Don't Just Blame Over-parametrization for Over-confidence: Theoretical Analysis of Calibration in Binary Classification

Yu Bai [1] **Song Mei** [2] **Huan Wang** [1] **Caiming Xiong** [1]

Abstract

Modern machine learning models with high accuracy are often miscalibrated—the predicted top probability does not reflect the actual accuracy, and tends to be *over-confident*. It is commonly believed that such over-confidence is mainly due to *over-parametrization*, in particular when the model is large enough to memorize the training data and maximize the confidence.

In this paper, we show theoretically that over-parametrization is not the only reason for over-confidence. We prove that *logistic regression is inherently over-confident*, in the realizable, under-parametrized setting where the data is generated from the logistic model, and the sample size is much larger than the number of parameters. Further, this over-confidence happens for general well-specified binary classification problems as long as the activation is symmetric and concave on the positive part. Perhaps surprisingly, we also show that over-confidence is not always the case—there exists another activation function (and a suitable loss function) under which the learned classifier is *under-confident* at some probability values. Overall, our theory provides a precise characterization of calibration in realizable binary classification, which we verify on simulations and real data experiments.

1. Introduction

Modern machine learning models such as deep neural networks with high accuracy tend to be miscalibrated: The predicted top probability (*confidence*) does not reflect the actual accuracy of the model, and tends to be *over-confident*. For example, a WideResNet 32 on CIFAR100 has on average a predicted top probability of 87%, while the actual test accuracy is only 72% (Guo et al., 2017). As the confidence is often comprehended as an estimate of the true accuracy, such over-confidence could be dangerous, especially in risk-sensitive domains such as medical AI (Begoli et al., 2019), self-driving cars (Michelmore et al., 2018), and so on. To address this issue, there is a growing line of research on improving the calibration of models, by either performing *recalibration* of well-trained models to adjust the confidence scores (Platt et al., 1999; Zadrozny & Elkan; Naeini et al., 2015; Guo et al., 2017), or by averaging the predictions over multiple models to make the confidence scores more accurate (Lakshminarayanan et al., 2016; Gal & Ghahramani, 2016). These methods in general can reduce the over-confidence and improve the calibration of the model, while preserving (or even improving) the model's accuracy (Ovadia et al., 2019).

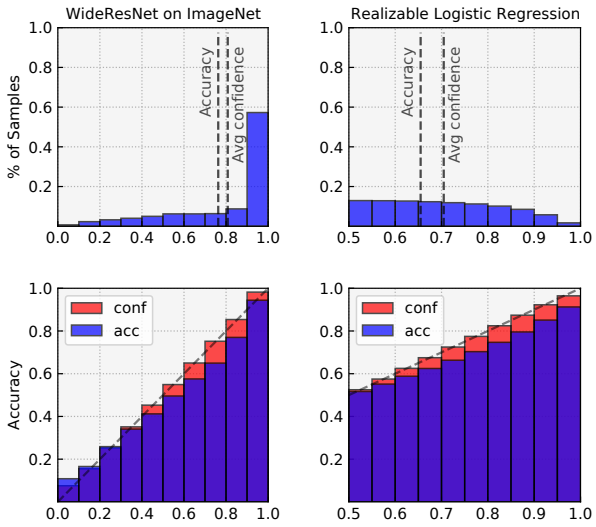

Figure 1. Reliability diagrams for calibration: Over-parametrized deep network vs. well-specified, under-parametrized logistic regression. The x-axes denote the confidences (predicted top probabilities) of the models. Left: 50-layer WideResNet on ImageNet. Right: Binary logistic regression on simulated data with $n = 2000$ and $d = 100$.

Despite these progresses, the more fundamental question of *why* such over-confidence happens for vanillaly trained

[1]Salesforce Research [2]University of California, Berkeley. Correspondence to: Yu Bai <yu.bai@salesforce.com>, Song Mei <songmei@berkeley.edu>.

Proceedings of the 38^{th} International Conference on Machine Learning, PMLR 139, 2021.

models remains not satisfactorily understood. One common understanding is that over-confidence is a result of *over-parametrization*: Models such as deep neural networks are large enough to memorize the entire training dataset, and are encouraged to magnify its weights and maximize the confidence so as to minimize the training loss (Mukhoti et al., 2020). Guo et al. (2017) also observed that increasing the depth and width makes the over-confident more severe, even when this improves the accuracy. However, so far it is unclear whether over-parametrization is the only reason, or whether there are other intrinsic reasons leading to over-confidence.

In this paper, we show that over-confidence is not just a result of over-parametrization and is more inherent. We conduct a precise theoretical study on the calibration in binary classification problems. Our main result shows that *standard logistic regression is also over-confident*, even in the well-specified, under-parametrized scenario where the model is correct (data generated from a linear logistic model), and there is abundant data (number of samples n much greater than number of parameters d).

Figure 1 illustrates our main finding via simulation: Similar to an over-parametrized neural network, the empirical risk minimizer of logistic regression is also over-confident at all confidence levels. Note that these two models have rather different behaviors in terms of the distribution of confidences, yet their over-confidence behaviors are similar.

Our contributions are summarized as follows:

- We show that *well-specified logistic regression is inherently over-confident*: Conditioned on the model predicting $p > 0.5$, the actual probability of the label being one is lower by an amount of $\Theta(d/n)$, in the limit of $n, d \to \infty$ proportionally and n/d is large (Section 3). In other words, the calibration error is always in the over-confident direction. We also show that the overall Calibration Error (CE) of the logistic model is $\Theta(d/n)$ in this limiting regime.

- We identify sufficient conditions for over- and under-confidence in general binary classification problems, where the data is generated from an arbitrary nonlinear activation, and we solve a well-specified empirical risk minimization (ERM) problem with a suitable loss function (Section 4). Our conditions imply that any symmetric, monotone activation $\sigma : \mathbb{R} \to [0, 1]$ that is *concave* at all $z > 0$ will yield a classifier that is over-confident at any confidence level.

- Another perhaps surprising implication is that *over-confidence is not universal*: We prove that there exists an activation function for which under-confidence can happen for a certain range of confidence levels.

- We perform simulation and real data experiments to test our theory (Section 5). Our experiments suggest that the over-confidence of logistic regression happens broadly in a variety of under-parametrized settings, within or beyond our theory's assumptions. We also verify that under-confidence can indeed happen in simulations with the activation function constructed above.

- On the technical end, our analysis develops a precise understanding of the high-dimensional proportional limit of ERM in the *sufficient data* regime (n/d is large) by rigorously establishing the first-order behavior of the solution to the characterizing system of nonlinear equations (Section 6), which may be of broader interest.

1.1. Related work

Algorithms for model calibration Practitioners have observed and dealt with the over-confidence of logistic regression long ago. *Recalibration algorithms* fix this by adjusting the output of a well-trained model, and dates back to the classical methods of Platt scaling (Platt et al., 1999), histogram binning (Zadrozny & Elkan) and isotonic regression (Zadrozny & Elkan, 2002). Platt et al. (1999) also uses a particular kind of label smoothing as a way of mitigating the over-confidence in logistic regression. Guo et al. (2017) show that temperature scaling, a simple method that learns a rescaling factor for the logits, is a competitive method for calibrating neural networks. A number of recent recalibration methods further improve the performances over these approaches (Kull et al., 2017; 2019; Ding et al., 2020; Rahimi et al., 2020; Zhang et al., 2020).

Another line of work improves calibration by aggregating the probabilisitic predictions over multiple models, using either an ensemble of models (Lakshminarayanan et al., 2016; Malinin et al., 2019; Wen et al., 2020; Tran et al., 2020), or randomized predictions such as Bayesian neural networks (Gal & Ghahramani, 2016; Gal et al., 2017; Maddox et al., 2019; Dusenberry et al., 2020). Finally, there are techniques for improving the calibration of a single neural network during training (Thulasidasan et al., 2019; Mukhoti et al., 2020; Liu et al., 2020).

Theoretical analysis of calibration Kumar et al. (2019) show that continuous rescaling methods such as temperature scaling is less calibrated than reported, and proposed a method that combines temperature scaling and histogram binning. Gupta et al. (2020) study the relationship between calibration and other notions of uncertainty such as confidence intervals. Shabat et al. (2020); Jung et al. (2020) study the sample complexity of estimating the multicalibration error (group calibration). A related theoretical result to ours is (Liu et al., 2019) which shows that the calibration error

of any classifier is upper bounded by its square root excess logistic loss over the Bayes classifier. This result can be translated to a $O(\sqrt{d/n})$ upper bound for well-specified logistic regression, whereas our main result implies $\Theta(d/n)$ calibration error in our high-dimensional limiting regime (with input distribution assumptions).

High-dimensional behaviors of empirical risk minimization There is a rapidly growing literature on limiting characterizations of convex optimization-based estimators in the $n \propto d$ regime (Donoho et al., 2009; Bayati & Montanari, 2011; El Karoui et al., 2013; Karoui, 2013; Stojnic, 2013; Thrampoulidis et al., 2015; 2018; Mai et al., 2019; Sur & Candès, 2019; Candès et al., 2020). Our analysis builds on the characterization for unregularized convex risk minimization problems (including logistic regression) derived in Sur & Candès (2019).

2. Preliminaries

In this paper we consider binary classification problems, where we observe n data points $\{(\mathbf{x}_i, y_i)\}_{i=1}^n \overset{\text{iid}}{\sim} P$ for some distribution P on $\mathbb{R}^d \times \{0,1\}$.

2.1. Calibration

Let $\widehat{f} : \mathbb{R}^d \to [0,1]$ be a (probabilistic) classifier. $\widehat{f}$ is said to be perfectly calibrated if $\mathbb{P}(Y = 1 | \widehat{f}(\mathbf{X}) = p) = p$ for all $p \in [0,1]$, that is, the actual probability of $Y = 1$ conditioned on $\widehat{f}$ predicting p is exactly p. In reality, we cannot hope for obtaining perfect calibration, and would rather desire ways of measuring the calibration error.

A standard metric is the Calibration Error (CE), which measures the difference between the prediction and the conditional mean of Y given the prediction (Guo et al., 2017):

$$\mathrm{CE}(\widehat{f}) := \mathbb{E}_{(\mathbf{X},Y)\sim P}\Big[\Big|\widehat{f}(\mathbf{X}) - \mathbb{E}[Y \mid \widehat{f}(\mathbf{X})]\Big|\Big]. \tag{1}$$

Notably, CE is the population (unbinned) version of the Expected Calibration Error (ECE), a commonly used calibration metric in recent work (Naeini et al., 2015; Guo et al., 2017; Ovadia et al., 2019; Nixon et al., 2019).

In this paper, we consider the *calibration error of $\widehat{f}$ at level p*:

$$\Delta_p^{\mathsf{cal}}(\widehat{f}) := p - \mathbb{P}_{(X,Y)\sim P}\Big(Y = 1 \mid \widehat{f}(\mathbf{X}) = p\Big) \tag{2}$$

for all $p \in (0,1)$. Note that $\Delta_p^{\mathsf{cal}}(\widehat{f})$ is the quantity inside the expectation in (1), and provides a more fine-grained characterization of the calibration error by specifying which p we are interested in.

Over-confidence and under-confidence The *confidence* of $\widehat{f}$ at $\mathbf{x}$ is the predicted top probability, i.e. $\max\{\widehat{f}(\mathbf{x}), 1 - \widehat{f}(\mathbf{x})\}$ for binary problems. In particular, when $\widehat{f}(\mathbf{x}) > 0.5$, the confidence is equal to $\widehat{f}(\mathbf{x})$. We say that the model is *over-confident* when the confidence is higher than the actual accuracy: For example, when the model predicts $\widehat{f}(\mathbf{x}) = 0.9$, but we have $\mathbb{E}[Y|\widehat{f}(\mathbf{x}) = 0.9] = 0.8$, then $\widehat{f}$ is over-confident at level $p = 0.9$. Note that in this case the calibration error at level 0.9 is positive: $\Delta_{0.9}^{\mathsf{cal}}(\widehat{f}) = 0.1 > 0$. In other words, over- or under-confidence is determined by the **sign** of the calibration error $\Delta_p^{\mathsf{cal}}(\widehat{f})$ in definition (2):

For any $p \in (0.5, 1)$:

- $\Delta_p^{\mathsf{cal}}(\widehat{f}) > 0$: $\widehat{f}$ is **over-confident** at level p;
- $\Delta_p^{\mathsf{cal}}(\widehat{f}) < 0$: $\widehat{f}$ is **under-confident** at level p.

We remark that we only state results for $p > 0.5$ in this paper; all the results also hold for $p \in (0, 0.5)$ by symmetry.

Extension to multi-class problems In our experiments we also consider multi-class classification problems, for which there is a standard generalization of definitions (2) and (1) (Guo et al., 2017): Given a multi-class predictor $\widehat{F} : \mathbb{R}^d \to \Delta_K$ where $K \geq 2$ is the number of classes, we replace Y with the indicator of correct prediction: $\mathbf{1}\Big\{Y = \arg\max_k \widehat{F}(x)_k\Big\}$, and replace $\widehat{f}(x)$ with the confidence $\max_k \widehat{F}(x)_k$. Thus the calibration error of $\widehat{F}$ at level $p \in [1/K, 1]$ is $\Delta_p^{\mathsf{cal}}(\widehat{F}) := p - \mathbb{P}\Big(Y = \arg\max_k \widehat{F}(\mathbf{X})_k \mid \max_k \widehat{F}(\mathbf{X})_k = p\Big)$.

2.2. Model and data distribution

We consider the following data distribution where $\mathbf{X}$ is standard Gaussian and $Y|\mathbf{X}$ follows a binary linear model with activation function $\sigma : \mathbb{R} \to [0,1]$:

$$P: \quad \mathbf{X} \sim \mathsf{N}(0, \mathbf{I}_d), \quad \mathbb{P}(Y = 1 \mid \mathbf{X} = \mathbf{x}) = \sigma(\mathbf{w}_\star^\top \mathbf{x}), \tag{3}$$

where $\mathbf{w}_\star \in \mathbb{R}^d$ is the ground truth coefficient vector. (This is also known as generalized linear models with link function σ (McCullagh, 2018)). We make the Gaussian input assumption as our analysis requires a precise limiting calculation; however, our real data experiments in Section 5.2 suggest that the implications of our theory may hold more broadly without such distributional assumptions.

Realizable logistic regression Our primary focus is *realizable logistic regression*, in which $\sigma(z) = \frac{1}{1+e^{-z}}$ is the logistic (sigmoid) activation, and we solve the unregularized

ERM (empirical risk minimization) problem

$$\begin{aligned}\widehat{\mathbf{w}} &= \arg\min_{\mathbf{w}} \widehat{R}_n(\mathbf{w}) \\ &:= \frac{1}{n}\sum_{i=1}^{n}\left[\log(1+\exp(\mathbf{w}^\top \mathbf{x}_i)) - y_i \mathbf{w}^\top \mathbf{x}_i\right].\end{aligned} \tag{4}$$

Let $R(\mathbf{w}) := \mathbb{E}[\widehat{R}_n(\mathbf{w})]$ denote the expected (population) risk. It is a classical result that $\arg\min_{\mathbf{w}} R(\mathbf{w}) = \mathbf{w}_\star$, i.e. logistic regression is well-specified when data comes from the logistic model (Hastie et al., 2009).

Extension to general activations We also consider generalizations where σ is a general monotone activation function, and we wish to learn a linear classifier $\widehat{\mathbf{w}}$ that is close to $\mathbf{w}_\star$. In this case, we consider solving the general ERM

$$\text{minimize } \widehat{R}_n(\mathbf{w}) := \frac{1}{n}\sum_{i=1}^{n}\rho(\mathbf{w}^\top \mathbf{x}_i) - y_i \mathbf{w}^\top \mathbf{x}_i, \tag{5}$$

where $\rho : \mathbb{R} \to \mathbb{R}$ is a loss function. Let $R(\mathbf{w}) := \mathbb{E}[\widehat{R}_n(\mathbf{w})]$ denote the expected (population) risk.

To make sure the problem is well-specified, we choose ρ to be the (integrated) convex loss associated with σ: $\rho(z) = \int_0^z \sigma(u)du + C$ for some constant C; in other words $\rho'(z) = \sigma(z)$. It is known that for such a choice of ρ we have $\arg\min_{\mathbf{w}} R(\mathbf{w}) = \mathbf{w}_\star$ (Kakade et al., 2011). (For completeness we also provide a proof in Appendix A.3.)

We require the following assumption on the activation function σ along with the loss function ρ, which only requires the activation to be smooth along with some basic properties, such as monotonicity and symmetry around 0.

Assumption A (Smooth activation)**.** *The loss function $\rho : \mathbb{R} \to \mathbb{R}$ is strictly convex and four-times continuously differentiable with uniformly bounded $\{1, 2, 3, 4\}$-th derivatives. The activation function $\sigma = \rho'$ is strictly increasing, and satisfies $\sigma(0) = 1/2$, $\lim_{z\to-\infty}\sigma(z) = 0$, $\lim_{z\to\infty}\sigma(z) = 1$, and $\sigma'(z) = \sigma'(-z) > 0$ for all $z \in \mathbb{R}$.*

3. Logistic regression is over-confident

As a warm-up, consider running unregularized (linear) logistic regression in the over-parametrized setting where $n < d$ and the data is separable. In this case, it is known that the (ERM) solution to the logistic regression (4) does not exist (Albert & Anderson, 1984; Candès et al., 2020); the gradient descent path will also diverge to infinity norm (Soudry et al., 2018). Using an approximate solution $\widehat{\mathbf{w}}$ with a high norm will cause the learned classifier $\sigma(\widehat{\mathbf{w}}^\top \mathbf{x})$ to be nearly a step function (outputs are close to either 0 or 1). Such classifiers are clearly over-confident whenever the true conditional distribution $Y|\mathbf{X}$ is not approximately deterministic.

We are now ready to present our main result, which states that even in the most vanilla setting (well-specified, under-parametrized), logistic regression is still over-confident.

Theorem 1 (Well-specified logistic regression is over-confident)**.** *Consider the classifier $\widehat{f}(\mathbf{x}) = \sigma(\widehat{\mathbf{w}}^\top \mathbf{x})$ obtained from logistic regression (4), where the data is generated from the logistic model (3). Then we have the following.*

- *In the limit of $n, d \to \infty$[1] and $d/n \to \kappa$, where $\kappa \in (0, \kappa_0]$ for some constant $\kappa_0 > 0$ (which only depends on $\|\mathbf{w}_\star\|$), for any $p \in (0.5, 1)$, almost surely, we have*

  $$\Delta_p^{\mathsf{cal}}(\widehat{f}) \to C_{p,\kappa} \quad \text{for some } C_{p,\kappa} > 0.$$

  In words, logistic regression gives inherently over-confident estimates of the actual probabilities.

- *We have, for small enough $\kappa > 0$,*

  $$C_{p,\kappa} = C_p \cdot \kappa + o(\kappa).$$

  In words, as the sample size $n/d = 1/\kappa$ becomes large, the over-confidence effect becomes weaker. The scaling of this over-confidence effect is roughly $C_p \cdot d/n$.

Over-confidence is inherent for logistic regression Theorem 1 considers the under-parametrized setting, as we allow $d/n = \kappa$ to be any *small* value, thus the sample size n can be arbitrarily higher than the dimension d. It thus suggests that over-confidence of logistic regression a rather fundamental property, and challenges the common belief that over-confidence mostly comes from over-parametrization. Furthermore, even though $\Delta_p^{\mathsf{cal}}(\widehat{f})$ becomes smaller as the sample size increases (κ becomes lower), Theorem 1 still asserts the sign of $\Delta_p^{\mathsf{cal}}(\widehat{f})$ being always positive in the proportional limit of $n, d \to \infty$, $d/n \to \kappa$. This result perhaps unveils another source of over-confidence in real-world machine learning models beyond linear logistic models.

Furthermore, Theorem 1 shows that logistic regression is over-confident at all $p \in (0.5, 1)$. This suggests that the over-confidence in every confidence bin, as an empirical observation in well-trained neural networks (Guo et al., 2017), holds for logistic regression as well and is not unique to large over-parametrized models.

Regularization; comparison with classical asymptotics We remark that our result only holds for *unregularized* logistic regression, while it is known that various regularization can improve calibration (Gal & Ghahramani, 2016; Thulasidasan et al., 2019). Indeed, in our model, applying regularization (e.g. an L2 regularizer) will in general reduce the calibration error, as long as the regularization reduces

[1] We assume $\|\mathbf{w}_\star\|$ is the same for all (n, d).

the norm of $\widehat{\mathbf{w}}$ and does not hurt its correlation with $\mathbf{w}_\star$ too much. However, we intentionally focus on unregularized logistic regression which resembles practical setups such as neural networks in the memorizing regime. We also note that, in general, the best regularization strength for the optimal accuracy and the optimal calibration may be different.

We also briefly remark that our setting of $d, n \to \infty$, $d/n \to \infty$ is different from classical asymptotic statistics (which considers fixed d and $n \to \infty$) (Van der Vaart, 2000). Classical asymptotics would imply $\sqrt{n}\Delta_p^{\mathsf{cal}}(\widehat{f}) \xrightarrow{d} \mathsf{N}(0, V^2)$ for some V^2, and thus $\Delta_p^{\mathsf{cal}}(\widehat{f})$ has about equal chance to be positive or negative; in contrast, we show that $\Delta_p^{\mathsf{cal}}(\widehat{f})$ has a positive bias in the proportional limit, a regime arguably more realistic than classical asymptotics.

CE of logistic regression Theorem 1 further implies a result on the calibration error (CE) of logistic regression.

Corollary 2 (Asymptotics of calibration error)**.** *In the same setting as Theorem 1, as $d, n \to \infty$, $d/n \to \kappa$, the* CE *of the logistic regression solution $\widehat{f}$ satisfies*

$$\mathrm{CE}(\widehat{f}) \to C_\kappa,$$

almost surely, where for small enough κ we have $C_\kappa = C\kappa + o(\kappa)$ for some absolute constant $C > 0$.

Corollary 2 implies that, in the limiting regime, the CE of logistic regression is $O(\kappa) = O(d/n)$. This improves over the results of Liu et al. (2019) in certain aspects. First, Liu et al. (2019, Corollary 2.4) showed that the CE of any classifier is bounded by the square root excess logistic loss over the Bayes classifier. This implies the CE of well-specified logistic regression is bounded by $\sqrt{d/n}$. Here we show the CE has a better rate $\Theta(d/n)$ at small d/n in our limiting regime[2]. Second, our Theorem 1 determines the sign of the calibration error (confidence $>$ accuracy), which is not implied by their results.

The proof of Corollary 2 follows directly from Theorem 1 by integrating $\Delta_p^{\mathsf{cal}}(\widehat{f})$ over $p \in (0, 1)$ (with p distributed as $\widehat{f}(\mathbf{x})$ for $\mathbf{x} \sim P$). The proof can be found in Appendix D.3.

3.1. Proof sketch of Theorem 1

We now provide a high-level overview of the proof of Theorem 1. A more detailed overview of the most technical steps is deferred to Section 6, and the full proofs can be found in Appendix C & D.

[2]We remark that Corollary 2 does not readily imply a $\Theta(d/n)$ result in the non-asymptotic setting. However, we believe a similar result (with additional terms such as $1/\sqrt{n}$) holds and can be established via a more refined analysis.

Closed-form expression for calibration error Recall that

$$\Delta_p^{\mathsf{cal}}(\widehat{f}) = p - \mathbb{E}_{\mathbf{x}}\big[\sigma(\mathbf{w}_\star^\top \mathbf{x}) \mid \sigma(\widehat{\mathbf{w}}^\top \mathbf{x}) = p\big].$$

As $\mathbf{x}$ is standard Gaussian, the conditional distribution of $\mathbf{x}|\sigma(\widehat{\mathbf{w}}^\top \mathbf{x}) = p$ can be characterized precisely in terms of the projection of $\mathbf{x}$ onto the direction $\widehat{\mathbf{w}}$ and its orthogonal complement subspace. Standard calculation then yields the closed form expression

$$\begin{aligned}&\Delta_p^{\mathsf{cal}}(\widehat{f}) \\ &= p - \mathbb{E}_Z\left[\sigma\left(\frac{\|\mathbf{w}_\star\|}{\|\widehat{\mathbf{w}}\|}\cos\widehat{\theta}\cdot\sigma^{-1}(p) + \sin\widehat{\theta}\,\|\mathbf{w}_\star\|\,Z\right)\right],\end{aligned} \tag{6}$$

where $\cos\widehat{\theta} = \frac{\widehat{\mathbf{w}}^\top \mathbf{w}_\star}{\|\widehat{\mathbf{w}}\|\|\mathbf{w}_\star\|}$ is the angle between $\widehat{\mathbf{w}}$ and $\mathbf{w}_\star$, and $Z \sim \mathsf{N}(0, 1)$. (See Lemma B.1 for the detailed statement and proof.)

Concentration of $\widehat{\mathbf{w}}$ In the second step, we apply results from recent advances in high-dimensional convex risk minimization (Sur & Candès, 2019; Taheri et al., 2020) to show that $\widehat{\mathbf{w}}$ concentrates around fixed values in the high-dimensional limit, in terms of its norm and cosine angle with $\mathbf{w}_\star$. These results show that, in the limit of $d, n \to \infty$ and $d/n \to \kappa$, the following concentration happens almost surely:

$$\begin{aligned}\|\widehat{\mathbf{w}}\| &\to R_\star = R_\star(\kappa, \|\mathbf{w}_\star\|),\\ \cos\widehat{\theta} &\to c_\star = c_\star(\kappa, \|\mathbf{w}_\star\|),\end{aligned} \tag{7}$$

Above, $R_\star$ and $c_\star$ are determined by the solutions of a system of nonlinear equations with three variables $(\alpha, \sigma, \lambda)$ (see Section 6 and Theorem C.1 for the formal statement).

The concentration directly implies that $\Delta_p^{\mathsf{cal}}(\widehat{f})$ converges to the following limiting calibration error (Corollary C.1):

$$\begin{aligned}&\Delta_p^{\mathsf{cal}}(\widehat{f}) \to C_{p,\kappa}\\ &:= p - \mathbb{E}_Z\left[\sigma\left(\frac{\|\mathbf{w}_\star\|}{R_\star}c_\star\cdot\sigma^{-1}(p) + \sqrt{1-c_\star^2}\,\|\mathbf{w}_\star\|\,Z\right)\right].\end{aligned} \tag{8}$$

This expression hints on potential sources of over- or under-confidence: (1) $R_\star$ and $c_\star$ will affect the "multiplier" $\|\mathbf{w}_\star\|\,c_\star/R_\star$ in front of $\sigma^{-1}(p)$, drifting the expectation away from p; (2) $c_\star$ also affects the expectation over the $\sqrt{1-c_\star^2}\,\|\mathbf{w}_\star\|\,Z$ term. This term itself has mean zero, but can affect the overall expectation through the nonlinear activation function σ.

Calculating the limiting calibration error The final part, as a technical crux of the proof, calculates the limiting calibration error (8) by precisely analyzing the interplay between the concentration values $R_\star$, $c_\star$ and the activation function σ. This is achieved by a novel analysis on the solutions of the aforementioned system of equations at small

κ. In particular, we show that $C_{p,\kappa} = C_p\kappa + o(\kappa)$ for small κ, and $C_p > 0$ is positive, thereby establishing Theorem 1. We present a more detailed description of this analysis in Section 6.

4. Over-confidence is not universal

It is natural to ask—based on Theorem 1—whether over-confidence is true in other well-specified problems as well, or is due to some specific property about logistic regression. This section makes steps towards this by looking at the generalized problem (5) where σ is an arbitrary activation and we solve the corresponding convex ERM.

Our main result in this section is the following characterization of sufficient conditions for whether over- or under-confidence happens in the general convex ERM (5). The proof of this result can be found in Appendix D.1.

Theorem 3 (Sufficient conditions for over- and under-confidence)**.** *In the same setting as Theorem 1 except that the activation function σ is general and satisfies Assumption A, let $\widehat{f}(\mathbf{x}) = \sigma(\widehat{\mathbf{w}}^\top \mathbf{x})$ be the classifier obtained from the convex ERM (5). We have simultaneously for any $p \in (0.5, 1)$ that, almost surely in the limit of $d, n \to \infty$, $d/n \to \kappa$,*

$$\Delta_p^{\mathsf{cal}}(\widehat{f}) \to C_{p,\kappa}(\sigma) = C_p(\sigma)\kappa + o(\kappa). \tag{9}$$

Further, we have the following sufficient conditions for the sign of $C_p(\sigma)$: For any $p \in (0.5, 1)$,

(a) If σ is concave at $\sigma^{-1}(p)$, i.e.,

$$\sigma''(\sigma^{-1}(p)) \le 0, \tag{10}$$

then $C_p(\sigma) > 0$, and $\widehat{f}$ is over-confident at this p for all sufficiently small κ.

(b) Conversely, if

$$\mathbb{E}_{Q_1 \sim \mathsf{N}(0, \|\mathbf{w}_\star\|^2)}[Q_1 \sigma''(Q_1)] > 0, \quad \text{and} \tag{11}$$

$$\sigma''(\sigma^{-1}(p)) - 2\sigma'(\sigma^{-1}(p)) \cdot \sigma^{-1}(p) / \|\mathbf{w}_\star\|^2 > 0, \tag{12}$$

then $C_p(\sigma) < 0$, and $\widehat{f}$ is under-confident at this p for all sufficiently small κ.

Interpretations Theorem 3 suggests that the *curvature* of the activation function σ is critical for determining its over- or under-confidence. We parse these sufficient conditions as follows:

- *Concavity of $\sigma(z)|_{z>0}$ implies over-confidence*: By part (a), at any p where $\sigma''(\sigma^{-1}(p)) \ge 0$, $\widehat{f}$ will be over-confident at that p. Moreover, any σ that is concave on the entire positive part ($\sigma''(z) \le 0$ for all $z > 0$) will result in over-confident at every $p > 0.5$. This strictly generalizes Theorem 1, and suggests that over-confidence is a common mode, as any σ that is monotone and bounded must have some concave regions on the positive part.
- *Under-confidence is possible but cannot hold at every p*: Part (b) suggests that under-confidence may be possible, provided we design σ that is sufficiently *convex* at $\sigma^{-1}(p)$ (to counteract the other term in (12)), and that additional condition (11) holds. However, as $\sigma''(z) > 0$ cannot happen for all $z > 0$, under-confidence cannot happen at every $p \in (0.5, 1)$.

Is the sufficient condition for under-confidence in Theorem 3(b) indeed possible? We give an affirmative answer.

Corollary 4 (Under-confidence can happen)**.** *There exists an activation function σ satisfying Assumption A, such that $C_p(\sigma) < 0$ for some $p \in (0.5, 1)$ and $\|\mathbf{w}_\star\| > 0$. At these p, the convex ERM (5) is under-confident in the limit of $d/n \to \kappa$ for all small κ.*

The activation we find in Corollary 4 is very close to the following activation function (up to minor tweaks in order to satisfy Assumption A):

$$\sigma_{\text{underconf}}(z) = \begin{cases} 0, & z < -2\pi, \\ \dfrac{1}{2} + \dfrac{1}{4\pi}(z - \sin z), & |z| \le 2\pi, \\ 1, & z > 2\pi. \end{cases} \tag{13}$$

(See Figure 2a for the plot of this activation function.) The unique feature about this $\sigma_{\text{underconf}}$ is that, unlike the logistic activation, this function is convex at all small values of $z > 0$. This leads to both the convexity condition (12) as well as the expectation condition (11) (which roughly requires the positive part in the expectation of $Q_1\sigma''(Q_1)$ to supercede the negative part).

To the best of our knowledge, this is the first known case of under-confidence for a well-specified classification problem, though we remark this under-confidence effect is weak and restricted to only a small region of p (see Figure 2c for simulation results using this activation).

5. Experiments

5.1. Simulations

We test our theories via simulations on well-specified under-parametrized logistic regression, as well as general convex ERM with the under-confident activation $\sigma_{\text{underconf}}$ (13).

For both activations, we generate data $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$ from the realizable model (3), where we fix $d = 100$, $\|\mathbf{w}_\star\| = 1$, and vary $d/n \in \{0.01, 0.05, 0.10, 0.25\}$. For each (d, n),

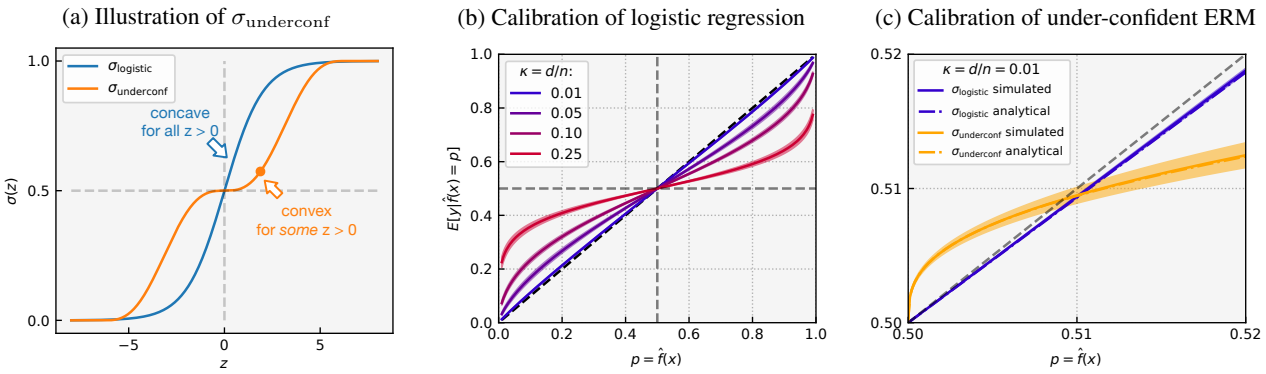

Figure 2. Binary classification simulations on realizable data. **(a)** Illustration of the activation function $\sigma_{\text{underconf}}$ constructed in Corollary 4 (cf. (13)), against the logistic (sigmoid) activation σ_{logistic}. **(b)** Calibration curves for simulated logistic regression, with $d = 100$ and $d/n \in \{0.01, 0.05, 0.1, 0.25\}$. **Logistic regression is over-confident** (prediction is higher than actual probability when prediction> 0.5) at all d/n. **(c)** Zoomed-in calibration curves for simulated realizable ERM with the $\sigma_{\text{underconf}}$ activation. In contrast to logistic regression, $\sigma_{\text{underconf}}$ **leads to under-confidence** for $p \in (0.5, 0.51)$, verifying our Theorem 3 and Corollary 4. Here "analytical" refers to our theoretical prediction $p - C_{p,\kappa}(\sigma)$ from Theorem 3. (b)(c): Shaded area are one-std error bars over 5 runs.

we generate 5 problem instances, solve the ERM problem on each instance, and plot the "calibration curves" (where the x-axis is p and y-axis is the average probability given the prediction: $\mathbb{E}[y|\widehat{f}^{(i)}(\mathbf{x}) = p] = p - \Delta_p^{\text{cal}}(\widehat{f}^{(i)})$), visualizing their mean and one-standard-deviation error bar. Notice that by the closed-form expression (6), we are able to compute $\Delta_p^{\text{cal}}(\widehat{f})$ exactly (using Gaussian integration) without needing to introduce a test set.

In addition to the simulated calibration curves, we also plot the limiting calibration curve suggested by Theorem 1 & 3, in which we compute the concentration values $R_\star, c_\star$ analytically by solving its defining equations (cf. Appendix C.1), and plug these values into the closed-form expression (8). This yields a curve of p against $p - C_{p,\kappa}(\sigma)$, which we compare against our simulated curves.

Results Figure 2 shows the results of our simulations. We find logistic regression indeed yields over-confident calibration curves (Figure 2b): $\mathbb{E}[y|\widehat{f}(\mathbf{x}) = p] = p - \Delta_p^{\text{cal}}(\widehat{f})$ is less than p for $p > 0.5$ (and greater than p for $p < 0.5$). Further, notice that the gap Δ_p^{cal} increases as we increase κ. This agrees with our intuition that over-confidence is more severe when d/n increases (effective sample size gets lower), and further suggests that the conclusion of our theory holds more broadly than its assumptions: κ can be as large as $\kappa = 0.25$ and d can be as low as 100, both being realistic values for modeling practice.

We also find that the under-confidence shown in Corollary 4 does show up in the simulations: With the activation $\sigma_{\text{underconf}}$, $\mathbb{E}[y|\widehat{f}(\mathbf{x}) = p]$ is *higher* than p for $p \in (0.5, 0.51)$, although this range of p is fairly narrow (Figure 2c).

Finally, we observe that our theoretical prediction $C_{p,\kappa}$ closely matches the simulation: the analytical calibration curve $p - C_{p,\kappa}(\sigma)$ and the mean simulated curve are almost identical for both activations, which further confirms our theory even at a realistic $d = 100$.

5.2. CIFAR10 with pseudo labels

We further test the generality of our theory beyond the Gaussian input assumption and the binary classification setting. We run multi-class logistic regression on the first 5 classes of CIFAR10, which contains $n = 25000$ training images and 5000 test images, and each image has $d = 3072$ features. We perform logistic regression on two kinds of labels:

- The true label $y^{\text{true}} \in \{0, 1, 2, 3, 4\}$.
- The pseudo-label $y^{\text{pseudo}} \in \{0, 1, 2, 3, 4\}$ generated as follows: After fitting the logistic classifier $\widehat{\mathbf{W}} \in \mathbb{R}^{3072\times 5}$ on the true labels, we generate pseudo-labels y_i^{pseudo} from the multi-class logistic (softmax) model

$$\mathbb{P}\Big(y_i^{\text{pseudo}} = k \mid \mathbf{x}_i\Big) = \frac{\exp(\widehat{\mathbf{W}}_k^\top \mathbf{x}_i)}{\sum_{k'} \exp(\widehat{\mathbf{W}}_{k'}^\top \mathbf{x}_i)}.$$

The motivation for the pseudo-labels is to construct a well-specified problem (labels do come from a linear softmax model) and remove the potential effect of model-misspecification with the true labels. Note that this problem is still in the under-parametrized setting as $d < n$.

As the exact conditioning $\widehat{f}(\mathbf{x}) = p$ is no longer computable on finite data, we compute the average confidence and accuracy on the test set using binning (10 equally spaced confidence bins in $[0.2, 1.0]$), similar as in the standard practice for evaluating the ECE (Guo et al., 2017). Additional experimental details are provided in Appendix E.2.

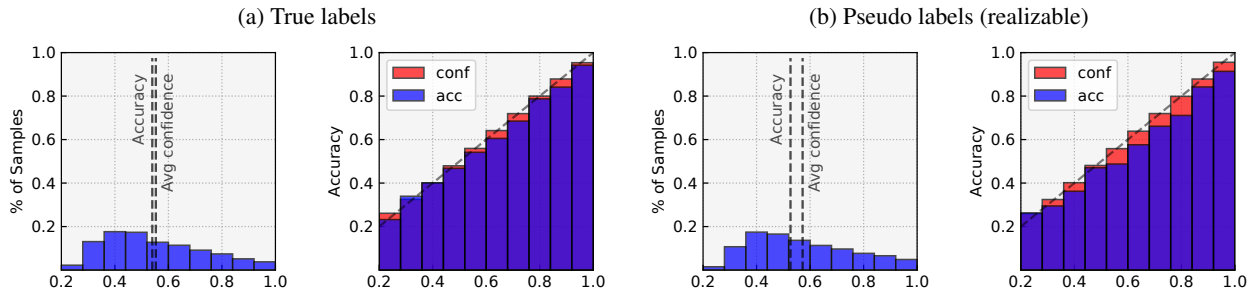

Figure 3. Calibration of multi-class logistic regression on CIFAR10's first 5-classes. The x-axes denote the confidences (predicted top probabilities) of the models. **(a)(b)**: Left: Confidence distribution across bins; Right: Average confidence against average accuracy within each bin (right); both evaluated on the test set. **(a)** Logistic regression on the true labels. **(b)** Logistic regression on pseudo-labels generated from the fitted logistic model (realizable setting) from step (a). Observe that **over-confidence happens for both the pseudo-labels generated from a multi-class logistic model, and the true labels.**

Results We find that logistic regression on (the 5-class subset of) CIFAR10 is over-confident on both the pseudo-labels and true labels (Figure 3). A closer look reveals that the over-confidence is more severe on the pseudo-labels than the true labels, yet both tasks exhibit a reasonable level of over-confidence (especially in the high confidence bins). This suggests our result that logistic regression is inherently over-confident may hold more broadly for other under-parametrized problems without strong assumptions on the input distribution, or even when the labels are not necessarily realizable by the model.

6. Overview of analysis

This section provides an overview of the two novel proof steps for our results in Section 3 & 4: (1) Characterization of the high-dimensional limit (concentration value) of logistic regression (4) and the general convex ERM (5) at small $\kappa = d/n$. (2) Determining the sign of the limiting calibration error based on the above characterization, filling in the (abbreviated) last part of the proof sketch in Section 3.1.

6.1. Local linear analysis at small κ

Let $\gamma := \|\mathbf{w}_\star\|$. By the results of Sur & Candès (2019), the values $R_\star, c_\star$ in (7) have the form $R_\star = \sqrt{\alpha_\star^2 + \kappa\sigma_\star^2}$ and $c_\star = (1 + \kappa\sigma_\star^2/\alpha_\star^2\gamma^2)^{-1/2}$, where $(\alpha_\star, \sigma_\star, \lambda_\star)$ are the solutions to the following system of nonlinear equations in three variables $(\alpha, \sigma, \lambda)$:

$$\begin{cases} \sigma^2 = \dfrac{1}{\kappa^2}\mathbb{E}\big[2\rho'(Q_1)\lambda^2\rho'(\mathsf{prox}_{\lambda\rho}(Q_2))^2\big], \\ 0 = \mathbb{E}\big[\rho'(Q_1)Q_1\lambda\rho'(\mathsf{prox}_{\lambda\rho}(Q_2))\big], \\ 1 - \kappa = \mathbb{E}\big[2\rho'(Q_1)/(1 + \lambda\rho''(\mathsf{prox}_{\lambda\rho}(Q_2)))\big]. \end{cases}$$

Above, (Q_1, Q_2) has a bivariate normal distribution with covariance depending on $(\alpha, \sigma, \kappa, \gamma)$, and prox is the prox operator. (See Theorem C.1 and Appendix C.1 for a formal statement.) These solutions are guaranteed to uniquely exist for small enough κ. However, they are only implicitly defined without closed-form expressions for these solutions, which prohibits us from analyzing their behaviors.

We overcome this issue by performing a local analysis of the solutions at small κ. We prove that, for small enough κ, we have the local linear approximation

$$\begin{aligned} \alpha_\star &= \alpha_\star(\kappa) = 1 + \bar{\alpha}_0\kappa + O(\kappa^2), \\ \sigma_\star &= \sigma_\star^2(\kappa) = \bar{\sigma}_0^2 + O(\kappa), \\ \lambda_\star &= \lambda_\star(\kappa) = \bar{\lambda}_0\kappa + O(\kappa^2), \end{aligned}$$

with closed-form expressions for $(\bar{\alpha}_0, \bar{\sigma}_0, \bar{\lambda}_0)$. For example, we have $\bar{\sigma}_0^2 = \mathbb{E}[\rho'(Q_1)\rho'(-Q_1)]/(\mathbb{E}[\rho''(Q_1)])^2$ where $Q_1 \sim \mathsf{N}(0, \gamma^2)$. (See Lemma C.1 for the formal statement.) These approximations imply similar approximations for $R_\star, c_\star$, which allows us to analyze the behavior of the limiting calibration error (8) locally at small κ.

6.2. Determining sign of the limiting calibration error

Towards proving Theorem 3 & 1, it remains for us to derive the sufficient conditions for the sign of $C_{p,\kappa}(\sigma)$. Using the above local linear approximation for $(R_\star, c_\star)$ and performing first-order calculus, we obtain

$$\begin{aligned} \lim_{\kappa\to 0}\frac{C_{p,\kappa}(\sigma)}{\kappa} &= \sigma'(\sigma^{-1}(p))\cdot\sigma^{-1}(p)\cdot\big(\bar{\alpha}_0 + \bar{\sigma}_0^2/\gamma^2\big) \\ &\quad - \frac{1}{2}\sigma''(\sigma^{-1}(p))\cdot\bar{\sigma}_0^2. \end{aligned}$$

(Lemma D.1). We prove that $\bar{\alpha}_0 + \bar{\sigma}_0^2/\gamma^2 > 0$ always holds regardless of the activation,γ, and p. This implies that, as long as $\sigma''(\sigma^{-1}(p)) \le 0$, the right-hand side in the equation above is positive. This gives part (a) of Theorem 3. On the other hand, if $\sigma''(\sigma^{-1}(p)) > 2\sigma'(\sigma^{-1}(p))\sigma^{-1}(p)/\gamma^2$ and $\bar{\alpha}_0 < 0$, then the right-hand side above is negative. These are exactly the sufficient conditions we required in part (b) of Theorem 3.

7. Conclusion

This paper provides a precise theoretical study of the calibration error of logistic regression and a class of general binary classification problems. We show that logistic regression is inherently over-confident by $\Theta(d/n)$ when n/d is large, and establish sufficient conditions for the over- or under-confidence of unregularized ERM for general binary classification. Our results reveal that (1) Over-confidence is not just a result of over-parametrization; (2) Over-confidence is a common mode but not universal. We believe our work opens up a number of future questions, such as the interplay between calibration and model training (or regularization), or theoretical studies of calibration on nonlinear models.

Acknowledgment

We thank Yuxin Chen for the insightful discussions on the system of nonlinear equations. We thank all anonymous reviewers for the many helpful feedback on this paper.

References

Albert, A. and Anderson, J. A. On the existence of maximum likelihood estimates in logistic regression models. *Biometrika*, 71(1):1–10, 1984.

Bayati, M. and Montanari, A. The dynamics of message passing on dense graphs, with applications to compressed sensing. *IEEE Transactions on Information Theory*, 57 (2):764–785, 2011.

Begoli, E., Bhattacharya, T., and Kusnezov, D. The need for uncertainty quantification in machine-assisted medical decision making. *Nature Machine Intelligence*, 1(1):20–23, 2019.

Candès, E. J., Sur, P., et al. The phase transition for the existence of the maximum likelihood estimate in high-dimensional logistic regression. *The Annals of Statistics*, 48(1):27–42, 2020.

Ding, Z., Han, X., Liu, P., and Niethammer, M. Local temperature scaling for probability calibration. *arXiv preprint arXiv:2008.05105*, 2020.

Donoho, D. L., Maleki, A., and Montanari, A. Message-passing algorithms for compressed sensing. *Proceedings of the National Academy of Sciences*, 106(45):18914–18919, 2009.

Dusenberry, M., Jerfel, G., Wen, Y., Ma, Y., Snoek, J., Heller, K., Lakshminarayanan, B., and Tran, D. Efficient and scalable bayesian neural nets with rank-1 factors. In *International conference on machine learning*, pp. 2782–2792. PMLR, 2020.

El Karoui, N., Bean, D., Bickel, P. J., Lim, C., and Yu, B. On robust regression with high-dimensional predictors. *Proceedings of the National Academy of Sciences*, 110 (36):14557–14562, 2013.

Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.

Gal, Y., Hron, J., and Kendall, A. Concrete dropout. *arXiv preprint arXiv:1705.07832*, 2017.

Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. *arXiv preprint arXiv:1706.04599*, 2017.

Gupta, C., Podkopaev, A., and Ramdas, A. Distribution-free binary classification: prediction sets, confidence intervals and calibration. *arXiv preprint arXiv:2006.10564*, 2020.

Hastie, T., Tibshirani, R., and Friedman, J. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.

Jung, C., Lee, C., Pai, M. M., Roth, A., and Vohra, R. Moment multicalibration for uncertainty estimation. *arXiv preprint arXiv:2008.08037*, 2020.

Kakade, S., Kalai, A. T., Kanade, V., and Shamir, O. Efficient learning of generalized linear and single index models with isotonic regression. *arXiv preprint arXiv:1104.2018*, 2011.

Karoui, N. E. Asymptotic behavior of unregularized and ridge-regularized high-dimensional robust regression estimators: rigorous results. *arXiv preprint arXiv:1311.2445*, 2013.

Kull, M., Silva Filho, T., and Flach, P. Beta calibration: a well-founded and easily implemented improvement on logistic calibration for binary classifiers. In *Artificial Intelligence and Statistics*, pp. 623–631. PMLR, 2017.

Kull, M., Perello-Nieto, M., Kängsepp, M., Song, H., Flach, P., et al. Beyond temperature scaling: Obtaining well-calibrated multiclass probabilities with dirichlet calibration. *arXiv preprint arXiv:1910.12656*, 2019.

Kumar, A., Liang, P., and Ma, T. Verified uncertainty calibration. *arXiv preprint arXiv:1909.10155*, 2019.

Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*, 2016.

Liu, J. Z., Lin, Z., Padhy, S., Tran, D., Bedrax-Weiss, T., and Lakshminarayanan, B. Simple and principled uncertainty estimation with deterministic deep learning via distance awareness. *arXiv preprint arXiv:2006.10108*, 2020.

Liu, L. T., Simchowitz, M., and Hardt, M. The implicit fairness criterion of unconstrained learning. In *International Conference on Machine Learning*, pp. 4051–4060. PMLR, 2019.

Maddox, W. J., Izmailov, P., Garipov, T., Vetrov, D. P., and Wilson, A. G. A simple baseline for bayesian uncertainty in deep learning. *Advances in Neural Information Processing Systems*, 32:13153–13164, 2019.

Mai, X., Liao, Z., and Couillet, R. A large scale analysis of logistic regression: Asymptotic performance and new insights. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3357–3361. IEEE, 2019.

Malinin, A., Mlodozeniec, B., and Gales, M. Ensemble distribution distillation. *arXiv preprint arXiv:1905.00076*, 2019.

McCullagh, P. *Generalized linear models*. Routledge, 2018.

Michelmore, R., Kwiatkowska, M., and Gal, Y. Evaluating uncertainty quantification in end-to-end autonomous driving control. *arXiv preprint arXiv:1811.06817*, 2018.

Mukhoti, J., Kulharia, V., Sanyal, A., Golodetz, S., Torr, P. H., and Dokania, P. K. Calibrating deep neural networks using focal loss. *arXiv preprint arXiv:2002.09437*, 2020.

Naeini, M. P., Cooper, G., and Hauskrecht, M. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.

Nixon, J., Dusenberry, M. W., Zhang, L., Jerfel, G., and Tran, D. Measuring calibration in deep learning. In *CVPR Workshops*, volume 2, 2019.

Ovadia, Y., Fertig, E., Ren, J., Nado, Z., Sculley, D., Nowozin, S., Dillon, J., Lakshminarayanan, B., and Snoek, J. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems*, pp. 13991–14002, 2019.

Platt, J. et al. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. 1999.

Rahimi, A., Shaban, A., Cheng, C.-A., Hartley, R., and Boots, B. Intra order-preserving functions for calibration of multi-class neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.

Shabat, E., Cohen, L., and Mansour, Y. Sample complexity of uniform convergence for multicalibration. *arXiv preprint arXiv:2005.01757*, 2020.

Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.

Stojnic, M. A framework to characterize performance of lasso algorithms. *arXiv preprint arXiv:1303.7291*, 2013.

Sur, P. and Candès, E. J. A modern maximum-likelihood theory for high-dimensional logistic regression. *Proceedings of the National Academy of Sciences*, 116(29):14516–14525, 2019.

Taheri, H., Pedarsani, R., and Thrampoulidis, C. Fundamental limits of ridge-regularized empirical risk minimization in high dimensions. *arXiv preprint arXiv:2006.08917*, 2020.

Thrampoulidis, C., Oymak, S., and Hassibi, B. Regularized linear regression: A precise analysis of the estimation error. In *Conference on Learning Theory*, pp. 1683–1709. PMLR, 2015.

Thrampoulidis, C., Abbasi, E., and Hassibi, B. Precise error analysis of regularized m-estimators in high dimensions. *IEEE Transactions on Information Theory*, 64(8):5592–5628, 2018.

Thulasidasan, S., Chennupati, G., Bilmes, J., Bhattacharya, T., and Michalak, S. On mixup training: Improved calibration and predictive uncertainty for deep neural networks. *arXiv preprint arXiv:1905.11001*, 2019.

Tran, L., Veeling, B. S., Roth, K., Swiatkowski, J., Dillon, J. V., Snoek, J., Mandt, S., Salimans, T., Nowozin, S., and Jenatton, R. Hydra: Preserving ensemble diversity for model distillation. *arXiv preprint arXiv:2001.04694*, 2020.

Van der Vaart, A. W. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Wen, Y., Tran, D., and Ba, J. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=Sklf1yrYDr.

Zadrozny, B. and Elkan, C. Obtaining calibrated probability estimates from decision trees and naive bayesian classifiers. Citeseer.

Zadrozny, B. and Elkan, C. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 694–699, 2002.

Zagoruyko, S. and Komodakis, N. Wide residual networks. *arXiv preprint arXiv:1605.07146*, 2016.

Zhang, J., Kailkhura, B., and Han, T. Y.-J. Mix-n-match: Ensemble and compositional methods for uncertainty calibration in deep learning. In *International Conference on Machine Learning*, pp. 11117–11128. PMLR, 2020.