

No Need for Ad-hoc Substitutes: The Expected Cost is a Principled All-purpose Classification Metric

Anonymous authors

Paper under double-blind review

Abstract

The expected cost (EC) is one of the main classification metrics introduced in statistical and machine learning books. It is based on the assumption that, for a given application of interest, each decision made by the system has a corresponding cost which depends on the true class of the sample. An evaluation metric can then be defined by taking the expectation of the cost over the data. Two special cases of the EC are widely used in the machine learning literature: the error rate (one minus the accuracy) and the balanced error rate (one minus the balanced accuracy or unweighted average recall). Other instances of the EC can be useful for applications in which some types of errors are more severe than others, or when the prior probabilities of the classes differ between the evaluation data and the use-case scenario. Surprisingly, the more general form for the EC is rarely used in the machine learning literature. Instead, alternative ad-hoc metrics like the F-beta score and the Matthews correlation coefficient (MCC) are used for many applications. In this work, we argue that the EC is superior to these alternative metrics, being more general, interpretable, and adaptable to any application scenario. We provide both theoretically-motivated discussions as well as examples to illustrate the behavior of the different metrics.

1 Introduction

During the development of machine learning systems, the model architecture, hyperparameters, training approaches, and other system characteristics, are determined based on the performance measured on a validation set. After development, the performance of the resulting system is reported using a separate evaluation set. In both of these steps it is essential that the metric used to evaluate performance be reflective of the needs of the application of interest. In this work, we focus on the study and comparison of metrics that are designed to evaluate classification systems. While systems trained for classification can sometimes be used for other purposes, like ranking of samples or producing an interpretable score of each class, the metrics needed to evaluate such systems are out of the scope of this work. Our goal here is to discuss metrics specifically designed for the evaluation of the quality of classification systems that output a single categorical decision for each input sample based only on the information extracted from that sample¹.

One of the most widely used metrics for the evaluation of categorical decisions is the total error rate or, equivalently, the accuracy, which is equal to one minus the total error rate. These metrics assume that all errors are equally costly regardless of the true class of the sample. In some applications, though, like those involving medical (Ashby & Smith, 2000; Kornak & Lu, 2011), biometric (Van Leeuwen & Brümmer, 2007; Gonzalez-Rodriguez, 2014), or business and economic decisions (Harsanyi, 1978; Berger, 2013), the costs may be extremely different across error types. In those cases, alternative ad-hoc metrics like the F-beta score (Rijsbergen, 1979) or the Matthews correlation coefficient (MCC) (Matthews, 1975), are often used for evaluation. In this work, we argue that the expected cost (EC), a decades-old generalization of the error rate (Savage, 1972), can be successfully used in all classification scenarios where categorical decisions are

¹Ranking systems that eventually produce categorical decisions by labelling the top M samples as being from the target class are also outside of the scope of this paper, since the resulting label for each sample is not based only on that sample but on the scores for other samples under evaluation.

made and where the severity of an error depends only on the true class of the sample and the decision made². Other metrics designed for this scenario, like the F-beta score or the MCC can be seen as ad-hoc substitutes for the EC, offering no advantages and various disadvantages over this metric.

The EC is based on a simple general assumption: given a certain use-case scenario a numerical value can be assigned to each combination of true class h and decision d , reflecting the cost incurred by the user of the system when making decision d for a sample of class h (Elkan, 2001; Berger, 2013; Savage, 1972). Once costs are defined for the application of interest, a principled way to measure the performance of a classification system is to take the expectation of the cost over the data. This process results in the EC metric, which is also sometimes called expected prediction error (Hastie et al., 2001), or expected loss (Bishop, 2006), and is equivalent to the negative expected utility (DeGroot, 1970; Bernardo & Smith, 1994). Note that, while the EC definition is very simple, the set of costs it relies upon are sometimes difficult to define given a certain application of interest (Kahneman, 2011). Yet, as we discuss in this work, costs are also implicitly involved in all alternative classification metrics. The benefit of the EC is that the costs are selected explicitly rather than being obscured within the metric’s definition.

The EC is used in statistical learning as the loss to be minimized when making decisions based on posterior probabilities using Bayes decision theory (Hastie et al., 2001; Bishop, 2006). Yet, this metric can be used to evaluate classification systems regardless of how decisions are made (Hernández-Orallo et al., 2012). The error rate and the balanced error rates which correspond to one minus the accuracy and balanced accuracy (or unweighted average recall), respectively, are special cases of the EC. The general form of the EC with unequal costs has been widely adopted by the speech processing community as an evaluation metric for the tasks of speaker verification and language recognition (see, e.g., Brümmer, 2010; Brümmer et al., 2021; Van Leeuwen & Brümmer, 2007; Greenberg et al., 2020), where it is called detection cost function (DCF). The DCF is the metric of choice for the periodic evaluations organized by NIST.³ Other than for these tasks, though, we are not aware of other applications for which the EC is used as standard practice. While some papers stand out proposing the use of EC for medical tasks (Ashby & Smith, 2000; Kornak & Lu, 2011; Godau et al., 2023), they are exceptions in a literature where various other metrics are used instead (Hicks et al., 2022). Interestingly, the EC has been recently highlighted as a “so far uncommon” metric with many desirable features in (Maier-Hein et al., 2022), a new paper on metrics for medical image analysis written by an international consortium of experts.

In this work, we review and compare the EC with other more widely-used classification metrics in the machine learning literature. We describe the EC in detail, explaining how it can be used under various different scenarios, including the case of unequal error severity, the case where the class prior in the evaluation data differs from the one expected in practice, and the case where the decisions do not coincide with the classes (e.g., when the system has an “abstain from choosing a class” option). We discuss the implicit assumptions that are made about the costs when computing error rate and balanced error rate. We also introduce the normalized EC (NEC), a more interpretable version of the EC obtained by dividing the EC by the EC of the best naive system. Finally, we describe alternative classification metrics, like the F-beta score and the MCC, and compare them with the EC, arguing that the EC is superior to those metrics, being more intuitive and flexible, having a theoretical framework for making optimal decisions, and having useful properties that facilitate its interpretation. We support our claims both with theoretical and empirical arguments. To our knowledge, no prior work has compared the expressions of the various classification metrics used in the literature with that of the EC. The code used to compute all metrics and generate all plots and tables in this paper is available as supplemental material. We hope this work will encourage the machine learning community to embrace this valuable but underused metric.

2 The expected cost (EC)

The EC is defined as a generalization of the probability of error for cases in which errors cannot be considered to have equally-severe consequences. In the context of decision theory, the EC is used as the function to

²In some scenarios, the cost of an error depends on characteristics other than just the true class of the sample (e.g., Zadrozny & Elkan, 2001), or are determined by the end user after deployment. Such scenarios are out of the scope of this paper.

³<http://www.nist.gov/itl/iad/mig/sre.cfm>

be minimized in order to make optimal decisions when the system outputs posterior probabilities. See, for example, the work by Bishop (2006, Section 1.5), Hastie et al. (2001, Section 2.4), Duda et al. (2001, Section 2.2), Elkan (2001), DeGroot (1970), and Savage (1972). The EC is the expectation of a cost function which assigns a numerical penalty to each combination of the true class of the sample, h , and decision made by the system, d . Since we assume that both h and d are categorical, this cost function can be specified by a matrix with components $C(h, d)$, which represent the cost we believe that the system should incur for making decision d when the true class of the samples was h .

The assignment of costs for each type of error is not necessarily trivial (Kahneman, 2011). For example, it may not be easy to decide exactly how much more costly it should be to treat a person for a certain pathology they do not have than for a person suffering from the pathology to go untreated. Selecting the costs usually requires knowledge of the application and should be done in consultation with field experts. Yet, while this may seem a disadvantage of the EC, it is important to note that the trade-off between different types of error is encoded in every classification metric (this issue is discussed in Sections 3 and 4). The advantage of the EC over other metrics is that the trade-off is explicit and transparent. Hence, instead of letting the metric choose the trade-off for us, we can make our best effort in choosing costs that are, at least approximately, relevant for our task of interest.

In some scenarios, the selection of costs may be facilitated by allowing the cost matrix to include a combination of positive and negative costs (a negative cost would be equivalent to a utility) or to define, instead, utility matrices, where bad decisions correspond to negative values and good decisions to positive ones. Utility matrices can be converted into cost matrices by multiplying them by minus one. Hence, the analysis can be restricted to cost matrices without loss of generality. Further, for every cost matrix with entries c'_{ij} , a new cost matrix equivalent to the original one can be created with entries given by $c_{ij} = c'_{ij} - \min_j c'_{ij}$. The equivalence is in the sense that systems would be ranked in the same way by both costs matrices (DeGroot, 1970; Brümmer, 2010). The normalized costs are non-negative and the best decision for each class has a cost of zero, which often facilitates the analyses. In this work it is assumed that all costs matrices have been normalized in this way.

In this work, as generally done in the machine learning literature, the expectation is computed with respect to the empirical distribution in the test data, i.e., as an average of the cost over the test samples. Given a test set $\{(h_1, d_1), \dots, (h_N, d_N)\}$, where $h_t \in \mathcal{H} = \{H_1, \dots, H_K\}$ are the true labels and $d_t \in \mathcal{D} = \{D_1, \dots, D_M\}$ are the system's decisions for sample t , the EC is given by:

$$\text{EC} = \frac{1}{N} \sum_{t=1}^N C(h_t, d_t) \quad (1)$$

where $C(h_t, d_t)$ is the cost incurred on sample t with true class h_t and system's decision d_t . We can rewrite this expression as:

$$\text{EC} = \frac{1}{N} \sum_{i=1}^K \sum_{j=1}^M c_{ij} N_{ij} = \sum_{i=1}^K \sum_{j=1}^M c_{ij} P_i R_{ij} \quad (2)$$

where $c_{ij} = C(H_i, D_j)$, $R_{ij} = N_{ij}/N_{i*}$ is the fraction of samples from class H_i for which the system made decision D_j , with $N_{i*} = \sum_j N_{ij}$ being the number of samples of class H_i in the test set, and where $P_i = N_{i*}/N$ is the empirical estimate of the prior probability of class H_i in the evaluation data. The first expression in this equation is derived by using the fact that the cost is the same for each combination of class H_i and decision D_j . Hence, the summation over the samples in Equation (1) can be converted to a (double) summation over every combination of i and j by multiplying the corresponding cost by the number of samples for that combination of class and decision, N_{ij} .

The second expression for the EC in Equation (2) allows for a useful generalization. While P_i is originally given by N_{i*}/N , it can potentially be changed arbitrarily. This allows us to set the priors so that they coincide with those we expect to see when the system is deployed which, in many applications, do not necessarily coincide with those observed in the evaluation data. Rather than forcing the evaluation data to have the priors we expect to see in practice, which would imply downsampling or upsampling the data, we can simply

use those priors when computing the cost and use the evaluation data as-is to compute the R_{ij} values. This advantage of the EC has been highlighted by (Godau et al., 2023), where the authors show the benefits of using an EC metric with the priors determined based on the deployment data rather than on the evaluation data.

Further, note that in the definition of the EC the set of possible decisions do not need to coincide with the set of classes. Rather than being the predicted class labels, the decisions may be the actions that could be taken based on the information in the input sample (Duda et al., 2001, Section 1.3.12). Hence, the EC allows for the evaluation of more general classification systems than other classification metrics like the F-beta score or the MCC, which assume that the set of decisions corresponds to the set of possible class labels. Consider, for example, a (simplified) medical scenario where a doctor needs to make a decision based on an MRI of the brain of a patient that may or may not have a tumor. The decisions made by the doctor do not necessarily need to be one of the two classes (patient does/does not have a tumor). They could instead be the actions that the doctor could take based on the MRI image which could be, for example: “perform surgery”, “send home”, “do more tests”. A cost for each combination of true class and decision can then be selected based on the severity of each type of error. Internally, the system can make decisions in order to optimize the resulting EC, as explained in Section 2.3. The “do more test” decision would be taken in cases in which the system was not certain enough about the class of the sample to make a decision to operate or send the patient home.

2.1 Normalized expected cost

As explained above, any cost or utility matrix can be mapped to an equivalent cost matrix with $c_{ij} \geq 0$, and at least one $c_{ij} = 0$ for each class H_i . The minimum value for any EC with these properties is 0, attained for a perfect system that always chooses a decision with cost 0 for each sample. The maximum value that EC can take, though, depends on the costs and priors used to compute it. This makes it hard to assess whether a certain value of EC is good or bad. For this reason, it is convenient to compute a normalized version of the EC (NEC) given by:

$$\text{NEC} = \frac{\text{EC}}{\text{EC}_{\text{naive}}} \quad \text{where} \quad \text{EC}_{\text{naive}} = \min_j \sum_{i=1}^K c_{ij} P_i \quad (3)$$

The value used for normalization is the EC for the best naive system that does not have access to the input sample and, hence, always makes the same decision \hat{j} . The system is free to choose \hat{j} to optimize the EC for the chosen costs and priors. The expression for EC_{naive} can be easily derived by considering that, for a system that always makes the same decision \hat{j} , R_{ij} is 0 for all $j \neq \hat{j}$ and 1 for $j = \hat{j}$. Plugging those values of R_{ij} in Equation (2) we find that the EC for that system is $\sum_{i=1}^K c_{i\hat{j}} P_i$. The minimum over all possible decisions \hat{j} is the EC of the best naive system.

For the special case where there are two possible true labels ($K = 2$) and two possible decisions ($M = 2$), and we assume that $c_{11} = c_{22} = 0$, the NEC reduces to:

$$\text{NEC} = \frac{c_{12}P_1R_{12} + c_{21}P_2R_{21}}{\min(c_{12}P_1, c_{21}P_2)} = \begin{cases} \alpha R_{12} + R_{21} & \text{if } \alpha \geq 1 \\ R_{12} + \alpha^{-1} R_{21} & \text{otherwise} \end{cases} \quad \text{where } \alpha = \frac{c_{12}P_1}{c_{21}P_2} \quad (4)$$

Hence, in this scenario, there is only one free parameter to choose in the NEC. All combinations of priors and costs that lead to the same α correspond to the same metric.

The NEC metric has an essential property: While its value can be larger than 1, this only happens for systems that are worse than a naive system. In particular, a system with NEC larger than 1 can be trivially improved upon by replacing it with one that always makes the same least-costly decision \hat{j} . Yet, discarding the system is usually not necessary. For systems that first generate a score for each class and then make decision based on those scores, a NEC larger than 1 indicates that the decision stage was not properly optimized. In the binary case, optimization of the decision stage can be achieved by simply tuning the decision threshold. The problem of optimizing the decision stage for a certain EC of choice is discussed in Section 2.3.

2.2 Error rate and accuracy

The standard error rate (also called total error or probability of error) is given by:

$$\text{ER} = \frac{1}{N} \sum_{i=1}^K \sum_{j=1|j \neq i}^K N_{ij} \quad (5)$$

Comparing this equation with Equation (2) we can see that the ER is a particular case of EC obtained when $M = K$, and $c_{ij} = 1$ for $i \neq j$ and $c_{ij} = 0$ for $i = j$, which we will call the “0-1 cost matrix”. Further, since the accuracy is given by one minus the ER, we can see that accuracy is also just a trivial function of one specific case of EC.

In the ER computation, errors from all classes are weighted equally. Hence, when the classes are highly imbalanced, the ER and the accuracy become somewhat insensitive to the performance in the minority classes. For this reason, when errors in the detection of the minority classes are considered more severe than those in the majority classes, the balanced ER (BER), also known as weighted ER, is used as metric instead of the ER. Alternatively, the balanced accuracy, given by one minus the balanced ER, sometimes called unweighted average recall (UAR), is also commonly used in these scenarios. The BER is defined as the average of the ER values per class. That is:

$$\text{BER} = \frac{1}{K} \sum_{i=1}^K \frac{1}{N_{i*}} \sum_{j=1|j \neq i}^K N_{ij} = \frac{1}{K} \sum_{i=1}^K \sum_{j=1|j \neq i}^K R_{ij}. \quad (6)$$

This expression coincides with the EC when $M = K$ and the costs are selected such that $c_{ij} = N/(KN_{i*}) = 1/(KP_i)$ for $i \neq j$ and $c_{ij} = 0$ for $i = j$. As a consequence of setting the costs this way, all error rates (R_{ij} for $i \neq j$) have the same influence on the final metric. Note that by deriving the balanced error rate as a special case of the EC we can see exactly what is being assumed when we compute this metric: that the cost for each type of error is given by the inverse of the frequency of the true class of the sample.

The general form of EC offers greater flexibility than these two special cases of the standard and balanced error rates, allowing us to accommodate scenarios in which the errors are neither all equally costly nor as costly as the inverse of the prior for each class. The EC allows us to think of each type of error independently and set their cost according to their severity.

2.3 Optimal decisions

Most modern classifiers are composed of two stages, one that generates posterior probabilities for the classes and a second one that makes the categorical decisions based on those posteriors. A unique advantage of the EC and NEC is that the optimal decisions for such systems can be obtained in closed form without having to tune decision thresholds (when $K = M = 2$) or regions (for any other case). The decisions that theoretically optimize the EC are called Bayes decisions (Hastie et al., 2001; Bishop, 2006). Given an EC determined by a specific cost matrix c_{ij} the Bayes decision for a sample x is given by:

$$d_B(x) = \arg \min_j \sum_{i=1}^K c_{ij} P(H_i|x) \quad (7)$$

where $P(H_i|x)$ is the posterior probability for class H_i given sample x as produced in an internal stage of our classifier.

One important condition needs to be satisfied for Bayes decisions to be optimal: the posteriors used to make the decisions need to be well calibrated. Loosely speaking, for well-calibrated posteriors, 90% of the times the posterior for a certain class h is 0.9, the true class is h . For a formal definition of calibration, please refer, for example, to the works by DeGroot & Fienberg (1983); Brümmer (2010); Filho et al. (2023). Unfortunately, good calibration is not always achieved by modern classifiers mostly due to overfitting of the parameters to the training data, which results in overconfident posteriors (Guo et al., 2017). Fortunately,

though, calibration quality can usually be easily fixed, if needed, by doing post-hoc calibration on the output of the classifier (Filho et al., 2023). The topic of calibration is out of the scope of this paper, but it is covered at length in the publications cited above and many others. The repository provided with this work includes tools for assessing and fixing calibration.

Assuming our classifier is calibrated, either inherently so or after adding a post-hoc calibration stage, the EC provides us with an elegant approach for making optimal decisions, even for the multi-class case. In fact, as discussed in Appendix A, the argmax decisions so widely used in the machine learning literature correspond to the Bayes decisions for the EC defined with the 0-1 cost matrix, i.e., the error rate (notably, this method is generally used without verifying whether the calibration condition is satisfied). Making Bayes decisions for a general EC of choice is as simple as implementing Equation (7). In Section 4, we show examples on making Bayes decisions for various EC metrics.

3 Other classification metrics

In this section we cover two of the most common binary classification metrics in the machine learning literature and discuss their disadvantages compared to the EC metric, the F-score and the Matthews correlation coefficient. We also briefly describe a common way to use binary classification metrics for the multi-class case. A discussion on three other metrics, net-benefit, likelihood ratio of positive samples, and specificity at a target sensitivity is included in Appendix C.

The area under the ROC curve (AUC), another widely used binary classification metric, is not strictly a classification metric in the sense considered in this work, since it does not evaluate the quality of categorical decisions but of the scores that would eventually be used to make such decisions after choosing one specific threshold. As such, it does not directly reflect the performance that a classification system would have in practice, once deployed with one specific decision threshold. Hence, we do not include AUC in this section, since it belongs to a different family of metrics from the ones analyzed in this work. Nevertheless, Appendix C includes a brief discussion of the AUC and another metric derived from the ROC-curve, the equal error rate (EER), and their relationship with the EC.

3.1 F-beta score

The F-beta score, here F_β , is defined as one minus a metric called effectiveness, which was first introduced in (Rijsbergen, 1979). The F_β is probably one of the most widely used metrics for binary classification in modern machine learning literature, after the accuracy and the error rate. For example, in the SuperGLUE benchmark⁴ of natural language processing (NLP) tasks, out of 10 tasks, 3 use the F_β (one uses MCC and the rest use accuracy). The SQuAD benchmark⁵ for question-answering has two metrics, one of which is F_β . The IEMOCAP benchmark⁶ for emotion recognition is evaluated with (weighted, micro or macro) F_β as well as accuracy. The slot-filling SLURP benchmark⁷ for spoken language understanding uses F_β . The F_β is also widely used in medical imaging, where it is called Dice coefficient (Taha & Hanbury, 2015), in opinion mining (Wang et al., 2019), and for epileptic seizure detection (Siddiqui et al., 2020). These are just a few examples from a long list of tasks where the F_β is selected for evaluation of performance, highlighting the wide spread use of this metric.

The F_β assumes that there are only two classes, that the set of decisions coincides with the set of classes ($K = M = 2$ and $\mathcal{D} = \mathcal{H}$), and that the problem is not symmetric; one of the classes is taken as the class of interest to be detected. Here, we will take the class of interest to be class H_2 . F_β is defined as follows:

$$F_\beta = (1 + \beta^2) \frac{\text{Precision} \text{ Recall}}{\beta^2 \text{ Precision} + \text{Recall}} \quad \text{where} \quad \begin{aligned} \text{Precision} &= N_{22}/N_{*2} \\ \text{Recall} &= N_{22}/N_{2*} \end{aligned} \quad (8)$$

⁴<https://super.gluebenchmark.com/leaderboard>

⁵<https://rajpurkar.github.io/SQuAD-explorer/>

⁶<https://paperswithcode.com/sota/emotion-recognition-in-conversation-on>

⁷<https://paperswithcode.com/sota/slot-filling-on-slurp>

Replacing those values for the precision and recall, we get:

$$F_\beta = \frac{(1 + \beta^2)N_{22}}{(1 + \beta^2)N_{22} + \beta^2 N_{21} + N_{12}} \quad (9)$$

F_β takes values between 0 and 1. Larger values indicate better performance, contrary to the EC for which larger values indicate worse performance. For this reason, in order to compare F_β with EC, it is convenient to work with $1 - F_\beta$, which is given by:

$$1 - F_\beta = \frac{\beta^2 N_{21} + N_{12}}{(1 + \beta^2)N_{22} + \beta^2 N_{21} + N_{12}} = \frac{\beta^2 R_{21}P_2 + R_{12}P_1}{\beta^2 P_2 + R_{*2}} \quad (10)$$

where $R_{*2} = N_{*2}/N$ is the fraction of samples in the evaluation set that are labelled as class H_2 . The numerator in this expression is the EC for the binary case when costs for mistakes are given by $c_{12} = 1$ and $c_{21} = \beta^2$, which we will call EC_{β^2} :

$$EC_{\beta^2} = \beta^2 R_{21}P_2 + R_{12}P_1. \quad (11)$$

We can now express $1 - F_\beta$ as a function of EC_{β^2} and its normalized version:

$$1 - F_\beta = \frac{EC_{\beta^2}}{\beta^2 P_2 + R_{*2}} = \min(\beta^2 P_2, P_1) \frac{NEC_{\beta^2}}{\beta^2 P_2 + R_{*2}} \quad (12)$$

That is, $1 - F_\beta$ is proportional to EC_{β^2} with a scaling factor given by the inverse of $\beta^2 P_2 + R_{*2}$. The first term in that expression is independent of the system, it only depends on the prior of class H_2 in the evaluation dataset, P_2 . When comparing different systems on the same dataset (or on different datasets with the same class priors) P_2 is fixed.⁸ The second term in the denominator in Equation (12), R_{*2} , is the percentage of samples that are labeled by the system as being of class H_2 . Given two systems with the same value for EC_{β^2} , F_β will favor the one that detects more samples as being from class H_2 , regardless of whether they are correctly or incorrectly classified (examples of this behavior will be given in Section 4). This is an odd thing to reward: having more samples from class H_2 detected is not beneficial in itself. If, for the application of interest, detecting the samples from class H_2 is more important than detecting the samples of class H_1 , the principled approach for defining a metric is to use an EC which penalizes errors in samples from class H_2 more than errors in samples of class H_1 (i.e., setting $c_{21} > c_{12}$).

As for the EC, we can compute the best F_β that would correspond to a naive system that always outputs the same decision. If the system always chooses class H_1 , $F_\beta = 0$ since $N_{22} = 0$. On the other hand, if the system always chooses class H_2 , the recall is 1 and the precision is equal to the prior of class H_2 , P_2 , which results in $F_\beta = (1 + \beta^2)P_2/(\beta^2 P_2 + 1)$. This is the value of F_β for the best naive system. Any system with an F_β worse than this value should be considered ineffective. A normalized version of the F_β is not used in this work because this is not usual practice in the community.

3.2 Matthews correlation coefficient

The Matthews correlation coefficient (MCC) was first introduced by Matthews (1975) for comparison of chemical structures, later proposed as a metric for binary classification by Baldi et al. (2000), and finally generalized to the multi-class case by Gorodkin (2004). In this section we will consider the binary version since this is enough to show its flaws as classification metric. Chicco & Jurman (2020) argue that MCC is the most informative single score to establish the quality of a binary classifier that outputs hard decisions. EC is not considered as an option in that paper. As we discuss below, we believe EC is superior to MCC as a metric for classification performance.

⁸Note that values of F_β across datasets with different priors are not comparable to each other. In particular, the F_β computed on a certain evaluation dataset is not a good predictor of the F_β that will be obtained during deployment if the priors between the two sets differ. This scenario can be addressed in a principled way with the EC by setting the priors in Equation (2) to those that are expected during deployment (Godau et al., 2023).

As the F_β , the MCC assumes that the set of decisions is the same as the set of classes. For the binary case MCC is defined as follows:

$$\text{MCC} = \frac{N_{11}N_{22} - N_{12}N_{21}}{\sqrt{(N_{11} + N_{21})(N_{11} + N_{12})(N_{22} + N_{12})(N_{22} + N_{21})}} \quad (13)$$

After some manipulations and using $R_{ij} = N_{ij}/N_{i*}$, we get

$$\text{MCC} = \sqrt{\frac{N_{1*}N_{2*}}{N_{*1}N_{*2}}} (1 - (R_{12} + R_{21})) = \sqrt{\frac{P_2(1 - P_2)}{R_{*2}(1 - R_{*2})}} (1 - \text{NEC}_b) \quad (14)$$

where $\text{NEC}_b = R_{12} + R_{21}$ is the normalized BER (this can be seen by dividing Equation (6) by the BER of the naive systems which is 0.5 for the binary case).

From this equation we see that when the NEC_b is larger than 1, which, as explained in Section 2.1, happens when the system is worse than the best naive system, the MCC is negative. Further, we can see that MCC and $1 - \text{NEC}_b$ differ by the square root factor. The numerator in this factor is fixed, it does not depend on the system, only on the priors (see footnote 8, which also applies to the MCC). The denominator, on the other hand, does depend on the system and is largest when $R_{*2} = 0.5$ and smaller as the decisions become more imbalanced (i.e., as R_{*2} approaches 0 or 1). Hence, the factor will grow as the decisions become more imbalanced, making MCC larger in absolute value. For two systems with the same $\text{NEC}_b < 1$, MCC will favor the one with larger imbalance in the decisions. As discussed with respect to F_β , this behavior does not seem reasonable since the frequency with which a system makes a certain decision is not a good or bad quality in itself. As in the F_β case, the MCC rewards behaviors of the system that do not directly reflect what is relevant for a task.

3.3 Multi-class metrics

The most common metrics for the multi-class case are accuracy and balanced accuracy (unweighted average recall) which, as we saw before, are one minus (a special case of) an EC. Another metric sometimes used for the multi-class case is a generalization of the F score called the macro F score. This metric is computed using Equation (8) with precision and recall calculated as averages over the per-class precision and recall (Grandini et al., 2020). The resulting metric offers no control over the importance of each type of error and no theoretical method for making decisions to optimize it.

One of the important advantage of the EC is that it can be used for the multi-class case without changes. Once the costs for each combination of true class and decision are defined, the metric is ready for use. In addition, as explained in Section 2.3, the EC comes with an elegant decision theory for making optimal decisions based on classifiers that output posterior probabilities, which is particularly useful for the multi-class case (or when the number of decisions is larger than two) where threshold tuning is not an option. We illustrate the use of the EC for the multi-class case in Section 4.2.

4 Empirical comparison

In this section we show an analysis of the metrics discussed in the previous section using two simulated datasets with 2 and 10 classes. The purpose of this section is to illustrate and further discuss the concepts explained in the previous sections. The procedure for generating the scores is described in Appendix D. Briefly, 100,000 samples are generated with an imbalanced prior distribution, with $P_1 = 0.8$ and equal prior for the rest of the classes. Samples for each class are drawn from a Gaussian distribution where the mean are separated by 1.0 and the standard deviation is 0.2, unless otherwise indicated. Then, perfectly calibrated posteriors are obtained using the known distribution and priors for each class.

4.1 Comparison of metrics for a two-class dataset

Figure 1 shows results for the 2-class dataset for F_1 , MCC, and three different NEC with $c_{12} = 1$ and $c_{21} = c$ set to 1, 4, and 16. Since $P_1 = 0.8$, the NEC with $c = 4$ corresponds to the normalized balanced error




Figure 1: Comparison of metrics for a simulated dataset with Gaussian per-class distributions and $P_1 = 0.8$. The subindex in the NEC values indicates the value of c_{21} , while c_{12} is 1.0 in all cases. The first plot shows the distribution of the log odds for each class scaled by the prior probability of the class. The other three plots show various metrics as a function of the log odds decision threshold. The dashed lines are included only to aid visualization of the various points highlighted in the discussion.

rate, $\text{NEC}_b = R_{12} + R_{21}$ (this can be verified by plugging these costs and priors in Equation (4)). The left plot in the figure shows the distribution of log posterior odds, $\log(P(H_2|x)/P(H_1|x))$, for each of the two classes. The rest of the plots show the behaviour of different metrics as a function of the decision threshold. The decisions are made by comparing the log odds score for each sample with the corresponding threshold, labelling the sample as class H_2 if the score is larger than the threshold and as class H_1 otherwise. For each threshold we can then obtain the N_{ij} counts needed to compute all metrics. Dashed vertical lines identify the value of different metrics for the same threshold, and, as a consequence, the same set of decisions.

The second plot in Figure 1 shows a comparison between one minus F_1 and NEC_1 , which are related by Equation (12). The fraction of samples labelled as class H_2 for each threshold, R_{*2} , is also shown. The figure highlights two thresholds: the one that minimizes $1 - F_1$ and another one that results in the same value of NEC_1 . As discussed in Section 3.1, between two systems with the same NEC_1 , the F_1 prefers the one with a larger R_{*2} , effectively giving more cost to the errors in class H_2 by favoring systems that produce more detections of this class. Concretely, given the threshold t_o for which F_1 is minimized in a certain dataset, an equivalent NEC for which the Bayes threshold is t_o can be defined by using Equation (20) in the Appendix, setting $c_{12} = 1$ and $c_{21} = e^{-t_o}$. For the example in Figure 1, the NEC that would lead to the same optimal threshold as the F_1 ($t_o = -0.60$) is the one with $c_{21} = 1.8$. That is, in this case, optimizing F_1 is equivalent to optimizing NEC assuming that errors in samples of class H_2 are 1.8 times more costly than those in samples of class H_1 . The problem is that this cost is not explicitly selected by the user of the metric and, instead, depends on the distribution of the scores under evaluation.

Figure 2 illustrates the dependency on the scores of the effective costs corresponding to F_1 . The figure shows the F_1 and two NEC for two different datasets: (1) the same dataset used for Figure 1, and (2) a dataset obtained using the same procedure as for the first but changing the standard deviation for the features from 0.20 to 0.06, making the task easier. We can see that the F_1 is optimized at a different threshold for both datasets, corresponding to an equivalent c_{21} of 1.8 for the original dataset, as discussed above, and an equivalent cost of 1.0 for the easier dataset (with c_{12} always fixed at 1.0). Hence, the F_1 effectively optimizes for a different trade-off between error types depending on the score distribution. On the other hand, since both datasets are well-calibrated, the Bayes threshold for each of the NECs (Appendix A.2) optimizes the corresponding NEC. This, we believe, is one of the strongest arguments to prefer the EC over the F_β (the analysis above was done for F_1 but holds for any F_β). For EC the costs can be explicitly selected and the optimal threshold for calibrated posteriors can be determined theoretically as a function of the costs. On the other hand, for F_β the costs are implicit and depend on factors that are outside of our control resulting on the optimal threshold and the effective costs varying across datasets.

The third plot in Figure 1 shows one minus MCC and the NEC value with $c = 4$, along with R_{*2} . The MCC and this NEC are related by Equation (14). The figure highlights two thresholds: the one that optimizes MCC and a second one that results in the same value of NEC_4 as the first one. Comparing the values of




Figure 2: F_1 and two NEC metrics with $c_{12} = 1$ and c_{21} indicated as a subindex, for two different datasets. In the left plot, the metrics are computed for the same dataset as in Figure 1. In the right plot they are computed on an easier dataset (note the change in the limits in the y axis). The two NEC are selected to use the effective cost corresponding to F_1 for each of the two datasets. For the easier dataset (right), the best threshold for F_1 coincides with the Bayes threshold for $NEC_{1.0}$. For the harder dataset (left), the best threshold for F_1 coincides with the Bayes threshold for $NEC_{1.8}$.

MCC across these two threshold, we can observe the behavior explained in Section 3.2: among two thresholds that are equivalent in terms of NEC_4 , the MCC prefers the one that results in a more imbalanced R_{*2} . Again, as for F_β , the optimal threshold for MCC is determined, in part, by the value of R_{*2} . As a consequence, also as for F_β , the equivalent cost implicit in the MCC depends on the dataset under evaluation giving us no explicit control over the trade-off between types of errors.

Finally, the fourth plot in Figure 1 compares the two NECs from the middle plots with each other and with a third NEC with a larger c_{21} . The Bayes threshold for NEC_c is given by $\log(1/c)$ (see Equation (20) in the Appendix): larger values of c_{21} move the threshold to the left decreasing the error rate for class H_2 and increasing it for class H_1 . Note that, as mentioned in Section 2.3, the Bayes threshold is only optimal if the scores are well-calibrated. If this is not the case, the threshold can be tuned empirically, as usually done for any other metric. Alternatively, a calibration stage can be added to the system enabling the use of Bayes thresholds (Filho et al., 2023). In any case, the EC allows us to explicitly select the relative cost between the different error types which directly and transparently affects the resulting metric values and, when the scores are calibrated, the corresponding optimal threshold.

4.2 Comparison of ECs for a 10-class dataset

In this section we illustrate the use of the EC for a 10-class problem. We compare EC and NEC obtained with the following cost matrices: 1) the 0-1 matrix (**C01**) which corresponds to the error rate, 2) a matrix with $c_{ij} = 1/(KP_i)$ for $i \neq j$ (**CinvP**) which corresponds to the balanced error rate, 3) the C01 matrix with the last row multiplied by 100, representing a use case where errors in the last class are much more costly than errors in other classes (**Cimb**), 4) the C01 matrix with an additional column corresponding to an abstention decision, with cost of 0.05 for all classes (**Cabs1**), and 5) as Cabs1 but with an abstention cost of 0.3 for all classes (**Cabs2**). Table 1 shows the EC and NEC results for these cost matrices for three methods for making categorical decisions: 1) choosing always the best naive decision for each cost matrix, 2) making “argmax” decisions, i.e., choosing the class with the largest score for each sample, regardless of the cost matrix (for this row, decisions are the same for all cost matrices), and 3) making the Bayes decisions corresponding to each cost matrix using Equation (7).

Comparing the last two lines in this table we see that, as theory shows, Bayes decisions result in better performance than argmax decisions, except for the C01 matrix for which Bayes decisions are the argmax decisions and, hence, those two results are the same. Further, the table highlights the benefit of normalization. While the EC might be misleadingly low for some cost matrices due to the imbalanced priors and the selected cost values, the NEC solves this problem by using the EC of the best naive system as baseline with which to normalize, resulting in a more interpretable metric.

The best naive decision depends on the cost matrix as determined by Equation (3). For C01, the class with the largest prior, H_1 , is the least-costly naive decision. For CinvP, all decisions are equally good due to the

Table 1: EC and NEC for various costs matrices defined in Section 4.2 for systems that make the best naive decision (the decision made by this system is indicated in parenthesis under the cost matrix name), argmax decisions, and Bayes decisions. The accuracy and balanced accuracy can be obtained as one minus the EC for C01 and CinvP, respectively. For the two cost matrices with an abstain decision, the percentage of abstain decisions made by the system is also listed (abs).

Naive Decision → Decision algorithm	C01 (class H_1)		CinvP (any class)		Cimb (class H_{10})		Cabs1 (abstain)			Cabs2 (class H_1)		
	EC	NEC	EC	NEC	EC	NEC	EC	NEC	abs	EC	NEC	abs
Naive	0.20	1.00	0.90	1.00	0.98	1.00	0.05	1.00	100	0.20	1.00	0
Argmax	0.06	0.32	0.28	0.31	0.36	0.37	0.06	1.29	0	0.06	0.32	0
Bayes	0.06	0.32	0.23	0.26	0.08	0.08	0.02	0.35	25	0.06	0.28	7

way the costs are defined. For Cimb, the least costly decision is H_{10} due to the large cost assigned to errors in that class. Finally, the least-costly decision for Cabs1 is to abstain, while for Cabs2 it is H_1 (see Appendix B for the expression for Bayes decisions and for the best naive decisions for the Cabs cost matrices). Looking at the percentage of abstention decisions made for those two matrices using the Bayes decision algorithm, for Cabs1 which has a lower abstention cost than Cabs2, the algorithm makes more abstention decisions. Finally, we can see that while the lower abstention cost in Cabs1 results in a lower EC value than for Cabs2, the NEC shows the opposite trend. This happens because the best naive system is different for both costs resulting in different baselines for both ECs.

Note that the scores used for these simulated experiments are, by design, perfectly calibrated, so Bayes decisions are optimal in this dataset. In practice, this might not be the case for a certain dataset of interest. Unlike in the scenario of the previous section where we could simply tune a threshold instead of using Bayes decision theory, for the multi-class case or, more generally, when the number of decisions is larger than 2, tuning the decision regions is not trivial (though attempts have been made in this direction (e.g., O'Brien et al., 2008)). In those scenarios, we recommend making sure that the system outputs are well-calibrated (Filho et al., 2023) and then applying Bayes decision theory as done in this section.

5 Conclusions

In this work we focus on the problem of evaluating the quality of categorical decisions generated by a machine learning classifier. We argue that the expected cost (EC), a classic evaluation function proposed decades ago, should be the primary metric to assess the performance of categorical decisions for a certain application of interest. The EC is based on the simple assumption that each decision made by the system can be assigned a cost which depends on the true class of the sample. The expectation of the cost over the data can then be used as a metric to assess the performance of a system. The EC can be customized to any application for which categorical decisions are required by adapting the costs, and optionally the class priors, according to the needs of the application.

While two special cases of the EC, the error rate and the balanced error rate are widely used in the machine learning literature, its general form is only rarely leveraged. Instead, metrics like the F-beta score or the MCC are used for applications in which the error rate is deemed inadequate for performance assessment. This is often the case in applications with large class imbalance and where errors in the minority class are considered more important than those in the majority class. In this work, we argue that the EC is a better metric than these alternatives for having the following important characteristics:

- The EC is defined based on the simple and general assumption that each combination of true class and decision can be assigned a corresponding fixed cost. In contrast, the trade-off between different types of error is obscured in other classification metrics.

- The class prior distribution can be considered a parameter of the metric. This is useful in cases in which the priors in the evaluation data do not coincide with those expected at deployment.
- Both binary and multi-class classification systems can be evaluated in a principled way with the EC, using the same general expression.
- Unlike other classification metrics, the EC does not assume that the set of possible decisions is the set of classes. Instead, the decisions can be the actions that may be taken by the user of the system.
- When normalized by the EC of the best naive systems for the selected costs and priors, the resulting metric is easily interpretable.
- Optimization of categorical decisions can be done using Bayes decision theory as long as the system scores are well-calibrated, eliminating the need for ad-hoc solutions.

The need to choose costs may seem like a drawback of the EC compared to other metrics like the F1 score or the MCC which have no parameters to select. Yet, it is important to note that, while these alternative metrics do not explicitly require the selection of costs, they effectively impose a trade off between different types of error. This is easy to see for the binary case, where decisions are made by thresholding, by noting that the threshold that optimizes a metric for a certain set of scores is the Bayes threshold for some EC. Hence, optimizing the original metric is equivalent to optimizing the related NEC. Effective costs are implicitly considered in the optimization. Yet, those costs are selected without our control and, to make things worse, they depend on the score distribution. The EC, on the other hand, requires an explicit choice of costs. Once the costs are chosen, the trade-off between different types of errors is fixed and transparent and does not depend on the scores under evaluation. While the costs may be hard to specify in some cases, we argue that having control over them is better than letting the metric implicitly choose them depending on the score distribution.

In conclusion, we argue that the EC, or its normalized version, should be the preferred metric for evaluating categorical decisions for a given application of interest. This work is accompanied by a python library of methods to compute the EC, along with utility methods for making Bayes decisions and calibrating scores. The code includes notebooks that produce the results in this paper. We hope this work will facilitate the wider adoption by the machine learning community of this classic, simple, and well-motivated metric.

References

- D. Ashby and A. F. M. Smith. Evidence-based medicine as bayesian decision-making. *Statistics in Medicine*, 19(23):3291–3305, 2000.
- P. Baldi, S. Brunak, Y. Chauvin, C. A. F. Andersen, and H. Nielsen. Assessing the accuracy of prediction algorithms for classification: an overview. *Bioinformatics*, 05 2000.
- J. O. Berger. *Statistical decision theory and Bayesian analysis*. Springer Science & Business Media, 2013.
- J. M. Bernardo and A. F. M. Smith. *Bayesian Theory*. John Wiley & Sons, 1994.
- C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
- A. P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997.
- N. Brümmer. *Measuring, Refining and Calibrating Speaker and Language Information Extracted from Speech*. PhD thesis, Stellenbosch University, 2010.
- N. Brümmer, L. Ferrer, and A. Swart. Out of a Hundred Trials, How Many Errors Does Your Speaker Verifier Make? In *Proc. Interspeech*, Brno, Czech Republic, Sept. 2021.
- D. Chicco and G. Jurman. The advantages of the matthews correlation coefficient (MCC) over F1 score and accuracy in binary classification evaluation. *BMC Genomics*, 21(1), Jan 2020.

- M. H. DeGroot. *Optimal Statistical Decisions*. McGraw-Hill, 1970.
- M. H. DeGroot and S. E. Fienberg. The comparison and evaluation of forecasters. *Journal of the Royal Statistical Society. Series D (The Statistician)*, 32(1/2):12–22, 1983.
- R. Duda, P. Hart, and D. Stork. *Pattern Classification*. Wiley, 2001.
- C. Elkan. The foundations of cost-sensitive learning. *Proceedings of the Seventeenth International Conference on Artificial Intelligence: 4-10 August 2001; Seattle*, 1, 05 2001.
- T. S. Filho, H. Song, M. Perello-Nieto, R. Santos-Rodriguez, M. Kull, and P. Flach. Classifier calibration: How to assess and improve predicted class probabilities: a survey. *Machine Learning*, 112, 2023.
- T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 2007.
- P. Godau, P. Kalinowski, E. Christodoulou, A. Reinke, M. Tizabi, L. Ferrer, P. Jäger, and L. Maier-Hein. Deployment of image analysis algorithms under prevalence shifts. In *Proc. of MICCAI*, 2023.
- J. Gonzalez-Rodriguez. Evaluating automatic speaker recognition systems: An overview of the NIST speaker recognition evaluations (1996-2014). *Loquens*, 1(1), 2014.
- J. Gorodkin. Comparing two k-category assignments by a k-category correlation coefficient. *Computational Biology and Chemistry*, 2004.
- M. Grandini, E. Bagli, and G. Visani. Metrics for multi-class classification: an overview. arXiv 2008.05756, 2020.
- C. S. Greenberg, L. P. Mason, S. O. Sadjadi, and D. A. Reynolds. Two decades of speaker recognition evaluation at the National Institute of Standards and Technology. *Computer Speech and Language*, 60, 2020.
- C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger. On calibration of modern neural networks. In *Proc. of the 34th International Conference on Machine Learning*, Sydney, Australia, 2017.
- J. C. Harsanyi. Bayesian decision theory and utilitarian ethics. *The American Economic Review*, 68(2): 223–228, 1978.
- T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning*. Springer-Verlag, 2001.
- J. Hernández-Orallo, P. Flach, and C. Ferri. A unified view of performance metrics: Translating threshold choice into expected classification loss. *Journal of Machine Learning Research*, 13(91), 2012.
- S. A. Hicks, I. Strümke, V. Thambawita, M. Hammou, M. A. Riegler, P. Halvorsen, and S. Parasa. On evaluation metrics for medical applications of artificial intelligence. *Scientific Reports*, 12(1), apr 2022.
- D. Kahneman. *Thinking, fast and slow*. 1st ed. New York : Farrar, Straus and Giroux, 2011.
- J. Kornak and Y. Lu. Bayesian decision analysis for choosing between diagnostic/prognostic prediction procedures. *Statistics and Its Interface*, 4(1):27–36, 2011.
- L. Maier-Hein, A. Reinke, P. Godau, and et.al. Metrics reloaded: Recommendations for image analysis validation. *arxiv:2206.01653*, 06 2022.
- B. W. Matthews. Comparison of the predicted and observed secondary structure of t4 phage lysozyme. In *Biochim Biophys Acta - Protein Struct.*, 1975.
- D. B. O'Brien, M. R. Gupta, and R. M. Gray. Cost-sensitive multi-class classification from probability estimates. In *Proceedings of the 25th international conference on Machine learning - ICML '08*. ACM Press, 2008.

- S. G. Pauker and J. P. Kassirer. Therapeutic decision making: A cost-benefit analysis. *New England Journal of Medicine*, 293(5):229–234, Jul 1975.
- C. J. V. Rijsbergen. *Information Retrieval*. Butterworths, 1979.
- L. J. Savage. *The foundations of statistics*. Courier Corporation, 1972.
- M. K. Siddiqui, R. Morales-Menendez, X. Huang, and N. Hussain. A review of epileptic seizure detection using machine learning classifiers. *Brain Informatics*, 7(1), may 2020.
- A. A. Taha and A. Hanbury. Metrics for evaluating 3d medical image segmentation: analysis, selection, and tool. *BMC Medical Imaging*, 15(1), aug 2015.
- D. A. Van Leeuwen and N. Brümmer. An introduction to application-independent evaluation of speaker recognition systems. In *Speaker classification I: Fundamentals, Features, and Methods*. Springer-Verlag, 2007.
- A. J. Vickers, B. Van Calster, and E. W. Steyerberg. Net benefit approaches to the evaluation of prediction models, molecular markers, and diagnostic tests. *BMJ*, 352, 2016.
- R. Wang, D. Zhou, M. Jiang, J. Si, and Y. Yang. A survey on opinion mining: From stance to product aspect. *IEEE Access*, 7:41101–41124, 2019.
- B. Zadrozny and C. Elkan. Learning and making decisions when costs and probabilities are both unknown. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '01, pp. 204–213, New York, NY, USA, 2001. Association for Computing Machinery.

A Special cases of Bayes decisions theory

In this section we discuss two special cases of Bayes decision theory, the argmax rule with is obtained when the cost function is given by the 0-1 cost matrix, and the thresholding rule obtained for the binary case with only two decisions.

A.1 The 0-1 cost matrix

A special scenario of Bayes decision theory applied to the EC arises when the cost matrix is the 0-1 square matrix given by one minus the identity matrix. In this case, the EC reduces to the error rate or one minus the accuracy (see Section 2.2). The Bayes decisions for this EC can be obtained by replacing $c_{ij} = 0$ for $i = j$ and $c_{ij} = 1$ for $i \neq j$ in Equation (7) in Section 2.3, which gives:

$$d_B(x) = \arg \min_{\hat{j}} \sum_{i=1|i \neq \hat{j}}^K P(i|x) \quad (15)$$

$$= \arg \min_{\hat{j}} 1 - P(\hat{j}|x) \quad (16)$$

$$= \arg \max_{\hat{j}} P(\hat{j}|x). \quad (17)$$

This is the standard argmax decision rule used in most of the literature in machine learning for making decisions in a multi-class setting. We can now see that this rule is optimal for the 0-1 cost matrix. For any other cost matrix, the general expression in Equation (7) should be used to make decisions.

A.2 Binary classification with two decisions

Another special case of Bayes decision theory corresponds to binary classification when the cost matrix is square (i.e., when there is only two possible decisions) with zeros in the diagonal. This is the simplest and

most common case in which there is one correct zero-cost decision for each class. In that case, Bayes decisions reduce to:

$$d_B(x) = \begin{cases} H_2, & \text{if } c_{12}P(H_1|x) < c_{21}P(H_2|x) \\ H_1, & \text{otherwise} \end{cases} \quad (18)$$

Using the fact that $P(H_2|x) = 1 - P(H_1|x)$, this rule becomes:

$$d_B(x) = \begin{cases} H_2, & \text{if } P(H_2|x) > t \\ H_1, & \text{otherwise} \end{cases} \quad \text{where } t = c_{12}/(c_{12} + c_{21}) = 1/(1 + c_{21}/c_{12}) \quad (19)$$

Sometimes we prefer to work with posterior log odds because the logarithm mitigates numerical precision problems and because their distributions are nicer for plotting (see an example in Figure 1 of the main text). When using log odds as scores, the decision rule becomes:

$$d_B(x) = \begin{cases} H_2, & \text{if } \log(P(H_2|x)/P(H_1|x)) > t_{\text{logodds}} \\ H_1, & \text{otherwise} \end{cases} \quad \text{where } t_{\text{logodds}} = \log(c_{12}/c_{21}) \quad (20)$$

In summary, for binary classification with one zero-cost decision per class, Bayes decisions simply consist of comparing the posterior for one of the classes, or some function of the posteriors, like the logodds, with a specific threshold which is a function of the ratio between the two costs. While to optimize metrics like the F_β or MCC we need to sweep a threshold and choose the best one for that metric, for the EC we simply need to plug the costs into Equation (19) or (20) to get the optimal threshold.

Note that, as discussed in the main text, Bayes decisions are optimal as long as the classifier is calibrated for the selected cost. If this is not the case, for the binary case above, the threshold can be tuned as for any other metric. For the general case, including multi-class cases or binary cases with more than two decisions (non-square cost matrix), the posteriors can be calibrated to enable the use of Bayes decision theory (Filho et al., 2023).

B Cost matrices with an abstention option

Finally, we derive the expression for the Bayes decisions when the set of decisions is given by $\{H_1, \dots, H_K, \text{abstain}\}$, the list of classes plus an abstain decision, and the cost matrix has the following form:

$$\mathbf{C} = [1 - I \quad \alpha \mathbf{1}] \quad (21)$$

where I is the identity matrix, so that $1 - I$ is the 0-1 cost matrix with a cost of 0 in the diagonal and a cost of 1 elsewhere, and where $\alpha \mathbf{1}$ is a vector of ones multiplied by a scalar α which determines the cost of abstention, set to be equal for all classes. Replacing these costs in Equation (7) we get the following expression:

$$d_B(x) = \arg \min_j \sum_{i=1}^K c_{ij} P(H_i|x) = \arg \min_j \begin{cases} \sum_{i=1}^K P(H_i|x) = 1 - P(H_j|x), & \text{if } \hat{j} \neq \text{abstain} \\ \sum_{i=1}^K \alpha P(H_i|x) = \alpha, & \text{if } \hat{j} = \text{abstain} \end{cases} \quad (22)$$

Or, equivalently:

$$d_B(x) = \arg \max_j \begin{cases} P(H_j|x), & \text{if } \hat{j} \neq \text{abstain} \\ 1 - \alpha, & \text{if } \hat{j} = \text{abstain} \end{cases} \quad (23)$$

Hence, the system will abstain if all the posteriors output by the system are smaller than $1 - \alpha$.

We can now easily derive the expression for the best naive decision which is the Bayes decision when the system outputs always the prior probability of each class in place of the posteriors⁹. Doing that, we obtain:

$$d_{\text{naive}} = \arg \max_j \begin{cases} P_j, & \text{if } \hat{j} \neq \text{abstain} \\ 1 - \alpha, & \text{if } \hat{j} = \text{abstain} \end{cases} \quad (24)$$

⁹Note that this same procedure can be used to arrive at the general expression for the EC of the best naive system, Equation (3) in Section 2.1, which was derived in a different way in that section.

Hence, the best naive decision is to abstain when all priors are smaller than $1 - \alpha$. For our multiclass experiments in Section 4 where $P_1 = 0.8$ and $P_i = 0.022$ for $i \neq 1$, the best naive decision is to abstain for Cabs1 where $1 - \alpha = 0.95$, but not for Cabs2, where $1 - \alpha = 0.7$. In that case, the best decision is H_1 , which has a prior of 0.8.

C Other classification metrics

In this section we describe other additional common metrics for classification besides the two discussed in Section 3 (F-beta score, and MCC) and compare them with the EC.

C.1 Net benefit

Net benefit, a metric used in binary classification for some medical applications (Vickers et al., 2016), is defined as:

$$\text{NB} = \frac{N_{22}}{N} - \frac{p_t}{1 - p_t} \frac{N_{12}}{N} \quad (25)$$

where p_t is a parameter of the metric called threshold probability for reasons that are explained below. We can rewrite this expression as follows:

$$\text{NB} = P_2 R_{22} - \frac{p_t}{1 - p_t} P_1 R_{12} = P_2 - (P_2 R_{21} + \frac{p_t}{1 - p_t} P_1 R_{12}) \quad (26)$$

where we used the following equalities: $P_i R_{ij} = N_{ij}/N$ and $R_{22} + R_{21} = 1$. The term in parenthesis in this equation is the EC when the number of labels and decisions is 2 ($K = M = 2$) and $c_{21} = 1$, and $c_{12} = p_t/(1 - p_t)$. We can then express NB as a function of the corresponding normalized EC, which we will call NEC_{p_t} :

$$\text{NB} = P_2 - \min \left(P_2, \frac{p_t}{1 - p_t} P_1 \right) \text{NEC}_{p_t} \quad (27)$$

Hence, while NB is an affine function of a NEC, making them equivalent for ranking systems, the NB loses one of the attractive qualities of the NEC: its ease of interpretation. A NEC value of 1.0 indicates that the system has the same performance as a naive system that does not use the input samples to make its decisions. For this reason, we believe normalized EC is preferable to NB.

NB's single parameter p_t is defined in such a way that it coincides with the optimal threshold for this metric. We can see this by first noting that the optimal threshold for NB is the same as the optimal threshold for NEC_{p_t} (since NB reaches its maximum value when NEC_{p_t} reaches its minimum). The threshold for NEC_{p_t} is given by Equation (19) with $c_{21} = 1$, and $c_{12} = p_t/(1 - p_t)$ resulting in $t = p_t$, which explains the fact that this parameter is called threshold probability.

Selecting a reasonable threshold may be easier than selecting costs for some applications. When this is the case, the implicit cost matrix corresponding to the selected threshold can be easily found using Equation (19), setting one of the costs to 1.0 and solving for the other. Setting one of the costs to 1.0 does not restrict the family of NECs we are able to obtain, since, as shown in Equation (4) in the main document, the NEC value, as the Bayes threshold, only depends on the ratio between the two costs.

C.2 LR+

The likelihood ratio for positive results, usually called LR+, is commonly used in diagnostic testing (Pauker & Kassirer, 1975). As for F_β and NB, it is used for non-symmetric binary classification problems where one of the classes is the class of interest. It is given by:

$$\text{LR+} = \frac{\text{sensitivity}}{1 - \text{specificity}} \quad (28)$$

where, if we assume class H_2 is the class of interest, sensitivity = R_{22} and specificity = R_{11} . LR+ is always positive and the larger its value, the better the system.

We can write LR+ as a function of a NEC value using the fact that $R_{22} + R_{21} = 1$, and $R_{11} + R_{12} = 1$:

$$\text{LR+} = \frac{1 - R_{21}}{R_{12}} = \frac{1 - \text{NEC}_b}{R_{12}} + 1 \quad (29)$$

where $\text{NEC}_b = R_{12} + R_{21}$ as defined in Section 3.2. So, given two systems with the same NEC_b , the LR+ will favor the one with lower R_{12} . We see no reason to prefer this metric over the EC where, if favoring lower values of R_{12} is desired, one can directly and explicitly set $c_{12} > c_{21}$ to the actual costs needed for the application.

C.3 Area under the curve and equal error rate

Two very common metrics used to evaluate binary classification systems are the area under the ROC curve (AUC) (Bradley, 1997) and the equal error rate (EER) (Brümmer et al., 2021). These metrics, unlike all other metrics covered in this paper which evaluate categorical decisions, evaluate the quality of the scores. The assumption made by these metrics is that the decisions will be made by thresholding, but no commitment to a specific threshold is made. Under this assumption, two types of curves can be created by sweeping a threshold, making categorical decisions for each case, and computing the resulting R_{ij} values:

- Receiver Operating Characteristic (ROC) curves, which correspond to R_{22} versus R_{12}
- Recall vs Precision (PR) curve, which correspond to R_{22} vs N_{22}/N_{*2} .

For each of these two curves, the area under the curve (AUC) can be obtained as a summary metric. Further, from the ROC we can find the threshold for which R_{12} equals $R_{21} = 1 - R_{22}$. The value of R_{12} at this threshold is called Equal Error Rate (EER). Interestingly, there is a very close tie between the EER and the total error rate (EC for the 0-1 cost matrix) computed over Bayes decisions: this EC is upper bounded by $\min(\text{EER}, P_1, P_2)$ (Brümmer et al., 2021), as long as calibration is perfect. If calibration is not perfect, this EC can grow without bound, while EER will not.

The ROC and PR curves and the resulting AUCs are immune to any monotonic transformation of the scores. They implicitly assume that, if hard decisions need to be made, the selection of the threshold can always be done optimally. This may or may not be a reasonable assumption in practice, depending on the task, the application and the available data. Further, the AUCs do not reflect the performance that a classification system used to make categorical decisions will have in practice since such systems need to commit to one specific threshold. Hence, their performance should be measured for the selected threshold rather than integrating over all possible thresholds.

Hence, while AUC may be a useful metric during the early stages of development, at some point a developer working on a system that will be used to make categorical decisions for a certain application should commit to a specific set of costs that reflect the needs of that application. If the application for which the system will be used is unknown during development or the costs cannot be determined ahead of deployment, then the system needs to be optimized for any possible threshold. In this case, using AUC as a metric would be appropriate, knowing that the actual threshold for each application would have to be determined empirically using some development data. Alternatively, the quality of the scores can be assessed using strictly proper scoring rules (SPSRs) which evaluate the quality of the scores as posterior probabilities (Gneiting & Raftery, 2007). A system developed to have low expected SPSR can be thresholded with the Bayes threshold without the need to tune the threshold for each new set of costs.

Finally, note that the AUC and EER metrics cannot be naturally generalized to the multi-class case, or to binary cases with more than two decisions. Hence, when the task has more than two classes or decisions, these metrics are not appropriate. While ad-hoc solutions are sometimes used to leverage metrics designed for two classes for the multi-class case, these solutions are far from ideal since they map the multi-class classification problem into a set of binary problems (Grandini et al., 2020). The resulting metrics do not

reflect the performance as will be perceived by the final user of the system. Instead, for those cases, we can use the EC as long as we can define a single cost matrix of interest. Again, in cases in which we would like to assess performance for a range of different cost matrices assuming Bayes decision theory will be used for making categorical decisions, expected SPSRs should be used as metric.

D Simulated dataset for experiments

Here we describe the procedure used to create simulated datasets for the experiments in this paper. Given a number of classes K , a prior for class H_1 of P_1 , and a per-class standard deviation σ , which are taken as parameters of the simulation, we proceed as follows:

1. Set the class priors to P_1 for class H_1 , and $P_i = (1 - P_1)/(K - 1)$ for classes 2 through K . In our experiments, P_1 is set to 0.8.
2. Set the total number of samples, $N = 100,000$ and determine the number of samples for each class, N_i as the closest integer to NP_i .
3. Generate N_i samples using a unidimensional Gaussian distribution $\mathcal{N}(\mu_i, \sigma^2)$, with mean $\mu_i = i - 1$ and standard deviation σ . We take this to be the x_t , the input features for each sample.
4. Compute the likelihoods for each class for each generated sample according to the class distributions used to draw these samples, i.e., $p(x|H_i) \sim \mathcal{N}(\mu_i, \sigma^2)$.
5. Finally, compute the posteriors as:

$$P(H_i|x) = \frac{p(x|H_i) P_i}{p(x)} = \frac{p(x|H_i) P_i}{\sum_j p(x|H_j) P_j}. \quad (30)$$

where the $p(x|h)$ are the likelihoods computed in step 4 and P_i is the prior for class H_i defined in step 1. With this procedure, the posteriors are perfectly calibrated for the test data.

The value of σ is set to 0.20 for all results in this work, except for those in Figure 2 where two datasets are used, one with $\sigma = 0.20$ and an easier one with $\sigma = 0.05$.