Adaptive covariate acquisition for minimizing total cost of classification

Daniel Andrade¹ · Yuzuru Okajima²

Received: 21 February 2020 / Revised: 26 January 2021 / Accepted: 10 February 2021 /
Published online: 18 April 2021
© The Author(s) 2021

Abstract

In some applications, acquiring covariates comes at a cost which is not negligible. For example in the medical domain, in order to classify whether a patient has diabetes or not, measuring glucose tolerance can be expensive. Assuming that the cost of each covariate, and the cost of misclassification can be specified by the user, our goal is to minimize the (expected) total cost of classification, i.e. the cost of misclassification plus the cost of the acquired covariates. We formalize this optimization goal using the (conditional) Bayes risk and describe the optimal solution using a recursive procedure. Since the procedure is computationally infeasible, we consequently introduce two assumptions: (1) the optimal classifier can be represented by a generalized additive model, (2) the optimal sets of covariates are limited to a sequence of sets of increasing size. We show that under these two assumptions, a computationally efficient solution exists. Furthermore, on several medical datasets, we show that the proposed method achieves in most situations the lowest total costs when compared to various previous methods. Finally, we weaken the requirement on the user to specify all misclassification costs by allowing the user to specify the minimally acceptable recall (target recall). Our experiments confirm that the proposed method achieves the target recall while minimizing the false discovery rate and the covariate acquisition costs better than previous methods.

Keywords Bayesian decision framework · Generalized additive models · Cost-sensitive classification · Recall guarantees

Editor: Jesse Davis.

✉ Daniel Andrade
andrade@hiroshima-u.ac.jp
Yuzuru Okajima
y-okajima@nec.com

¹ Education and Research Center for Artificial Intelligence and Data Innovation, Hiroshima University, Hiroshima, Japan

² Central Research Laboratories, NEC, Kawasaki, Japan

1 Introduction

In some applications, acquiring covariates comes at a cost which is not negligible. For example, in the medical domain, in order to classify whether a patient has diabetes or not, measuring glucose tolerance can be expensive. On the other hand, glucose tolerance can be a good indicator for diabetes, i.e. it increases our chance of predicting diabetes (or its absence) correctly.

The example illustrates that in the medical domain we often have to strike a balance between classification accuracy and the cost of acquiring the covariates. A rational criteria to decide on the best trade-off is to minimize the *expected total cost of classification*: expected cost of misclassification plus the cost of acquired covariates.

In the first part of this article, we formalize the optimization of the expected total cost of classification using the (conditional) Bayes risk and describe the optimal solution using a recursive procedure. However, it turns out that the procedure is computationally infeasible due to basically two factors: (1) calculating the Bayes risk requires to estimate a high dimensional integral, (2) the number of different covariate acquisition paths is exponential in the number of covariates.

As a consequence, we introduce two assumptions: (1) the optimal classifier can be represented by a generalized additive model (GAM), (2) the optimal sets of covariates are limited to a sequence of sets of increasing size. We show that under these two assumptions, a computationally efficient solution exists.

Our framework requires that the user can specify the cost of misclassification: the false positive cost (cost of wrongly classifying a healthy person as having diabetes), and false negative cost (cost of classifying a diabetes patient as healthy). However, we show that the requirement on the user to specify the false negative cost can be replaced by the specification of a lower bound on the recall. This is motivated by the medical domain where it is more common to specify the minimally acceptable recall (target recall), rather than specifying the false negative cost.

Our main contributions are as follows:

1. We describe the optimal solution for minimizing the expected total cost of classification, which has not been clarified in previous works like Dulac-Arnold et al. (2012) and Shim et al. (2018).
2. We prove that for a GAM, the estimation of the (conditional) Bayes risk reduces to a one dimensional density estimation and integral which can be solved computationally efficiently.
3. We propose an effective heuristic to estimate an optimal monotone increasing sequence of covariate sets by learning the regression coefficients of GAM with a group lasso penalty.
4. We prove that our framework can be used to guarantee a user-specified recall level, like 95% which is common in the medical domain.
5. We show on four medical datasets that the proposed method can lead to lower total cost of classification than the previous works in Ji and Carin (2007), Dulac-Arnold et al. (2012), Xu et al. (2012), Nan and Saligrama (2017) and Shim et al. (2018). Furthermore, evaluation under the requirement of a target recall shows that the proposed method achieves the target recall while minimizing the remaining costs and false discovery rate (FDR) better than previous works.

This article extends our preliminary work in Andrade and Okajima (2019) as follows:

- Replacing the linear classifier with GAM to allow for non-linear relations between the covariates and the response (see Sect. 4.1).¹
- Allowing for the specification of a minimal recall (see Sect. 6).
- Determining the sequence of covariate sets by the solution path of a group lasso penalized convex optimization problem (see Sect. 5.2).²
- Additional experimental evaluations on two more medical datasets, and additional analysis in terms of false positive costs, recall, weighted average, and runtime (see Sect. 7). Furthermore, instead of symmetric misclassification costs, which were used in Andrade and Okajima (2019), our evaluation here focuses on asymmetric misclassification costs that are more realistic in practice.

In the next section, we formalize the optimal decision procedure to achieve, in expectation, the lowest total cost of classification. In Sect. 3, we introduce a (non-adaptive) method that minimizes an upper of the lowest achievable total cost, which we extend in Sect. 4 to an adaptive method. In Sect. 5, we explain two approximations for finding a sequence of monotone increasing covariate sets that is used by the proposed method. In Sect. 6, we show how the proposed framework can also be used for guaranteeing a target recall. Extensive empirical evaluations of our proposed method and previous methods are provided in Sect. 7. In Sect. 8, we give a concise review of related work. Finally, in Sect. 9, we summarize our findings.

2 A cost rational selection criteria

In this section, we clarify and formally describe the optimal decision for acquiring covariates and classification. For that purpose we appropriately extend the definitions of a Bayes procedure which are known to be optimal for classification, see e.g. Anderson (2003), Berger (2013).

Let $L := \{l_1, \dots, l_c\}$ denote the set of class labels, and c_{y,y^*} the cost of classifying a sample as class y^* , when the true label is y . A decision procedure $\delta^* : \mathbb{R}^p \rightarrow L$ for which

$$\forall \delta : \mathbb{E}_{\mathbf{x},y}[c_{y,\delta(\mathbf{x})}] \geq \mathbb{E}_{\mathbf{x},y}[c_{y,\delta^*(\mathbf{x})}]$$

is called a Bayes procedure. The following procedure δ^* is a Bayes procedure [for a proof see, for example, Theorem 6.7.1 in Anderson (2003)]:

$$\begin{aligned} \delta^*(\mathbf{x}) &= \arg \min_{y^* \in L} \sum_{y \in L} p(y|\mathbf{x}) \cdot c_{y,y^*} \\ &= \arg \min_{y^* \in L} \mathbb{E}_y[c_{y,y^*}]. \end{aligned} \quad (1)$$

¹ In Andrade and Okajima (2019), β was the regression coefficient vector. Here β takes now the form of a matrix where each column specifies the coefficients of the basis functions corresponding to one covariate.

² In Andrade and Okajima (2019) we introduced only the greedy covariate acquisition approach which is explained here in Sect. 5.1.

The expected misclassification cost of the Bayes procedure, i.e. $\mathbb{E}_{\mathbf{x},y}[c_{y,\delta^*(\mathbf{x})}]$, is called the Bayes risk.

Let us denote by $V := \{1, \dots, p\}$ the index set of covariates with $V \cap L = \emptyset$. We denote the Bayes procedure for classifying a sample based only on the covariates $S \subseteq V$ by $\delta_S^* : \mathbb{R}^{|S|} \rightarrow L$. That means

$$\delta_S^*(\mathbf{x}_S) = \arg \min_{y^* \in L} \sum_{y \in L} p(y|\mathbf{x}_S) \cdot c_{y,y^*}. \quad (2)$$

When it is clear from the context, we drop the index on δ_S^* , and just write $\delta^*(\mathbf{x}_S)$ instead of $\delta_S^*(\mathbf{x}_S)$.³

2.1 Optimal procedure

The classical definition of Bayes procedure does not consider the cost of covariate acquisition, and assumes that all covariates are acquired at once. Therefore, let us first formally extend the definition appropriately.

We use the following definition of a decision procedure.

Definition 1 A function of the form

$$\pi : \mathbb{R}^p \times 2^V \rightarrow L \cup V,$$

which fulfills, $\forall \mathbf{x} \in \mathbb{R}^p, S \subseteq V$:

$$\pi(\mathbf{x}, S) = \pi(\mathbf{x} \odot \mathbf{1}_S, S), \quad (3)$$

$$\pi(\mathbf{x}, S) \in L \cup (V \setminus S), \quad (4)$$

is called a decision procedure.⁴

The condition in Eq. (3) means that a decision procedure uses only the covariates that are indexed by S ; the condition in Eq. (4) means that a decision procedure cannot select a covariate that is already in S . In summary, the decision procedure $\pi(\mathbf{x}, S)$ either classifies the current sample, or selects a new covariate based on the observations $\mathbf{x}_S$. To simplify the notation, we write $\pi(\mathbf{x}_S)$ instead of $\pi(\mathbf{x}, S)$. Furthermore, we denote the cost of acquiring covariate i by c_i .

Given a sample $\mathbf{x}$ with class label y , we denote the loss of a decision procedure π as $l((\mathbf{x}, y), \pi)$. The loss can be computed recursively as follows. Let $l((\mathbf{x}, y), \pi, \emptyset) := l((\mathbf{x}, y), \pi, \emptyset)$, with

$$l((\mathbf{x}, y), \pi, S) = \begin{cases} c_{y,\pi(\mathbf{x}_S)} & \text{if } \pi(\mathbf{x}_S) \in L, \\ c_{\pi(\mathbf{x}_S)} + l((\mathbf{x}, y), \pi, S \cup \{\pi(\mathbf{x}_S)\}) & \text{else.} \end{cases} \quad (5)$$

If not stated otherwise, we assume that all costs are non-negative, i.e. $c_i \geq 0$, and $c_{y,y'} \geq 0$.

³ Remark about our notation: we denote by bold font a column vector, e.g. $\mathbf{x} \in \mathbb{R}^p$, and a column vector indexed by a set $A \subseteq V$ denotes the corresponding sub-vector, e.g. $\mathbf{x}_A \in \mathbb{R}^{|A|}$.

⁴ $\odot$ denotes the Hadamard product, and $\mathbf{1}_S \in \mathbb{R}^p$ is the vector that is one in all positions indexed by S , and zero otherwise.

Theorem 1 Let π^* be a decision procedure, such that for all $S \subseteq V$, we have

$$\pi^*(\mathbf{x}_S) = \arg \min_{i \in L \cup (V \setminus S)} \begin{cases} \mathbb{E}_y[c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_S] & \text{else.} \end{cases} \quad (6)$$

Then π^* is a Bayes procedure. That means for any other decision procedure π we have

$$\mathbb{E}_{\mathbf{x}, y} [l((\mathbf{x}, y), \pi^*)] \leq \mathbb{E}_{\mathbf{x}, y} [l((\mathbf{x}, y), \pi)].$$

The proof is given in the “Appendix”. We note that, if the covariates are discrete, we can formulate the problem as a stationary Markov decision process (MDP) where every policy leads to a terminal state (Zubek et al., 2004; Bayer-Zubek 2004). The Bayes procedure from Theorem 1 is then equivalent to the optimal policy defined by the Bellman updates with the discounting factor set to 1 (Russell & Norvig, 2003).

For continuous covariates, implementing the exact decision procedure π^* is, in general, intractable.⁵ The reason is that in order to recursively evaluate the loss, we need to evaluate a sequence of interchanging minimizations and expectations. To illustrate this, consider the following example with only two covariates, $V = \{1, 2\}$, we then have

$$\begin{aligned} \mathbb{E}_{\mathbf{x}, y} [l((\mathbf{x}, y), \pi^*)] = \min \bigg\{ & c_1 + \mathbb{E}_{\mathbf{x}_1} \left[\min \left\{ c_2 + \mathbb{E}_{\mathbf{x}_2} [C(\mathbf{x}_{\{1,2\}}) | \mathbf{x}_1], C(\mathbf{x}_1) \right\} \right], \\ & c_2 + \mathbb{E}_{\mathbf{x}_2} \left[\min \left\{ c_1 + \mathbb{E}_{\mathbf{x}_1} [C(\mathbf{x}_{\{1,2\}}) | \mathbf{x}_2], C(\mathbf{x}_2) \right\} \right], \\ & C(\mathbf{x}_\emptyset) \bigg\}, \end{aligned} \quad (7)$$

where the first line corresponds to first acquiring covariate 1 and then either classifying or acquiring covariate 2, the second line corresponds to first acquiring covariate 2, and the last line corresponds to not acquiring any covariate. Note that to simplify notation in the above example, we defined $C(\mathbf{x}_S)$ to be the expected cost of (mis-)classification when given $\mathbf{x}_S$, i.e.

$$C(\mathbf{x}_S) := \min_{i \in L} \mathbb{E}_y [c_{y,i} | \mathbf{x}_S].$$

Therefore, we propose two relaxations and corresponding methods named Cost-sensitive Covariate Selection (COS) and Adaptive Cost-sensitive Forward Selection (AdaCOS) in Sects. 3 and 4, respectively.

3 Cost-sensitive covariate selection (COS)

Our first relaxation is to pull-out all minimizations in the recursion of Eq. (6). Also see the example in Eq. (7) which explicitly writes out all minimizations. This leads to the following upper bound:

⁵ Note that our proposed general framework is applicable to both continuous and discrete variables, where the expectations are then either evaluated via integrals or sums. Though, the specific Gaussian approximation in Sect. 4.1, will be more appropriate for continuous variables.

Fig. 1 Shows all possible decisions, for an example with $V = \{1, 2\}$. Each edge represents one decision: either asking for the value of a covariate, or classifying the sample based on the observed covariates so far using the Bayes procedure. Each leaf shows the upper bound $\mathcal{U}$ of the optimal expected total cost of classification when using the covariates that were selected on the path from the root to the leaf

$$\mathbb{E}_{\mathbf{x}, y}[\mathcal{U}((\mathbf{x}, y), \pi^*)] \leq \min_{S \subseteq V} \left(\mathbb{E}_{\mathbf{x}_S, y}[c_{y, \delta^*(\mathbf{x}_S)}] + \sum_{i \in S} c_i \right).$$

In the following, we denote this upper bound by $\mathcal{U}$. However, directly trying to minimize $\mathcal{U}$ is still computationally difficult due to the exponential number of possible sets S . Note that this is similar to covariate selection in logistic regression like e.g. in Tibshirani (1996), O'Hara et al., (2009). Two important differences are that, in general, costs associated with covariates can be different from each other, and that the Bayes risk needs to be evaluated for all possible subsets $S \subseteq V$. The situation is illustrated in Fig. 1.

We denote the method selecting the set S^* that minimizes $\mathcal{U}$, by Cost-sensitive Covariate Selection (COS). In order to approximately find the set S^* we will use the methods described in Sect. 5. One disadvantage of COS is that it always selects the same set of covariates S^* for any sample, though for some samples less/more covariates might be sufficient/necessary for good classification accuracy.

4 Adaptive cost-sensitive forward selection

The method COS proposed in the previous section is not adaptive, i.e. it does not take into account the actually observed covariates in order to decide whether to proceed the acquisition of additional covariates, or whether to classify based on the covariates observed so far. However, as discussed in Sect. 2.1, without any additional assumptions estimating the optimal procedure π^* from Theorem 1 is computationally infeasible. We therefore introduce two assumptions:

1. The optimal set S of acquired covariates belongs to $\mathfrak{S} = \{S_1, S_2, \dots, S_q\}$, where $S_1 \subset S_2 \subset S_3 \dots S_q \subseteq V$. An approximation method for finding the optimal sequence of subsets $\mathfrak{S}$ will be described in Sect. 5.
2. The conditional class probability $p(y|\mathbf{x}_S)$, for $S \in \mathfrak{S}$, belongs to a logistic generalized additive model.

Fig. 2 Shows an example where $S_1 = \{2\}$, and $S_2 = \{2, 1\}$. Assumes we have not acquired any covariates, yet. In order to decide whether to start covariate acquisition, AdaCOS compares the current expected cost of misclassification marked in green with the expected future costs when acquiring covariates $\{x_2\}$ or $\{x_2, x_1\}$, marked in blue and orange, respectively

Fig. 3 Shows an example where $S_1 = \{3\}$, $S_2 = \{3, 1\}$, $S_3 = \{3, 1, 2\}$, and $S_4 = \{3, 1, 2, 4\}$. Assumes we have already acquired the covariates S_2 . In order to decide whether to proceed acquisition, AdaCOS compares the current expected cost of misclassification marked in green with the expected future costs when additionally acquiring covariates $\{x_2\}$ or $\{x_2, x_4\}$, marked in blue and orange, respectively

Before we proceed, let us introduce our definition of future costs. Let $A \subseteq V$ and $S \subseteq V \setminus A$, then we define

$$F_{\mathbf{x}_A}(S) := \overbrace{\mathbb{E}_{\mathbf{x}_S, y}[c_{y,\delta^*}(\mathbf{x}_{A \cup S}) | \mathbf{x}_A]}^{\text{(conditional) Bayes risk}} + \overbrace{\sum_{i \in S} c_i}^{\text{covariate costs}}. \quad (8)$$

$F_{\mathbf{x}_A}(S)$ is the expected total additional cost of classification when we have already acquired the covariates A , and are planning to acquire additionally the covariates S before classifying. In particular, the upper bound $\mathcal{U}$ can be expressed as $\min_{S \subseteq V} F_{\mathbf{x}_0}(S)$.

Our approximation of the Bayes procedure π^* from Theorem 1 is given in Algorithm 1. First, we acquire all covariates indexed by S_1 , and then check whether acquiring any additional covariates from $S_2 \setminus S_1, \dots, S_q \setminus S_1$ reduces the total cost of classification in expectation. If that is the case, we acquire the covariates in $S_2 \setminus S_1$, and proceed analogously. If the total cost of classification is not expected to decrease with more covariates, we stop and classify based on the covariates acquired so far. An example of the procedure, where $V = \{1, 2\}$, is shown in Fig. 2. Note the example in Fig. 2, corresponds to the right most path from the tree in Fig. 1. Another, different example, with $V = \{1, 2, 3, 4\}$, is shown in Fig. 3.

Algorithm 1: Adaptive Cost-sensitive Forward Selection (AdaCOS)
for classifying a test sample.

Input: $S_1, \dots, S_q$
 $S_0 := \emptyset$
for $i \in \{1, \dots, q-1\}$ **do**
 acquire $\mathbf{x}_{S_i \setminus S_{i-1}}$
 if $\forall j \in \{i+1, \dots, q\} : F_{\mathbf{x}_{S_i}}(S_j \setminus S_i) \geq F_{\mathbf{x}_{S_i}}(\emptyset)$ **then**
 output class $\delta^*(\mathbf{x}_{S_i})$
 end
end

The algorithm is adaptive in the sense that the expected future costs $F_{\mathbf{x}_A}(S)$ depend on the covariates $\mathbf{x}_A$ observed so far. Therefore, we see that the effectiveness of the algorithm hinges on the non-trivial task of calculating $F_{\mathbf{x}_A}(S)$.

Finally, we note that in contrast to AdaCOS, the method COS does not use Algorithm 1, but selects only one fixed set of covariates $\hat{S} := \arg \min_{S \subseteq V} F_{\mathbf{x}_0}(S)$, which is used for all samples at test time.

4.1 Bayes risk estimation

The main challenge in evaluating the future costs is to estimate the multi-dimensional integral in $\mathbb{E}_{\mathbf{x}_S, y} [c_{y, \delta^*(\mathbf{x}_{AUS})} | \mathbf{x}_A]$. By assuming that the conditional class probability $p(y | \mathbf{x}_{AUS})$ can be modeled by a logistic generalized additive model, we will show that it is possible to reduce the multi-dimensional integral into a one-dimensional with an effective approximation.

The logistic generalized additive model is defined as follows. Given the regression coefficients $\boldsymbol{\beta} = (\boldsymbol{\beta}_1, \boldsymbol{\beta}_2, \dots, \boldsymbol{\beta}_p) \in \mathbb{R}^{s \cdot p}$, and intercept τ , the conditional class probability $p(y = 1 | \mathbf{x}, \boldsymbol{\beta}, \tau)$ is modeled as

$$p(y = 1 | \mathbf{x}_{AUS}, \boldsymbol{\beta}, \tau) = g\left(\tau + \sum_{i \in AUS} \boldsymbol{\beta}_i^T f_i(x_i)\right),$$

where g denotes the sigmoid function, and the non-linear transformations $f_i : \mathbb{R} \rightarrow \mathbb{R}^s$ are learned from the training data using penalized B splines, where s is the number of splines (Hastie et al., 2009).⁶

Furthermore, we assume that there are only two class labels $\{0, 1\}$, and that the costs of misclassification are larger than the cost of correct classification.⁷ We then have

⁶ For learning the B splines we make use of pyGAM available at <https://pygam.readthedocs.io/en/latest/>

⁷ This allows also for negative costs. In particular, a negative cost for $c_{y,y}$ means a benefit for correct classification.

$$\begin{aligned}
& \mathbb{E}_{\mathbf{x}_S, y} [c_{y, \delta^*(\mathbf{x}_{AUS})} | \mathbf{x}_A] \\
&= \mathbb{E}_{\mathbf{x}_S} \left[\sum_y c_{y, \delta^*(\mathbf{x}_{AUS})} p(y | \mathbf{x}_{AUS}) | \mathbf{x}_A \right] \\
&= \mathbb{E}_{\mathbf{x}_S} [c_{0, \delta^*(\mathbf{x}_{AUS})} p(y = 0 | \mathbf{x}_{AUS}) | \mathbf{x}_A] \\
&\quad + \mathbb{E}_{\mathbf{x}_S} [c_{1, \delta^*(\mathbf{x}_{AUS})} p(y = 1 | \mathbf{x}_{AUS}) | \mathbf{x}_A].
\end{aligned}$$

Since

$$\begin{aligned}
\delta^*(\mathbf{x}_{AUS}) &= \arg \min [p(y = 1 | \mathbf{x}_{AUS}) c_{1,0} + p(y = 0 | \mathbf{x}_{AUS}) c_{0,0}, \\
&\quad p(y = 0 | \mathbf{x}_{AUS}) c_{0,1} + p(y = 1 | \mathbf{x}_{AUS}) c_{1,1}],
\end{aligned}$$

we have, where we defined $p := p(y = 1 | \mathbf{x}_{AUS})$,

$$\begin{aligned}
\delta^*(\mathbf{x}_{AUS}) = 1 &\Leftrightarrow p c_{1,0} + (1 - p) c_{0,0} \geq (1 - p) c_{0,1} + p c_{1,1} \\
&\Leftrightarrow p(c_{1,0} - c_{1,1}) \geq (1 - p)(c_{0,1} - c_{0,0}) \\
&\Leftrightarrow \frac{p(c_{1,0} - c_{1,1})}{(1 - p)(c_{0,1} - c_{0,0})} \geq 1 \\
&\Leftrightarrow \frac{g(\sum_{i \in AUS} \beta_i^T f_i(x_i) + \tau) \cdot (c_{1,0} - c_{1,1})}{(1 - g(\sum_{i \in AUS} \beta_i^T f_i(x_i) + \tau)) \cdot (c_{0,1} - c_{0,0})} \geq 1 \\
&\Leftrightarrow e^{\sum_{i \in AUS} \beta_i^T f_i(x_i) + \tau} \geq \frac{c_{0,1} - c_{0,0}}{c_{1,0} - c_{1,1}} \\
&\Leftrightarrow \sum_{i \in AUS} \beta_i^T f_i(x_i) \geq \log\left(\frac{c_{0,1} - c_{0,0}}{c_{1,0} - c_{1,1}}\right) - \tau \\
&\Leftrightarrow \sum_{i \in S} \beta_i^T f_i(x_i) \geq \log\left(\frac{c_{0,1} - c_{0,0}}{c_{1,0} - c_{1,1}}\right) - \tau - \sum_{i \in A} \beta_i^T f_i(x_i) \\
&\Leftrightarrow z \geq z^*,
\end{aligned}$$

where we defined $z := \sum_{i \in S} \beta_i^T f_i(x_i)$, and $z^* := \log\left(\frac{c_{0,1} - c_{0,0}}{c_{1,0} - c_{1,1}}\right) - \tau - \sum_{i \in A} \beta_i^T f_i(x_i)$. We see that $\delta^*(\mathbf{x}_{AUS})$ depends only on z (random variable) and z^* (fixed). In the following, to simplify notation, let us denote by $h(z)$ the conditional distribution $p(z | \mathbf{x}_A)$. We thus have

$$\begin{aligned}
& \mathbb{E}_{\mathbf{x}_S} [c_{1, \delta^*(\mathbf{x}_{AUS})} p(y = 1 | \mathbf{x}_{AUS}) | \mathbf{x}_A] \\
&= \mathbb{E}_{\mathbf{x}_S} [c_{1, \delta^*(\mathbf{x}_{AUS})} g(\sum_{i \in S} \beta_i^T f_i(x_i) + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) | \mathbf{x}_A] \\
&= \mathbb{E}_z [c_{1, \delta^*(z, z^*)} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) | \mathbf{x}_A] \\
&= \int c_{1, \delta^*(z, z^*)} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) h(z) dz \\
&= c_{1,0} \int_{-\infty}^{z^*} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) h(z) dz \\
&\quad + c_{1,1} \int_{z^*}^{\infty} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) h(z) dz.
\end{aligned}$$

Analogously, we have

$$\begin{aligned}
& \mathbb{E}_{\mathbf{x}_S} [c_{0,\delta^*(\mathbf{x}_{AUS})} p(y=0|\mathbf{x}_{AUS})|\mathbf{x}_A] \\
&= \mathbb{E}_{\mathbf{x}_S} [c_{0,\delta^*(\mathbf{x}_{AUS})} (1 - g(\sum_{i \in S} \beta_i^T f_i(x_i) + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau))|\mathbf{x}_A] \\
&= \mathbb{E}_z [c_{0,\delta^*(z,z^*)} (1 - g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau))|\mathbf{x}_A] \\
&= c_{0,0} \int_{-\infty}^{z^*} (1 - g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau)) h(z) dz \\
&\quad + c_{0,1} \int_{z^*}^{\infty} (1 - g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau)) h(z) dz \\
&= c_{0,0} \int_{-\infty}^{z^*} h(z) dz - c_{0,0} \int_{-\infty}^{z^*} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) h(z) dz \\
&\quad + c_{0,1} \int_{z^*}^{\infty} h(z) dz - c_{0,1} \int_{z^*}^{\infty} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) h(z) dz.
\end{aligned}$$

Thus the remaining task is to evaluate the following integral

$$\begin{aligned}
& \int_{a'}^{b'} g(z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau) h(z) dz \\
&= \int_{a' + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau}^{b' + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau} g(u) h(u - \sum_{i \in A} \beta_i^T f_i(x_i) - \tau) du.
\end{aligned} \tag{9}$$

We assume that $h(z) = p(z|\mathbf{x}_A)$ can be well approximated by a normal distribution with mean μ_z and variance σ^2 . We defer the explanation of how to estimate μ_z and σ^2 to Sect. 4.1.1.

The integral in Eq. (9) has no analytic solution. One popular strategy is to approximate the sigmoid function g by the cumulative distribution function of the standard normal distribution Φ , as in Gaussian process classification (Rasmussen & Williams, 2006). However, it turns out that this approximation is not applicable here, since a' or b' is a finite real number in our case. Instead, we use here the fact that the sigmoid function can be well approximated with only a few number of linear functions. In order to facilitate notation, let us introduce the following constants:

$$\begin{aligned}
a &:= a' + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau, \\
b &:= b' + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau, \\
\mu &:= \mu_z + \sum_{i \in A} \beta_i^T f_i(x_i) + \tau.
\end{aligned}$$

Then we can write the integral in Eq. (9) as

$$\int_a^b g(u) \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(u-\mu)^2} du. \tag{10}$$

Let us define the following piece-wise linear approximation of the sigmoid function:

Fig. 4 Comparison of Sigmoid function approximations. For both, the linear function approximation (proposed) and the discrete bin approximation (Ji and Carin 2007), we set $\xi = 40$. For the normal CDF approximation we use $\Phi(\sqrt{\frac{\pi}{8}}u)$. Note that the proposed linear function approximation is almost indistinguishable from the Sigmoid function. In contrast the error of the normal approximation is large for around -3.0 and around 3.0 . The discrete approximation is obviously the worst

$$g(u) \approx \sum_{t=1}^{\xi+2} \left(1_{[b_{t-1}, b_t]}(u) (m_t u + v_t) \right),$$

where for $1 \leq t \leq \xi + 1$, we set $b_t := -10 + \frac{20}{\xi}(t - 1)$, and for $1 \leq t \leq \xi$, we set

$$m_{t+1} := \frac{g(b_{t+1}) - g(b_t)}{b_{t+1} - b_t}, \quad v_{t+1} := g(b_t) - m_{t+1} b_t,$$

and

$$b_0 := -\infty, m_1 := 0, v_1 := g(b_1), \\ b_{\xi+2} := +\infty, m_{\xi+2} := 0, v_{\xi+2} := g(b_{\xi+1}),$$

and ξ is the number of linear approximations, which is, for example, set to 40. A comparison with the approximation $\Phi(\sqrt{\frac{\pi}{8}}u)$ is shown in Fig. 4. That means for a relatively few number of linear approximations, we can achieve an approximation that is more accurate than the Φ -approximation. More importantly, as we show below, this allows for a tractable calculation of the integral in Equation (10), which is *not* the case when using the Φ -approximation. Then we have

$$\begin{aligned}
& \int_a^b g(u) \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(u-\mu)^2} du \\
&= \int_a^b \sum_{i=1}^{\xi+2} \left(1_{[b_{i-1}, b_i]}(u) (m_i u + v_i) \right) \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(u-\mu)^2} du \\
&= \sum_{i=1}^{\xi+2} m_i \int_{\max(a, b_{i-1})}^{\min(b, b_i)} u \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(u-\mu)^2} du \\
&\quad + v_i \Phi_{\max(a, b_{i-1})}^{\min(b, b_i)},
\end{aligned}$$

where we define $\Phi_l^o := \int_l^o \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(u-\mu)^2} du$, which can be well approximated with standard numerical libraries. The remaining integral can also be expressed by Φ using the substitution $u - \mu := r$, we have

$$\begin{aligned}
& \int_l^o u \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}(u-\mu)^2} du \\
&= \int_{l-\mu}^{o-\mu} r \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}r^2} dr + \mu \int_{l-\mu}^{o-\mu} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2\sigma^2}r^2} dr \\
&= \frac{\sigma}{\sqrt{2\pi}} \left(e^{-\frac{1}{2\sigma^2}(l-\mu)^2} - e^{-\frac{1}{2\sigma^2}(o-\mu)^2} \right) + \mu \Phi_l^o.
\end{aligned}$$

4.1.1 Estimation of μ_z and σ^2

Recall that we assumed that $p(z|\mathbf{x}_A)$ can be well approximated by a normal density with mean μ_z and variance σ^2 . In order to estimate μ_z and σ^2 , we propose to model z given $\mathbf{x}_A$ as a regression problem with additive noise, where z is the response variable, and $\mathbf{x}_A$ are the explanatory variables. In detail, for learning the regression model from the training data $\{\mathbf{x}^{(k)}\}_{k=1}^n$, we prepare a collection of response and explanatory variable pairs of the form $\{(z^{(k)}, \mathbf{x}_A^{(k)})\}_{k=1}^n$, where $z^{(k)} = \sum_{i \in S} \beta_i^T f_i(x_i^{(k)})$. We note that for training the regression model, we could additionally exploit unlabeled data as follows.⁸ Once we learned β_i and f_i from training data with class label information y , we can use all data, with and without class label, to create the data $\{(z^{(k)}, \mathbf{x}_A^{(k)})\}_{k=1}^{n+n_u}$, where n_u is the number of unlabeled samples.

For our experiments, we use a standard Bayesian linear regression model with a scaled inverse χ^2 distribution prior on the noise variance (Gelman et al., 2013).

⁸ Though, for the medical datasets we consider in Sect. 7, no unlabeled training data was available.

5 Cost-aware non-linear covariate selection

In the previous section, we assumed that the covariates are acquired in a specific sequence. In this section, we discuss two different approximation strategies for finding the optimal sequence of subsets $\mathfrak{S} = \{S_1, S_2, \dots, S_q\}$, where $S_1 \subset S_2 \subset S_3 \dots S_q \subseteq V$, such that the expected total cost of classification tends to minimal for a set $S \in \mathfrak{S}$.

5.1 Forward selection

We suggest to set $q := p + 1$, and use greedy forward selection as outlined in Algorithm 2.

Algorithm 2: Forward selection for finding subsets $S_1 \subset S_2, \dots, S_{p+1} \subseteq V$.

Input: $\{(y^{(k)}, \mathbf{x}^{(k)})\}_{k=1}^n$
 $S_1 := \emptyset$
for $i \in \{1, \dots, p\}$ **do**
 $S_{i+1} := \arg \min_{j \in V \setminus S_i} \mathbb{E}_{\mathbf{x}_{S_i}} [F_{\mathbf{x}_{S_i}}(\{j\})]$
end

Note that from the definition in Eq. 8, we have

$$\mathbb{E}_{\mathbf{x}_S} [F_{\mathbf{x}_S}(\{j\})] = \mathbb{E}_{\mathbf{x}_S} [\mathbb{E}_{x_j, y} [c_{y, \delta^*(\mathbf{x}_{S \cup \{j\}})} | \mathbf{x}_S]] + \sum_{i \in S} c_i.$$

We estimate $\mathbb{E}_{\mathbf{x}_S} [F_{\mathbf{x}_S}(\{j\})]$ using 10-fold cross-validation of the labeled training data $\{(y^{(k)}, \mathbf{x}^{(k)})\}_{k=1}^n$. In particular, for one fold $\mathcal{A}_f \subseteq \{1, \dots, n\}$, we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_S} [\mathbb{E}_{x_j, y} [c_{y, \delta^*(\mathbf{x}_{S \cup \{j\}})} | \mathbf{x}_S]] \\ &= \mathbb{E}_{\mathbf{x}_{S \cup \{j\}}, y} [c_{y, \delta^*(\mathbf{x}_{S \cup \{j\}})}] \\ &\approx \frac{1}{|\mathcal{A}_f|} \sum_{k \in \mathcal{A}_f} c_{y^{(k)}, \delta_f(\mathbf{x}_{S \cup \{j\}}^{(k)})}, \end{aligned}$$

where the model for the conditional probability $p(y | \mathbf{x}_{S \cup \{j\}})$ used by δ_f is trained using the samples in $\{1, \dots, n\} \setminus \mathcal{A}_f$. The final estimate is acquired by averaging over all folds.

An advantage of the above forward-selection procedure is that it uses an unbiased estimate of $\mathbb{E}_{\mathbf{x}_S} [F_{\mathbf{x}_S}(\{j\})]$, and assuming the variance is not too large, we can expect to find a good local minima.

However, there are several disadvantages. First, if the variance of the estimator is high, we might get stuck in a bad local minima. Second, the forward-selection procedure is extremely computationally expensive, and, as a consequence, unfeasible if p is large. Finally, a more subtle disadvantage is that it requires the full specification of the misclassification costs, i.e. the specification of $c_{0,1}$ and $c_{1,0}$. As a consequence, it is not applicable when we are provided only with $c_{0,1}$, which we will discuss in Sect. 6.

5.2 Group Lasso penalty

Some of the disadvantages of the feed-forward selection method can be overcome by jointly training the model for the conditional probability $p(y|\mathbf{x})$ with a sparsity-enforcing penalty on the regression coefficients. Here, this is possible since we assume a generalized additive model for $p(y|\mathbf{x})$.

In particular, we propose to acquire the sets $S_1 \subset S_2 \subset S_3 \dots S_q$ by using the search path of a penalized logistic loss function, subject to a post-hoc correction to ensure strict monotonicity. In detail, for different values of λ , we solve the following convex optimization problem

$$\underset{\beta, \tau}{\text{minimize}} - \sum_{k=1}^n \log p(y = y^{(k)} | \mathbf{x}^{(k)}, \beta, \tau) + \lambda \sum_{i=1}^p c_i \|\beta_i\|_2, \quad (11)$$

where $\beta = (\beta_1, \beta_2, \dots, \beta_p) \in \mathbb{R}^{s \cdot p}$. The group lasso penalty ensures that the regression coefficients β_i are either all zero or all non-zero (Hastie et al., 2015). Note that in Eq. (11) each group is scaled by c_i which ensures that the regression coefficient of covariates with high cost are penalized more.⁹ As a consequence, in order to be included into the final model, covariates with high cost are required to lower the negative log-likelihood term more than covariates with low costs. By inspecting the search path for different values of $\lambda_1 > \lambda_2 > \dots > \lambda_q$, we acquire the corresponding sets $S_1, S_2, \dots, S_q$. For λ , we choose 100 values evenly spaced on the log-scale between 0.0001 and 0.02, that is $q = 100$. In more detail, after training with λ_j , covariate i will be included in S_j , if and only if $\beta_i \neq 0$. In general, after removing duplicate sets, this will lead to $S_1 \subset S_2 \subset S_3 \dots S_{q'} \subset V$, where $q' \leq q$. Furthermore, if not already present, we add the empty $\emptyset$, and the full covariate set V . In the rare case, where one set violates the strict monotonicity, we remove the set. For example, for the Breast Cancer data, group lasso finds the following sets:

$$\begin{aligned} S_1 &= \{\} \\ S_2 &= \{6\} \\ S_3 &= \{2, 6\} \\ S_4 &= \{2, 3, 6\} \\ S_5 &= \{3, 6, 8\} \\ S_6 &= \{2, 3, 6, 8\} \\ &\dots \end{aligned}$$

and therefore remove S_5 .

The conditional class probability $p(y = 1 | \mathbf{x}, \beta, \tau)$ is modeled as

$$p(y = 1 | \mathbf{x}, \beta, \tau) = g(\tau + \sum_{i=1}^p \beta_i^T f_i(x_i)). \quad (12)$$

The non-linear transformations $f_i : \mathbb{R} \rightarrow \mathbb{R}^s$, where s is the number of splines, are learned from the training data using penalized B splines (Hastie et al., 2009).

⁹ In order to make this type of penalty meaningful we ensure that each $f_i(x_i)$ has mean 0 and standard deviation 1.

6 Extension to classification with recall guarantees

So far, we assumed that both misclassification costs $c_{0,1}$ and $c_{1,0}$ are given. Arguably, the false positive cost $c_{0,1}$ is relatively easy to specify. For example, in the medical domain, this might correspond to the price of a medicine which was unnecessarily prescribed to a healthy patient.

On the other hand, the specification of $c_{1,0}$ is more difficult. For example, specifying the cost of a dead patient (that might have been rescued) is difficult. Therefore, in the medical domain, it is more common to try to make a guarantee on the recall.¹⁰ In particular, it is common practice to require that the recall is 95% (Kanao et al., 2009).

In the following, we show how to estimate the false negative cost $c_{1,0}$, given the false positive cost $c_{0,1}$ and the requirement that the recall is larger or equal to some value r .

In the following, we denote by $\mathbb{1}_M$ the indicator function which is 1 if expression M is true and otherwise 0.

Given a distribution over $(y, \mathbf{x})$ such that $\mathbb{E}[1_{y=1}] > 0$, the recall of a decision procedure δ is defined as:

$$R_\delta := \int p(\mathbf{x}|y=1) \cdot \mathbb{1}_{\delta(\mathbf{x})=1} d\mathbf{x}.$$

Assuming that δ is a Bayes procedure, we have

$$\begin{aligned} \delta(\mathbf{x}) = 1 &\Leftrightarrow p(y=1|\mathbf{x}) \cdot c_{1,0} \geq p(y=0|\mathbf{x}) \cdot c_{0,1} \\ &\Leftrightarrow p(y=1|\mathbf{x}) \cdot c_{1,0} \geq (1 - p(y=1|\mathbf{x})) \cdot c_{0,1} \\ &\Leftrightarrow p(y=1|\mathbf{x}) \geq \frac{c_{0,1}}{c_{1,0} + c_{0,1}}. \end{aligned}$$

Setting

$$t := \frac{c_{0,1}}{c_{1,0} + c_{0,1}},$$

we have that

$$R_\delta = \int p(\mathbf{x}|y=1) \cdot \mathbb{1}_{p(y=1|\mathbf{x}) \geq t} d\mathbf{x}.$$

In order to emphasize the dependence on t , we write in the following R_t instead of R_δ . In particular, let t^* be chosen such that

$$r = R_{t^*} \tag{13}$$

where r is, for example, 0.95. Then the implicitly defined cost $c_{1,0}$ is given by

$$c_{1,0} = \frac{1 - t^*}{t^*} c_{0,1}. \tag{14}$$

¹⁰ Here we use the terminology from the machine learning literature, though, in the medical literature the term “sensitivity” is more common than “recall”.

It remains to show how t^* can be estimated. In general, Equation (13) does not have a solution (in terms of t). We therefore solve the following problem

$$\max_t t, \text{ subject to } r \leq R_t, \quad (15)$$

which has always a solution (since $t = 0$ fulfills the constraint). Since $p(\mathbf{x}|y = 1)$ is unknown, we use the empirical training data distribution to estimate R_t :

$$\begin{aligned} R_t &= \int p(\mathbf{x}|y = 1) \cdot \mathbb{1}_{p(y=1|\mathbf{x}) \geq t} d\mathbf{x} \\ &\approx \frac{1}{n_1} \sum_{k: y^{(k)}=1} \mathbb{1}_{p^{(-k)}(y=1|\mathbf{x}^{(k)}) \geq t}, \end{aligned} \quad (16)$$

where n_1 is the number of true samples (i.e. label $y = 1$), and $p^{(-k)}(y = 1|\mathbf{x}^{(k)})$ is the estimate of $p(y = 1|\mathbf{x}^{(k)})$ of the classifier that was trained without sample k . In practice, since this type of leave-one-out estimation is computationally expensive, we use instead 10-fold cross-validation.

Since the expression in Eq. (16) is a monotone decreasing step function in t , we can easily solve the problem in (15) by sorting $\{p^{(-k)}(y = 1|\mathbf{x}^{(k)})\}_{k=1}^{n_1}$ in decreasing order.

6.1 Adaptive covariate acquisition with recall guarantees

So far, we discussed how to estimate $c_{1,0}$ in the situation where only one classifier based on $p(y = 1|\mathbf{x})$ is used. However, in the adaptive acquisition setting, the situation is more complicated since, in general, for different observed sets of variables, the conditional class probabilities are different. In particular, let $S \subset S' \subseteq V$, where V is the set of all variables. Then, in general, we have

$$p(y = 1|\mathbf{x}_S) \neq p(y = 1|\mathbf{x}_{S'}),$$

which means that, in general, the optimal threshold t^* which guarantees recall $\geq r$ is different for different sets of observed variables S . Furthermore, in the adaptive setting, estimating the recall using Equation (16) is not valid anymore since the distribution of the samples, with label $y = 1$, and for which we select the variable set S' is, in general, different from $p(\mathbf{x}|y = 1)$, i.e. $p(\mathbf{x}|y = 1, \text{acquired } S') \neq p(\mathbf{x}|y = 1)$.

Nevertheless, we show in the following that it is possible to define the cost $c_{1,0}$ such that the recall requirement is fulfilled. First, let us introduce the following notations. Let $S_1 \subset S_2 \subset S_3 \dots S_q$, be the sets of variables that are considered for adaptive variable acquisition, i.e. first we acquire S_1 , and then we decide whether to classify or whether we acquire additionally the variables in $S_2 \setminus S_1$, and so forth. Moreover, let $\delta_{t,S}$ be the classifier based on the observed variables S and using threshold t , i.e.

$$\delta_{t,S}(\mathbf{x}_S) := \begin{cases} 1 & \text{if } p(y = 1|\mathbf{x}_S) \geq t, \\ 0 & \text{else.} \end{cases}$$

To simplify the notation, we write in the following $\delta_{t,S}$ short for $\delta_{t,S}(\mathbf{x}_S)$.

Strict control of the recall can be achieved by requiring that

$$p(\delta_{t_1,S_1} = 1, \delta_{t_2,S_2} = 1, \dots, \delta_{t_q,S_q} = 1 | y = 1) \geq r. \quad (17)$$

This can be seen as follows. Assume that $y = 1$ and any classifier δ_{t_j, S_j} outputs label 0, then an adversarial selection strategy will select this classifier. Otherwise, if all classifiers output label 1, then even an adversarial selection strategy needs to select a classifier δ_{t_j, S_j} for which the output is 1. By the requirement of Inequality (17), the latter will happen with probability of at least r .

If we require that all thresholds are the same, i.e.

$$t = t_1 = t_2 = \dots = t_q,$$

then we can proceed as before. That means, based on Inequality (17), we first calculate t^* , and then specify the false negative cost $c_{1,0}$ using Equation (14). Analogously to before for checking Inequality (17), we use the empirical training data estimate:

$$\begin{aligned} p(\delta_{t, S_1} = 1, \delta_{t, S_2} = 1, \dots, \delta_{t, S_q} = 1 | y = 1) \\ \approx \frac{1}{n_1} \sum_{k: y^{(k)} = 1} \mathbb{1}_{p^{(-k)}(y=1 | \mathbf{x}_{S_1}^{(k)}) \geq t} \cdot \mathbb{1}_{p^{(-k)}(y=1 | \mathbf{x}_{S_2}^{(k)}) \geq t} \cdots \mathbb{1}_{p^{(-k)}(y=1 | \mathbf{x}_{S_q}^{(k)}) \geq t}. \end{aligned}$$

7 Experiments

We evaluate our proposed method on four real datasets from the medical domain which are frequently used for cost-sensitive classification: Pima Diabetes dataset ($p = 8$, $n = 768$, positive class ratio ≈ 0.35), the Wisconsin Breast Cancer dataset ($p = 10$, $n = 683$, positive class ratio ≈ 0.35), Heart-disease dataset ($p = 13$, $n = 303$, positive class ratio ≈ 0.46), and the PhysioNet dataset ($p = 30$, $n = 12000$, positive class ratio ≈ 0.14). The first three datasets are all available at the UCI Machine Learning repository,¹¹ the PhysioNet data is available at <https://archive.physionet.org/pn3/challenge/2012/>.

Note that the PhysioNet data (Goldberger et al., 2000) contains for each patient several health check measures like cholesterol, taken at different times during their stay in the intensive care unit. As in Shim et al., (2018), for each patient we use the last recorded value of each attribute to predict death ($y = 1$) or survival ($y = 0$). After filtering attributes which are mostly missing, we acquire a data set with 12000 patients and 30 attributes.

For Diabetes and Heart-disease we use the covariate costs as defined in Ji and Carin (2007), and Turney (1994), respectively. For the other datasets, we set the covariate costs uniformly to one.

Note that the Heart-disease and PhysioNet data contain missing values. For methods which cannot handle missing values (including our proposed methods) we assume that all covariates are jointly distributed according to a multivariate normal distribution, where the covariance matrix is estimated from all samples (including missing values) using the method from Lounici et al. (2014).

We compare the proposed method AdaCOS to the following methods:

COS The proposed method but fixing the covariate set $S \in \{S_1, S_2, \dots, S_q\}$ to the one which minimizes the total costs in expectation, i.e.

¹¹ <https://archive.ics.uci.edu/ml/index.html>.

$$\mathbb{E}_{\mathbf{x}_S, y} [c_{y, \delta^*(\mathbf{x}_S)}] + \sum_{i \in S} c_i,$$

which is estimated using 10-fold crossvalidation as in Sect. 5.1.¹²

Full Model The logistic generalized additive model which always acquires (and uses) all covariates.

Shim2018 The cost-sensitive classification method based on deep reinforcement learning as proposed in Shim et al. (2018).¹³

GreedyMiser The cost-sensitive tree construction method proposed in Xu et al. (2012).¹⁴

AdaptGbrt The method proposed in Nan and Saligrama (2017), which requires the specification of a high accuracy classifier for which we use the Full Model.¹⁵

For all methods we estimate the hyperparameters with 10-fold crossvalidation, except where this is too computationally expensive: for Shim2018 we use the hold-out data split as in their provided implementation, for AdaptGbrt we use 5-fold crossvalidation.

As evaluation measure, we use the average total cost of classification, defined as

$$\text{avg total cost} := \frac{1}{n_t} \sum_{k=1}^{n_t} (c_{y^{(k)}, y_*^{(k)}} + \sum_{i \in S^{(k)}} c_i),$$

where n_t is the number of test samples; $y^{(k)}$ and $y_*^{(k)}$ is the k -th true test class and predicted test class, respectively; $S^{(k)}$ is the set of covariates that were used by the prediction model for classifying the k -th sample.

7.1 Results

For each dataset we use 5-fold cross-validation and report mean and standard deviation of the total costs. We evaluate all methods on two settings:

- user-specified false positive and false negative misclassification costs.
- user-specified false positive misclassification cost and target recall.

If not stated otherwise, we use group lasso, as explained in Sect. 5.2, for acquiring the sets $S_1 \subset S_2, \dots, S_q$.

7.1.1 User-specified misclassification costs

In the first setting, we assume that the user specifies the false positive cost in $\{1, 5, 10, 50, 100, 500, 1000\}$. The false negative cost is set to be 10 times the false positive cost, which reflects that it is more important to detect infected patients than avoiding wrongly classifying healthy patients.

The total cost of misclassification is shown in the top plot of Figs. 5, 6, 7, and 8 for Diabetes, Breast Cancer, PhysioNet and Heart-disease, respectively. We observe that with

¹² In case where the target recall is specified, we first estimate the false negative cost as in Sect. 6.1, and then proceed as before.

¹³ Available at <https://github.com/OpenXAIPProject/Joint-AFA-Classification>.

¹⁴ Available at <http://kilian.cs.cornell.edu/code/html>.

¹⁵ Available at <https://github.com/fnan/AdaptApprox>.

Fig. 5 Results on Pima Diabetes dataset with user-specified false positive cost in {1, 5, 10, 50, 100, 500, 1000}. The false negative cost is set to be 10 times the false positive cost

respect to minimizing the total cost of classification (top plots), our proposed method AdaCOS performs better than all previously proposed methods.

In each of those figures, in the middle and bottom plot, we also show the weighted accuracy and the number of acquired covariates each plotted against the false positive cost (which is set by the user), respectively. Since we assume that false negative classification

Fig. 6 Results on Breast Cancer dataset with user-specified false positive cost in {1, 5, 10, 50, 100, 500, 1000}. The false negative cost is set to be 10 times the false positive cost

have 10 times higher cost than false positive classification, we use the weighted accuracy defined by

$$\text{weighted accuracy} = \frac{\text{true positives} \cdot 10 + \text{true negatives}}{\text{number of true labels} \cdot 10 + \text{number of false labels}}.$$

Fig. 7 Results on PhysioNet dataset with user-specified false positive cost in $\{1, 5, 10, 50, 100, 500, 1000\}$. The false negative cost is set to be 10 times the false positive cost

From the bottom plots, as expected, we see that all methods start acquiring more covariates as the user-specified false positive cost increases. At the same time, all methods, except Shim2018, show an increase in (weighted) accuracy. In particular, Shim2018 underperforms on the smaller datasets Diabetes, Breast Cancer, and Heart-disease, which is likely

Fig. 8 Results on Heart-disease dataset with user-specified false positive cost in $\{1, 5, 10, 50, 100, 500, 1000\}$. The false negative cost is set to be 10 times the false positive cost

due to the difficulty of adjusting the hyper-parameters of their deep neural network classifier on small hold-out validation data.

In terms of (weighted) accuracy, Full Model performs always optimal, i.e. even for small datasets we do not find any gains in predictive accuracy by using a sparser model. As a conclusion, if the covariate costs are zero or negligible, we might just opt for the full

Fig. 9 Shows the weighted accuracy vs the number of acquired covariates. Results for where the number of acquired covariates was larger than any of the proposed methods are not shown. (All results were acquired with user-specified false positive cost in $\{1, 5, 10, 50, 100, 500, 1000\}$, and false negative cost 10 times the false positive cost.)

Fig. 10 Comparison of the proposed method AdaCOS when defining the covariates sets $S_1, S_2, \dots, S_q$ using either the group lasso penalty or forward selection

model to get the lowest total costs. On the other hand, if the ratio of false-positive cost to covariate cost is less than around 100, the full model performs considerably worse than the proposed method in terms of total cost.

Table 1 Shows the average training time in minutes per 1000 samples (standard deviation in brackets) for forward selection and group lasso

	Pima	Breast Cancer	Heart-disease	PhysioNet
Forward Selection	16.92 (1.12)	39.09 (1.03)	225.76 (4.66)	–
Group Lasso	6.17 (0.57)	11.0 (1.17)	29.66 (1.24)	20.1 (4.45)
Full Model	1.54 (0.18)	2.7 (0.14)	21.95 (1.27)	2.72 (0.78)
Shim2018	609.67 (9.69)	693.07 (9.56)	2278.36 (15.42)	46.55 (0.12)
AdaptGBRT	694.83 (11.96)	735.0 (4.93)	1585.2 (20.53)	351.99 (2.05)
GreedyMiser	1.35 (0.02)	1.87 (0.15)	4.43 (0.12)	0.41 (0.03)

For reference, we also show the training time of all other methods. Datasets are sorted by number of variables from smallest to largest ($p = 8, 10, 13, 30$)

Table 2 Shows the average test time in minutes per 1000 samples (standard deviation in brackets) for all methods

	Pima	Breast Cancer	Heart-disease	PhysioNet
AdaCOS	0.59 (0.02)	0.46 (0.04)	0.5 (0.04)	0.79 (0.07)
COS	0.01 (4.4e-04)	0.01 (8.2e-04)	0.01 (1.7e-03)	0.03 (8.2e-04)
Full Model	2.e-03 (1.2e-04)	1.1e-03 (4.6e-05)	5.e-03 (1.8e-03)	6.0e-04 (3.8e-05)
Shim2018	0.01 (1.9e-03)	0.01 (1.3e-03)	0.02 (1.7e-03)	3.9e-03 (4.3e-04)
AdaptGBRT	1.8e-04 (3.5e-06)	2.e-04 (1.9e-06)	3.6e-04 (6.1e-06)	4.0e-05 (7.5e-07)
GreedyMiser	2.4e-04 (2.e-05)	2.7e-04 (3.1e-05)	5.9e-04 (8.0e-05)	2.2e-05 (1.3e-06)

Datasets are sorted by number of variables from smallest to largest ($p = 8, 10, 13, 30$)

The (weighted) accuracy of AdaCOS and COS are similar, while the former achieves the same accuracy with less covariates. This demonstrates the effectiveness of estimating the expected cost of misclassification depending on what we have observed so far using the conditional Bayes risk as in Eq. (8).

Additionally, in Fig. 9, we show how the (weighted) accuracy changes as the number of acquired covariates increases.¹⁶ The plot confirms that AdaCOS consistently improves over COS.

Comparison of Group Lasso and Forward Selection Next, we compare the forward selection strategy (Sect. 5.1) and group lasso (Sect. 5.2) when used for acquiring the covariate sets $S_1 \subset S_2, \dots, S_q$. For Diabetes, Breast Cancer, and Heart-disease, the total costs of the proposed method AdaCOS with forward selection and group lasso are shown in Fig. 10. Due to the high computational costs for large p , it was not feasible to apply forward selection to the PhysioNet dataset. We find that, except for the Breast Cancer dataset, both covariate acquisition strategies lead to similar results. For Breast Cancer, forward selection appears to be superior to group lasso.

¹⁶ Note that each point displayed in the plot corresponds to one result when the false positive cost was set to one value in $\{1, 5, 10, 50, 100, 500, 1000\}$. However, points are overlapping when the acquired covariates (and also the weighted accuracy) were exactly the same for two different false positive costs.

Table 3 Shows the total cost of misclassification under the same cost setting as in Ji and Carin (2007), Dulac-Arnold et al. (2012): user-specified misclassification costs are symmetric (either 400 or 800), cost of correct classification equals -50 . The results for the methods DWSM and POMDP are taken from Dulac-Arnold et al. (2012) and Ji and Carin (2007), respectively

	400	800
AdaCOS (forward selection)	68.83 (12.43)	157.35 (24.34)
COS (forward selection)	70.81 (16.75)	161.18 (29.58)
AdaCOS (group lasso)	78.1 (5.96)	180.64 (14.12)
COS (group lasso)	82.55 (12.84)	171.71 (16.1)
Full Model	99.52 (8.67)	191.2 (16.38)
Shim2018	113.82 (17.87)	246.71 (26.71)
AdaptGBRT	87.41 (12.1)	175.8 (17.76)
GreedyMiser	91.36 (14.43)	200.96 (31.93)
DWSM	74.0 (–)	181.0 (–)
POMDP	75.0 (–)	180.0 (–)

Lowest costs are given in bold

Runtime Comparison In Table 1, first two rows, we compare the training time of AdaCOS with forward selection and AdaCOS with group lasso (Sect. 5.2).¹⁷ We clearly, confirm that group lasso is considerably faster than forward selection, allowing us to apply our proposed method to the PhysioNet dataset.

For reference, in Table 1, row 3 to 6, we also show the training time of all other methods, but we note that a fair comparison is difficult due to different types of implementations and the Matlab license conditions. In particular, the experiments of our proposed methods and Shim2018, were all run using a Python environment with 88 Intel Xeon 2.20GHz CPUs, whereas GreedyMiser and AdaptGbrt run on Matlab with 8 Intel Xeon 3.50 GHz CPUs.¹⁸

Finally, in Table 2, we also show the runtime at test-time of all methods. Though AdaCOS is slower than all other methods, the proposed method is still practically fast, requiring less than 2 min to process all test samples for PhysioNet.

Symmetric misclassification costs on Diabetes dataset In order to compare to the results reported in Ji and Carin (2007), Dulac-Arnold et al. (2012), we also evaluate on the Diabetes dataset with symmetric misclassification costs (i.e. false negative and false positive costs are the same), and the cost for correct classification set to -50 . The results, shown in Table 3, suggest that also in this setting the proposed method can have an advantage over previously proposed methods. In particular, the proposed method AdaCOS with forward selection has the lowest total costs, though, when using group lasso the proposed method underperforms.

7.1.2 User-specified target recall

Finally, we investigate the setting where the user specifies target recall r instead of false negative costs. Here, we show the results for $r = 0.95$, the results for target recall $r = 0.99$ are similar and given in the supplement material.

¹⁷ Note that the training of AdaCOS and COS is basically the same. AdaCOS and COS differ only at test time.

¹⁸ Due to the Matlab's license conditions we could not run the experiments on the same server as the other methods.

Fig. 11 Actually observed recall on test data when threshold on class probabilities was adjusted on validation data to have recall ≥ 0.95 .

Fig. 12 Results on Pima Diabetes dataset with user-specified false positive cost in $\{5, 10, 50, 100, 500, 1000\}$, and target recall ≥ 0.95 .

For this setting, we do not consider the method AdaptGbrt, since it does not allow the output of class probabilities or scores. For Shim2018 and GreedyMiser, we found that simply using the class probabilities/scores from the validation data to learn thresholds with recall $\geq r$, tended to lead to recall less than r on the test data, as shown in Fig. 11.

Fig. 13 Results on Breast Cancer dataset with user-specified false positive cost in {1, 5, 10, 50, 100, 500, 1000}, and target recall ≥ 0.95 .

Therefore, in order to make all results comparable, we show the results for Shim2018 and GreedyMiser at the same recall level as the proposed method AdaCOS.

The recall of the proposed method AdaCOS on the test data, as shown in Fig. 11, never violates the target recall of 0.95.

Fig. 14 Results on PhysioNet dataset with user-specified false positive cost in $\{1, 5, 10, 50, 100, 500, 1000\}$, and target recall ≥ 0.95 .

Since no false negative costs are provided, we cannot evaluate in terms of total costs anymore. Instead, we evaluate in terms of average operation costs, defined as the average cost of false positives plus the costs for covariate acquisition:

Fig. 15 Results on Heart-disease dataset with user-specified false positive cost in $\{1, 5, 10, 50, 100, 500, 1000\}$, and target recall ≥ 0.95

$$\text{avg operation costs} := \frac{1}{n_t} \sum_{k=1}^{n_t} ((1 - y^{(k)}) \cdot y_*^{(k)} \cdot c_{0,1} + \sum_{i \in S^{(k)}} c_i).$$

The results for all datasets are shown in Figures 12, 13, 14, and 15. For completeness, on each of the those figures, we also plot the false discovery rate (FDR) against the covariate costs (bottom plots), where the crosses of different sizes mark the standard deviation of FDR and covariates costs.

For all datasets, except Heart-Disease, the proposed method AdaCOS has the smallest operation costs. Furthermore, AdaCOS tends to achieve a lower false discovery rate with less covariates used.

8 Related work

Here, we briefly summarize various previous works for cost-sensitive classification.

Cost-sensitive Classification Without Covariate Costs The theory about optimal decision making given a (mis-)classification cost matrix and conditional class probabilities $p(y|x)$ has been extensively studied in the statistical literature, see for example (Berger 2013), and the machine learning community, see for example (Zadrozny & Elkan, 2001; Zadrozny et al., 2003; Elkan, 2001). However, the costs of acquiring covariates are not addressed in these early works. The work in Zadrozny and Elkan (2001) also addresses the problem of an undefined cost-matrix. However, in contrast to our work, they are concerned with sample dependent costs, and do not consider recall constraints which we address.

Markov Decision Process (MDP) Framework The MDP formulation and solution using an action-utility representation (Q-learning) in Zubeck et al. (2004), Bayer-Zubeck (2004) is closest to our approach. Their method also leads to a Bayes procedure. However, they do not provide a formal proof and consider only discrete covariates. The work in Dulac-Arnold et al. (2011), Dulac-Arnold et al. (2012), Karayev et al. (2013) also uses the MDP framework. However, their proposed method cannot incorporate the uncertainty about the covariate distributions. The work in Ji and Carin (2007) tries to model such uncertainties by modeling the cost-sensitive classification problem as a partial observable Markov decision process (POMDP). However, their POMDP formulation can lead to repeatedly selecting the same covariates, and as a consequence they need to adapt the stopping criteria. In contrast, our proposed method, based on the extended definition of a Bayes procedure, leads to a solution for deciding when to stop asking for new covariates that has provably the lowest total costs in expectation.

Reinforcement Learning Approaches Janisch et al. (2017), Shim et al. (2018) suggest to use deep reinforcement learning with Q-learning. In contrast to MDP, a discriminative decision maker is learned which does not require an environmental model. Their method performs promising in the domain where huge amounts of labeled training data is available. Alternatively, the work in Benbouzid et al. (2012) suggests the use of SARSA. The method in Contardo et al. (2016) also addresses this problem with reinforcement learning. However, methods based on reinforcement learning are hard to optimize, and can lead to suboptimal behavior as our experimental comparison to Shim et al. (2018) demonstrated.

Discriminative Decision Approach The work in Wang et al. (2015) proposes an intriguing method for finding a decision procedure that is guaranteed to converge to the Bayes risk given sufficient enough training data. Their idea is to create a Bayes optimal classifier for all possible subsets of covariates, and a directed a-cyclic graph that connects them. They formulate the problem as an empirical risk minimization (ERM) problem, and show that with infinitely many training samples the loss at each node converges to the Bayes risks. However, in order to allow for scalability their method requires to acquire covariates in

batches. The work in Trapeznikov and Saligrama (2013), Wang et al. (2014b) uses a similar framework but is restricted to a fixed sequential order, which is different from our method that also learns the sequential order. The work in Kusner et al. (2014), Xu et al. (2013), Wang et al. (2014a) also uses the ERM framework together with budget constraints, but it does not allow to incorporate asymmetric misclassification costs.

Cost-sensitive Tree Construction The works in Xu et al. (2012), Nan et al. (2015), Nan et al. (2016), Nan and Saligrama (2017), Peter et al. (2017) learn decision trees subject to budget constraints on the features. In particular, the methods in Nan and Saligrama (2017), Peter et al. (2017) are considered state of the art for this task. Their usage of gradient boosted decision trees (Friedman 2001) makes them in particular effective for very large training data. Cost-sensitive decision trees for discrete covariates are also considered in Sheng and Ling (2006), and extended to Bayesian Networks in Bilgic and Getoor (2007). In contrast to our work, these works do not explicitly optimize the expected total costs of misclassification, which can result in inferior performance, as we confirmed in Sect. 7.

Entropy-Based Approaches The work in Kanani and Melville (2008), Gao and Koller (2011), Kapoor and Horvitz (2009), Gong et al. (2019) optimizes a criteria that combines the costs of features with an estimate of the class entropy of the resulting classifier. As such their objective function is different from ours.

Density Estimation via Autoencoders The work in Kachuee et al. (2018) suggests to acquire the covariate which has the highest sensitivity to the output prediction y . In order to account for different covariate acquisition costs the sensitivity scores are re-scaled appropriately. The sensitivity scores are estimated using a denoising autoencoder. Similarly, the work in Ma et al. (2019) uses as objective function the expected Shannon information, which is estimated via a variational autoencoder. Both objective functions are not related to the minimization of the expected total cost.

Others The work in Greiner et al. (2002) extends the Probably Approximately Correct (PAC) framework to prove the existence of a cost-sensitive classifier that is with high probability optimal in the sense of providing minimal average total costs. However, they assume a probability distribution over only discrete covariates. The method in Lakkaraju and Rudin (2017) is additionally focused on interpretability, and, as a consequence, optimizes an objective function that is different from ours. Imitation learning is also applied to this task by He et al. (2012), but their definition of loss is different from minimizing the total classification costs that we consider here. The work in Nan et al. (2014) assumes a margin-based classifier and uses a k-nearest neighbor approach to estimate the accuracy of the classifier in each step before deciding whether to acquire the next feature. However, their stopping criteria is solely based on an estimate of the classification accuracy, and as such, does not allow for a decision which is cost optimal.

9 Conclusions

In this article, we addressed the problem of cost-sensitive classification where the goal is to minimize the total costs, defined as the expected cost of misclassification plus the cost for covariate acquisition.

In Sect. 2, we rigorously formalized this goal as the minimization of the (conditional) Bayes risk which can change after the acquisition of a new covariate. However, solving this minimization problem is hard. First, the evaluation of the conditional Bayes risk requires to estimate and integrate over a high dimensional density. Second, the Bayes risk must be

evaluated for all combinations of covariate sets which is exponential in p the number of covariates.

In order to overcome the computational difficulties, we introduced two working assumptions:

1. The optimal classifier can be expressed as a generalized additive model (GAM).
2. The optimal sets of covariates can be expressed as a sequence of sets that are monotone increasing, namely $S_1 \subset S_2 \dots \subset S_q$.

Using the first assumption, we showed, in Sect. 4, that the evaluation of the conditional Bayes risk reduces to a one dimensional density estimation and integration problem which can be efficiently estimated.

Furthermore, we showed that the sequence $S_1 \subset S_2 \dots \subset S_q$ can be computationally efficiently acquired by inspecting the regression coefficient path when penalizing GAM with group lasso.

Our experiments suggest that our proposed method AdaCOS achieves in most situations the lowest total costs of classification, when compared to the previous methods POMDP, DWSM, GreedyMiser, AdaptGbrt, and Shim2018 (Ji & Carin, 2007; Dulac-Arnold et al., 2012; Xu et al., 2012; Nan & Saligrama, 2017; Shim et al., 2018).

We note that some previous methods like Shim2018 (Shim et al., 2018) do not share our working assumptions, and instead use a very flexible classifier (deep neural network) and covariate acquisition strategy based on reinforcement learning. However, for small ($n < 1000$) datasets, and even for medium-sized ($n > 10000$) datasets like PhysioNet, we found that a generalized additive model is competitive or even better than a neural network classifier, and the flexibility of the reinforcement learning seems to suffer from high variance.

Finally, we considered the situation where not all misclassification costs are specified by the user. In particular, we considered the situation where the user specifies a target recall instead of the cost of false negative classification. We showed that it is possible to apply the proposed method by estimating the implicitly defined false negative cost. Our experiments showed that the resulting method indeed achieves the desired minimum recall, while minimizing the false discovery rate and covariate acquisition cost.

The source code of the proposed method and for reproducing all results is available at https://github.com/andrade-stats/AdaCOS_public.

Appendix

Here, we prove Theorem 1, which states that the procedure π^* , as defined in Equation (6), is a Bayes procedure. That means we need to show that for any other decision procedure π we have

$$\mathbb{E}_{\mathbf{x},y}[l((\mathbf{x}, y), \pi^*)] \leq \mathbb{E}_{\mathbf{x},y}[l((\mathbf{x}, y), \pi)].$$

Proof Let $S \subseteq V$ be the set of already observed covariates, then the expected remaining costs for a decision procedure π is given by

$$\mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi, S) | \mathbf{x}_S].$$

We will prove by induction that for any $S \subseteq V$, and any decision procedure π , we have

$$\mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S] \leq \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi, S) | \mathbf{x}_S].$$

The claim then follows by setting $S := \emptyset$.

Base case: $S = V$. We have

$$\pi^*(\mathbf{x}_S) = \arg \min_{i \in L} \mathbb{E}_y [c_{y,i} | \mathbf{x}_S].$$

Therefore $\pi^*(\mathbf{x}_S)$ is a Bayes procedure, and as a consequence

$$\mathbb{E}_y [c_{y, \pi^*(\mathbf{x}_S)} | \mathbf{x}_S] \leq \mathbb{E}_y [c_{y, \pi(\mathbf{x}_S)} | \mathbf{x}_S].$$

And therefore

$$\mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S] \leq \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi, S) | \mathbf{x}_S].$$

(Since $S = V$, and we have

$$\mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S] = \mathbb{E}_y [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S] = \mathbb{E}_y [c_{y, \pi^*(S)} | \mathbf{x}_S],$$

and the same analogously for π .)

Induction step: $S \subset V$. Assume that for all $S \cup \{i\}$, where $i \in V \setminus S$, the induction assumptions holds, that is

$$\mathbb{E}_{\mathbf{x}_{V \setminus (S \cup \{i\})}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_{S \cup \{i\}}] \leq \mathbb{E}_{\mathbf{x}_{V \setminus (S \cup \{i\})}, y} [l((\mathbf{x}, y), \pi, S \cup \{i\}) | \mathbf{x}_{S \cup \{i\}}].$$

Let $\hat{\pi}$ denote a Bayes procedure. Using the structure of the loss function as defined in Eq. (5), we have

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \hat{\pi}, S) | \mathbf{x}_S] \\ &= \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} \left[\begin{cases} c_{y, \hat{\pi}(\mathbf{x})} & \text{if } \hat{\pi}(\mathbf{x}) \in L, \\ c_{\hat{\pi}(\mathbf{x})} + l((\mathbf{x}, y), \hat{\pi}, S \cup \{\hat{\pi}(\mathbf{x})\}) & \text{else.} \end{cases} \middle| \mathbf{x}_S \right] \\ &\geq \min_{i \in LU(V \setminus S)} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} \left[\begin{cases} c_{y,i} & \text{if } i \in L, \\ c_i + l((\mathbf{x}, y), \hat{\pi}, S \cup \{i\}) & \text{else.} \end{cases} \middle| \mathbf{x}_S \right] \\ &= \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \hat{\pi}, S \cup \{i\}) | \mathbf{x}_S] & \text{else.} \end{cases} \\ &= \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{x_i} \left[\mathbb{E}_{\mathbf{x}_{V \setminus (S \cup \{i\})}, y} [l((\mathbf{x}, y), \hat{\pi}, S \cup \{i\}) | \mathbf{x}_{S \cup \{i\}}] \middle| \mathbf{x}_S \right] & \text{else.} \end{cases} \\ &\stackrel{(1)}{\geq} \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{x_i} \left[\mathbb{E}_{\mathbf{x}_{V \setminus (S \cup \{i\})}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_{S \cup \{i\}}] \middle| \mathbf{x}_S \right] & \text{else.} \end{cases} \\ &= \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_S] & \text{else.} \end{cases} \\ &= \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S], \end{aligned}$$

where in the line marked by (1) we used the induction assumption. The last line follows from Lemma 1. Since $\hat{\pi}$ is a Bayes procedure, we must have equality in the second and fifth line. Therefore π^* is also a Bayes procedure. $\square$

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S] \\ &= \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_S] & \text{else.} \end{cases} \end{aligned}$$

Lemma 1

$$\begin{aligned} & \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S) | \mathbf{x}_S] \\ &= \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} \left\{ \begin{array}{ll} c_{y, \pi^*(\mathbf{x}_S)} & \text{if } \pi^*(\mathbf{x}_S) \in L, \\ c_{\pi^*(\mathbf{x}_S)} + l((\mathbf{x}, y), \pi^*, S \cup \{\pi^*(\mathbf{x}_S)\}) & \text{else.} \end{array} \right. | \mathbf{x}_S] \\ &= \begin{cases} \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [c_{y, \pi^*(\mathbf{x}_S)} | \mathbf{x}_S] & \text{if } \pi^*(\mathbf{x}_S) \in L, \\ c_{\pi^*(\mathbf{x}_S)} + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{\pi^*(\mathbf{x}_S)\}) | \mathbf{x}_S] & \text{else.} \end{cases} \\ &= \begin{cases} \mathbb{E}_y [c_{y, \pi^*(\mathbf{x}_S)} | \mathbf{x}_S] & \text{if } \pi^*(\mathbf{x}_S) \in L, \\ c_{\pi^*(\mathbf{x}_S)} + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{\pi^*(\mathbf{x}_S)\}) | \mathbf{x}_S] & \text{else.} \end{cases} \end{aligned}$$

Proof

1. Case: $\pi^*(\mathbf{x}_S) \in L$. Then because of the definition of π^* , we have

$$\mathbb{E}_y [c_{y, \pi^*(\mathbf{x}_S)} | \mathbf{x}_S] = \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_y [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_S] & \text{else.} \end{cases}$$

2. Case: $\pi^*(\mathbf{x}_S) \notin L$. Then because of the definition of π^* , we have

$$\begin{aligned} & c_{\pi^*(\mathbf{x}_S)} + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{\pi^*(\mathbf{x}_S)\}) | \mathbf{x}_S] \\ &= \min_{i \in LU(V \setminus S)} \begin{cases} \mathbb{E}_y [c_{y,i} | \mathbf{x}_S] & \text{if } i \in L, \\ c_i + \mathbb{E}_{\mathbf{x}_{V \setminus S}, y} [l((\mathbf{x}, y), \pi^*, S \cup \{i\}) | \mathbf{x}_S] & \text{else.} \end{cases} \end{aligned}$$

$\square$

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10994-021-05958-z>.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson, T. W. (2003). *An introduction to multivariate statistical analysis* (Vol. 2). Wiley.
- Andrade, D., & Okajima, Y. (2019). Efficient bayes risk estimation for cost-sensitive classification. In *The 22nd international conference on artificial intelligence and statistics* (pp. 3372–3381).
- Bayer-Zubek, V. (2004). Learning diagnostic policies from examples by systematic search. In *Proceedings of the 20th conference on uncertainty in artificial intelligence* (pp. 27–34). AUAI Press.
- Benbouzid, D., Busa-Fekete, R., & Kégl, B. (2012). Fast classification using sparse decision dags. In *Proceedings of the 29th international conference on machine learning* (pp. 747–754).
- Berger, J. O. (2013). *Statistical decision theory and Bayesian analysis*. Springer.
- Bilgic, M., & Getoor, L. (2007). Voila: Efficient feature-value acquisition for classification. In *Proceedings of the national conference on artificial intelligence* (Vol. 22, p. 1225). AAAI Press, Menlo Park.
- Contardo, G., Denoyer, L., & Artières, T. (2016). Sequential cost-sensitive feature acquisition. In *International symposium on intelligent data analysis* (pp. 284–294). Springer.
- Dulac-Arnold, G., Denoyer, L., Preux, P., & Gallinari, P. (2011). Datum-wise classification: a sequential approach to sparsity. In *Joint European conference on machine learning and knowledge discovery in databases* (pp. 375–390). Springer.
- Dulac-Arnold, G., Denoyer, L., Preux, P., & Gallinari, P. (2012). Sequential approaches for learning datum-wise sparse representations. *Machine Learning*, 89(1–2), 87–122.
- Elkan, C. (2001). The foundations of cost-sensitive learning. *International Joint Conference on Artificial Intelligence*, 17, 973–978.
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics*, 29, 1189–1232.
- Gao, T., & Koller, D. (2011). Active classification based on value of classifier. In *Advances in neural information processing systems* (pp 1062–1070).
- Gelman, A., Stern, H. S., Carlin, J. B., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*. Chapman and Hall/CRC.
- Goldberger, A. L., Amaral, L. A., Glass, L., Hausdorff, J. M., Ivanov, P. C., Mark, R. G., Mietus, J. E., Moody, G. B., Peng, C. K., & Stanley, H. E. (2000). Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *Circulation*, 101(23), e215–e220.
- Gong, W., Tschitschek, S., Nowozin, S., Turner, R.E., Hernández-Lobato, J. M., & Zhang, C. (2019). Ice-breaker: Element-wise efficient information acquisition with a bayesian deep latent gaussian model. In *Advances in neural information processing systems* (pp 14791–14802).
- Greiner, R., Grove, A. J., & Roth, D. (2002). Learning cost-sensitive active classifiers. *Artificial Intelligence*, 139(2), 137–174.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction*. Springer.
- Hastie, T., Tibshirani, R., & Wainwright, M. (2015). *Statistical learning with sparsity: The lasso and generalizations*. Chapman and Hall/CRC.
- He, H., Eisner, J., & Daume, H. (2012). Imitation learning by coaching. In *Advances in neural information processing systems* (pp 3149–3157).
- Janisch, J., Pevný, T., & Lisý, V. (2017). Classification with costly features using deep reinforcement learning. arXiv preprint [arXiv:1711.07364](https://arxiv.org/abs/1711.07364).
- Ji, S., & Carin, L. (2007). Cost-sensitive feature acquisition and classification. *Pattern Recognition*, 40(5), 1474–1485.
- Kachuee, M., Darabi, S., Moatamed, B., & Sarrafzadeh, M. (2018). Dynamic feature acquisition using denoising autoencoders. *IEEE Transactions on Neural Networks and Learning Systems*, 30(8), 2252–2262.
- Kanani, P., & Melville, P. (2008). Prediction-time active feature-value acquisition for cost-effective customer targeting. In *Advances in neural information processing systems (NIPS)*.
- Kanao, K., Komori, O., Nakashima, J., Ohigashi, T., Kikuchi, E., Miyajima, A., Nakagawa, K., Eguchi, S., & Oya, M. (2009). PSA cut-off nomogram that avoid over-detection of prostate cancer in elderly men. *The Journal of Urology*, 4(181), 748.
- Kapoor, A., & Horvitz, E. (2009). Breaking boundaries: Active information acquisition across learning and diagnosis. In *Proceedings of the 22nd international conference on neural information processing systems* (pp. 898–906). Curran Associates Inc.
- Karayev, S., Fritz, M. J., & Darrell, T. (2013). Dynamic feature selection for classification on a budget. In *International conference on machine learning (ICML): Workshop on prediction with sequential models*.

- Kusner, M. J., Chen, W., Zhou, Q., Xu, Z. E., Weinberger, K. Q., & Chen, Y. (2014). Feature-cost sensitive learning with submodular trees of classifiers. In *AAAI* (pp. 1939–1945).
- Lakkaraju, H., & Rudin, C. (2017). Learning cost-effective and interpretable treatment regimes. In *Artificial intelligence and statistics* (pp. 166–175).
- Lounici, K. (2014). High-dimensional covariance matrix estimation with missing observations. *Bernoulli*, 20(3), 1029–1058.
- Ma, C., Tschitschek, S., Palla, K., Hernandez-Lobato, J. M., Nowozin, S., & Zhang, C. (2019). Eddi: Efficient dynamic discovery of high-value information with partial VAE. In *International conference on machine learning* (pp. 4234–4243).
- Nan, F., & Saligrama, V. (2017). Adaptive classification for prediction under a budget. In *Advances in neural information processing systems* (pp. 4730–4740).
- Nan, F., Wang, J., Trapeznikov, K., & Saligrama, V. (2014). Fast margin-based cost-sensitive classification. In *2014 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (pp. 2952–2956). IEEE.
- Nan, F., Wang, J., & Saligrama, V. (2015). Feature-budgeted random forest. In *International conference on machine learning* (pp. 1983–1991).
- Nan, F., Wang, J., & Saligrama, V. (2016). Pruning random forests for prediction on a budget. In *Advances in neural information processing systems* (pp. 2334–2342).
- O'Hara, R. B., & Sillanpää, M. J. (2009). A review of Bayesian variable selection methods: What, how and which. *Bayesian Analysis*, 4(1), 85–117.
- Peter, S., Diego, F., Hamprecht, F. A., & Nadler, B. (2017). Cost efficient gradient boosting. In *Advances in neural information processing systems* (pp. 1550–1560).
- Rasmussen, C. E., & Williams, C. K. (2006). *Gaussian processes for machine learning*. MIT Press.
- Russell, S., & Norvig, P. (2003). *Artificial intelligence: A modern approach*. Pearson.
- Sheng, V. S., & Ling, C. X. (2006). Feature value acquisition in testing: a sequential batch test algorithm. In *Proceedings of the 23rd international conference on Machine learning* (pp. 809–816). ACM.
- Shim, H., Hwang, S. J., & Yang, E. (2018). Joint active feature acquisition and classification with variable-size set encoding. In *Advances in neural information processing systems* (pp. 1368–1378).
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B (Methodological)*, 58, 267–288.
- Trapeznikov, K., & Saligrama, V. (2013). Supervised sequential classification under budget constraints. In *Artificial intelligence and statistics* (pp. 581–589).
- Turney, P. D. (1994). Cost-sensitive classification: Empirical evaluation of a hybrid genetic decision tree induction algorithm. *Journal of Artificial Intelligence Research*, 2, 369–409.
- Wang, J., Bolukbasi, T., Trapeznikov, K., & Saligrama, V. (2014a). Model selection by linear programming. In *European conference on computer vision* (pp. 647–662). Springer.
- Wang, J., Trapeznikov, K., & Saligrama, V. (2014b). An LP for sequential learning under budgets. In *Artificial intelligence and statistics* (pp. 987–995).
- Wang, J., Trapeznikov, K., & Saligrama, V. (2015). Efficient learning by directed acyclic graph for resource constrained prediction. In *Advances in neural information processing systems* (pp. 2152–2160).
- Xu, Z., Kusner, M., Weinberger, K., & Chen, M. (2013). Cost-sensitive tree of classifiers. In *International conference on machine learning* (pp. 133–141).
- Xu, Z., Weinberger, K. Q., & Chapelle, O. (2012). The greedy miser: learning under test-time budgets. In *Proceedings of the 29th international conference on international conference on machine learning* (pp. 1299–1306). Omnipress.
- Zadrozny, B., & Elkan, C. (2001). Learning and making decisions when costs and probabilities are both unknown. In *Proceedings of the seventh ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 204–213).
- Zadrozny, B., Langford, J., & Abe, N. (2003). Cost-sensitive learning by cost-proportionate example weighting. In *Third IEEE international conference on data mining* (pp. 435–442). IEEE.
- Zubek, V. B., & Dietterich, T. G. (2002). Pruning improves heuristic search for cost-sensitive learning. In *International conference on machine learning* (pp. 19–26).