

Personalized Cardiovascular Disease Risk Mitigation via Longitudinal Inverse Classification

Michael T. Lash
Business Analytics Area
University of Kansas
 Lawrence, USA
 michael.lash@ku.edu

W. Nick Street
Business Analytics Department
University of Iowa
 Iowa City, USA
 nick-street@uiowa.edu

Abstract—Cardiovascular disease (CVD) is a serious illness affecting millions world-wide and is the leading cause of death in the US. Recent years, however, have seen tremendous growth in the area of personalized medicine, a field of medicine that places the patient at the center of the medical decision-making and treatment process. Many CVD-focused personalized medicine innovations focus on genetic biomarkers, which provide person-specific CVD insights at the genetic level, but do not focus on the practical steps a patient could take to mitigate their risk of CVD development. In this work we propose longitudinal inverse classification, a recommendation framework that provides personalized lifestyle recommendations that minimize the predicted probability of CVD risk. Our framework takes into account historical CVD risk, as well as other patient characteristics, to provide recommendations. Our experiments show that earlier adoption of the recommendations elicited from our framework produce significant CVD risk reduction.

I. INTRODUCTION

Cardiovascular disease (CVD) is a serious illness that affects millions in both the United States and across the world. In 2016, CVD was the leading cause of death in the US, responsible for more than 840,000 deaths [1], thus accounting for one in three of all US deaths [2]. Moreover, over 1 million individuals experienced a CVD event in 2019. Comprehensively, these statistics paint a troubling picture of the current state of cardiovascular health.

Encouragingly, however, over the past four to five decades cardiovascular disease mortality has been on the decline [3]. This observed decline is attributable to many advances in modern medicine, ranging from surgical bypass innovations to preventive care and education. Most recently, however, researchers have begun to worry that CVD mortality may be plateauing and the decline even reversing among certain populations [3].

Advances in precision medicine may be the key to encouraging continued CVD mortality reduction. Precision medicine methodologies operate by placing the individual at the center of the medical decision-making process to find the interventions or treatments that maximize the probability of a good outcome. Therefore, these methodologies consider person-specific characteristics that have a bearing on the medical issue at hand. Such characteristics may be incredibly fine-grained, defined in terms of one's genetic make-up, or more course-grained, such as demographic information, and may

also include factors that pertain to one's lifestyle, such as exercise and eating habits.

A particularly promising stream of precision medicine research is in the area of machine learning. Broadly speaking, machine learning methods induce (i.e., train) a model that learns a mapping from features (i.e., variables) to some outcome of interest, either continuously valued (regression) or discrete (classification) using historical data. In the context of CVD precision medicine, the features represent patient characteristics (demographics, lifestyle), which are mapped to CVD outcome, and the historical data are historical medical records. The trained model can then be used to make predictions as to the disease outcome for new patients.

Trained machine learning models are incapable of providing personalized (i.e., precision) recommendations that mitigate the probability of developing the disease they've learned to predict, however. In other words, machine learning models are strictly diagnostic. To turn predictions into precision prescriptions additional innovations are needed.

Inverse classification turns predictions into personalized prescriptions by finding the instance-specific (patient-specific) feature value perturbations that minimize the probability of an undesirable classification. These perturbations correspond to manipulations of the patient's characteristics. Intuitively, not all characteristics can be changed. For instance, it would make little sense to perturb age or genetic characteristics because the patient cannot change her age or genetics. On the other hand, she can change her eating habits and exercise levels. Therefore, perturbations are only made to lifestyle or similarly manipulable feature values. The perturbations are moderated by feasibility constraints that can be tailored to the individual and are implemented to avoid infeasible recommendations (e.g., run 1000 miles/day). Ultimately, the perturbations obtained from application of inverse classification represent personalized recommendations that optimally minimized the probability (risk) of CVD.

In this work we propose a *longitudinal* inverse classification approach to CVD risk minimization. Longitudinal, in this context, indicates that we will observe the same patients over several time periods. By devising a longitudinal approach to the inverse classification problem we can observe how inverse classification-elicited recommendations affect CVD risk over

multiple, temporal periods. In other words, we can explore how personalized recommendations adopted earlier vs. later affect future CVD risk over several periods of time. The recommendations elicited from the proposed inverse classification process will be in the form of lifestyle adjustments, such as changes to diet and exercise. The proposed longitudinal approach will also allow us to more realistically model CVD development by taking into account past patient history.

Our innovations in the areas of CVD risk minimization, precision medicine, and inverse classification can be summarized as follows: **(1)** We develop an updated inverse classification framework to accommodate longitudinal data, thereby creating a more realistic formulation that creates temporal linkages between each observation period. **(2)** Using our proposed longitudinal inverse classification framework we show how early adoption of the elicited precision recommendations continues to lower CVD development probability during subsequent periods. In other words, we show how lifestyle adjustments made earlier continue to reduce CVD probability during later observation periods. **(3)** We explore and solve an intermediary “missing variable” problem that frequently arises when longitudinal data are used. Namely, we propose and evaluate the use of so-called “missing value estimators” that provide estimates for features that are not measured during some time periods in the longitudinal study.

The paper proceeds by first disclosing related work in Section 2, followed by our *longitudinal inverse classification* formulation in Section 3. Section 4 discloses our data, experiments and results and Section 5 concludes the paper.

II. RELATED WORK

We decompose our discussion of related works according to four categories: precision and personalized medicine, personalized medicine applied to CVD, machine learning with longitudinal data, and inverse classification.

Personalized and precision medicine are relatively new avenues of research that have been gaining popularity in recent years. Personalized medicine differs from so-called “traditional” forms of one-size-fits-all medicine by placing the individual at the center of medical decision-making process to find a medical course of action that is specific to the individual [4]. Personalized medicine may operate at a very fine-grained genetic level to discover new drugs [5], for instance, or more course-grained demographic and lifestyle levels to prevent and mitigate chronic diseases [6]. This study is concerned with the latter, where we are attempting to mitigate the risk of CVD development.

Personalized medicine research in the area of CVD is somewhat understudied, particularly in regard to technologies that have been deployed and are being used in practice [7]. Nevertheless, recent years have witnessed a growth in personalized medicine for CVD [8], much of which is focused on genomic research with particular emphasis being placed on the identification of biomarkers indicative of the disease [8]–[11]. In contrast, this work focuses on the creation of a framework that can immediately be used to make personalized

lifestyle recommendations that mitigate long-term CVD risk, thus bridging the gap between traditional preventive lifestyle-focused CVD research [1], [2] and personalized medicine.

Emergent data mining and machine learning research involving **longitudinal data** is focused on methodologically leveraging such data, as well as the specific domains in which such methods can be employed. In [12] the authors explore deep neural networks that, with minimal preprocessing, learn a mapping from patients’ lab tests to over 130 diseases (multi-task learning). In [13] unsupervised learning methods are applied to longitudinal health data to learn a disease progression model. The model can subsequently be used to aid patients in making long-term treatment decisions. These works exemplify the way in which models can be learned to aid in predicting disease [12] and in forecasting disease progression [13]. In this work we examine how (1) coupling longitudinal data with predictive models can make disease risk estimation more accurate and (2) how predictive models that incorporate historical risk and past behavior can be used to make recommendations that optimally minimize the likelihood of developing a certain disease.

Inverse classification methods are varied in their approach to finding optimal recommendations, either adopting a greedy [14]–[17] or non-greedy formulation [18]–[21]. Past works also vary in their implementation of constraints that lead to more realistic recommendations, either being completely unconstrained [14], [15], [17], or constrained [16], [18]–[21]. In this work, we adopt the formulation and framework related by [20], [21] which accounts for (a) the features that can and cannot be changed (e.g. age cannot be changed, but exercise levels can), (b) varying degrees of change difficulty (feature-specific costs) and (c) a restriction on the cumulative amount of change (budget). As in [21], we implement a method that avoids making greedy recommendations while still accounting for (a), (b) and (c).

III. LONGITUDINAL INVERSE CLASSIFICATION

In this section we begin by providing some preliminary notation followed by a definition and formulation of past CVD risk estimators and a definition and formulation of so-called missing feature value estimators, which are necessarily incorporated into our formulation of a longitudinal inverse classification framework.

A. Preliminaries

Let $\{(id_v^{(i)}, \mathbf{x}_v^{(i)}, y_{v+1}^{(i)})\}_{i=1}^{n_v}$ be a dataset of n instances, where $id_v^{(i)}$ is a value that uniquely identifies patient i , $\mathbf{x}_v^{(i)}$ denotes patient i ’s feature vector, $y_{v+1}^{(i)} \in \{0, 1\}$ is the known CVD outcome of patient i with 1 indicating that the patient developed CVD and 0 indicating that the patient did not develop CVD. Furthermore, $v = 1, \dots, V$ denotes the discrete temporal unit in which the $\{\mathbf{x}_v^{(i)}\}_{i=1}^{n_v}$ patient characteristics were measured. On the other hand, the outcome of interest $y_{v+1}^{(i)}$ is observed at the next (i.e., immediately preceding) temporal period, hence the $v + 1$ notation. Note that, for the sake of notational convenience, we assume that there is a way

to determine $\{y_{v+1}^{(i)}\}_{i=1}^{n_v}$ when $v = V$. Furthermore, let $f_v(\cdot)$ denote a classifier induced on training data from visit (i.e., temporal period) v ; $f : \mathbb{R}^{n_v \times p_v} \rightarrow [0, 1]$, where n_v and p_v denote the number of instances and number of features measured, respectively, at visit v . We assume that $f_v(\cdot)$ is a classification function that produces probability estimates of CVD outcome, hence the bounded output domain of $[0, 1]$. This is not a limiting assumption since any classifier that does not natively produce probability estimates can be “trained” to produce such estimates using Platt Scaling [22].

In this setting, an assumption made about the longitudinal datasets is that of *instance continuity*, where the same instances are represented in each of the defined $v = 1, \dots, V$ datasets. In other words, given a set of numbers uniquely identifying each instance in a dataset at visit v , denoted ID_v , the following holds

$$ID_{v+1} = \{id_{v+1}^{(i)}\}_{i=1}^{n_{v+1}} \subseteq ID_v = \{id_v^{(i)}\}_{i=1}^{n_v} : \text{for } v = 1, \dots, V. \quad (1)$$

The specification given in (1) allows us to follow the progression of individual instances over time. However, to insure that the classification function induced on $\{\mathbf{x}_v^{(i)}\}_{i=1}^n$ learns a consistent representation of the outcome, namely $\{y_v^{(i)}\}_{i=1}^n$, the instances who experience the outcome of interest at v are removed from the $v + 1, \dots, V$ datasets. In other words, the following holds

$$\{id_{v+1}^{(i)}\}_{i=1}^{n_{v+1}} \cap \{id_v^{(i)} | y_v^{(i)} = 1\}_{i=1}^{n_v} = \emptyset : \text{for } v = 1, \dots, V. \quad (2)$$

Let \mathcal{F}_v denote the feature set at visit v and U, I, D denote index sets corresponding to certain feature categories. The U index set denotes “unchangeable” features, such as age and height that cannot be changed by the patient and for which we cannot provide recommendations – e.g., $\mathcal{F}_U = \{\text{age, height, } \dots\}$. The D index set, conversely, denotes features that can be changed “directly” and represent the features for which we can provide recommendations. The features include lifestyle attributes such as diet and exercise – e.g., $\mathcal{F}_D = \{\text{diet, exercise, } \dots\}$. The I index set indicates so-called “indirectly” changeable features. This set of features correspond to those variables that cannot be changed directly, but may change as a consequence to changes made in the \mathcal{F}_D feature set. Such variables might include blood pressure and blood glucose, among others – e.g., $\mathcal{F}_I = \{\text{blood pressure, blood glucose, } \dots\}$.

Since the \mathcal{F}_I features change as a consequence of manipulations made to the \mathcal{F}_D , which are the features we will ultimately be making recommendations about (i.e., making changes to), we wish to model this dependence explicitly. However, the \mathcal{F}_I features also depend upon the \mathcal{F}_U , which will remain static. Therefore, let $H : \mathbb{R}^{|U|+|D|} \rightarrow \mathbb{R}^{|I|}$ denote an *indirectly changeable feature estimator* that takes the \mathbf{x}_U and \mathbf{x}_D feature values as input and provides estimates estimates for the \mathbf{x}_I feature values. Therefore, we can use $H(\cdot)$ to track how changes made to \mathbf{x}_D influence the \mathcal{F}_I feature values. This indirect feature value estimator will be substituted in place

of the \mathbf{x}_I feature values when we formulate our longitudinal inverse classification framework in the proceeding subsection.

B. Historical Risk Estimates Features

Since we observe the same instances in the datasets specified over the $v = 1, \dots, V$ visits, we can also observe how the outcome of interest temporally develops. Moreover, we can induce models that take into account *outcome development* (i.e., disease progression) when learning a mapping to the outcome of interest.

To induce such models some notion of outcome development (i.e., progression) must exist somewhere in the data. Initially, however, the datasets are temporally “unlinked”. In other words no connection exists to link a patient’s observed features values at an earlier time to those at later time. A naive way to create a temporal link would be to carry-forward a patient’s observed feature values at each of the $v' = 1, \dots, v - 1$ time periods and add them as features in the v period. This would, of course, create increasingly large feature vectors and is not practical.¹ Therefore, we propose using CVD *risk* estimates – the predicted probability of CVD occurring – as features in “future” datasets. The addition of such features act as temporal links between previous datasets and future datasets, providing a notion of how likely an instance was to develop CVD during previous visits, thus providing a more realistic picture of outcome development. Furthermore, CVD risk estimates are represented as single values and are therefore more practical to use than the aforementioned naive procedure.

To formalize this notion, let

$$risk_v^{(i)} = f_v(\mathbf{x}_v^{(i)}) \quad (3)$$

denote an estimate of risk, obtained from classification function $f_v(\cdot)$ at visit v , for patient $id_v^{(i)}$ having feature values $\mathbf{x}_v^{(i)}$. To create a vector of *past* risk, corresponding to an arbitrary number of past visits, we can write

$$\mathbf{risk}_v^{(i)} = (f_{v-(v-1)}(\mathbf{x}_{v-(v-1)}^{(i)}), f_{v-(v-2)}(\mathbf{x}_{v-(v-2)}^{(i)}), \dots, f_{v-1}(\mathbf{x}_{v-1}^{(i)})). \quad (4)$$

Then, for each subsequent temporal period $v' = v + 1, \dots, V$ we construct a dataset $\{(id_{v'}^{(i)}, \ddot{\mathbf{x}}_{v'}^{(i)}, y_{v'+1}^{(i)})\}_{i=1}^{n_{v'}}$, where $\ddot{\mathbf{x}}_{v'}^{(i)} = (\mathbf{x}_{v'}^{(i)}, \mathbf{risk}_{v'-1}^{(i)})$. A simple example of this process, for a single risk estimate, is illustrated in Figure 1.

Correspondingly, classifiers $f_{v'}(\cdot) : \text{for } v' = v + 1, \dots, V$ can be induced on the risk estimate-containing datasets.

C. Missing Feature Estimators

One issue with longitudinal data is that features measured during earlier time periods may not be measured at subsequent periods. This presents a slight problem in the context of inverse classification because we would like to induce classifiers on the same set of features from one time period to the next.

¹We do, however, compare to this naive temporal linkage strategy in our experiments.

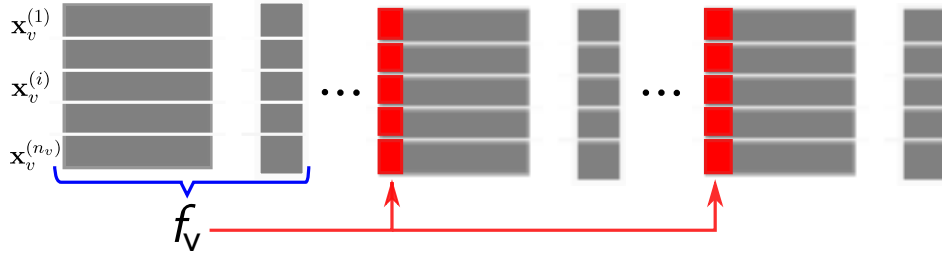


Fig. 1: An illustration of the addition of past risk estimates as features in future datasets, using estimates from a single period in time only. The risk estimate features are illustrated in red.

More importantly, we would like to optimize over the same set of \mathcal{F}_D features at each time period, observing how such optimization effects risk estimates in the future; thus we require that classifier induction take place on the same set of \mathcal{F}_D features.

Therefore, we propose to use so-called *missing feature-value estimators* to estimate the missing feature values at future time periods. Such estimators will be induced using the known features at visit v to induce classifiers or regressors, depending upon whether the missing feature is binary or continuous. Therefore, we will define the full set of features \mathcal{F} at visit $v = 1$ to ensure there is at least one time period from which we can always estimate future feature values. Using these estimators, we will estimate the instances' missing feature values at visit v . Our proposition to use such estimators is explored in the experiments section.

To formalize this notion, let $\mathcal{F}_{v=1}$ denote the feature set for the first time period in the longitudinal study. With this in mind, the following relation holds for the data in our longitudinal context:

$$\{\mathcal{F}_{v=\tilde{v}} \subseteq \mathcal{F}_{v=1} : \tilde{v} = 2, \dots, V\}. \quad (5)$$

The specification in (5) indicates that the data at visit $v = 1$ is sufficiently capable of providing estimates for the missing feature values at any subsequent time period because the features used at all future time periods are subsets or proper subsets (indicating no missing features) of the features defined at $v = 1$.

With this in mind, let M_v denote the set of missing features at visit v , where $M_v = \{\mathcal{F}_{v=1} \setminus \mathcal{F}_v\}$. Permit the instances' subvectors at $v = 1$ to be indexed using M_v , i.e., $\{\mathbf{x}_{M_v, v=1}^{(i)}\}_{i=1}^{n_{v=1}}$. Then, we define the estimation of missing features, using missing feature-value estimators, as follows:

$$x_{m,v}^{(i)} = \varepsilon_m(\mathbf{x}_v^{(i)}) : \text{for } m \in M_v \quad (6)$$

where $\varepsilon_m(\cdot)$ denotes a classifier or regressor induced on $\{\mathbf{x}_{\mathcal{F}_v, v=1}^{(i)}\}_{i=1}^{n_{v=1}}$ to learn the mapping to each $m \in M_v$, i.e., $\varepsilon_m : \{\mathbf{x}_{\mathcal{F}_v, v=1}^{(i)}\}_{i=1}^{n_{v=1}} \rightarrow \{x_{m,v}^{(i)}\}_{i=1}^{n_{v=1}}$.

Subsequently, let

$$\check{\mathbf{x}}_v^{(i)} = \left(\check{\mathbf{x}}_v^{(i)}, \left(x_{m,v}^{(i)} \right)_{m \in M_v} \right) \quad (7)$$

where $\check{\mathbf{x}}_v^{(i)}$ indicates an instance at visit v composed of the original feature values, missing feature values, and historical risk estimates. Note that we can now refer to $\check{\mathbf{x}}_v^{(i)}$ using the index sets U, I, D defined in a previous subsection.

Similarly, we can induce classifiers $\check{f}_{\tilde{v}}(\cdot) : \text{for } \tilde{v} = 2, \dots, V$ on data containing both past risk estimates and missing feature-value estimates.

D. Inverse Classification Formulation

With the discussion of past risk estimate and missing feature-values resolved, we can formulate our longitudinal inverse classification framework as:

$$\begin{aligned} \min_{\check{\mathbf{x}}_{D,v}} \check{f}_v(\check{\mathbf{x}}_{U,v}, H(\check{\mathbf{x}}_{U,v}, \check{\mathbf{x}}_{D,v}), \check{\mathbf{x}}_{D,v}) \quad (8) \\ \text{s.t. } \mathbb{C}_v(\check{\mathbf{x}}_{D,v} - \bar{\mathbf{x}}_{D,v}) \leq B_v \\ l_{j,v} \leq \check{x}_{j,v} \leq u_{j,v}, \text{ for } j \in D, \end{aligned}$$

where $\mathbb{C}(\mathbf{z}) = \sum_{j \in D} c_j^+(z_j)_+ + c_j^-(z_j)_-$, such that $(z)_+ = \max\{0, z\}$ and $(z)_- = \max\{0, -z\}$, is a cost function that measures the deviation of the directly changeable feature values $\check{\mathbf{x}}_{D,v}$ from their original values $\bar{\mathbf{x}}_{D,v}$ and applies a user-specified cost $c_j : j \in D$ to these deviations. The costs c_j allow for a patient to specify which feature values they would prefer to change (e.g., willing to exercise more than a change to diet). These cost-changes are summed together and subjected to a budget constraint B_v that controls the extent of changes recommended, thus avoiding over-radical recommendations (e.g., exercise 1000 hours in a week) and further allows patients and clinicians to express preferences. The $l_{j,v}(u_{j,v}) : j \in D$ are lower (upper) bounds that ensure a recommendation makes sense in the real world – e.g., prevents negative dietary intake from being recommended. Note that in (8) $H(\check{\mathbf{x}}_{U,v}, \check{\mathbf{x}}_{D,v})$ is used instead of $\check{\mathbf{x}}_{I,v}$ to account for how changes made to the directly changeable features affect the indirectly changeable features.

IV. EXPERIMENTS AND RESULTS

In this section we propose and execute several experiments that assess the viability of our proposed use of past risk estimates, missing feature value estimators, and our specified longitudinal inverse classification framework disclosed in (8) in being able to reduce CVD risk.

To such an end, we first disclose our proposed experiments and subsequently describe our CVD data. Next, we provide some details on the specific design choices made for our classifier, optimizer, and indirect feature value estimator. Following these disclosures, we report and discuss the results of our experiments.

A. Evaluation Specifics

1) *Experiments*: We propose the following three experiments to assess the viability of our proposed methods and corresponding inverse classification framework: **(1)** To assess the viability of our proposed use of missing feature value estimators we compare the performance of various classifiers and regressors to that of a simple carry-forward procedure (baseline). The carry-forward procedure simply entails imputing the missing feature values using the instances’ known feature values at $v = 1$. In other words, the values are carried forward from $v = 1$ to v . **(2)** We explore our proposed use of previous risk estimates as features in subsequent visits by inducing classifiers on data containing such estimates, inducing classifiers on data containing carried-forward past-visit instance feature vectors, and comparing the predictive performance between the two methods. The carry-forward procedure replaces our original definition of risk in (4) by $(\mathbf{x}_{v-(v-1)}^{(i)}, \dots, \mathbf{x}_{v-1}^{(i)}, \mathbf{x}_v^{(i)})$. **(3)** We examine the effects of applying our longitudinal inverse classification procedure to different temporal observation periods on future CVD risk. Particularly, we examine the effect of a single application of inverse classification applied early vs repeated applications of inverse classification on predicted CVD risk observed in future periods.

2) *Dataset Description*: To perform the experiments outlined in the previous subsection we use ARIC (atherosclerosis risk in communities study) data, a freely available dataset (subject to approval) from BioLINCC. Derivation of longitudinal data from this study produced $V = 3$ visit datasets having the statistics outlined in Table I, with visit four being used to define the CVD outcome (y) associated with the third visit ($v = 3$). The features defined in $v = 1$, along with additional

Dataset	Instances	Feats/Missing	$y_{v+1} = 1$
$v = 1$	12223	122/0	232
$v = 2$	11057	98/24	249
$v = 3$	9883	74/48	231

TABLE I: Dataset descriptors.

dataset details, such as the $\mathcal{F}_U, \mathcal{F}_I, \mathcal{F}_D$ features can be found in the Appendix section at the end of this work.

3) *Classifier, Optimizer, and Indirect Feature Value Estimator*: Our longitudinal inverse classification framework disclosed in (8) can be optimized by applying various gradient descent methods if $\hat{f}_v(\cdot)$ is differentiable and has an L -Lipschitz continuous gradient. Otherwise, various heuristics such as genetic algorithms, simulated annealing, etc. can be applied. In our inverse classification experiments we use support vector machines (SVM) for $\hat{f}_v(\cdot)$, which is both differentiable and has an L -Lipschitz continuous gradient. Thus, we use projected

gradient descent (PGD) [23], [24] to optimize the objective function in (8) to obtain recommendations. We elect to use kernel regression as our indirect feature estimator since it is also differentiable and provided good predictive performance experimentally.

B. Experiment 1: Missing Feature Estimators

To justify our use of missing feature-value estimators, we select three known features at $v = 2$ at random and treat them as if they are missing. Subsequently, we induce missing feature-value estimators with $v = 1$ data, using the remaining known features at $v = 2$ (i.e., sans the three randomly selected features). We then estimate the three features at $v = 2$ using the induced estimators and derive predictive performance measures by comparing the estimates to the actual values (AUC for categorical features and MSE for continuous). As stated during our description of this experiment, we compare the estimators to a “carry-forward” procedure. The results of this experiment are disclosed in Table II.

Feat/type \ Method	Alcohol/cont	Statin Use/bin	Hematocrit/cont
Carry	50.55	.579	7.55
RBF SVM	47.65	.50	9.31
Lin SVM	47.65	.50	9.31
CART	30.97	.569	2.37
kNN	29.37	.50	12.81
Log Reg	NA	.984	NA
Ridge	29.42*	NA	1.24

TABLE II: Missing feature estimation results given in MSE for continuous features and AUC for binary.

Table II first shows that, for each of the three features, there was at least one estimator that outperformed the “carry-forward” procedure, thus justifying the use of estimators over the naive carry-forward procedure. Second, we observe that ridge regression is the best performing continuous value estimator for “hematocrit” and a very close second for “alcohol consumed”. Therefore, in our inverse classification experiments, we used ridge regression to estimate continuous features. Finally, we found that logistic regression was best able to estimate the categorical feature “statin use” and thus use logistic regression to estimate categorical features in our inverse classification experiments.

C. Experiment 2: Risk Estimates vs. Feature Vectors

To create a temporal linkage between longitudinal datasets, we proposed to use past risk estimates as features in future visit datasets. To assess the viability of this method we induce classifiers on datasets containing these features, also inducing classifiers on datasets containing feature vectors that are carried forward from past visits. We then compare the predictive performance between the two methods to assess the viability of our proposed strategy. The results of this comparison-based experiment are presented in Figure 2.

Figure 2 shows that the predictive performance of the two methods is incredibly similar. At $v = 3$, the carry-forward

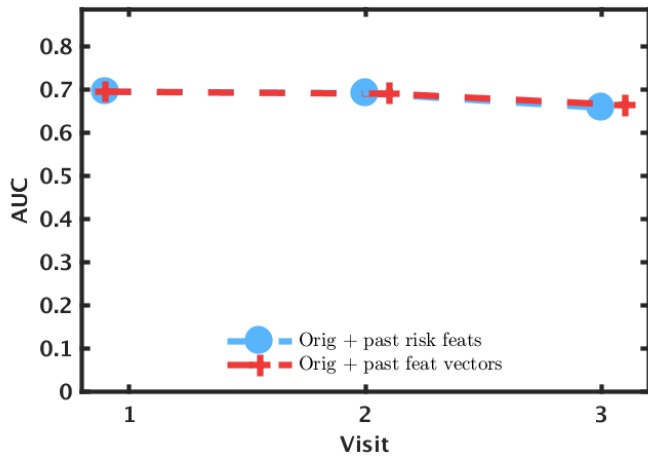


Fig. 2: Performance of classifiers induced using past risk estimates vs. classifiers induced using historical feature vectors, by visit.

procedure enjoys a $< .01$ AUC benefit over our proposed past risk method. Since the two methods perform virtually identical we elect to adopt our proposed past risk method, which provides a much more compact representation than the carry-forward procedure (i.e., past risk represents each past dataset as a single value, whereas the carry-forward procedure uses the entire feature vector of each past visit).

D. Experiment 3: Longitudinal Inverse Classification

To examine the effects of longitudinal inverse classification, we execute the inverse classification procedure in two different ways:

- We execute the inverse classification procedure at $v = 1$. We then use these optimized D feature-values as the D feature values at $v = 2$ and $v = 3$ and predict their probability of CVD at each of these visits. This represents instances' continued implementation of their personalized recommendation from $v = 1$. Additionally, we use estimates of historical risk at $v = 2$ and $v = 3$ based on the optimized D values at $v = 1$. We report the results in terms of average probability of CVD.
- We execute the same procedure as in (a.), above, but also apply a similar procedure at $v = 2$ using the newly optimized D feature-values as the D feature-values for the corresponding instances in the $v = 3$ dataset. This represents instances' that make initial improvements toward a beneficial outcome, but then also make additional, follow-up improvements in an attempt to further reduce their probability of CVD. We also report these results in terms of average probability of CVD.

Figure 3 shows the results of (a.) in black and (b.) in blue, which can be compared to the original predicted probability of CVD in red.

Examining the results of both (a.) and (b.) at $v = 2$, we can see that both were able to reduce the average probability

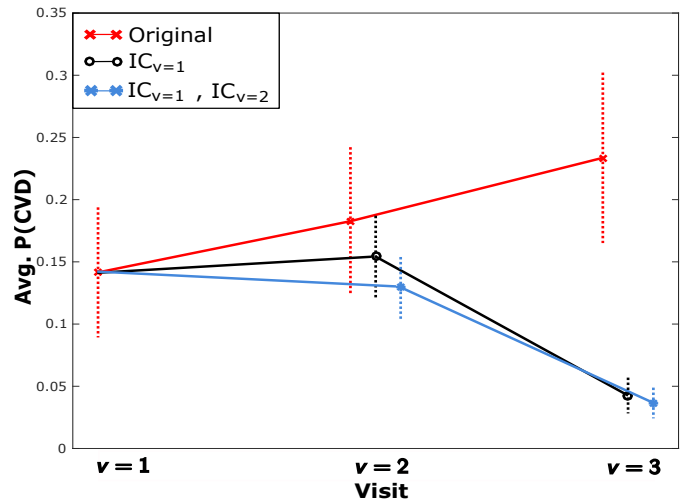


Fig. 3: Longitudinal inverse classification results by visit with a budget $B = 2$. Red represents average probability of CVD without applying inv. Black indicates the average probability results after applying inverse classification at $v = 1$. Blue indicates the average probability results after applying inverse classification at $v = 1$ and $v = 2$.

of CVD; the result of (b.), expressed in blue, suggests that a second application of the inverse classification process further reduces the average probability of CVD.

Interestingly, however, the results of (a.) and (b.) at $v = 3$ show that, while (a.) and (b.) kept the average predicted probability of CVD low, their difference in average predicted outcome is very comparable ((b.) is slightly less than (a.)). This is surprising because it suggests that applying a sufficient amount of initial effort to improve ones probability of a good outcome has a comparable “long-term” effect to application of additional CVD-mitigating effort later on, perhaps indicating/reinforcing the notion that interventions are best made earlier rather than later (by virtue of the “diminishing effect” observed).

V. CONCLUSIONS

Cardiovascular disease is a major health concern world wide. Personalized medicine provides an avenue by which CVD risk can be mitigated, thus potentially improving the health outlook for millions. Much of the on-going CVD research on personalized medicine focuses on the discovery of genetic biomarkers, but does not address personalized lifestyle adjustments that may reduce CVD risk. In this work, we proposed longitudinal inverse classification that makes personalized lifestyle recommendations that mitigate the risk of CVD, also addressing the sub-problems of missing features and the creation of temporal linkages among longitudinal datasets. We show that our proposed methods to address these two subproblems are reasonable and that longitudinal inverse classification can be used to reduce the long-term risk of CVD. Furthermore, we found that personalized lifestyle adjustments

adopted earlier are approximately as beneficial as repeated lifestyle adjustment made over several time periods.

REFERENCES

[1] E. J. Benjamin, P. Muntner, and M. S. Bittencourt, "Heart disease and stroke statistics-2019 update: a report from the american heart association," *Circulation*, vol. 139, no. 10, pp. e56–e528, 2019.

[2] A. H. Association *et al.*, "Heart disease and stroke statistics 2019 at-a-glance," *online*, 2019.

[3] G. A. Mensah, G. S. Wei, P. D. Sorlie, L. J. Fine, Y. Rosenberg, P. G. Kaufmann, M. E. Mussolino, L. L. Hsu, E. Addou, M. M. Engelgau *et al.*, "Decline in cardiovascular mortality: possible causes and implications," *Circulation research*, vol. 120, no. 2, pp. 366–380, 2017.

[4] N. J. Schork, "Personalized medicine: time for one-person trials," *Nature*, vol. 520, no. 7549, pp. 609–611, 2015.

[5] L. R. Cardon and T. Harris, "Precision medicine, genomics and drug discovery," *Human molecular genetics*, vol. 25, no. R2, pp. R166–R172, 2016.

[6] D. M. Minich and J. S. Bland, "Personalized lifestyle medicine: relevance for nutrition and lifestyle recommendations," *The Scientific World Journal*, vol. 2013, 2013.

[7] "Taking personalized medicine to heart," *Nature Medicine*, vol. 24, no. 2, p. 113, 2018.

[8] M.-S. Lee, A. J. Flammer, L. O. Lerman, and A. Lerman, "Personalized medicine in cardiovascular diseases," *Korean circulation journal*, vol. 42, no. 9, pp. 583–591, 2012.

[9] J. A. Leopold and J. Loscalzo, "Emerging role of precision medicine in cardiovascular disease," *Circulation research*, vol. 122, no. 9, pp. 1302–1315, 2018.

[10] G. Currie and C. Delles, "Precision medicine and personalized medicine in cardiovascular disease," in *Sex-Specific Analysis of Cardiovascular Function*. Springer, 2018, pp. 589–605.

[11] R. Kahn, R. M. Robertson, R. Smith, and D. Eddy, "The impact of prevention on reducing the burden of cardiovascular disease," *Circulation*, vol. 118, no. 5, pp. 576–585, 2008.

[12] N. Razavian, J. Marcus, and D. Sontag, "Multi-task prediction of disease onsets from longitudinal lab tests," *arXiv preprint arXiv:1608.00647*, 2016.

[13] X. Wang, D. Sontag, and F. Wang, "Unsupervised learning of disease progression models," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2014, pp. 85–94.

[14] C. C. Aggarwal, C. Chen, and J. Han, "The inverse classification problem," *Journal of Computer Science and Technology*, vol. 25, no. May, pp. 458–468, 2010.

[15] C. L. Chi, W. N. Street, J. G. Robinson, and M. A. Crawford, "Individualized patient-centered lifestyle recommendations: An expert system for communicating patient specific cardiovascular risk information and prioritizing lifestyle options," *Journal of Biomedical Informatics*, vol. 45, no. 6, pp. 1164–1174, 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.jbi.2012.07.011>

[16] M. V. Mannino and M. V. Koushik, "The cost minimizing inverse classification problem : A algorithm approach," *Decision Support Systems*, vol. 29, no. 3, pp. 283–300, 2000.

[17] C. Yang, W. N. Street, and J. G. Robinson, "10-year CVD risk prediction and minimization via inverse classification," in *Proceedings of the 2nd ACM SIGHIT symposium on International health informatics - IHI '12*, 2012, pp. 603–610.

[18] D. Barbella, S. Benzaid, J. Christensen, B. Jackson, X. V. Qin, and D. Musicant, "Understanding support vector machine classifications via a recommender system-like approach," in *Proceedings of the International Conference on Data Mining*, 2009, pp. 305–11.

[19] P. C. Pendharkar, "A potential use of data envelopment analysis for the inverse classification problem," *Omega*, vol. 30, no. 3, pp. 243–248, 2002.

[20] M. T. Lash, Q. Lin, W. N. Street, J. G. Robinson, and J. Ohlmann, "Generalized inverse classification," in *2017 SIAM International Conference on Data Mining (SDM)*. SIAM, 2017, pp. 162–170.

[21] M. T. Lash, Q. Lin, W. N. Street, and J. G. Robinson, "A budget-constrained inverse classification framework for smooth classifiers," in *Data Mining Workshops (ICDMW), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1184–1193.

[22] J. Platt, "Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods," *Advances in Large Margin Classifiers*, vol. 10, no. 3, pp. 61–74, 1999.

[23] Y. Nesterov, "Gradient methods for minimizing composite objective function," *Mathematical Programming, Series B*, vol. 140, pp. 125–161, 2013.

[24] S. Ghadimi and G. Lan, "Stochastic first- and zeroth-order methods for nonconvex stochastic programming," *SIAM Journal on Optimization*, vol. 23, pp. 2341–2368, 2013.

APPENDIX

The following tables enumerate the variables used in this study at $v = 1$.

Feature Name
Insulin (uu-ml), Height (cm), Age, Peripheral Artery Disease, Peripheral Artery Disease (definition 2), Plaque/shadowing in either internal, Plaque in either internal carotid, Cholesterol lowering med (last 2 weeks), Hypertension (definition 5), Education level, Diabetes, Age when menopause began, Menopause status, Ever smoked cigarettes, High blood pressure med (past 2 weeks), Agina-chest pain med (past 2 weeks), Heart rhythm control med (past 2 weeks), Heart failure med (past 2 weeks), Blood thinning med (past 2 weeks), Blood sugar med (past 2 weeks), Stroke med (past 2 weeks), Walking leg pain med (past 2 weeks), Headache or cold med (past 2 weeks), Pain meds (past 2 weeks), Gender, Race, Years smoked cigarettes

ST. I: Unchangeable features \mathcal{F}_U for the ARIC CVD dataset.

Feature Name: σ
BMI (Body Mass Index), Recalibrated HDL cholesterol (mg/dl), Re-calibrated LDL cholesterol (mg/dl), Total cholesterol (mmol/L), Total triglycerides (mmol/L), 2nd and 3rd systolic blood pressure (avg.), 2nd and 3rd systolic blood pressure (avg.) Num 2, Waist girth (cm), Hip girth (cm), Heart rate, White blood count, Apolipoprotein AI(mg-dl), Apolipoprotein B (mg-dl), Apolp(A) Data (ug-ml), Ankle-brachial index (Def 4), FV(1)/FVC Predicted (%), FEV(1) (L), FVC (L), Hematocrit, Hemaglobin, Platelet count, Neutrophils, Neutrophil bands, Lymphocytes, Monocytes, Eosinophils, Basophils, APTT Value, VIII: C Value, Fibrinogen Value, VII Value, ATIII Value, Protein: C Value, VWF Value, Cornell voltage (uV), Waist-hip ratio, Vegetable fat (% kcal), Carbs (% kcal), Alcohol (% kcal), Omega fatty acid (g), Calf girth (cm), Subcaps measure 2 (mm), Triceps measure 2 (mm), Uric acid (mg-dl), Total protein (gm-dl), Albium (gm-dl), Phosphorus (mg-dl), Magnesium (meq-l), Calcium (mg-dl), Urea nitrogein (mg-dl), Potassium (mmol-l), Sodium (mmol-l), Creatinine (mg-dl), Weight (lb), Total fat (% kcal), Saturate fatty acid (% kcal), Protein (% kcal), Polyunsaturated fatty acid (% kcal), Monounsaturated fatty acid (% kcal), Total fat (g)

ST. II: Indirectly changeable features \mathcal{F}_I for the ARIC CVD dataset.

c^+ / c^-	Feature:Cost
c^+	Dark or grain breads: 3, Peanut butter: 4, Nuts: 5, Other(prunes,avocado): 5, Vegetables: 6, Fruit: 6, Fiber: 7, Vegetable fat: 5, Polyunsaturated fat: 5
c^-	Liver: 8, White carbs: 6, Fish: 9, Cereal: 4, Cigarettes: 9, Caffeine: 7, Carbs: 7, Cholesterol: 6, Sodium: 7, Animal fat: 7, Saturated fat: 6
c^+ / c^-	Exercise hours: 10, Alcohol: 9

ST. III: Directly changeable features \mathcal{F}_D for the ARIC CVD dataset.