

The VOROS: Lifting ROC Curves to 3D to Summarize Unbalanced Classifier Performance

Christopher Ratigan¹, Lenore Cowen^{1,2}

¹ Department of Mathematics, Tufts University, Medford, MA 02155 USA

² Department of Computer Science, Tufts University, Medford, MA 02155 USA
christopher.ratigan@tufts.edu, cowen@cs.tufts.edu

Abstract

While the area under the ROC curve is perhaps the most common measure that is used to rank the relative performance of different binary classifiers, longstanding field folklore has noted that it can be a measure that ill-captures the benefits of different classifiers when either the actual class values or misclassification costs are highly unbalanced between the two classes. We introduce a new ROC surface, and the VOROS, a volume over this ROC surface, as a natural way to capture these costs, by lifting the ROC curve to 3D. Compared to previous attempts to generalize the ROC curve, our formulation also provides a simple and intuitive way to model the scenario when only ranges, rather than exact values, are known for possible class imbalance and misclassification costs.

Code — <https://github.com/ChristopherRatigan/VOROS>

Introduction

The ROC curve, constructed by plotting the False Positive Rate (FPR) against the True Positive Rate (TPR), is the most canonical way to measure the performance of a binary classifier.

Definition 1. *Given an ROC curve $y = f(x)$, the area under the ROC curve, henceforth AUROC, is defined by*

$$\int_0^1 f(x) dx$$

As has long been known (Hand 2009), automatically choosing the binary classifier with the higher AUROC makes perfect sense when two rather strong conditions are met: 1) the expected ratio of positives to negatives in the dataset where the classifier is to be deployed is balanced or nearly balanced, and 2) the cost of misclassifying a positive example as negative is equal or nearly equal to the cost of misclassifying a negative example as positive, i.e. that $C_0/C_1 \approx 1$ and $|\mathcal{P}|/|\mathcal{N}| \approx 1$, where we denote by \mathcal{P} the set of observations that are positive and \mathcal{N} denotes the set of observations that are negative. C_0 and C_1 denote the cost of misclassifying a case from \mathcal{N} as belonging to \mathcal{P} and the converse, respectively. The area under the ROC curve of a binary classifier \mathcal{F} represents the probability that given two observations $\rho \in \mathcal{P}$

and $\eta \in \mathcal{N}$, \mathcal{F} assigns a higher score to ρ than it does to η (Bradley 1997).

But, if the costs or classes are not balanced, then it is not always sensible to directly compare the AUROCs in this pairwise manner. If one class has many more cases than the other, comparing a random positive and negative example is unlikely to occur in practice. Likewise, if the two types of errors differ in individual costs then it may not be as important that it is possible to separate a randomly chosen positive example from a randomly chosen negative example.

Observe that when exact information is available on the ratios C_0/C_1 and $|\mathcal{P}|/|\mathcal{N}|$, it is always possible to find a binary classifier (point) taken from the ROC curve of classifier \mathcal{F}_1 that minimizes the total misclassification cost, and a binary classifier taken from the ROC curve of classifier \mathcal{F}_2 that minimizes the total misclassification cost, and then choose whichever classifier has lower total misclassification cost. However, as we (see example datasets, below) and others demonstrate (Provost, Fawcett, and Kohavi 1998; Hand 2009; Drummond and Holte 2000)), such a cost comparison no longer corresponds to a comparison of the AUROCs. In fact, classifiers with better AUROCs can perform much worse in total misclassification cost in these settings.

Much previous work has attempted to deal with the shortcomings of the AUROC by looking at different notions of a weighted AUROC (McClish 1989; Maurer and Pontil 2020; Li and Fine 2010); however, most of these formulations require fixing the cost distribution and class imbalance as a prior. (Two notable exceptions are (Hernández-Orallo, Flach, and Ramirez 2011) and the recent paper of (Shao et al. 2024), that takes an interesting (and complementary approach) to solving the problem we address in this paper, namely simultaneously representing different weighted costs)). Versions of weighted AUROC that apply for a fixed cost distribution and class imbalance do not easily extend to model uncertainty in the cost distribution or only bounds on the class imbalance (though we note that a previous paper of Guo et al. did define a “3D ROC histogram,” with a general 3rd axis that could represent either classifier confidence or cost; where the cost version is closest in spirit to the present work (Guo, Shen, and Zhang 2019)).

By instead making something similar to the weighted AUROC a slice of a ROC volume, lifted to 3D, we are able to model *ranges* of misclassification costs and class imbalance.

We introduce a new ROC surface, where the volume integrated over this ROC surface becomes a generalization of the AUROC where misclassification costs are also naturally represented. Besides mathematical elegance, the advantage of our VOROS to previous weighted AUROC approaches is that we also have a very natural way to ask for the better classifier in a regime where the application domain gives only bounds on the expected class imbalances or misclassification costs, and does not assume these values are known exactly. We show that the VOROS is efficient to compute, and show in several benchmark datasets, that it will choose better classifiers, according to how costs and class imbalances are estimated to occur in the environment in which the classifier will be deployed.

We demonstrate the utility of this approach on two standard benchmark classification datasets, the UCI Wisconsin Breast Cancer Dataset (William et al. 1995) and BUSI (Al-Dhabyani et al. 2020), as well as a highly imbalanced credit card fraud dataset.

Background and Definitions

Definition 2 (ROC Space). Fix a dataset \mathcal{X} consisting of observations assigned the actual class labels positive (1) and negative (0). For a binary classifier \mathcal{F} , the performance of \mathcal{F} on \mathcal{X} is represented as a point in ROC space where \mathcal{F} is mapped to the point (x, y) with x the false positive rate and y the true positive rate of \mathcal{F} on the dataset \mathcal{X} .

An intuitive rule in ROC space is that one classifier does better than another classifier if it is above and to the left as the following definition makes precise.

Definition 3. A point (x_1, y_1) dominates another point (x_2, y_2) in ROC space if it has a lower false positive rate and higher true positive rate, i.e. if $x_1 \leq x_2$ and $y_1 \geq y_2$.

There are a few points in ROC space that are important for our analysis. A perfect classifier, which correctly classifies all observation in \mathcal{X} has ROC coordinates $(0, 1)$. Likewise, the point $(1, 0)$ represents a model which misclassifies every case. We complete the unit square in ROC space by defining the baseline classifiers.

Definition 4. The baseline classifiers \mathcal{B}_0 and \mathcal{B}_1 are the classification models that label all observations as negative or positive respectively.

There is a traditional way of going from points in ROC space to a curve for classifiers that can output a numerical score interpreted as the probability that an observation belongs to the positive class. Simply vary a threshold s which divides the scores of the observations into positive and negative classes based on whether the score of a particular example is greater than or less than s (Metz 1978; Swets 1988).

Definition 5. Given a classifier $\mathcal{F} : \mathcal{X} \rightarrow [0, 1]$, let $\delta_{\mathcal{N}}(s)$ be the cumulative distribution of observations in the negative class and $\delta_{\mathcal{P}}(s)$ be that of the positive class. Then, the observed ROC curve for \mathcal{F} is a plot of $(1 - \delta_{\mathcal{N}}(s), 1 - \delta_{\mathcal{P}}(s))$.

ROC curves inherit the following natural notion of dominance from the domination of individual points.

Definition 6. Given two ROC curves f and g , f dominates g if for every point P on the graph of g there is some point Q on the graph of f such that Q dominates P and no point on the graph of f is strictly dominated by any point on the graph of g .

Definition 7. The Upper Convex Hull of a collection of n points $\mathcal{C} = \{(x_i, y_i)\}_{i=1}^n$, in ROC space is the boundary of the convex hull of \mathcal{C} together with the additional points $(0, 0)$, $(1, 1)$ and $(1, 0)$ running from $(1, 1)$ to $(0, 0)$ oriented counterclockwise.

While the definition of ROC curve intuitively creates a continuous curve in ROC space for continuous distributions $\delta_{\mathcal{N}}$ and $\delta_{\mathcal{P}}$, it is common to take the discrete set of classifiers given by varying the threshold parameter on a finite dataset, connect them piecewise by lines, and call the resulting graph a “curve” (Drummond and Holte 2006). In what follows, we will focus on this discretized case, though much of this generalizes to continuous “fitted” ROC curves.

We note that for such piecewise linear ROC curves f and g , the vertices of f dominate the vertices of g if and only if the graph of f lies above the graph of g over the entire interval $[0, 1]$. In this way, domination in ROC space is equivalent to the usual domination of real valued functions.

Lemma 8. The Upper Convex Hull of an ROC curve dominates the ROC curve.

Proof. A full proof is given in (Provost and Fawcett 2001). \square

It has long been known that in ROC space, it is possible to compare the expected costs of classification models to decide which point on an ROC curve, or a collection of such curves is optimal for a given weighting of the relative numbers and costs of false positives and false negatives. The following definition helps to quantify this notion of cost.

Definition 9. Let $\mathcal{X} = \mathcal{P} \sqcup \mathcal{N}$ be the set of observations to be classified, where \mathcal{P} denotes observations whose actual label is positive and likewise, \mathcal{N} denotes negative observations. We define C_0 and C_1 to be the costs of an individual false positive and false negative respectively. Additionally, given a point (x, y) in ROC space, we define the expected cost of (x, y) to be

$$C_0|\mathcal{N}|x + C_1|\mathcal{P}|(1 - y)$$

The point in ROC space with the maximal expected cost is $(1, 0)$ which misclassifies every case. We can scale our cost by this maximum so that costs lie in the unit interval putting the costs of classifiers on a normalized scale.

Definition 10. The normalized expected cost of a point $(x, y) = (FPR, TPR)$ is given by

$$Cost(x, y, t) = tx + (1 - t)(1 - y)$$

where $t = \frac{C_0|\mathcal{N}|}{C_0|\mathcal{N}| + C_1|\mathcal{P}|}$ is the portion of cost borne by the false positives.

This formulation of cost is well-known, though there has not been a universally accepted way to apply it to ROC curves in their entirety (Provost, Fawcett, and Kohavi 1998).

Lemma 11. Let $P(x_1, y_1)$ and $Q(x_2, y_2)$ be two points in ROC space and let $t \in [0, 1]$ be fixed. Then, $\text{Cost}(x_1, y_1, t) = \text{Cost}(x_2, y_2, t)$ if and only if the slope of the line between P and Q is $\frac{t}{1-t}$.

Proof. The relevant equation is

$$tx_1 + (1-t)(1-y_1) = tx_2 + (1-t)(1-y_2)$$

Rearranging yields

$$t(x_1 - x_2) = (1-t)(y_1 - y_2)$$

Dividing by $(1-t)(x_1 - x_2)$ on both sides yields the desired result

$$m = \frac{t}{1-t} = \frac{y_1 - y_2}{x_1 - x_2}$$

Where if $t = 1$, we let $m = \infty$ define a vertical line. \square

Definition 12. Given a fixed value of $t \in [0, 1]$ an iso-performance line through (x_1, y_1) is the line through (x_1, y_1) in ROC space with slope $\frac{t}{1-t}$ representing all points with the same cost as (x_1, y_1) .

Intuitively, iso-performance lines are the equivalence classes of points in ROC space with the same cost for a fixed value of t . These lines instill a linear order which divides a classifier from better performing classifiers above and to the left of the line (but not necessarily the point) and worse performing classifiers below and to the right of the iso-performance line.

Lemma 13. Let $t \in [0, 1]$ be fixed and let $y = f(x)$ be an ROC curve. A point $(h, f(h))$ has the minimum cost of any point on the ROC curve if the upper convex hull of the curve lies below and to the right of the iso-performance line through $(h, f(h))$ with slope $m = \frac{t}{1-t}$.

Proof. Let $P(h, f(h))$ be such as point, then points with a lower cost than P lie above the line through P with slope $\frac{t}{1-t}$, but if P achieves the minimal cost of any point on the ROC curve, all points on the ROC curve must lie below and to the right of the iso-performance line as required. \square

The term for these extremal iso-performance lines in convex geometry is “supporting lines” and the space of all such lines is the ROC Surface as explained below.

Definition 14. Given a region R in the plane, a supporting line ℓ of R is a line intersecting R such that all of R lies on the same side of ℓ .

Definition 15. Let $t \in [0, 1]$, an optimal point on an ROC curve is a point on the upper convex hull that has a supporting line of slope $\frac{t}{1-t}$.

Note, it is possible for an ROC curve to have multiple optimal points. Such points occur when the upper convex hull contains a line segment, any point of which is optimal for t such that $\frac{t}{1-t}$ is the slope of the segment.

Not only can supporting lines be used to pick out the optimal classifier for a given cost parameter t , they also partition ROC space according to cost. The area of classifiers which cost more provides a natural ordering which is explicitly related to the expected cost. This stands in contrast to the area under the ROC curve defined in the introduction.

Main Contribution

To connect our notion back to the area under the ROC curve, we now generalize this notion to a cost-aware space by defining the area of lesser classifiers.

Definition 16. Given $t \in [0, 1]$, and a point (h, k) in ROC space, the area of lesser classifiers of (h, k) , dependent on t , denoted $A_t(h, k)$ is the area of the points (h', k') in ROC space, such that (h', k') has a greater expected cost than (h, k) . Define such (h', k') to be lesser classifiers to (h, k) in this case.

Note if we choose (h, k) to be an optimal point on our curve, then the region containing lesser classifiers will always contain the area under the ROC curve, but it also typically includes additional area in the region above the ROC curve. As we see with the Credit Card Fraud dataset analyzed below, a classifier with a worse AUROC can have a better area of lesser classifiers value for ranges of unbalanced classification class sizes or misclassification costs.

We are now ready to define the ROC surface, which will allow us to look at the area of lesser classifiers in particular ranges of class imbalance or cost imbalance.

Definition 17. The ROC surface, ROS, over a point (h, k) in ROC space is the surface, in (x, y, t) -space, given by

$$y = \frac{t}{1-t}(x - h) + k$$

where $t \in [0, 1]$

The ROS parameterizes all lines through (h, k) with non-negative (possibly infinite) slope.

Lemma 18. The ROC surface over a point (h, k) is a saddle surface with the equation

$$t = \frac{Y}{X + Y}$$

where $Y = y - k$ and $X = x - h$ are simply horizontal shifts of the ROC surface over the origin.

Proof. The proof follows by solving for t in the formula for the ROC surface. \square

The ROC surface provides a natural coordinate system for assessing relative costs over a range of parameter values t since if t is a fixed constant, then for the ROC points (x_1, y_1) and (x_2, y_2) , we need only compare the parallel lines

$$y = \frac{t}{1-t}(x - x_1) + y_1 \text{ and } y = \frac{t}{1-t}(x - x_2) + y_2$$

In particular, setting $x = 0$ yields $y_i - \frac{t}{1-t}x_i$. The classifier with the lower cost at this threshold then has a higher value of $y_i - \frac{t}{1-t}x_i$.

Definition 19. The ROC Surface associated to an ROC curve $y = f(x)$, is the surface in (x, y, t) -space which for each value of t is the maximum of the lines in ROC space through points $(p, f(p))$ with slope $\frac{t}{1-t}$.

This surface divides the unit cube, $[0, 1]^3$, into two sets. When viewed from the vertical line $x = 0, y = 1$, points $P(x, y, t)$ behind the ROC surface map to points (x, y) in ROC space such that the cost of (x, y) is greater than that of the optimal point on the ROC curve at that t value. Similarly, points $P(x, y, t)$ in front of the ROC surface represent points which perform better than the optimal point on the ROC curve at the associated t value.

This gives a natural notion of performance for the whole ROC curve: the volume over the ROC surface.

Definition 20. The volume over the ROC Surface, *VOROS*, is defined by

$$VOROS(f(x)) = \int_0^1 \max_{x \in [0,1]} (A_t(x, f(x))) dt$$

Where $\max(A_t(x, f(x)))$ is the maximal area of lesser classifiers for any point on the ROC curve taken for each t .

One of the useful properties of *VOROS* is that it can be readily modified to handle ranges and distributions over the relative costs of misclassifications. This is useful because in many applications, it is reasonable to assume one might have bounds on the expected cost imbalances, but not know them exactly. We can capture this uncertainty in our *VOROS* measure as follows:

Definition 21. Let μ be a probability measure for $t \in [0, 1]$, let $[a, b] \subset [0, 1]$ and let $A_t(x, f(x))$ be the area of lesser classifiers associated to a point on the ROC curve $y = f(x)$, then

$$\begin{aligned} VOROS(f(x), [a, b], \mu) \\ = \frac{1}{\mu([a, b])} \int \max_{x \in [0,1]} (A_t(x, f(x))) d\mu(t) \end{aligned}$$

This is a generalization of our initial definition of *VOROS* (Definition 20) since $VOROS(f(x), [0, 1], \mu)$ is simply $VOROS(f(x))$ if μ is the standard measure on \mathbb{R} . In what follows we omit μ assuming it to be the standard measure on \mathbb{R} .

Lemma 22. Given an ROC point (h, k) , and fixed costs $(t, 1 - t)$, if (h, k) outperforms both baseline classifiers, then the area of lesser classifiers for (h, k) , $A_t(h, k)$, has the following form.

$$1 + \frac{(1-k)^2}{2} + \frac{h^2}{2} - h(1-k) - \frac{(1-k)^2}{2t} - \frac{h^2}{2(1-t)}$$




Figure 1: The area of Lesser Classifiers for the point (h, k) lies below the iso-performance line with slope $m = \frac{t}{1-t}$.

Proof. Let $m = \frac{t}{1-t}$ be the slope of the iso-performance line. First, note that since (h, k) outperforms the baseline classifiers, the area of lesser classifiers will be of the general form of the shaded region in Figure 1.

This is because the line bounding this region is above and to the left of the parallel lines through $(0, 0)$ and $(1, 1)$. Thus, the dividing line for lesser classifiers must intersect the y -axis above $(0, 0)$ and intersect the line $y = 1$ to the left of $(1, 1)$.

As such, the area of this region is given by one minus the area of the white triangle. The horizontal base of the triangle is $B = \frac{1-k}{m} + h$ while the height of the triangle is $H = 1 - k + mh$. Hence, the area of the shaded region is

$$\begin{aligned} A_t(h, k) &= 1 - \frac{1}{2}BH \\ &= 1 - \frac{1}{2} \left(\frac{1-k}{m} + h \right) (1 - k + mh) \\ &= 1 - \frac{1}{2} \left(\frac{(1-k)^2}{m} + 2(1-k)h + mh^2 \right) \\ &= 1 - \frac{(1-k)^2}{2m} - h(1-k) - \frac{mh^2}{2} \end{aligned}$$

Since $m = \frac{t}{1-t} = -1 + \frac{1}{1-t}$, the area of lesser classifiers is then

$$1 + \frac{(1-k)^2}{2} + \frac{h^2}{2} - h(1-k) - \frac{(1-k)^2}{2t} - \frac{h^2}{2(1-t)}$$

□

Integrating this last expression yields the following result

Lemma 23. Given a vertex $v = (h, k)$ from the upper convex hull of an ROC curve $\{x_i, y_i\}_{i=0}^n$. If v is an optimal point on the cost interval $[a, b] \subset [0, 1]$, then, the volume of lesser classifiers for the ROC curve on the interval (a, b) is given by

$$\frac{1}{b-a} \left(1 + \frac{(1-k)^2}{2} + \frac{h^2}{2} - h(1-k) \right) + \frac{(1-k)^2}{2} \ln \left(\frac{1-b}{1-a} \right) - \frac{h^2}{2} \ln \left(\frac{b}{a} \right)$$

Note that Lemma 23 gives an explicit way to compute the VOROS between points in the upper convex hull of ROC space, which will contain at most n points. Thus the VOROS can be calculated explicitly in time on the order of the time it takes to calculate the convex hull of n points. In our implementation we use the Quickhull algorithm (Bykat 1978; Eddy 1977) which runs in $O(n \log n)$ expected time.

There is a direct relationship between the minimum normalized expected cost of a classifier and the volume integrated over the ROC surface.

Theorem 24. *Let (h, k) be the optimal point on an ROC curve for $t \in [a, b]$. If $c_t(h, k)$ is the normalized expected cost of this point, then the volume over the ROC surface is simply*

$$\frac{1}{b-a} \int_a^b 1 - \frac{(c_t(h, k))^2}{2t(1-t)} dt$$

Proof. Since the volume over the ROC surface is simply the integral of the area bounded by a classifier at each threshold, all we need to show is that for each value of t , we have the formula

$$A_t(h, k) = 1 - \frac{c_t(h, k)^2}{2t(1-t)}$$

From earlier, we know that the left hand side is

$$1 + \frac{(1-k)^2}{2} + \frac{h^2}{2} - h(1-k) + \frac{(1-k)^2}{2t} - \frac{h^2}{2(1-t)}$$

The right hand side is simply,

$$\begin{aligned} 1 - \frac{c_t(h, k)^2}{2t(1-t)} &= 1 - \frac{(th + (1-t)(1-k))^2}{2t(1-t)} \\ &= 1 - \frac{t^2h^2 + 2t(1-t)h(1-k) + (1-t)^2(1-k)^2}{2t(1-t)} \\ &= 1 - \left(\frac{th^2}{2(1-t)} + h(1-k) + \frac{(1-t)(1-k)^2}{2t} \right) \\ &= 1 - \left(\frac{h^2}{2(1-t)} + h(1-k) + \frac{(1-k)^2}{2t} - \frac{(1-k)^2}{2} \right) \end{aligned}$$

where in the last equation we used that $\frac{t}{1-t} = -1 + \frac{1}{1-t}$. From here distributing the negative and rearranging terms yields the desired result. \square




Figure 2: The two possibilities for the area of lesser classifiers for the better of the baselines.




Figure 3: (Left) Graph of $A_t(\{(0,0), (1,1)\})$. (Right) The Volume bounded by these areas.

This theorem shows an important feature of the VOROS. Given two ROC Surfaces S_1 and S_2 associated to ROC curves f_1 and f_2 , if for all $t \in [a, b] \subseteq [0, 1]$ we have that S_1 has a lower cost than S_2 , then $\text{VOROS}(f_1, [a, b], \mu) > \text{VOROS}(f_2, [a, b], \mu)$ since the associated integrands have the same inequality.

There is a natural probabilistic interpretation of the VOROS: it is the probability that choosing $x, y \in [0, 1]$ uniformly and independently at random and t independently according to μ , the randomly chosen point (x, y) has a higher cost than an optimal point of the ROC curve. This is particularly useful for the problem of comparing performance of classifiers in cost or class imbalanced problems because if we can estimate the cost or class imbalance, we get a more accurate measure of classifier performance than from the AUROC alone.

For example, the AUROC of the hybrid classifier consisting of the two baselines \mathcal{B}_0 and \mathcal{B}_1 is $\frac{1}{2}$, but picking the better of the baseline classifiers will always yield a classification model which has less than $\frac{1}{2}$ the maximal expected cost so long as false positives and false negatives are unequal in aggregate cost.

Example 25. *Consider the ROC curve consisting of the vertices $\{(0,0), (1,1)\}$. The Volume over the ROC surface for this curve is $\frac{3}{2} - \ln(2) \approx .807$.*

Set $m = \frac{t}{1-t}$. Then picking the better classifier gives us either of the areas in Figure 2. Either way, the area of lesser classifiers is one minus the area of the white triangle.

$$A_t = 1 - \frac{1}{2}BH = \begin{cases} 1 - \frac{1}{2m} & m > \frac{1}{2} \\ 1 - \frac{m}{2} & m < \frac{1}{2} \end{cases}$$

	VOROS([0,1])	AUROC
Log	99.8%	99.3%
Forest	99.8%	99.0%
Naive	98.4%	95.1%
Baseline	80.7%	50%

Table 1: Table of VOROS and AUROC values for classifiers trained on the Wisconsin Breast Cancer dataset.




Figure 4: (a) A plot of the ROC curves for the Wisconsin Breast Cancer dataset. (b) A plot of the ROC curves for the BUSI dataset. The solid curves are from Logistic Regression, the dashed curves are from Random Forests and the dotted curves are from Naive Bayes.

Substituting the expression $m = \frac{t}{1-t}$ gives us the integral.

$$\int_0^1 A_t dt = \int_0^{\frac{1}{2}} 1 - \frac{t}{2(1-t)} dt + \int_{\frac{1}{2}}^1 1 - \frac{(1-t)}{2t} dt$$

A graph of A_t ($\{(0,0), (1,1)\}$) as a function of t next to the volume as a region in (x, y, t) -space is shown in Figure 3. The minimum value of the graph of A_t is 0.5, which occurs when the positive and negative classes have equal aggregate cost, i.e. $t = 1 - t = \frac{1}{2}$. The value of VOROS as a performance measure stems from the way in which the graph of A_t approaches 1 on both ends of the interval $[0, 1]$: The more imbalanced the classification problem, the less room there is to improve over the better of the two baselines.

Application to the Wisconsin Breast Cancer Dataset

We first show the application of the VOROS to the Wisconsin Breast Cancer data (William et al. 1995). This is a standard binary classification benchmark dataset with 357 positive observations and 212 negative examples. We directly apply logistic regression, naive Bayes and random forest classifiers (all using sklearn default parameters) to a stratified train/test split. Table 1 summarizes the overall AUROC and VOROS of each classifier without any known bounds on the aggregate cost ratio on the test set, where Figure 4(a) gives the AUROCs and Figure 5 gives each VOROS.

The area under the ROC curve is always less than that of the corresponding volume over the ROC surface. While naively, it may be tempting to claim that both VOROS and AUROC rank the classifiers as Log > Forest > Naive, the situation is a bit more subtle than this as shown in Figure 4.




Figure 5: (a) The VOROS for the Logistic Regression classifier. (b) The VOROS for the Bayes classifier. (c) The VOROS for the Random Forest classifier. Built on the Wisconsin Breast Cancer dataset.

	VOROS([0,0.25])	VOROS(0.75,1)
Log	99.8%	99.9%
Forest	99.9%	99.6%
Baseline	92.5%	92.5%

Table 2: Table summarizing the VOROS for two different cost ranges on the Wisconsin Breast Cancer dataset.

Here, we see that our logistic regression and random forest classifier each dominate the naive Bayes model in ROC space, and both logistic regression and random forest do better than naive Bayes for all choices of cost/class imbalance. However, logistic regression and random forest cannot be ranked so definitively since the ROC curves cross (Hand and Till 2001). Which classifier does best will depend on the costs.

Thus, in ROC space neither curve dominates the other. This can be addressed by VOROS if we are given bounds on the parameter t . For example, if we take t to be in the range $0 \leq t \leq 0.25$, then we get that random forest outperforms logistic regression. On the other hand, if t were in the range $0.75 \leq t \leq 1$, then we get the opposite ranking as shown in Table 2. Recall that these percentages represent the portion of the ROC points which are lesser classifiers in this range of cost values. For this range, picking the better of the two baselines already outperforms over 92% of ROC points.

Application to BUSI

We also consider a more modern breast cancer dataset, the Breast UltraSound Images dataset, henceforth BUSI (Al-Dhabyani et al. 2020). This dataset consists of full ultrasound breast images, from 600 women patients. It is actually a 3-class dataset. So, to make it a binary classification task, we consider both “normal” and “benign” to be the negative class and “malignant” to be the positive class.

We encode the data using Pytorch’s default vision trans-

	VOROS([0,1])	AUROC
Log	97.2%	91.8%
Forest	97.5%	93.3%
Naive	95.4%	88.3%
Baseline	80.7%	50%

Table 3: Table summarizing the VOROS and AUROC values for classifiers trained on the BUSI dataset.

	V([0,1/3])	V(1/3,2/3)	V(2/3,1)
Log	96.5%	96.7%	98.4%
Forest	98.0%	96.7%	98.0%
Naive	96.4%	93.5%	96.1%
Baseline	89.2%	63.7%	89.2%

Table 4: Table summarizing the VOROS over different ranges of costs for the BUSI dataset. Note we abbreviate VOROS to simply V due to space constrains.

former after cropping the images and scaling them to be the same size. The resulting numerical features are then used to fit logistic regression (with $C = 1000$ and $max_iter = 25$), naive Bayes (sklearn default), and random forest classifiers (sklearn default).

Table 3 summarizes the results. As we have seen before, the area under the ROC curve is less than that of the corresponding volume over the ROC surface. The ranking of non-baseline classifiers according to AUROC is Forest > Log > Naive, which is the same ranking for the VOROS over the entire interval $[0, 1]$, however, plotting the ROC curves, a different story emerges.

In Figure 4(b) all three curves intersect, hence no single model dominates, and which of these three models performs better will depend on the specific cost and class imbalances, as captured in the VOROS (see Figure 6, where Table 4 summarizes the VOROS for a few example choices of cost intervals).

Table 4 shows that the logistic regression model performs better for cases where the aggregate cost of false positives is much higher than that of false negatives. This can be confirmed by a careful inspection of Figure 4 (b) where the solid curve crosses paths with the dotted curve. There is a range of steeply sloped lines where the solid curve will perform better and these contribute to its better performance in the VOROS for the rightmost column above.

Overall, we see that how classifiers should be ranked depends strongly on the relative aggregate costs of false positives and false negatives on these data.

Application to a Credit Fraud Dataset

The previous two examples have both very good absolute performance of the classifiers considered, plus the class imbalance is not so extreme. The biggest strength of the VOROS, however, is when there is extreme imbalance in class membership and misclassification costs, which we demonstrate in an example credit card fraud dataset. We use a processed

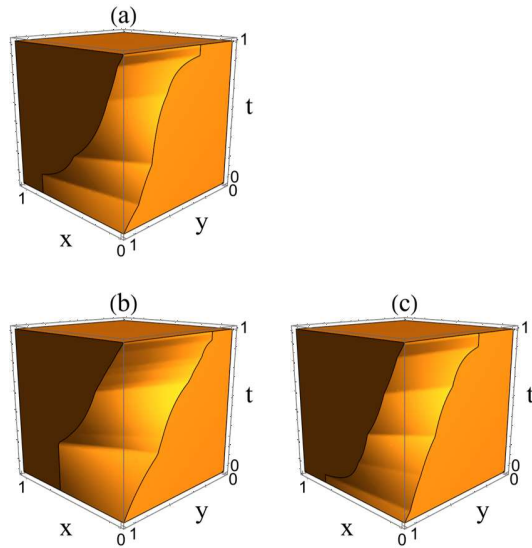


Figure 6: (a) The VOROS for the Logistic Regression classifier. (b) The VOROS for the Bayes classifier. (c) The VOROS for the Random Forest classifier. Built on the BUSI dataset.




Figure 7: A plot of the ROC curves for the Dal Pozzolo Credit Fraud dataset. The solid line connects the optimal points on each curve that have the same cost for a specific value of t .

dataset derived from a set of transactions made by credit cards in September 2013 by European cardholders over the course of two days. (Dal Pozzolo et al. 2015, 2014, 2017; Dal Pozzolo 2015; Carcillo et al. 2018a,b; Lebichot et al. 2020; Carcillo et al. 2021; Le Borgne et al. 2022; Lebichot et al. 2021) There are a total of 492 frauds (positive class) out of 284,807 transactions, so the fraud class accounts for only about .172% of the transactions. The dataset is published with derived features which are a set of principal components from PCA processing the original features of the dataset (since the original features could not be published for confidentiality concerns).

We build logistic regression models: one using just the first principal component (M_1) and the other using just the second principal component (M_2) to model whether transactions were fraudulent. This results in the ROC curves shown in Figure 7.

Since these curves intersect, which curve performs better

	$V([999/5999, 99/399])$	AUROC
M_1	90.1%	79.5%
M_2	89.6 %	85.3%
Base	86.9%	50%

Table 5: Table summarizing the VOROS vs. AUROC in the chosen cost range for the Dal Pozzolo Credit Fraud data. Note we abbreviate VOROS to simply V due to space constrains.

depends on the aggregate misclassification costs. Figure 7 shows the resulting ROC convex hulls together with a line showing the optimal points on each curve that have the same expected cost for a particular choice of aggregate class imbalance, Figure 8 shows each VOROS. Here, $t \approx 0.2165$ results in the these optimal points having the same normalized expected costs. This is the dividing threshold between M_1 performing better and M_2 performing better.

There is no guidance given in the dataset for the costs of a false positive versus a false negative, but if we assume that the cost of incorrectly flagging a true transaction as fraud is small (requiring more testing and verification, or customer annoyance to get a flagged card unfrozen) compared to the cost of missing fraud, then the misclassification cost of not detecting fraud could vary depending on both the number of fraudulent transactions and the dollar amount charged. Suppose we want to choose between these two classifiers, M_1 and M_2 , when we have only bounds on these values in the real world: for example, we assume that between 0.1% and 1% of transactions are fraudulent, and that the cost of the false negatives is between 500 and 5000 times the cost of a false positive. Using the AUROC, we would pick M_2 , but using the VOROS we find that $\frac{1-t}{t} = \frac{C_0|N|}{C_1|P|} = \frac{C_0}{C_1} \frac{|N|}{|P|}$ is the product of the cost and class ratios. Which satisfy the following bounds

$$99 \leq \frac{|N|}{|P|} \leq 999 \text{ and } \frac{1}{5000} \leq \frac{C_0}{C_1} \leq \frac{1}{500}$$

These bounds then translate into the following inequality

$$\frac{99}{500} \leq \frac{t}{1-t} \leq \frac{999}{5000}$$

Which since $t = \frac{\frac{t}{1-t}}{\frac{t}{1-t} + 1}$, induces the limits

$$\frac{99}{599} \leq t \leq \frac{999}{5999}$$

Taking the Volume over each ROC surface yields the results in Table 5. This VOROS comparison immediately shows us facts about classifier performance in our estimated interval that cannot be obtained from comparing the AUROCs. First, we learn that the baseline classifier, that flags *all* transactions as fraud actually is not doing so badly in this range by comparing its VOROS (86.9%) to the values for M_1 and M_2 (where the strength of a baseline classifier for problems with very rare or very costly classes is high, but is not immediately obvious from looking at ROC curves on their own. In this




Figure 8: The VOROS for logistic regression classifiers (a) using the first principal component, M_1 , and (b) the second principal component M_2 .

scenario, most ML papers shift to AUPRC or other measures of classifier performance as a result, so as to not recommend a classifier that is worse than a baseline). Second, comparing the AUROCs in Table 5, one would expect classifier M_2 to dominate, but in this cost range, M_1 is better, and since the VOROS is larger in this region, for this range we should pick classifier M_1 instead of either M_2 or the baseline. This agrees with a cost analysis which shows that if the slope of the supporting line to the dotted and dashed curves in Figure 7 are between $\frac{99}{500}$ and $\frac{999}{5000}$, then the supporting line for the dotted curve is higher than that of the dashed curve.

Discussion

We introduced the VOROS, a natural and easily-computable measure that generalizes the AUROC. We showed both theoretically and in example datasets that it has utility for comparing classifiers when classes or costs are unbalanced. We do note that an effective use of the VOROS is limited to domains where subject matter experts can put some bounds on expected misclassification costs. So far, our work handles only binary classification tasks. In fact, it is important to remark that the Volume *under* the ROC surface, or VUS, studied by (Ferri, Hernández-Orallo, and Salido 2003) and others (He and Frey 2008; Waegeman, De Baets, and Boullart 2008; Kang and Tian 2013) is an entirely different concept, generalizing the AUROC to multiple classes, rather than generalizing binary classifiers to different cost and class imbalances. Indeed, a generalization our VOROS work to also handle multiple classes, perhaps by increasing the dimension of the surface, remains a possible topic of future research, as is a generalization to incorporate varied misclassification costs by instance as in (Fawcett 2006).

Ethics Statement

In many application domains (including breast cancer screening), there are ethical implications for how one measures the cost of a false positive versus a false negative. Once the costs, including ethical costs, of misclassification have been considered by experts, we can analyze the results in our framework.

References

Al-Dhabyani, W.; Gomaa, M.; Khaled, H.; and Fahmy, A. 2020. Dataset of breast ultrasound images. *Data in Brief*, 28:

104863.

- Bradley, A. P. 1997. The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7): 1145–1159.
- Bykat, A. 1978. Convex hull of a finite set of points in two dimensions. *Information Processing Letters*, 7(6): 296–298.
- Carcillo, F.; Dal Pozzolo, A.; Le Borgne, Y.-A.; Caelen, O.; Mazzer, Y.; and Bontempi, G. 2018a. SCARFF: a scalable framework for streaming credit card fraud detection with spark. *Information Fusion*, 41: 182–194.
- Carcillo, F.; Le Borgne, Y.-A.; Caelen, O.; and Bontempi, G. 2018b. Streaming active learning strategies for real-life credit card fraud detection: assessment and visualization. *International J. of Data Science and Analytics*, 5: 285–300.
- Carcillo, F.; Le Borgne, Y.-A.; Caelen, O.; Kessaci, Y.; Oblé, F.; and Bontempi, G. 2021. Combining unsupervised and supervised learning in credit card fraud detection. *Information sciences*, 557: 317–331.
- Dal Pozzolo, A. 2015. *Adaptive machine learning for credit card fraud detection*. Ph.d. thesis, Université libre de Bruxelles.
- Dal Pozzolo, A.; Boracchi, G.; Caelen, O.; Alippi, C.; and Bontempi, G. 2017. Credit card fraud detection: a realistic modeling and a novel learning strategy. *IEEE Transactions on Neural Networks and Learning Systems*, 29(8): 3784–3797.
- Dal Pozzolo, A.; Caelen, O.; Johnson, R. A.; and Bontempi, G. 2015. Calibrating probability with undersampling for unbalanced classification. In *2015 IEEE Symposium Series on Computational Intelligence*, 159–166. IEEE.
- Dal Pozzolo, A.; Caelen, O.; Le Borgne, Y.-A.; Waterschoot, S.; and Bontempi, G. 2014. Learned lessons in credit card fraud detection from a practitioner perspective. *Expert Systems with Applications*, 41(10): 4915–4928.
- Drummond, C.; and Holte, R. C. 2000. Explicitly representing expected cost: An alternative to ROC representation. In *Proceedings of the Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 198–207.
- Drummond, C.; and Holte, R. C. 2006. Cost curves: An improved method for visualizing classifier performance. *Machine Learning*, 65: 95–130.
- Eddy, W. F. 1977. A new convex hull algorithm for planar sets. *ACM Transactions on Mathematical Software (TOMS)*, 3(4): 398–403.
- Fawcett, T. 2006. ROC graphs with instance-varying costs. *Pattern Recognition Letters*, 27(8): 882–891.
- Ferri, C.; Hernández-Orallo, J.; and Salido, M. A. 2003. Volume under the ROC surface for multi-class problems. In *European Conference on Machine Learning*, 108–120. Springer.
- Guo, R.; Shen, X.; and Zhang, X. 2019. 3D ROC histogram: A new ROC analysis tool incorporating information on instances. *IEEE Access*, 7: 183396–183404.
- Hand, D. J. 2009. Measuring classifier performance: a coherent alternative to the area under the ROC curve. *Machine Learning*, 77(1): 103–123.
- Hand, D. J.; and Till, R. J. 2001. A simple generalisation of the area under the ROC curve for multiple class classification problems. *Machine Learning*, 45: 171–186.
- He, X.; and Frey, E. C. 2008. The meaning and use of the volume under a three-class ROC surface (VUS). *IEEE Transactions on Medical Imaging*, 27(5): 577–588.
- Hernández-Orallo, J.; Flach, P. A.; and Ramirez, C. F. 2011. Brier Curves: a New Cost-Based Visualisation of Classifier Performance. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, 585–592.
- Kang, L.; and Tian, L. 2013. Estimation of the volume under the ROC surface with three ordinal diagnostic categories. *Computational Statistics & Data Analysis*, 62: 39–51.
- Le Borgne, Y.-A.; Sibli, W.; Lebichot, B.; and Bontempi, G. 2022. Reproducible machine learning for credit card fraud detection—Practical handbook. *Université Libre de Bruxelles*.
- Lebichot, B.; Le Borgne, Y.-A.; He-Guelton, L.; Oblé, F.; and Bontempi, G. 2020. Deep-learning domain adaptation techniques for credit cards fraud detection. In *Recent Advances in Big Data and Deep Learning: Proceedings of the INNS Big Data and Deep Learning Conference INNSBDDL2019*, 78–88. Springer.
- Lebichot, B.; Paldino, G. M.; Sibli, W.; He-Guelton, L.; Oblé, F.; and Bontempi, G. 2021. Incremental learning strategies for credit cards fraud detection. *International Journal of Data Science and Analytics*, 12(2): 165–174.
- Li, J.; and Fine, J. P. 2010. Weighted area under the receiver operating characteristic curve and its application to gene selection. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 59(4): 673–692.
- Maurer, A.; and Pontil, M. 2020. Estimating weighted areas under the ROC curve. *Advances in Neural Information Processing Systems*, 33: 7733–7742.
- McClish, D. K. 1989. Analyzing a portion of the ROC curve. *Medical Decision Making*, 9(3): 190–195.
- Metz, C. E. 1978. Basic principles of ROC analysis. In *Seminars in Nuclear Medicine*, volume 8, 283–298. Elsevier.
- Provost, F.; and Fawcett, T. 2001. Robust classification for imprecise environments. *Machine learning*, 42: 203–231.
- Provost, F. J.; Fawcett, T.; and Kohavi, R. 1998. The case against accuracy estimation for comparing induction algorithms. In *Proceedings of the 15th International Conference on Machine Learning (ICML-98)*, 445–453.
- Shao, H.; Xu, Q.; Yang, Z.; Wen, P.; Peifeng, G.; and Huang, Q. 2024. Weighted ROC curve in cost space: Extending AUC to cost-sensitive learning. *Advances in Neural Information Processing Systems*, 36.
- Swets, J. A. 1988. Measuring the accuracy of diagnostic systems. *Science*, 240(4857): 1285–1293.
- Waegeman, W.; De Baets, B.; and Boullart, L. 2008. ROC analysis in ordinal regression learning. *Pattern Recognition Letters*, 29(1): 1–9.
- William, W.; Olvi, M.; Nick, S.; and W., S. 1995. Breast Cancer Wisconsin (Diagnostic). UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C5DW2B>.