
On the Interpretability of Conditional Probability Estimates in the Agnostic Setting

Yihan Gao

Aditya Parameswaran

Jian Peng

University of Illinois at Urbana-Champaign

Abstract

We study the interpretability of conditional probability estimates for binary classification under the agnostic setting or scenario. Under the agnostic setting, conditional probability estimates do not necessarily reflect the true conditional probabilities. Instead, they have a certain calibration property: among all data points that the classifier has predicted $\mathcal{P}(Y = 1|X) = p$, p portion of them actually have label $Y = 1$. For cost-sensitive decision problems, this calibration property provides adequate support for us to use Bayes Decision Theory. In this paper, we define a novel measure for the calibration property together with its empirical counterpart, and prove an uniform convergence result between them. This new measure enables us to formally justify the calibration property of conditional probability estimations, and provides new insights on the problem of estimating and calibrating conditional probabilities.

1 Introduction

Many binary classification algorithms, such as naive Bayes and logistic regression, naturally produce confidence measures in the form of conditional probability of labels. These confidence measures are usually interpreted as the conditional probability of the label $y = 1$ given the feature x . An important research question is how to justify these conditional probabilities, i.e., how to prove the trustworthiness of such results.

In classical statistics, this question is usually studied under the realizable assumption, which assumes that the true underlying probability distribution has the

same parametric form as the model assumption. More explicitly, statisticians usually construct a parametric conditional distribution $\mathcal{P}(Y|X, \theta)$, and assume that the true conditional distribution is also of this form (with unknown θ). The justification of conditional probabilities can then be achieved by using either hypothesis testing or confidence interval estimation on θ .

However, in modern data analysis workflows, the realizable assumption is often violated, e.g. data analysts usually try out several off-the-shelf classification algorithms to identify those that work the best. This setting is often called *agnostic* — essentially implying that we do not have any knowledge about the underlying distribution. Under the agnostic setting, conditional probability estimates can no longer be justified by standard statistical tools, as most hypothesis testing methods are designed to distinguish two parameter areas in the hypothesis space (e.g., $\theta < \theta_0$ v.s. $\theta \geq \theta_0$), and confidence intervals require realizable assumption to be interpretable.

In this paper, we study the interpretability of conditional probabilities in binary classification in the agnostic setting: what kind of guarantees can we have without making any assumption on the underlying distribution? Justifying these conditional probabilities is important for applications that explicitly utilize the conditional probability estimates of the labels, including medical diagnostic systems (Cooper, 1984) and fraud detection (Fawcett and Provost, 1997). In such applications, the misclassification loss function is often asymmetric (i.e., false positive and false negative incur different loss), and accurate conditional probability estimates are crucial empirically. In particular, in medical diagnostic systems, a false positive means additional tests are needed, while a false negative could potentially be fatal.

Summary of Notation

We focus on the binary classification problem in this paper. Let us first define some notations here that will be used throughout the paper:

Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS) 2017, Fort Lauderdale, Florida, USA. JMLR: W&CP volume 54. Copyright 2017 by the author(s).

- \mathcal{X} denotes the discrete feature space and $\mathcal{Y} = \{\pm 1\}$ denotes the label space.
- \mathcal{P} denotes the underlying distribution over $\mathcal{X} \times \mathcal{Y}$ that governs the generation of datasets.
- $D = \{(X_1, Y_1), \dots, (X_n, Y_n)\}$ denotes a set of i.i.d. data points from \mathcal{P} .
- A fuzzy classifier is a function from \mathcal{X} to $[0, 1]$ where the output denotes the estimated conditional probability of $\mathcal{P}(Y = 1|X)$.

Interpretations of Conditional Probability Estimates

Ideally, we hope that our conditional probability estimates can be interpreted as the true conditional probabilities. This interpretation is justified if we can prove that the conditional probability estimates are close to the true values. Let $l_1(f, \mathcal{P})$ be the l_1 distance between the true distribution and the estimated distribution as a measure of the ‘‘correctness’’ of conditional probability estimates:

$$l_1(f, \mathcal{P}) = \mathbb{E}_{X \sim \mathcal{P}} |f(X) - \mathcal{P}(Y = 1|X)|$$

Here X is a random variable representing the feature vector of a sample data point, Y is the label of X and $f(X)$ is a fuzzy classifier that estimates $\mathcal{P}(Y = 1|X)$. If we can prove that $l_1(f, \mathcal{P}) \leq \epsilon$ for some small ϵ , then the output of f can be approximately interpreted as the true conditional probability.

Unfortunately, as we will show in this paper, it is impossible to guarantee any reasonably small upper bound for $l_1(f, \mathcal{P})$ under the agnostic assumption. In fact, as we will demonstrate in this paper, for the cases where we have to make the agnostic assumption, the estimated conditional probabilities are usually no longer close to the true values in practice.

Therefore, instead of trying to bound the l_1 distance, we develop an alternative interpretation for these conditional probability estimates. We introduce the following calibration definition for fuzzy classifiers:

Definition 1. Let \mathcal{X} be the feature space, $\mathcal{Y} = \{\pm 1\}$ be the label space and \mathcal{P} be the distribution over $\mathcal{X} \times \mathcal{Y}$. Let $f : \mathcal{X} \rightarrow [0, 1]$ be a fuzzy classifier, then we say f is calibrated if for any $p_1 < p_2$, we have:

$$\mathbb{E}_{X \sim \mathcal{P}} [\mathbb{1}_{p_1 < f(X) \leq p_2} f(X)] = \mathcal{P}(p_1 < f(X) \leq p_2, Y = 1)$$

Intuitively, a fuzzy classifier is calibrated if its output correctly reflects the relative frequency of labels among instances they believe to be similar. For instance, suppose the classifier output $f(X) = p$ for n data points, then roughly there are np data points with label $Y = 1$. We also define a measure of how close f is to be calibrated:

Definition 2. A fuzzy classifier f is ϵ -calibrated if

$$c(f) = \sup_{p_1 < p_2} |\mathcal{P}(p_1 < f(X) \leq p_2, Y = 1) - \mathbb{E}_{X \sim \mathcal{P}} [\mathbb{1}_{p_1 < f(X) \leq p_2} f(X)]| \leq \epsilon$$

f is ϵ -empirically calibrated with respect to D if

$$c_{emp}(f, D) = \frac{1}{n} \sup_{p_1 < p_2} \left| \sum_{i=1}^n \mathbb{1}_{p_1 < f(X_i) \leq p_2, Y_i=1} - \sum_{i=1}^n \mathbb{1}_{p_1 < f(X_i) \leq p_2} f(X_i) \right| \leq \epsilon$$

where $D = \{(X_i, Y_i), \dots, (X_n, Y_n)\}$ is a size n dataset consisting of i.i.d. samples from \mathcal{P} .

Note that the empirical calibration measure $c_{emp}(f, D)$ can be efficiently computed on a finite dataset. We further prove that under certain conditions, $c_{emp}(f, D)$ converges uniformly to $c(f)$ over all functions f in a hypothesis class. Therefore, the calibration property of these classifiers can be demonstrated by showing that they are empirically calibrated on the training data.

The calibration definition is motivated by analyzing the properties of commonly used conditional probability estimation algorithms: many such algorithms will generate classifiers that are naturally calibrated. Our calibration property justifies the common practice of using calibrated conditional probability estimates as true conditional probabilities: we show that if the fuzzy classifier is calibrated and the output of the classifier is the only source of information, then the optimal strategy is to apply Bayes Decision Rule on the conditional probability estimates.

The uniform convergence result of $c_{emp}(f, D)$ and $c(f)$ has several applications. First, it can be directly used to prove a fuzzy classifier is (almost) calibrated, which is necessary for the conditional probability estimates to be interpretable. Second, it suggests that we need to minimize the empirical calibration measure to obtain calibrated classifiers, which is a new direction for designing conditional probability estimation algorithms. Finally, taking an uncalibrated conditional probability estimates as input, we can calibrate them by minimizing the calibration measure. In fact, one of the most well-known calibration algorithm, the isotonic regression algorithm, can be interpreted this way.

Paper Outline

The rest of this paper is organized as following. In Section 2, we argue that the l_1 distance cannot be provably bounded under the agnostic assumption (Theorem 1) and then motivate our calibration definition.

In Section 3 we present the uniform convergence result (Theorem 2) and discuss the potential applications. In Section 4, we conduct experiments to illustrate the behavior of our calibration measure on several common classification algorithms.

Related Work

Our definition of calibration is similar to the definition of calibration in prediction theory (Foster and Vohra, 1998), where the goal is also to make predicted probability values match the relative frequency of correct predictions. In prediction theory, the problem is formulated from a game-theoretic point of view: the sequence generator is assumed to be malevolent, and the goal is to design algorithms to achieve this calibration guarantee no matter what strategy the sequence generator uses.

To the best of our knowledge, there is no other work addressing the interpretability of conditional probability estimates in agnostic cases. Our definition of calibration is also connected to the problem of calibrating conditional probability estimates, which has been studied in many papers (Zadrozny and Elkan, 2002) (Platt, 1999).

2 The Calibration Definition: Motivation & Impossibility Result

2.1 Impossibility result for l_1 distance

Recall that the l_1 distance between f and \mathcal{P} is defined as:

$$l_1(f, \mathcal{P}) = \mathbb{E}_{X \sim \mathcal{P}} |f(X) - \mathcal{P}(Y = 1|X)|$$

Suppose f is our conditional probability estimator that we learned from the training dataset. We attempt to prove that the l_1 distance between f and \mathcal{P} is small under the agnostic setting. With the agnostic setting, we do not know anything about \mathcal{P} , and the only tool we can utilize is a validation dataset D_{val} that consists of i.i.d. samples from \mathcal{P} . Therefore, our best hope would be a prover $A_f(D)$ that:

- Returns 1 with high probability if $l_1(f, \mathcal{P})$ is small.
- Returns 0 with high probability if $l_1(f, \mathcal{P})$ is large.

The following theorem states that no such prover exists, and the proof can be found in the appendix.

Theorem 1. *Let \mathcal{Q} be a probability distribution over \mathcal{X} , and $f : \mathcal{X} \rightarrow [0, 1]$ be a fuzzy classifier. Define B_f as:*

$$B_f = \mathbb{E}_{X \sim \mathcal{Q}} \min(f(X), 1 - f(X))$$

If we have that $\forall x \in \mathcal{X}, \mathcal{Q}(x) < \frac{1}{10n^2}$, then there is no prover $A_f : \{\mathcal{X} \times \mathcal{Y}\}^n \rightarrow \{0, 1\}$ for f satisfying the following two conditions:

For any \mathcal{P} over $\mathcal{X} \times \mathcal{Y}$ such that $\mathcal{P}_X = \mathcal{Q}$ (i.e., $\forall x \in \mathcal{X}, \sum_{y \in \mathcal{Y}} \mathcal{P}(x, y) = \mathcal{Q}(x)$), suppose $D_{val} \in \{\mathcal{X} \times \mathcal{Y}\}^n$ is a validation dataset consisting of n i.i.d. samples from \mathcal{P} :

1. *If $l_1(f, \mathcal{P}) = 0$, then $\mathbf{P}_{D_{val}}(A_f(D_{val}) = 1) > \frac{2}{3}$.*
2. *If $l_1(f, \mathcal{P}) > \frac{B_f}{2}$, then $\mathbf{P}_{D_{val}}(A_f(D_{val}) = 1) < \frac{1}{3}$.*

We made the assumption in Theorem 1 to exclude the scenario where a significant amount of probability mass concentrates on a few data points so that their corresponding conditional probability can be estimated via repeated sampling. Note that the statement is not true in the extreme case where all probability mass concentrates on one single data point (i.e., $\exists x \in \mathcal{X}, \mathcal{Q}(x) = 1$). The assumption is true when the feature space \mathcal{X} is large enough such that it is almost impossible for any data point to have significant enough probability mass to get sampled more than once in the training dataset.

The significance of Theorem 1 is that any attempt to guarantee a small upper bound of $l_1(f, \mathcal{P})$ would definitely fail. Thus, we can no longer interpret the conditional probability estimates as the true conditional probabilities under the agnostic setting. This result motivates us to develop a new measure of “correctness” to justify the conditional probability estimates.

2.2 $l_1(f, \mathcal{P})$ in practice

The fact that we cannot guarantee an upper bound of the l_1 distance is not merely a theoretical artifact. In fact, in the cases where we need to make the agnostic assumption, the value of $l_1(f, \mathcal{P})$ is often very large in practice. Here we use the following document categorization example to demonstrate this point.

Example 1. *Denote Z to be the collection of all English words. In this problem the feature space $\mathcal{X} = Z^*$ is the collection of all possible word sequences, and \mathcal{Y} denotes whether this document belongs to a certain topic (say, football). Denote \mathcal{P} as the following data generation process: X is generated from the Latent Dirichlet Allocation model (Blei et al., 2003), and Y is chosen randomly according to the topic mixture.*

We use logistic regression, which is parameterized by a weight function $w : Z \rightarrow \mathbb{R}$, and two additional parameters a and b . For each document $X = z_1 z_2 \dots z_k$, the output of the classifier is:

$$f(X) = \frac{1}{1 + \exp(-a \sum_{i=1}^k w(z_i) - b)}$$

The reason that we are using automatically generated documents instead of true documents here is that

the conditional probabilities $P(Y|X)$ are directly computable (otherwise we cannot evaluate $l_1(f, \mathcal{P})$ and other measures). We conducted an experimental simulation for this example, and the experimental details can be found in the appendix. Here we summarize the major findings: the logistic regression classifier has very large l_1 error, which is probably due to the discrepancy between the logistic regression model and the underlying model. However, the logistic regression classifier is almost naturally calibrated in this example. This is not a coincidence, and we will discuss the corresponding intuition in Section 2.3.

2.3 The Motivation of the Calibration Measure

Let us revisit Example 1. This time, we fix the word weight function w . In this case, every document X can be represented using a single parameter $w(X) = \sum_i w(z_i)$, and we search for the optimal a and b such that the log-likelihood is maximized. This is illustrated in Figure 1.




Figure 1: Illustration of Example 1

Now, intuitively, to maximize the log-likelihood, we need the sigmoid function $\frac{1}{1+\exp(-aw(X)-b)}$ to match the conditional probability of Y conditioned on $w(X)$: $\mathcal{P}(Y = 1|w(X))$. Therefore, for the optimal a and b , we could say that the following property is roughly correct:

$$\mathcal{P}(Y = 1|w(X)) \approx \frac{1}{1 + \exp(-aw(X) - b)}$$

In other words,

$$\forall 0 \leq p \leq 1, \mathbb{E}[\mathcal{P}(Y = 1|X)|f(X) = p] \approx p$$

Let us examine this example more closely. The reason why the logistic regression classifier tells us that $f(X) \approx p$ is because of the following: among all the documents with similar weight $w(X)$, about p portion of them actually belong to the topic in the training dataset. This leads to an important observation: logistic regression estimates the conditional probabilities by computing the relative frequency of labels among documents it believes to be similar.

This behavior is not unique to logistic regression. Many other algorithms, including decision tree classifiers, nearest neighbor (NN) classifiers, and neural networks, exhibit similar behavior:

- In decision trees, all data points reaching the same decision leaf are considered similar.
- In NN classifiers, all data points with the same nearest neighbors are considered similar.
- In neural networks, all data points with the same output layer node values are considered similar.

We can abstract the above conditional probability estimators as the following two-step process:

1. Partition the feature space \mathcal{X} into several regions.
2. Estimate the relative frequency of labels among all data points inside each region.

The definition of the calibration property follows easily from the above two-step process. We can argue that the classifier is approximately calibrated, if for each region S in the feature space \mathcal{X} , the output conditional probability of data points in S is close to the actual relative frequency of labels in S . The definition for the calibration property then follows from the fact that all data points inside each region have the same output conditional probabilities.

$$\begin{aligned} \forall p_1 < p_2, \quad & \mathcal{P}(p_1 < f(X) \leq p_2, Y = 1) \\ & = \mathbb{E}_{X \sim \mathcal{P}}[\mathbb{1}_{p_1 < f(X) \leq p_2} f(X)] \end{aligned}$$

Using Calibrated Conditional Probabilities in Decision Making

The calibration property justifies the common practice of using estimated conditional probabilities in decision making. Consider the binary classification problem with asymmetric misclassification loss: we lose a points for every false positive and b points for every false negative. In this case, the best decision strategy is to predict 1 if $\mathcal{P}(Y = 1|X) \geq \frac{a}{a+b}$ and predict -1 otherwise. Now consider the case when we do not know $\mathcal{P}(Y = 1|X)$, but only know the value of $f(X)$ instead. If we can only use $f(X)$ to make decision, and f is calibrated, then the best strategy is to use $f(X)$ in the same way as $\mathcal{P}(Y = 1|X)$ (the proof can be found in the appendix):

Claim 1. *Suppose we are given a calibrated fuzzy classifier $f : \mathcal{X} \rightarrow [0, 1]$, we need to make decisions solely based on the output of f . Denote our decision as $g : [0, 1] \rightarrow \{\pm 1\}$ (i.e., our decision for X is $g(f(X))$).*

Then the optimal strategy g^* to minimize the expected loss is the following:

$$g^*(x) = \begin{cases} 1 & x \geq \frac{a}{a+b} \\ -1 & x < \frac{a}{a+b} \end{cases}$$

3 Uniform Convergence of the Calibration Measure

3.1 The Uniform Convergence Result

Let \mathcal{G} be a collection of functions from $\mathcal{X} \times \mathcal{Y}$ to $[0, 1]$, the Rademacher Complexity (Bartlett and Mendelson, 2003)¹ of \mathcal{G} with respect to D is defined as (Shalev-Shwartz and Ben-David, 2014):

$$R_D(\mathcal{G}) = \frac{1}{n} \mathbb{E}_{\sigma \sim \{\pm 1\}^n} \left[\sup_{g \in \mathcal{G}} \sum_{i=1}^n \sigma_i g(x_i, y_i) \right]$$

Then we have the following result:

Theorem 2. *Let \mathcal{F} be a set of fuzzy classifiers, i.e., functions from \mathcal{X} to $[0, 1]$. Let \mathcal{H} be the set of binary classifiers obtained by thresholding the output of fuzzy classifiers in \mathcal{F} :*

$$\mathcal{H} = \{ \mathbb{1}_{p_1 < f(x) \leq p_2} : p_1, p_2 \in \mathbb{R}, f \in \mathcal{F} \}$$

Suppose the Rademacher Complexity of \mathcal{H} satisfies:

$$\mathbb{E}_D R_D(\mathcal{H}) + \sqrt{\frac{2 \ln(8/\delta)}{n}} < \frac{\epsilon}{2}$$

Then,

$$\Pr_D \left(\sup_{f \in \mathcal{F}} |c(f) - c_{emp}(f, D)| > \epsilon \right) < \delta$$

The proof of this theorem, together with a discussion on the hypothesis class \mathcal{H} , can be found in the appendix.

3.2 Applications of Theorem 2

3.2.1 Verifying the calibration of classifier

The first application of Theorem 2 is that we can verify whether the learned classifier f is calibrated. For simple hypothesis spaces \mathcal{F} (e.g., logistic regression), the corresponding hypothesis space \mathcal{H} has low Rademacher Complexity. In this case, Theorem 2 naturally guarantees the generalization of calibration measure.

There are also cases where the Rademacher Complexity of \mathcal{H} is not small. One notable example is SVM classifiers with Platt Scaling (Platt, 1999):

¹Our definition of Rademacher Complexity comes from Shalev-Shwartz and Ben-David’s textbook (2014), which is slightly different from the original definition in Bartlett and Mendelson’s paper (2003).

Claim 2. *Let $\mathcal{X} \subseteq \mathbb{R}^d$ and $\forall x \in \mathcal{X}, \|x\|_2 \leq 1$. Let \mathcal{F} be the following hypothesis class:*

$$\mathcal{F} = \left\{ x \rightarrow \frac{1}{1 + \exp(aw^T x + b)} : w \in \mathbb{R}^d, \|w\|_2 \leq B, a, b \in \mathbb{R} \right\}$$

If the training data size $n < d$ and the training data X_i are linearly independent, then $R_D(\mathcal{H}) = \frac{1}{2}$.

The proof can be found in the appendix. In the case of SVM, the dimensionality of the feature space is usually much larger than the training dataset size (this is especially true for kernel-SVM). In this situation, we can no longer verify the calibration property using only the training data, and we have to keep a separate validation dataset to calibrate the classifier (as suggested by Platt (1999)). When verifying the calibration of classifier on a validation dataset. The hypothesis class $\mathcal{F} = \{f\}$, and it is easy to verify that $\mathbb{E}_D R_D(\mathcal{H})$ is $O(\sqrt{\log n/n})$ in this case. Therefore, with enough validation data, we can still bound the calibration measure.

3.2.2 Implications on Learning Algorithm Design

Standard conditional probability estimation usually maximizes the likelihood to find the best classifier within the hypothesis space. However, since we can only guarantee the conditional probability estimates to be calibrated under the agnostic assumption, any calibrated classifier is essentially as good as the maximum likelihood estimation in terms of interpretability. Therefore, likelihood maximization is not necessarily the only method for estimating conditional probabilities.

There are other loss functions that are already widely used for binary classification. For example, hinge loss is at the foundation of large margin classifiers. Based on our discussion in this paper, we believe that these loss functions can also be used for conditional probability estimation. For example, Theorem 2 suggests the following constrained optimization problem:

$$\min \mathcal{L}(f, D) \quad \text{s.t.} \quad c_{emp}(f, D) = 0$$

where $\mathcal{L}(f, D)$ is the loss function we want to minimize. By optimizing over the space of empirically calibrated classifiers, we can ensure that the resulting classifier is also calibrated with respect to \mathcal{P} .

In fact, the conditional probability estimation algorithm developed by Kakade et al. (2011) already implicitly follows this framework (more elaboration on this point can be found in the appendix). We believe that many more interesting algorithms can be developed along this direction.

3.2.3 Connection to the Calibration Problem

Suppose that we are given an uncalibrated fuzzy classifier $f_0 : \mathcal{X} \rightarrow [0, 1]$, and we want to find a function g from $[0, 1]$ to $[0, 1]$, so that $g \circ f_0$ presents a better conditional probability estimation. This is the problem of classifier calibration, which has been studied in many papers (Zadrozny and Elkan, 2002) (Platt, 1999).

Traditionally, calibration algorithms find the best link function g by maximizing likelihood or minimizing squared loss. In this paper, we suggest a different approach to the calibration problem. We can find the best g by minimizing the empirical calibration measure $c_{\text{emp}}(g \circ f_0)$. Let us assume w.l.o.g. that the training dataset $D = \{(x_1, y_1), \dots, (x_n, y_n)\}$ satisfies

$$g(f_0(x_1)) \leq \dots \leq g(f_0(x_n))$$

Then we have,

$$\begin{aligned} c_{\text{emp}}(g \circ f_0, D) &= \frac{1}{n} \sup_{p_1, p_2} \left| \sum_{i=1}^n \mathbb{1}_{p_1 < g(f_0(x_i)) \leq p_2} (\mathbb{1}_{y_i=1} - g(f_0(x_i))) \right| \\ &\leq \frac{1}{n} \max_{a, b} \left| \sum_{a < i \leq b} (\mathbb{1}_{y_i=1} - g(f_0(x_i))) \right| \end{aligned}$$

This expression can be used as the objective function for calibration: we search over the space of hypotheses \mathcal{G} to find a function g that minimizes this objective function. Compared to other loss functions, the benefits of minimizing this objective function is that the resulting classifier is more likely to be calibrated, and therefore provides more interpretable conditional probability estimates.

In fact, one of the most well-known calibration algorithms, the isotonic regression algorithm, can be viewed as minimizing this objective function:

Claim 3. *Let \mathcal{G} be the set of all continuous non-decreasing functions from $[0, 1]$ to $[0, 1]$. Then the optimal solution that minimizes the squared loss $\min_{g \in \mathcal{G}} \sum_{i=1}^n (\mathbb{1}_{y_i=1} - g(f_0(x_i)))^2$ also minimizes $\min_{g \in \mathcal{G}} \max_{a, b} \left| \sum_{a < i \leq b} (\mathbb{1}_{y_i=1} - g(f_0(x_i))) \right|$*

The proof can be found in the appendix. Using this connection we proved several interesting properties of the isotonic regression algorithm, which can also be found in the appendix.

4 Empirical behavior of the calibration measure

In this section, we conduct some preliminary experiments to demonstrate the behavior of the calibration

measure on some common algorithms. We use two binary classification datasets from the UCI Repository²: ADULT³ and COVTYPE⁴. COVTYPE has been converted to a binary classification problem by treating the largest class as positive and the rest as negative. Five algorithms have been used in these experiments: naive Bayes(NB), boosted decision trees, SVM⁵, logistic regression(LR), random forest(RF).




Figure 2: The empirical calibration error

Figure 2 shows the empirical calibration error c_{emp} on test datasets for all methods. From the experimental results, it appears that Logistic Regression and Random Forest naturally produce calibrated classifiers, which is intuitive as we discussed in the paper. The calibration measure of Naive Bayes seems to be depending on the dataset. For large margin methods (SVM and boosted trees), the calibration measures are high, meaning that they are not calibrated (on these two datasets).

There is also an interesting connection between the calibration error and the benefit of applying a calibration algorithm, which is illustrated in Figure 3. In this experiment, we used a loss parameter p to control the asymmetric loss: each false negative incurs $1 - p$ cost and each false positive incurs p cost. All the algorithms are first trained on the training dataset, then calibrated on a separate validation set of size 2000 using isotonic regression. For each algorithm, we compute the prior-calibration and post-calibration average losses on the testing dataset using the following decision rule: For each data point X , we predict $Y = 1$ if and only if we predict that $\Pr(Y = 1|X) \geq p$. Finally, we report the ratio between two losses:

$$\text{loss ratio} = \frac{\text{the average loss after calibration}}{\text{the average loss before calibration}}$$

²These datasets are chosen from the datasets used in Niculescu-Mizil and Caruana’s work (2005). We only used two datasets because the experiments are only explorative (i.e., identifying potential properties of the calibration measure). More rigorous experiments are needed to formally verify these properties.

³<https://archive.ics.uci.edu/ml/datasets/Adult>

⁴<https://archive.ics.uci.edu/ml/datasets/Covertypes>

⁵For SVM and boosting, we rescale the output score to $[0, 1]$ by $(x - \min)/(\max - \min)$ as in Niculescu-Mizil and Caruana’s paper (2005)

As we can see in the Figure 3, the calibration proce-

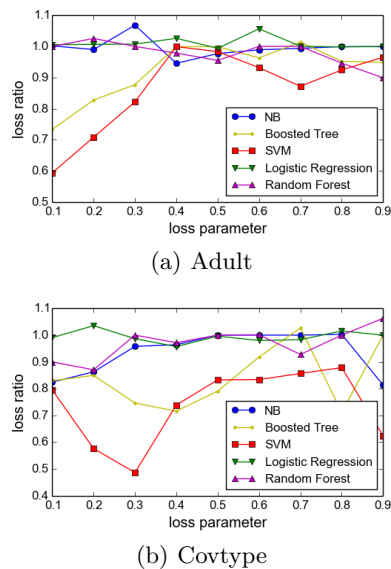


Figure 3: The loss ratio on two datasets

dures on average reduces the cost by 3%-5% for naive Bayes and random forest, 20% for SVM, 12% for boosted trees, and close to 0% for logistic regression. Comparing with the results in Figure 2, the two algorithms that benefit most from calibration (i.e., SVM and boosted trees) also has high empirical calibration error. This result suggests that if an algorithm already has a low calibration error to begin with, then it is not likely to benefit much from the calibration process. This finding could potentially help us decide whether we need to calibrate the current classifier using isotonic regression (Niculescu-Mizil and Caruana, 2005).

5 Conclusion

In this paper, we discussed the interpretability of conditional probability estimates under the agnostic assumption. We proved that it is impossible to upper bound the l_1 error of conditional probability estimates under such scenario. Instead, we defined a novel measure of calibration to provide interpretability for conditional probability estimates. The uniform convergence result between the measure and its empirical counterpart allows us to empirically verify the calibration property without making any assumption on the underlying distribution: the classifier is (almost) calibrated if and only if the empirical calibration measure is low. Our result provides new insights on conditional probability estimation: ensuring empirical calibration is already sufficient for providing interpretable conditional probability estimates, and thus many other loss functions (e.g., hinge loss) can also be utilized for es-

timating conditional probabilities.

Acknowledgement

We thank the anonymous reviewers for their valuable feedback. We acknowledge support from grant IIS-1513407 and IIS-1633755 awarded by the NSF, grant 1U54GM114838 awarded by NIGMS and 3U54EB020406-02S1 awarded by NIBIB through funds provided by the trans-NIH Big Data to Knowledge (BD2K) initiative (www.bd2k.nih.gov), and funds from Adobe, Google, the Sloan Foundation, and the Siebel Energy Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the funding agencies and organizations.

References

- [1] P. L. Bartlett and S. Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *The Journal of Machine Learning Research*, 3:463–482, 2003.
- [2] P. L. Bartlett and A. Tewari. Sparseness vs estimating conditional probabilities: Some asymptotic results. *The Journal of Machine Learning Research*, 8:775–790, 2007.
- [3] P. N. Bennett. Assessing the calibration of naive bayes posterior estimates. Technical report, DTIC Document, 2000.
- [4] C. M. Bishop et al. *Pattern recognition and machine learning*, volume 4. springer New York, 2006.
- [5] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.
- [6] G. F. Cooper. Nestor: A computer-based medical diagnostic aid that integrates causal and probabilistic knowledge. Technical report, DTIC Document, 1984.
- [7] T. Fawcett and F. Provost. Adaptive fraud detection. *Data mining and knowledge discovery*, 1(3):291–316, 1997.
- [8] D. P. Foster and R. V. Vohra. Asymptotic calibration. *Biometrika*, 85(2):379–390, 1998.
- [9] R. L. Graham. An efficient algorithm for determining the convex hull of a finite planar set. *Information processing letters*, 1(4):132–133, 1972.

- [10] S. M. Kakade, V. Kanade, O. Shamir, and A. Kalai. Efficient learning of generalized linear and single index models with isotonic regression. In *Advances in Neural Information Processing Systems*, pages 927–935, 2011.
- [11] L. Lee. Measures of distributional similarity. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*, pages 25–32. Association for Computational Linguistics, 1999.
- [12] A. Niculescu-Mizil and R. Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pages 625–632. ACM, 2005.
- [13] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [14] J. Platt et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. *Advances in large margin classifiers*, 10(3):61–74, 1999.
- [15] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press, 2014.
- [16] V. N. Vapnik and A. Y. Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability & Its Applications*, 16(2):264–280, 1971.
- [17] V. N. Vapnik and V. Vapnik. *Statistical learning theory*, volume 1. Wiley New York, 1998.
- [18] B. Zadrozny and C. Elkan. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. In *ICML*, volume 1, pages 609–616. Citeseer, 2001.
- [19] B. Zadrozny and C. Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 694–699. ACM, 2002.