

KEYWORDS

Bayesian optimal design, data collection, follow-up study, longitudinal measurements, study design

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Statistica Neerlandica* published by John Wiley & Sons Ltd on behalf of Netherlands Society for Statistics and Operations Research.

1 | INTRODUCTION

Longitudinal covariate measurements are often carried out in follow-up studies when the covariates are time varying. These measurements give useful information about the trajectories of the covariates. Frequent remeasurements provide more information than infrequent ones, but in practice, limited resources may restrict the number of measurements and researchers have to consider how to design the study cost efficiently.

We study optimal design in a scenario where we cannot afford to remeasure the entire cohort but can only select a subset of the cohort, called a subcohort. The goal is to estimate the effects of the covariates on survival as precisely as possible. The study is designed sequentially, which here means that the subcohorts are selected just before the measurement times and all information collected prior to a new measurement is utilized. This kind of design procedure is realizable if researchers can clearly define beforehand the purpose of data collection, that is, the parameters of interest to be estimated from the data. In addition, the outcome data must already be available during the follow-up. The proposed method requires especially that up-to-date survival information can be obtained when needed. This is possible, for example, in Finland, where data on mortality and hospitalizations are available from administrative registries.

Use of a Bayesian version of D_s -optimality, an optimality criterion based on Fisher information, is proposed for the selection of the subcohort. In addition to the D_s -criterion, there are also other optimality criteria, which were originally developed for design of experiments (Atkinson, Donev, & Tobias, 2007; Pukelsheim, 1993), but can also be applied in observational studies. For example, Karvanen, Kulathinal, and Gasbarra (2009) considered optimal subset selection for genotyping in a follow-up study, Buzoianu and Kadane (2009) investigated selection of patients for a diagnostic test and Mehtälä, Auranen, and Kulathinal (2015) studied optimal time spacings for observations of a multistate Markov process.

Reinikainen, Karvanen, and Tolonen (2016) studied the same problem with a frequentist approach that was restricted to one time-varying covariate and one remeasurement after the baseline. The limitation of this approach is that there is no straightforward way to generalize it to more realistic scenarios where several covariates and measurement points are allowed. The Bayesian approach presented in this paper offers flexibility that makes it possible to handle these challenges.

In a multivariate setting with multiple longitudinal measurements, several computational challenges are faced. Because of the very large discrete design space, it is practically impossible to evaluate all possible designs. Thus, some heuristic method is needed to find an approximate solution. Bayesian analysis with Markov chain Monte Carlo (MCMC) offers flexibility in the modeling, but requires computational resources. Selecting only subcohorts for remeasurement creates a large amount of missing data, which further complicates the data analysis.

The data missing by design have previously (Reinikainen et al., 2016) been handled with multiple imputation and a likelihood-based approach with numerical integration. These approaches do not look promising for the generalized problem because multiple imputation of covariates conditioned on survival data would be complicated and numerical integration would become infeasible as increasing the number of covariates increases the dimension of the integral. Here, we handle missing data by using Bayesian data augmentation, which is expected to be a more flexible method when the number of measurement points and covariates is increased. Parallel computing is utilized in a simulation study.

The optimal subcohort selection is studied using simulated and real data. The real data consist of the Finnish cohorts from the Seven Countries Study (Keys, 1970), an international epidemiologic study characterized by a long follow-up time and several longitudinal covariate measurements. We use body mass index and smoking as time-varying covariates and all-cause mortality as the outcome. With these data, the proposed selection procedure is compared with simple random sampling (SRS) of individuals to be remeasured and with a case where the entire cohort is selected for remeasurement.

2 | SURVIVAL MODEL

In this section, we introduce the notation for our study design and survival model. These are later used to present our optimal subcohort selection procedure. Many of the following assumptions are made to be suitable for the real data example of Section 6, but the general idea is applicable to other designs and models as well.

Let us consider a follow-up study in which survival time is the response variable and M longitudinal measurements are carried out for the time-varying covariates after the baseline measurement. We denote the covariate values by x_{mjh} for the measurement $m = 0, \dots, M$, the individual $j = 1, \dots, N$ and the covariate $h = 1, \dots, H$. The corresponding random variables are denoted by X_{mjh} and for all the covariates and individuals shortly by X_m . The measurements for the covariates are carried out at time points τ_0, \dots, τ_M in calendar time for the individuals who are alive and have been selected to be measured. The individuals in the cohort may be of different ages at the time of the baseline measurement and the start of the follow-up τ_0 . The follow-up has a predetermined length ending at the time τ_{M+1} .

The exact survival times are available from a registry. The observed survival information at the time of m th remeasurement for the individual j is denoted by $y_{mj} = (t_{mj}, \delta_{mj})$, where t_{mj} is the continuously measured survival time with age as the time scale and δ_{mj} is the status indicator telling the vital status at time τ_m ($\delta_{mj} = 1$ for an event and $\delta_{mj} = 0$ for censoring). Thus, if $\delta_{mj} = 1$, t_{mj} is the age at the event and if $\delta_{mj} = 0$, t_{mj} is the age at τ_m .

Although the survival time is observed continuously, the piece-wise modeling approach uses separate time and status variables for each measurement time interval. The individual j has time-to-event variables t_{1j}, t_{2j}, \dots for each part of the follow-up where they are still alive. We denote the random variables related to survival information by Y_{mj} and for all the individuals by Y_m .

We use notation $\mathbf{x}_m = (x_{m1}, \dots, x_{mH})^T$ and $\boldsymbol{\beta} = (\beta_1, \dots, \beta_H)^T$ and assume that the covariate effects fulfill the Markov property

$$\lambda(t_{m+1} | \mathbf{x}_0, \dots, \mathbf{x}_m, \delta_m = 0) = \lambda(t_{m+1} | \mathbf{x}_m, \delta_m = 0).$$

We continue by assuming that the covariates are related to the hazard of the event through the proportional hazards model

$$\lambda(t_{m+1} | \mathbf{x}_m, \delta_m = 0) = \lambda_0(t_{m+1} | \delta_m = 0) \exp(\boldsymbol{\beta}^T \mathbf{x}_m). \quad (1)$$

Conditioning on $\delta_m = 0$ means that only those individuals contribute here who have not died before the m th remeasurement. As the covariate information is updated during the follow-up, also the hazard of an individual is assumed to change correspondingly. The survival times are

assumed to follow the Weibull distribution, when the baseline hazard function has the form

$$\lambda_0(t_{m+1}|\delta_m = 0) = \frac{a}{b} \left(\frac{t_{m+1}}{b} \right)^{a-1},$$

where a is the shape parameter and b is the scale parameter. Then model (1) becomes a parametric form of the time-dependent Cox model (Therneau & Grambsch, 2000). Other distributions than the Weibull could also be used for the survival times.

Now, we can write the survival function and the density function as

$$\begin{aligned} S(t_{m+1}|\mathbf{x}_m, \delta_m = 0) &= S_0(t_{m+1}|\delta_m = 0)^{\exp(\beta^T \mathbf{x}_m)} \text{ and} \\ p(t_{m+1}|\mathbf{x}_m, \delta_m = 0) &= \lambda(t_{m+1}|\mathbf{x}_m, \delta_m = 0)S(t_{m+1}|\mathbf{x}_m, \delta_m = 0), \end{aligned}$$

where $S_0(t_{m+1}|\delta_m = 0)$ is the baseline survival function. The modeling is carried out piecewisely in time, because the covariate information changes at the measurement points. For this reason we have to deal with left-truncated and possibly right-censored Weibull distributions in the time intervals $(\tau_0, \tau_1]$, $(\tau_1, \tau_2]$, Survival times are left-truncated at the lower limit of a time interval by scaling the likelihood with $S(t_{mj}|\mathbf{x}_{mj}, \delta_m = 0)$ because an observed survival time $t_{m+1,j}$ cannot be smaller than t_{mj} . This is a sequential conditioning where the likelihood is scaled with the survival probability at age t_{mj} . In addition, left-truncation is necessary because the age of an individual is used as the time scale and we do not assume that the follow-up would begin at time zero, which would be the time of birth. Thus, we have to take into account that persons who have died before the start of the follow-up did not have an opportunity to be included in the cohort. The likelihood contribution for the individual j for the parameters $\beta_1, \dots, \beta_H, a$ and b is

$$L_j(\beta_1, \dots, \beta_H, a, b) = \prod_{m=0}^{m'_j} \left(\frac{p(t_{m+1,j}|\mathbf{x}_{mj}, \delta_m = 0)}{S(t_{mj}|\mathbf{x}_{mj}, \delta_m = 0)} \right)^{\delta_{m+1,j}} \left(\frac{S(t_{m+1,j}|\mathbf{x}_{mj}, \delta_m = 0)}{S(t_{mj}|\mathbf{x}_{mj}, \delta_m = 0)} \right)^{1-\delta_{m+1,j}}, \quad (2)$$

where $m'_j = \max\{0, \dots, M : \delta_{mj} = 0\}$.

3 | OPTIMAL SUBCOHORT SELECTION

If the entire cohort cannot be remeasured because of financial limitations, we have to select a subcohort, which we can afford to measure. The optimal selection aims to make the estimates of the parameters of interest as precise as possible subject to financial constraints. In this paper, we focus on the estimation of regression parameters β .

Assume that baseline covariate measurements (and possibly some longitudinal measurements) have been carried out and that continuous survival information can be obtained during the study for all individuals. When we want to carry out the next longitudinal measurement for a subcohort, we proceed by taking the following general steps:

1. Just before the new longitudinal measurement, use the data already collected to obtain prior information about the parameters of interest
2. Define the optimality criterion as an expectation over the prior distribution
3. Maximize the optimality criterion and select an optimal subcohort for the remeasurement
4. Remeasure the covariates for the selected subcohort

The rest of this section and Section 4.1 provides a description of one possible way to perform steps 2. and 3. Section 4.2 focuses on step 1.

3.1 | Bayesian optimal design in a general form

An optimal design problem can be seen as a problem of maximizing the expected utility $U(\xi)$ for a design ξ from a design space Ξ (Chaloner & Verdinelli, 1995). See the seminal paper by Lindley (1956) on Bayesian optimal design and Woods, Overstall, Adamou, and Waite (2017) for a recent review on the topic. In our problem, ξ is an indicator matrix with individuals on rows and measurement times on columns, where an element (j, m) is 1 if individual j has been selected for measurement m and 0 otherwise. The constraint of limited resources means here that the column sums are fixed in ξ . The column sums need not be the same, because we may have different amount of resources for different re-examinations.

Data $\mathbf{w} = (\mathbf{x}, \mathbf{y})$, where \mathbf{x} corresponds to covariate data and \mathbf{y} to survival outcome, come from a sample space \mathcal{W} . The outcome data \mathbf{y} and the baseline covariate measurements are assumed to be available on all individuals, whereas longitudinal covariate data is collected according to the design ξ . The data are assumed to follow a model $p(\mathbf{w}|\theta)$, where parameters θ belong to the parameter space Θ .

The fully Bayesian solution for the optimal design problem would involve integrating a measure of observed utility over data \mathbf{w} and the posterior distribution of parameters θ . Instead of this, we use a common approach (Atkinson et al., 2007; Chaloner & Verdinelli, 1995) where the integration is done over the prior density distribution of θ and the utility is defined as a function of the expected information. With this notation, the expected utility can be written as

$$U(\xi) = \int_{\Theta} g \left[E_{\mathbf{w}|\theta, \xi} \{ I_{\mathbf{w}}(\theta) \} \right] p(\theta) d\theta, \quad (3)$$

where g is a function such as determinant (D-optimality), $E_{\mathbf{w}|\theta, \xi} \{ I_{\mathbf{w}}(\theta) \}$ is an expectation of the Fisher information of the parameters θ over data \mathbf{w} and $p(\theta)$ is the prior density distribution of θ .

Here, we use the D_s -criterion (Atkinson et al., 2007), which is a special case of the widely used D-optimality criterion. The D-optimal design maximizes the determinant of the Fisher information matrix or equivalently minimizes the determinant of the covariance matrix. D_s -optimality considers only a subset of s parameters. If the parameter vector $\theta = (\theta_1, \dots, \theta_s, \dots, \theta_p)^T$ includes first the s parameters of interest and then $p - s$ nuisance parameters, D_s -optimal design minimizes the determinant of the $s \times s$ upper left submatrix of $I(\theta)^{-1}$. In our case, s is the number of β -parameters in the survival model, and thus we will call the criterion D_{β} -optimality.

3.2 | Selection criterion for the Weibull proportional hazards model

Let us consider the Fisher information matrix of the model introduced in Section 2 including parameters $\theta^* = (\beta_1, \dots, \beta_H, a, b)$:

$$I_{X,Y}(\theta^*) = -E \left(\frac{\partial^2 \log p(X_0, \dots, X_M, Y_1, \dots, Y_{M+1})}{\partial \theta^{*2}} \right), \quad (4)$$

where $Y_m = (T_m, \delta_m)$. The parameter vector θ^* includes the parameters of the survival model but not the parameters of covariate processes. We assume Markov properties $p(Y_{m+1}|X_0, \dots, X_m, Y_1, \dots, Y_m) = p(Y_{m+1}|X_m, Y_m)$ and $p(X_m|X_0, \dots, X_{m-1}, Y_1, \dots, Y_m) = p(X_m|X_{m-1})$, and decompose the logarithmic joint distribution as $\log p(X_0, \dots, X_M, Y_1, \dots, Y_{M+1}) = \log p(X_0) + \log p(Y_1|X_0) + \dots + \log p(X_M|X_{M-1}) + \log p(Y_{M+1}|X_M, Y_M)$, which allows us to decompose $I_{X,Y}(\theta^*)$ similarly. The Markov assumption is made here to avoid unnecessarily complicating the description of the method. Other dependency structures could be assumed as well, depending on the data for which the method is applied.

When we are selecting individuals for the m th remeasurement, we have measured covariates X_0, \dots, X_{m-1} , which may include missing values and survival information is observed up to the time of the m th remeasurement, that is, Y_m is known. Therefore, the selection for the m th remeasurement is based on the expectations of X_m and Y_{m+1} utilizing the previously observed data. Only those individuals who have not yet died or been censored can be considered as candidates for the remeasurement. That is, optimization is carried out by selecting individuals who are still in the risk set for further examination. Individuals whose covariates have not been measured previously are also candidates for the remeasurement. Due to the Markov assumptions, it is sufficient that the selection is based only on the expectations with respect to the next unobserved part and not of all the forthcoming parts of the follow-up.

Now, using the above-mentioned decomposition in matrix (4), the information matrix used in the selection for the m th remeasurement can be written as

$$I_{X,Y}^m(\theta^*) = -E \left[\frac{\partial^2}{\partial \theta^{*2}} \left\{ \log p(X_0) + \log p(Y_1|X_0) + \dots \right. \right. \\ \left. \left. + \log p(X_{m-1}|X_{m-2}) + \log p(Y_m|X_{m-1}, Y_{m-1}) \right\} \right. \\ \left. + E \left\{ \left(\frac{\partial^2}{\partial \theta^{*2}} \log p(X_m|X_{m-1}) + \frac{\partial^2}{\partial \theta^{*2}} \log p(Y_{m+1}|X_m, Y_m) \right) \middle| X_0, \dots, X_{m-1}, Y_m \right\} \right].$$

In the presentation of this matrix, the last term includes unobserved values, which is why an additional conditional expectation is taken on it. Above, the terms which do not include variable Y vanish, because they do not include the survival model parameters θ^* (parameters of the Weibull proportional hazards model). This leads to

$$I_{X,Y}^m(\theta^*) = E \left[-\frac{\partial^2}{\partial \theta^{*2}} \left\{ \log p(Y_1|X_0) + \dots + \log p(Y_m|X_{m-1}, Y_{m-1}) \right\} \right] \\ + E \left[E \left\{ -\frac{\partial^2}{\partial \theta^{*2}} \log p(Y_{m+1}|X_m, Y_m) \middle| X_0, \dots, X_{m-1}, Y_m \right\} \right] \\ = I_{Y_1|X_0}(\theta^*) + \dots + I_{Y_m|X_{m-1}, Y_{m-1}}(\theta^*) + E\{I_{Y_{m+1}|X_m, Y_m}(\theta^*)\}, \quad (5)$$

where the outer expectation of the last term is with respect to unobserved data $Y_{m+1}|X_m, Y_m$.

As values Y_1, \dots, Y_m and X_0, \dots, X_{m-1} are already observed, the first m terms in (5) are replaced by the observed information $J(\theta^*)$. The information matrix is then a mixture of observed and expected information and its element in row i and column k is

$$\Psi_{X,Y}^m(\theta^*)_{i,k} = J_{Y_1|X_0}(\theta^*)_{i,k} + \dots + J_{Y_m|X_{m-1}, Y_{m-1}}(\theta^*)_{i,k} + E\{I_{Y_{m+1}|X_m, Y_m}(\theta^*)\}_{i,k} \\ = -\sum_{j=1}^N \left[\frac{\partial^2}{\partial \theta_i^* \partial \theta_k^*} \log p(y_{1j}|\mathbf{x}_{0j}) \right]$$

$$\begin{aligned}
& - \cdots - \sum_{j=1}^{N_{m-1}} \left[\frac{\partial^2}{\partial \theta_i^* \partial \theta_k^*} \log p(y_{mj} | \mathbf{x}_{m-1,j}, y_{m-1,j}) \right] \\
& - \sum_{j=1}^{n_m} \left[E \left\{ \frac{\partial^2}{\partial \theta_i^* \partial \theta_k^*} \log p(Y_{m+1,j} | X_{mj}, y_{mj}) \middle| \mathbf{x}_{0,j}, \dots, \mathbf{x}_{m-1,j}, y_{mj} \right\} \right], \quad (6)
\end{aligned}$$

where N_{m-1} is the number of individuals who have not had an event or been censored before the measurement $m - 1$ and n_m is the number of individuals to be selected for the m th measurement. Note that when only a subcohort has been selected for the measurements at a time point $\tau_{m'}, m' \in \{1, \dots, m-1\}$, then covariates are missing for the individuals not selected. We describe in Section 4.2, how these missing data are handled. The subcohort selection is carried out just before the new measurement, so Y_{m+1} and X_m are not observed for anyone. The expectation can be calculated by Monte Carlo integration.

The calculation of $\Psi_{X,Y}^m(\theta^*)$ requires the second-order partial derivatives of $\log p(y_1 | \mathbf{x}_0), \dots, \log p(y_{m+1} | \mathbf{x}_m)$. The value of the D_β -criterion is obtained by taking the determinant of the $H \times H$ upper left submatrix of $\Psi_{X,Y}^m(\theta^*)^{-1}$, where $\theta^* = (\beta_1, \dots, \beta_H, a, b)$. We denote this value of the criterion by $D_\beta^m(\xi^m, \theta^*)$, where ξ^m is an indicator matrix describing which individuals have been measured at the time points τ_0, \dots, τ_m and θ^* emphasizes that the criterion depends on the parameters.

3.3 | Bayesian selection

Now, we combine the Bayesian optimal design theory introduced in Section 3.1 and the criterion derived in Section 3.2. Here, the optimal selection is presented in a general form and the practical solution for the search problem is given in Section 4.1. Fisher information matrices of nonlinear models usually depend on model parameters (Chaloner & Verdinelli, 1995), which is also the case in our application. Therefore, some prior information about the parameters is needed in order to use the D_β -criterion. In the Bayesian approach the information obtained from the data already collected during the follow-up and/or from previous studies can be used to provide prior distributions of the parameters when applying the optimality criterion. In other words, we use informative priors, which are actually posteriors from the data already collected, in the subcohort selections. In the selection for the m th remeasurement, the prior probability density distribution $p^m(\theta^*)$ is, in fact, the posterior $p(\theta^* | x_0, \dots, x_{m-1}, y_1, \dots, y_m)$.

To minimize the D_β -criterion, we specify the last column of ξ^m so that the expected utility

$$U(\xi^m) = - \int_{\theta^*} D_\beta^m(\xi^m, \theta^*) p^m(\theta^*) d\theta^*, \quad (7)$$

will be maximized. Above, $D_\beta^m(\xi^m, \theta^*)$ is the value of the D_β -criterion for the m th remeasurement depending on the design ξ^m and the parameters θ^* . This is a specific form of the Equation (3).

4 | COMPUTATIONAL IMPLEMENTATION

4.1 | Search for optimal design

The integral in (7) can be approximated by sampling parameter values from the multivariate prior probability density distribution $p^m(\theta^*)$, generating data (X_m, Y_{m+1}) given the parameters and then

replacing the integral with a mean $\frac{1}{q} \sum_{l=1}^q D_{\beta}^m(\xi, \theta_l^*)$, where q is the number of realizations sampled from $p^m(\theta^*)$ (Atkinson, Demetrio, & Zocchi, 1995). The priors become more informative during the follow-up as the amount of collected data increases.

In addition to the model parameters and the new data, the missing covariate values are also treated as unknown parameters. The predictive distributions of the missing values are used as the prior distributions in the selections. This means that if the previous measurements include missing data, the criterion (7) averages also over these informative prior distributions of the missing values. We draw q realizations from the priors of the missing values and use them similarly to the realizations from $p^m(\theta^*)$.

In practice, the number of different subcohorts that could be selected for the new measurement is easily so large that it is computationally impossible to go through each of them. Therefore some heuristic method is needed. We use a so-called greedy method (Wright & Bailer, 2006), also known as sequential search (Dykstra, 1971), to find an approximately optimal subcohort. This method selects n individuals sequentially one by one: when $k - 1$ individuals ($0 < k < n$) have been selected for the subcohort, the k th selection is made so that the expected utility is maximized on the condition that information from previously selected $k - 1$ individuals is included in the calculation of the criterion. The procedure goes on similarly by selecting the next individual to be included in the subcohort so that the expected utility is maximized taking into account the information obtained from the previously selected individuals.

From the beginning of the selection procedure, the information matrix (6) includes all the information already collected during the follow-up. All the individuals who have not had an event are considered as candidates for the new measurement. The procedure continues by testing which candidate should be included in the expectation part on the last row of (6) in order to obtain the minimum value of the Bayesian optimality criterion $D_{\beta}^m(\xi^m, \theta^*)$. If there are two or more individuals who would minimize the criterion, the selection between them can be done randomly. The expected information of the selected individual is then included permanently into (6) and the procedure continues until the subcohort has reached the predetermined size. The covariate measurements are carried out after the whole subcohort has been selected.

4.2 | Parameter estimation

The estimation of the parameters θ^* of the survival model (2) is needed before each subcohort selection and finally when the follow-up study has ended. It should be emphasized that the model used in the subcohort selections need not be the same as the model used in the final analysis of the data, although the subcohort will be optimal with respect to the selection model. Once the data have been collected, the validity of the model assumptions can be reassessed and the model can be changed if needed. If the variables are the same in both models, the selected subcohort is likely to be better than a simple random sample, although the final analysis model could be more complex than the selection model. This is exemplified in Section 6.

In the subcohort selections, the estimated posterior distributions of the parameters are used as informative prior density distributions. The parameters are estimated using a Bayesian approach with MCMC sampling. We use noninformative priors for the parameters: $\beta \sim N(\mathbf{0}, 10^4 \cdot \mathbf{I})$, $r \sim \text{Gamma}(1, 0.0001)$ and $\alpha \sim N(0, 10^4)$, where r and α are reparameterized Weibull parameters, so that $r = a$ and $\alpha = -a \log b$, $\mathbf{0}$ is a zero vector and \mathbf{I} is an $H \times H$ identity matrix.

The subcohort selection may lead to a large amount of data “missing by design,” which we handle by Bayesian data augmentation (Tanner & Wong, 1987). The data missing by design are

missing at random because the selection can depend only on variables that have already been measured. Bayesian data augmentation has previously been used for data missing by design by for example, Kulathinal and Arjas (2006).

We consider continuous and binary covariates. The covariate models are estimated as a part of the entire Bayesian model. Recall the previously made Markov assumption

$$p(x_m | x_0, \dots, x_{m-1}, y_1, \dots, y_m) = p(x_m | x_{m-1}, \delta_m = 0).$$

If x_1 is a continuous covariate (body mass index in our application), the m th remeasurement of the covariate x_1 of individual j is modeled with a linear regression

$$x_{mj1} = c + \gamma x_{m-1,j,1} + \varepsilon_{mj},$$

where c is a constant and $\varepsilon_{mj} \sim N(0, v)$. The constant c , the coefficient γ and the error variance v could also be different at each measurement time if the structure of the covariate process should be allowed to vary over time. We assign priors $c \sim N(0, 100)$, $\gamma \sim N(0, 100)$ and $v \sim \text{Gamma}(1, 0.01)$, which are uninformative (vague) with respect to the scale of the covariate x_1 in the application.

If x_2 is a binary covariate, we use a logistic regression model

$$\begin{aligned} x_{mj2} &\sim \text{Bernoulli}(\pi) \\ \text{logit}(\pi) &= d_0 + d_1 x_{m-1,j,2}. \end{aligned}$$

Above, all the parameters remain constant in time, but this assumption could be relaxed if required. Uninformative priors $(d_0, d_1)^T \sim N(\mathbf{0}, 10^4 \cdot \mathbf{I})$ are used. Note that the parameters of the covariate process do not include those of the survival model and the priors of these two models are independent. Usually, there is some prior information available about the relationships of the variables and thus the use of informative priors would be justified. When applying these missing data models with MCMC estimation, it is recommended to center the covariates to improve the convergence (Lunn, Jackson, Best, Thomas, & Spiegelhalter, 2012).

5 | SIMULATION STUDY

5.1 | Description of the simulation study

Simulation studies were performed to illustrate what kind of subcohort should be selected according to the D_β -criterion and what is the benefit of using it. We compared the use of D_β -optimal subcohort selection with SRS and evaluated how much precision is lost in the estimation when compared with measuring the entire cohort.

We considered a setting with two independent continuous covariates and a 30-year follow-up with three measurement times at time points 0, 10, and 20 years. The size of the cohort was 1,500 individuals and the baseline ages were generated from the uniform distribution with the range from 45 to 65 years. Baseline measurements of continuous covariates x_0 and z_0 were made to follow $N(0, 1)$ distribution, first remeasurements x_1 and z_1 were drawn from $N(\gamma x_0, \sigma_\varepsilon^2)$ and $N(\gamma z_0, \sigma_\varepsilon^2)$ and second remeasurements x_2 and z_2 from $N(\gamma x_1, \sigma_\varepsilon^2)$ and $N(\gamma z_1, \sigma_\varepsilon^2)$, where $\gamma = 0.5$ and $\sigma_\varepsilon^2 = 0.75$. These parameter values lead to serial correlation of 0.5 between consecutive measurements and to a constant variance of the covariates at each measurement. The covariates were generated independently of each other and independently of age.

Survival times, that is, the ages at the event, were simulated from the Weibull distribution conditioned on the covariates through a time-dependent Weibull proportional hazards model. The shape parameter of the Weibull distribution was set to $a = 6.3$ and the scale parameter to $b = 27900$ (in days), which roughly equal the parameters estimated from the real data used in Section 6. To investigate if the magnitude of hazard ratios of the covariates has some effect on the selection, we used different regression coefficients for the two covariates. The coefficients were $\beta_x = 0.1$ ($e^{\beta_x} = 1.11$) and $\beta_z = 0.4$ ($e^{\beta_z} = 1.49$), respectively, for covariates x and z . If the event had not occurred at the end of the follow-up (30 years after the baseline), the survival time was censored. The measurement times were the same for all the individuals in calendar time, but as the individuals were of different ages, the measurements were not carried out at the same time points in age.

The simulation was repeated 100 times on a supercomputer of CSC-IT Center for Science Ltd by utilizing parallel computing. The three design approaches (D_β , SRS and full cohort) were applied to each simulated dataset. In the beginning of the follow-up, the cohort included 1,500 individuals, at the time of the second measurement, on average, 1,173 individuals were alive, at the time of the third measurement, on average, 712 individuals were alive and at the end of the follow-up, on average, 299 individuals were alive. Model parameters were estimated using the OpenBUGS version 3.2.3 (Lunn, Spiegelhalter, Thomas, & Best, 2009) and the rest of the calculations were carried out using the R statistical software version 3.1.1 (R Core Team, 2014).

5.2 | Subcohort selection

The subcohort selections are carried out using the Weibull proportional hazards model as the underlying model with age as the time scale. The left panels of Figure 1 show what kind of individuals are selected by the D_β -criterion for the second measurement. The selections for the second and third measurements are carried out up to 600 individuals sequentially one by one, but it is worth noting that the order is irrelevant when analyzing the data. The older individuals are clearly preferred in the selection, which may arise from the fact that older individuals are more likely to have an event during the next part of the follow-up and therefore provide more information than those who are likely to be censored. On the other hand, individuals with extreme covariate values are selected first. This is a reflection of the result that extreme selection is optimal for first-order linear regression models (Elfving, 1952).

The individuals are plotted separately by baseline covariates x and z , but there seems not to be a clear difference in the selection patterns, despite the different regression coefficients. The bottom left panel of Figure 1 shows the selection order against the sum of the absolute values, $|x| + |z|$. This is presented to illustrate better how extremity of the combination of the two independent covariates are preferred.

The selection of individuals for the third measurement, seen in the right panels of Figure 1, seems to be quite similar to the previous selection. The individuals first selected for the third measurement have also been selected for the second measurement. After approximately 200 individuals, the selection also includes individuals who were not selected for the second measurement. In the figure, the values plotted for missing measurements represent averages of 100 independent values generated for each missing value in an MCMC estimation. The preference for higher age in the selection is also seen with these individuals.

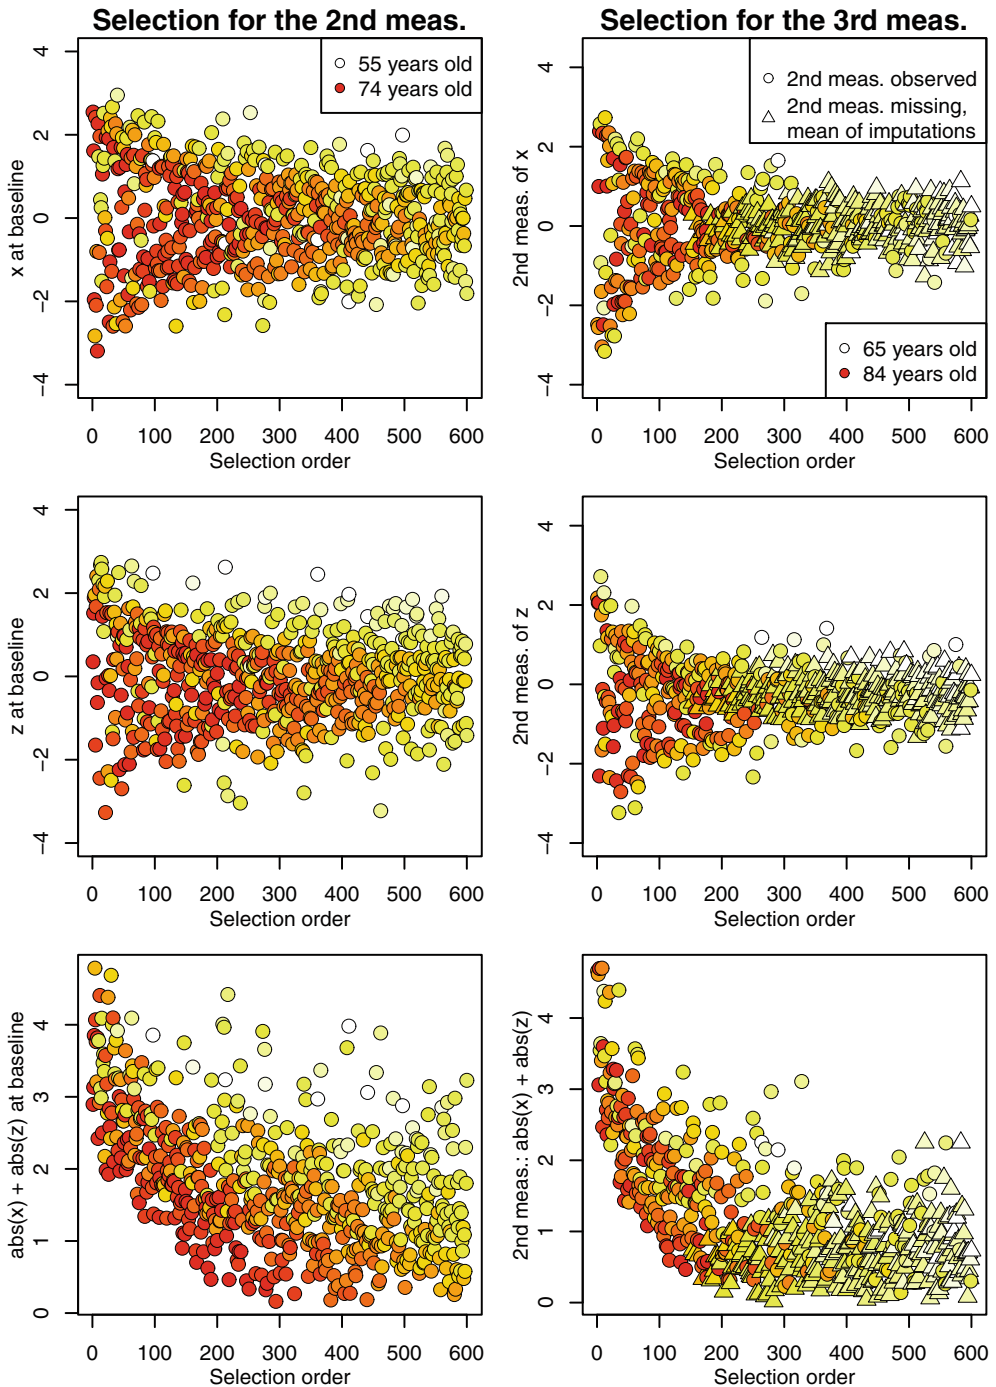


FIGURE 1 Selection orders of individuals for the second and third measurements for any n up to 600 using D_β -optimality and simulated data. Each point corresponds to one individual: the color shows the age of the individual at the time of the selection, the vertical axis shows the value of the covariate in the previous measurement (x in the uppermost, z in the middle and $|x| + |z|$ in the lowest panel) and the horizontal axis shows the round when the individual was selected in the greedy algorithm

TABLE 1 Simulation results for different designs from 100 simulation runs and for different sizes of the subcohorts (n). The entire cohort was measured at baseline. $\bar{\beta}_x$ and $\bar{\beta}_z$ indicate the means of the posterior means of β_x and β_z . SD is the standard deviation of the posterior means and Mean (SE) is the mean of the standard errors estimated from the Markov chain Monte Carlo (MCMC) chains

		Covariate x ($\beta_x = 0.1$)			Covariate z ($\beta_z = 0.4$)		
	Design	$\bar{\beta}_x$	SD	Mean (SE)	$\bar{\beta}_z$	SD	Mean (SE)
$n = 600$	Full cohort	0.096	0.031	0.029	0.40	0.028	0.030
	SRS	0.096	0.036	0.032	0.40	0.031	0.033
	D_β	0.095	0.033	0.031	0.40	0.032	0.032
$n = 500$	SRS	0.096	0.036	0.033	0.40	0.033	0.035
	D_β	0.095	0.034	0.032	0.40	0.032	0.034
$n = 400$	SRS	0.098	0.041	0.036	0.40	0.042	0.039
	D_β	0.094	0.036	0.034	0.40	0.034	0.035
$n = 300$	SRS	0.093	0.043	0.043	0.42	0.063	0.059
	D_β	0.095	0.037	0.036	0.40	0.040	0.038

Abbreviation: SRS, simple random sampling.

5.3 | Design comparisons

The analysis was carried out using the SRS-designs, D_β -designs and the entire simulated data set without subcohort selections. The sizes of the subcohorts varied from 300 to 600. In each design, the entire cohort was measured at baseline and the same subcohort size was used in both the second and the third measurements. We compared the bias and SEs of β_x and β_z between the D_β -design and SRS.

The results in Table 1 show that there is no considerable bias in the estimates when only a subcohort is selected for remeasurements. According to sds of the estimates and mean SEs, the D_β -design seems to lead almost consistently to more precise estimation of the coefficients than the SRS-design. The difference becomes more prominent when the subcohort size decreases. We achieve virtually the same precision when comparing the D_β -design with subcohort size 400 to the SRS-design with subcohort size 500, or when comparing the D_β -design with subcohort size 300 to the SRS-design with subcohort size 400.

An important observation is that the precision does not decrease dramatically although only 300 individuals are remeasured. 300 is only on average 26% of the individuals alive at the time of the second measurement and on average 42% of the individuals alive at the time of the third measurement.

6 | RESULTS FOR THE EAST-WEST STUDY

Next, we will present an application to data from a real follow-up study, the East–West study (Reinikainen, Laatikainen, Karvanen, & Tolonen, 2015). The East–West study was started as the Finnish part of the international Seven Countries Study (Keys, 1970) initiated in the late 1950s to investigate cardiovascular diseases and their risk factors across different countries and cultures. The Finnish cohorts consist of all men born between 1900 and 1919 and

living in two geographically defined areas in Eastern Finland and in South-Western Finland ($N = 1711$). The data include baseline measurements carried out in 1959 and longitudinal measurements in 1964, 1969, 1974, 1984, 1989, 1994, 1999 and follow-up for mortality until the end of 2011.

Our analysis with the East–West data is an example of using a binary covariate and a continuous covariate with nonlinear effect. For this example only a part of the data is used. We consider the measurement in the year 1964 as the baseline measurement, measurements in 1974 and 1984 as remeasurements and 1994 as the end of the follow-up, when censoring is carried out. After removing individuals who died before 1964, we had 1,594 individuals. In the setting of this example, at the time of the second measurement 1,225 individuals were alive, at the time of the third measurement 766 individuals were alive and at the end of the follow-up 320 individuals were alive.

All-cause mortality is the outcome in the analyses and age is chosen as the time scale. Smoking status is used as a binary covariate and body mass index (BMI) as a continuous covariate, whose effect on survival is assumed to be quadratic. Some studies have reported U-shaped associations between BMI and all-cause mortality (Corrada, Kawas, Mozaffar, & Paganini-Hill, 2006; Zhao et al., 2014) and this was also found in the East–West cohorts. The quadratic effect is implemented in the survival model simply by adding a squared term in BMI. Note that the final analysis of the data could still be carried out by using other methods, for example, splines (Therneau & Grambsch, 2000), even if polynomials were used in subcohort selection. Before the second measurement, there was no information in the data about the changes in smoking, so we assumed in the calculation of the expectation in (6) that a smoker at the baseline will be a nonsmoker in second measurement with probability 0.4 and a nonsmoker will become a smoker with probability 0.1. In the selection for the third measurement, these probabilities were estimated from the data. The correlation structures of the data implicated that AR(1) models can be used for the covariate processes.

TABLE 2 Results for the East–West data for different sizes of the subcohorts (n). The entire cohort was measured at baseline. For simple random sampling (SRS), $\bar{\beta}_1$, $\bar{\beta}_2$, $\bar{\beta}_3$ and SE are means of the posterior means and standard errors estimated from the Markov chain Monte Carlo chains from 1,000 analyses

		BMI (linear)	BMI (quadratic)	Smoking
Design		$\bar{\beta}_1$ (SE)	$\bar{\beta}_2$ (SE)	$\bar{\beta}_3$ (SE)
$n = 600$	Full cohort	−0.043 (0.0086)	0.0065 (0.0011)	0.39 (0.069)
	SRS	−0.045 (0.0138)	0.0071 (0.0017)	0.45 (0.075)
	D_β	−0.041 (0.0087)	0.0063 (0.0011)	0.45 (0.067)
$n = 500$	SRS	−0.049 (0.0176)	0.0078 (0.0022)	0.47 (0.080)
	D_β	−0.043 (0.0091)	0.0062 (0.0011)	0.42 (0.071)
$n = 400$	SRS	−0.052 (0.0218)	0.0086 (0.0022)	0.48 (0.083)
	D_β	−0.042 (0.0093)	0.0071 (0.0012)	0.45 (0.073)
$n = 300$	SRS	−0.044 (0.0232)	0.0094 (0.0017)	0.47 (0.080)
	D_β	−0.040 (0.0110)	0.0058 (0.0015)	0.48 (0.075)

Abbreviation: BMI, body mass index.

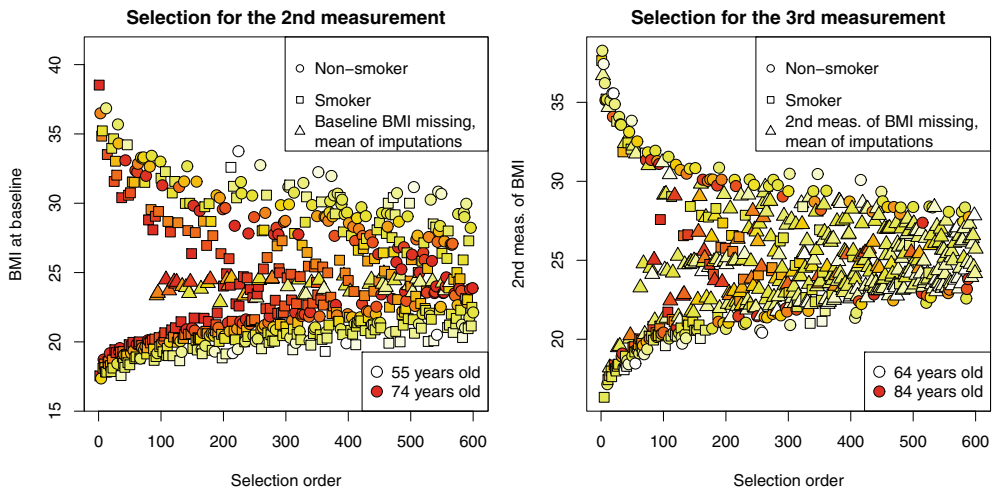


FIGURE 2 Selection order of individuals for the second measurement (left panel) and for the third measurement (right panel) for any n up to 600 using D_β -optimality and the East-West data. Each point corresponds to one individual: the color shows the age of the individual at the time of the selection, the vertical axis shows the value of body mass index (BMI) in the previous measurement and the horizontal axis shows the round when the individual was selected in the greedy algorithm. Individuals with missing previous BMI are indicated by triangle symbols regardless of the missingness of the smoking status. Usually, the missingness of BMI means also the missingness of the smoking status in the data

In order to improve the convergence in MCMC estimation, all the covariates were centered before analysis by subtracting the means of observed baseline values from the corresponding covariate values. The mean of baseline BMI observations was 24.25 and the centered baseline BMI ranged from -9.25 to 21.62 . The estimates of the analysis with the full cohort (Table 2) correspond to an upward opening parabola, which has minimum risk with centered BMI of 3.31.

The selection order of 600 individuals for the second measurement according to D_β can be seen in the left panel of Figure 2. The preference for extreme BMI values is clear. Although baseline measurements are missing for some individuals in the original data, some of these are still selected into the subcohort. At a fixed BMI level, older individuals seem to be selected before younger. We have a quadratic model in BMI, so one could have expected that in addition to extreme values also average values of BMI would have been preferred. The average values are, however, not so important in the second measurement, because we have baseline measurements for all individuals and thus have observed average values already much more than extreme values. The selected subcohort includes 327 individuals who were smokers at the baseline and 232 baseline non-smokers. At the baseline, there were also clearly more smokers than nonsmokers. Apparently the selection procedure tries to balance the expected number of smokers and nonsmokers measured altogether.

Figure 2 (right panel) shows the subcohort selection for the third measurement. Similar patterns can be observed here as in the previous selection. One big difference is the large number of individuals who have missing previous measurements. This can be explained by the fact that many old individuals who were measured in the second measurement are already dead at the time of the third measurement. The effect of age does not seem to be so strong, but this comes mainly from the mixing of smokers and nonsmokers in the plot. These are separated in Figure 3, which




FIGURE 3 Selection order of individuals for the third measurement using D_β -optimality and the East–West data. These panels represent the right panel of Figure 2 decomposed into smokers, nonsmokers and those who have missing value of body mass index (BMI) in the second measurement. Individuals with missing BMI are indicated by triangle symbols regardless of the missingness of the smoking status. Usually, the missingness of BMI means also the missingness of the smoking status in the data

TABLE 3 Results for comparisons of using different models in the subcohort selections and in the final analysis using the East–West data. 500 individuals were selected for the second and third measurements. For simple random sampling (SRS), $\bar{\beta}_1$, $\bar{\beta}_2$ and SE are means of the posterior means and standard errors estimated from the Markov chain Monte Carlo chains from 1,000 analyses

Analysis model	Selection	BMI (linear)	BMI (quadratic)
		$\bar{\beta}_1$ (SE)	$\bar{\beta}_2$ (SE)
Quadratic	Full cohort	−0.059 (0.0082)	0.0073 (0.0011)
	SRS	−0.059 (0.0096)	0.0079 (0.0013)
	D_β (quadr. model)	−0.059 (0.0085)	0.0077 (0.0011)
	D_β (lin. model)	−0.059 (0.0087)	0.0073 (0.0011)
Linear	Full cohort	−0.038 (0.0077)	
	SRS	−0.033 (0.0084)	
	D_β (quadr. model)	−0.038 (0.0084)	
	D_β (lin. model)	−0.037 (0.0084)	

Abbreviation: BMI, body mass index.

reveals that there are less smokers than nonsmokers in the selected subcohort. In fact, there were 104 smokers, 203 nonsmokers and 459 individuals with missing smoking status among the candidates for the third measurement, of which all 104 smokers, 138 nonsmokers and 358 individuals with missing smoking status were selected for the third measurement.

Table 2 shows that there is clear benefit of using the D_β -design instead of the SRS. All the standard errors in the D_β -design are smaller than in the SRS-design for each subcohort size used. The D_β -selection leads usually to estimates closer to those obtained using the full cohort, than the SRS. Surprisingly, both selection methods seem to lead to greater estimates of the effect of smoking than the full cohort.

The standard errors of the quadratic term of BMI and smoking, obtained using SRS, are smaller when $n = 300$ than when $n = 400$, which is an unexpected result. Three hundred individuals are only 24% of the individuals alive at the time of the second measurement and 39% of those alive at the time of the third measurement, which may be too small proportions in a real study to obtain reliable estimates. Estimation may become sensitive to model misspecification when the proportion of missing data becomes large (Saarela, Kulathinal, & Karvanen, 2012).

In practice, an analyst would not necessarily like to use the same model in optimal selection and in the final analysis. Table 3 shows the results from an example of using different models in the selections and final analysis. Here, we use the same data as in the previous example, but (centered) BMI as the only covariate. When the analysis model is quadratic, the quadratic selection model leads to slightly better precision than the linear selection model but both selection models still outperform SRS. When the analysis model is linear, the results do not deteriorate even if the quadratic selection model is used.

7 | DISCUSSION

The cost-efficiency of a follow-up study can be improved by careful planning of longitudinal measurements. The present paper considered the scenario where we can afford measuring the time-varying covariates only for a subset of the cohort. We proposed using a Bayesian approach in optimal subcohort selection with a Fisher information based D_β -optimality. Our work overcomes the limitations of the previous paper (Reinikainen et al., 2016), where a simple scenario with only one covariate and one re-measurement was considered.

The estimates and their precision corresponding to the D_β -selection and SRS were compared. The use of the D_β -optimality led to more precise estimates and the precision was seen to remain satisfactory compared with the full cohort design.

The results indicated that in order to obtain estimates as precise as possible for regression parameters of the survival model, old individuals with extreme covariate values should be preferred, which is consistent with our previous results (Reinikainen et al., 2016). A similar result was obtained when we used a covariate with quadratic effect (BMI) in the real follow-up data example. Another covariate used in this example was smoking status as a binary variable. Optimal selection seemed to balance the expected number of smokers and nonsmokers measured.

The general idea of measuring only a subcohort is applied in many epidemiological study designs, like case-control and case-cohort designs and their variants (Keogh & White, 2013; Kulathinal et al., 2007; Sun, Joffe, Chen, & Brunelli, 2010). However, in our application the setup is different. We use an approach with an explicit utility function describing the goal of the study. The presented approach borrows elements from optimal design of experiments and applies them to the design of an observational study.

We recommend the use of a Bayesian approach in this kind of sequential study design problem. Some prior knowledge is always required in nonlinear design problems, because optimal designs depend on model parameters and so a Bayesian approach with informative priors is a natural way to incorporate this knowledge into design optimization. Bayesian data augmentation appears here to be a flexible method for the handling of missing data when we increase the number of measurement points and covariates.

The proposed approach can be adopted to other scenarios where repeated measurements are conducted. However, it is difficult to provide a general purpose software implementation because the details of Bayesian models differ case by case and the forms of survival model and optimality

criterion should be selected to fit the available data and the aims of the study. We recommend that researchers test the performance of their implementations in simulations similar to those presented in Table 1 before applying them in real study design.

In experiments, the optimal design for a linear model might consist of only two design points, which would give no power for detecting nonlinear effects. In principle, the same applies also in our setup, but in practice, the problem is not realized in observational studies with continuous covariates and moderate sample size. The reason for this is that only a few individuals with the optimal covariate values are available in the cohort and after they are selected, the selection procedure must rely on individuals with a wider variety of covariate values.

Although the covariate processes were assumed to be piece-wise constant in the subcohort selections, a joint model of survival and longitudinal data (Rizopoulos, 2012) could be considered in the final analysis for more realistic treatment of time-varying covariates. In the general case with multiple longitudinal covariates, joint modeling is, however, computationally very demanding and would be further complicated in our approach with a large amount of missing data.

We considered the design optimization with respect to only one model at a time. However, if two or more models or utility functions would be of interest, compound design criteria could be applied (Atkinson et al., 2007). Then, the optimality criterion should include the parameters of all models of interest. If we later want to use the collected data to some purpose not addressed in subcohort selection, the optimality does not hold anymore, and in an extreme case the selected subcohort could perform even worse than SRS. This situation is unlikely to occur, if a new outcome variable has the same covariates as the one used in optimal selection, or if new covariates are correlated with those used in optimization. The situation is similar to case-control studies where the controls for a specific outcome can be used for other outcomes (Saarela et al., 2012; Saarela, Kulathinal, Arjas, & Läärä, 2008) although they are not optimal.

The selection procedure presented in this paper requires that up-to-date survival information is already available during the study. The information of the measured covariates is also needed during the study if it is used in the optimal selection. The proposed approach is also applicable to some retrospective designs. For instance, consider a study where blood samples or other biological specimen are collected and stored for all individuals and years later some biomarkers are measured from the stored sample. Selecting only a subcohort for these measurements may be a reasonable option if the extraction of the biomarkers is expensive. Then an approach similar to one presented in this paper could be used to optimally select the subcohort.

The proposed selection method is nonrandom in the sense that individuals are selected deterministically according to the selection criterion, but we do not see this as a disadvantage when the parameters of a survival model are of interest. Distributions of the covariates or absolute risks of the outcome in the population can also be assessed using the estimated analysis model. If it is important to assess the distributions of the covariates without relying on model assumptions, the subcohort should be selected as a random sample (Kulathinal et al., 2007). In this setting a randomized version of the sequential design construction could be considered (Atkinson & Biswas, 2014). Our method may introduce some selection bias, if the effect of a covariate changes with age and this has not been taken into account in the selection. The procedure could also be developed to include the costs of data collection in the optimization (Fedorov & Leonov, 2013).

ACKNOWLEDGEMENTS

The research of the first author was supported by the Emil Aaltonen Foundation. Hanna Tolonen from Finnish Institute for Health and Welfare is acknowledged for providing the data of the

East–West study. The authors thank CSC – IT Center for Science Ltd. for providing computing resources and Marie Reilly, Kari Auranen and Juho Kopra for helpful comments.

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available on request from the Finnish Institute for Health and Welfare. The data are not publicly available due to privacy or ethical restrictions.

ORCID

Jaakko Reinikainen  <https://orcid.org/0000-0002-7825-5359>

REFERENCES

- Atkinson, A., Demetrio, C., & Zocchi, S. (1995). Optimum dose levels when males and females differ in response. *Journal of the Royal Statistical Society Series C (Applied Statistics)*, 44, 213–226.
- Atkinson, A. C., & Biswas, A. (2014). *Randomised response-adaptive designs in clinical trials*. Boca Raton, FL: CRC Press.
- Atkinson, A. C., Donev, A. N., & Tobias, R. D. (2007). *Optimum experimental designs, with SAS*. Oxford, UK: Oxford University Press.
- Buzoianu, M., & Kadane, J. B. (2009). Optimal Bayesian design for patient selection in a clinical study. *Biometrics*, 65, 953–961.
- Chaloner, K., & Verdinelli, I. (1995). Bayesian experimental design: A review. *Statistical Science*, 10, 273–304.
- Corrada, M. M., Kawas, C. H., Mozaffar, F., & Paganini-Hill, A. (2006). Association of body mass index and weight change with all-cause mortality in the elderly. *American Journal of Epidemiology*, 163, 938–949.
- Dykstra, O. (1971). The augmentation of experimental data to maximize $|X'X|$. *Technometrics*, 13, 682–688.
- Elfving, G. (1952). Optimum allocation in linear regression theory. *The Annals of Mathematical Statistics*, 23, 255–262.
- Fedorov, V. V., & Leonov, S. L. (2013). *Optimal design for nonlinear response models*. Boca Raton, FL: CRC Press.
- Karvanen, J., Kulathinal, S., & Gasbarra, D. (2009). Optimal designs to select individuals for genotyping conditional on observed binary or survival outcomes and non-genetic covariates. *Computational Statistics & Data Analysis*, 53, 1782–1793.
- Keogh, R. H., & White, I. R. (2013). Using full-cohort data in nested case–control and case–cohort studies by multiple imputation. *Statistics in Medicine*, 32, 4021–4043.
- Keys, A. (1970). Coronary heart disease in seven countries. *Circulation*, 41, 186–195.
- Kulathinal, S., & Arjas, E. (2006). Bayesian inference from case-cohort data with multiple end-points. *Scandinavian Journal of Statistics*, 33, 25–36.
- Kulathinal, S., Karvanen, J., Saarela, O., Kuulasmaa, K., & MORGAM Project. (2007). Case-cohort design in practice – Experiences from the MORGAM project. *Epidemiological Perspectives & Innovations*, 4, 15.
- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *The Annals of Mathematical Statistics*, 27, 986–1005.
- Lunn, D., Jackson, C., Best, N., Thomas, A., & Spiegelhalter, D. (2012). *The BUGS book: A practical introduction to Bayesian analysis*. Boca Raton, FL: CRC Press.
- Lunn, D., Spiegelhalter, D., Thomas, A., & Best, N. (2009). The BUGS project: Evolution, critique and future directions. *Statistics in Medicine*, 28, 3049–3067.
- Mehtälä, J., Auranen, K., & Kulathinal, S. (2015). Optimal observation times for multistate Markov models – applications to pneumococcal colonization studies. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64, 451–468.
- Pukelsheim, F. (1993). *Optimal design of experiments*. New York, NY: Wiley.
- R Core Team. (2014). *R: A language and environment for statistical computing*. Vienna, Austria: Retrieved from: R Foundation for Statistical Computing. <http://www.R-project.org/>
- Reinikainen, J., Karvanen, J., & Tolonen, H. (2016). Optimal selection of individuals for repeated covariate measurements in follow-up studies. *Statistical Methods in Medical Research*, 25, 2420–2433.

- Reinikainen, J., Laatikainen, T., Karvanen, J., & Tolonen, H. (2015). Lifetime cumulative risk factors explain cardiovascular disease mortality in a 50-year follow-up study in Finland. *International Journal of Epidemiology*, 44, 108–116.
- Rizopoulos, D. (2012). *Joint models for longitudinal and time-to-event data: With applications in R*. Boca Raton, FL: CRC Press.
- Saarela, O., Kulathinal, S., Arjas, E., & Läärä, E. (2008). Nested case-control data utilized for multiple outcomes: A likelihood approach and alternatives. *Statistics in Medicine*, 27, 5991–6008.
- Saarela, O., Kulathinal, S., & Karvanen, J. (2012). Secondary analysis under cohort sampling designs using conditional likelihood. *Journal of Probability and Statistics*, 2012, 1–37.
- Sun, W., Joffe, M. M., Chen, J., & Brunelli, S. M. (2010). Design and analysis of multiple events case-control studies. *Biometrics*, 66, 1220–1229.
- Tanner, M. A., & Wong, W. H. (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association*, 82, 528–540.
- Therneau, T. M., & Grambsch, P. M. (2000). *Modeling survival data: Extending the Cox model*. New York, NY: Springer.
- Woods, D. C., Overstall, A. M., Adamou, M., & Waite, T. W. (2017). Bayesian design of experiments for generalized linear models and dimensional analysis with industrial and scientific application. *Quality Engineering*, 29, 91–103.
- Wright, S. E., & Bailer, A. J. (2006). Optimal experimental design for a nonlinear response in environmental toxicology. *Biometrics*, 62, 886–892.
- Zhao, W., Katzmarzyk, P. T., Horswell, R., Wang, Y., Li, W., Johnson, J., ... Hu, G. (2014). Body mass index and the risk of all-cause mortality among patients with Type 2 diabetes. *Circulation*, 130, 2143–2151.

How to cite this article: Reinikainen, J., & Karvanen, J. (2022). Bayesian subcohort selection for longitudinal covariate measurements in follow-up studies. *Statistica Neerlandica*, 76(4), 372–390. <https://doi.org/10.1111/stan.12264>