INSTITUTE
OF ECONOMIC STUDIES
Faculty of Social Sciences
Charles University

ACQUISITION OF COSTLY INFORMATION IN DATA-DRIVEN DECISION MAKING

Lukas Janasek

IES Working Paper 10/2022

Institute of Economic Studies,
Faculty of Social Sciences,
Charles University in Prague

[UK FSV – IES]

Opletalova 26
CZ-110 00, Prague
E-mail : ies@fsv.cuni.cz
http://ies.fsv.cuni.cz

Institut ekonomických studií
Fakulta sociálních věd
Univerzita Karlova v Praze

Opletalova 26
110 00 Praha 1

E-mail : ies@fsv.cuni.cz
http://ies.fsv.cuni.cz

Bibliographic information:
Janasek L. (2022): "Acquisition of Costly Information in Data-Driven Decision Making" IES Working Papers 10/2022. IES FSV. Charles University.

Acquisition of Costly Information in Data-Driven Decision Making

Lukas Janasek[a]

[a]Institute of Economic Studies, Charles University & Institute of Information Theory and Automation, Czech Academy of Sciences, Prague, Czech Republic
E-mail: lukas.janasek@fsv.cuni.cz

May 2022

Abstract:
This paper formulates and solves an economic decision problem of the acquisition of costly information in data-driven decision making. The paper assumes an agent predicting a random variable utilizing several costly explanatory variables. Prior to the decision making, the agent learns about the relationship between the random variables utilizing its past realizations. During the decision making, the agent decides what costly variables to acquire and predicts using the acquired variables. The agent's utility consists of the correctness of the prediction and the costs of the acquired variables. To solve the decision problem, we split the decision process into two parts: acquisition of variables and prediction using the acquired variables. For the prediction, we propose an approach for training a single predictive model accepting any combination of acquired variables. For the acquisition, we propose two methods using supervised machine learning models: a backward estimation of the expected utility of each variable and a greedy acquisition of variables based on a myopic estimate of the expected utility. We evaluate the methods on two medical datasets. The results show that the methods acquire the costly variables efficiently.

JEL: C44, C45, C52, C73, D81, D83
Keywords: costly information, data-driven decision-making, machine learning

Acknowledgements: I am gratefull to Josef Barunik for many useful comments and discussion. I gratefully acknowledge the support from the Czech Science Foundation under the EXPRO GX19-28231X project.

1 Introduction

In an information-rich world, understanding what information to acquire and internalize is of crucial importance. In theory, more information implies lower uncertainty and better decisions of economic agents. In practice, however, the acquisition and internalization of information can be costly and the decision making of an agent complicates. Not only has the agent to decide how to act upon information but also upon *what* information to act. When considering what information to acquire and internalize, the agent faces the following trade-off. On one side, additional information increases the agent's expected utility by lowering uncertainty, on the other side, it decreases the agent's utility through the incurred costs. Consequently, the agent balances the costs of information against the reduction in uncertainty provided by the information.

A standard decision theory (Raiffa and Schlaifer, 1961) addresses decision making with costs stemming from the acquisition of information. The theory introduces a concept of the value of information and concludes that an agent acquires information if its benefits exceed its costs. A recently growing literature on the theory of rational inattention (Sims, 2003, 2010; Maćkowiak et al., 2018) focuses on the internalization of information. The theory argues that the internalization of information consumes agent's attention, which is a limited resource. As a result, the agent trades off the scarce attention against the benefits of internalized information. Both the standard decision theory and the theory of rational inattention provide valuable insights into the agent's decision making and answer why economic agents act upon limited information even if the full information is available. The theories, however, rely on a subtle assumption: the agent knows the distribution of her uncertainty.

In contrast to the theory, in many real-life applications, the decisions of agents are data-driven and imposing assumptions on its distribution could be troublesome. Without knowing the data distribution, the standard theory fails to describe the agent's behaviour. We can find several domains where the decisions are data-driven with the acquisition and processing of the data being costly. A first such domain is a medical diagnosis utilizing costly tests or examinations. For example, Turney (1995), Kachuee et al. (2019), Karkkainen et al. (2019), Raykar et al. (2010), or Shim et al. (2018) studied the trade-off between the

correctness of data-driven medical diagnosis and the costs of the medical tests or examinations. Krause and Guestrin (2014), Viola and Jones (2001), Wang et al. (2014), Chen et al. (2012), and Xu et al. (2013) studied systems of sensors and recognition systems making a trade-off between the correct prediction and the cost of the input data. Another domain is computer adaptive testing and questionnaires studied by Chen et al. (2015) with the costs being time or impatience of respondents. Finally, Early et al. (2016) studied the data-driven forecast of future energy consumption by prospective tenants, while Kanani and Melville (2008) described a system for targeting customers and business partners containing a costly variable provided by an external vendor.

In parallel to the economic literature, there is a recently growing literature on cost-sensitive machine learning addressing the data-driven decision making. The literature usually focuses on the task of a prediction with costly variables and omits distributional assumptions about uncertainty. It provides valuable methods in the domain of data-driven decision making with costly data, but it usually does not relate to the economic theory. Viola and Jones (2001), Xu et al. (2013), and Nan and Saligrama (2017) proposed methods that rely on a collection of supervised models each operating on a different subset of costly variables. Wang et al. (2015) proposed to train a collection of models for all combinations of acquired costly variables, which is a close approach to the work in this paper. Another stream of literature focused on the (deep) reinforcement learning Janisch et al. (2017), Shim et al. (2018), and Kachuee et al. (2019) trained a deep Q-network serving as an acquisition policy. The last category of methods relies on a heuristic evaluation of all variables and selection of the next best variable (see Sheng and Ling (2006), Early et al. (2016), Kanani and Melville (2008), Karkkainen et al. (2019), or Vivar et al. (2020))

In contrast to cost-sensitive machine learning literature, we focus on the acquisition of costly information in data-driven decision making formulated as an economic decision problem. We assume an agent predicting a value of a random variable using several costly explanatory variables. The distribution of the random variables is unknown. Prior to the decision making, the agent can learn about the relationship between the predicted variable and the explanatory variables from past realizations of the variables. During the decision making, the agent chooses what costly variables to acquire and predicts based on the acquired variables. The agent's utility reflects both the prediction correctness and the costs

of the acquired variables.

We consider the formulated model to belong to the family of rational inattention models by its core nature. The agent actively decides about the information to act upon. Nevertheless, two main differences distinguish the model from the general model of choice under rational inattention (Maćkowiak et al. (2018)). First, the information structure the agent acquire is not a homogeneous continuous signal. Instead, the agent acquires measurable variables with a specific meaning. Hence, the information structure in our model is heterogeneous. The heterogeneity allows us to describe decision making in situations when different pieces of information contain a different value for the decision making and, similarly, when different pieces of information are more or less difficult to process (and thus are more or less costly). Second, the agent has to learn about the uncertainty from the past observations (which we can perceive as experience) and does not have full control over the uncertainty. The agent only controls what variables to acquire. In contrast, the agent sets the uncertainty to any desired level in the general model of choice under rational inattention. These two modification allows us to describe agent's behaviour in real-life situations where the variables and their costs are measurable.

Finally, the contribution of the paper is the following. First, the paper proposes data-driven methods which serve as an approximation of behaviour of economic agents who face the trade-off between prediction uncertainty and variable costs while utilizing the past observations with an unknown distribution. Second, the paper provides solution methods usable in any area where predictions require explanatory variables whose acquisition might be costly. Such areas are, for instance, medical diagnosis, computer adaptive testing, various areas of computer science, or credit and insurance scoring, where the similar trade-off exists as well.

This paper is structured as follows. In section 2, we formulate the problem of acquisition of costly information in data-driven decision making. In section 3, we separate the problem into two parts: optimal prediction and acquisition of costly variables. We address each part of the problem separately and propose a full and an approximate solution to both parts following the expected utility approach. In section 4, we evaluate the proposed methods on two medical datasets and discuss the resulting behaviour.

2 Acquisition of costly information in data-driven decision making

Assume an agent whose goal is to predict a value of random variable y. Further, assume the agent can observe a vector of random variables $x = (x_1, x_2, ..., x_k)$ where individual variables x_i provide information about the value of y. To acquire a value of some of the variables in x, the agent has to exert some effort or spend money making the acquisition of variables costly. Random variables y and x come from unknown distribution $\mathbf{D}$. We assume that the distribution cannot be expressed analytically and estimated. The agent is endowed with a sample of n independent observations realized from the distribution $\mathbf{D}$:

$$(\boldsymbol{y}, \boldsymbol{X})$$

The agent decides which of costly variables $x_1, x_2, ..., x_k$ to acquire and what prediction $\widehat{y}$ to set. We assume that the agent has no other prior knowledge of variables y and x and relies on the observed sample only. Hence, we can refer to the decision problem as *acquisition of costly information in data-driven decision making*.

2.1 Preferences

First, we assume the agent receives utility based on prediction $\widehat{y}$ and the true value of y. The agent decides about the prediction and its correctness contains some value for the agent. To capture the agent's preferences about prediction $\widehat{y}$ and the true value of y, we assume utility function u_0:

$$u_0\left[\widehat{y}, y\right]$$

We will refer to u_0 as prediction utility since the utility stems purely from the correctness of prediction of y.

The specific form of prediction utility depends on the nature of the decision problem. In this paper, we will focus on a multiple choice problem where y belongs to one of K classes from a discrete set, and the goal is to predict the correct class. We could extend the problem statement and the solution to other types of decision making as well (e.g regression, counts, ordered choice, etc.), but we omit this possibility.

For the K-class classification, we can think of the prediction utility as a matrix of real

numbers $\boldsymbol{U} \in \mathbb{R}^{K \times K}$:

$$\boldsymbol{U} = \begin{pmatrix} u_{11} & u_{12} & \dots & u_{1K} \\ u_{21} & u_{22} & \dots & u_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ u_{K1} & u_{K2} & \dots & u_{KK} \end{pmatrix}$$

The rows of matrix $\boldsymbol{U}$ refer to utilities from predicted class $\widehat{y}$ and the columns refer to utilities from specific true class y. When the agent predicts class A while y belongs to class B, the prediction utility is:

$$u_0[\widehat{y} = A, y = B] = u_{AB} \quad \text{for} \quad A, B \in \{1, 2, \dots, K\}$$

Second, the acquisition of variables $x_1, x_2, \dots, x_k$ is costly. To reflect the costs in the agent's preferences, we assume that each variable x_j a has fixed cost of acquisition c_j. The costs might stem from financial expenditures, time requirements, inconvenience, or any other burden associated with the acquisition. To indicate which variables the agent has acquired, we will use binary indicator vector z with $z_j = 1$ if variable x_j has been acquired and $z_j = 0$ otherwise. The overall costs incurred to all acquired variables are a sum of the individual costs:

$$\sum_{j:z_j=1} c_j$$

Finally, the agent's overall utility is the difference between the prediction utility and the overall costs:

$$u\left[\widehat{y}, y, z\right] = u_0\left[\widehat{y}, y\right] - \sum_{j:z_j=1} c_j$$

The overall utility depends on both the prediction and the acquired variables and implies *a trade-off between an improvement of prediction and costs incurred to the acquisition of variables*. The agent maximizes the expectation of the overall utility over the unknown distribution **D**:

$$\mathbb{E}_{x,y \sim D(x,y)} \left\{ u_0[\widehat{y}, y] - \sum_{j:z_j=1} c_j \right\}$$

There are two aspects related to the utility form. First, the prediction utility and the costs are additive. Second, there are no dependencies between the costs of the variables. We assume that the acquisition cost of a variable is not related to the cost of any other

variable.

2.2 Decision framework

Knowing the agent's preferences, we can turn to the decision process itself. At the beginning, pair (y, x) is sampled from distribution **D**. The agent does not know the sampled value of any of the variables $x_1, \ldots, x_k$ or target variable y. The set of known values of costly variables is empty and $z = (0, 0, \ldots, 0)$. To represent the agent's current knowledge about x, we will denote the set of acquired observations as x^z. E.g. when the agent acquired x_1 whose value is 1 and x_2 whose value is 0.5, we have:

$$x^z = \{x_1 = 1, x_2 = 0.5\}$$

We do not assume any constraints to the acquisition of costly variables. The agent is allowed to acquire any costly variable $x_1, x_2, \ldots, x_k$ and set prediction $\widehat{y}$ to any admissible value. Since there are no costs associated with postponing the acquisition, the optimal strategy is to acquire variables sequentially. The agent can condition the next decision on the observed value of the first variable and make a more informed decision about further acquisition and prediction. At worst, the agent will acquire the second variable and receive the same utility as when acquiring both variables at once. For this reason, we focus on a sequential acquisition of variables:

At the beginning the agent decides between $k + 1$ options: predict target variable y directly or acquire one of k costly variables $x_1, x_2, \ldots, x_k$. The direct prediction results in utility $u = u_0[\widehat{y}, y]$, where prediction $\widehat{y}$ is made without any of the costly variables, and the decision making ends. When acquiring variable x_j, the agent pays cost c_j, and agent's current observation is $x^z = \{x_j\}$. Given the observation, the agent either decides to predict based on the observed variable x_j or acquires one of $k - 1$ remaining variables. The process proceeds until the agent decides to predict target variable y. Clearly, once the agent acquires all of the variables $x_1, x_2, \ldots, x_k$, the only remaining option is to predict. The sequential decision process for three costly variables is depicted in Figure 1. For simplicity, we will refer to the acquisition of information about variable x_j as action a_j and the decision to predict target variable y (and end the decision process) as action a_0.

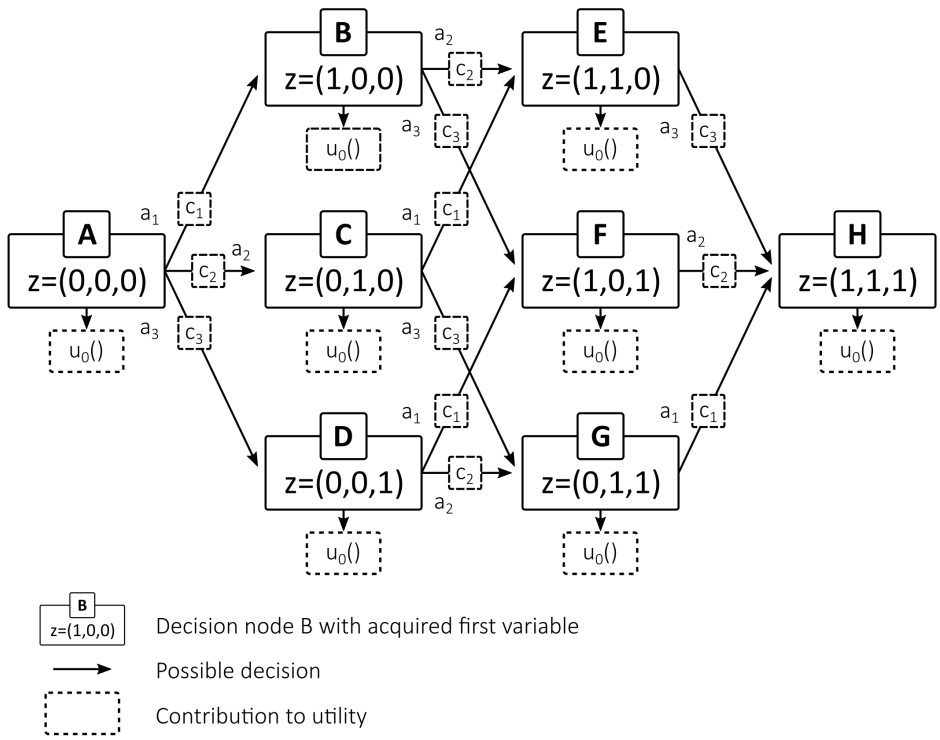

Figure 1. A scheme of sequential decision process for three costly variables

The decision process starts in decision node A with the agent not knowing any of the costly variables. The agent can terminate in decision node A by a direct prediction of target variable y and receive the prediction utility u_0. Or, the agent can decide to acquire information about a new variable. For instance, by acquiring variable x_1 (action a_1), the agent moves to decision node B, and the agent's utility decreases by c_1. From the node B, the agent can move to nodes E and F, etc. The utility is then given by the prediction utility in the node where the agent predicts minus the sum of costs on the path to the node.

3 Directed acyclical graph of supervised models

In this part, we propose to use supervised machine learning models for the approximation of the agent's optimal behaviour. We train a collection of supervised models on the observed sample $(\boldsymbol{y}, \boldsymbol{X})$ which performs both the acquisition of variables and the prediction of target variable y for any new example. To train the collection, we split the problem into two parts:

- Static prediction based on acquired costly variables - represented by a predictive model denoted as f
- Acquisition of costly variables - represented by an acquisition policy denoted as π

Predictive model f passively excepts the acquired costly variables serving as inputs to the model and predicts the target variable y. We train model f to return a prediction that maximizes the expected utility conditioned on the values of the acquired variables. The policy passively excepts the acquired costly variables and returns one of the available actions. The process repeats until the policy decides to predict with model f. We train the policy to select an action associated with the highest expected utility assuming the optimal decision making in the subsequent steps. The sequential nature of the decision problem allows us to train the policy in a backward fashion.

3.1 Predictive model

Now, we focus on the static prediction performed by predictive model f. The goal is to train model f (its weights) so that it returns optimal prediction $\widehat{y}^*$ for any combination of acquired variables indicated by z:

$$\widehat{y}^* = f(x^z, z)$$

We consider two approaches towards the model structure.

3.1.1 Predictive model for each decision-node

The first and straightforward structure of model f is a collection of individual models $f_z : \mathcal{X}^z \to \mathcal{Y}$ each operating on single decision node z:

$$\widehat{y}^* = f(x^z, z) = f_z(x^z)$$

where $\mathcal{X}^z$ denotes the space of values of acquired variables and $\mathcal{Y}$ denotes the space of possible values of y.

Obviously, the choice of model f_z and its training depends on nature of the task. In our case, we focus on the K-class classification. For both binary and multi-class classification, there is a plethora of supervised machine learning models (see Bishop, 2006). We will consider only models estimating the probability that a class is correct rather than directly classifying an example. Using the estimated probability of each class conditioned on current observations x^z, we select a class associated with the highest expected utility.

Hence, model f_z represents two steps: we predict the probability of a class and select the class associated with the highest expected utility. The first step is performed by the supervised classification model trained on the data containing acquired variables only. The second step is an evaluation of the expected utility of each class and the selection of the best class. Thus, the output of the model f_z is a class maximizing the expected utility.

Training one model for each possible combination of variables can be though very impractical. The number of required models grows exponentially with the number of costly variables. The exponential increase is not an issue for a lower number of costly variables, but training a collection of models for a higher number of costly variables is intractable. For that reason, we propose a second approach towards the structure of f.

3.1.2 A single predictive model on extended examples

We propose to train a single predictive model excepting any combination of acquired variables. To train the model, we adjust the structure of the data on which the model is trained. We extend each example by the indicator variable z and fill unknown values of x with zero. The indicator variable z ensures that the "true" zero and the zero standing for a missing value are distinguishable (see Dulac-Arnold et al., 2011; Shim et al., 2018). We refer to this process as "masking". The model then takes the masked extended examples as its inputs.

To generate the learning data, we duplicate each example in the original training data d-times. For each duplicate, we sample indicator z uniformly without replacement from the set of admissible indicators, mask the variables according to sampled z (we consider the variables as not acquired), and append z to the masked example. The higher d, the

more information is kept in the original data and usually the higher the performance of the predictive model. When $d = 2^k$, then each example is masked with all possible masks and the extended data contain full information provided by the original training data. As a consequence of the random sampling, the lack of a variable is not informative. Generally, training a model on more examples implies higher performance in inference. Since we do not know which variables will be acquired by the policy, the uniform sampling ensures that the model performs comparably across all combinations of potentially missing variables.

Notice that the selected model needs to have a sufficient capacity to be able to learn from such a data structure. For instance, we might consider tree-based models such as random forest or gradient boosted trees or a simple artificial neural network. Simple linear models such as logistic regression are not suitable for the encoded data.

The idea of pre-training the predictive model is similar to pre-training a neural network for prediction actions introduced by Janisch et al. (2017) in the context of deep reinforcement learning. The sampling procedure of the data is though different and the pre-trained neural network is supposed to change during the training. In contrast, our model f is fixed and independent from acquisition policy π.

3.2 Acquisition policy

Once we trained the predictive model f, we can find a suitable policy π conducting the acquisition of costly variables.

3.2.1 Objective and structure

Denote $z(\pi)$ as the indicator variable for variables acquired by policy π. Our goal is to find policy π^* that maximizes average utility on the observed sample:

$$\pi^* = \underset{\pi}{\operatorname{argmax}} \frac{1}{n} \sum_{i=1}^{n} \left\{ u_0 \left[f\left(\boldsymbol{X}_i^{z(\pi)}, z(\pi)\right), \boldsymbol{y}_i \right] - \sum_{j:z(\pi)_j=1} c_j \right\}$$

The assumed policy is dependent on the values of each example $\boldsymbol{X}_i$ and allow a fully flexible acquisition of variables conditioned on the already acquired variables. As a result, the policy might acquire a different combination of costly variables for each example i.

The structure of the acquisition policy is as follows. We train individual "node policies" π_z each operating on a given combination of acquired variables according to z. Node

policy takes as inputs the acquired variables x^z and returns one of the available actions:

$$\pi_z : \mathcal{X}^z \rightarrow \{a_0, a_j : z_j = 0\}$$

Each node policy consists of several supervised models that approximate the expected utility of available actions in the decision node. Policy π is then represented as a collection of node policies π_z that consist of several supervised models.

3.2.2 Full backward solution

We train the individual node policies in a backward fashion. We start by training the node policy in the last decision node and proceed sequentially until we train the node policy in the first decision node. The learning process is depicted in Figure 7. The "learn" function performs the learning process of a single noted policy. The function takes as inputs the predictions of model f_z and outputs of the already trained policies and trains a new node policy.

Assume we are training policy π_z^* in decision node z with l acquired variables ($0 \leq l \leq k$), and we have already trained the previous node policies. We split the training of the node policy into the evaluation of individual available actions. To evaluate a_0, we perform

Figure 2. A scheme for learning an optimal policy backwards

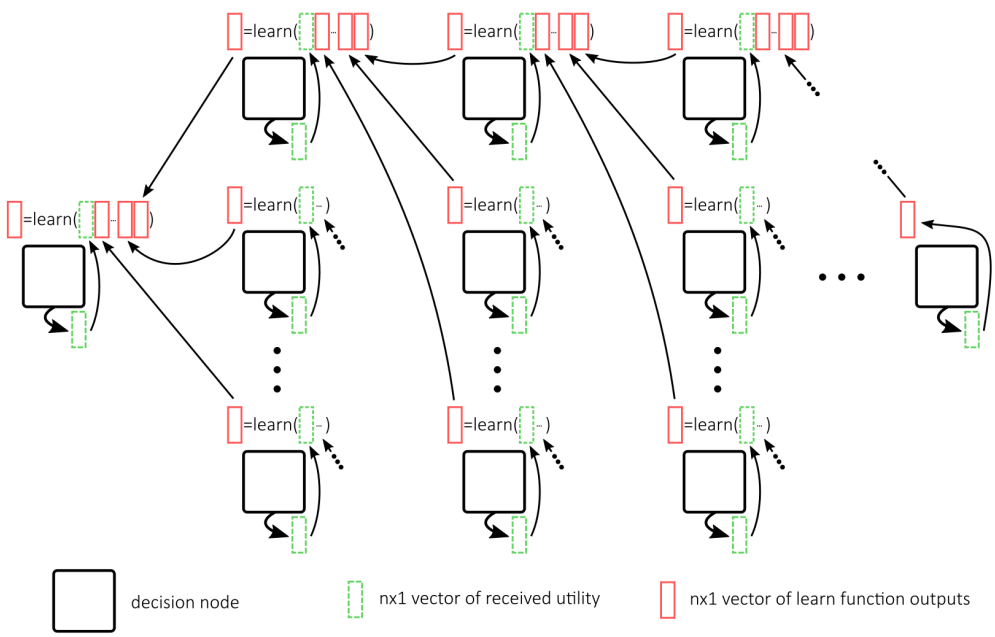

the following steps:

1. We perform prediction using the classification model (the first step in f) for all examples in training data $\boldsymbol{X}^z$. We receive $n \times K$ matrix of probabilities $\boldsymbol{P}$.

2. We evaluate the (estimated) expected utility of each class and each example as $\boldsymbol{P} \times \boldsymbol{U}'$ minus incurred costs on each example. We obtain an $n \times K$ matrix from which we select the maximum of each row. We obtain $n \times 1$ vector of (estimated) expected utilities from the best class for each example $\widehat{\boldsymbol{q}}[z, a_0]$.

When we are in a decision node with all the acquired variables and there is no other available action, we set the evaluation of the node directly as:

$$\widehat{v}[z] = \widehat{\boldsymbol{q}}[z, a_0]$$

and we always select action a_0. The evaluation of decision node z represents the expected utility when being in decision node z and deciding optimally thereafter.

Figure 3. A scheme for training a node policy as an action with the maximum expected utility

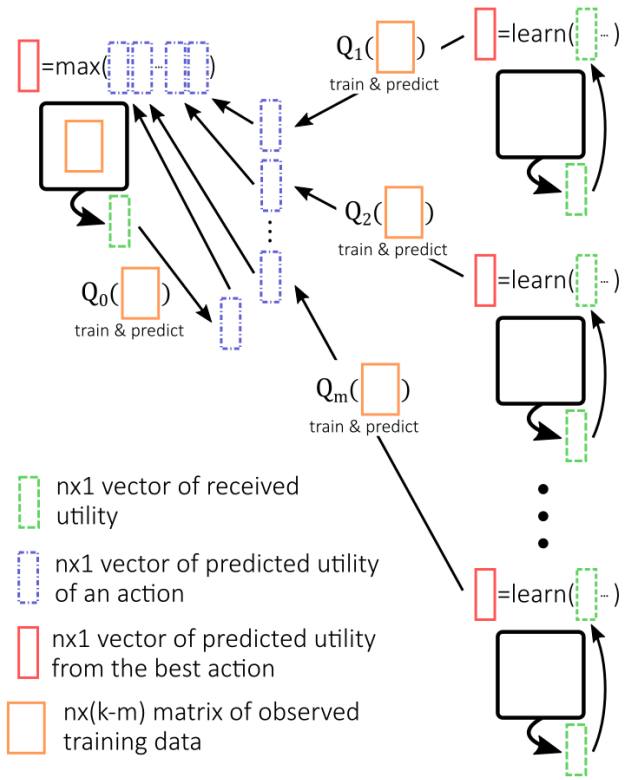

To evaluate action a_m from the remaining actions, we perform the following steps:

1. We take the evaluation $\widehat{v}[z(m)]$ of adjacent decision node $z(m)$ to which we get by performing action a_m in decision node z.

2. We train regression model $Q_m : \mathcal{X}^z \to \mathbb{R}$ on inputs $\boldsymbol{X}^z$ with target $\widehat{v}[z(m)]$.

3. We predict with regression model Q_m on inputs $\boldsymbol{X}^z$ and obtain an estimate of expected utility when performing action a_m for given x^z. We denote the vector of predictions as $\widehat{\boldsymbol{q}}[z, a_m]$.

We evaluate the decision node z with the expected utility for each example i as:

$$\widehat{v}[z]_i = \max_a \left\{ \widehat{\boldsymbol{q}}[z, a]_i \right\}$$

During prediction phase, we evaluate each of the available actions a_m with the trained regression model Q_m and select the action with the highest expected utility for each example in testing data. The node policy π_z^* is thus given by the collection of the regression models and the selection of an action with the highest fitted value. The choice of regression model Q_m depends on the expected data structure. Generally, it can be any suitable regression model. For an overview of possible models see e.g. Bishop (2006). A scheme showing the training process is depicted in Figure 3.

3.2.3 Greedy approach

The backward training of acquisition policy is an exhaustive approach that comes at a high computational cost. Its advantage is that it approximates the optimal policy directly, though it requires training exponentially many (2^k) node policies. For a higher number of costly variables, the method in not computationally feasible. In this section we propose a less accurate but computationally more efficient alternative. Instead of training a node policy on the expected utility estimate from the full depth of the decision graph, we train it using the expected utility estimate restricted to an arbitrary depth m. In other words, we allow the agent to evaluate the expected utility from the acquisition of at most m additional costly variables.

In each node, we construct a vector of expected utilities for each action $\widehat{\boldsymbol{q}}^m[z, a]$ estimated from the limited depth m. Then, based on the estimates $\widehat{\boldsymbol{q}}^m[z, a]$, we train node policy π_z which acquires the next variable for each example. Then we estimate the vector

of expected utilities $\widehat{\boldsymbol{q}}^m[z, a]$ in the subsequent decision node(s) z. We proceed until the policy selects prediction for all the examples.

We compute the m-depth estimate of expected utility as follows. Consider a subset of n' examples from training data $\boldsymbol{X}^z$ for which we observed costly variables given by z. Our goal is to construct $n' \times 1$ vector $\widehat{\boldsymbol{q}}^m[z, a]$ of utility estimates for each of the examples. We proceed similarly as in the case of the full solution but instead of starting in the last decision node with all acquired variables we start in all nodes with m additionally acquired variables. In the case that $m = 1$, we consider only nodes adjacent to node z. In all of the nodes, we exclude the possibility of acquisition of further variables and set the evaluation of each node z' directly as the utility from action a_0:

$$\widehat{\boldsymbol{v}}[z'] = \widehat{\boldsymbol{q}}[z', a_0]$$

This allows us to ignore the utility stream from more distant decision nodes at the expense of an imprecise estimate.

For the evaluation of the predict action a_0, we can either utilize a pre-trained predictive model on the extended dataset or train predictive model $f_{z'}$ whenever the predictive model is needed. The learning process is hereafter identical as for the full solution. An illustration of an estimate from depth $m = 2$ is depicted in Figure 4.

Having $n' \times 1$ vector $\widehat{\boldsymbol{q}}^m[z, a]$ of utility estimates, we train node policy π_z on the utility estimate serving as the target and subset of n' examples of training data $\boldsymbol{X}^z$ serving as inputs. Then we apply the policy on the same examples and select a new action for each example. Since we select a different action for different examples, we move simultaneously to different decision nodes z. The process results in a gradual partition of training data based on the observed values of the costly variables. Whenever the policy selects action a_0 for an example, we disregard the example from the future acquisition. Next, for each of the selected decision nodes, we calculate an m-step estimate and repeat the whole process. Since the structure of the decision graph is not a tree, it might happen that examples in a decision node come from different paths. For this reason, we always evaluate decision nodes with the lowest number of acquired variables first. Then we can train the node policy on all the examples coming from parent decision nodes.

Figure 4. m depth estimate of utility from actions

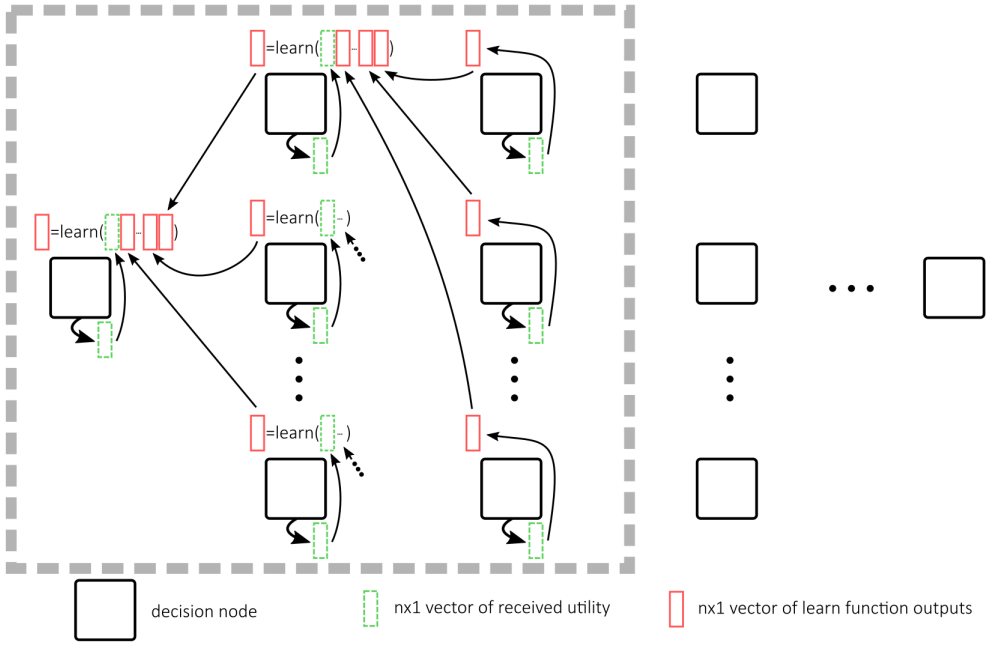

During the prediction phase, we disregard node policies trained as a part of m-step estimation and consider only those trained on the estimated utility $\widehat{\boldsymbol{q}}^m[z, a]$. Additionally, we prevent the examples to follow a new path that was not selected during the training phase. Otherwise, it could happen that we would not have a trained node policy selecting the next actions.

4 Experiments on medical datasets

In this section, we evaluate the proposed methods on two publicly available medical datasets. In the first dataset, euthyroid, the goal is to predict whether a patient suffers from euthyroid sick syndrome utilizing several examinations and other information about the patient. In the second dataset, NHANES diabetes, the goal is to predict one of three patient's health conditions: normal, prediabetes, and diabetes utilizing demographic information and results obtained from examinations, laboratory tests, and questionnaires.

4.1 Euthyroid

The euthyroid dataset is publicly available on `http://archive.ics.uci.edu/ml/machine-learning-d` `thyroid-disease/`. We filter all examples with missing values since the missing values cannot be acquired during the experiments. The resulting dataset contains 2000 observations

with 1762 negative examples and 238 positive examples. Next, it contains 17 variables each associated with an acquisition cost. For the experiment, we consider only costs associated with medical tests and ignore unit costs associated with acquisition of variables such as patient's sex or age. The costs of medical tests are following:

test	TSH	T3	TT4	T4U
cost	22.78	11.41	14.51	11.41

The utility function stemming from diagnosis is unknown, but we set it as:

$$\boldsymbol{U} = \beta \begin{pmatrix} 0 & -10 \\ -1 & 0 \end{pmatrix}$$

We assume that classifying a patient as healthy when having the disease (false negative) is 10 times more costly than classifying a patient as ill when being healthy (false positive). Furthermore, we assume that there are no costs associated with a correct diagnosis. We parametrize the utility by coefficient β and experiment with its several values. Generally, higher β implies lower relative costs of the medical tests. We use random forest for both the classification and the node policy, but any combination of suitable supervised models could be used. We train the classification model and the policy on 75% of examples and keep 25% of examples for testing.

The experimental results are depicted in Figure 5, Figure 6, and Figure 7. Figure 5 compares mean (dis)utility attained on the testing sample for various values of β. We can see that for $\beta = 20$ the selected methods achieve a similar utility as a static prediction with free variables only. The static prediction with all variables provides a very low utility due to the high costs of the medical tests forming most of the disutility. With higher β the gap in utilities of the selected methods and static prediction with free variables widens. For $\beta = 100$ the evaluated methods achieve a considerably higher utility compared to static predictions with all and free variables.

The graph illustrates that *the potential gain from the dynamic acquisition of variables depends on the trade-off between costs and predictions*. If the costs are high (β is low) then the optimal policy is to acquire almost no costly variables, and there is nothing to gain over the static classification with free variables. With higher values of β the gain over static prediction

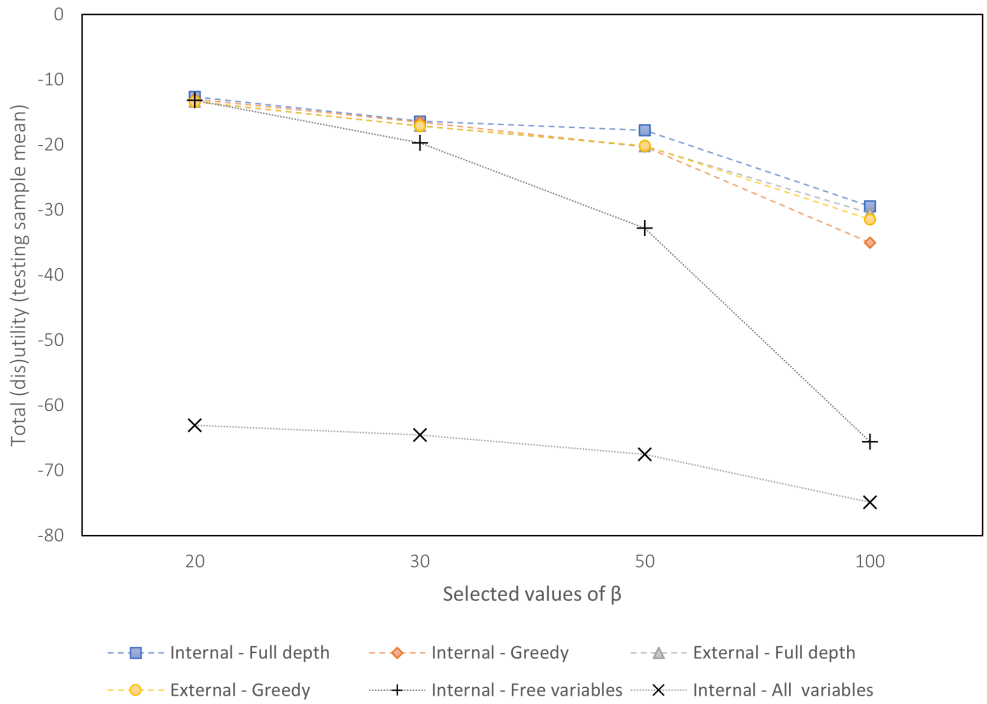

Figure 5. A comparison of methods for the euthyroid dataset in mean (dis)utility
A comparison of methods for the euthyroid dataset in mean (dis)utility on the testing sample for selected betas. "Internal" refers to a classifier trained for each subset of costly variables during the acquisition process, "External" refers to a classifier pre-trained on an extended dataset prior to the acquisition. "Full depth" refers to exhaustive backward training, "Greedy" refers to the greedy acquisition of costly variables with $m = 1$.

with free variables increases since the variables are relatively less costly and provide a higher prediction utility. The opposite holds for the gain over the static classification with all variables. With higher values of β the optimal policy acquires more and more costly variables. In an extreme case, the optimal policy could acquire all the costly variables reaching the utility of the static classification with all costly variables.

Next, the graph compares the selected methods themselves. The classifier trained on each subset of costly variables combined with the exhaustive backward training of the policy (blue square) outperforms the three remaining methods for all values of β. The high performance is expected since the method does not involve any simplifications neither in classification nor in the policy training. The difference between the greedy acquisition of variables and the exhaustive backward training is visible for the collection of classifiers only. For the pre-trained classifier, the greedy acquisition and the backward training perform comparably. Nonetheless, the methods are *data-driven*, and their performance is very much dependent on the specific data.

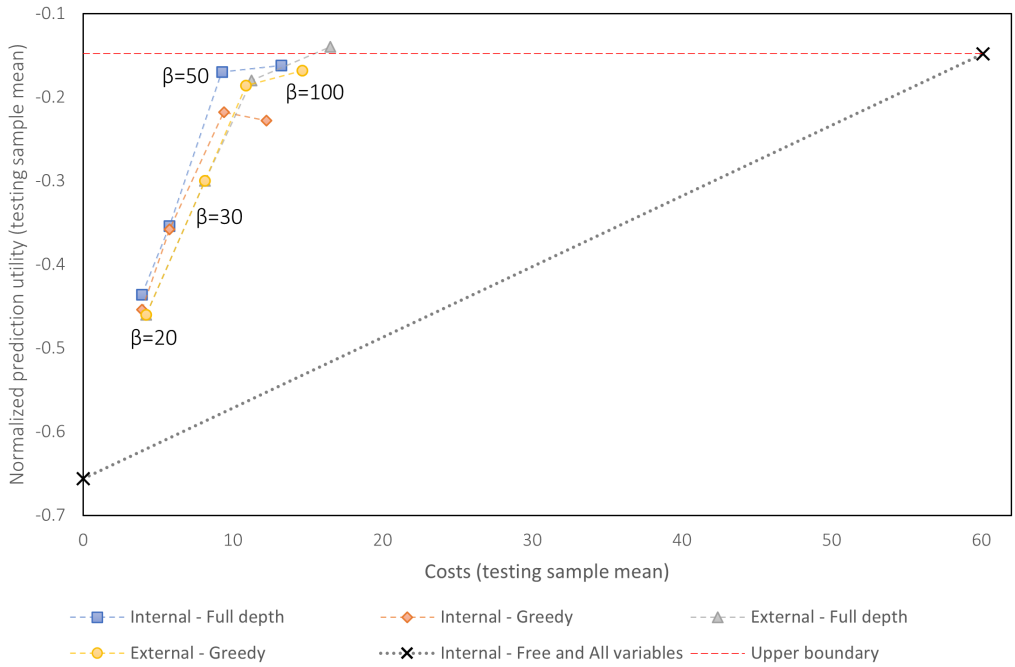

Figure 6. A comparison of methods for the euthyroid dataset by the trade-off costs and prediction utility

A comparison of methods for the euthyroid dataset by the trade-off between costs and prediction utility for selected values of β. The attained prediction utility is scaled down by corresponding β to make the prediction utilities comparable across various values of β. The legend matches the legend in Figure 5

Figure 6 depicts the trade-off between the prediction utility and the costs of variables. On the vertical axis, we have a normalized prediction utility which refers to the prediction utility when $\beta = 1$. The prediction performance of each method is then comparable across various values of β. On the horizontal axis, we have the costs incurred to the medical tests. When predicting without any medical test, the total costs are zero (a black cross in the left bottom corner). When predicting with all the medical tests, we achieve the highest possible prediction utility (we have maximal information) but at the expense of the maximal costs (a black cross in the right upper corner). The highest prediction utility is depicted by the horizontal dashed red line and constitutes an upper boundary for the prediction utility attained by other methods. We can see that there is a potential to improve the prediction utility by the acquisition of costly medical tests.

Obviously, we aim at achieving the highest prediction utility while incurring the lowest costs. We can see that all the selected methods acquire medical tests efficiently. With higher β the prediction utility increases sharply, while the incurred costs on the medical

tests grow relatively slowly. The most efficient method is the backward training combined with the classifier trained for each subset of the medical tests. With exception of $\beta = 100$, it achieves the highest prediction utility, while incurring the lowest costs. Recall that the methods are *data-driven*. Not only depends *the potential efficiency* of the methods on values of β, but it also *depends on the informativeness of individual medical tests*. For instance, if there is one cheap and very informative medical test, then the prediction utility increases sharply at the low costs. The other tests are not acquired and the prediction is highly cost-efficient. On the other hand, if all the tests provide only partial information and are moderately costly (the net utility gain of each test is slightly above zero in general), then we cannot expect a sharp increase in the prediction utility. The costs of such medical tests offset the prediction utility and the added value of the dynamic acquisition is low. It does not matter what (and if at all) the medical tests are acquired. The informativeness of medical tests is embedded in the externally given data structure, and there is nothing we can do about it. Nevertheless, the proposed acquisition methods approximate the optimal acquisition regardless of its potential improvements over the static classification.

Figure 7 depicts the acquisition of medical tests on testing data for the backward train-

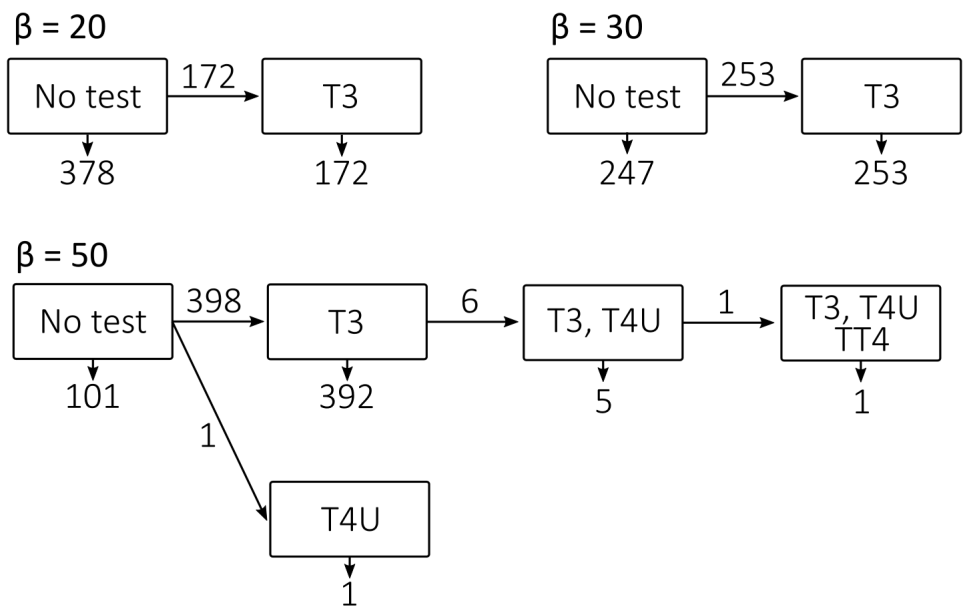

Figure 7. Acquisition paths on the euthyroid dataset

The euthyroid dataset: Acquisition paths for selected values of β on the testing sample for "Internal - Full depth" method. A patient is either diagnosed directly or other tests are requested. The numbers represent the number of patients in the testing sample for which the action was selected. Down arrows represent prediction.

ing with a classifier trained on all subsets of medical tests. For $\beta = 20$, we diagnose 378 patients directly without requesting additional medical tests. For the remaining 172 patients we request medical test "T3" before their diagnosis. For $\beta = 30$, we value the correct diagnosis more and request the medical test "T3" for 253 of the patients. For $\beta = 50$, the examination process is more complex. We request test "T3" for most of the patients and request additional medical tests for few other patients. For one patient we request test "T4U" as the first test instead of test "T3". The acquisition of medical tests is consistent with theory. With a higher emphasis on a correct diagnosis, we acquire more costly tests.

Additionally, there is an important caveat related to the three figures. It might be tempting to select β providing the most convenient results. E.g. $\beta = 50$ in Figure 6 provides an attractive trade-off between prediction utility and costs. However, selecting β contradicts the nature of the ACID problem. We assume an externally given utility function already containing the trade-off. Our objective is not to optimize the trade-off but the overall utility composed of prediction utility and costs. We have no free choice over β since β is embedded in the decision problem from the very beginning. On the other hand, determining a proper utility specification might be troublesome in many practical applications. Analyzing the trade-off between prediction utility and costs as depicted in Figure 6 might provide useful guidance.

4.2 NHANES diabetes

The NHANES diabetes dataset is accessible on `https://github.com/mkachuee/Opportunistic` and has been constructed by Kachuee et al. (2019). The dataset contains 92062 observations with 17% of healthy patients (marked as 0 in the dataset), 80% of patients with pre-diabetes (1) and 2.5% patients with diabetes (2). The dataset contains 45 costly variables in four categories: demographic, laboratory, examination, and questionnaire. For our experiments, we consider costs stemming from laboratory and examination variables only. There are 4 laboratory variables each with costs 9 and 13 examination variables with costs 5. Overall, we consider 17 costly variables and 38 free variables. The utility matrix from diagnosis is unknown, but we set it as:

$$\boldsymbol{U} = \beta \begin{pmatrix} 0 & -10 & -20 \\ -1 & 0 & -10 \\ -2 & -1 & 0 \end{pmatrix}$$

which approximately reflects the idea that overlooking more severe health conditions is more costly than a wrong diagnosis of less severe health conditions. Because of the high number of costly variables, we consider the greedy acquisition only. We parametrize the utility matrix by coefficient β and use a random forest for classification and the node policy. We train the classification model and the policy on 70% of examples and keep 30% of examples for testing.

The outcomes of the experiment are depicted in Figure 8 and Figure 9. Figure 8 compares the mean (dis)utility attained on the testing examples similarly as Figure 5 for the experiment on the euthyroid dataset. We observe that classification with all variables results in a very high disutility compared to classification with free variables only. The reason is that the total costs of the 17 costly variables are 185. The utility attained from the greedy acquisition is comparable to utility from classification with free variables for all values of β. The graph suggests that most of the variables are too costly and provide low information. Figure 9 compares the trade-off between prediction utility and costs for the methods similarly as Figure 6. The figure is truncated to the maximal costs of 10. The horizontal dashed red line represents the prediction utility attained for classification with all variables

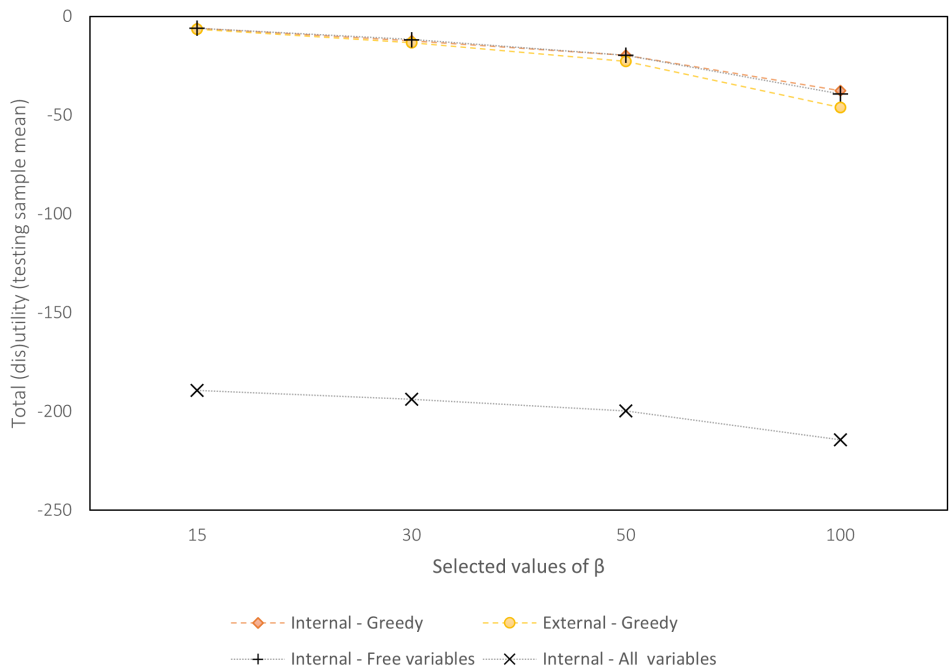

Figure 8. A comparison of methods for the NHANES diabetes dataset in mean (dis)utility
A comparison of methods for the NHANES diabetes dataset in mean (dis)utility on the testing sample for selected betas. The legend matches the legend in Figure 5.

and constitutes the upper boundary for the prediction utility. The black cross in the left bottom part of the graph corresponds to classification with free variables only. The upper boundary and the (almost) horizontal dotted black line intersect outside of the graph when costs are equal to 185. When $\beta = 15$ the classifier trained for all acquired subsets overlaps with the classification with free variables, but the pre-trained classifier achieves lower prediction utility. The reason is that the pre-trained classifier has lower classification performance in general (because of the computation constraints we duplicated the dataset only 10 times).

For higher values of β, the greedy acquisition combined with a collection of classifiers acquires several variables very efficiently. When $\beta = 100$ the greedy acquisition closes half of the gap in prediction utility between classification with all and costly variables using variables with only around 2% of the overall costs. The result suggests that only very few variables provide sufficient information compensating their costs. On the other hand, the greedy acquisition combined with the pre-trained models does not seem to utilize the

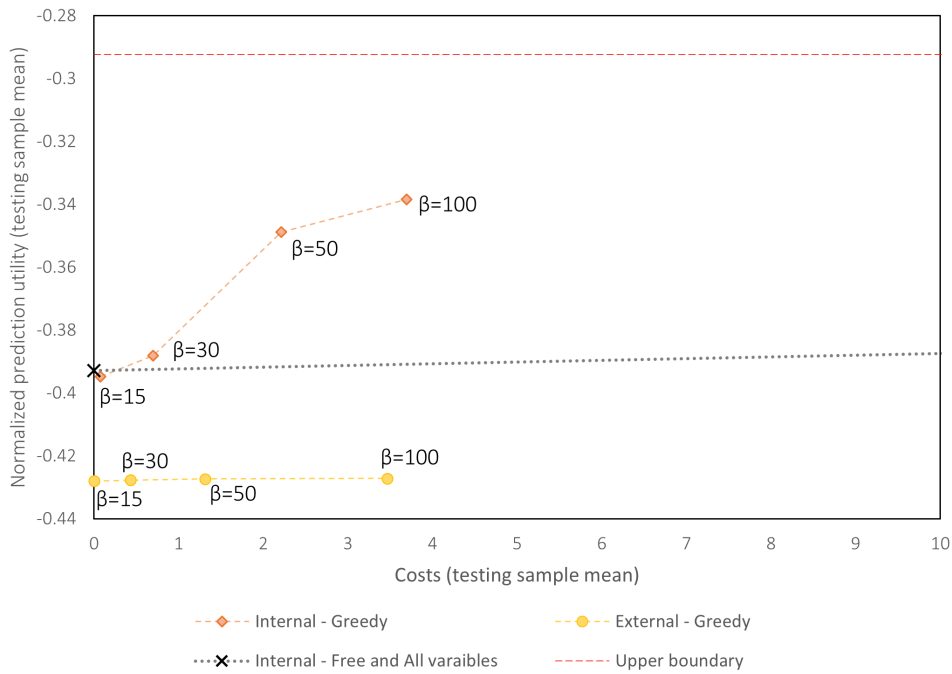

Figure 9. A comparison of methods for the NHANES diabetes dataset by the trade-off between costs and prediction utility

A comparison of methods for the NHANES diabetes dataset by the trade-off between costs and prediction utility for selected values of β. The attained prediction utility is scaled down by corresponding β to make the prediction utilities comparable across various values of β. The legend matches the legend in Figure 5.

acquired variables that efficiently. The reason might be the weaker pre-trained classifier. Overall, the experiment on the NHANES data suggests that for the prediction of diabetes the demographic and questionnaire data provide cost-effective information and only a few additional examinations have to be requested.

5 Conclusion

In a world abundant with information, agents face two decision problems: what information to acquire and internalize and how to act upon the information. For the acquisition and internalization of information, the economic decision theory describes a trade-off between costs of information and benefits of the information assuming that agents are fully aware of their uncertainty. In contrast to the economic theory, this paper introduces an economic decision problem where an agent does not know the uncertainty but has to learn about it from past observations. The paper assumes an agent whose goal is to predict a value of a random variable utilizing several costly explanatory variables. The agent does not know the distribution of the variables, but, prior to the prediction, the agent can learn about the relationship between the explanatory variables and the predicted variable from the past observations. During the prediction, the agent chooses what costly variables to acquire and trades off the costs of the variables and the information the variables provide.

The generality of the problem statement allows for an application to any area where the prediction requires explanatory variables whose acquisition is costly. The direct real-life applications are in medical diagnosis, credit scoring, computer adaptive testing, and several areas of computer science.

The paper presents a solution to the problem using machine learning methods that allow learning from the data with unknown distribution. The paper divides the solution into two rather separate parts: acquisition of costly variables and prediction based on costly variables. For the prediction, the paper presents an approach for training a single predictive model accepting any combination of acquired variables. The approach allows avoiding a training of an exponential number of predictive models for each combination of acquired variables.

For the acquisition of costly variables, the paper presents two solution methods estimating the expected utility: backward training of supervised machine learning models

forming a directed acyclical graph and a greedy acquisition serving as a computationally feasible yet less accurate alternative. Both methods allow for a fully flexible acquisition of costly variables conditioned on the already acquired variables. The paper evaluates the methods on two medical datasets and shows that the methods acquire costly variables efficiently.

References

Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer Science+Business Media.

Chen, M., Z. Xu, K. Weinberger, O. Chapelle, and D. Kedem (2012, 21–23 Apr). Classifier cascade for minimizing feature evaluation cost. In N. D. Lawrence and M. Girolami (Eds.), *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, Volume 22 of *Proceedings of Machine Learning Research*, pp. 218–226. PMLR.

Chen, S., A. Choi, and A. Darwiche (2015). Computer adaptive testing using the same-decision probability. In *Proceedings of the Twelfth UAI Conference on Bayesian Modeling Applications Workshop - Volume 1565*, pp. 34–43. CEUR-WS.org.

Dulac-Arnold, G., L. Denoyer, P. Preux, and P. Gallinari (2011, 08). Datum-wise classification: A sequential approach to sparsity. Volume 6911.

Early, K., S. Fienberg, and J. Mankoff (2016, 09). Test time feature ordering with focus: interactive predictions with minimal user burden. pp. 992–1003.

Janisch, J., T. Pevný, and V. Lisý (2017). Classification with costly features using deep reinforcement learning. *CoRR abs/1711.07364*.

Kachuee, M., O. Goldstein, K. Kärkkäinen, S. Darabi, and M. Sarrafzadeh (2019). Opportunistic learning: Budgeted cost-sensitive learning from data streams. *CoRR abs/1901.00243*.

Kachuee, M., K. Karkkainen, O. Goldstein, D. Zamanzadeh, and M. Sarrafzadeh (2019). Nutrition and health data for cost-sensitive learning. *CoRR abs/1902.07102*.

Kanani, P. and P. Melville (2008, 01). Prediction-time active feature-value acquisition for cost-effective customer targeting.

Karkkainen, K., M. Kachuee, O. Goldstein, and M. Sarrafzadeh (2019). Cost-sensitive feature-value acquisition using feature relevance.

Krause, A. and C. Guestrin (2014). Optimal value of information in graphical models. *CoRR abs/1401.3474*.

Maćkowiak, B., F. Matejka, and M. Wiederholt (2018). Survey: Rational inattention, a disciplined behavioral model. CEPR Discussion Papers 13243, C.E.P.R. Discussion Papers.

Nan, F. and V. Saligrama (2017). Adaptive classification for prediction under a budget.

Raiffa, H. and R. Schlaifer (1961). *Applied Statistical Decision Theory* (1 ed.). Harvard University Press.

Raykar, V., B. Krishnapuram, and S. Yu (2010, 06). Designing efficient cascaded classifiers: Tradeoff between accuracy and cost. pp. 853–860.

Sheng, V. S. and C. X. Ling (2006). Feature value acquisition in testing: A sequential batch test algorithm. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, New York, NY, USA, pp. 809–816. Association for Computing Machinery.

Shim, H., S. J. Hwang, and E. Yang (2018). Joint active feature acquisition and classification with variable-size set encoding. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Eds.), *Advances in Neural Information Processing Systems*, Volume 31. Curran Associates, Inc.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics 50*(3), 665–690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

Sims, C. A. (2010). Chapter 4 - rational inattention and monetary economics. Volume 3 of *Handbook of Monetary Economics*, pp. 155–181. Elsevier.

Turney, P. D. (1995). Cost-sensitive classification: Empirical evaluation of a hybrid genetic decision tree induction algorithm. *CoRR cs.AI/9503102*.

Viola, P. and M. Jones (2001). Rapid object detection using a boosted cascade of simple features. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, Volume 1, pp. I–I.

Vivar, G., K. Mullakaeva, A. Zwergal, N. Navab, and S.-A. Ahmadi (2020). Peri-diagnostic decision support through cost-efficient feature acquisition at test-time.

Wang, J., T. Bolukbasi, K. Trapeznikov, and V. Saligrama (2014). Model selection by linear programming. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars (Eds.), *Computer Vision – ECCV 2014*, Cham, pp. 647–662. Springer International Publishing.

Wang, J., K. Trapeznikov, and V. Saligrama (2015). Efficient learning by directed acyclic graph for resource constrained prediction.

Xu, Z., M. Kusner, K. Weinberger, and M. Chen (2013, 17–19 Jun). Cost-sensitive tree of classifiers. Volume 28 of *Proceedings of Machine Learning Research*, Atlanta, Georgia, USA, pp. 133–141. PMLR.

Appendix

A Source code

The code used for the experiments is available in the following GitHub repository:

`https://github.com/OrangePiano/acquisition_of_costly_variables`

The repository contains all the information necessary for the replication of the results.

IES Working Paper Series

2022

1. Klara Kantova: *Parental Involvement and Education Outcomes of Their Children*
2. Gabriel Nasser, Doile de Doyle, Paulo Rotella Junior, Luiz Célio Souza Rocha, Priscila França Gonzaga Carneiro, Rogério Santana Peruchi, Karel Janda, Giancarlo Aquila: *Impact of Regulatory Changes on Economic Feasibility of Distributed Generation Solar Units*
3. Paulo Rotella Junior, Luiz Célio Souza Rocha, Rogério Santana Peruchi, Giancarlo Aquila, Karel Janda, Edson de Oliveira Pamplona: *Robust Portfolio Optimization: A Stochastic Evaluation of Worst-Case Scenarios*
4. Adam Kučera, Evžen Kočenda, Aleš Maršál: *Yield Curve Dynamics and Fiscal Policy Shocks*
5. Karel Janda, Ladislav Kristoufek, Barbora Schererova, David Zilberman: *Price Transmission and Policies in Biofuels-Related Global Networks*
6. Daniel Kolář: *Wealth Survey Calibration: Imposing Consistency with Income Tax Data*
7. Michal Hlaváček, Ilgar Ismayilov: *Meta-analysis: Fiscal Multiplier*
8. Salim Turdaliev, Karel Janda: *Increasing Block Tariff Electricity Pricing and the Propensity to Purchase Dirty Fuels: Empirical Evidence from a Natural Experiment*
9. Vojtěch Mišák: *Crime and weather. Evidence from the Czech Republic.*
10. Lukas Janasek: *Acquisition of Costly Information in Data-Driven Decision Making*

All papers can be downloaded at: http://ies.fsv.cuni.cz

Univerzita Karlova v Praze, Fakulta sociálních věd
Institut ekonomických studií [UK FSV – IES] Praha 1, Opletalova 26
E-mail : ies@fsv.cuni.cz http://ies.fsv.cuni.cz