**Everything has its Price: Foundations of Cost-Sensitive Learning and its
Application in Psychology**

Philipp Sterner¹, David Goretzko¹, and Florian Pargent¹

¹ LMU Munich

Author Note

Philipp Sterner <https://orcid.org/0000-0002-8856-4868>

David Goretzko <https://orcid.org/0000-0002-2730-6347>

Florian Pargent <https://orcid.org/0000-0002-2388-553X>

LMU Munich, Department of Psychology, Leopoldstrasse 13, D-80802 Munich.

Acknowledgements: We thank Anja Betz for checking the reproducibility of our manuscript and online materials.

Parts of the current article were presented at a conference talk by the first author at the 2021 Conference of the German Psychological Society - Personality Psychology and Psychological Diagnostics (DPPD) Section. A preprint of the article was published on PsyArXiv. The analyses scripts and supplementary material supporting this article are openly available on the Open Science Framework at <https://osf.io/cvks7/>.

The authors made the following contributions. Philipp Sterner: Conceptualization, Formal Analysis, Writing - Original Draft Preparation, Writing - Review & Editing; David Goretzko: Conceptualization, Formal Analysis, Writing - Review & Editing; Florian Pargent: Conceptualization, Formal Analysis, Supervision, Writing - Review & Editing.

Correspondence concerning this article should be addressed to Philipp Sterner, Leopoldstrasse 13, D-80802 Munich. E-mail: p.sterner@psy.lmu.de

Abstract

Psychology has seen an increase in machine learning (ML) methods. In many applications, observations are classified into one of two groups (binary classification). Off-the-shelf classification algorithms assume that the costs of a misclassification (false-positive or false-negative) are equal. Because this is often not reasonable (e.g., in clinical psychology), cost-sensitive learning (CSL) methods can take different cost ratios into account. We present the mathematical foundations and introduce a taxonomy of the most commonly used CSL methods, before demonstrating their application and usefulness on psychological data, i.e., the drug consumption dataset ($N = 1885$) from the UCI Machine Learning Repository. In our example, all demonstrated CSL methods noticeably reduce mean misclassification costs compared to regular ML algorithms. We discuss the necessity for researchers to perform small benchmarks of CSL methods for their own practical application. Thus, our open materials provide R code, demonstrating how CSL methods can be applied within the `mlr3` framework (<https://osf.io/cvks7/>).

Keywords: cost-sensitive learning, machine learning, misclassification costs, tutorial

Everything has its Price: Foundations of Cost-Sensitive Learning and its Application in Psychology

Psychological science has seen an increase in the use of machine learning (ML) techniques to analyze data. More focus has been put on *predicting* future behavior or outcomes instead of mostly *explaining* psychological phenomena (Orrù et al., 2020; Shmueli, 2010; Yarkoni & Westfall, 2017). In supervised ML, there are two important prediction settings: regression and classification (James et al., 2021). In regression, the dependent variable is continuous whereas in classification it is categorical. In the special case of binary classification the dependent variable is dichotomous: observations either belong to the so-called positive or to the negative class, with different conventions on how to assign those labels. The main goal of ML algorithms is to maximize accuracy in the prediction of new, unseen data. Regression models minimize some difference between predicted and true responses whereas classification models minimize the amount of misclassifications (i.e., predicting that an observation x belongs to class i when in reality it belongs to class j ; Elkan, 2001; Ling & Sheng, 2008). Classification is often less straightforward and specific problems arise. Off-the-shelf algorithms assume that all forms of misclassification errors – in binary classification: *false positives* (FP) and *false negatives* (FN) – are equally “bad” or, in more technical terms, have the same misclassification costs (MC). Yet in reality, this is often not the case (Domingos, 1999; Elkan, 2001; Zadrozny & Elkan, 2001a). In medical or psychological assessment, for example, classifying an ill person as healthy is often regarded as worse (“costlier”) compared to classifying a healthy person as ill. However, if ill persons are much rarer than healthy ones, a cost-insensitive algorithm could increase its prediction accuracy by classifying all persons as healthy (Elkan, 2001). Needless to say, the consequences of not treating anyone would be detrimental. To counteract this tendency to stubbornly maximize accuracy, researchers have developed multiple techniques that take different MC into account. In the special field of ML called *cost-sensitive learning* (CSL; Elkan, 2001), the goal of an algorithm shifts

toward minimizing the total MC. Many domains are already using CSL to build applicable models that put a premium on different MC, for example, business (predicting customer churn; Coussement, 2013), medicine (cancer detection; Freitas, 2011; Lu et al., 2019), and computer science (software defect detection; Siers & Islam, 2018). However, the usefulness of different CSL methods has not yet been discussed for psychological science. This is surprising because in many psychological applications, unequal MC are rather the norm than the exception. It can be argued that in suicide prediction, for example, missing a person who is at risk of attempting to commit suicide is far worse than wrongfully taking a second look at a person who actually is not. Unequal MC could also be assumed for drug consumption risk or employee turnover, depending on the diagnostic goal of the application.

The aim of this paper is to introduce psychologists to the basic concepts and methods of CSL. We want to demonstrate a selection of the most commonly used CSL methods on a psychological dataset to propagate the usefulness of CSL for psychological classification problems.¹ For our empirical demonstration, we use the openly available drug consumption dataset (Fehrmann et al., 2017) from the UCI Machine Learning Repository (Dua & Graff, 2019). This dataset contains records of usage of 18 legal and illegal drugs for 1,885 individuals.

Machine Learning in Psychological Science

Traditionally, psychological science has been relying on classical statistical methods such as linear regression models. Unfortunately, their performance in predicting new, unseen data is often poor for various reasons (Menton, 2020; Yarkoni & Westfall, 2017). This lack of generalizability often stems from the fact that models with a large number of predictors fitted to relatively small datasets tend to represent the sample data too closely to enable accurate predictions on new data. This tendency of a model to simply

¹ Our paper focuses solely on binary classification problems; for an introduction to multi-class CSL, see Zhou and Liu (2010).

“remember” the sample data is called *overfitting* (Menton, 2020). The basic approach to detect overfitting in ML is to split the sample data into two parts: a training set whose observations are used to train the algorithm and a test set whose observations are predicted to estimate the performance of the trained algorithm on new, unseen data (Yarkoni & Westfall, 2017). Since going into detail about the terminology and foundations of ML would be beyond the scope of this paper, we refer newcomers to Menton (2020), Orrù et al. (2020), or Yarkoni and Westfall (2017) and Pargent and Albert-von der Gönna (2018) for introductions more tailored to psychologists; for a special focus on personality research and assessment, see Stachl et al. (2020); for a special focus on clinical psychology, see Dwyer et al. (2018).

In psychology, the performance of classical statistical models is hardly ever evaluated on new data (Yarkoni & Westfall, 2017). Thus, reported performance estimates (e.g., in-sample R^2 for the linear models dominant in psychology) sometimes poorly quantify the ability to predict new data from the same population. However, because predicting future experiences and behaviors is the goal of many psychological applications, different methods are needed to estimate predictive performance. For this reason, ML algorithms are trained on one part of the sample and evaluated on another (Menton, 2020). The most common performance measure to evaluate classification models is the so-called accuracy (ACC; Chawla et al., 2002). The ACC of a model is the sum of correct predictions divided by all predictions (i.e., for binary classification, $\frac{\#TP + \#TN}{\#TP + \#FN + \#TN + \#FP}$; $\#TP$, $\#FN$, $\#TN$, $\#FP$ = number of true positives, false negative, true negative, false positive predictions). Most state of the art binary classification algorithms do not directly predict class membership but instead produce some probability that an observation belongs to each class. Non-CSL algorithms usually assume a probability threshold of 0.5. That is, an observation is predicted to belong to that class with estimated class probability greater than 0.5 (Elkan, 2001). However, when MC of FPs and FNs differ, the interpretation of ACC can be misleading because it does not provide information about the cost-optimality

of predictions. A prediction can be optimal from the probabilistic perspective but not from the cost perspective (Chawla et al., 2002; Viaene & Dedene, 2005).

This can be demonstrated with the famous *Receiver Operating Characteristic* (ROC) curve (Majnik & Bosnic, 2013; Maloof, 2003). ROC analysis is based on two additional performance measures: sensitivity (SENS; also true positive rate or recall; $\frac{\#TP}{\#TP+\#FN}$) and specificity (SPEC; also true negative rate $\frac{\#TN}{\#TN+\#FP}$). The ROC curve for a classification algorithm plots 1 - specificity on the x-axis against sensitivity on the y-axis, while varying the probability threshold used to make class predictions. Each point on the curve corresponds to a certain prediction threshold, ranging from an estimated probability of 1 on the left to 0 on the right. The ROC curve of an optimal classifier intersects the point (0,1), which would indicate that the algorithm classifies all positive observations as positive and no negative observation as positive (Viaene & Dedene, 2005). The line passing through the origin with $y = x$ is equivalent to guessing the class at random. ROC curves for different ML algorithms or CSL methods can be compared and point-wise error bounds can be used to visualize the variability of out-of-sample performance estimates. To compare different algorithms when MC are unknown, the area under the ROC curve (AUC) can be used as a performance measure which is independent of the decision criterion or any prior probabilities of classes. The AUC is interpreted as the probability that the positive observation is ranked higher than the negative observation when such a pair is chosen at random (Chawla et al., 2002; Viaene & Dedene, 2005). However, when there is information on different MC, the algorithm with the highest AUC can be suboptimal for the appropriate cost trade-off in the concrete application.

Figure 1 shows an exemplary ROC curve for a single classification algorithm. When ignoring different MC, the resulting classifier would achieve comparable values for SENS and SPEC when choosing the standard prediction threshold of 0.5 marked in the figure. However, depending on the application, high SENS could be valued higher and the SENS/SPEC trade-off resulting from the marked prediction threshold of 0.2 might be

favored. In this way, ROC can be used by practitioners to heuristically find a prediction threshold which will result in a trade-off which incorporates an appropriate implicit weighting of MC based on their expert knowledge (Rousson & Zumbunn, 2011). In our paper, we will present the reverse approach: specifying concrete MC first and then use different CSL strategies like automatically adjusting prediction thresholds to achieve cost-optimal predictive performance. When concrete MC are available, it is possible to define a natural cost-sensitive performance measure: For the Mean Misclassification Cost (MMC) we add the MC for each prediction and divide this sum by the total number of predictions. We will use MMC as our primary performance measure to evaluate the predictive performance of different CSL methods in our empirical demonstration.

The performance measures presented above all quantify the *outer loss* of a predictive model. They are simply definitions of when we colloquially consider some predictions as “good,” “precise,” or “accurate.” That is, we evaluate the performance of an algorithm to generalize on new data by its ability to optimize an appropriate external performance measure determined by the concrete application. This outer loss needs to be distinguished from the *inner loss* function which is internally optimized by the ML algorithm during training (Bischl & Moosbauer, 2019). As an example, logistic regression internally minimizes the so-called binomial loss (which we will discuss later). It would make intuitive sense if the outer and inner loss were equal. However, this is often not feasible because many ML algorithms rely on specific inner loss functions for efficient estimation of model parameters. Thus, although different MC could in theory be taken into account both in the inner and in the outer loss function, most algorithms cannot directly incorporate different MC when estimating model parameters.

Principles of Cost-Sensitive Learning

Before presenting the most common CSL methods, we give an introduction to the mathematical foundations of CSL. As mentioned earlier, most algorithms treat the two

Figure 1

Exemplary ROC curve; Sensitivity = True Positive Rate; 1 - Specificity = False Positive Rate; the diagonal is equivalent to guessing the prediction of classes; the grey line is the performance of the algorithm; the two points marked on the grey line correspond to thresholds of 0.5 and 0.2.

misclassification errors as if the associated costs are equal. Consequently, algorithms are trained to maximize ACC (using an implicit prediction threshold of 0.5) and minimize the amount of misclassifications. However, given unequal MC, it can be optimal to classify an observation as belonging to one class even though another class is more probable. A predictor vector $x \in \mathbb{R}^P$ should be classified as belonging to the class $i \in \{1, \dots, I\}$ associated with the minimum expected cost (Elkan, 2001; Viaene & Dedene, 2005). Posterior probabilities of a class given a predictor vector can be computed. The optimal prediction for an observation v with predictor values x_v is that class i among all I classes which has the lowest value

$$\sum_{j=1}^I P(j \mid x_v) C(i, j). \quad (1)$$

$P(j \mid x_v)$ is the conditional probability of class j estimated by the classifier for a given observation v with predictor vector x_v . $C(i, j)$ defines the cost of predicting class i when the actual class for observation v is j . Note the fundamental principle of CSL: we might act as if i is the true class even though class j is more probable (e.g., considering the MC, it can be optimal to classify a person as a drug user even though he or she most likely is not; Elkan, 2001).

The costs used in equation (1) are assumed to be given and can be contrasted in a cost matrix C . In C , the rows depict the predicted classes whereas the columns depict the true classes. Table 1 shows the special case of binary classification. In our manuscript, we use the convention to label the positive class with 1 and the negative class with 0. Thus, the entry c_{10} (c_{01}) denotes the cost of a FP (FN) prediction. In most cases, classifying an observation correctly has zero costs ($c_{00} = c_{11} = 0$). Elkan (2001) defines two reasonableness conditions: First, the cost of an incorrect prediction should always be higher than the cost of a correct prediction, that is $c_{10} > c_{00}$ and $c_{01} > c_{11}$. Second, no row m dominates row n (with $m \neq n$), that means there is no m for which $C(m, j) \geq C(n, j)$, for all j . If this was actually the case, it would be optimal to always predict n , regardless of the actual class j (Margineantu, 2000).

Table 1

Cost matrix C with all entries (i, j) and their associated costs

	actual positive	actual negative
predicted positive	$C(1, 1) = c_{11}$	$C(1, 0) = c_{10}$
predicted negative	$C(0, 1) = c_{01}$	$C(0, 0) = c_{00}$

Three important properties of the entries in C need to be mentioned: First, the costs

do not necessarily have to be monetary. For example, costs can also be the time that was spent or the severity of a disease (Elkan, 2001; Ting, 2002). Second, in the most frequent case of binary classification with assumed zero costs for correct classifications, not the absolute value of costs is important but only their ratio. Thus, one MC can be fixed as the value 1 and the other one can be set in relation (Domingos, 1999; Ting, 2002). Third, all costs of entries (i, j) in C are independent of observation v . That is, for every v classified as belonging to class i even though j is the correct class, the same costs apply. We refer readers interested in dealing with example-dependent costs to Zadrozny and Elkan (2001a).

Taking into account all of the above, we can use equation (1) to derive a rule for optimal decision making:

$$P(j = 0 \mid x_v)c_{10} + P(j = 1 \mid x_v)c_{11} \leq P(j = 0 \mid x_v)c_{00} + P(j = 1 \mid x_v)c_{01}. \quad (2)$$

That is, the best prediction is class 1 (i.e., the positive class) if the expected cost of this prediction (left side of equation (2)) is less or equal to the expected cost of the prediction of class 0 (right side of equation (2)) (Elkan, 2001).

Given $p = P(j = 1 \mid x_v)$, equation (2) is equivalent to

$$(1 - p)c_{10} + pc_{11} \leq (1 - p)c_{00} + pc_{01}. \quad (3)$$

If both sides of equation (3) are exactly equal, both class predictions would be optimal. We call the optimal threshold p^* , thus

$$(1 - p^*)c_{10} + p^*c_{11} = (1 - p^*)c_{00} + p^*c_{01}. \quad (4)$$

Solving equation (4) for p^* gives

$$p^* = \frac{c_{10} - c_{00}}{c_{10} - c_{00} + c_{01} - c_{11}}. \quad (5)$$

Putting everything together and keeping in mind the reasonableness conditions (so that the denominator of equation (5) is nonzero), the optimal prediction is class 1 if $p \geq p^*$ (Elkan, 2001).

Methods for Cost Sensitive Learning

In the Introduction, we explained that in applications with unequal MC, solely relying on ACC is not enough and performance measures like MMC should be considered to select an appropriate model. CSL usually goes one step further by trying to produce models with superior cost performance instead of only choosing the best among a set of cost-insensitive models. Before describing the most important methods, we present a new taxonomy of CSL in Table 2.

Table 2

Taxonomy of CSL methods

1. Cost-Sensitive Hyperparameter Tuning
2. Cost-Sensitive Meta-Learning
2.1 Adjusting the Prediction Threshold
Theoretical Threshold (Coussement, 2013)
Empirical Threshold (Sheng & Ling, 2006)
2.2 Adjusting the Training Data
Under-/Oversampling (Elkan, 2001)
SMOTE (Chawla et al., 2002; Fernández et al., 2018)
MetaCost (Domingos, 1999)
3. Cost-Sensitive Model Training
3.1 Algorithm Specific Weighting Function
Theoretical Weights (Zadrozny et al., 2003)
Empirical Weights (Zadrozny et al., 2003)
3.2 Custom Inner Loss-Functions (Wang et al., 2020)

For good predictive performance, advanced ML algorithms strongly rely on hyperparameters. A simple way for CSL when working with such algorithms is to tune

those hyperparameters based on MMC. However, this approach will not work well for algorithms which are relatively insensitive to the choice of hyperparameter values or do not have hyperparameters at all. Cost-sensitive meta-learning refers to a group of approaches that use model-agnostic *wrapper* methods to make any ML algorithm cost-sensitive. These meta-methods can be further divided into methods that adjust the threshold used to make concrete predictions based on estimated prediction probabilities and into methods that adjust the training data. While hyperparameter tuning and meta-methods both lead to cost-sensitive algorithms only indirectly, the most sophisticated approaches aim at directly incorporating classification costs into the training phase of the ML algorithm. This can be achieved by using the algorithm-specific weighting functions provided by most influential ML libraries or by directly implementing customized inner loss functions for libraries that enable this functionality.

Cost-Sensitive Hyperparameter Tuning

In contrast to regular model parameters like regression coefficients, hyperparameters are second-order parameters of ML algorithms. They cannot be automatically learned by the model on the basis of the training data and need to be set before training the algorithm. Hyperparameters often exert a lot of influence on the learning process and performance of the final model. Because well-performing default values are only available for some ML algorithms (e.g., random forests; Bernard et al., 2009; Probst et al., 2019), it is often recommendable to *tune* hyperparameters based on resampling methods like cross-validation (CV). For each training set, the algorithm is trained repeatedly, each time with a different set of hyperparameter values. Then, the algorithm is retrained on the combined training set with the combination of hyperparameter values that achieved the best predictive performance averaged across all test sets. When estimating the predictive performance of the final model resulting from such a tuning regime, it is crucial to apply *nested resampling* (Bischl et al., 2012) to avoid overoptimistic performance estimates. The

outer resampling loop is used to evaluate predictive performance (similar to the case without any hyperparameters) while an inner resampling loop is created in which the tuning procedure as described above is repeated within each of the outer training sets. A more detailed description of nested resampling can be found in Michael Lang et al. (2020) or Pargent and Albert-von der Gönna (2018). If ACC is used in the inner tuning loop, costs are not taken into account when selecting optimal hyperparameter values, even if the performance of the tuned algorithms is evaluated using MMC in the outer performance evaluation loop. However, it is easy to also apply the same cost measure in the inner resampling loop to prefer hyperparameter settings that ensure good performance when costs are taken into account. Note that the cost-sensitive hyperparameter tuning approach cannot only be used for hyperparameters of the ML algorithm but also for hyperparameters belonging to a cost-sensitive meta-learning method applied either before or after the cost-insensitive ML algorithm. When both algorithm and meta-learning hyperparameters are present, they can all be tuned at the same time. We will use these strategies in several of the CSL methods described below.

Cost-Sensitive Meta-Learning

Some authors divide CSL into two categories: On the one hand, algorithms exist that are in themselves cost-sensitive and naturally take into account different MC when classifying new data (Turney, 1995). However, transforming each ML algorithm into a cost-sensitive version can be time-consuming, requires a lot of technical knowledge, and a rewrite of software (Domingos, 1999). This is why researchers have developed so called *wrapper* methods that either modify the predictions or the training data of a classification algorithm to consider different MC. This form of CSL is also referred to as cost-sensitive meta-learning (Ling & Sheng, 2008; Zadrozny et al., 2003).

The following two meta-learning strategies work by adjusting the prediction threshold used to transform probability estimates into concrete predictions. They can be

used in combination with most cost-insensitive ML algorithms.

Theoretical Threshold. We already noted that most state of the art classification algorithms produce estimates of $P(j \mid x_v)$ for an observation v based on its predictor values x_v . Non-CSL algorithms try to maximize ACC, usually by implicitly assuming a probability threshold of 0.5. That is, if $P(j = 1 \mid x_v) \geq 0.5$ in binary classification, the prediction is class 1, otherwise it is class 0 (Elkan, 2001). CSL methods base this translation of estimated probabilities into class predictions on a threshold p^* that usually differs from 0.5. The optimal class of an observation can be directly determined by making use of the criterion defined in equation (1). That is, each observation is classified as belonging to the class that yields the lowest expected cost according to the Bayes optimal prediction, even though the other class might be more probable. Obviously, a reliable estimation of the class probabilities $P(j \mid x_v)$ for each observation v is important.

If a classifier does yield class probabilities, we can draw on the decision rule introduced in equation (5) and classify an observation v as belonging to the positive class if

$$P(j = 1 \mid x_v) > \frac{c_{10} - c_{00}}{c_{10} - c_{00} + c_{01} - c_{11}}. \quad (6)$$

Note that in our case in which costs are example-independent (i.e., every misclassification to a certain class has a fixed cost, independent of v), equation (6) specifies a stable classification threshold (Viaene & Dedene, 2005). This approach is known as Theoretical Thresholding (TT) or Direct Minimum Expected Cost Classification (DMECC; Viaene & Dedene, 2005).

Empirical Threshold. While almost all modern ML algorithms do produce class probability estimates $P(j \mid x_v)$, those estimates can be biased and badly calibrated (Van Calster et al., 2019; Zadrozny & Elkan, 2001a). Strongly biased estimates can lead to a non-optimal TT (Domingos, 1999; Ling & Sheng, 2008).

Sheng and Ling (2006) developed a method that also relies on posterior probability estimates, but instead of the exact estimate only the correct order of probabilities is required. This method is called *empirical thresholding (ET)*. ET finds the best prediction threshold for an algorithm with regard to the MC based on the training data. Mathematically, the MC become a function of the threshold p : $MC = f(p)$. By calculating the MC for every possible p (i.e., all probabilities between 0 and 1), the optimal threshold p^* can be determined by choosing the global minimum of the resulting curve. It is important to note that using the complete training data to find the optimal threshold might lead to overfitting. Thus, ET determines the optimal threshold via k -fold CV by predicting the probability estimates on the inner test sets. In this way the threshold p^* is simply treated as a hyperparameter similar to the approach described in the earlier section on cost-sensitive tuning.

The following three cost-sensitive meta learning strategies work by adjusting the training data and can thus be used in combination with any cost-insensitive ML algorithm.

Under- and Oversampling. An alternative to TT and ET is to implicitly achieve the optimal threshold by altering the frequency of the positive and negative class in the training data (in our example, the distribution of drug users and non-users).

The intuition is that depending on the cost ratio, we try to force the ML algorithm to place a higher focus on one of the two classes, by increasing the ratio of the more important class in the training data. This can be achieved by either removing observations from the less important class (Undersampling; US: discard observations at random) or duplicating observations from the more important class (Oversampling; OS: draw observations with replacement).

In the frequent case of binary classification with an internal probability threshold of the cost-insensitive ML algorithm of 0.5 and $c_{00} = c_{11} = 0$, the class adjustment ratio is simply $\frac{c_{01}}{c_{10}}$ (Elkan, 2001). OS requires to increase the size of the positive class by a factor of

$\frac{c_{01}}{c_{10}}$ (if $\frac{c_{01}}{c_{10}} > 1$) or to increase the size of the negative class by a factor of $\frac{c_{10}}{c_{01}}$ (if $\frac{c_{01}}{c_{10}} < 1$). US requires to decrease the size of the negative class by a factor of $\frac{c_{10}}{c_{01}}$ (if $\frac{c_{01}}{c_{10}} > 1$) or to decrease the size of the positive class by a factor of $\frac{c_{01}}{c_{10}}$ (if $\frac{c_{01}}{c_{10}} < 1$). All mathematical details can be found in Elkan (2001).

It is not obvious whether US or OS should be preferred: With US, valuable data is discarded (Viaene & Dedene, 2005; Weiss et al., 2007). Every psychologist who has ever experienced the effort of data collection will surely flinch just thinking about throwing data away. Additionally, with sample sizes in psychology often being smaller than what is recommended for proper sizes of training and test data (Orrù et al., 2020), US appears to be a less favorable option than OS. However, because OS adds data to the existing sample, this increases the computation time needed to train the algorithm. Moreover, OS might lead to overfitting because observations are duplicated (Domingos, 1999; Seiffert et al., 2008).

Synthetic Minority Oversampling Technique. US and OS just discard or duplicate data at random. To counteract some disadvantages, Chawla et al. (2002) introduced a more sophisticated way of duplicating data, called *Synthetic Minority Oversampling TEchnique* (SMOTE). Instead of randomly duplicating observations from the training data, SMOTE adds synthetic “new” data. By creating synthetic data close to already existing observations from the more important class, algorithms are provided with more observations to be trained with.

The steps needed to synthetically oversample observations are performed by a k -nearest neighbor approach. Each observation from the more important class is connected to its k -nearest neighbors from the same class via a line segment. When plotting only the predictor values, a line segment would be the straight path between the observation and the next closest one. Synthetic observations are created along these line segments, either on all or on a specific number of segments, depending on the desired amount of

oversampling. For example, Chawla et al. (2002) set $k = 5$ and show that for an oversampling of the positive class by 200%, two of the five neighbors are chosen and one synthetic observation is created on each of the two lines connecting the original positive observation to its two chosen neighbors. The original SMOTE paper (Chawla et al., 2002) has initiated a whole literature of more complex variants (for a review, see Fernández et al., 2018). To keep it simple, we will only apply the original and most widely applied version in our empirical demonstration. Note that the SMOTE implementation in R only allows OS for the smaller class (i.e. the smaller class must have the higher MC).

MetaCost. Domingos (1999) pioneered CSL by introducing *MetaCost*, the first real cost-sensitive method beyond US and OS. The initial goal was to create a method that converted existing algorithms into CSL ones. As with US and OS, MetaCost is a CSL method that adjusts the training data, in this case the target label for some training observations. The core assumption of MetaCost is that given a certain cost structure, we do not know the optimal class of an observation, not even in the training data. Thus, the training data should be relabeled so that each observation in the training data has the label (i.e., is classified into the class) with the lowest conditional risk according to the MC. Even if the new label differs from the original one that was obtained when gathering the data, it would still be changed. Only after the process of relabeling, the classifier should be trained and used to predict the unadjusted target labels in the test data. We have already introduced the formula to calculate the conditional risk in equation (1):

$$R(i \mid x_v) = \sum_j P(j \mid x_v) C(i, j) \quad (7)$$

More specifically, the conditional risk $R(i \mid x_v)$ is the expected cost of classifying observation v as belonging to class i . Just as in equation (1), the goal is to find the Bayes optimal prediction by minimizing equation (7). In our example, we would first ignore whether a person did or did not use drugs and calculate the “new” class for each person

according to equation (7).

The crucial step in cost-optimally relabeling the training observations is estimating the conditional class probabilities $P(j | x)$. Although most algorithms produce these estimations in the training phase, their performance in doing so can be poor (Zadrozny & Elkan, 2001a, 2001b). For this case –and because many older ML algorithms could not produce class probabilities directly– MetaCost draws on an ensemble learning method called *bagging* (bootstrap aggregation; Breiman, 1996). In detail, let s be the size of the training data. A bootstrap sample is drawn m times by picking s observations with replacement from the training data, and the ML algorithm is trained to each dataset. The resulting ensemble aggregates the results across the m models. In case the models cannot predict class probabilities, the ensemble classifies an observation as belonging to the class that was predicted by most of the m models (so called *majority voting*). If class probabilities can be estimated in the training phase, the aggregation of results is the unweighted average of the class’s probability given a specific observation. Now that we have obtained estimates of the class probabilities, we can use equation (7) to relabel the observations in the training data as belonging to the estimated cost-optimal class (i.e., minimizing equation (7) for each observation v). The relabeled training data can then be used to train the classifier.

Cost-Sensitive Model Training

The remaining CSL methods will focus on directly estimating model parameters (e.g., regression coefficients) in a cost-sensitive way by somehow adjusting the inner loss function which is used to train the ML algorithm. These methods are more complex and we need exemplary equations to illustrate the main concepts.

As mentioned in the introduction, logistic regression minimizes the so-called binomial (or cross entropy) loss. The regression coefficients are estimated so that the following expression returns the smallest possible value

$$-\sum_{v=1}^N y_v \ln(\hat{p}_v) + (1 - y_v) \ln(1 - \hat{p}_v) \quad (8)$$

on the training set, where $\hat{p}_v$ is the estimated probability $P(y_v = 1|x_v)$ that observation v belongs to the positive class. The sum is the log-likelihood of a logistic regression for which high values are preferred: for observations from the positive class ($y_v = 1$), we want $\hat{p}_v$ to be close to 1; for observations from the negative class ($y_v = 0$), we want $\hat{p}_v$ to be close to 0. Because it is tradition to minimize the loss (instead of maximizing the likelihood), we simply reverse the sign by defining the binomial loss as the negative log-likelihood. The binomial loss is symmetric, thus, it is cost-insensitive and does not favor any of the two classes. If observation v belongs to the positive class ($y_v = 1$), the term $\ln(\hat{p}_v)$ is added to the total loss; otherwise ($y_v = 0$) the term $\ln(1 - \hat{p}_v)$ is added. This mechanism can be used for CSL by deliberately adjusting the influence of loss components resulting from false positive and false negative predictions. The binomial loss function is also often used in advanced ML algorithms like gradient boosting. While formula (8) stays the same, the only difference between algorithms is how the estimated probabilities $\hat{p}_v$ depend on the structure of the model parameters, i.e. how to compute predictions for the respective ML algorithm.

Theoretical Weights. The *theoretical weighting* (TW) method is closely related to US and OS with the convenient advantage that the distribution of training data is not altered. Hence, no data are discarded or duplicated. The factor with which training observations would be multiplied when performing OS is now assigned as a weight to each of these observations (Elkan, 2001). As a result, the algorithm puts a premium on the class with a high weight/cost ratio during training and learns to avoid errors associated with high MC.

Conveniently, most common packages for ML algorithms (e.g., logistic regressions, random forests, and gradient boosting) support observation weights w_v that influence the

impact of observation v during the training process (Michael Lang et al., 2020). How weighting is achieved mostly depends on the inner loss function employed by the respective ML algorithm. For our example of binomial loss (e.g., in logistic regression or gradient boosting), the implementation could be

$$-\sum_{v=1}^N w_v \cdot \{y_v \ln(\hat{p}_v) + (1 - y_v) \ln(1 - \hat{p}_v)\} \quad (9)$$

which reduces to equation (8) in the case of $w_v = 1$, *for all* v . Imagine $w_z = 5$ for some observation z . This would result in the same total loss as if we duplicated observation z five times. For CSL with example-dependent costs (Zadrozny & Elkan, 2001a) w_v could be different for each v . However, in most applications we are only interested in class-specific costs. When giving the same weight to all observations that belong to the same class, we can simplify the above formula by introducing the class weight w_{pos} :

$$-\sum_{v=1}^N w_{pos} \cdot y_v \ln(\hat{p}_v) + (1 - y_v) \ln(1 - \hat{p}_v) \quad (10)$$

With $w_{pos} > 1$, we would punish false negatives more strongly than false positives. Again, using $w_{pos} = 5$ should be comparable to OS of the positive class by a factor of 5 (Elkan, 2001).

Empirical Weights. Empirical weighting (EW) is similar in concept to ET: the class weight parameter w_{pos} is treated as a hyperparameter and tuned by CV.

Custom Inner Loss Functions. In addition to the algorithm-specific TW approach provided by most popular ML libraries, some packages enable the user to forward a custom inner loss function which is directly used to train the algorithm (Wang et al., 2020). While this feature gives ultimate freedom to use cost-sensitive loss functions designed for a specific application, the approach also comes with disadvantages: First, deriving and correctly implementing a valid loss function requires a deep mathematical

understanding of the respective ML algorithm which cannot be expected from most psychologists. Second, running the algorithm with a custom loss function is usually much slower than with the computationally optimized weighting scheme provided by the package developer. For example, Wang et al. (2020) implemented two custom loss functions for gradient boosting with the xgboost package: *weighted binomial loss* which is the same as equation (9) and *focal loss* (which we will not discuss here). What is not mentioned in their paper is that weighted binomial loss is mathematically identical to the weighting function already provided by the xgboost package, which is much faster and can be used without additional programming. Other work on custom objective functions exists (e.g., on cost-sensitive support vector machines; P. Cao et al., 2013; Iranmehr et al., 2019) but, again, extensive mathematical knowledge is required to understand let alone build these algorithms. For these reasons, we decided against using custom inner loss functions in our empirical demonstration and only demonstrate weighting.

Which CSL Method Should You Use?

Having presented the most common CSL methods, the question naturally arises which one performs best. Unfortunately, theoretical and empirical work suggests that no single method is always superior. For US and OS, various studies showed that while US can perform well (especially for higher cost ratios between FP and FN misclassifications), random OS almost never seems to be a preferable choice (Domingos, 1999; Seiffert et al., 2008; Weiss et al., 2007). In fact, due to the aforementioned disadvantages of US and OS, Coussement (2013) did not even include them in his benchmark experiment. Zhou and Liu (2005) compared the two methods (and SMOTE) to thresholding and found the latter to outperform the sampling methods for two-class classification tasks. In contrast, Weiss et al. (2007) mentioned that US and OS might be the only options due to a lack of direct cost-sensitive methods. This might have been true at that time. However, many algorithms now have direct cost-sensitive implementations (see, e.g., C. Cao & Wang, 2018; Cheng et

al., 2017; Li et al., 2015) and additional CSL methods are also available. As a matter of fact, non-sampling methods (e.g., thresholding, weighting) consistently perform well with differing “rankings” between studies. Coussement (2013) found that, for low cost ratios, ET and MetaCost are the best choices whereas TW is preferable for higher cost ratios. Another study, however, showed that ET and TW perform equally well over various and increasing cost ratios (Seiffert et al., 2008). SMOTE achieved a similar performance, especially at lower cost ratios. Finally, Sheng and Ling (2006) conducted a benchmark experiment and came to the conclusion that ET and TW are on par with each other, both outperforming MetaCost. Additionally, ET was the more robust method for high cost ratios.

Existing benchmark studies demonstrate that the performance of different CSL methods depend on at least the following factors: (1) the cost ratio between FP and FN misclassifications; (2) the ML algorithms with which the CSL methods are used (most benchmark studies used decision trees or logistic regression); (3) the frequency of positive and negative observations; (4) additional dataset characteristics. Thus, it is extremely difficult to transfer any results from CSL benchmark studies to your own application. Benchmark studies can never test all possible cost ratios which always strongly depend on each individual application. Moreover, there is only a comparably small number of openly available datasets appropriate for large-scale CSL benchmarks because differing MC typically arise in applications with strong individual or institutional data protection concerns (e.g., medical diagnoses, credit default). Thus, most benchmark studies rather focus on classification problems with imbalanced data instead of cost ratios that are derived from theoretical considerations. It is not clear whether the ranking of different CSL methods in these benchmarks can be transferred one to one to cases where different MC are investigated. As a consequence, researchers often have no choice but to compare several CSL methods for their own application, instead of relying on existing benchmarks. With this demonstration of applying a selection of CSL methods to a single openly available dataset using a predefined exemplary cost ratio, we want to assist psychological researchers

in constructing a similar comparison for their own application. Our paper deliberately is not an additional benchmark study and the results on our demonstration dataset should not be generalized to other applications.

Tell Me Who You Are And I Tell You What You Smoke

After we gave an overview of CSL methods we now apply them to an exemplary dataset and discuss their potential usefulness for the psychological domain. Even though our goal is not to build a concrete model that can be applied in psychological practice, we shortly introduce our data example:

The relation between personality traits and illicit substance use has been well-researched for many decades. Sieber (1981) was one of the first to investigate substance use over longer periods of time. More specifically, he found that in young men illicit drug use was positively related to high values in *Depressiveness* and negatively related to high values in *Dominance* and *Excitability*. Decades later, recent studies still support the notion that personality traits are related to drug use. In a longitudinal study, Allen and Laborde (2020) showed that high values in *Neuroticism*, *Extraversion*, and *Openness*, as well as low values in *Agreeableness* and *Conscientiousness* are associated with subsequent illicit substance use. In addition to the Big Five of personality, *Impulsivity* and *Sensation-Seeking* were found to influence substance use (Huertas et al., 2019; Ruiz-Olivares et al., 2019) and might even share a genetic base with it (Skóra et al., 2020).

Thanks to Fehrmann et al. (2017) making their dataset on drug consumption and personality publicly available on the UCI Machine Learning Repository (Dua & Graff, 2019), another body of research has emerged over the last few years: Many studies have been conducted with the goal of predicting drug consumption on the basis of the Big Five as well as Impulsivity and Sensation Seeking (Kumari et al., 2018; Liu, 2019; Nath et al., 2017; Qiao et al., 2019). In these studies, different kinds of ML algorithms (e.g., random

forests or artificial neural networks) were used to predict the usage of various licit or illicit drugs. Moreover, the duration of substance use (Adinugroho et al., 2019) or even the opposite, that is, personality profiles on the basis of drug consumption history, were predicted (Ferwerda & Tkalcic, 2020). The majority of these studies showed that substance use can be predicted on the basis of personality with satisfying accuracy. In fact, the features related to the aforementioned personality traits proved to be highly important for the prediction of substance use in all of these studies.

However, to the best of our knowledge, no study so far has taken into account different MC when training algorithms or interpreting results. Admittedly, to validly determine the different MC in the area of drug use prediction, domain expertise (which we do not have) would be needed. From a mathematical perspective only the ratio of costs is important, not their actual values (Domingos, 1999; Ting, 2002). Thus, it is common practice in benchmark studies on CSL methods to set the less costly misclassification error to 1 and adjust the value for the costlier error accordingly (Coussement, 2013; Weiss et al., 2007).

We decided to set the cost ratio between a FP (i.e., predicting that a person will use drugs when the person actually will not) and a FN misclassification (i.e., predicting that a person will not use drugs when the person actually will) to 1:5. That is, we assess the MC of a FN prediction five times as high as the MC of a FP. We want to briefly explain our rationale behind this decision: By no means do we want to imply that a FP error in this case is not associated with bad consequences (i.e., costs). Multiple studies showed that drug usage and addiction are highly stigmatized in the general public (Barry et al., 2014; Janulis et al., 2013; McCradden et al., 2019). This might manifest itself in social avoidance of people with drug addiction (Corrigan et al., 2009) or the perception in the general public that people who suffer from drug addiction are the most dangerous among all people with mental illnesses (Parcesepe & Cabassa, 2013). The stigmatization, in turn, might result in a self-stigmatization by the people affected, leading to feelings of shame and social

withdrawal (Matthews et al., 2017; McCradden et al., 2019). By setting our FP:FN cost ratio to 1:5 we do not want to downplay these possible negative consequences of a FP prediction. One could argue, though, that a FN is associated with higher costs than a FP, especially since the predicted variable is drug *use*, not *addiction*. Substance use is a major global health problem and its prevention is a key component of facing this issue (Allen & Laborde, 2020; Huertas et al., 2019). Especially the immediate medical consequences potentially arising from only one-time use can be detrimental and their prevention should be a main goal (C.-Y. Chen & Lin, 2009; Devlin & Henry, 2008). Thus, we assume that it would be worse to “miss” a future drug user (FN) than to “unnecessarily” take a second look at a non-user (FP).

Method

Data

The drug consumption dataset was collected and made publicly available by Fehrmann et al. (2017) on the UCI Machine Learning Repository (Dua & Graff, 2019). The data were gathered between March 2011 and March 2012 and comprise drug usage records of $N = 1885$ individuals ($n_{female} = 942, n_{male} = 943$). Age ranges in years were 18-24 ($n = 643$; 34.11%), 25-34 ($n = 481$; 25.52%), 35-44 ($n = 356$; 18.89%), 45-54 ($n = 294$; 15.60%), 55-64 ($n = 93$; 4.93%), and 65+ ($n = 18$; 0.95%).

Participants reported their drug usage history concerning 18 different licit and illicit drugs: alcohol, amphetamines, amyl nitrite, benzodiazepines, cannabis, chocolate, cocaine, caffeine, crack, ecstasy, heroin, ketamine, legal highs, LSD, methadone, mushrooms, nicotine, volatile substance abuse (VSA) and one fictitious drug, Semeron, which was included to identify participants who over-claim drug use. Participants indicated on a 7-point Likert-scale whether they have “never used” the specific drug, used it “over a decade ago,” or “within the last decade,” “year,” “month,” “week,” or “day.” Apart from

demographic variables and drug usage history, personality traits were gathered to be used as features for drug usage prediction. The *Revised NEO Five-Factor Inventory* (NEO-FFI-R; Costa & McCrae, 1992) was used to assess the Big Five factors of personality: Neuroticism, Extraversion, Openness, Conscientiousness, and Agreeableness. Each dimension was assessed using 12 items on a 5-point Likert-scale, resulting in a total of 60 items. The *Barratt Impulsiveness Scale* (BIS-11; Stanford et al., 2009) was employed to evaluate participants' Impulsivity. The BIS-11 contains 30 self-report items on a 4-point Likert-scale. It is divided into the three subscales *motor impulsiveness*, *attentional impulsiveness*, and *non-planning*, with ten items per subscale. Finally, participants filled out the *Impulsiveness Sensation-Seeking* (ImpSS) scale to assess Sensation-Seeking on the trait level. The scale comprises 19 items on a true-false response format of which eight items measure Impulsivity and 11 items measure Sensation-Seeking. Despite the name and subscales, the ImpSS is regarded as a tool to assess general Sensation-Seeking (Zuckerman, 1994). Unfortunately, Fehrman et al. (2017) only provided participants' aggregated scores on these scales, preventing us from including all items as separate predictors (Pargent & Albert-von der Gönna, 2018).

Conveniently, Fehrman et al. (2017) transformed all features into continuous variables. That is, they quantified all ordinal (e.g., the answers on the Likert-scales) and all nominal (e.g., gender or ethnicity) features so that ML algorithms can be easily applied. The mathematical approach of this input feature transformation is described in Fehrman et al. (2017). As features we used "Age," "Gender," "Education," "Country," "Ethnicity," "Neuroticism," "Extraversion," "Openness," "Agreeableness," "Conscientiousness," "Impulsivity," and "Sensation Seeking." Inquiring drug use history from "never used" to "used within the last day" allows for a more detailed investigation of drug usage. However, we dichotomized the sample into non-users ("never used" and "used over a decade ago") and users ("used within the last decade" to "used within the last day"). This way we were able to work on a binary classification problem. According to Fehrman et al. (2017),

participants who used a drug over a decade ago cannot be validly considered to be a drug user. Hence, they recommend researchers interested in solving binary classification problems to split the data like we just described. Furthermore, we only included the drugs of the *heroin pleiad* (cf. Fehrmann et al., 2017) in our analyses. That is, we only paid attention to and, thus, predicted the usage history of crack, cocaine, heroin, or methadone. The reasons for this were twofold: First, if we included all 18 drugs, the distribution between users and non-users would have been highly skewed ($n = 467$ non-users and $n = 1,418$ users). By only including drugs of the heroin pleiad, we obtained a more balanced distribution, yielding $n = 1053$ (55.86%) participants who indicated to have never used any of the drugs (or to have used it over a decade ago) and $n = 832$ (44.14%) participants who indicated to have used at least one of the drugs (at least once within the last decade). Second, since we wanted to investigate costly misclassifications, it would not have been reasonable to include legal drugs like chocolate. The non-medical use of all drugs of the heroin pleiad was illegal in the country the sample was gathered and even one-time use could lead to detrimental health consequences (C.-Y. Chen & Lin, 2009).

Algorithms

We applied three different algorithms for classification: logistic regression (LR), random forest (RF), and gradient boosting (GB). This provided a more comprehensive overview with regard to the performance of different CSL methods on algorithms with varying complexity.

Logistic Regression

LR is a classical statistical method for binary classification. It models the probability of observation v belonging to the positive class ($y_v = 1$) given predictor values $x_{1v}, \dots, x_{kv}$. Note that if we used a regular linear regression, our model could output

probabilities smaller than 0 or greater than 1. To ensure that the range of possible values for this probability lies within (0,1), a logistic function is used (James et al., 2021):

$$p_v = P(y_v = 1|x_v) = \frac{e^{\beta_0 + \beta_1 x_{v1} + \dots + \beta_k x_{vk}}}{1 + e^{\beta_0 + \beta_1 x_{v1} + \dots + \beta_k x_{vk}}} \quad (11)$$

Estimated probabilities $\hat{p}_v$ can then be converted into class memberships. Usually, the following criterion applies: if $\hat{p}_v \geq 0.5$, v is classified as belonging to the positive class, otherwise as belonging to the negative class (Elkan, 2001). Because the coefficients $\beta_0, \beta_1, \dots, \beta_k$ are unknown, they must be estimated using the training data. We seek parameter estimates so that $\hat{p}_v$ is near 0 for individuals who indicated not to have used drugs and near 1 for individuals who indicated to have used drugs. The most commonly used method to fit the model is *maximum likelihood estimation* (James et al., 2021) which is equivalent to minimizing the binomial loss described earlier. LR does not have any hyperparameters and thus does not require additional tuning.

Random Forest

A RF is an ensemble ML algorithm based on decision trees. A bootstrap sample is drawn *num.trees* times from the training data and a deep decision tree is built on each of the bootstrap samples. At every node of a tree only a random subset of *mtry* predictor variables is being considered for splitting (Breiman, 2001). For each tree, $\hat{p}_v$ is equal to the relative frequency of the positive class in the end node of that tree in which observation v belongs. The RF algorithm then averages $\hat{p}_v$ over all trees to stabilize the predictions. By only considering a reduced number of predictors at every node, the different trees are less correlated. Hence, averaging their results leads to a more effective reduction of variance (James et al., 2021). We used the efficient *ranger* algorithm (Wright & Ziegler, 2017) to train RFs. Although several hyperparameters like *num.trees* and *mtry* could be tuned, the RF has been shown to perform very well with default values (Probst et al., 2019) and is

often considered the best off-the-shelf algorithm (Fernández-Delgado et al., 2014). Thus, we did not use hyperparameter tuning and instead used the default values $num.trees = 500$ and $mtry = \sqrt{P}$ (with P being the total number of predictor variables).

Gradient Boosting

GB is a more complex ensemble ML method also based on decision trees (James et al., 2021). In contrast to the RF in which decision trees are independently grown on bootstrap samples and later combined into a final prediction, GB sequentially grows shallow decision trees without drawing bootstrap samples. Each additional tree is fitted to the residuals from a model including all previous trees (Friedman et al., 2000). Thus, observations that were difficult to classify correctly receive a higher implicit weight in the next tree. In contrast to RF, GB often requires excessive tuning to achieve good predictive performance, which increases both complexity and computation time. For example, when increasing the size of the ensemble unchecked, GB will inevitably overfit to the training data due to its aggressive loss minimization (Bischl & Moosbauer, 2019). GB implementations greatly differ between packages. We used the xgboost algorithm (T. Chen & Guestrin, 2016) and tuned the following subset of hyperparameters (we used default values for all others): (1) The number of trees ($nrounds$) in the ensemble on a range from 1 to 500, (2) the learning rate (η) on a range from 0.001 to 0.01, (3) the maximum depth of each tree (max_depth) on a range from 1 to 10, (4) the regularization parameter (λ) on a range from 0.001 to 0.1, and (5) the proportion of randomly sampled predictor variables within each tree ($colsample_bytree$) on a range from 0.5 to 1.

Benchmark Design

We compared different algorithms and CSL methods on the drug consumption dataset with a cost ratio of 1:5 between FP and FN predictions. The CSL methods TT, ET, OS, SMOTE, MetaCost, TW, and EW were applied in combination with all three ML

algorithms. For gradient boosting we also added a separate condition using only cost sensitive hyperparameter tuning. CSL methods were used with the following specifications: The threshold in TT was set to $\frac{1}{6}$ (according to equation (6), we obtain a theoretically cost-optimal threshold of $p^* = \frac{c_{10}-c_{00}}{c_{10}-c_{00}+c_{01}-c_{11}} = \frac{1-0}{1-0+5-0} = \frac{1}{6}$). The threshold in ET was tuned on a range between 0.01 and 0.99. The oversampling rate in OS was set to 5 ($\frac{c_{01}}{c_{10}} = 5$). In SMOTE, we set the number of neighbors to 5 and altered the data so that in the training sets of CV, there were five times as many users as in the regular dataset. The theoretical threshold from TT was used for MetaCost. We implemented the simple version of MetaCost where for each observation, all models of the ensemble are used for relabeling (in contrast to only using models whose training data did not contain the observation to be relabeled; cf. Domingos, 1999). We always used RF as the algorithm within the MetaCost ensemble with an ensemble size of 50 to reduce complexity. The class weights in TW were set to 5 and 1. The class weights in EW were tuned on a range between 1 and 50. MMC was used as inner performance measure for ET, EW, and for cost-sensitive hyperparameter tuning with GB. In conditions with GB, hyperparameters of ML algorithm and CSL method were tuned jointly. When tuning hyperparameters, 5-fold CV was used in the inner loop of nested CV to compare predictive performance of 1000 randomly sampled (combinations of) hyperparameter values (i.e. random search; Bergstra & Bengio, 2012). To evaluate the usefulness of CSL, our benchmark also contained non-CSL versions of all three ML algorithms and a featureless baseline model. In the non-CSL version of GB, ACC was used as performance measure when tuning hyperparameters. The featureless baseline model did not use any predictor variables but randomly predicted each class with a probability equal to the relative frequency of that class in the training set. In the outer loop of our nested resampling scheme, the predictive performance of all algorithms was estimated via 3-times repeated 5-fold CV. We calculated MMC, SENS, SPEC, ACC, and AUC averaged across test sets with the mean. For each performance measure, we also calculated the standard deviation across all test sets. Note that for the CSL methods TT

and ET no AUC can be computed because the prediction threshold is fixed by the method.

Software and Code

All analyses were conducted in the open-source software *R* (R Core Team, 2020). The paper was written with the *papaja* package (Aust & Barth, 2020). The *data.table* package (Dowle & Srinivasan, 2021) was used for data pre-processing. All basic ML operations like training algorithms and predictions were conducted with the *mlr3* (Michel Lang et al., 2019) and its companion packages (see Michel Lang & Schratz, 2021). CSL methods were employed using *mlr3pipelines* (Binder et al., 2020). RFs were fitted by the *ranger* package (Wright & Ziegler, 2017) and GB was performed with the *xgboost* package (T. Chen & Guestrin, 2016). The benchmark analysis was run on the RStudio Server cluster of the Leibniz Supercomputing Centre (LRZ). The entire code including some additional analyses and tutorials can be found in our OSF repository (<https://osf.io/cvks7/>).

Results

The predictive performance of all CSL methods for LR in contrast to the non-CSL condition on the drug consumption dataset is shown in Table 3. Results for GB and RF are displayed in Tables 4 and 5, respectively. The same overall pattern emerged for each ML algorithm: While the non-CSL condition had the best ACC, all CSL methods had better cost-sensitive performance. MMC was roughly halved when using CSL methods. Our selected cost ratio was supposed to penalize FNs more strongly than FPs. Accordingly, all CSL methods shifted towards higher SENS and lower SPEC compared to the non-CSL condition. No differences in AUC could be observed between CSL methods. Considering the standard deviation across test sets, cost-sensitive performance based on the MMC was highly similar for all CSL methods with each ML algorithm. However, for each ML algorithm there was always at least one CSL method with inferior cost-sensitive

Table 3*Performance of different CSL methods with Logistic Regression*

Method	MMC	ACC	SENS	SPEC	AUC
Non-CSL	0.94(0.13)	0.71(0.03)	0.63(0.04)	0.77(0.03)	0.78(0.03)
TT	0.49(0.02)	0.56(0.03)	0.97(0.01)	0.24(0.03)	NA(NA)
ET	0.47(0.03)	0.62(0.03)	0.95(0.02)	0.37(0.06)	NA(NA)
OS	0.83(0.09)	0.71(0.03)	0.69(0.03)	0.72(0.03)	0.78(0.03)
SMOTE	0.47(0.03)	0.61(0.03)	0.96(0.01)	0.33(0.03)	0.78(0.03)
MetaCost	0.48(0.02)	0.62(0.02)	0.95(0.01)	0.36(0.03)	0.78(0.03)
TW	0.48(0.03)	0.59(0.03)	0.96(0.01)	0.3(0.03)	0.78(0.03)
EW	0.46(0.03)	0.62(0.02)	0.95(0.01)	0.35(0.03)	0.78(0.03)

Note. Cost ratio between false positive and false negative = 1:5. MMC = Mean misclassification costs. ACC = Accuracy. SENS = Sensitivity = True positive rate. SPEC = Specificity = True negative rate. AUC = Area Under the ROC Curve. TT = Theoretical thresholding. ET = Empirical thresholding. OS = Oversampling. SMOTE = Synthetic minority oversampling technique. TW = Theoretical weights. EW = Empirical weights. For TT and ET, no AUC can be calculated. Values in parentheses are the standard deviations across all test sets.

performance: OS showed inferior performance with all ML algorithms. SMOTE had inferior performance with RF. Cost-sensitive hyperparameter tuning with GB performed much worse compared to the other CSL methods and more similar to the non-CSL condition. While the non-CSL condition of all ML algorithms had highly similar ACC, they differed with respect to MMC. Cost-sensitive performance of the well performing CSL methods was highly similar for all ML algorithms on this dataset.

Discussion

CSL in Practice

In this paper we described the most common CSL methods and applied them in combination with different ML algorithms on a psychological dataset classifying drug consumption based on psychological variables like personality scores. Psychologists regularly make complex diagnostic decisions. Recently, psychology started implementing

Table 4*Performance of different CSL methods with Gradient Boosting*

Method	MMC	ACC	SENS	SPEC	AUC
Non-CSL	0.88(0.15)	0.7(0.03)	0.67(0.05)	0.73(0.02)	0.78(0.03)
CS_Hyper	0.86(0.11)	0.71(0.02)	0.68(0.04)	0.73(0.03)	0.77(0.02)
TT	0.56(0.02)	0.44(0.02)	1(0)	0(0)	NA(NA)
ET	0.48(0.03)	0.61(0.04)	0.95(0.02)	0.34(0.07)	NA(NA)
OS	0.63(0.11)	0.68(0.03)	0.83(0.06)	0.57(0.07)	0.76(0.03)
SMOTE	0.48(0.03)	0.62(0.03)	0.94(0.01)	0.37(0.06)	0.77(0.03)
MetaCost	0.49(0.03)	0.61(0.02)	0.95(0.01)	0.35(0.04)	0.76(0.03)
TW	0.47(0.03)	0.62(0.03)	0.95(0.02)	0.37(0.07)	0.78(0.02)
EW	0.48(0.04)	0.61(0.02)	0.95(0.02)	0.34(0.04)	0.78(0.02)

Note. Cost ratio between false positive and false negative = 1:5. MMC = Mean misclassification costs. ACC = Accuracy. SENS = Sensitivity = True positive rate. SPEC = Specificity = True negative rate. AUC = Area Under the ROC Curve. TT = Theoretical thresholding. ET = Empirical thresholding. OS = Oversampling. SMOTE = Synthetic minority oversampling technique. TW = Theoretical weights. EW = Empirical weights. For TT and ET, no AUC can be calculated. Values in parentheses are the standard deviations across all test sets.

ML algorithms to assist in the diagnostic process; for example, in suicide prevention (Walsh et al., 2017), employee performance and turnover prediction (Sajjadi et al., 2019), or offender recidivism (Hamilton et al., 2015). In all of these cases, equal MC are implausible. Thus we are convinced that CSL can be of great value for psychological assessment. Had we not taken into account the differing MC in our empirical example, more than 30% of drug users would have remained undetected. In our diagnostic use case, this would be highly insufficient. All CSL methods led to better cost-sensitive performance than non-CSL methods. While most CSL methods performed equally well, single methods performed noticeably worse with certain ML algorithms on our dataset. GB, which is more complex to train and requires more computational resources due to hyperparameter tuning, did not outperform the untuned RF or even simple LR.

Our empirical demonstration showcases several common findings when working with

Table 5*Performance of different CSL methods with Random Forest*

Method	MMC	ACC	SENS	SPEC	AUC
Non-CSL	0.84(0.13)	0.71(0.03)	0.69(0.04)	0.73(0.03)	0.78(0.02)
TT	0.48(0.02)	0.58(0.02)	0.96(0.01)	0.29(0.03)	NA(NA)
ET	0.48(0.03)	0.62(0.04)	0.95(0.03)	0.36(0.08)	NA(NA)
OS	0.85(0.12)	0.71(0.03)	0.69(0.05)	0.72(0.03)	0.78(0.02)
SMOTE	0.61(0.06)	0.7(0.01)	0.83(0.03)	0.6(0.03)	0.77(0.02)
MetaCost	0.48(0.03)	0.61(0.02)	0.95(0.01)	0.35(0.03)	0.77(0.02)
TW	0.47(0.03)	0.66(0.02)	0.92(0.02)	0.46(0.03)	0.78(0.02)
EW	0.46(0.03)	0.63(0.02)	0.95(0.02)	0.39(0.05)	0.78(0.02)

Note. Cost ratio between false positive and false negative = 1:5. MMC = Mean misclassification costs. ACC = Accuracy. SENS = Sensitivity = True positive rate. SPEC = Specificity = True negative rate. AUC = Area Under the ROC Curve. TT = Theoretical thresholding. ET = Empirical thresholding. OS = Oversampling. SMOTE = Synthetic minority oversampling technique. TW = Theoretical weights. EW = Empirical weights. For TT and ET, no AUC can be calculated. Values in parentheses are the standard deviations across all test sets.

cost-sensitive prediction tasks: As presented in the Introduction, benchmark studies have repeatedly shown that no CSL method is always superior on all datasets and with all ML algorithms (Coussement, 2013; Seiffert et al., 2008; Sheng & Ling, 2006). At the same time, no ML algorithm is always superior. More complex algorithms do not always have an advantage over simple ones like logistic LR (Hamilton et al., 2015; Sanderson et al., 2020). We cannot stress enough how important it is not to generalize our empirical findings to other combinations of datasets $\times$ CSL methods $\times$ ML algorithms! As a result, practitioners that want to use CSL on their own dataset have no choice but to perform a small benchmark experiment to compare some of the most promising ML algorithms and CSL methods. Our online materials showcase how this can be done with the `mlr3` framework in R. In addition to the code used to compute the empirical example in our paper, we also provide a reduced version of our benchmark which is fast enough to be run on a simple laptop in a matter of minutes (<https://osf.io/cvks7/>). This should help readers

to familiarize themselves with running their own CSL benchmarks.

When comparing CSL methods, it is important to use appropriate performance evaluation based on resampling. While modern ML frameworks like `mlr3` make this as simple as possible and are built in a way that makes it difficult to make common mistakes, we sometimes encounter ML applications with invalid performance estimates. We briefly note two common errors which lead to a grave overestimation of predictive performance: 1) CSL methods that alter the training data like OS, SMOTE, or MetaCost are applied to the complete dataset and predictive performance is evaluated on altered data instead of using unaltered test sets. 2) Cost-sensitive hyperparameter tuning or tuning hyperparameters of a CSL method is performed on the whole dataset when evaluating predictive performance instead of repeating the tuning process in each resampling iteration. To make researchers more aware of these problems, our online materials contain a demonstration of making the above mistakes when analyzing the drug consumption dataset (<https://osf.io/cvks7/>).

How to Choose Misclassification Costs?

In our empirical example, all CSL methods decreased the MMC compared to the non-CSL version that did not take different MC into account. However, they always did so by sacrificing ACC. That is, more observations were classified incorrectly, at the same time, MMC decreased. This is not surprising because the very principle of CSL is to act as if one class is the true class of an observation even though the other class might be more probable. We saw that the trade-off between accuracy and MMC also reflects in a trade-off between SENS and SPEC. Since we put a premium on the classification of the positive class (i.e., the drug users), the algorithms more aggressively classified observations as users. Thus, SENS was increased by all CSL methods while SPEC was reduced.

These principles demonstrate that whether CSL methods produce appropriate predictions crucially depends on the selected MC. If the cost matrix is inappropriate for

the concrete application, CSL methods might easily yield worse predictions than the non-CSL version. So how do we find appropriate MC in practice? Unfortunately, there is no general answer for this question which strongly depends on each concrete application: In some cases, MC directly depend on the monetary costs of making a wrong decision. However, in most psychological applications, costs refer to psychological dimensions like decreased psychological well-being when receiving the wrong treatment or the anxiety from being faced with a stigmatizing diagnosis that might not be accurate. In these contexts, MC are often obtained either from the experiences of professional experts (e.g., therapists), or from people directly affected by the predictive model (e.g., patients). In many applications, costs might consist of a trade-off of both monetary and psychological aspects. While a detailed discussion on how to approach these rather complicated issues goes beyond the scope of this article, we want to suggest a few directions: While psychology has mostly neglected the topic of determining decision costs in a standardized way, other fields have worked with more formal frameworks for decades. For example, health economics deals with the problem of deciding which one of several health interventions should be funded by a public health system, when considering the monetary cost of an intervention, the probability of the intervention to achieve certain health states, as well as the expected duration and subjective preferences of those health states (Brazier et al., 2017; Drummond et al., 2015; Hunter et al., 2015). Standardized questionnaires like the EQ-5D (Stolk et al., 2019) have been developed, in which patients in medical trials rate their current health state on multiple health dimensions. Such questionnaires allow to quantify the preferences of all possible health states on a unidimensional scale based on normative value sets derived from a representative sample. To obtain such unidimensional preference scales, several techniques are used to elicit preferences for certain health states: The most common techniques are visual-analogue scales (simply rating a health state on a visual analogue scale ranging from worst to best imaginable health; Parkin & Devlin, 2006), Standard Gamble (weighing a health state against a gamble which could result in the best

possible health with probability p but could also result in death with probability $1 - p$; Gafni, 1994), Time Trade-Off (weighing living in a health state for t years against living in perfect health for $k < t$ years and then dying; Lugnér & Krabbe, 2020), and Discrete Choice Experiments (computing a ranking from multiple binary decisions between health states which differ with respect to their rating on multiple health dimensions; Bahrampour et al., 2020). We think that psychology could greatly profit from learning about these techniques (and perhaps avoid some of their weaknesses) which should in many circumstances promise advantages over cost ratios determined purely subjectively by some researcher developing a prediction model (like we did in our empirical example).

Limitations and Future Research

The predictive models built in this study are not readily applicable in practice. We did not assess the costs of misclassification errors based on clinical domain expertise but instead used a single ad-hoc cost ratio for our generic empirical example. Although our paper strives to give a thorough introduction on how psychologists can approach classification tasks in which expert knowledge on different MC are available, we did not cover other topics which are of utmost importance when building prediction models that can be responsibly applied in research and practice. The Covid-19 pandemic has produced many cautionary examples of clinical prediction models which are unusable in practice or even pose severe risk of negative consequences to human health and well-being (Roberts et al., 2021; Wynants et al., 2020). We appeal to our colleagues in psychology not to repeat mistakes by prematurely building or applying predictive models for high stakes situation without adhering to all scientific quality standards.

CSL methods are also often applied when dealing with an imbalanced dataset (Japkowicz & Stephen, 2002; Maloof, 2003; Nguyen et al., 2010). A dataset is imbalanced when one class is over- while the other class is underrepresented. Often times, the true class with the highest cost when classified incorrectly is also the underrepresented one,

which has the consequence that the algorithm has less observations of the more important class to learn from. Hence, a cost-insensitive algorithm will mostly predict the other class in order to maximize accuracy (McCarthy et al., 2005). We specifically chose the target variable in a way that the observations of users and non-users are almost balanced. In this, we sought to focus on the general usefulness of CSL methods for handling different MC. However, future research should investigate how CSL methods can aid us in the related case of imbalanced psychological datasets.

When using ML for assessment or clinical decision-making, properly estimated probabilities are crucial to avoid adverse consequences. This is even more important when MC are considered and CSL methods are applied. Hence, practitioners should not focus solely on common performance measures (e.g., ACC) and MC, but also on how well the applied algorithms are calibrated (e.g., Van Calster et al., 2019), i.e., how close the estimated probabilities for a specific class resemble the respective frequencies in the data. TRIPOD (Transparent Reporting of a multivariable prediction model for Individual Prognosis Or Diagnosis; Moons et al., 2015) and PROBAST (Prediction model Risk Of Bias Assessment Tool; Wolff et al., 2019) provide useful tools and resources for researchers and practitioners who seek to develop applicable prediction models.

Conclusion

The goal of this paper was to introduce psychologists to the principles of cost-sensitive machine learning. We presented the mathematical foundations of CSL and introduced a taxonomy for its most commonly used methods, namely cost-sensitive hyperparameter tuning, meta-learning, and model training. We applied these methods to a psychological dataset, predicting drug consumption based on personality characteristics, thereby demonstrating their application and usefulness in the psychological domain. CSL methods were combined with three ML algorithms of different complexity.

Because predicting future behavior is one of the main goals of psychology, supervised ML can be highly beneficial for our field. To further increase the potential of ML, CSL methods are needed that incorporate unequal misclassification costs. Although all CSL methods decreased MMC compared to a non-CSL algorithm in our example, no CSL method proved superior, which is in line with larger benchmark studies. Thus, it is inevitable for researchers to conduct small benchmarks on their own data. For this, our paper and online materials provide guidance in the form of reproducible and adjustable R code. We hope that future research will further investigate the usefulness of CSL in more psychological applications. This should contribute to an important discussion between researchers and practitioners on how to determine concrete misclassification costs, which are the basis for all CSL methods.

References

- Adinugroho, S., Sari, Y. A., & Hidayat, N. (2019). Drug usage duration classification using extreme learning machine based on personality features. *International Conference on Sustainable Information Engineering and Technology*, 33–37.
- Allen, M. S., & Laborde, S. (2020). A prospective study of personality and illicit drug use in Australian adults. *Personality and Individual Differences*, 163. <https://doi.org/10.1016/j.paid.2020.110048>
- Aust, F., & Barth, M. (2020). *papaja: Create APA manuscripts with R Markdown*. <https://github.com/crsh/papaja>
- Bahrampour, M., Byrnes, J., Norman, R., Scuffham, P. A., & Downes, M. (2020). Discrete choice experiments to generate utility values for multi-attribute utility instruments: A systematic review of methods. *The European Journal of Health Economics*, 21(7), 983–992. <https://doi.org/https://doi.org/10.1007/s10198-020-01189-6>
- Barry, C. L., McGinty, E. E., Pescosolido, B. A., & Goldman, H. H. (2014). Stigma, discrimination, treatment effectiveness, and policy: Public views about drug addiction and mental illness. *Psychiatric Services*, 65(10), 1269–1272. <https://doi.org/doi.org/10.1176/appi.ps.201400140>
- Bergstra, J., & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13, 281–305.
- Bernard, S., Heutte, L., & Adam, S. (2009). Influence of hyperparameters on random forest accuracy. In J. A. Benediktsson, J. Kittler, & F. Roli (Eds.), *Multiple classifier systems* (pp. 171–180). Springer Berlin Heidelberg.
- Binder, M., Pfisterer, F., Schneider, L., Bischl, B., Lang, M., & Dandl, S. (2020). *mlr3pipelines: Preprocessing operators and pipelines for 'mlr3'*. <https://CRAN.R-project.org/package=mlr3pipelines>

- Bischl, B., Mersmann, O., Trautmann, H., & Weihs, C. (2012). Resampling methods for meta-model validation with recommendations for evolutionary computation. *Evolutionary Computation*, 20(2), 249–275.
- Bischl, B., & Moosbauer, J. (2019). Predictive modelling. *Lecture Slides, Summerterm 2019, LMU Munich*.
- Brazier, J., Ratcliffe, J., Saloman, J., & Tsuchiya, A. (2017). *Measuring and valuing health benefits for economic evaluation*. OXFORD university press.
- Breiman, L. (1996). Bagging predictors. *Machine Learning*, 24, 123–140.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
<https://doi.org/10.1023/A:1010933404324>
- Cao, C., & Wang, Z. (2018). IMCStacking: Cost-sensitive stacking learning with feature inverse mapping for imbalanced problems. *Knowledge-Based Systems*, 150, 27–37. <https://doi.org/10.1016/j.knosys.2018.02.031>
- Cao, P., Zhao, D., & Zaiane, O. (2013). An optimized cost-sensitive SVM for imbalanced data learning. *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 280–292.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., & Kegelmeyer, W. P. (2002). SMOTE: Synthetic minority over-sampling technique. *Journal of Artificial Intelligence Research*, 16, 321–357.
- Chen, C.-Y., & Lin, K.-M. (2009). Health consequences of illegal drug use. *Current Opinion in Psychiatry*, 22, 287–292.
<https://doi.org/10.1097/YCO.0b013e32832a2349>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- Cheng, F., Zhang, J., Wen, C., Liu, Z., & Li, Z. (2017). Large cost-sensitive margin distribution machine for imbalanced data classification. *Neurocomputing*, 224,

- 45–57. <https://doi.org/10.1016/j.neucom.2016.10.053>
- Corrigan, P. W., Kuwabara, S. A., & O’Shaughnessy, J. (2009). The public stigma of mental illness and drug addiction. *Journal of Social Work*, 9(2), 139–147. <https://doi.org/10.1177/1468017308101818>
- Costa, P. T., & McCrae, R. R. (1992). Revised NEO personality inventory (NEO-PI-R) and NEO five-factor inventory (NEO-FFI): Professional manual. *Psychological Assessment Resources*.
- Coussement, K. (2013). Improving customer retention management through cost-sensitive learning. *European Journal of Marketing*, 48(4), 477–495. <https://doi.org/10.1108/EJM-03-2012-0180>
- Devlin, R. J., & Henry, J. A. (2008). Clinical review: Major consequences of illicit drug consumption. *Critical Care*, 12(1), 202–206. <https://doi.org/10.1186/cc6166>
- Domingos, P. (1999). MetaCost: A general method for making classifiers cost-sensitive. *Proceedings of the Fifth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 155–164.
- Dowle, M., & Srinivasan, A. (2021). *Data.table: Extension of ‘data.frame’*. <https://CRAN.R-project.org/package=data.table>
- Drummond, M. F., Sculpher, M. J., Claxton, K., Stoddart, G. L., & Torrance, G. W. (2015). *Methods for the economic evaluation of health care programmes*. Oxford University Press.
- Dua, D., & Graff, C. (2019). UCI Machine Learning Repository. In *University of California, Irvine, School of Information and Computer Sciences*. <http://archive.ics.uci.edu/ml>
- Dwyer, D. B., Falkai, P., & Koutsouleris, N. (2018). Machine learning approaches for clinical psychology and psychiatry. *Annual Review of Clinical Psychology*, 14, 91–118. <https://doi.org/10.1146/annurev-clinpsy-032816-%20045037>

- Elkan, C. (2001). The foundations of cost-sensitive learning. *International Joint Conference on Artificial Intelligence*, 17(1), 973–978.
- Fehrmann, E., Mirkes, E. M., Muhammad, A. K., Egan, V., & Gorban, A. N. (2017). *The Five Factor Model of personality and evaluation of drug consumption risk*. <https://arxiv.org/abs/1506.06297>
- Fernàndez, A., Garcìa, S., Herrera, F., & Chawla, N. V. (2018). SMOTE for learning from imbalanced data: Progress and challenges, marking the 15-year anniversary. *J. Artif. Intell. Res.*, 61, 863–905.
<https://doi.org/10.1613/jair.1.11192>
- Fernández-Delgado, M., Cernadas, E., Barro, S., & Amorim, D. (2014). Do we need hundreds of classifiers to solve real world classification problems? *Journal of Machine Learning Research*, 15, 3133–3181.
<http://jmlr.org/papers/v15/delgado14a.html>
- Ferwerda, B., & Tkalcic, M. (2020). Exploring the prediction of personality traits from drug consumption profiles. *Adjunct Publication of the 28th ACM Conference on User Modeling, Adaptation and Personalization*, 2–5.
<https://doi.org/10.1145/3386392.3397589>
- Freitas, A. (2011). Building cost-sensitive decision trees for medical applications. *AI Communications*, 24, 285–287. <https://doi.org/10.3233/AIC-2011-0490>
- Friedman, J., Hastie, T., & Tibshirani, R. (2000). Additive logistic regression: A statistical view of boosting (with discussion and a rejoinder by the authors). *Annals of Statistics*, 28(2), 337–407.
- Gafni, A. (1994). The standard gamble method: What is being measured and how it is interpreted. *Health Services Research*, 29(2), 207.
- Hamilton, Z., Neuilly, M.-A., Lee, S., & Barnoski, R. (2015). Isolating modeling effects in offender risk assessment. *Journal of Experimental Criminology*, 11(2), 299–318. <https://doi.org/10.1007/s11292-014-9221-8>

- Huertas, E., López-Moreno, J. A., Fernández, V., Echeverry-Alzate, V., & Bühler, K.-M. (2019). Associations between experimental substance use, FAAH-gene variations, impulsivity and sensation seeking. *Psicothema*, *31*(3), 239–245.
<https://doi.org/10.7334/psicothema2019.27>
- Hunter, R. M., Baio, G., Butt, T., Morris, S., Round, J., & Freemantle, N. (2015). An educational review of the statistical issues in analysing utility data for cost-utility analysis. *Pharmacoeconomics*, *33*(4), 355–366.
<https://doi.org/https://doi.org/10.1007/s40273-014-0247-6>
- Iranmehr, A., Masnadi-Shirazi, H., & Vasconcelos, N. (2019). Cost-sensitive support vector machines. *Neurocomputing*, *343*, 50–64.
- James, G., Witten, D., Hastie, T., & Tibshirani, R. (2021). *An Introduction to Statistical Learning: With Applications in R*. New York: Springer.
- Janulis, P., Ferrari, J. R., & Fowler, P. (2013). Understanding public stigma toward substance dependence. *Journal of Applied Social Psychology*, *43*, 1065–1072.
<https://doi.org/10.1111/jasp.12070>
- Japkowicz, N., & Stephen, S. (2002). The class imbalance problem: A systematic study. *Intelligent Data Analysis*, *6*, 429–449.
- Kumari, D., Kilam, S., Nath, P., & Swetapadma, A. (2018). Prediction of alcohol abused individuals using artificial neural network. *International Journal of Information Technology*, *10*(2), 233–237.
<https://doi.org/10.1007/s41870-018-0094-3>
- Lang, Michel, Binder, M., Richter, J., Schratz, P., Pfisterer, F., Coors, S., Au, Q., Casalicchio, G., Kotthoff, L., & Bischl, B. (2019). mlr3: A modern object-oriented machine learning framework in R. *Journal of Open Source Software*. <https://doi.org/10.21105/joss.01903>
- Lang, Michel, & Schratz, P. (2021). *mlr3verse: Easily install and load the 'mlr3' package family*. <https://CRAN.R-project.org/package=mlr3verse>

- Lang, Michael, Schratz, P., Binder, M., Pfisterer, F., Richter, J., Reich, N. G., & Bischl, B. (2020). *mlr3 book*. <https://mlr3book.mlr-org.com>
- Li, X., Zhao, H., & Zhu, W. (2015). A cost sensitive decision tree algorithm with two adaptive mechanisms. *Knowledge-Based Systems*, 88, 24–33. <https://doi.org/10.1016/j.knosys.2015.08.012>
- Ling, C. X., & Sheng, V. S. (2008). Cost-sensitive learning and the class imbalance problem. In *Encyclopedia of Machine Learning*. Springer.
- Liu, E. M. (2019). Predicting illicit drug use with artificial neural network. *European Journal of Humanities and Social Sciences*, 3, 131–137.
- Lu, H., Xu, Y., Ye, M., Yan, K., Gao, Z., & Jin, Q. (2019). Learning misclassification costs for imbalanced classification on gene expression data. *BMC Bioinformatics*, 20, 681–690. <https://doi.org/10.1186/s12859-019-3255-x>
- Lugnér, A. K., & Krabbe, P. F. (2020). An overview of the time trade-off method: Concept, foundation, and the evaluation of distorting factors in putting a value on health. *Expert Review of Pharmacoeconomics & Outcomes Research*, 20(4), 331–342. <https://doi.org/10.1080/14737167.2020.1779062>
- Majnik, M., & Bosnic, Z. (2013). ROC analysis of classifiers in machine learning: A survey. *Intelligent Data Analysis*, 17, 531–558. <https://doi.org/10.3233/IDA-130592>
- Maloof, M. A. (2003). Learning when data sets are imbalanced and when costs are unequal and unknown. *ICML Workshop on Learning from Imbalanced Data Sets II*.
- Margineantu, D. D. (2000). On class-probability estimates and cost-sensitive evaluation of classifiers. *Workshop Notes, Workshop on Cost-Sensitive Learning, International Conference on Machine Learning*.
- Matthews, S., Dwyer, R., & Snoek, A. (2017). Stigma and self-stigma in addiction. *Bioethical Inquiry*, 14, 275–286. <https://doi.org/10.1007/s11673-017-9784-y>

- McCarthy, K., Zabar, B., & Weiss, G. M. (2005). Does cost-sensitive learning beat sampling for classifying rare classes? *Proceedings of the 1st International Workshop on Utility-Based Data Mining*, 69–77.
- McCradden, M. D., Vasileva, D., Orchanian-Cheff, A., & Buchman, D. Z. (2019). Ambiguous identities of drugs and people: A scoping review of opioid-related stigma. *International Journal of Drug Policy*, 74, 205–215.
<https://doi.org/10.1016/j.drugpo.2019.10.005>
- Menton, W. H. (2020). Generalizability of statistical prediction from psychological assessment data: An investigation with the MMPI-2-RF. *Psychological Assessment*, 32(5), 473–492. <https://doi.org/10.1037/pas0000808>
- Moons, K. G., Altman, D. G., Reitsma, J. B., Ioannidis, J. P., Macaskill, P., Steyerberg, E. W., Vickers, A. J., Ransohoff, D. F., & Collins, G. S. (2015). Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): Explanation and elaboration. *Annals of Internal Medicine*, 162(1), W1–W73.
- Nath, P., Kilam, S., & Swetapadma, A. (2017). A machine learning approach to predict volatile substance abuse for drug risk analysis. *Third International Conference on Research in Computational Intelligence and Communication Networks*, 255–258.
- Nguyen, T.-N., Gantner, Z., & Schmidt-Thieme, L. (2010). Cost-sensitive learning methods for imbalanced data. *The 2010 International Joint Conference on Neural Networks*, 1–8.
- Orrù, G., Monaro, M., Conversano, C., Gemignani, A., & Sartori, G. (2020). Machine learning in psychometrics and psychological science. *Frontiers in Psychology*, 10, 2970. <https://doi.org/10.3389/fpsyg.2019.02970>
- Parcesepe, A. M., & Cabassa, L. J. (2013). Public stigma of mental illness in the United States: A systematic literature review. *Administration and Policy in*

- Mental Health and Mental Health Services Research*, 40(5), 384–399.
<https://doi.org/10.1007/s10488-012-0430-z>
- Pargent, F., & Albert-von der Gönna, J. (2018). Predictive modeling with psychological panel data. *Zeitschrift Für Psychologie*, 226(4), 246–258.
<https://doi.org/10.1027/2151-2604/a000343>
- Parkin, D., & Devlin, N. (2006). Is there a case for using visual analogue scale valuations in cost-utility analysis? *Health Economics*, 15(7), 653–664.
<https://doi.org/https://doi.org/10.1002/hec.1086>
- Probst, P., Wright, M. N., & Boulesteix, A.-L. (2019). Hyperparameters and tuning strategies for random forest. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 0(0), e1301. <https://doi.org/10.1002/widm.1301>
- Qiao, Z., Chai, T., Zhang, Q., Zhou, X., & Chu, Z. (2019). Predicting potential drug abusers using machine learning techniques. *International Conference on Intelligent Informatics and Biomedical Sciences*, 283–286.
- R Core Team. (2020). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing. <https://www.R-project.org/>
- Roberts, M., Driggs, D., Thorpe, M., Gilbey, J., Yeung, M., Ursprung, S., Aviles-Rivero, A. I., Etmann, C., McCague, C., Beer, L.others. (2021). Common pitfalls and recommendations for using machine learning to detect and prognosticate for COVID-19 using chest radiographs and CT scans. *Nature Machine Intelligence*, 3(3), 199–217.
<https://doi.org/10.1038/s42256-021-00307-0>
- Rousson, V., & Zumbrunn, T. (2011). Decision curve analysis revisited: Overall net benefit, relationships to ROC curve analysis, and application to case-control studies. *BMC Medical Informatics Decis. Mak.*, 11, 45.
<https://doi.org/10.1186/1472-6947-11-45>
- Ruiz-Olivares, R., Lucena, V., Raya, A. F., & Herruzo, J. (2019). Personality

- profiles and how they relate to drug consumption among young people in Spain. *Personality and Individual Differences*, 149, 291–295.
- Sajjadiani, S., Sojourner, A. J., Kammeyer-Mueller, J. D., & Mykerezzi, E. (2019). Using machine learning to translate applicant work history into predictors of performance and turnover. *Journal of Applied Psychology*, 104(10), 1207–1225. <https://doi.org/10.1037/apl0000405>
- Sanderson, M., Bulloch, A. G. M., Wang, J., Williamson, T., & Patten, S. B. (2020). Predicting death by suicide using administrative health care system data: Can recurrent neural network, one-dimensional convolutional neural network, and gradient boosted trees models improve prediction performance? *Journal of Affective Disorders*, 264, 107–114.
- Seiffert, C., Khoshgoftaar, T. M., Van Hulse, J., & Napolitano, A. (2008). A comparative study of data sampling and cost sensitive learning. *IEEE International Conference on Data Mining Workshops*, 46–52. <https://doi.org/10.1109/ICDMW.2008.119>
- Sheng, V. S., & Ling, C. X. (2006). Thresholding for making classifiers cost-sensitive. *AAAI*, 476–481.
- Shmueli, G. (2010). To Explain or to Predict? *Statistical Science*, 25(3), 289–310. <https://doi.org/10.1214/10-STS330>
- Sieber, M. F. (1981). Personality scores and licit and illicit substance use. *Personality and Individual Differences*, 2, 235–241.
- Siers, M. J., & Islam, M. Z. (2018). Novel algorithms for cost-sensitive classification and knowledge discovery in class imbalanced datasets with an application to NASA software defects. *Information Sciences*, 459, 53–70. <https://doi.org/10.1016/j.ins.2018.05.035>
- Skóra, M. N., Pattij, T., Beroun, A., Kogias, G., Mielenz, D., Vries, T. de, Radwanska, K., & Müller, C. P. (2020). Personality driven alcohol and drug

- abuse: New mechanisms revealed. *Neuroscience and Biobehavioral Reviews*, 116, 64–73. <https://doi.org/10.1016/j.neubiorev.2020.06.023>
- Stachl, C., Pargent, F., Hilbert, S., Harari, G. M., Schoedel, R., Vaid, S., Gosling, S. D., & Bühner, M. (2020). Personality research and assessment in the era of machine learning. *European Journal of Personality*. <https://doi.org/10.1002/per.2257>
- Stanford, M. S., Mathias, C. W., Dougherty, D. M., Lake, S. L., Anderson, N. E., & Patton, J. H. (2009). Fifty years of the Barratt Impulsiveness Scale: An update and review. *Personality and Individual Differences*, 47(5), 385–395.
- Stolk, E., Ludwig, K., Rand, K., Hout, B. van, & Ramos-Goñi, J. M. (2019). Overview, update, and lessons learned from the international EQ-5D-5L valuation work: Version 2 of the EQ-5D-5L valuation protocol. *Value in Health*, 22(1), 23–30. <https://doi.org/https://doi.org/10.1016/j.jval.2018.05.010>
- Ting, K. M. (2002). An instance-weighting method to induce cost-sensitive trees. *IEEE Transactions on Knowledge and Data Engineering*, 14(3), 659–665.
- Turney, P. D. (1995). Cost-sensitive classification: Empirical evaluation of a hybrid genetic decision tree induction algorithm. *Journal of Artificial Intelligence Research*, 2, 369–409.
- Van Calster, B., McLernon, D. J., Van Smeden, M., Wynants, L., & Steyerberg, E. W. (2019). Calibration: The achilles heel of predictive analytics. *BMC Medicine*, 17(1), 1–7. <https://doi.org/10.1186/s12916-019-1466-7>
- Viaene, S., & Dedene, G. (2005). Cost-sensitive learning and decision making revisited. *European Journal of Operational Research*, 166, 212–220. <https://doi.org/10.1016/j.ejor.2004.03.031>
- Walsh, C. G., Ribeiro, J. D., & Franklin, J. C. (2017). Predicting risk of suicide attempts over time through machine learning. *Clinical Psychological Science*, 5(3), 457–469. <https://doi.org/10.1177/2167702617691560>

- Wang, C., Deng, C., & Wang, S. (2020). Imbalance-XGBoost: Leveraging weighted and focal losses for binary label-imbalanced classification with XGBoost. *Pattern Recognition Letters*, 136, 190–197.
<https://doi.org/https://doi.org/10.1016/j.patrec.2020.05.035>
- Weiss, G. M., McCarthy, K., & Zabar, B. (2007). Cost-sensitive learning vs. Sampling: Which is best for handling unbalanced classes with unequal error costs? *Proceedings of the 2007 International Conference on Data Mining*.
- Wolff, R. F., Moons, K. G., Riley, R. D., Whiting, P. F., Westwood, M., Collins, G. S., Reitsma, J. B., Kleijnen, J., & Mallett, S. (2019). PROBAST: A tool to assess the risk of bias and applicability of prediction model studies. *Annals of Internal Medicine*, 170(1), 51–58.
- Wright, M. N., & Ziegler, A. (2017). ranger: A fast implementation of random forests for high dimensional data in C++ and R. *Journal of Statistical Software*, 77(1), 1–17. <https://doi.org/10.18637/jss.v077.i01>
- Wynants, L., Van Calster, B., Collins, G. S., Riley, R. D., Heinze, G., Schuit, E., Bonten, M. M. J., Dahly, D. L., Damen, J. A., Debray, T. P. A., Jong, V. M. T. de, De Vos, M., Dhiman, P., Haller, M. C., Harhay, M. O., Henckaerts, L., Heus, P., Kammer, M., Kreuzberger, N., . . . Smeden, M. van. (2020). Prediction models for diagnosis and prognosis of covid-19: Systematic review and critical appraisal. *BMJ*, 369. <https://doi.org/10.1136/bmj.m1328>
- Yarkoni, T., & Westfall, J. (2017). Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6), 1100–1122. <https://doi.org/10.1177/1745691617693393>
- Zadrozny, B., & Elkan, C. (2001a). Learning and making decisions when costs and probabilities are both unknown. *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 204–213.
- Zadrozny, B., & Elkan, C. (2001b). Obtaining calibrated probability estimates from

- decision trees and naive Bayesian classifiers. *Eighteenth International Conference on Machine Learning*.
- Zadrozny, B., Langford, J., & Abe, N. (2003). Cost-sensitive learning by cost-proportionate example weighting. *Proceedings of the Third IEEE International Conference on Data Mining*, 435–442.
- Zhou, Z.-H., & Liu, X.-Y. (2005). Training cost-sensitive neural networks with methods addressing the class imbalance problem. *IEEE Transactions on Knowledge and Data Engineering*, 18(1), 63–77.
- Zhou, Z.-H., & Liu, X.-Y. (2010). On multi-class cost-sensitive learning. *Computational Intelligence*, 26(3), 232–257.
- Zuckerman, M. (1994). *Behavioral expressions and biosocial bases of sensation seeking*. New York: Cambridge University Press.