MEASURING UNCERTAINTY CALIBRATION

Anonymous authors

Paper under double-blind review

ABSTRACT

We make two contributions to the problem of estimating the L_1 calibration error of a binary classifier from a finite dataset. First, we provide an upper bound for any classifier where the calibration function has bounded variation. Second, we provide a method of modifying any classifier so that its calibration error can be upper bounded efficiently without significantly impacting classifier performance and without any restrictive assumptions. All our results are non-asymptotic and distribution-free. We conclude by providing advice on how to measure calibration error in practice. Our methods yield practical procedures that can be run on real-world datasets with modest overhead.

1 Introduction

Machine Learning models are increasingly used to drive decision making in many tasks of practical importance. This process is notoriously susceptible (Cohen & Goldszmidt, 2004) to how well model outputs match probabilities of events in the real world, a property known as calibration. An important pre-requisite for ensuring good model calibration is the ability to measure it. This paper focuses precisely on this problem.

The most common approach to calibration is to place model outputs into discrete buckets (Zadrozny & Elkan, 2001; Naeini et al., 2015; Kumar et al., 2019). While this makes the calibration error possible to estimate using elementary tools, it introduces a dilemma. If one treats the estimate of the calibration error obtained by bucketing as a proxy for the calibration error of the classifier before the bucketing, the method is notoriously unreliable, producing different answers depending on the bucketing scheme used (Arrieta-Ibarra et al., 2022). On the other hand, if one treats the bucketing as part of the classifier, classification performance is likely to suffer since the bucketing is unknown to the training process.¹

Another approach is to frame the problem as a frequentist hypothesis test with the null hypothesis of zero calibration error (Arrieta-Ibarra et al., 2022; Tygert, 2023). This type of method is statistically powerful in detecting deviations from perfect calibration. However, because it focuses on a 'zero-error' null hypothesis, it is primarily designed to decide whether a model is (nearly) perfectly calibrated—not to quantitatively compare degrees of miscalibration between models.² Moreover, the proposed cumulative score is not intuitively interpretable and the formal derivation requires the asymptotic regime i.e. a large-enough sample size.³

Contributions We make two contributions to the problem of estimating the L^1 calibration error of a binary classifier from finite data:

- Certified bounds under bounded variation. We show that if the calibration function has bounded variation, then one can obtain a distribution-free, non-asymptotic *upper bound* on calibration error using a variant of total variation denoising. This provides the first finitesample guarantees under such a weak structural assumption.
- 2. **Certified bounds via perturbation to enforce smoothness.** When no assumption on bounded variation is acceptable, we propose a simple perturbation of classifier outputs that guarantees that the calibration function has bounded derivatives. This modification leaves

¹For example if the classifier is a neural network, one cannot back-propagate through the bucketed outputs.

²One might of course compare test statistics or p-values informally.

³Arrieta-Ibarra et al. (2022) do provide an experimental evaluation in the finite-sample regime.

classification performance essentially unchanged, yet enables a kernel-based estimator that yields tighter finite-sample bounds on calibration error.

Both approaches are non-asymptotic and fully distribution-free. We complement the theory with experiments on synthetic and real datasets, and conclude with actionable advice on how to measure calibration error in practice.

2 PRELIMINARIES

Calibration Function Consider a function $\eta:[0,1]\to [0,1]$. Denote by p(s) the arbitrary distribution of design points (scores). Denote by p(y|s) the Bernoulli distribution with mean $\eta(s)$. We have $\eta(s)=\mathbb{E}[y|s]$, and the function η is called the *calibration function*.

Datasets Consider an i.i.d dataset of tuples $\{(s_i, y_i)\}_i \sim p(s)p(y|s)$. The dataset is split into two index sets generated independently: training T and validation V, where we use the notations $Y_T = \{y_i : i \in T\}$, $Y_V = \{y_i : i \in V\}$, $S_T = \{s_i : i \in T\}$ and $S_V = \{s_i : i \in V\}$. Both these sets are generated independently from any data used to train the classifier.

Main Objective The main objective of this paper is to bound the quantity

$$CE \triangleq \mathbb{E}_s \left[\left| s - \eta(s) \right| \right], \tag{1}$$

known (Lee et al., 2023) as the L_1 expected calibration error, using data from finite (training and validation) datasets. In the next section, we describe the proof approach and structural assumptions on η that make this possible.

3 STRUCTURE OF RESULTS

The Need For Assumptions We begin by noting that estimating the quantity 1 is impossible without making structural assumptions on η . At a very minimum, for the expectation in equation 1 to be well-defined, one needs measurability. However, measurability alone is not sufficient to ensure the ability to estimate η from a finite dataset. In fact, Lee et al. (2023) have shown that even continuity (which implies measurability) is not sufficient, i.e. that there exist choices of p and q such that any estimator of the quantity 1 gives an answer wrong by a constant even with infinite data. We therefore need assumptions of a different kind. In this paper, we focus on two alternatives for what they can be.

Bounded Variation and Bounded Derivatives Assumptions First, in section 4, we assume that η has *bounded variation*, which turns out to be enough to bound the calibration error in a theoretically principled way using a variant of bucketing (with a special process for constructing buckets). However, because bounded variation is still a relatively weak assumption, it allows for many nearly pathological functions which the bound has to take into account, adversely impacting sample efficiency. This motivates us to look for stronger assumptions. Specifically, in section 5, we choose to assume η has two bounded derivatives. While this assumption ostensibly looks hard to justify, we do in fact show that we can always achieve it by applying a small perturbation to the classifier outputs. We demonstrate in section 7 that this only affects classifier performance in a minimal way in practice.

Lower vs Upper Bounds As a side note we remark that, while the techniques we develop work for both upper and lower bounds (with high probability), we focus the presentation on upper bounds since we typically wish to certify that the calibration error is small, not that it is large. We also note that all our results are fully distribution-free i.e. the distribution of scores p(s) can be discrete or continuous or a mixture, without any assumptions.

Main Proof Technique Our proof technique consists in constructing a surrogate for η . We then bound the calibration error as a sum of the calibration error wrt. the surrogate and the error of the surrogate construction process. The construction of the surrogate is different depending on the

assumption we are willing to make. Specifically, we use Total Variation denoising in case of the bounded variation assumption (section 4) and a kernel smoother in case of the bounded-derivatives assumption (section 5).

Concentration In our proof, we also need to quantify concentration of the empirical mean of bounded random variables. We use Bernstein's inequality to do that. Because there are several versions in literature, we quote the one we used. This version matches the work of Mnih et al. (2008). For a sum of n i.id, bounded random variables with mean μ and support on [0,1], we have $\mathbb{P}\left[\mu \leq \bar{X}_n + \mathrm{BB}(n,\delta,\widehat{\sigma}_n^2)\right] \geq 1-\delta$, where $\bar{X}_n := \frac{1}{n}\sum_{i=1}^n X_i, \widehat{\sigma}_{X_i}^2 := \frac{1}{n}\sum_{i=1}^n \left(X_i - \bar{X}_n\right)^2$ and

$$\mathrm{BB}(n,\delta,\widehat{\sigma}_{X_i}^2) = \sqrt{\frac{2\,\widehat{\sigma}_{X_i}^2\,\ln\!\left(3/\delta\right)}{n}} + \frac{3\,\ln\!\left(3/\delta\right)}{n}.$$

Here, we used the symbol BB to denote the concentration term in the Bernstein bound. We could have equally used Hoeffding inequality here, but Bernstein is slightly sharper.

4 MEASURING CALIBRATION UNDER BOUNDED VARIATION

Bounded Variation To formalize our assumption on η , we will need the concept of a Bounded Variation space (DeVore & Lorentz, 1993). For a function f, we can define total variation on the interval I as

$$TV(f, I) := \sup \left\{ \sum_{i=1}^{n-1} |f(s_{i+1}) - f(s_i)| \right\},\,$$

where the supremum is taken over the set of all finite sequences s_1, s_2, \ldots, s_n contained in the interval I. We say that the function f is of bounded variation, written $f \in BV([0,1])$ when $TV(f,[0,1]) < \infty$. Our theoretical work in this section is based on the assumption that the calibration function η satisfies $TV(\eta,[0,1]) \leq V$ for some known constant V.

TV Denoising Under the bounded variation assumption, we want to reconstruct η from noisy observations. To to so, it is useful to consider the TV denoising estimate $\hat{\eta}_T$. It is obtained by solving the optimization problem

$$\hat{\eta}_T = \arg\min_{v \in [0,1]^{|T|}} \frac{1}{2|T|} \|y_T - v\|_2^2 + \lambda \|Dv\|_1$$
(2)

where y_T is a vector of labels in the training set and $\hat{\eta}_T$ is our approximation of η on the training set. The matrix $D \in \mathbb{R}^{(|T|-1)\times |T|}$ is defined as:

$$(Dv)_i = v_{i+1} - v_i, \quad i = 1, 2, \dots, |T| - 1$$
 (3)

and $\lambda>0$ is a regularization parameter. Mammen & Van De Geer (1997) were the first to have shown results about reconstruction error of similar algorithms, i.e. the difference between η_T and $\hat{\eta}_T$. For our purposes, we rely on the fact that, if we set $\lambda=\sqrt{\frac{1}{6|T|}\ln\frac{4(|T|-1)}{\delta'}}$ in the optimization problem (2), then with probability $1-\delta'$ we have

$$\frac{1}{n} \|\eta_T - \hat{\eta}_T\|_1 \le \text{TVB}(\delta'). \tag{4}$$

Here, we denoted by η_T the vector containing the values of η on the training set. The term $\text{TVB}(\delta')$ diminishes if we increase the size of the training set. This result is proven (and a formula for $\text{TVB}(\delta')$ is provided) in Appendix A as Corollary 1, using proof techniques developed by Hütter & Rigollet (2016). Equation 4 is useful because it measures how close the reconstruction $\hat{\eta}$ is to the true η on the training set. However, in order to bound the calibration error, we need a similar guarantee on the population level, i.e. for scores sampled from p(s). To do this, we first define $\hat{\eta}$ for arbitrary scores from its values on the training set

$$\hat{\eta}(s) = \{\hat{\eta}_T\}_{i_\star} \quad \text{where} \quad i_\star = \arg\max_i \{i \in T : s_i \le s\}.$$

We can then write a bound similar to 4, but in expectaiton over s.

$$\mathbb{E}_{s}\left[|\eta(s) - \hat{\eta}(s)|\right] \le B(\delta) + (V + \hat{V})\sqrt{\frac{\log(2/\delta)}{2|T|}}.$$
(5)

See Corollary 2 in Appendix A for the proof.

TV Smoothing as Bucketing The function $\hat{\eta}$ obtained by TV denoising is piecewise constant. Therefore, a legitimate view of the smoothing process is to consider it as a very special example of a bucketing scheme, where the buckets correspond to the constant pieces of $\hat{\eta}$.

Bounding the Calibration Error To upper-bound the calibration error, we rely on a simple but powerful idea. First, imagine that the function η in the definition of the calibration error (1) is known. In that case, we could rely on the Bernstein bound to compute the expectation, using the fact that the absolute value in the definition of the calibration error is bounded in [0,1]. However, of course in practice η is not known directly and is observable only via noisy observations. This motivates us to proceed by first obtaining an approximation $\hat{\eta}$ to η by solving the optimization problem (2) on the training set and then apply Bernstein on the validation set. Note that (2) (and hence $\hat{\eta}$) is computable from the dataset. Intuitively, we can then use the approximation $\hat{\eta}$ in place of η . Of course, this leads to additional approximation error, which has to be added to the upper bound. Fortunately, the approximation error can be bounded by using equation (4). We now proceed to make this formal.

Proposition 1 (CE Upper Bound under Bounded Variation). *Construct the surrogate calibration function* $\hat{\eta}$ *by solving the optimization problem* (2). *With probability* $1 - \delta$, *we have*

$$CE \le \frac{1}{n} \sum_{i \in V} |s_i - \hat{\eta}(s_i)| + BB(|V|, \delta_1, \hat{\sigma}_{\{|s_i - \hat{\eta}(s_i)|\}_{i \in V}}^2) + TVB(\delta_2) + (V + \hat{V}) \sqrt{\frac{\log(2/\delta_3)}{2|T|}},$$

where $\delta_1 + \delta_2 + \delta_3 = \delta$ and $\hat{\sigma}^2_{\{|s_i - \hat{\eta}(s_i)|\}_{i \in V}}$) is the empirical variance of the term $|s_i - \hat{\eta}(s_i)|$ on the validation set.

Proof. We write

$$\mathbb{E}_{s}[|s - \eta(s)|] \le \underbrace{\mathbb{E}_{s}[|\hat{\eta}(s) - s|]}_{\text{term I}} + \underbrace{\mathbb{E}_{s}[|\hat{\eta}(s) - \eta(s)|]}_{\text{term II}}$$
(6)

and bound term I and II separately. We begin by bounding term I. Since the term under expectation is bounded in [0,1] and the validation set is independent of the training set used to obtain $\hat{\eta}$, we use Bernstein inequality with probability δ_1 . We bound term II using equation 5. The result follows by picking the delta terms so that $\delta_1 + \delta_2 + \delta_3 = \delta$.

Proposition 1 is important because it gives us a handle on bounding the calibration error of any classifier with a bounded variation calibration function in a form computable from observable quantities. Note that the term \hat{V} (the total variation of the surrogate) is computable from data (since the TV solution is piecewise constant).

When is Bounded Variation Reasonable? Verifying that a function is of bounded variation is impossible to do from finite samples only, even if the samples had not been noisy. However, the result in this section is still useful. This is because, in scenarios of practical relevance, η is not a completely arbitrary function but one that was obtained from training a classifier. Since the classifier is trying to distinguish positive form negative examples, it is not unreasonable to expect that η should be monotone increasing, i.e. that higher scores correspond to a higher chance of seeing a positive label. It is well known that all monotone functions mapping to [0,1] have total variation bounded by 1, i.e. we can use V=1.

5 Measuring Calibration Under Bounded Derivatives

In certain scenarios the bounded variation assumption from the previous section is not suitable. This can be for two reasons—either we don't know the total variation bound V or we want to achieve better sample efficiency. In this section, we address this need by providing a method of obtaining guarantees on calibration for a broad class of classifiers. The main idea is to start with a completely arbitrary classifier and modify it in a way which ensures the ability to bound calibration error but does not harm classification performance.

Ensuring Smoothness By Perturbation The modification amounts to perturbing the probability output of the classifier by a small amount, specified by the perturbation bandwidth. This can be done either only at inference time or at both inference time and training time, to ensure better performance. We experimentally evaluate the effect of the perturbation in section 7 and conclude that there is almost no loss of performance relative to an unmodified classifier for realistic choices of the bandwidth. This allows us to write the calibration function of the perturbed classifier as

$$\eta(s) = \frac{\int_0^1 \eta_{\text{orig}}(s_{\text{orig}}) k(s \mid s_{\text{orig}}) dp(s_{\text{orig}})}{\int_0^1 k(s \mid s_{\text{orig}}) dp(s_{\text{orig}})}.$$
 (7)

See Appendix C for the derivation of this formula. This expression is a kernel-weighted average of the original calibration function $\eta_{\text{orig}}(s_{\text{orig}})$, smoothened with the kernel k.

Kernel Choice Denote by h the amount of the perturbation (the bandwidth). Formally, we define the score s output by the perturbed classifier as a sample from the probability distribution with the PDF defined by the equation

$$k(s \mid s_{\text{orig}}) = \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s_{\text{orig}} - s}{h}\right), \tag{8}$$

where sech is the hyperbolic secant function and the normalizer $Z(s_{\rm orig},h)$ is chosen so that the PDF integrates to one over [0,1]. By construction, the perturbed score is in the interval [0,1]. It turns out that using a perturbation of the form given in equation 8 has significant benefits over using a more conventional choice like the truncated Gaussian. We defer a discussion of our motivation for choosing the perturbation to appendix C.2.

Bounded Derivatives From Perturbation Crucially, the following Lemma, proved in appendix C, shows that the calibration function of the modified classifier has two bounded derivatives, regardless of the properties of the calibration function before the perturbation.

Lemma 1 (Bounded Derivatives From Perturbation). For any classifier that perturbs its output scores by replacing an output s_{orig} with a perturbed output s sampled according the PDF in equation 8, the calibration function η is twice differentiable. Moreover its first derivative is uniformly bounded by $\frac{1}{2h}$ and its second derivative is uniformly bounded by $\frac{3}{2}\frac{1}{h^2}$.

Surrogate of Calibration Function Via Nadaraya-Watson Smoothing Lemma 1 gives us the derivative bounds that allow us to approximate η in a principled way, i.e. bound the smoothing error of a kernel smoother applied to the dataset. We now define $\hat{\eta}$ as

$$\hat{\eta}(s') = \sum_{i \in T} w_i(s') y_i, \quad w_i(s') = \frac{k'(s'|s_i)}{\sum_{j \in T} k'(s'|s_j)}.$$

Then we have:

$$\mathbb{E}_{Y_T}\left[|\hat{\eta}(s') - \eta(s')|S_T\right] \le g_T(s'). \tag{9}$$

where $g_T(s')$ is fully computable from data and diminishes for s' in the support of p(s). See Appendix B for proof and the precise formula for $g_T(s')$. Unlike the choice of the perturbation kernel k, the choice of the smoothing kernel k' used to construct the surrogate is an implementation detail. We use Epanechnikov weights with nearest-neighbor fallback and boundary renormalization, tempered by an exponent $\tau=1.2$ before re-normalizing; since only the "weights sum to one" property matters. One could also explore alternatives to using a kernel smoother. For example, we can see that Lemma 1 implies that η is Lipschitz with constant $\frac{1}{2h}$ and use that to bound the error of a bucketing scheme. We do in fact compare to this approach as part of our experiments.

Bounding The Calibration Error We now use equation 9 to bound the calibration error. Again, we stress that all terms in the bound are possible to estimate from data.

Proposition 2 (CE Upper Bound under Bounded Derivatives). Assuming η is twice differentiable with bounds on the first and second derivative, we have

$$CE \leq \frac{1}{|V|} \sum_{i \in V} |\hat{\eta}(s_i) - s_i| + \frac{1}{|V|} \sum_{i \in V} g_T(s_i) + BB(|V|, \delta/2, \hat{\sigma}^2_{\{|s_i - \hat{\eta}(s_i)|\}_{i \in V}}) + RBB(|V|, \delta/2, \hat{\sigma}^2_{\{\frac{1}{B}g_T(s_i)\}_{i \in V}})$$

with probability $1 - \delta$, with the randomness taken over the validation set, conditional on the training set, for a constant R that depends⁴ on the kernel k.

Proof. We begin by introducing conditioning on the training set. Because the quantity $|s - \eta(s)|$ doesn't depend on the training set, we can write.

$$\mathbb{E}_s[|s - \eta(s)|] = \mathbb{E}_s[|s - \eta(s)||S_T, Y_T]$$

Introduce the notation

$$h(s') = \mathbb{E}_{Y_T} [|\hat{\eta}(s') - \eta(s')|S_T] = \mathbb{E}_{Y_T'} [|\hat{\eta}(s') - \eta(s')|S_T],$$

where the last equality holds because we can always rename a bound variable. We can now apply the triangle inequality to get

$$\mathbb{E}_{s}[|s - \eta(s)||S_{T}, Y_{T}] \leq \mathbb{E}_{s}[|s - \hat{\eta}(s)||S_{T}, Y_{T}] + \mathbb{E}_{s}[h(s)|S_{T}, Y_{T}]$$

$$= \mathbb{E}_{s}[|s - \hat{\eta}(s)||S_{T}, Y_{T}] + \mathbb{E}_{s}[h(s)|S_{T}], \qquad (10)$$

where the last equality is because h(s') doesn't depend on Y_T . We can now apply the Bernstein inequality twice, once for each term on the right hand side of equation (10). Note that both terms under the expectation are bounded in [0,1]. Note also that the validation set is drawn i.i.d..

For the first term, invoking Bernstein (with half the probability budget) conditionally on S_T, Y_T . This gives us

$$\mathbb{E}_{s}\left[|s-\hat{\eta}(s)||S_{T},Y_{T}\right] \leq \frac{1}{|V|} \sum_{i \in V} |\hat{\eta}(s_{i}) - s_{i}| + BB(|V|, \delta/2, \hat{\sigma}_{\{|s_{i}-\hat{\eta}(s_{i})|\}_{i \in V}}^{2}). \tag{11}$$

For the second term in (10), we proceed analogously as before, this time conditioning on S_T and invoking Lemma 6. In order to apply the Bernstein inequality, we introduce a constant R that ensures that $\frac{1}{R}g_T(s') \in [0,1]$ for all s'. We get

$$\mathbb{E}_{s,Y_T} [h(s)|S_T] \le \frac{1}{|V|} \sum_{i \in V} g_T(s_i) + R \operatorname{BB}(|V|, \delta/2, \hat{\sigma}^2_{\{\frac{1}{R}g_T(s_i)\}_{i \in V}}). \tag{12}$$

We get the result by plugging terms (11) and (12) into (10) and using the union bound. \Box

Implementation Detail (Cross-Fitting). The theory in the paper is written for a fixed split of the data indices into training and validation. In practice, we use K-fold cross-fitting: randomly partition the samples into disjoint folds, and for each fold fit a surrogate eta on the complement while evaluating only on the validation folds, then we aggregate across folds. This preserves the fixed-split assumption (every validation point is scored by a model that did not train on it) while reducing variance and making full use of the data.⁵

6 RELATED WORK

Bucketing Variants of bucketing were first introduced in meteorology (Murphy & Winkler, 1977) and have been used to approximate the calibration error ever since the ML community started paying attention to calibration (Zadrozny & Elkan, 2001; Niculescu-Mizil & Caruana, 2005; Naeini et al., 2015; Leathart et al., 2017; Guo et al., 2017; Vaicenavicius et al., 2019; Kumar et al., 2019). Nixon et al. (2019) has criticized vanilla bucketing and proposed alternative binning schemes.

⁴For the truncated Epanechnikov kernel, we set $R = b_1 h + \frac{1}{2} b_2 h^2 + \frac{1}{2}$.

⁵We apply such cross-fitting to all three methods we evaluate.

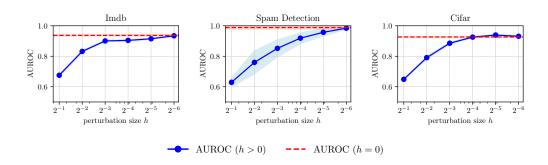


Figure 1: AUROC under Perturbation

Lipschitz Assumptions Vaicenavicius et al. (2019); Dimitriadis et al. (2023); Futami & Fujisawa (2024) used the Lipschitz assumption to get a handle on variants of the calibration error. However, the assumption was not justified. Our paper is more general in two ways: bounded variation (as in section 4) is a far weaker assumption then Lipschitz and our perturbation scheme in section 5 guarantees smoothness (and hence the Lipschitz property) rather than just assuming it. Zhang et al. (2020) assumes Hölder continuity, again without justifying it from first principles.

Kolmogoroff-Smirnoff and Kuiper Arrieta-Ibarra et al. (2022); Tygert (2023) propose a scheme for measuring calibration based on cumulative sum metrics. By relating these metrics to random walks, it is possible to derive highly efficient frequentist tests where the null hypothesis represents perfect calibration. Our work is different from these approaches in several ways. First, our measured quantity in equation 1 is more interpretable. Second, unlike them, we do not rely on asymptotics of any kind do our result stay theoretically valid for any sample size. Third, our technique can be used to compare two poorly calibrated models with each other while they only support comparisons against a baseline of perfect calibration. Overall, unlike KS-based metrics, which detect mis-calibration but do not yield interpretable bounds, our approach directly certifies an upper bound on CE.

Impossibility Results Arrieta-Ibarra et al. (2022) have argued that binning requires an infinite number of samples per bin to accomplish both asymptotic consistency and asymptotic power. However, these results were obtained without structural assumptions on the calibration function η . Moreover, Lee et al. (2023) has shown that assuming continuity of η is not enough, i.e. that there are continuous calibration functions which make calibration error impossible to estimate from finite samples. We circumvent these problems when we introduce our assumptions of bounded variation and continuity with bounded derivatives.

7 EXPERIMENTS

Overview of Experiments We approach the experimental evaluation of our method in two stages. First, we establish that we can afford to perturb classifier outputs by as much as 2^{-6} for realistic datasets, without significantly impacting AUROC. Second, we use Lemma 1 to certify this guarantees bounded derivatives, allowing us to obtain the calibration error up to about 0.02 using about 10^7 samples. While this may seem like a lot of data, we note that estimating the calibration error is an inherently hard problem, which scales poorly with the size of the dataset. In fact, to the best of our knowledge, our estimator is the best currently available for the setting we consider.

Perturbation vs AUROC In figure 1, we measure the effect of perturbing the output probabilities of a classifier on its classification performance, measured by AUROC. We examine three classification tasks: IMDB, Spam Detection, where the classifier is a fine-tuned version of BERT as well as CIFAR where we use a vision transformer. We note that perturbing by $h=2^{-6}$ costs us almost nothing in terms of AUROC for all three datasets. Note that, to obtain Figure 1, the neural network loss was modified to account for the fact that we are applying perturbation at inference time. Details of this are explained in appendix E. The modification is cheap, meaning that the cost of training the network with the modified loss is virtually the same as using the standard cross-entropy loss.

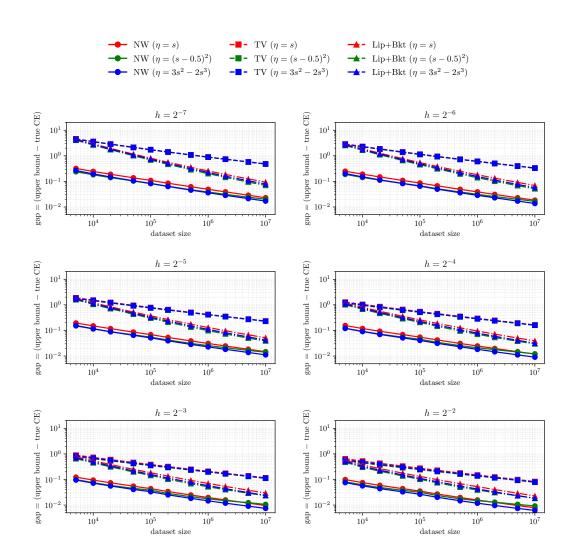


Figure 2: Number of Samples Needed to Achieve Gap (Synthetic Data).

Upper Bound Quality vs Sample Size The best way to evaluate an estimator is by comparing to the ground truth, which is unavailable for real-world datasets. We therefore introduced three synthetic examples of the calibration functions (see heading of Figure 2 for definitions). The figure also shows the sample efficiency of bounding calibration error, defined as the gap between the probabilistic upper bound and the ground truth of the calibration error. Since we used synthetic datasets for this experiment, the ground truth is known and we can evaluate the gap exactly. We evaluate three methods: our proposed kernel estimator (denoted as NW in the plot), Lipschitz bucketing (denoted as Lip+Bkt in the plot) and TV denoising (denoted as TV).

Real Data Experiment We also wanted to establish how the three bounding techniques work on real data. To do this, we used the Amazon Polarity dataset. Since the true calibration error is unknown, we plot the upper bounds on the calibration error directly. It can be seen that NW smoothing again gives the tightest bounds.

Note On Repeats Figures 2 and 3 do not contain confidence bars because they are small. We repeated the all experiments on these plots 64 times and report mean performance.

Empirical Rates We can see that the log-log curves in Figure 2 are straight lines. We can interpret their slopes as empirical rates. We did that obtaining the following results.

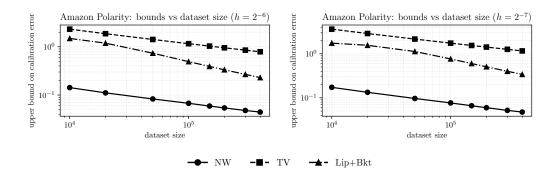


Figure 3: Upper Bounds vs Dataset Size (Real Data).

Method	Empirical Rate ⁶	Theoretical Rate
NW	[-0.381, -0.231]	-1/3
TV	[-0.377, -0.136]	-1/4
Lip+Bkt	[-0.514, -0.302]	-1/3

We can see that the rates obtained in practice match the theory. We note that, while we achieve the same rate as Lipschitz bucketing, our constants are better, meaning we produce significantly tighter bounds.

8 LARGE LANGUAGE MODEL USAGE

Large Language Models (ChatGPT 4, ChatGPT 5 and Gemini) were used extensively at all stages in the writing of the paper. Both the theoretical results and the code were verified by the authors. We describe the details of LLM usage in Appendix G.

9 PRACTICAL ADVICE

The preferred technique is to apply a small perturbation and use Proposition 2 (works whether or not training is aware of the perturbation). If perturbations are impossible, assume bounded variation and use Proposition 1 (less sample-efficient). Without either assumption, the problem is intractable in practice.

10 CONCLUSIONS

We have proposed two new methods for measuring the calibration error of a binary classifier and described their tradeoffs. We empirically demonstrated it is possible to measure calibration error on a real task of practical importance. We concluded by giving practical advice on how to measure calibration error in practice. While we focused on binary classifiers, we conjecture that our perturbation technique extends naturally to multiclass calibration (future work).

REPRODUCIBILITY STATEMENT

We provided full proofs for all our claims. We commit to release the source code that fully reproduces our figures upon acceptance.

ETHICS STATEMENT

Our theoretical results make abstract statements about the quality of estimators. For the experiments, we used readily available public datasets. Neither gives rise to ethics concerns.

REFERENCES

- Imanol Arrieta-Ibarra, Paman Gujral, Jonathan Tannen, Mark Tygert, and Cherie Xu. Metrics of calibration for probabilistic predictions. *Journal of Machine Learning Research*, 23(351):1–54, 2022.
- Ira Cohen and Moises Goldszmidt. Properties and benefits of calibrated classifiers. In *European conference on principles of data mining and knowledge discovery*, pp. 125–136. Springer, 2004.
- CM Da Fonseca and J Petronilho. Explicit inverses of some tridiagonal matrices. *Linear Algebra and its Applications*, 2001.
- Ronald A DeVore and George G Lorentz. *Constructive approximation*, volume 303. Springer Science & Business Media, 1993.
- Timo Dimitriadis, Lutz Dümbgen, Alexander Henzi, Marius Puke, and Johanna Ziegel. Honest calibration assessment for binary outcome predictions. *Biometrika*, 110(3):663–680, 2023.
- Futoshi Futami and Masahiro Fujisawa. Information-theoretic generalization analysis for expected calibration error. *Advances in Neural Information Processing Systems*, 37:84246–84297, 2024.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pp. 1321–1330. PMLR, 2017.
- Jan-Christian Hütter and Philippe Rigollet. Optimal rates for total variation denoising. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir (eds.), 29th Annual Conference on Learning Theory, volume 49 of Proceedings of Machine Learning Research, pp. 1115–1146, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. URL https://proceedings.mlr.press/v49/huetter16.html.
- Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009.
- Ananya Kumar, Percy S Liang, and Tengyu Ma. Verified uncertainty calibration. *Advances in neural information processing systems*, 32, 2019.
- Tim Leathart, Eibe Frank, Geoffrey Holmes, and Bernhard Pfahringer. Probability calibration trees. In *Asian conference on machine learning*, pp. 145–160. PMLR, 2017.
- Donghwan Lee, Xinmeng Huang, Hamed Hassani, and Edgar Dobriban. T-cal: An optimal test for the calibration of predictive models. *Journal of Machine Learning Research*, 24(335):1–72, 2023. URL http://jmlr.org/papers/v24/22-0320.html.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL http://www.aclweb.org/anthology/P11-1015.
- Enno Mammen and Sara Van De Geer. Locally adaptive regression splines. *The Annals of Statistics*, 25(1):387–413, 1997.
- Julian John McAuley and Jure Leskovec. From amateurs to connoisseurs: modeling the evolution of user expertise through online reviews. In *Proceedings of the 22nd international conference on World Wide Web*, pp. 897–908, 2013.
- Volodymyr Mnih, Csaba Szepesvári, and Jean-Yves Audibert. Empirical bernstein stopping. In *Proceedings of the 25th international conference on Machine learning*, pp. 672–679, 2008.
- Allan H Murphy and Robert L Winkler. Can weather forecasters formulate reliable probability forecasts of precipitation and temperature. *National weather digest*, 2(2):2–9, 1977.
 - Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 29, 2015.

Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pp. 625–632, 2005.

Jeremy Nixon, Michael W Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. Measuring calibration in deep learning. In *CVPR workshops*, volume 2, 2019.

Walter Rudin. Principles of Mathematical Analysis. McGraw-Hill, New York, 3rd edition, 1976.

David Schulte, Felix Hamborg, and Alan Akbik. Less is more: Parameter-efficient selection of intermediate tasks for transfer learning. In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 9431–9442, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.529. URL https://aclanthology.org/2024.emnlp-main.529/.

Mark Tygert. Calibration of p-values for calibration and for deviation of a subpopulation from the full population. *Advances in Computational Mathematics*, 49(5):70, 2023.

Juozas Vaicenavicius, David Widmann, Carl Andersson, Fredrik Lindsten, Jacob Roll, and Thomas Schön. Evaluating model calibration in classification. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 3459–3467. PMLR, 16–18 Apr 2019. URL https://proceedings.mlr.press/v89/vaicenavicius19a.html.

Bianca Zadrozny and Charles 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*, pp. 204–213, 2001.

Jize Zhang, Bhavya Kailkhura, and T Yong-Jin Han. Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In *International conference on machine learning*, pp. 11117–11128. PMLR, 2020.

A TOTAL VARIATION DENOISING

A.1 PROBLEM FORMULATION

Given a set of training points $s_1, s_2, \ldots, s_{|T|}$, we introduce the vector notations

$$\{\eta_T\}_i = \eta(s_i), \ \{\hat{\eta}_T\}_i = \hat{\eta}(s_i).$$

We consider a variant of the classical TV denoising problem where we observe binary training data $y_T \in \{0,1\}^{|T|}$ generated according to:

$$y_{T_i} \sim \text{Bernoulli}(\eta_{T_i}), \quad i = 1, 2, \dots, |T|.$$
 (13)

We can always decompose y_T as:

$$y_T = \eta_T + \epsilon_T \tag{14}$$

where $\epsilon_T \in \mathbb{R}^n$ is the noise term representing the deviation of our observations from the expected values.

The TV denoising estimate $\hat{\eta}_T$ is obtained by solving:

$$\hat{\eta}_T = \arg\min_{\eta_T \in [0,1]^{|T|}} \frac{1}{2|T|} \|y_T - \eta_T\|_2^2 + \lambda \|D\eta_T\|_1$$
(15)

where $D \in \mathbb{R}^{(|T|-1)\times |T|}$ is the finite difference operator defined as:

$$(D\eta_T)_i = \eta_{T_{i+1}} - \eta_{T_i}, \quad i = 1, 2, \dots, |T| - 1$$
(16)

Since the operator D is linear, it can be represented as a matrix. We abuse notation slightly, also calling this matrix D. The variable $\lambda > 0$ in equation 15 is a regularization parameter. We will later derive a way to set λ optimally as a function of observable quantities.

We note that by assumption, the total variation of the true parameter vector is bounded: $||D\eta_T||_1 \le V$ for some known constant V > 0.

A.2 AUXILIARY LEMMAS

We first need to establish two useful Lemmas.

Lemma 2. Each row of the matrix $D^{+\top}$ has squared 2-norm at most $\frac{1}{3}|T|$. Here, $D^{+\top}$ is the transpose of the Moore-Penrose pseudo-inverse of the matrix D.

Proof. Because D is of full row rank, we have $D^{+\top} = (D^{\top}(DD^{\top})^{-1})^{\top} = (DD^{\top})^{-1}D$. Let us first examine the elements of the matrix $(DD^{\top})^{-1}$. The matrix DD^{\top} is tri-diagonal and its inverse can be computed using the formula derived by Da Fonseca & Petronilho (2001). Using the formula and simplifying gives

$$\{(DD^{\top})^{-1}\}_{ij} = \frac{\min(i,j)(|T| - \max(i,j))}{|T|}.$$

Multiplying by D and again simplifying terms, we get

$$\{(D)^{+\top}\}_{jk} = \{(DD^{\top})^{-1}D\}_{jk} = \begin{cases} \frac{k}{|T|} & \text{for } k < j\\ \frac{k-|T|}{|T|} & \text{for } k \ge j \end{cases}.$$

The squared-2 norm of row j simplifies as:

$$\sum_{k=1}^{j-1} \left(\frac{k}{|T|}\right)^2 + \sum_{k=j}^{|T|-1} \left(\frac{k-|T|}{|T|}\right)^2 = \frac{1}{|T|^2} \left(\sum_{k=1}^{j-1} k^2 + \sum_{k=j}^{|T|-1} (|T|-k)^2\right)$$

$$\leq \frac{1}{|T|^2} \left(\sum_{k=1}^{j-1} k^2 + \sum_{k=1}^{|T|-j} k^2\right)$$

$$\leq \frac{1}{|T|^2} \left(\sum_{k=1}^{|T|-1} k^2\right)$$

$$= \frac{(|T|-1)|T|(2|T|-1)}{6|T|^2}$$

$$= \frac{2|T|^2 - 3|T| + 1}{6|T|}$$

$$= \frac{2|T| - 3 + \frac{1}{|T|}}{6} \leq \frac{2|T|}{6} = \frac{1}{3}|T|$$

Lemma 3. Let noise ϵ be defined as in equation (14). Then we simultaneously have

$$\|\Pi\epsilon\|_2 \le t_1, \|D^{+\top}\epsilon\|_{\infty} \le t_2$$

with probability $1 - \delta$, where t_1 and t_2 are defined as in Lemma 4.

Proof. We begin with the bound on $\|\Pi\epsilon\|_2$. First we use the fact that $\Pi = \frac{1}{|T|} \mathbf{1} \mathbf{1}^{\top}$ to write

$$\|\Pi\epsilon\|_{2} = \left\| \frac{1}{|T|} \left(\sum \epsilon_{j} \right) 1 \right\|_{2} = \frac{1}{|T|} \left| \sum \epsilon_{j} \right| \|1\|_{2} = \frac{1}{\sqrt{|T|}} \left| \sum \epsilon_{j} \right|.$$

Now, ϵ_j are centered Bernoulli random variables, hence they are sub-Gaussian with parameter $\frac{1}{4}$. Hence $\sum \epsilon_i$ is sub-Gaussian with parameter $\frac{|T|}{4}$. This allows us to write.

$$P(\left|\sum \epsilon_j\right| \geq t) \leq 2 \exp\left\{-\frac{t^2}{2^{\frac{|T|}{4}}}\right\} = 2 \exp\left\{-\frac{2t^2}{|T|}\right\}.$$

Setting the right-hand side to equal δ' , we get $t = \sqrt{\frac{|T|}{2} \ln \frac{2}{\delta'}}$. Hence, with probability $1 - \delta'$, we have

$$\|\Pi\epsilon\|_2 \le \sqrt{\frac{1}{2}\ln\frac{2}{\delta'}}.\tag{17}$$

We now move to bound the term $\|D^{+\top}\epsilon\|_{\infty}$. Invoking Lemma 2, we know that every row of $D^{+\top}$ has squared 2-norm upper bounded by $\frac{1}{3}|T|$. Also, like before, each element of ϵ is sub-Gaussian with parameter $\frac{1}{4}$. Combining these insights, and using the fact that noise terms are independent, the element j of the vector $D^{+\top}\epsilon$ is sub-Gaussian with parameter $\frac{|T|}{3}\frac{1}{4}=\frac{|T|}{12}$. This allows us to write.

$$P(\left|\left\{D^{+\top}\epsilon\right\}_{j}\right| \ge t) \le 2\exp\left\{-\frac{t^{2}}{2\frac{|T|}{12}}\right\} = 2\exp\left\{-\frac{6t^{2}}{|T|}\right\}.$$

By the union bound, we have

$$P(\left\|D^{+\top}\epsilon\right\|_{\infty} \geq t) = P(\max_{j} \left|\left\{D^{+\top}\epsilon\right\}_{j}\right| \geq t) \leq 2(|T|-1)\exp\left\{-\frac{6t^{2}}{|T|}\right\}.$$

Setting the right-hand side to δ'' , we obtain $t = \sqrt{\frac{|T|}{6} \ln \frac{2(|T|-1)}{\delta''}}$. This means that with probability $1 - \delta''$, we have

$$||D^{+\top}\epsilon||_{\infty} \le \sqrt{\frac{|T|}{6} \ln \frac{2(|T|-1)}{\delta''}}$$

The result follows by using (17) and (A.2) together, setting $\delta' = \delta'' = \delta/2$ and invoking the union bound.

A.3 PROOF OF MAIN RESULT ON TV DENOISING

Lemma 4. Define $e = \|\eta_T - \hat{\eta}_T\|_2$. If we set $\lambda = \frac{t_2}{|T|}$ in the optimization problem (2), then with probability $1 - \delta$ we have

$$\frac{1}{|T|}e^2 \le \frac{1}{|T|} \left[2t_1^2 + 2t_1\sqrt{t_1^2 + 4t_2V} + 4t_2V \right],$$

where
$$t_1 = \sqrt{\frac{1}{2} \ln \frac{4}{\delta}}$$
 and $t_2 = \sqrt{\frac{|T|}{6} \ln \frac{4(|T|-1)}{\delta}}$.

The proof follows the technique proposed by Hütter & Rigollet (2016). We make all constants explicit.

Proof. Using equation (2) and optimality, we get.

$$\frac{1}{2n} \|y_T - \hat{\eta}_T\|_2^2 + \lambda \|D\hat{\eta}_T\|_1 \le \frac{1}{2|T|} \|y_T - \eta_T\|_2^2 + \lambda \|D\eta_T\|_1$$

Using the fact that $y_T = \eta_T + \epsilon_T$, we get

$$\frac{1}{2|T|}\|\hat{\eta}_T - \eta_T - \epsilon_T\|_2^2 + \lambda \|D\hat{\eta}_T\|_1 \le \frac{1}{2|T|} \|\epsilon_T\|_2^2 + \lambda \|D\eta_T\|_1.$$

Expanding the squared term on the left hand side, we get

$$\frac{1}{2|T|}\|\hat{\eta}_T - \eta_T\|_2^2 - \frac{1}{|T|}(\hat{\eta}_T - \eta_T)^\top \epsilon_T + \frac{1}{2|T|}\|\epsilon_T\|_2^2 + \lambda \|D\hat{\eta}_T\|_1 \le \frac{1}{2|T|}\|\epsilon_T\|_2^2 + \lambda \|D\eta_T\|_1.$$

Canceling the term $\frac{1}{2|T|} \|\epsilon_T\|_2^2$, we get

$$\frac{1}{2|T|} \|\hat{\eta}_T - \eta_T\|_2^2 - \frac{1}{|T|} (\hat{\eta}_T - \eta_T)^\top \epsilon_T + \lambda \|D\hat{\eta}_T\|_1 \le \lambda \|D\eta_T\|_1.$$

Rearranging terms, we get

$$\frac{1}{|T|} \|\hat{\eta}_T - \eta_T\|_2^2 \le \frac{2}{|T|} (\hat{\eta}_T - \eta_T)^\top \epsilon - 2\lambda \|D\hat{\eta}_T\|_1 + 2\lambda \|D\eta_T\|_1.$$
 (18)

We now focus on the term $(\hat{\eta}_T - \eta_T)^{\top} \epsilon_T$. Denote by Π the projection on the null-space of D (the space spanned by the vector of all ones). Noting that $D^{\top}D^{+\top} = (I - \Pi)$, we write

$$(\hat{\eta}_{T} - \eta_{T})^{\top} \epsilon = (\hat{\eta}_{T} - \eta_{T})^{\top} (I - \Pi) \epsilon + (\hat{\eta}_{T} - \eta_{T})^{\top} \Pi \epsilon_{T}$$

$$= (\hat{\eta}_{T} - \eta_{T})^{\top} D^{\top} D^{+\top} \epsilon_{T} + (\hat{\eta}_{T} - \eta_{T})^{\top} \Pi \epsilon_{T}$$

$$\leq \|D(\hat{\eta}_{T} - \eta_{T})\|_{1} \|D^{+\top} \epsilon_{T}\|_{\infty} + \|(\hat{\eta}_{T} - \eta_{T})\|_{2} \|\Pi \epsilon_{T}\|_{2}.$$

Here, the last inequality follows by using Hölder inequality twice. Invoking Lemma 3, we get

$$(\hat{\eta}_T - \eta_T)^{\top} \epsilon_T \le ||D(\hat{\eta}_T - \eta_T)||_1 t_2 + ||(\hat{\eta}_T - \eta_T)||_2 t_1$$

Plugging back into equation (18), we get

$$\frac{1}{|T|} \|\hat{\eta}_T - \eta_T\|_2^2 \le \frac{2}{|T|} \|D(\hat{\eta}_T - \eta_T)\|_1 t_2 + \frac{2}{|T|} \|(\hat{\eta}_T - \eta_T)\|_2 t_1 - 2\lambda \|D\hat{\eta}_T\|_1 + 2\lambda \|D\eta_T\|_1.$$

Rearranging terms, we get

$$\frac{1}{|T|}\|\hat{\eta}_T - \eta_T\|_2^2 - \frac{2}{|T|}\|(\hat{\eta}_T - \eta_T)\|_2 t_1 \leq \frac{2}{|T|}\|D(\hat{\eta}_T - \eta_T)\|_1 t_2 - 2\lambda\|D\hat{\eta}_T\|_1 + 2\lambda\|D\eta_T\|_1.$$

Using the fact that $||D(\hat{\eta}_T - \eta_T)||_1 \le ||D\hat{\eta}_T||_1 + ||D\eta_T||_1$, we get

$$\frac{1}{|T|}\|\hat{\eta}_T - \eta_T\|_2^2 - \frac{2}{|T|}\|(\hat{\eta}_T - \eta_T)\|_2 t_1 \le \frac{2}{|T|}(\|D\hat{\eta}_T\|_1 + \|D\eta_T\|_1)t_2 - 2\lambda\|D\hat{\eta}_T\|_1 + 2\lambda\|D\eta_T\|_1.$$

Setting $\lambda = \frac{t_2}{|T|}$ and canceling terms, we get

$$\frac{1}{|T|} \|\hat{\eta}_T - \eta_T\|_2^2 - \frac{2}{|T|} \|(\hat{\eta}_T - \eta_T)\|_2 t_1 \le \frac{2}{|T|} \|D\eta_T\|_1 t_2 + 2\frac{t_2}{|T|} \|D\eta_T\|_1 = \frac{4t_2}{|T|} \|D\eta_T\|_1.$$

Using the assumption of bounded total variation, we get $||D\eta_T||_1 \leq V$, so that

$$\frac{1}{|T|} \|\hat{\eta}_T - \eta_T\|_2^2 - \frac{2}{|T|} \|(\hat{\eta}_T - \eta_T)\|_2 t_1 \le \frac{4t_2}{|T|} V.$$

Denoting $e = \|\hat{\eta}_T - \eta_T\|_2$, we have

$$e^2 - 2et_1 - 4t_2V < 0.$$

This implies

$$e \le t_1 + \sqrt{t_1^2 + 4t_2V}.$$

Squaring and dividing by |T|, we get

$$\frac{1}{|T|}e^2 \le \frac{1}{|T|} \left[t_1^2 + 2t_1 \sqrt{t_1^2 + 4t_2 V} + t_1^2 + 4t_2 V \right]$$
$$= \frac{1}{|T|} \left[2t_1^2 + 2t_1 \sqrt{t_1^2 + 4t_2 V} + 4t_2 V \right].$$

A.4 MOVING TO THE L1 NORM

Using Lemma 4 and the Cauchy-Schwarz inequality, we can make the following corollary.

Corollary 1. If we set $\lambda = \frac{t_2}{|T|}$ in the optimization problem (2), then with probability $1 - \delta$ we have

$$\frac{1}{|T|}\|\eta_T - \hat{\eta}_T\|_1 \leq \sqrt{\frac{1}{|T|}\left[2t_1^2 + 2t_1\sqrt{t_1^2 + 4t_2V} + 4t_2V\right]},$$

where
$$t_1 = \sqrt{\frac{1}{2} \ln \frac{4}{\delta}}$$
 and $t_2 = \sqrt{\frac{|T|}{6} \ln \frac{4(|T|-1)}{\delta}}$.

A.5 BOUNDED VARIATION POPULATION TRANSFER

One limitation of Corollary 1 is that it only holds on the training set. We now use bounded variation to relate this to a population-based expectation. We start by proving the following Lemma.

Lemma 5. Consider a function $g:[0,1] \to \mathbb{R}$ of bounded variation V_g . Consider a sample s_1, \ldots, s_n sampled from a distribution with CDF F and a empirical CDF F_n , we have

$$\left| \mathbb{E}_{s}[g(s)] - \frac{1}{n} \sum_{i=1}^{n} g(s_{i}) \right| \leq V_{g} ||F - F_{n}||_{\infty}$$

Proof. Consider the integral

$$\int gdF - \int df F_n = \int gd(\underbrace{F - F_n}_{H}),$$

where we introduced the notation $H = F - F_n$. Integrating by parts, we obtain

$$\int gdH = -\int Hdg + \underbrace{g(1)H(1) + g(0)H(0)}_{0},$$

where the last two terms vanish because H(0) = H(1) = 0. This gives

$$\int gd(F - F_n) = -\int (F - F_n)dg.$$

Taking absolute values, we get

$$\left| \int gd(F - F_n) \right| = \left| \int (F - F_n)dg \right|.$$

Using the total variation assumption we have,

$$\left| \int (F - F_n) dg \right| \le V_g ||F - F_n||_{\infty}$$

We now combine Lemma 5 with corollary 1, obtaining a bound useful for our problem.

Corollary 2. With probability $1 - \delta$

$$\mathbb{E}_{s}[|\eta(s) - \hat{\eta}(s)|] \le \frac{1}{|T|} \sum_{i \in T} |\eta(s_i) - \hat{\eta}(s_i)| + (V + \hat{V}) \sqrt{\frac{\log(2/\delta)}{2|T|}}.$$

Proof. Apply Lemma 5 to the function $g(s) = |\eta(s) - \hat{\eta}(s)|$. Apply the DKW theorem to bound the term $||F - F_n||_{\infty}$ with probability $1 - \delta$. Apply the property that $V_g \leq V + \hat{V}$ by construction (where V is the total variation of $\hat{\eta}$ and \hat{V} is the total variation of $\hat{\eta}$).

B FINITE DERIVATIVE RESULTS

B.1 BOUNDING THE SMOOTHING ERROR UNDER CONTINUITY

Lemma 6 (Smoothing Error). Assume η is twice differentiable with the first and second derivatives uniformly bounded by b_1 and b_2 respectively. Consider a fixed design (set of scores $s_1, \ldots, s_{|T|}$). Denote by k_{s_i} a nonnegative kernel. Denote by $\hat{\eta}$ the Nadaraya-Watson smoothed approximation to η obtained from a dataset of |T| samples, defined as

$$\hat{\eta}(s') = \sum_{i \in T} w_i(s') y_i, \quad w_i(s') = \frac{k_{s_i}(s')}{\sum_{j \in T} k_{s_j}(s')}.$$

Then we have:

$$\mathbb{E}_{Y_T}\left[|\hat{\eta}(s') - \eta(s')|S_T\right] \le g_T(s'),$$

where

$$g_T(s') = b_1 \sum_{i \in T} w_i(s') |s' - s_i| + b_2 \frac{1}{2} \sum_{i \in T} w_i(s') (s' - s_i)^2 + \frac{1}{2} \sqrt{\sum_{i \in T} w_i^2(s')}.$$

Proof. We begin by decomposing the term we want to bound.

$$\hat{\eta}(s') - \eta(s') = \underbrace{\left(\sum_{i \in T} w_i(s')\eta(s_i)\right) - \eta(s')}_{\text{bias}} + \underbrace{\sum_{i \in T} w_i(s')(y_i - \eta(s_i))}_{\text{poise}}$$
(19)

We will bound the bias and noise term separately. We start with the bias term. Taking the Taylor expansion of η around training point s_i , evaluated at s', we get

$$\eta(s') = \eta(s_i) + \eta'(s_i)(s' - s_i) + \frac{1}{2}\eta''\xi_i(s')(s' - s_i)^2.$$

for some $\xi_i(s')$. Subtracting $\eta(s_i)$ and summing across the training points, we get

$$\sum_{i \in T} w_i(s')(\eta(s') - \eta(s_i)) = \sum_{i \in T} w_i(s')\eta'(s_i)(s' - s_i) + \frac{1}{2} \sum_{i \in T} w_i(s')\eta''(\xi_i(s'))(s' - s_i)^2.$$

Bounding, we get

$$\left| \sum_{i \in T} w_i(s') (\eta(s') - \eta(s_i)) \right| \le b_1 \sum_{i \in T} w_i(s') |s' - s_i| + b_2 \frac{1}{2} \sum_{i \in T} w_i(s') (s' - s_i)^2.$$

We now bound the noise term. Applying Jensen's inequality we have

$$\mathbb{E}_{Y_T} \left[\left| \sum_{i \in T} w_i(s')(y_i - \eta(s_i)) \right| \middle| S_T \right] \le \sqrt{\mathbb{E}_{Y_T} \left[\left(\sum_{i \in T} w_i(s')(y_i - \eta(s_i)) \right)^2 \middle| S_T \right]}.$$

The term $y_i - \eta(s_i)$ is zero mean and sub-Gaussian with variance at most 1/4. Hence we have

$$\sqrt{\mathbb{E}_{Y_T} \left[\left(\sum_{i \in T} w_i(s') (y_i - \eta(s_i)) \right)^2 \middle| S_T \right]} \le \frac{1}{2} \sqrt{\sum_{i \in T} w_i^2(s')}.$$

Putting both terms from (19) together, we get the result.

C Score Perturbation

C.1 BOUNDING THE DERIVATIVES OF THE CALIBRATION FUNCTION OF THE MODIFIED CLASSIFIER

We have a dataset of tuples (s_{orig}, y) where the components are generated as follows.

$$s_{
m orig} \sim p(s_{
m orig})$$

$$y \sim \text{Bernoulli}(\eta_{\text{orig}}(s_{\text{orig}}))$$

Here, $p(s_{\text{orig}})$ is an arbitrary probability measure for the scores $s_{\text{orig}} \in [0,1]$ (it can be discrete, continuous, or a mixture). The function $\eta_{\text{orig}}(s_{\text{orig}}) = \mathbb{P}(Y=1|S=s_{\text{orig}})$ is the calibration function of the original classifier (before perturbation).

Now, we create a new dataset of tuples (s_{orig}, s, y) where the components are generated by perturbing the original scores:

$$s_{\rm orig} \sim p(s_{\rm orig})$$

$$s \sim k(s \mid s_{ ext{orig}}) \ y \sim ext{Bernoulli}(\eta_{ ext{orig}}(s_{ ext{orig}}))$$

Here, for each original score s_{orig} , $k(s \mid s_{\text{orig}})$ is a probability density function (PDF) on [0, 1], which acts as a smoothing kernel. The joint distribution is given by

$$\begin{split} p(s_{\text{orig}}, s, y = 1) &= p(s_{\text{orig}}) \underbrace{\underbrace{k(s \mid s_{\text{orig}})}_{p(s \mid s_{\text{orig}})} \underbrace{\eta_{\text{orig}}(s_{\text{orig}})}_{p(y = 1 \mid s_{\text{orig}})} } \\ p(s_{\text{orig}}, s, y = 0) &= p(s_{\text{orig}}) \underbrace{k(s \mid s_{\text{orig}})}_{p(s \mid s_{\text{orig}})} \underbrace{(1 - \eta_{\text{orig}}(s_{\text{orig}}))}_{p(y = 0 \mid s_{\text{orig}})} \end{split}$$

We are interested in the calibration function of this new perturbed classifier, which is the conditional probability p(y=1|s). To find it, we first compute the marginals:

$$\begin{split} p(y=1,s) &= \int_0^1 k(s\mid s_{\text{orig}}) \eta_{\text{orig}}(s_{\text{orig}}) \, dp(s_{\text{orig}}) \;, \\ p(y=0,s) &= \int_0^1 k(s\mid s_{\text{orig}}) (1-\eta_{\text{orig}}(s_{\text{orig}})) \, dp(s_{\text{orig}}) \;, \\ p(s) &= p(y=1,s) + p(y=0,s) = \int_0^1 k(s\mid s_{\text{orig}}) (\eta_{\text{orig}}(s_{\text{orig}}) + 1-\eta_{\text{orig}}(s_{\text{orig}})) \, dp(s_{\text{orig}}) \\ &= \int_0^1 k(s\mid s_{\text{orig}}) \, dp(s_{\text{orig}}) \;. \end{split}$$

This allows us to compute the calibration function of the perturbed classifier, which we denote by $\eta(s)$:

$$\eta(s) = \frac{p(y=1,s)}{p(s)} = \frac{\int_0^1 \eta_{\text{orig}}(s_{\text{orig}}) k(s \mid s_{\text{orig}}) dp(s_{\text{orig}})}{\int_0^1 k(s \mid s_{\text{orig}}) dp(s_{\text{orig}})}, \ s \in [0,1].$$
 (20)

This expression is a kernel-weighted average of the original calibration function $\eta_{\text{orig}}(s_{\text{orig}})$.

C.1.1 HYPERBOLIC SECANT KERNEL

We use the hyperbolic secant function as our kernel. For a fixed bandwidth h > 0 and for each $s_{\text{orig}} \in [0, 1]$, we define the *sech* kernel $k(s \mid s_{\text{orig}})$:

$$Z(s_{\text{orig}}, h) = \int_0^1 \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) ds, \tag{21}$$

$$k(s \mid s_{\text{orig}}) = \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) = \frac{1}{Z(s_{\text{orig}}, h)} \frac{2}{e^{\frac{s - s_{\text{orig}}}{h}} + e^{-\frac{s - s_{\text{orig}}}{h}}}.$$
 (22)

First, we find a closed form equation for $Z(s_{\text{orig}}, h)$ and a lower bound of this quantity, which we will use in later proofs.

Lemma 7 (Closed form and lower bound for $Z(s_{\text{orig}}, h)$). For any h > 0 and $s_{\text{orig}} \in [0, 1]$, the normalization constant in (21) admits the closed form

$$Z(s_{orig}, h) = h \left[\arctan \left(\sinh \left(\frac{1 - s_{orig}}{h} \right) \right) + \arctan \left(\sinh \left(\frac{s_{orig}}{h} \right) \right) \right].$$
 (23)

Moreover, $Z(s_{orig}, h)$ is symmetric about $s_{orig} = \frac{1}{2}$, strictly maximized at $s_{orig} = \frac{1}{2}$, and attains its minimum at the endpoints $s_{orig} \in \{0, 1\}$. In particular,

$$\min_{s_{orig} \in [0,1]} Z(s_{orig}, h) = Z(0, h) = Z(1, h) = h \arctan\left(\sinh\left(\frac{1}{h}\right)\right). \tag{24}$$

Proof. With the substitution $u=(s-s_{\text{orig}})/h$ (so $ds=h\,du$), the limits $s\in[0,1]$ become $u\in[-s_{\text{orig}}/h,\,(1-s_{\text{orig}})/h]$, hence

$$Z(s_{\mathrm{orig}},h) = \int_0^1 \mathrm{sech}\Big(\frac{s-s_{\mathrm{orig}}}{h}\Big) \, ds = h \int_{-s_{\mathrm{orig}}/h}^{(1-s_{\mathrm{orig}})/h} \mathrm{sech}(u) \, du.$$

Using $\frac{d}{du}$ $\arctan(\sinh u) = \operatorname{sech} u$, we obtain

$$\begin{split} Z(s_{\text{orig}},h) &= h \Big[\arctan\big(\sinh u\big)\Big]_{u=-s_{\text{orig}}/h}^{u=(1-s_{\text{orig}})/h} \\ &= h \Big(\arctan\Big(\sinh\Big(\frac{1-s_{\text{orig}}}{h}\Big)\Big) + \arctan\Big(\sinh\Big(\frac{s_{\text{orig}}}{h}\Big)\Big)\Big), \end{split}$$

which is (23) (we used the oddness of sinh and arctan).

Now, we will find a lower bound for Z. First, differentiate (23) with respect to s_{orig} :

$$\frac{\partial}{\partial s_{\text{orig}}} Z(s_{\text{orig}}, h) = h \left(-\frac{1}{h} \operatorname{sech}\!\left(\frac{1 - s_{\text{orig}}}{h}\right) + \frac{1}{h} \operatorname{sech}\!\left(\frac{s_{\text{orig}}}{h}\right) \right) = \operatorname{sech}\!\left(\frac{s_{\text{orig}}}{h}\right) - \operatorname{sech}\!\left(\frac{1 - s_{\text{orig}}}{h}\right)$$

Since $x\mapsto \mathrm{sech}\,x$ is strictly decreasing on $[0,\infty)$, we have $\partial_{s_\mathrm{orig}}Z(s_\mathrm{orig},h)>0$ for $s_\mathrm{orig}<\frac{1}{2}$, $\partial_{s_\mathrm{orig}}Z(s_\mathrm{orig},h)=0$ at $s_\mathrm{orig}=\frac{1}{2}$, and $\partial_{s_\mathrm{orig}}Z(s_\mathrm{orig},h)<0$ for $s_\mathrm{orig}>\frac{1}{2}$. Hence Z is strictly increasing on $[0,\frac{1}{2}]$ and strictly decreasing on $[\frac{1}{2},1]$, so its maximum is at $s_\mathrm{orig}=\frac{1}{2}$ and its minimum is attained at the endpoints $s_\mathrm{orig}\in\{0,1\}$. Evaluating (23) at $s_\mathrm{orig}=0$ or $s_\mathrm{orig}=1$ yields

$$Z(0,h) = Z(1,h) = h \arctan\left(\sinh\left(\frac{1}{h}\right)\right),$$

which is strictly positive for every h > 0, proving the uniform lower bound (24).

Now, we show that $k(s \mid s_{\text{orig}})$ defined in Eq. (22) is a valid probability density function over [0, 1] in the following Lemma.

Lemma 8. For any given $s_{orig} \in [0,1]$ and h > 0, the function $k(s \mid s_{orig})$ defined in Eq. (22) is a valid probability density function over [0,1].

Proof. We must verify two conditions: non-negativity and that the integral over [0,1] is unity.

(1) Non-negativity. Recall that for any $t \in \mathbb{R}$, $\operatorname{sech}(t) \in (0,1]$ is continuous and strictly positive. Therefore the map $s \mapsto \operatorname{sech}\left((s-s_{\operatorname{orig}})/h\right)$ is continuous and strictly positive on [0,1]. Regarding the normalization constant, we showed in Lemma 7 that $Z(s_{\operatorname{orig}},h)>0$. Since the numerator in $k(s\mid s_{\operatorname{orig}})$ and the denominator $Z(s_{\operatorname{orig}},h)$ are both positive, we have

$$k(s \mid s_{\text{orig}}) = \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) \ge 0 \quad \text{for all } s \in [0, 1].$$

(2) Unit integral. Using the definition of $Z(s_{\text{orig}}, h)$,

$$\int_0^1 k(s\mid s_{\mathrm{orig}})\,ds = \frac{1}{Z(s_{\mathrm{orig}},h)}\int_0^1 \mathrm{sech}\Big(\frac{s-s_{\mathrm{orig}}}{h}\Big)\,ds = \frac{Z(s_{\mathrm{orig}},h)}{Z(s_{\mathrm{orig}},h)} \ = \ 1.$$

Since $k(\cdot \mid s_{\text{orig}}) \geq 0$ on [0,1] and integrates to 1, it is a valid probability density function on [0,1].

Substituting this specific kernel into our general formula (20) gives the explicit sech-perturbed calibration:

$$\eta(s) = \frac{\int_0^1 \eta_{\text{orig}}(s_{\text{orig}}) \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) dp(s_{\text{orig}})}{\int_0^1 \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) dp(s_{\text{orig}})}, \ s \in [0, 1].$$
(25)

C.1.2 A UNIFORM BOUND FOR THE FIRST DERIVATIVE

To analyze the smoothness of $\eta(s)$, we compute its first derivative. Let's define the numerator and denominator of (25) as separate functions:

$$N(s) = \int_{[0,1]} \eta_{\text{orig}}(s_{\text{orig}}) \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) dp(s_{\text{orig}}),$$

$$D(s) = \int_{[0,1]} \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) dp(s_{\text{orig}}),$$

so that $\eta(s) = N(s)/D(s)$. To ensure $\eta(s)$ is well-defined, we must show $D(s) \neq 0$. We show this in the following Lemma.

Lemma 9 (Positivity of the denominator). For the sech kernel, the denominator D(s) is strictly positive for all $s \in [0, 1]$ and any probability measure p on [0, 1].

Proof. Fix $s \in [0,1]$ and h > 0. For every $s_{\text{orig}} \in [0,1]$ we have $\operatorname{sech}((s - s_{\text{orig}})/h) > 0$ and $Z(s_{\text{orig}}, h) > 0$ (Lemma 7), hence

$$f(s_{\text{orig}}) \coloneqq \frac{1}{Z(s_{\text{orig}}, h)} \operatorname{sech}\left(\frac{s - s_{\text{orig}}}{h}\right) > 0 \quad \text{for all } s_{\text{orig}}.$$

Therefore $D(s) = \int_{[0,1]} f(s_{\text{orig}}) dp(s_{\text{orig}}) > 0$, since the integral of a strictly positive function with respect to a probability measure is strictly positive.

Here, we provide the derivative of the sech kernel.

Lemma 10 (Derivative of the sech kernel). For every $s_{orig} \in [0, 1]$, the partial derivative of $\operatorname{sech}\left(\frac{s-s_{orig}}{h}\right)$ with respect to s is

$$\frac{\partial}{\partial s}\operatorname{sech}\left(\frac{s-s_{orig}}{h}\right) = -\frac{1}{h}\operatorname{sech}\left(\frac{s-s_{orig}}{h}\right)\operatorname{tanh}\left(\frac{s-s_{orig}}{h}\right), \quad for \ s \in [0,1].$$
 (26)

Proof. Let $u := (s - s_{\text{orig}})/h$. Since $\operatorname{sech} u = 1/\cosh u$ and $(\cosh u)' = \sinh u$,

$$\frac{d}{du}\operatorname{sech} u = -\frac{\sinh u}{\cosh^2 u} = -\operatorname{sech} u \tanh u.$$

By the chain rule with $\partial u/\partial s = 1/h$,

$$\frac{\partial}{\partial s} \operatorname{sech}\!\left(\frac{s-s_{\operatorname{orig}}}{h}\right) = -\frac{1}{h} \operatorname{sech}\!\left(\frac{s-s_{\operatorname{orig}}}{h}\right) \tanh\!\left(\frac{s-s_{\operatorname{orig}}}{h}\right),$$

which proves (26) for all $s \in [0, 1]$.

To compute the derivatives of N(s) and D(s), we need to differentiate under the integral sign. The following Lemma provides the justification.

Lemma 11 (Differentiation under the integral sign). The derivatives of N(s) and D(s) can be computed as:

$$N'(s) = \int_{[0,1]} \eta_{orig}(s_{orig}) \frac{1}{Z(s_{orig}, h)} \frac{\partial}{\partial s} \operatorname{sech}\left(\frac{s - s_{orig}}{h}\right) dp(s_{orig}), \tag{27}$$

$$D'(s) = \int_{[0,1]} \frac{1}{Z(s_{orig}, h)} \frac{\partial}{\partial s} \operatorname{sech}\left(\frac{s - s_{orig}}{h}\right) dp(s_{orig}). \tag{28}$$

Proof. Let us start by focusing on N(s). First, we need to find $\frac{\partial}{\partial s}N(s)$. Let's represent N(s) as $N(s) = \int_{[0,1]} f(s_{\text{orig}}, s) dp(s_{\text{orig}})$, where

$$f(s_{\rm orig},s) = \eta_{\rm orig}(s_{\rm orig}) \, \frac{1}{Z(s_{\rm orig},h)} \, {\rm sech}\bigg(\frac{s-s_{\rm orig}}{h}\bigg) \, .$$

Now recall by Lemma 10 (using the oddness of tanh),

$$\frac{\partial}{\partial s} \operatorname{sech} \left(\frac{s - s_{\operatorname{orig}}}{h} \right) \; = \; \frac{1}{h} \operatorname{sech} \left(\frac{s - s_{\operatorname{orig}}}{h} \right) \tanh \left(\frac{s_{\operatorname{orig}} - s}{h} \right), \quad \text{for } s \in [0, 1].$$

Therefore, we notice that for every s we have

$$\begin{split} \left| \frac{\partial}{\partial s} f(s_{\text{orig}}, s) \right| &= \left| \eta_{\text{orig}}(s_{\text{orig}}) \, \frac{1}{Z(s_{\text{orig}}, h)} \frac{1}{h} \, \text{sech} \bigg(\frac{s - s_{\text{orig}}}{h} \bigg) \, \text{tanh} \bigg(\frac{s_{\text{orig}} - s}{h} \bigg) \right| \\ &\leq \frac{1}{h Z(s_{\text{orig}}, h)} =: g(s_{\text{orig}}) \; . \end{split}$$

Recall that $Z(s_{\text{orig}},h)$ is strictly positive and lower bounded $Z(s_{\text{orig}},h) \geq h \arctan\left(\sinh\left(\frac{1}{h}\right)\right)$ from Lemma 7. Because of this and the fact that h>0, we have $\sup_{s_{\text{orig}}\in[0,1]}1/hZ(s_{\text{orig}},h)<\infty$. Hence g is bounded, so we can apply the dominated convergence theorem (Theorem 11.32 in (Rudin, 1976)), and we may pass the derivative through the integral. An analogous reasoning easily applies to D(s) too.

By the quotient rule, the derivative of $\eta(s)$ is:

$$\eta'(s) = \frac{N'(s)D(s) - N(s)D'(s)}{D(s)^2}. (29)$$

Alternative Representation via Covariance. The derivative has an elegant alternative representation. Let's define a new probability measure $w_s(ds_{\text{orig}})$ on [0,1] that depends on s. This measure re-weights the original measure $p(s_{\text{orig}})$ based on the kernel:

$$w_s(ds_{\text{orig}}) = \frac{\frac{1}{Z(s_{\text{orig}},h)} \operatorname{sech}\left(\frac{s-s_{\text{orig}}}{h}\right)}{D(s)} dp(s_{\text{orig}}).$$
(30)

Note that
$$\int_0^1 w_s(ds_{\text{orig}}) = \frac{1}{D(s)} \int_0^1 \frac{1}{Z(s_{\text{orig}},h)} \operatorname{sech}\left(\frac{s-s_{\text{orig}}}{h}\right) dp(s_{\text{orig}}) = \frac{D(s)}{D(s)} = 1.$$

Now, we use this alternative representation to prove that the derivative of $\eta(s)$ can be expressed as a covariance under w_s in the following Proposition.

Lemma 12. The derivative of the perturbed calibration function can be expressed as a covariance under the measure $w_s(ds_{orig})$:

$$\eta'(s) = \frac{1}{h} \operatorname{Cov}_{S \sim w_s} \left(\eta_{orig}(S), \tanh\left(\frac{S-s}{h}\right) \right), \quad s \in [0, 1].$$

Proof. Starting from (29) and writing $T(u) := \tanh\left(\frac{s-u}{h}\right)$, Lemma 11 together with Lemma 10 gives

$$\begin{split} \frac{N'(s)}{D(s)} &= \frac{1}{D(s)} \int_{[0,1]} \eta_{\text{orig}}(s_{\text{orig}}) \, \frac{1}{Z(s_{\text{orig}},h)} \frac{\partial}{\partial s} \, \text{sech}\!\!\left(\frac{s-s_{\text{orig}}}{h}\right) dp(s_{\text{orig}}) \\ &= -\frac{1}{h} \, \mathbb{E}_{w_s} \big[\eta_{\text{orig}}(S) \, T(S)\big], \end{split}$$

and

$$\frac{D'(s)}{D(s)} = \frac{1}{D(s)} \int_{[0,1]} \frac{1}{Z(s_{\text{orig}},h)} \frac{\partial}{\partial s} \operatorname{sech}\!\!\left(\frac{s-s_{\text{orig}}}{h}\right) dp(s_{\text{orig}}) = -\frac{1}{h} \,\, \mathbb{E}_{w_s}\!\left[T(S)\right].$$

Since $\eta(s) = \frac{N(s)}{D(s)} = \mathbb{E}_{w_s}[\eta_{\text{orig}}(S)]$, the quotient rule (29) yields

$$\begin{split} \eta'(s) &= \frac{N'(s)}{D(s)} - \frac{N(s)}{D(s)} \frac{D'(s)}{D(s)} \\ &= -\frac{1}{h} \; \mathbb{E}_{w_s} \big[\eta_{\text{orig}}(S) \, T(S) \big] \; + \; \frac{1}{h} \; \mathbb{E}_{w_s} [\eta_{\text{orig}}(S)] \; \mathbb{E}_{w_s} [T(S)] \\ &= -\frac{1}{h} \Big(\, \mathbb{E}_{w_s} [\eta_{\text{orig}}(S) \, T(S)] - \mathbb{E}_{w_s} [\eta_{\text{orig}}(S)] \; \mathbb{E}_{w_s} [T(S)] \Big) \\ &= -\frac{1}{h} \; \text{Cov}_{S \sim w_s} \Big(\eta_{\text{orig}}(S), \; \text{tanh} \Big(\frac{s-S}{h} \Big) \Big). \end{split}$$

Because tanh is odd, this is equivalently

$$\eta'(s) = \frac{1}{h} \operatorname{Cov}_{S \sim w_s} \left(\eta_{\operatorname{orig}}(S), \tanh \left(\frac{S-s}{h} \right) \right).$$

Using this alternative representation as a covariance, we provide a uniform bound for the first derivative in the following Corollary.

Corollary 3 (Uniform Lipschitz bound). For the sech kernel and any bandwidth h > 0, the perturbed calibration function η is Lipschitz continuous with a constant independent of p and η :

$$\sup_{s \in [0,1]} |\eta'(s)| \le \frac{\tanh(1/h)}{2h} \le \frac{1}{2h}.$$

Proof. By Lemma 12, $\eta'(s) = \frac{1}{h} \operatorname{Cov}_{w_s} \left(\eta_{\operatorname{orig}}(S), T \right)$ with $T = \tanh \left(\frac{S-s}{h} \right)$. By Cauchy–Schwarz, we know that

$$|\operatorname{Cov}(X,Y)| \le \sqrt{\mathbb{V}(X)\,\mathbb{V}(Y)},$$

where with $\mathbb{V}(\cdot)$ we indicate the variance of a random variable.

Also, Popoviciu's inequality says that, if $Z \in [a, b]$, then we have a bound for its variance $\mathbb{V}(Z)$:

$$\mathbb{V}(Z) \le \frac{(b-a)^2}{4}.$$

Since $\eta_{\text{orig}}(S) \in [0,1]$ and $T \in [-\tanh(1/h), \tanh(1/h)]$, Cauchy–Schwarz plus Popoviciu's inequality gives

$$\frac{1}{h}|\text{Cov}(\eta_{\text{orig}}(S), T)| \leq \frac{1}{h}\sqrt{\frac{1}{4} \cdot \frac{(2\tanh(1/h))^2}{4}} = \frac{1}{h}\sqrt{\frac{\tanh^2(1/h)}{4}} = \frac{\tanh(1/h)}{2h}.$$

Noticing that $tanh(1/h) \le 1$ yields the claim.

C.1.3 SECOND DERIVATIVE AND A UNIFORM BOUND

Let us focus now on bounding the second derivative. First, we will find the second derivative of the kernel in the next Lemma.

Lemma 13 (Second derivative of the kernel). For every $s_{orig} \in [0, 1]$ and h > 0,

$$\frac{\partial^2}{\partial s^2} \operatorname{sech}\left(\frac{s-s_{orig}}{h}\right) = \frac{1}{h^2} \operatorname{sech}\left(\frac{s-s_{orig}}{h}\right) \left(2 \tanh^2\left(\frac{s-s_{orig}}{h}\right) - 1\right), \quad s \in [0,1].$$

Proof. Let $u = (s - s_{\text{orig}})/h$. Using $\frac{d}{du} \operatorname{sech} u = -\operatorname{sech} u \tanh u$ and $\frac{d}{du} \tanh u = \operatorname{sech}^2 u$,

$$\frac{d^2}{du^2}\operatorname{sech} u = -\frac{d}{du}(\operatorname{sech} u \, \tanh u) = \operatorname{sech} u \, \tanh^2 u - \operatorname{sech} u \, \operatorname{sech}^2 u = \operatorname{sech} u \, (2 \tanh^2 u - 1).$$

Chain rule gives the claim.

Now, we justify differentiation under the integral sign for second-order derivatives.

Lemma 14 (Differentiation under the integral sign (second order)). For fixed h > 0 and all $s \in [0, 1]$,

$$N''(s) = \int_{[0,1]} \eta_{orig}(s_{orig}) \frac{1}{Z(s_{orig}, h)} \frac{\partial^2}{\partial s^2} \operatorname{sech}\left(\frac{s - s_{orig}}{h}\right) dp(s_{orig}),$$

$$D''(s) = \int_{[0,1]} \frac{1}{Z(s_{orig}, h)} \frac{\partial^2}{\partial s^2} \operatorname{sech}\left(\frac{s - s_{orig}}{h}\right) dp(s_{orig}).$$

Proof. Similarly to the proof of Lemma 11, we want to apply the dominated convergence theorem. By Lemma 13, the integrands are bounded in absolute value by $\frac{1}{h^2Z(s_{\text{orig}},h)}$. Also, we know from Lemma 7 that $\min_{s_{\text{orig}} \in [0,1]} Z(s_{\text{orig}},h) = h \arctan\left(\sinh\left(\frac{1}{h}\right)\right) > 0$. Thus $\frac{1}{h^2Z(s_{\text{orig}},h)}$ is an integrable dominator, so the dominated convergence theorem applies (Theorem 11.32 from (Rudin, 1976)), allowing differentiation under the integral.

Similarly to what we shown for the first derivative, we show that the second derivative can be represented as a covariance under w_s in the following Lemma.

Lemma 15 (Second derivative in covariance form). Let $T(S) = \tanh\left(\frac{S-s}{h}\right)$ and expectation/covariance under $S \sim w_s$, where w_s is defined as in Eq. 30. We have that:

$$\eta''(s) = \frac{1}{h^2} \Big(\operatorname{Cov}_{w_s} \big(\eta_{orig}(S), \, 2T(S)^2 - 1 \big) - 2 \, \mathbb{E}_{w_s} [T(S)] \, \operatorname{Cov}_{w_s} \big(\eta_{orig}(S), \, T(S) \big) \Big).$$

Proof. First, notice that, since D(s) > 0 (Lemma 9), all ratios are well-defined. By Lemma 14 and Lemma 13, using that $T(S)^2 = \tanh^2\left(\frac{S-s_{\text{orig}}}{h}\right) = \tanh^2\left(\frac{s_{\text{orig}}-S}{h}\right)$,

$$\frac{N''(s)}{D(s)} = \frac{1}{h^2} \; \mathbb{E}_{w_s} \big[\eta_{\text{orig}}(S) \, (2T(S)^2 - 1) \big], \qquad \frac{D''(s)}{D(s)} = \frac{1}{h^2} \; \mathbb{E}_{w_s} \big[2T(S)^2 - 1 \big].$$

From Lemma 11 and Lemma 10,

$$\frac{N'(s)}{D(s)} = \frac{1}{h} \ \mathbb{E}_{w_s}[\eta_{\mathrm{orig}}(S)T(S)], \qquad \frac{D'(s)}{D(s)} = \frac{1}{h} \ \mathbb{E}_{w_s}[T(S)], \qquad \eta(s) = \frac{N(s)}{D(s)} = \mathbb{E}_{w_s}[\eta_{\mathrm{orig}}(S)].$$

We now use the quotient rule for the second derivative:

$$\eta'' = \frac{N''D - ND''}{D^2} - 2\,\frac{(N'D - ND')D'}{D^3} = \left(\tfrac{N''}{D} - \eta\,\tfrac{D''}{D}\right) - 2\left(\tfrac{N'}{D} - \eta\,\tfrac{D'}{D}\right)\tfrac{D'}{D}.$$

Substituting the four identities above yields

$$\eta''(s) = \frac{1}{h^2} \Big(\mathbb{E}_{w_s} [\eta_{\text{orig}}(S)(2T(S)^2 - 1)] - \mathbb{E}_{w_s} [\eta_{\text{orig}}(S)] \mathbb{E}_{w_s} [2T(S)^2 - 1] \Big)$$
$$- \frac{2}{h^2} \Big(\mathbb{E}_{w_s} [\eta_{\text{orig}}(S)T(S)] - \mathbb{E}_{w_s} [\eta_{\text{orig}}(S)] \mathbb{E}_{w_s} [T(S)] \Big) \mathbb{E}_{w_s} [T(S)].$$

Recognizing covariances,

$$\mathbb{E}[\eta(2T(S)^2 - 1)] - \mathbb{E}[\eta_{\text{orig}}(S)] \,\mathbb{E}[2T(S)^2 - 1] = \text{Cov}(\eta_{\text{orig}}(S), \, 2T(S)^2 - 1),$$

$$\mathbb{E}[\eta_{\text{orig}}(S)T(S)] - \mathbb{E}[\eta_{\text{orig}}(S)] \,\mathbb{E}[T(S)] = \text{Cov}(\eta_{\text{orig}}(S), \, T(S)),$$

we obtain

$$\eta''(s) = \frac{1}{h^2} \Big(\operatorname{Cov}_{w_s} \big(\eta_{\operatorname{orig}}(S), \, 2T^2(S) - 1 \big) - 2 \, \mathbb{E}_{w_s}[T(S)] \, \operatorname{Cov}_{w_s} \big(\eta_{\operatorname{orig}}(S), \, T(S) \big) \Big).$$

Corollary 4 (Uniform second-derivative bound). For the sech kernel and any bandwidth h > 0,

$$\sup_{s \in [0,1]} |\eta''(s)| \ \leq \ \frac{3}{2} \, \frac{\tanh^2(1/h)}{h^2} \ \leq \ \frac{3}{2} \, \frac{1}{h^2}.$$

Proof. Let $\tau := \tanh(1/h)$. Then $T(S) := \tanh\left(\frac{S-s}{h}\right) \in [-\tau, \tau]$. Hence, by Popoviciu's inequality,

$$\mathbb{V}_{w_s}(T(S)) \leq \frac{(2\tau)^2}{4} = \tau^2, \quad 2T(S)^2 - 1 \in [-1, \, 2\tau^2 - 1] \ \Rightarrow \ \mathbb{V}_{w_s}(2T(S)^2 - 1) \leq \frac{(2\tau^2)^2}{4} = \tau^4.$$

Since $\eta_{\text{orig}}(S) \in [0, 1], \mathbb{V}(\eta_{\text{orig}}(S)) \leq 1/4$. Using Cauchy–Schwarz,

$$\begin{split} |\mathrm{Cov}(\eta, 2T^2 - 1)| &\leq \sqrt{\mathbb{V}(\eta)\,\mathbb{V}(2T^2 - 1)} \leq \tfrac{1}{2}\,\tau^2, \qquad |\mathrm{Cov}(\eta, T)| \leq \sqrt{\mathbb{V}(\eta)\,\mathbb{V}(T)} \leq \tfrac{1}{2}\,\tau, \\ \text{and } |\mathbb{E}[T]| &\leq \tau. \text{ Therefore, Lemma 15 gives} \end{split}$$

$$|\eta''(s)| \le \frac{1}{h^2} \left(\frac{1}{2}\tau^2 + 2\tau \cdot \frac{1}{2}\tau\right) = \frac{3}{2}\frac{\tau^2}{h^2} \le \frac{3}{2}\frac{1}{h^2}.$$

C.2 WHY NOT A TRUNCATED GAUSSIAN?

Another plausible choice for the kernel would have been the (truncated) Gaussian. Let $\phi(t)=\frac{1}{\sqrt{2\pi}}e^{-t^2/2}$ and $\Phi(t)=\int_{-\infty}^t\phi(u)\,du$ be the standard normal pdf and cdf. Fix a bandwidth h>0. For each $s_{\rm orig}\in[0,1]$ define the *truncated Gaussian* density on [0,1] centered at $s_{\rm orig}$,

$$Z(s_{\text{orig}}, h) = \Phi\left(\frac{1 - s_{\text{orig}}}{h}\right) - \Phi\left(\frac{-s_{\text{orig}}}{h}\right) \in (0, 1), \tag{31}$$

$$k(s \mid s_{\text{orig}}) = \frac{1}{h Z(s_{\text{orig}}, h)} \phi\left(\frac{s - s_{\text{orig}}}{h}\right). \tag{32}$$

However, in this paper we decided not to use such a kernel because of the rate of its derivative. In particular, we have that:

$$\frac{\partial}{\partial s} k(s \mid s_{\mathrm{orig}}) \; = \; - \; \frac{s - s_{\mathrm{orig}}}{h^2} \; k(s \mid s_{\mathrm{orig}}) = \mathcal{O}\left(\frac{1}{h^2}\right).$$

By following a similar reasoning to the one done for the sech kernel (using the dominated convergence theorem, representing the derivative of η as a covariance, and bounding the covariance), we would have got a bound in the order of $\mathcal{O}\left(\frac{1}{h^2}\right)$ for the first derivative, which is worse than the one we obtained for the sech kernel $(\mathcal{O}\left(\frac{1}{h}\right))$. For this theoretical reason, we opted for the more suitable sech kernel.

D IMPLEMENTATION DETAILS

Plug-in bandwidth For the kernel k' we pick the bandwidth by minimizing a one-dimensional bias-variance surrogate obtained from the h-dependent parts of the bound. With the Epanechnikov kernel, one can derive

$$a = \frac{3}{8}b_1, \ b = \frac{1}{10}b_2, \ c = \frac{1.15}{2\sqrt{2n_{\text{tr}}}},$$

so that optimizing $2bt^5 + at^3 - \frac{1}{2}c = 0$, gives the optimal $h = t^2$. We solve by a few Newton steps and clip to $[10^-4, 1/4]$. This choice is an implementation detail; the theory only requires convex weights summing to one.

Validation subsampling For speed we evaluate the held-out quantities on a uniformly sampled subset $V_k' \subseteq V_k$ per fold, with size $\min\{15000, \max(5000, 0.05 n)\}$. Each validation point is still scored by a model trained without it, so the conditional independence and i.i.d. assumptions required by the empirical-Bernstein step remain valid; subsampling only increases variance (slightly loosening the bound) and does not alter δ -allocation or correctness.

Shift aggregation in Bucketing. We evaluate the Lipschitz+bucketing bound for every candidate partition defined by a bucket count B and a fractional shift $r \in \{0, \dots, R-1\}$ (offset r/B); using a per-candidate confidence level $\delta/(BR)$, we take the *minimum* bound over all (B,r). By a union bound, this aggregate is valid at overall level δ .

E IMPLEMENTATION OF THE PERTURBATION

One advantage of our perturbation method is that, unlike bucketing, it is compatible with training by back-prop. This means we can perturb during classifier training, increasing resilience to perturbation at test time. To do that we change the standard cross entropy loss to.

$$L = -\frac{1}{N} \left[\sum_{i=1}^{N} y_i \mathbb{E} \log s_{1,i} + (1 - y_i) \mathbb{E} \log s_{0,i} \right] + L_m, \tag{33}$$

Where y_i are the ground-truth binary labels, $\mathbb{E} \log s_{1,i}$ and $\mathbb{E} \log s_{0,i}$ are the expectation of the two log perturbed probabilities, computed using Gauss Legendre quadrature in practice.

The loss L_m directly compares the logits of pairs of samples across classes, ensuring class separation. Specifically, for each example i in a batch, we first consider the logit difference

$$z_i = \hat{l}_1 - \hat{l}_0,$$

where \hat{l}_1 and \hat{l}_0 denote the logits for the positive and negative classes, respectively. We now restrict attention to near pairs (i,j), defined by the condition $|s^i-s^j|<2h$ i.e. where the probabilities are close. For each such pair, the margin gap is given by

$$g_{ij} = z^i - z^j.$$

The loss penalizes cases where this margin is smaller than the target 2h, using a smooth softplus penalty. Formally, the loss is defined as

$$L_m = \frac{1}{|C|} \sum_{(i,j)\in\mathcal{N}} \log(1 + \exp(2h - g_{ij})),$$

where C is the set of near positive–negative pairs. This formulation encourages the logit difference of positive samples relative to negatives to be at least 2h, thereby pushing the classifier towards maintaining a larger decision margin.

F DATASETS

We perform perturbation experiments on three binary classification tasks:

- **IMDb:** A text data set of movie reviews with associated binary labels for sentiment (positive and negative) (Maas et al., 2011). For this task we fined-tuned a pre-trained BERT back-bone with a sequence classification head. We used 1,000 training examples, 250 validation examples and 500 test examples to compute reported metrics.
- **Spam Detection:** A text dataset of emails labeled as either "spam" or "not-spam" Schulte et al. (2024). As for the IMDb data set, we fined-tuned a pre-trained BERT back-bone with a sequence classification head, using 1,000 training examples, 250 validation examples and 500 test examples.
- **CIFAR-10:** An image data set containing 10 distinct classes: airplane, automobile, bird, cat, deer, dog, frog, horse, ship and truck (Krizhevsky, 2009). We make a binary task on this dataset by grouping means of transport (airplane, automobile, ship, truck) and animals (bird, cat, deer, dog, frog, horse). We fine-tune a pre-trained visual transformer (ViT) with a classification head. We use 2,500 training examples, 250 validation examples and 500 test examples.
- Amazon Polarity: A large-scale text dataset of Amazon product reviews labeled with binary sentiment (positive and negative) (McAuley & Leskovec, 2013).

G LLM USAGE

In accordance with the ICLR LLM usage policy, we disclose that large language models (LLMs) were used as tools during this project. First, LLMs were employed to draft concise *code prototypes from mathematical descriptions* (e.g., turning a specification of the bounds into runnable Python). All such code was subsequently scrutinized by the authors; the final implementations and experimental results are the product of human verification. Second, LLMs were used in an *iterative ideation loop for conceptualizing proofs*. The final version of formal statements, proofs, as well as the entirety of the paper's exposition were written by humans. Importantly, the paper's *central idea*—that injecting noise yields a *smooth* calibration function enabling tighter bounds—was conceived by a human author; LLMs did not originate this insight. LLMs were also used to polish the writing and as tools to find relevant related work.