
Certifying Out-of-Domain Generalization for Blackbox Functions

Maurice Weber¹ Linyi Li² Boxin Wang² Zhikuan Zhao¹ Bo Li² Ce Zhang¹

Abstract

Certifying the robustness of model performance under bounded data distribution drifts has recently attracted intensive interest under the umbrella of *distributional robustness*. However, existing techniques either make strong assumptions on the model class and loss functions that can be certified, such as smoothness expressed via Lipschitz continuity of gradients, or require to solve complex optimization problems. As a result, the wider application of these techniques is currently limited by its *scalability* and *flexibility* — these techniques often do not scale to large-scale datasets with modern deep neural networks or cannot handle loss functions which may be non-smooth such as the 0-1 loss. In this paper, we focus on the problem of certifying distributional robustness for *blackbox models and bounded loss functions*, and propose a novel certification framework based on the Hellinger distance. Our certification technique scales to ImageNet-scale datasets, complex models, and a diverse set of loss functions. We then focus on one specific application enabled by such scalability and flexibility, i.e., certifying out-of-domain generalization for large neural networks and loss functions such as accuracy and AUC. We experimentally validate our certification method on a number of datasets, ranging from ImageNet, where we provide the first non-vacuous certified out-of-domain generalization, to smaller classification tasks where we are able to compare with the state-of-the-art and show that our method performs considerably better.

1. Introduction

The wide application of machine learning models in the real world brings an emerging challenge of understanding the

¹Department of Computer Science, ETH Zurich
²UIUC, USA. Correspondence to: Maurice Weber <maurice.weber@inf.ethz.ch>, Ce Zhang <ce.zhang@inf.ethz.ch>.

performance of a machine learning model under different data distributions — ML systems operating autonomous vehicles which are trained based on data collected in the northern hemisphere might fail when deployed in desert-like environments or under different weather conditions (Volk et al., 2019; Dai & Van Gool, 2018), while recognition systems have been shown to fail when deployed in new environments (Beery et al., 2018). Similar concerns also apply to many mission-critical applications such as medicine and cyber-security (Koh et al., 2021; AlBadawy et al., 2018; Gulrajani & Lopez-Paz, 2021). In all these applications, it is imperative to have a sound understanding of the model’s robustness and possible failure cases in the presence of a shift in the data distribution, and to have corresponding guarantees on the performance.

Recently, this problem has attracted intensive interest under the umbrella of *distributional robustness* (Scarf, 1958; Ben-Tal et al., 2013; Gao & Kleywegt, 2016; Kuhn et al., 2019; Blanchet & Murthy, 2019; Duchi et al., 2021). Specifically, let P be a joint data distribution over features $X \in \mathcal{X}$ and labels $Y \in \mathcal{Y}$, and let $h_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ be a machine learning model parameterized by θ . For a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$, we hope to compute

$$\mathcal{R}_\theta(\mathcal{U}_P) := \sup_{Q \in \mathcal{U}_P} \mathbb{E}_{(X,Y) \sim Q}[\ell(h_\theta(X), Y)] \quad (1)$$

where $\mathcal{U}_P \subseteq \mathcal{P}(\mathcal{Z})$ is a set of probability distributions on \mathcal{Z} , called the *uncertainty set*. Intuitively, this measures the *worst-case risk* of h_θ when the data distribution drifts from P to another distribution in \mathcal{U}_P .

Providing a technical solution to this problem has gained increased attention over the years, as summarized in Table 1. However, most — if not all — existing approaches, impose strong constraints such as bounded Lipschitz gradients on both h and ℓ and rely on expensive certification methods such as direct minimax optimization. As a result, these methods have been applied only to small-scale datasets and ML models.

In this paper, we consider the case that both h and ℓ can be *non-convex and non-smooth* — h can be a full-fledged neural network, e.g., ImageNet-scale EfficientNet-B7 (Tan & Le, 2019), and ℓ can be a general non-smooth loss function such as the 0-1 loss. We provide, to our best knowledge, the first practical method for blackbox functions that scales to

Ref.	Assumptions on ℓ	Assumption on h	Distance	Largest Dataset
(Gao & Kleywegt, 2016)	Generalised Lipschitz Continuity		Wasserstein	–
(Sinha et al., 2018)	Bounded, Smoothness	Smoothness	Wasserstein	MNIST
(Staib & Jegelka, 2019)	Bounded, Continuous	Kernel Methods	MMD	–
(Shafieezadeh-Abadeh et al., 2019)		Lipschitz Continuity	Wasserstein	–
(Blanchet & Murthy, 2019)	Bounded, Smoothness	Smoothness	Wasserstein	–
(Cranko et al., 2021)	Generalised Lipschitz Continuity		Wasserstein	–
Our Method	Bounded	any Blackbox	Hellinger	ImageNet

Table 1. Current landscape of certified distributional robustness.

real-world, ImageNet-scale neural networks and datasets. Our key innovation is a novel algorithmic framework that arises from bounding inner products between elements of a suitable Hilbert space. Specifically, we can characterize the upper bound of the performance of h on any Q within the uncertainty set as a function of the Hellinger distance, a specific type of f -divergence, and the expectation and variance of the loss of h on P .

We then apply our framework to the problem of certifying the out-of-domain generalization performance of a given classifier, taking advantage of its scalability and flexibility. Specifically, let P be the *in-domain distribution*, and h_θ a classifier. Then, to reason about the performance of h_θ on shifted distributions Q , we provide a certificate in the following form:

$$\forall Q: \text{dist}(Q, P) \leq \rho \implies \mathbb{E}_{(X, Y) \sim Q}[\ell(h(X), Y)] \leq C_\ell(\rho, P) \quad (2)$$

where C_ℓ is a bound which depends on the distance ρ and the distribution P . This requires several nontrivial instantiations of our framework with careful practical considerations. To this end, we first develop a certification algorithm that relies only on a finite set of samples from the in-domain distribution P . Moreover, we also instantiate it with different domain drifting models such as label drifting and covariate drifting, connecting the general Hellinger distance to the degree of domain drifting specific to these scenarios. We then consider a diverse range of loss functions, including JSD loss, 0-1 loss, and AUC. To the best of our knowledge, we provide the first certificate for such diverse realistic scenarios, which is able to scale to large problems.

Last but not least, we conduct intensive experiments verifying the efficiency and effectiveness of our result. Our method is able to scale to datasets and neural networks as large as ImageNet and full-fledged models like EfficientNet-B7 and BERT. We further apply our method on smaller-scale datasets, in order to compare with strong, state-of-the-art methods. We show that our method provides much tighter certificates.

Our contributions can be summarized as follows:

- We present a novel framework which provides a non-

vacuous, computationally tractable bound to the distributionally robust worst-case risk $\mathcal{R}_\theta(\mathcal{U}_P)$ for general bounded loss functions ℓ and models h .

- We apply this framework to the problem of certifying out-of-domain generalization for blackbox functions and provide a means to certify distributional robustness in specific scenarios such as label and covariate drifts.
- We provide an extensive experimental study of our approach on a wide range of datasets including the large scale ImageNet (Russakovsky et al., 2015) dataset, as well as NLP datasets with complex models.

2. Distributional Robustness for Blackbox Functions

In this section, we present our main results, namely, a computationally tractable upper bound to the worst-case risk (1) for uncertainty sets expressed in terms of Hellinger balls around the data-generating distribution P . The technique is based on the non-negativity of Gram matrices which, by expressing expectation values as inner products between elements of a suitable Hilbert space, can be leveraged to relate expectation values of a blackbox function under different probability distributions P and Q .¹ We describe the underlying technique leading to our main result in Theorem 2.2, which upper bounds the worst-case population loss using both the expectation and variance.

For the remainder of this section, to simplify notation and maintain generality, we consider generic loss functions $\ell: \mathcal{Z} \rightarrow \mathbb{R}_+$ which contain the model h and take inputs from a generic input space \mathcal{Z} . For example in the context of supervised learning, $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ can be the product space of features and labels and the loss $\ell(z) = \tilde{\ell}(h_\theta(x), y)$ can be seen as a composition of the loss function $\tilde{\ell}$ and the model h_θ . We denote the set of probability measures on the space \mathcal{Z} by $\mathcal{P}(\mathcal{Z})$. For two measures μ, ν on \mathcal{Z} , we say that ν is absolutely continuous with respect to μ , denoted $\nu \ll \mu$ if $\mu(A) = 0$ implies that $\nu(A) = 0$ for any

¹The idea behind our methods is inspired by how Gram matrices are used in quantum chemistry (Weinhold, 1968; Weber et al., 2021) to bound expectation values of quantum observables. However, the adaptation to machine learning is nontrivial and requires careful analysis.

measurable set $A \subseteq \mathcal{Z}$. Among the plethora of distances between probability measures, such as total variation and Wasserstein distance, a particularly popular choice is the family of f -divergences which has been extensively studied in the context of distributionally robust optimization (Ben-Tal et al., 2013; Lam, 2016; Duchi & Namkoong, 2019; Duchi et al., 2021). In this paper, we focus on the Hellinger distance, which is a particular type of f divergence.

Definition 2.1 (Hellinger-distance). Let $P, Q \in \mathcal{P}(\mathcal{Z})$ be probability measures on \mathcal{Z} that are absolutely continuous with respect to a reference measure μ with $P, Q \ll \mu$. The Hellinger distance between P and Q is defined as

$$H(P, Q) := \sqrt{\frac{1}{2} \int_{\mathcal{Z}} \left(\sqrt{p(z)} - \sqrt{q(z)} \right)^2 d\mu(z)} \quad (3)$$

where $p = \frac{dP}{d\mu}$ and $q = \frac{dQ}{d\mu}$ are the Radon-Nikodym derivatives of P and Q with respect to μ . The Hellinger distance is independent of the choice of the reference measure μ .

The Hellinger distance is bounded with values in $[0, 1]$, with $H(P, Q) = 0$ if and only if $P = Q$ and the maximum value of 1 attained when P and Q have disjoint support. Furthermore, H defines a metric on the space of probability measures and hence satisfies the triangle inequality. We will now show how the Hellinger distance can be expressed in terms of an inner product between elements of a suitable Hilbert space, which ultimately enables us to use the theory of Gram matrices to derive an upper bound on the worst-case population risk (1) for uncertainty sets given by Hellinger balls. Consider the Hilbert space $L_2(\mathcal{Z}, \Sigma, \mu)^2$ of square-integrable functions $f: \mathcal{Z} \rightarrow \mathbb{R}$, endowed with the inner product $\langle f, g \rangle = \int_{\mathcal{Z}} fg d\mu$. Within this space, we can identify any probability distribution $P \ll \mu$ with a unit vector $\psi_P \in L_2(\mathcal{Z}, \Sigma, \mu)$ via the square root of its Radon-Nikodym derivative $\psi_P := \sqrt{dP/d\mu}$. This mapping enables us to write the Hellinger distance and, more generally, expectation values in terms of inner products. To see this, note that for any two probability measures P, Q on \mathcal{Z} , it holds that

$$\langle \psi_P, \psi_Q \rangle = \int_{\mathcal{Z}} \sqrt{dP} \sqrt{dQ} = 1 - H^2(P, Q) \quad (4)$$

and similarly, for any essentially bounded function $f \in L_\infty$, we have

$$\mathbb{E}_P[f(Z)] = \int_{\mathcal{Z}} f(z) dP(z) = \langle \psi_P, f \cdot \psi_P \rangle \quad (5)$$

where the product $(f \cdot \psi_P)(z) = f(z) \cdot \psi_P(z)$ is to be understood as pointwise multiplication³. For $f \in L_\infty$,

²We take Σ to be the Borel σ -algebra on \mathcal{Z} , being the smallest σ -algebra containing all open sets on \mathcal{Z} .

³More precisely, every $f \in L_\infty(\mathcal{Z}, \Sigma, \mu)$ defines a bounded

linear operator $M_f: L_2 \rightarrow L_2$ acting on elements of L_2 via pointwise multiplication, $g \mapsto M_f(g) := f \cdot g$ with $(f \cdot g)(z) = f(z) \cdot g(z)$ for any $z \in \mathcal{Z}$.

$$G := \begin{pmatrix} 1 & \langle \psi_Q, \psi_P \rangle & \langle \psi_Q, f\psi_P \rangle \\ \langle \psi_Q, \psi_P \rangle & 1 & \langle \psi_P, f\psi_P \rangle \\ \langle f\psi_P, \psi_Q \rangle & \langle f\psi_P, \psi_P \rangle & \langle f\psi_P, f\psi_P \rangle \end{pmatrix}. \quad (6)$$

The crucial observation is that G is positive semidefinite and thus has a non-negative determinant which can be viewed as a second degree polynomial $\pi(x)$ evaluated at $x = \langle \psi_Q, f\psi_P \rangle$ and is given by

$$\det(G) =: \pi(x)|_{x=\langle \psi_Q, f\psi_P \rangle} \quad (7)$$

where $\pi(x) = ax^2 + bx + c$ is a polynomial with coefficients

$$\begin{aligned} a &= -1, & b &= 2 \cdot \langle \psi_P, \psi_Q \rangle \cdot \langle \psi_P, f\psi_P \rangle \\ c &= (1 - |\langle \psi_P, \psi_Q \rangle|^2) \langle f\psi_P, f\psi_P \rangle - \langle f\psi_P, \psi_P \rangle^2. \end{aligned} \quad (8)$$

The non-negativity of $\det(G)$ implies that $\pi(x = \langle \psi_Q, f\psi_P \rangle) \geq 0$ and thus effectively restricts the values which $\langle \psi_Q, f\psi_P \rangle$ can take to be bounded within the square roots of π so that

$$\frac{b}{2} - \sqrt{\frac{b^2}{4} + c} \leq \langle \psi_Q, f\psi_P \rangle \leq \frac{b}{2} + \sqrt{\frac{b^2}{4} + c}. \quad (9)$$

For positive functions $f \geq 0$, we can upper bound $\langle \psi_Q, f\psi_P \rangle$ via the Cauchy-Schwarz inequality and obtain a lower bound on the expectation of f under Q . Taking as our function f to be the loss function of interest $f := \ell$, under the assumption that $\sup_{z \in \mathcal{Z}} |\ell(z)| \leq M$ for some $M > 0$, we can finally recast this lower bound as an upper bound on the expectation of ℓ under Q . Taking the supremum with respect to Q leads to a bound on the worst-case risk (1). We remark that in this way, we obtain both lower and upper bounds on the expected loss. As we will show, these bounds can be used to bound useful statistics, such as the accuracy or the AUC score used in binary classification. In the following Theorem, we state our main result as an upper bound to the worst-case risk (1) and refer the reader to Appendix A.2 for the analogous lower bound.

Theorem 2.2. Let $\ell: \mathcal{Z} \rightarrow \mathbb{R}_+$ be a loss function and suppose that $\sup_{z \in \mathcal{Z}} |\ell(z)| \leq M$ for some $M > 0$. Then, for any probability measure P on \mathcal{Z} and $\rho > 0$ we have

$$\begin{aligned} \sup_{Q \in B_\rho(P)} \mathbb{E}_Q[\ell(Z)] &\leq \mathbb{E}_P[\ell(Z)] + 2C_\rho \sqrt{\mathbb{V}_P[\ell(Z)]} \\ &+ \rho^2(2 - \rho^2) \left[M - \mathbb{E}_P[\ell(Z)] - \frac{\mathbb{V}_P[\ell(Z)]}{M - \mathbb{E}_P[\ell(Z)]} \right] \end{aligned} \quad (10)$$

where $C_\rho = \sqrt{\rho^2(1 - \rho^2)^2(2 - \rho^2)}$ and $B_\rho(P) = \{Q \in \mathcal{P}(\mathcal{Z}): H(P, Q) \leq \rho\}$ is the Hellinger ball of radius ρ

centered at P . The radius ρ is required to be small enough such that

$$\rho^2 \leq 1 - \left[1 + \frac{(M - \mathbb{E}_P[\ell(Z)])^2}{\mathbb{V}_P[\ell(Z)]} \right]^{-1/2}. \quad (11)$$

We refer the reader to Appendix A.1 for a full proof and now make some general observations about this result. The bound (10) presents a *pointwise* guarantee in the sense that it upper bounds the distributional worst-case risk for a particular model $\ell(\cdot)$. This is in contrast to bounds which hold uniformly for an entire model class and introduce complexity measures such as covering numbers and VC-dimension which are hard to compute for many practical problems. Other techniques which yield a pointwise robustness certificate of the form (10), typically express the uncertainty set as a Wasserstein ball around the distribution P (Sinha et al., 2018; Shafieezadeh-Abadeh et al., 2019; Blanchet & Murthy, 2019; Cranko et al., 2021), and require the model ℓ to be sufficiently smooth. For example, the certificate presented in (Sinha et al., 2018) can only be tractably computed for small neural networks for which one can upper bound their smoothness by bounding the Lipschitz constant of their gradients. For more general and large-scale neural networks, these bounds quickly become intractable and/or lead to vacuous certificates. For example, it is known that computing the Lipschitz constant of neural networks with ReLU activations is NP-hard (Virmaux & Scaman, 2018). Secondly, we emphasize that our bound (10) is “faithful”, in the sense that, as the radius approaches zero, $\rho \rightarrow 0$, the bound converges towards the true expectation $\mathbb{E}_P[\ell(Z)]$. This is of course desirable for any such bound as it indicates that any intrinsic gap vanishes as the covered distributions become increasingly closer to the reference distribution P . A third observation is that the bound (10) is *monotonically increasing* in the variance, indicating that low-variance models exhibit better generalization properties, which can be seen in light of the bias-variance trade-off. More specifically, from the form our bound (10) takes, we see that minimizing the variance-regularized objective $\mathcal{L}(\theta) = \mathbb{E}_{Z \sim P}[\ell_\theta(Z)] + \lambda \mathbb{V}_{Z \sim P} \ell_\theta(Z)$, effectively amounts to minimizing an upper bound on the worst-case risk. Indeed, various recent works have highlighted the connection between variance regularization and generalization (Lam, 2016; Maurer & Pontil, 2009; Gotoh et al., 2018; Duchi & Namkoong, 2019) and our result provides further evidence for this observation.

3. Certifying Out-of-domain Generalization

Taking advantage of our weak assumptions on the loss functions and models, we now apply our framework to the problem of certifying the out-of-domain generalization performance of a given classifier, when measured in terms of different loss functions. In practice, one is typically only given

a finite sample Z_1, \dots, Z_n from the in-domain distribution P and the bound (10) needs to be estimated empirically. To address this problem, our next step is to present a finite sampling version of the bound (10) which holds with arbitrarily high probability over the distribution P . Second, we instantiate our results with specific distribution shifts, namely, shifts in the label distribution, and shifts which only affect the covariates. Finally, we highlight specific loss and score functions and show how our result can be applied to certify the out-of-domain generalization of these functions.

3.1. Finite Sample Results

Let $Z_1, \dots, Z_n \stackrel{iid}{\sim} P$ be an independent and identically distributed sample from the in-domain distribution P . One immediate way to use our bound would be to construct the empirical distribution \hat{P}_n and consider the worst-case risk over distributions $Q \in B_\rho(\hat{P}_n)$, while computing the bound on the right hand side of (10) with the empirical mean and unbiased sample variance. However, for $\rho < 1$, the Hellinger ball $B_\rho(\hat{P}_n)$ will in general only contain distributions with discrete support since any continuous distribution Q has distance 1 from \hat{P}_n . We therefore seek another path and make use of concentration inequalities for the population variance and mean, in order to get statistically sound guarantees which hold with arbitrarily high probability. To achieve this, we bound the expectation value via Hoeffding’s inequality (Hoeffding, 1963), and the population variance via a bound presented in (Maurer & Pontil, 2009). In the second step, we use the union bound as a means to bound both variance and expectation simultaneously with high probability. We leave the derivation and proof to Appendix 3.1. These ingredients lead to the finite sampling-based version of Theorem 2.2, which we state in the following Corollary.

Corollary 3.1 (Finite-sampling bound). *Let Z_1, \dots, Z_n be independent random variables drawn from P and taking values in \mathcal{Z} . For a loss function $\ell: \mathcal{Z} \rightarrow [0, M]$, let $\hat{L}_n := \frac{1}{n} \sum_{i=1}^n \ell(Z_i)$ be the empirical mean and $S_n^2 := \frac{1}{n(n-1)} \sum_{1 \leq i < j \leq n} (\ell(Z_i) - \ell(Z_j))^2$ be the unbiased estimator of the variance of the random variable $\ell(Z)$, $Z \sim P$. Then, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$*

$$\begin{aligned} \sup_{Q \in B_\rho(P)} \mathbb{E}_Q[\ell(Z)] &\leq \hat{L}_n + 2C_\rho \sqrt{S_n^2} + \Delta_{n,\rho} \\ &+ \rho^2(2 - \rho^2) \left[M - \hat{L}_n \right. \\ &\left. + \frac{S_n^2 + 2M \sqrt{\frac{2S_n^2 \ln 2/\delta}{n-1}} + \frac{2M^2 \ln 2/\delta}{n-1}}{\hat{L}_n - M \left(1 - \sqrt{\frac{\ln 2/\delta}{2n}} \right)} \right] \end{aligned} \quad (12)$$

where

$$\Delta_{n,\rho} = \left(\frac{2C_\rho}{\sqrt{n-1}} - \frac{\rho^2(2 - \rho^2)}{2\sqrt{n}} \right) M \sqrt{2 \ln 2/\delta} \quad (13)$$

and $C_\rho = \sqrt{\rho^2(1-\rho^2)^2(2-\rho^2)}$. The radius ρ is required to be small enough such that

$$\rho^2 \leq 1 - \left[1 + \left(\frac{\hat{L}_n - M \left(1 - \sqrt{\frac{\ln 2/\delta}{2n}} \right)}{\sqrt{S_n^2} + M \sqrt{\frac{2 \ln 1/\delta}{n-1}}} \right)^2 \right]^{-1/2}. \quad (14)$$

Thus, we have derived a certificate for out-of-domain generalization for general bounded loss functions and models h which can be efficiently estimated from finite data sampled from the distribution P .

3.2. Specific Distribution Shifts

We now consider specific distribution shifts and discuss our main results in light of shifts in the distributions of labels and covariates.

3.2.1. LABEL DISTRIBUTION SHIFTS

Shifts in the label distribution occur when, during deployment, an ML-system operates in an environment where the relative frequency of certain classes increases or decreases, compared to the training environment, or, as is common in practical applications, instances of previously unseen classes appear. This can potentially harm the model performance dramatically and can have severe implications, in particular in the context of fairness and ethics in machine learning. To investigate this type of distribution shift, we follow the common practice to assume that the distribution over covariates, conditioned on the labels, stays constant. Formally, here, we consider the distribution shift $P \rightarrow Q$ expressed via

$$p(x, y) = \pi(x|y)p(y) \mapsto q(x, y) = \pi(x|y)q(y) \quad (15)$$

where $\pi(x|y)$ is given by a fixed distribution over covariates, conditioned on labels. In this case, it can be shown that the Hellinger distance is equal to the L_2 norm between the square roots of the (label) probability vectors $p = (p(1), \dots, p(K))^T \in \mathbb{R}^K$ and $q = (q(1), \dots, q(K))^T \in \mathbb{R}^K$ where K is the number of classes, so that

$$H(P, Q) = \frac{1}{\sqrt{2}} \|\sqrt{p} - \sqrt{q}\|_2 \quad (16)$$

where the square root is applied to each element in the respective probability vector.

3.2.2. COVARIATE DISTRIBUTION SHIFTS

In contrast to label distribution shifts, here we consider shifts to the distribution of covariates. This models scenarios where the relative frequency of labels stays constant,

but environments change, for example the shift from day to night in autonomous driving or wildlife surveillance. Formally, we consider the shift $P \rightarrow Q$ with

$$p(x, y) = \pi(y|x)p(x) \mapsto q(x, y) = \pi(y|x)q(x) \quad (17)$$

where $\pi(y|x)$ is given by a fixed distribution on labels, conditioned on the covariates. In this scenario, the Hellinger distance between P and Q reduces to the distance between the marginals

$$H(P, Q) = \sqrt{\frac{1}{2} \int_{\mathcal{X}} \left(\sqrt{p(x)} - \sqrt{q(x)} \right)^2 dx}. \quad (18)$$

In principle, this quantity could be estimated from unlabeled samples of a target distribution Q , enabling one to reason about distributional robustness of a given model, by evaluating our bounds from Theorem 2.2 and Corollary 3.1. However, in practice, it is generally difficult to estimate f -divergences, and in particular the Hellinger distance, from data for practically relevant problem instances. Although first steps in this direction have been made (Nguyen et al., 2007; 2010; Sreekumar et al., 2021), it remains largely an open problem and a potential solution would give our approach additional ounces of practical significance. We view this problem as orthogonal to certifying out-of-domain generalization and believe that research efforts towards such an end-to-end solution pose an exciting future research direction.

Discussion. We notice that when considering label- or covariate distribution shifts, we are effectively interested in a subset of all probability distributions with a given predefined Hellinger distance. In other words, if the shift $P \rightarrow Q$ models the label distribution shift with distance $H(P, Q) \leq \rho$, then applying the certificate (10) with radius ρ also covers every other type distribution shift bounded by ρ and hence gives a more conservative view than desired. This is because, in general,

$$\sup_{\substack{Q: H(P, Q) \leq \rho \\ q(\cdot|y) \equiv p(\cdot|y)}} \mathbb{E}_Q[\ell(Z)] \leq \sup_{Q: H(P, Q) \leq \rho} \mathbb{E}_Q[\ell(Z)] \quad (19)$$

arising from the additional constraint that $q(x|y) = p(x|y)$ for all $x \in \mathcal{X}$. A similar argument can be made for covariate shifts. Naturally, this leads to an intrinsic gap between the actual and certified robustness, which we also observe in our experiments. Finally, it is worth pointing out the connection with generalization from finite amounts of data which can be seen as a specific instantiation of the worst-case risk (1) where the in-domain distribution corresponds to the empirical distribution \hat{P}_n and the radius ρ decays as $\mathcal{O}(1/n)$. In this sense, the distribution shift originates from the transition from the empirical to the true data distribution. This type of distribution shift has been analyzed in (Duchi & Namkoong, 2019) where further links to variance-based regularization have been established.

3.3. Specific Loss and Score Functions

We now turn our attention to specific loss functions and discuss, in particular, the Jensen-Shannon divergence loss, the classification error, and the AUC score.

3.3.1. JENSEN-SHANNON DIVERGENCE

The Jensen-Shannon Divergence is a particular type of loss function for classification models, and serves as a symmetric alternative to other common losses such as cross entropy. It has been observed that the JSD loss and its generalizations have favorable properties compared to the standard cross entropy loss, as it is bounded, symmetric, and its square root is a distance and hence satisfies the triangle inequality. In (Englesson & Azizpour, 2021) it has been observed that JSD loss can be seen as an interpolation between cross-entropy and mean absolute error and is particularly well suited for classification problems with noisy labels. Formally, the Jensen-Shannon divergence is defined as

$$D_{JS}(P, Q) := \frac{1}{2} \left(D_{KL}(P \| \mu) + D_{KL}(Q \| \mu) \right) \quad (20)$$

where D_{KL} is the Kullback-Leibler divergence and $\mu = \frac{1}{2}(P + Q)$. Since it is a bounded loss function, it is straightforward to apply our results to certify the out-of-domain generalization for the JSD loss and, due to its smoothness, allows for a principled comparison between our bound and the Wasserstein distance certificates proposed in (Sinha et al., 2018; Cranko et al., 2021).

3.3.2. CLASSIFICATION ERROR

The classification error is among the most popular choices for measuring the performance of classification models and serves as a means to assess how accurate a classifier is on a given data distribution. As it is a non-smooth function, existing approaches cannot in general certify distributional robustness for this function. In contrast, one can immediately instantiate our Theorem 2.2 (or the finite sampling version from Corollary 3.1) with this loss. Indeed, for a fixed model $h: \mathcal{X} \rightarrow \mathcal{Y}$, let $\epsilon_P := \mathbb{P}_{(X,Y) \sim P}[h(X) \neq Y]$ and analogously ϵ_Q . Then, in the infinite sampling regime, we immediately get an upper bound on the worst-case classification error from Theorem 2.2. Namely, for a sufficiently small radius $\rho^2 \leq 1 - \sqrt{\epsilon_P}$, we have

$$\sup_{Q \in B_\rho(P)} \epsilon_Q \leq \epsilon_P + 2C_\rho \sqrt{\epsilon_P(1 - \epsilon_P)} + \rho^2(2 - \rho^2)(1 - 2\epsilon_P) \quad (21)$$

where $C_\rho = \sqrt{\rho^2(1 - \rho^2)^2(2 - \rho^2)}$.

3.3.3. AUC SCORE

Among other uses, the Area under the ROC (AUC) score (Hanley & McNeil, 1982; Cl  men  on et al., 2008)

is a metric to measure the performance of binary classification models. Unlike the classification error, which captures the ability to classify a single randomly chosen instance, the AUC score provides a means to quantify the ability to correctly assigning to any positive instance a higher score than to a randomly chosen negative instance. For a binary classification model $h: \mathcal{X} \rightarrow \mathbb{R}$ that outputs the score of the positive class, the AUC score is defined as

$$\text{AUC}(h) = \mathbb{P}[h(X) \geq h(X') | Y = 1, Y' = -1] \quad (22)$$

where (X, Y) and (X', Y') are independent and identically distributed according to P . By introducing the notation $X_\pm := X | Y = \pm 1$, we can equivalently write the AUC score as an expectation value over the joint (conditional) distribution of $Z := (X_+, X_-)$

$$\text{AUC}(h) = \mathbb{E}_{(X_+, X_-) \sim P_Z} [\mathbb{1}_{\{h(X_+) \geq h(X_-)\}}]. \quad (23)$$

We notice that only distribution shifts on the covariates have an impact on the AUC score. For this reason, we consider a setting similar to the covariate shift setting of Sec. 3.2.2, although we consider shifts in the conditional distribution $p(x|y) \mapsto q(x|y)$ for each $y \in \{\pm 1\}$ in contrast to shifts in the marginals. Due to independence, the probability density function of $Z \sim P_Z$ can be written as

$$p_Z(x_+, x_-) = p(x|y = +1)p(x|y = -1) \quad (24)$$

and similarly for the shifted distribution Q . Thus, assuming that for both negative and positive samples a distribution drift with $H(P_{X|Y=y}, Q_{X|Y=y}) \leq \rho$ occurs, the squared Hellinger distance between P_Z and Q_Z is bounded by

$$H^2(P_Z, Q_Z) \leq \rho^2(2 - \rho^2). \quad (25)$$

Thus, for certifying out-of-domain generalization for the AUC score, we can apply our bound by instantiating it with Hellinger distance $\sqrt{\rho^2(2 - \rho^2)}$. We remark that for the AUC score, one is typically interested in *lower* bounding it under distribution shifts. To that end, we present a lower bound version of our Theorem 2.2 in Appendix A.2.

4. Experiments

We now experimentally validate our theoretical findings on a diverse collection of datasets and scenarios. We first provide certificates considering generic distribution shifts $P \rightarrow Q$ and then provide detailed analysis on the two specific scenarios described in Sections 3.2.1 and 3.2.2, namely, shifts in the label and in the covariate distributions. Finally, we construct a synthetic example that allows for a fair comparison of our bounds with the Wasserstein certificate of (Sinha et al., 2018), which indicates that in addition to favorable scalability properties, our bounds are also considerably tighter. We remark that all our bounds are computed

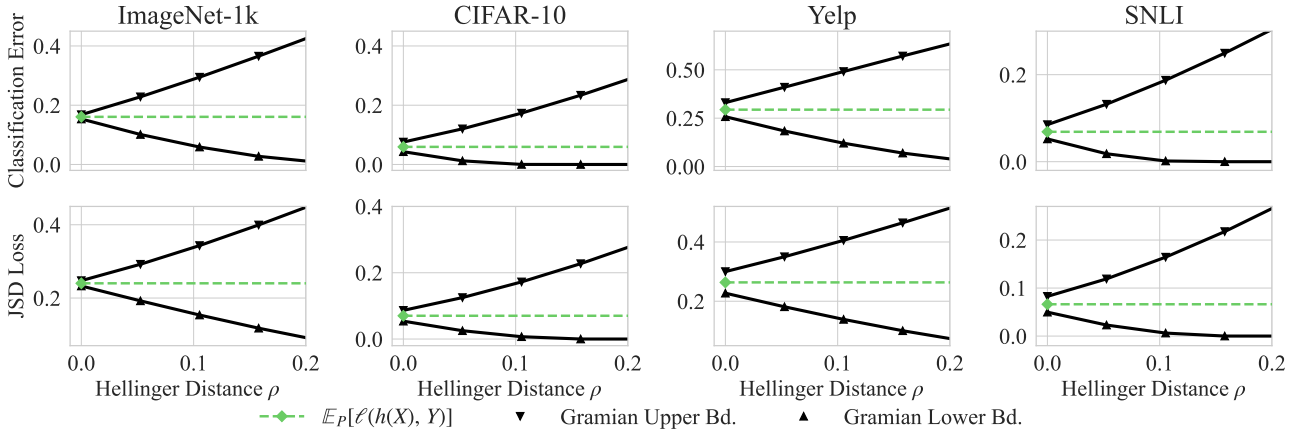


Figure 1. Distributional robustness certificates for generic distribution shifts on vision and NLP datasets for JSD and 0-1 loss.




Figure 2. Distributional robustness certificates for AUC against generic distribution shifts on binary ImageNet and CIFAR datasets.

using the finite sampling bounds presented in Corollary 3.1 and hold with 99% probability ($\delta = 0.01$).⁴

Datasets We certify out-of-domain generalization on two standard vision datasets: ImageNet-1k (Russakovsky et al., 2015) containing objects of 1,000 different classes; and CIFAR-10 (Krizhevsky, 2009), which contains natural images of 10 different classes. We also conduct experiments on the standard natural language processing (NLP) datasets Yelp (Challenge) and SNLI (Bowman et al., 2015). We follow Lin et al. (2017) to sample 2,000 examples for the Yelp test set and 10,000 examples for the SNLI test set.

Models For classification on ImageNet-1k, we use the EfficientNet-B7 (Tan & Le, 2019) architecture which we initialize with pre-trained weights; we use DenseNet-121 (Huang et al., 2017) for CIFAR-10. On Yelp, we use BERT (Devlin et al., 2018) and on SNLI we use a DeBERTa architecture (He et al., 2020).

Settings for AUC Scores When we consider AUC scores, we further constrain all multiclass datasets into a binary

⁴Our code is publicly available at <https://github.com/DS3Lab/certified-generalization>.




Figure 3. Certified Generalization for label distribution shifts. Each gray point corresponds to a randomly sampled label distribution with corresponding Hellinger distance and empirical loss.

version. To this end, on ImageNet, we randomly choose two classes and train a ResNet-152 architecture to discriminate between the two Synsets *n01601694* and *n04330267* (corresponding to the classes ‘water ouzel’ and ‘stove’). Similarly, on CIFAR-10 we also pick two classes at random and train a ResNet-110 classifier for the two classes ‘bird’ and ‘horse’.

4.1. Certifying Distribution Shifts

Figures 1 and 2 illustrate the certificates that we provide on a diverse range of datasets, considering three different scores: classification error, JSD loss, and AUC score. In all these figures, the x-axis corresponds to the degree of distribution drift, and the *Gramian Certificate* curves correspond to the lower and upper bound of these scores under distribution drifts. To our best knowledge, this is the first time that nonvacuous certificates are obtained on this diverse range of datasets, scores, and large-scale models.

Label Distribution Shifts To get a better indication of how well our certificates capture the true risk under label distribution shifts, we randomly generate 100,000 shifted class distributions on the CIFAR-10 and Yelp datasets by 1) subsampling existing classes, 2) removing the counts






Figure 4. Certificate against covariate shift on colored MNIST.

of existing classes, and 3) including new "unseen" classes. This allows us to empirically compute both the classification error and the Hellinger distance and enables us to compare the certificates to the actual loss on the shifted distribution. We can see from Figure 3 that our certificates indeed provide a valid upper and lower bound. Note that, given that all shifted class distributions are randomly sampled, we might not hit the true worst-case scenario, explaining the clear gap between the generalization certificates and the scores obtained from the randomly generated label distributions. Another reason for the gap can be attributed to the intrinsic gap for label and covariate shifts, discussed in Section 3.2. We refer the reader to Appendix F.1 for analogous figures with a larger set of model architectures on the CIFAR-10 dataset. Finally, we point out the difficulty in sampling these class distributions for datasets with a large number of classes and include analogous figures for ImageNet and the SNLI dataset in Appendix F.2.

Covariate Distribution Shifts. We now investigate our certificates in light of changes in the distribution of the covariates and consider the scenario described in Sec. 3.2.2. In this experiment, we use the binary Colored MNIST dataset (Kim et al., 2019; Arjovsky et al., 2019), which is constructed from the MNIST dataset by coloring the digits 0-4 in green and 5-9 in red for the training set, while flipping the coloring in the test set. The classifier is then trained to classify the digits into the two groups $\{0, 1, 2, 3, 4\}$ and $\{5, 6, 7, 8, 9\}$. In this setting, the classifier learns to perfectly distinguish the two classes in the training set, but fails on the testing set since the color is a stronger predictor than the shape of the digits. To investigate the space between these two extreme cases, we generate mixture distributions between training and test distribution in the following way. We set P to be the training distribution and Q the testing distribution (containing digits with flipped colors). Guided by a mixing parameter γ , we mix P and Q to obtain the mixture distribution $\Pi_\gamma := \gamma \cdot P + (1 - \gamma) \cdot Q$. Since P and Q have disjoint support, we compute the Hellinger distance between P and Π_γ as $H(P, \Pi_\gamma) = \sqrt{1 - \sqrt{\gamma}}$ as shown in Appendix E. Figure 4 illustrates our robustness certificates for the 0-1 loss and the AUC Score, as well as the empirical losses $\mathbb{E}_{\Pi_\gamma}[\ell(X, Y)]$ for different values of the mixture



(a) Robustness certificates for varying network widths w and $n_h = 2$ hidden layers.



(b) Robustness certificates for varying numbers of hidden layers n_h and fixed width $w = 4$.

Figure 5. Comparison of our approach with the Wasserstein-based certificates from (Sinha et al., 2018; Cranko et al., 2021) for varying levels of model complexity.

parameter γ . We see from the figure that our technique provides quite tight certificates for both classification error and AUC score.

4.2. Comparison with Wasserstein Certificates

We now construct a synthetic example that enables a fair comparison with two baseline certificates based on the Wasserstein distance. Namely, we compare our approach with 1) the certificate which uses the Lipschitz constant of the ML model, presented in (Cranko et al., 2021); and 2) with the pointwise robustness certificate derived in (Sinha et al., 2018) from the dual formulation of the worst-case risk. We remark that these certificates cannot be applied to our previous examples because of their prohibitive assumptions. To make the three techniques comparable, we consider a Gaussian mixture model and certify the Jensen-Shannon divergence loss, while modeling distribution shifts as dislocations, $X \mapsto X + \delta$ for a fixed perturbation vector δ . This allows us to parameterize the distribution shift via the

L_2 -norm of δ and obtain a one-to-one correspondence between our Hellinger distance and the Wasserstein distance, and enables a principled comparison. We describe the details of this synthetic dataset in Appendix C. To investigate how the techniques scale with increased model complexity, we use fully connected feedforward neural networks with varying depths and widths. In addition, to accommodate (Sinha et al., 2018)’s assumptions on smoothness, we use ELU activation functions on all layers. We remark that the bound in (Sinha et al., 2018) requires one to solve a complex maximization problem, which requires the composition of the loss function and the network to be sufficiently smooth. Furthermore, the concavity of the maximization problem hinges on knowledge of the Lipschitz constant of the gradient. For small examples, this Lipschitz constant can be obtained, as we show in Appendix D for the JSD loss function. As can be seen in Figure 5, all bounds converge to the expected loss $\mathbb{E}_P[\ell(X, Y)]$ as the perturbation goes to zero, $\|\delta\|_2 \rightarrow 0$. However, the certificate from (Sinha et al., 2018) quickly becomes vacuous as the perturbation magnitude increases. In addition, both baseline bounds become loose with increasing model complexity, while our bound is virtually agnostic to the model architecture as it only depends on the variance and expected loss on the distribution P .

5. Related Work

Distributionally robust optimization first appeared in the context of inventory management (Scarf, 1958) and has since been discovered by the machine learning community as a useful tool to train machine learning models which generalize better to new distributions (Ben-Tal et al., 2013; Gao & Kleywegt, 2016; Shafieezadeh-Abadeh et al., 2019). The uncertainty set occurring in the distributionally robust loss has been studied in terms of Wasserstein balls in (Gao & Kleywegt, 2016; Sinha et al., 2018; Shafieezadeh-Abadeh et al., 2019; Cranko et al., 2021; Lee & Raginsky, 2017; Cisse et al., 2017; Kuhn et al., 2019; Blanchet & Murthy, 2019), and f -divergence balls in (Ben-Tal et al., 2013; Duchi et al., 2021; Lam, 2016; Duchi & Namkoong, 2019; 2021). From a more general viewpoint, (Husain, 2020) connects integral probability metrics with distributional robustness in general and provides links with generative adversarial networks. In another vein, maximum mean discrepancy measures have been investigated in (Staub & Jegelka, 2019) for generalization in Kernel methods. (Sinha et al., 2018) propose a method to certify generalization by using the dual formulation of the Wasserstein worst-case risk. However, their approach requires the loss model and loss function to be smooth and relies on an estimate of the Lipschitz constant of gradients, which quickly becomes vacuous for large problem sizes. Related techniques based on Wasserstein distances (Gao & Kleywegt, 2016; Shafieezadeh-Abadeh et al., 2019; Blanchet & Murthy, 2019; Kuhn et al., 2019; Cranko

et al., 2021) make similarly prohibitive assumptions and generally fail to provide scalable alternatives. In contrast, we study uncertainty sets expressed as Hellinger balls and provide a model-specific distributional robustness guarantee which only makes minimal assumptions on the loss (namely, boundedness) and thus scales to large problems. The authors in (Subbaswamy et al., 2021) consider distributionally robust optimization under fine-grained shifts in the marginal distributions, and reason about the worst-case risk on sub-populations in the data distribution. Orthogonal to our work is the topic of certified adversarial robustness (Wong et al., 2018; Lecuyer et al., 2019; Cohen et al., 2019; Szegedy et al., 2014; Carlini & Wagner, 2017). This line of research seeks to reason about robustness at the instance level, while we aim to bound the worst-case risk over a set of distributions.

6. Conclusion

In this paper, we have studied the problem of certifying the out-of-domain generalization for blackbox functions. To that end, we have presented a framework to bound the worst-case population risk over an uncertainty set of probability distributions given by a Hellinger ball. In contrast to existing approaches, our framework is scalable since it treats the loss function together with the model as a blackbox and thus requires virtually no knowledge about the internals of, e.g., neural networks. We have provided experimental evidence that our technique can handle large models and datasets and provides, to the best of our knowledge, the first non-vacuous out-of-domain generalization bounds for problems as large as ImageNet with a full-fledged EfficientNet-B7. While our techniques provide a means to certify robustness against *general* distribution shifts, future research directions can potentially extensively study more specific distribution shifts. In addition, it will be interesting to link our results to related topics such as fairness in machine learning.

Acknowledgments

CZ and the DS3Lab gratefully acknowledge the support from the Swiss State Secretariat for Education, Research and Innovation (SERI)’s Backup Funding Scheme for European Research Council (ERC) Starting Grant TRIDENT (101042665), the Swiss National Science Foundation (Project Number 200021_184628, and 197485), InnoSuisse/SNF BRIDGE Discovery (Project Number 40B2-0_187132), European Union Horizon 2020 Research and Innovation Programme (DAPHNE, 957407), Botnar Research Centre for Child Health, Swiss Data Science Center, Alibaba, Cisco, eBay, Google Focused Research Awards, Kuaishou Inc., Oracle Labs, Zurich Insurance, and the Department of Computer Science at ETH Zurich. BL, LL and BW are supported by NSF grant No.1910100, NSF CNS 2046726, C3 AI, and the Alfred P. Sloan Foundation

References

- AlBadawy, E. A., Saha, A., and Mazurowski, M. A. Deep learning for segmentation of brain tumors: Impact of cross-institutional training and testing. *Medical physics*, 45(3):1150–1158, 2018.
- Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.
- Beery, S., Van Horn, G., and Perona, P. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 456–473, 2018.
- Ben-Tal, A., Den Hertog, D., De Waegenare, A., Melenberg, B., and Rennen, G. Robust solutions of optimization problems affected by uncertain probabilities. *Management Science*, 59(2):341–357, 2013.
- Blanchet, J. and Murthy, K. Quantifying distributional model risk via optimal transport. *Mathematics of Operations Research*, 44(2):565–600, 2019.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. A large annotated corpus for learning natural language inference. In *EMNLP*, 2015.
- Carlini, N. and Wagner, D. Adversarial examples are not easily detected: Bypassing ten detection methods. In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pp. 3–14, 2017.
- Challenge, Y. D. data retrieved from Yelp Dataset Challenge, <https://www.yelp.com/dataset/challenge>.
- Cisse, M., Bojanowski, P., Grave, E., Dauphin, Y., and Usunier, N. Parseval networks: Improving robustness to adversarial examples. In *International Conference on Machine Learning*, pp. 854–863. PMLR, 2017.
- Cléménçon, S., Lugosi, G., and Vayatis, N. Ranking and empirical minimization of u-statistics. *The Annals of Statistics*, 36(2):844–874, 2008.
- Cohen, J., Rosenfeld, E., and Kolter, Z. Certified adversarial robustness via randomized smoothing. In *International Conference on Machine Learning*, pp. 1310–1320. PMLR, 2019.
- Cranko, Z., Shi, Z., Zhang, X., Nock, R., and Kornblith, S. Generalised lipschitz regularisation equals distributional robustness. In *International Conference on Machine Learning*, pp. 2178–2188. PMLR, 2021.
- Dai, D. and Van Gool, L. Dark model adaptation: Semantic image segmentation from daytime to nighttime. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 3819–3824. IEEE, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Duchi, J. and Namkoong, H. Variance-based regularization with convex objectives. *The Journal of Machine Learning Research*, 20(1):2450–2504, 2019.
- Duchi, J. C. and Namkoong, H. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, 2021.
- Duchi, J. C., Glynn, P. W., and Namkoong, H. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 2021.
- Engleson, E. and Azizpour, H. Generalized jensen-shannon divergence loss for learning with noisy labels. In Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems*, 2021. URL <https://openreview.net/forum?id=TiwPYwg3IRf>.
- Gao, R. and Kleywegt, A. J. Distributionally robust stochastic optimization with wasserstein distance. *arXiv preprint arXiv:1604.02199*, 2016.
- Gotoh, J.-y., Kim, M. J., and Lim, A. E. Robust empirical optimization is almost the same as mean–variance optimization. *Operations research letters*, 46(4):448–452, 2018.
- Gulrajani, I. and Lopez-Paz, D. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=lQdXeXDoWtI>.
- Hanley, J. A. and McNeil, B. J. The meaning and use of the area under a receiver operating characteristic (roc) curve. *Radiology*, 143(1):29–36, 1982.
- He, P., Liu, X., Gao, J., and Chen, W. Deberta: Decoding-enhanced bert with disentangled attention. *arXiv preprint arXiv:2006.03654*, 2020.
- Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Husain, H. Distributional robustness with ipms and links to regularization and gans. *arXiv preprint arXiv:2006.04349*, 2020.

- Kim, B., Kim, H., Kim, K., Kim, S., and Kim, J. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9012–9020, 2019.
- Koh, P. W., Sagawa, S., Marklund, H., Xie, S. M., Zhang, M., Balsubramani, A., Hu, W., Yasunaga, M., Phillips, R. L., Gao, I., et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pp. 5637–5664. PMLR, 2021.
- Krizhevsky, A. Learning multiple layers of features from tiny images. Technical report, 2009.
- Kuhn, D., Esfahani, P. M., Nguyen, V. A., and Shafieezadeh-Abadeh, S. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations Research & Management Science in the Age of Analytics*, pp. 130–166. INFORMS, 2019.
- Lam, H. Robust sensitivity analysis for stochastic systems. *Mathematics of Operations Research*, 41(4):1248–1275, 2016.
- Lecuyer, M., Atlidakis, V., Geambasu, R., Hsu, D., and Jana, S. Certified robustness to adversarial examples with differential privacy. In *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 656–672. IEEE, 2019.
- Lee, J. and Raginsky, M. Minimax statistical learning with wasserstein distances. *arXiv preprint arXiv:1705.07815*, 2017.
- Lin, Z., Feng, M., Santos, C. N. d., Yu, M., Xiang, B., Zhou, B., and Bengio, Y. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130*, 2017.
- Maurer, A. and Pontil, M. Empirical bernstein bounds and sample variance penalization. In *Proceedings of the Twenty Second Annual Conference on Computational Learning Theory*, 2009.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Nonparametric estimation of the likelihood ratio and divergence functionals. In *2007 IEEE International Symposium on Information Theory*, pp. 2016–2020. IEEE, 2007.
- Nguyen, X., Wainwright, M. J., and Jordan, M. I. Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*, 56(11):5847–5861, 2010.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3): 211–252, 2015.
- Scarf, H. A min-max solution of an inventory problem. *Studies in the mathematical theory of inventory and production*, 1958.
- Shafieezadeh-Abadeh, S., Kuhn, D., and Esfahani, P. M. Regularization via mass transportation. *Journal of Machine Learning Research*, 20(103):1–68, 2019.
- Sinha, A., Namkoong, H., and Duchi, J. Certifiable distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=Hk6kPgZA->.
- Sreekumar, S., Zhang, Z., and Goldfeld, Z. Non-asymptotic performance guarantees for neural estimation of f-divergences. In *International Conference on Artificial Intelligence and Statistics*, pp. 3322–3330. PMLR, 2021.
- Staib, M. and Jegelka, S. Distributionally robust optimization and generalization in kernel methods. *Advances in Neural Information Processing Systems*, 32:9134–9144, 2019.
- Subbaswamy, A., Adams, R., and Saria, S. Evaluating model robustness and stability to dataset shift. In *International Conference on Artificial Intelligence and Statistics*, pp. 2611–2619. PMLR, 2021.
- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., and Fergus, R. Intriguing properties of neural networks. In *International Conference on Learning Representations*, 2014. URL <http://arxiv.org/abs/1312.6199>.
- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114. PMLR, 2019.
- Virmaux, A. and Scaman, K. Lipschitz regularity of deep neural networks: analysis and efficient estimation. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.
- Volk, G., Müller, S., von Bernuth, A., Hospach, D., and Bringmann, O. Towards robust cnn-based object detection through augmentation with synthetic rain variations. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 285–292. IEEE, 2019.
- Weber, M., Anand, A., Cervera-Lierta, A., Kottmann, J. S., Kyaw, T. H., Li, B., Aspuru-Guzik, A., Zhang, C., and Zhao, Z. Toward reliability in the nisq era: Robust interval guarantee for quantum measurements on approximate states. *arXiv preprint arXiv:2110.09793*, 2021.

Weinhold, F. Lower bounds to expectation values. *Journal of Physics A: General Physics*, 1(3):305, 1968.
URL <https://doi.org/10.1088/0305-4470/1/3/301>.

Wong, E., Schmidt, F., Metzen, J. H., and Kolter, J. Z. Scaling provable adversarial defenses. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

A. Proofs

A.1. Proof of Theorem 2.2

We begin the proof by stating a lemma which allows one to bound inner products between elements of a Hilbert space \mathcal{H} .

Lemma A.1. *Let \mathcal{H} be a Hilbert space with inner product $\langle \cdot, \cdot \rangle$, let $A \in \mathcal{B}(\mathcal{H})$ be a positive semidefinite bounded linear operator on \mathcal{H} and let $u, v \in \mathcal{H} \setminus \{0\}$ be such that*

$$\frac{|\langle u, v \rangle|^2}{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2)} \geq \frac{\Delta}{\langle v, Av \rangle^2}, \quad (26)$$

where

$$\Delta := (\langle v, v \rangle \langle Av, Av \rangle - \langle Av, v \rangle^2). \quad (27)$$

Then

$$\langle u, Au \rangle \geq \frac{|\langle u, v \rangle|^2 \langle v, Av \rangle}{\langle v, v \rangle^2} - \frac{2|\langle u, v \rangle| \sqrt{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}}{\langle v, v \rangle^2} + \frac{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}{\langle v, v \rangle^2 \langle v, Av \rangle} \quad (28)$$

Proof of Lemma A.1. In the following we denote by $\Re(z) := \frac{1}{2}(z + \bar{z})$ and $\Im(z) := \frac{1}{2i}(z - \bar{z})$ the real and imaginary parts of a complex number $z \in \mathbb{C}$. Let G be the Gram matrix of the vectors u, v, Av and recall that Gram matrices are positive semidefinite, $G \geq 0$. Since the determinant of a matrix is given by the product of its eigenvalues, it follows that

$$\begin{aligned} 0 \leq \det(G) &= \begin{vmatrix} \langle u, u \rangle & \langle u, v \rangle & \langle u, Av \rangle \\ \langle v, u \rangle & \langle v, v \rangle & \langle v, Av \rangle \\ \langle Av, u \rangle & \langle Av, v \rangle & \langle Av, Av \rangle \end{vmatrix} \\ &= [(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \langle Av, Av \rangle - \langle u, u \rangle |\langle Av, v \rangle|^2] \\ &\quad + 2\Re(\langle u, v \rangle \langle v, Av \rangle \langle Av, u \rangle) - \langle v, v \rangle |\langle Av, u \rangle|^2. \end{aligned} \quad (29)$$

Let $\phi \in \mathbb{R}$ be such that $e^{i\phi} \langle u, v \rangle = |\langle u, v \rangle|$ and let $\tilde{u} = e^{-i\phi} u$.⁵ Thus, we have

$$\begin{aligned} &[(\langle u, u \rangle \langle v, v - \langle u, v \rangle^2 \rangle) \langle Av, Av \rangle - \langle u, u \rangle |\langle Av, v \rangle|^2 - \langle v, v \rangle \Im(\langle Av, \tilde{u} \rangle)^2] \\ &\quad + 2|\langle u, v \rangle| \langle v, Av \rangle \Re(\langle Av, \tilde{u} \rangle) - \langle v, v \rangle \Re(\langle Av, \tilde{u} \rangle)^2 \geq 0 \end{aligned} \quad (30)$$

The LHS of this inequality can be seen as a quadratic polynomial in $\Re(\langle Av, \tilde{u} \rangle)$ and the non-negativity effectively constrains the values that $\Re(\langle Av, \tilde{u} \rangle)$ can take to be within the roots of the polynomial. Thus, we have, in particular,

$$\begin{aligned} \Re(\langle Av, \tilde{u} \rangle) &\geq \frac{|\langle u, v \rangle| \langle v, Av \rangle}{\langle v, v \rangle} - \frac{\sqrt{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta - \langle v, v \rangle \Im(\langle Av, \tilde{u} \rangle)^2}}{\langle v, v \rangle} \\ &\geq \frac{|\langle u, v \rangle| \langle v, Av \rangle}{\langle v, v \rangle} - \frac{\sqrt{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}}{\langle v, v \rangle} \end{aligned} \quad (31)$$

with $\Delta := (\langle v, v \rangle \langle Av, Av \rangle - \langle Av, v \rangle^2)$. Since A is positive semidefinite, it has a square root, i.e. there exists a linear operator $A^{1/2}$ with $A^{1/2} A^{1/2} = A$. It follows that

$$\begin{aligned} \Re(\langle Av, \tilde{u} \rangle) &\stackrel{(i)}{\leq} |\langle Av, \tilde{u} \rangle| = |\langle Av, u \rangle| \stackrel{(ii)}{=} |\langle A^{1/2} v, A^{1/2} u \rangle| \\ &\stackrel{(iii)}{\leq} \sqrt{\langle A^{1/2} u, A^{1/2} u \rangle} \sqrt{\langle A^{1/2} v, A^{1/2} v \rangle} \\ &= \sqrt{\langle u, Au \rangle} \sqrt{\langle v, Av \rangle} \end{aligned} \quad (32)$$

where in (i) we have used $\Re(z) \leq |z|$ for any $z \in \mathbb{C}$, in (ii) we have used that $A^{1/2}$ is self-adjoint and in (iii) we have used the Cauchy-Schwarz inequality. Combining this with (31) and dividing each side by $\sqrt{\langle v, Av \rangle}$ yields

$$\sqrt{\langle u, Au \rangle} \geq \frac{|\langle u, v \rangle| \sqrt{\langle v, Av \rangle}}{\langle v, v \rangle} - \frac{\sqrt{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}}{\langle v, v \rangle \sqrt{\langle v, Av \rangle}} \quad (33)$$

⁵We use the convention that inner products are linear in their second argument and conjugate linear in the first inner product.

The RHS in this inequality is non-negative as long as

$$|\langle u, v \rangle| \geq \frac{\sqrt{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}}{\langle v, Av \rangle}. \quad (34)$$

Thus, in this case, squaring both sides of (33) yields

$$\langle u, Au \rangle \geq \frac{|\langle u, v \rangle|^2 \langle v, Av \rangle}{\langle v, v \rangle^2} - 2 \frac{|\langle u, v \rangle| \sqrt{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}}{\langle v, v \rangle^2} + \frac{(\langle u, u \rangle \langle v, v \rangle - |\langle u, v \rangle|^2) \Delta}{\langle v, v \rangle^2 \langle v, Av \rangle} \quad (35)$$

which is the desired result. \square

We will now show how Lemma A.1 can be used to upper bound the worst-case risk (1). Let $Q \in \mathcal{P}(\mathcal{Z})$ be an arbitrary probability measure on \mathcal{Z} with $H(P, Q) \leq \rho$. Denote by ψ_P, ψ_Q the positive square roots of the Radon-Nikodym derivatives of P and Q , respectively, with respect to an arbitrary measure μ with $P, Q \ll \mu^6$

$$\psi_P := \sqrt{\frac{dP}{d\mu}} \quad \text{and} \quad \psi_Q := \sqrt{\frac{dQ}{d\mu}}. \quad (36)$$

Note that ψ_P and ψ_Q are square-integrable with respect to μ and real-valued, $\psi_P, \psi_Q \in L_2(\mathcal{Z}, \Sigma, \mu)$ where we set Σ to be the Borel σ -algebra on \mathcal{Z} and assume that L_2 contains only real-valued functions. It is well known that L_2 together with the inner product $\langle f, g \rangle_{L_2} := \int_{\mathcal{Z}} fg d\mu$ is a Hilbert space. Furthermore, the space of essentially bounded functions $L_\infty(\mathcal{Z}, \Sigma, \mu)$ is isometrically isomorphic to the set of linear bounded operators on \mathcal{Z} . That is, each $f \in L_\infty$ defines a linear operator M_f via pointwise multiplication, $L_2 \ni \psi \mapsto M_f \psi$ with $M_f \psi : z \mapsto (M_f \psi)(z) = f(z) \cdot \psi(z)$. It follows that for any $f \in L_\infty$, we can write its expectation with respect to P (and equivalently Q) in terms of the inner product on L_2

$$\mathbb{E}_{Z \sim P}[f(Z)] = \int_{\mathcal{Z}} f(z) dP(z) = \int_{\mathcal{Z}} f(z) \frac{dP}{d\mu}(z) d\mu(z) = \int_{\mathcal{Z}} \psi_P(z) f(z) \psi_P(z) d\mu(z) = \langle \psi_P, M_f \psi_P \rangle_{L_2} \quad (37)$$

Similarly, we can write the variance of $f(Z)$ with respect to P (and equivalently Q) in terms of inner products as

$$\mathbb{V}_{Z \sim P}[f(Z)] = \langle \psi_P, M_{f^2} \psi_P \rangle_{L_2} - \langle \psi_P, M_f \psi_P \rangle_{L_2}^2. \quad (38)$$

To simplify notation, we write $f \cdot \psi$ for the image of ψ under M_f for $f \in L_\infty$ and we drop the subscript in the inner product whenever it is clear from context. Recall that M is an upper bound on the loss function ℓ , so that $\sup_{z \in \mathcal{Z}} |\ell(z)| \leq M$. It follows that the function $f_\ell(\cdot) := M - \ell(\cdot)$ is essentially bounded with respect to μ and hence defines a bounded linear operator (which is also self-adjoint since we only consider real-valued functions in this work). Applying Lemma A.1 to the Hilbert space L_2 and identifying $u \equiv \psi_Q, v \equiv \psi_P$ and A with the operator defined by f_ℓ immediately yields the lower bound

$$\begin{aligned} \mathbb{E}_Q[M - \ell(Z)] &\geq |\langle \psi_P, \psi_Q \rangle|^2 \mathbb{E}_P[M - \ell(Z)] \\ &\quad - 2|\langle \psi_P, \psi_Q \rangle| \sqrt{(1 - |\langle \psi_P, \psi_Q \rangle|^2) \mathbb{V}_P[M - \ell(Z)]} + \frac{(1 - |\langle \psi_P, \psi_Q \rangle|^2) \mathbb{V}_P[M - \ell(Z)]}{\mathbb{E}_P[M - \ell(Z)]} \end{aligned} \quad (39)$$

Rearranging terms and noting that $\mathbb{V}_P[M - \ell(Z)] = \mathbb{V}_P[\ell(Z)]$ leads to

$$\begin{aligned} \mathbb{E}_Q[\ell(Z)] &\leq |\langle \psi_P, \psi_Q \rangle|^2 \mathbb{E}_P[\ell(Z) - M] + M \\ &\quad + 2|\langle \psi_P, \psi_Q \rangle| \sqrt{(1 - |\langle \psi_P, \psi_Q \rangle|^2) \mathbb{V}_P[\ell(Z)]} - \frac{(1 - |\langle \psi_P, \psi_Q \rangle|^2) \mathbb{V}_P[\ell(Z)]}{\mathbb{E}_P[M - \ell(Z)]}. \end{aligned} \quad (40)$$

Note that the inner product $\langle \psi_P, \psi_Q \rangle$ is known as the Hellinger affinity and related to the squared Hellinger distance between P and Q via

$$H^2(P, Q) = \frac{1}{2} \int_{\mathcal{Z}} (\psi_P - \psi_Q)^2 d\mu = 1 - \int_{\mathcal{Z}} \psi_P \psi_Q d\mu = 1 - \langle \psi_P, \psi_Q \rangle. \quad (41)$$

⁶Such a measure μ always exists as one can choose $\mu = P + Q$.

Thus, the requirement that the inner product $\langle \psi_P, \psi_Q \rangle$ satisfies is lower bounded by the quantity in (26) can be expressed as an upper bound on ρ^2 as

$$\rho^2 \leq 1 - \left[1 + \left(\frac{M - \mathbb{E}_P[\ell(Z)]}{\sqrt{\mathbb{V}_P[\ell(Z)]}} \right)^2 \right]^{-1/2}. \quad (42)$$

Finally, substituting $\langle \psi_P, \psi_Q \rangle = 1 - H^2(P, Q) = 1 - \rho^2$ in (40), setting $C_\rho = \sqrt{\rho^2(2 - \rho^2)(1 - \rho^2)^2}$ and rearranging terms yields

$$\begin{aligned} \mathbb{E}_Q[\ell(Z)] &\leq \mathbb{E}_P[\ell(Z)] + 2C_\rho \sqrt{\mathbb{V}_P[\ell(Z)]} \\ &\quad + \rho^2(2 - \rho^2) \left[M - \mathbb{E}_P[\ell(Z)] - \frac{\mathbb{V}_P[\ell(Z)]}{\mathbb{E}_P[M - \ell(Z)]} \right]. \end{aligned} \quad (43)$$

Since the choice of Q was arbitrary and the RHS in this inequality does not depend on Q , taking the supremum of the LHS over all Q with $H(P, Q) \leq \rho$ gives the desired result.

A.2. A lower bound version of Theorem 2.2

Given the proof of Theorem 2.2, it is straightforward to adapt it so as to yield a *lower* bound on expectation values using Lemma A.1. Indeed, by instantiating this Lemma with the function ℓ (instead of $f_\ell(\cdot) := M - \ell(\cdot)$) we obtain a lower bound by following the analogous, subsequent reasoning as in the proof of Theorem 2.2.

Theorem A.2 (Lower bound). *Let $\ell: \mathcal{Z} \rightarrow \mathbb{R}_+$ be a nonnegative function taking values in \mathcal{Z} . Then, for any probability measure P on \mathcal{Z} and $\rho > 0$ we have*

$$\inf_{Q \in B_\rho(P)} \mathbb{E}_Q[\ell(Z)] \geq \mathbb{E}_P[\ell(Z)] - 2C_\rho \sqrt{\mathbb{V}_P[\ell(Z)]} - \rho^2(2 - \rho^2) \left[\mathbb{E}_P[\ell(Z)] - \frac{\mathbb{V}_P[\ell(Z)]}{\mathbb{E}_P[\ell(Z)]} \right] \quad (44)$$

where $C_\rho = \sqrt{\rho^2(1 - \rho^2)^2(2 - \rho^2)}$ and $B_\rho(P) = \{Q \in \mathcal{P}(\mathcal{Z}) : H(P, Q) \leq \rho\}$ is the Hellinger ball of radius ρ centered at P . The radius ρ is required to be small enough such that

$$\rho^2 \leq 1 - \left[1 + \frac{\mathbb{E}_P[\ell(Z)]^2}{\mathbb{V}_P[\ell(Z)]} \right]^{-1/2}. \quad (45)$$

B. Finite Sampling Errors

Here we explain the reasoning behind the finite-sampling version of our main Theorem stated in Corollary 3.1. Let us first recall a version of Hoeffding's inequality, formulated in terms of our setting.

Theorem B.1 ((Hoeffding, 1963)). *Let Z_1, \dots, Z_n be independent random variables drawn from P and taking values in \mathcal{Z} . Let $\ell: \mathcal{Z} \rightarrow [0, M]$ be a loss function and let $\hat{L}_n := \frac{1}{n} \sum_{i=1}^n \ell(Z_i)$ be the mean under the empirical distribution \hat{P}_n . Then, for $\delta > 0$, with probability at least $1 - \delta$,*

$$\mathbb{E}_P[\ell(Z)] \leq \hat{L}_n + M \sqrt{\frac{\ln 1/\delta}{2n}}. \quad (46)$$

We remark that one could in principle different concentration inequalities at this stage which can potentially improve upon Hoeffding's inequality. For example, (Maurer & Pontil, 2009) present a finite sampling version of Bennett's inequality which is known to be an improvement over Hoeffding's inequality in the low variance regime. We leave such considerations for interesting future work. Recall that the certificate (10) is monotonically increasing in the variance. For this reason, we are interested in an upper bound on the population variance which can be computed from finite samples. To achieve this, we use the variance bound presented in Theorem 10 in (Maurer & Pontil, 2009) which we state here for completeness and adapt it to our use case.

Theorem B.2 ((Maurer & Pontil, 2009), Theorem 10). *Let Z_1, \dots, Z_n be independent random variables drawn from P and taking values in \mathcal{Z} . For a loss function $\ell: \mathcal{Z} \rightarrow [0, M]$, let $S_n^2 := \frac{1}{n(n-1)} \sum_{1 \leq i < j \leq n} (\ell(Z_i) - \ell(Z_j))^2$ be the unbiased estimator of the variance of the random variable $\ell(Z)$, $Z \sim P$. Then, for $\delta > 0$, with probability at least $1 - \delta$,*

$$\sqrt{\mathbb{V}_P[\ell(Z)]} \leq \sqrt{S_n^2} + M \sqrt{\frac{2 \ln 1/\delta}{n-1}} \quad (47)$$

Finally, we employ the union bound to upper bound both expectation and variance simultaneously with high probability. Thus, for any $\delta > 0$, we have with probability at least $1 - \delta$

$$\begin{aligned}\mathbb{E}_P[\ell(Z)] &\leq \hat{L}_n + M\sqrt{\frac{\ln 2/\delta}{2n}}, \\ \sqrt{\mathbb{V}_P[\ell(Z)]} &\leq \sqrt{S_n^2} + M\sqrt{\frac{2\ln 2/\delta}{n-1}}.\end{aligned}\tag{48}$$

Finally, plugging in these upper bounds for the population quantities in Theorem 2.2 leads to the desired finite sampling bound. Getting the finite sampling version of the lower bound in Theorem A.2 is analogous by using the corresponding lower bound variant of Hoeffding, but still the same upper bound for the variance.

C. Synthetic Dataset

We consider a binary classification task with covariates $X \in \mathbb{R}^2$ and labels $Y \in \pm 1$, where the data is distributed according to the Gaussian mixture

$$X|Y = y \sim \mathcal{N}(y \cdot \mu, \mathbb{1}_2).\tag{49}$$

with $p(y) = 1/2$ and $\mu = (2, 0)^T \in \mathbb{R}^2$. When considering the distribution shift $P \rightarrow Q$ arising from perturbations $X \mapsto X + \delta$ for a fixed $\delta \in \mathbb{R}^2$, both the Wasserstein distance and Hellinger distance can be evaluated as functions of the L_2 -norm of the perturbation:

$$W_2(P, Q) = \|\delta\|_2, \quad H(P, Q) = \sqrt{1 - e^{-\|\delta\|_2^2/8}}.\tag{50}$$

For our classification model, we use a small neural network with ELU activations and 2 hidden layers of size 4 and 2. The ELU activations, in combination with spectral normalization of the weights, enforce the model to be smooth and hence satisfy the assumptions required for the certificate from (Sinha et al., 2018).

D. Lipschitz Constant for Gradients of Neural Networks with Jensen-Shannon Divergence Loss

Let us first recall the dual reformulation of the Wasserstein worst-case risk, which is the central result that underpins the distributional robustness certificate presented in (Sinha et al., 2018).

Proposition D.1 ((Sinha et al., 2018), Proposition 1). *Let $\ell: \Theta \times \mathcal{Z} \rightarrow \mathbb{R}$ and $c: \Theta \times \mathcal{Z} \rightarrow \mathbb{R}_+$ be continuous, and let $\phi_\gamma(\theta; z_0) := \sup_{z \in \mathcal{Z}} \{\ell(\theta; z) - \gamma c(z, z_0)\}$. Then, for any distribution P and any $\rho > 0$,*

$$\sup_{Q: W_c(P, Q) \leq \rho} \mathbb{E}_Q[\ell(\theta; Z)] = \inf_{\gamma \geq 0} \{\gamma \rho + \mathbb{E}_P[\phi_\gamma(\theta; Z)]\}.\tag{51}$$

where $W_c(P, Q) := \inf_{\pi \in \Pi(P, Q)} \int_{\mathcal{Z}} c(z, z') d\pi(z, z')$ is the 1-Wasserstein distance between P and Q .

From this result, (Sinha et al., 2018) derive a robustness certificate which can be instantiated to hold uniformly over a function of families parametrized by $\theta \in \Theta$, but also a certificate that holds pointwise, that is, for a single model $\ell(\theta_0; \cdot)$. One requirement for this certificate to be tractable is that the surrogate function ϕ_γ be concave in z . As shown in (Sinha et al., 2018), this is the case when γ is larger than the Lipschitz constant L of the gradient of ℓ with respect to z . Thus one needs to compute L and choose $\gamma \geq L$ so that the inner maximization in (51) is guaranteed to converge and hence a robustness certificate can be calculated.

Here, we present the calculation of the Lipschitz constant for the gradient of the Jensen-Shannon divergence loss with respect to input features. For the remainder of this section, we set $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ with a binary label space $|\mathcal{Y}| = C = 2$. We will always write vectors in bold roman letters, for example $\mathbf{p} = (p_1, \dots, p_C) \in \mathbb{R}^C$ and $\mathbf{e}_y \in \mathbb{R}^C$ denotes a standard basis vector with zeros everywhere except 1 at position y . We consider a feedforward neural network with L layers and ELU activation functions, denoted by σ :

$$F_L(\theta; x) := \sigma(\theta_L \cdot \sigma_{L-1}(\theta_{L-1} \cdots \sigma(\theta_1 \cdot x) \cdots))\tag{52}$$

and we are interested in calculating $L > 0$ such that

$$\|\nabla \ell(F_L(\theta; x), y) - \nabla \ell(F_L(\theta; x'), y)\|_* \leq L \|x - x'\|_2\tag{53}$$

where the gradient is taken with respect to x and where ℓ is the Jensen-Shannon divergence. To achieve this, we apply Proposition 5 in (Sinha et al., 2018) which states that the Jacobian of F_L is $\beta_L(\theta)$ -Lipschitz with respect to the operator norm induced by $\|\cdot\|_2$ with

$$\beta_l(\theta) = \alpha_l(\theta) \sum_{j=1}^l \left\{ \frac{L_j^1}{(L_j^0)^2} \alpha_j(\theta) \right\}, \quad \alpha_l(\theta) = \prod_{j=1}^l L_j^0 \|\theta_j\|_{op}. \quad (54)$$

where L_j^0 is the Lipschitz constant of each activation function σ_j and L_j^1 is the Lipschitz constant of its Jacobian. It is useful to write this recursively as

$$\begin{aligned} \alpha_{l+1}(\theta) &= L_{l+1}^0 \|\theta_{l+1}\|_{op} \alpha_l(\theta), \\ \beta_{l+1}(\theta) &= L_{l+1}^0 \|\theta_{l+1}\|_{op} \beta_l(\theta) + L_{l+1}^1 \|\theta_{l+1}\|_{op}^2 \alpha_l(\theta)^2 \end{aligned} \quad (55)$$

In our case, since we have ELU activations, we have $L_j^0 = L_j^1 = 1$ for all j ((Sinha et al., 2018), example 3). Finally, viewing $\ell(\mathbf{p}, \mathbf{e}_y)$ as an $L + 1$ layer neural network with a single output dimension, we have that $\nabla_z \ell(\mathbf{p}(x), y)$ is L^* -Lipschitz continuous with constant

$$L^* = L_{L+1}^0 \beta_L(\theta) + L_{L+1}^1 \alpha_L(\theta)^2 \quad (56)$$

where we have used that $\|\theta_{L+1}\|_{op} = \|\mathbf{1}\|_{op} = 1$ and where L_{L+1}^0 is the Lipschitz constant of the function $z \mapsto \ell(\mathbf{p}(z), y)$ and L_{L+1}^1 is the Lipschitz constant of $z \mapsto \nabla_z \ell(\mathbf{p}(z), y)$ and $\mathbf{p}(z)$ is the softmax probability vector

$$\mathbf{p}(z) = \left(\frac{e^{z_1}}{\sum_j e^{z_j}}, \dots, \frac{e^{z_C}}{\sum_j e^{z_j}} \right)^T \in \mathbb{R}^C. \quad (57)$$

We now show the calculation of L_{L+1}^0 and L_{L+1}^1 . Fix $z \in \mathbb{R}^C$ and $y \in \mathcal{Y}$, and let $\mathbf{e}_y \in \mathbb{R}^K$ be the one hot encoded label vector with zero everywhere except at position y . The Jensen-Shannon divergence loss between a vector of predicted class probabilities \mathbf{p} and the class label \mathbf{e}_y is given by

$$\ell(\mathbf{p}, \mathbf{e}_y) = \frac{1}{2} (D_{KL}(\mathbf{p} \| \mathbf{m}) + D_{KL}(\mathbf{e}_y \| \mathbf{m})) \quad (58)$$

with $\mathbf{m} = \frac{1}{2}(\mathbf{p} + \mathbf{e}_y)$. The Kullback Leibler divergences are

$$\begin{aligned} D_{KL}(\mathbf{p} \| \mathbf{m}) &= 1 + p_y \log \left(\frac{p_y}{1 + p_y} \right) \\ D_{KL}(\mathbf{e}_y \| \mathbf{m}) &= 1 + \log \left(\frac{1}{1 + p_y} \right) \end{aligned} \quad (59)$$

where $\log = \log_2$ is the logarithm with base 2. The Jensen-Shannon divergence loss is thus given by

$$\ell(\mathbf{p}, \mathbf{e}_y) = 1 + \frac{1}{2} (p_y \log(p_y) - (1 + p_y) \log(1 + p_y)). \quad (60)$$

The gradient $\nabla_z \ell(\hat{p}, e_y)$ of the loss with respect to the input x is given by

$$\begin{aligned} \nabla_z \ell(\mathbf{p}, \mathbf{e}_y) &= \frac{1}{2} \nabla_z (p_y \log(p_y) - (1 + p_y) \log(1 + p_y)) \\ &= \frac{1}{2} \nabla_z (p_y \log(p_y)) - \frac{1}{2} \nabla_z ((1 + p_y) \log(1 + p_y)) \\ &= \frac{1}{2} (1 + \log(p_y)) \nabla_z p_y - \frac{1}{2} (1 + \log(1 + p_y)) \nabla_z p_y \end{aligned} \quad (61)$$

Noting that

$$\nabla_z p_y(x) = p_y(\mathbf{e}_y - \mathbf{p}) \quad (62)$$

yields the expression

$$\nabla_z \ell(\mathbf{p}, \mathbf{e}_y) = \frac{1}{2} \log \left(\frac{p_y}{1 + p_y} \right) p_y(\mathbf{e}_y - \mathbf{p}). \quad (63)$$

Thus,

$$\begin{aligned} L_{L+1}^0 &= \sup_z \|\nabla_z \ell(\mathbf{p}, \mathbf{e}_y)\|_2 = \frac{1}{2} \sup_z \left(\log \left(\frac{1+p_y}{p_y} \right) p_y \|\mathbf{e}_y - \mathbf{p}\|_2 \right) \\ &= \sup_z \|\nabla_z \ell(\mathbf{p}, \mathbf{e}_y)\|_2 = \frac{1}{\sqrt{2}} \sup_z \left(\log \left(\frac{1+p_y}{p_y} \right) p_y (1-p_y) \right) \approx 0.314568. \end{aligned} \quad (64)$$

We will now calculate $L_{L+1}^1 = \sup_x \|J\|_2$ where $J \equiv J_{\ell(\mathbf{p}, \mathbf{e}_y)}$, is the Jacobian of $\ell(\mathbf{p}, \mathbf{e}_y)$ and $\|J\|_2$ is given by the largest singular value of J . For ease of notation, let $f_i(x) \equiv (\nabla_z \ell)_i$ and recall that J is defined by

$$J = \begin{pmatrix} \nabla_z^T f_1 \\ \vdots \\ \nabla_z^T f_C \end{pmatrix}. \quad (65)$$

Note that

$$\begin{aligned} \nabla_z f_i &= \nabla_z \frac{1}{2} \log \left(\frac{p_y}{1+p_y} \right) p_y (\delta_{iy} - p_i) \\ &= \frac{1}{2} \left(\frac{1+p_y}{p_y} \left[\frac{\nabla_z p_y}{1+p_y} - \frac{p_y}{(1+p_y)^2} \nabla_z p_y \right] \right) p_y (\delta_{iy} - p_i) + \\ &\quad + \frac{1}{2} \log \left(\frac{p_y}{1+p_y} \right) (\delta_{iy} - p_i) \nabla_z p_y - \frac{1}{2} \log \left(\frac{p_y}{1+p_y} \right) p_y \nabla_z p_i \\ &= \frac{1}{2} \left(\frac{1}{1+p_y} + \log \left(\frac{p_y}{1+p_y} \right) \right) (\delta_{iy} - p_i) \nabla_z p_y - \frac{1}{2} \log \left(\frac{p_y}{1+p_y} \right) p_y \nabla_z p_i \end{aligned} \quad (66)$$

and hence, using $\nabla_z p_y = p_y(\mathbf{e}_y - \mathbf{p})$,

$$\nabla_z f_i = \frac{1}{2} \left(\frac{p_y}{1+p_y} + p_y \log \left(\frac{p_y}{1+p_y} \right) \right) (\delta_{iy} - p_i) (\mathbf{e}_y - \mathbf{p}) - \frac{1}{2} p_y \log \left(\frac{p_y}{1+p_y} \right) p_i (\mathbf{e}_i - \mathbf{p}) \quad (67)$$

It follows that the Jacobian is given by

$$J = \frac{1}{2} \left(\frac{p_y}{1+p_y} + p_y \log \left(\frac{p_y}{1+p_y} \right) \right) (\mathbf{e}_y - \mathbf{p}) \cdot (\mathbf{e}_y - \mathbf{p})^T + \frac{1}{2} p_y \log \left(\frac{1+p_y}{p_y} \right) (\text{diag}(\mathbf{p}) - \mathbf{p} \cdot \mathbf{p}^T). \quad (68)$$

Since we are only interested in the binary case $C = 2$, we see that

$$A := (\mathbf{e}_y - \mathbf{p}) \cdot (\mathbf{e}_y - \mathbf{p})^T = (1-p_y)^2 \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix} \quad (69)$$

with spectrum $\sigma(A) = \{0, 2(1-p_y)^2\}$. The eigenvalues of $\text{diag}(\mathbf{p})$ are p_i and hence $\lambda(\text{diag}(\mathbf{p})) \subseteq [0, 1]$, and $\sigma(\mathbf{p} \cdot \mathbf{p}^T) = \{0, \|\mathbf{p}\|_2^2\}$. It follows that $\sigma(\text{diag}(\mathbf{p}) - \mathbf{p} \cdot \mathbf{p}^T) \subseteq [-\|\mathbf{p}\|_2^2, 1]$. Thus, by Weyl's inequality and noting that the term in front of $(\mathbf{e}_y - \mathbf{p}) \cdot (\mathbf{e}_y - \mathbf{p})^T$ is always negative, we have for any eigenvalue λ of J that

$$(1-p_y)^2 \left(\frac{p_y}{1+p_y} + p_y \log \left(\frac{p_y}{1+p_y} \right) \right) - \frac{1}{2} p_y \log \left(\frac{1+p_y}{p_y} \right) \|\mathbf{p}\|_2^2 \leq \lambda \leq \frac{1}{2} p_y \log \left(\frac{1+p_y}{p_y} \right) \quad (70)$$

Note that J is symmetric, and hence its largest singular value is given by the largest absolute value of its eigenvalues. Taking the infimum (supremum) of the LHS (RHS) with respect to z yields the bounds

$$-\frac{1}{2} \leq \lambda \leq \frac{1}{2} \quad (71)$$

and hence

$$L_{L+1}^1 = \sup_z \|J\|_2 \leq \frac{1}{2}. \quad (72)$$

It follows that $\nabla_x \ell(\mathbf{p}(F_L(\theta; x)), y)$ is L^* -Lipschitz with

$$L^* = L_{L+1}^0 \beta_L(\theta) + \frac{1}{2} \alpha_L(\theta)^2 \quad (73)$$

and $L_{L+1}^0 = 0.314568$. Finally, choosing $\gamma \geq L^*$ in (51) makes the objective in the surrogate loss ϕ_γ concave and hence enables the certificate

$$\begin{aligned} \sup_{Q: W_c(P, Q) \leq \rho} \mathbb{E}_Q[\ell(\theta; Z)] &\leq \gamma\rho + \mathbb{E}_P[\phi_\gamma(\theta; Z)] \\ &= \gamma\rho + \mathbb{E}_{(X, Y) \sim P}[\sup_{x \in \mathcal{X}} \ell(F_L(\theta; x), Y) - \gamma \|x - X\|_2^2]. \end{aligned} \quad (74)$$

E. Hellinger distance for mixtures of distributions with disjoint support

Consider two joint (feature, label)-distributions $P, Q \in \mathcal{P}(\mathcal{X} \times \mathcal{Y})$ with densities f_P and f_Q with respect to a suitable measure. P and Q have disjoint support if

$$\forall x \in \mathcal{X}, y \in \mathcal{Y}: \quad f_Q(x, y) > 0 \iff f_P(x, y) = 0. \quad (75)$$

In this case, for $\gamma \in (0, 1)$, we define the mixture measure as $\Pi_\gamma := \gamma P + (1 - \gamma)Q$ with density

$$\pi_\gamma(x, y) = \gamma f_P(x, y) + (1 - \gamma)f_Q(x, y). \quad (76)$$

We can calculate the squared Hellinger distance between P and Π_γ as

$$\begin{aligned} H^2(P, \Pi_\gamma) &= 1 - \int \int_{\mathcal{X} \times \mathcal{Y}} \sqrt{f_P(x, y)} \sqrt{\gamma f_P(x, y) + (1 - \gamma)f_Q(x, y)} dx dy \\ &= 1 - \sqrt{\gamma} \int \int_{f_P > 0} f_P(x, y) \sqrt{1 + \frac{1 - \gamma}{\gamma} \frac{f_Q(x, y)}{f_P(x, y)}} dx dy \\ &= 1 - \sqrt{\gamma} \int \int_{f_P > 0} f_P(x, y) dx dy \\ &= 1 - \sqrt{\gamma}. \end{aligned} \quad (77)$$

F. Additional Experimental Results

F.1. Additional Model Architectures on CIFAR-10

Here, we present results for a diverse set of model architectures, evaluated on the CIFAR-10 dataset.

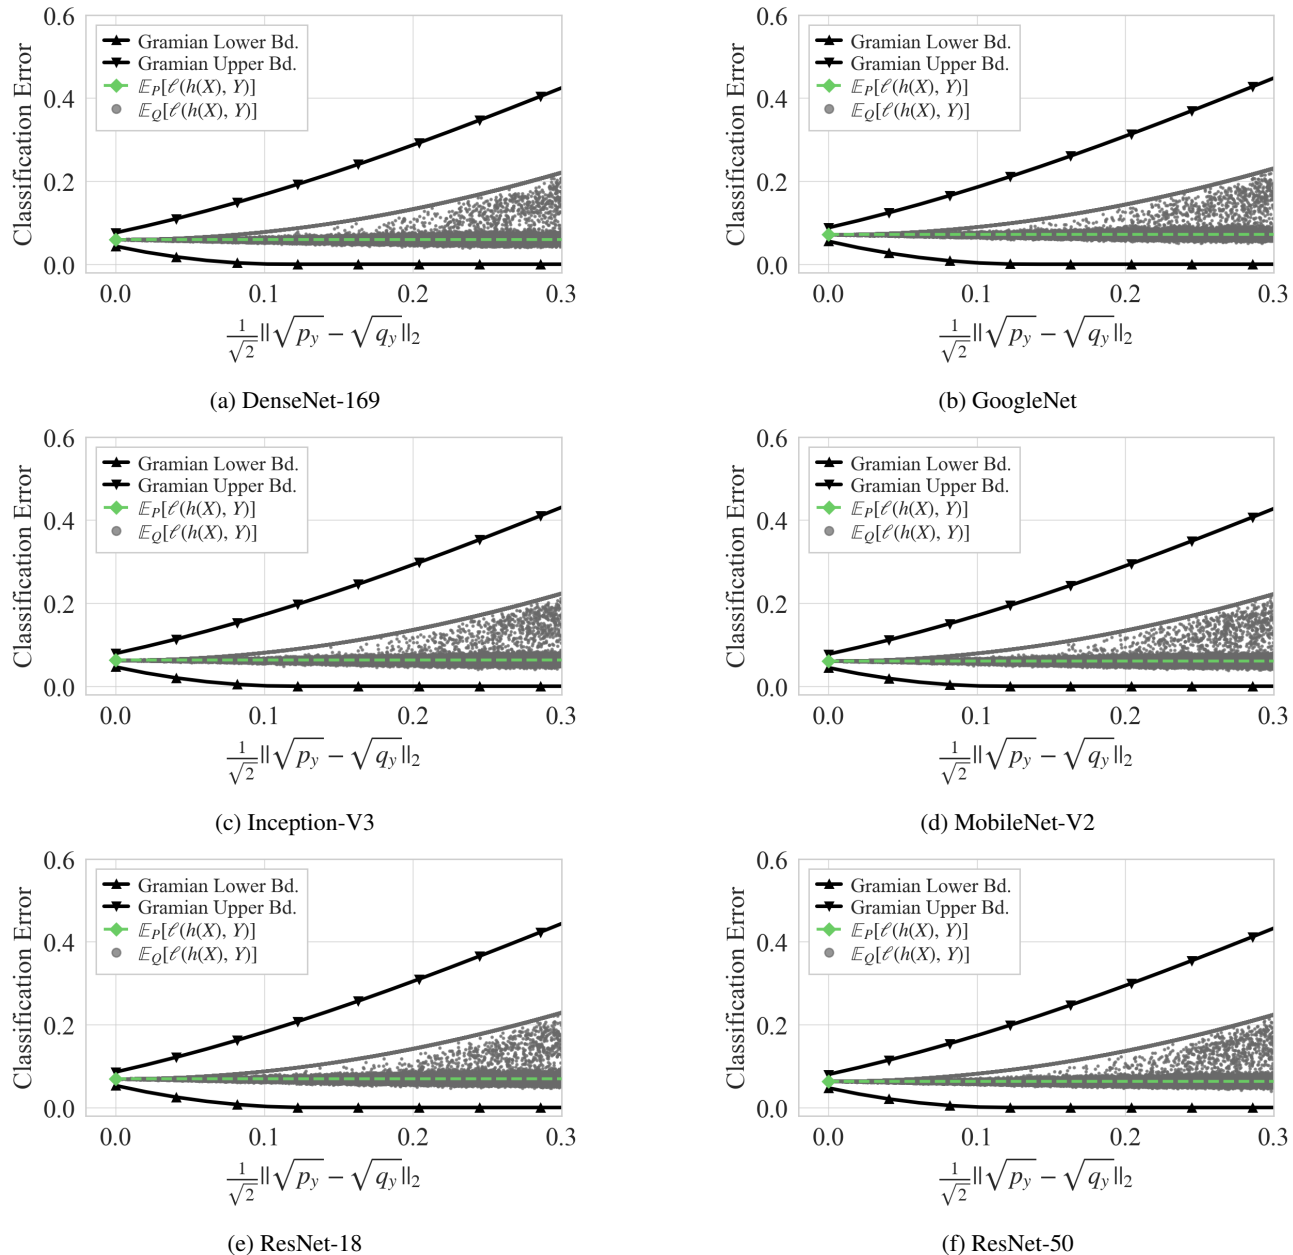


Figure 6. Certified classification error with label distribution shifts on CIFAR-10.

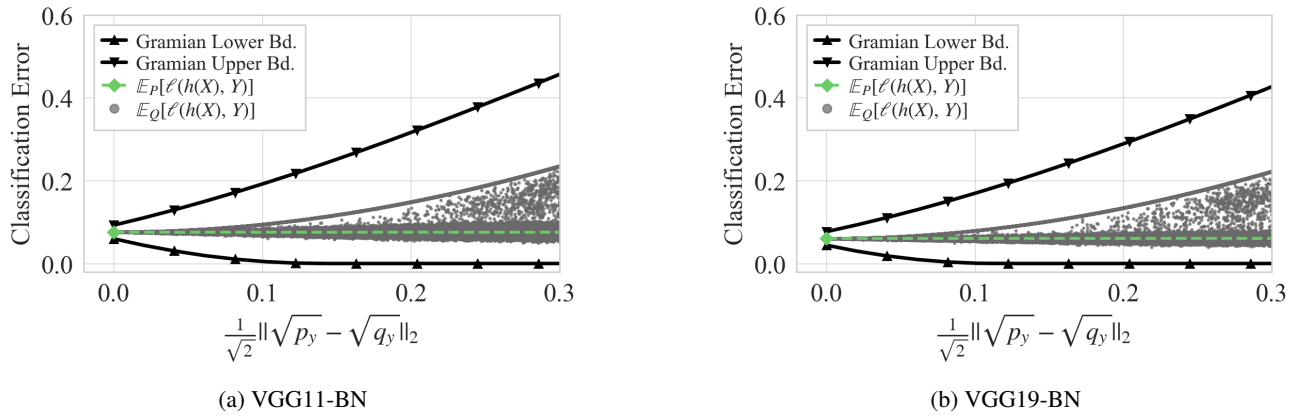


Figure 7. Certified classification error with label distribution shifts on CIFAR-10.

F.2. Results for Additional Loss and Score Functions

Here we present additional results for JSD loss, classification error and AUC score.

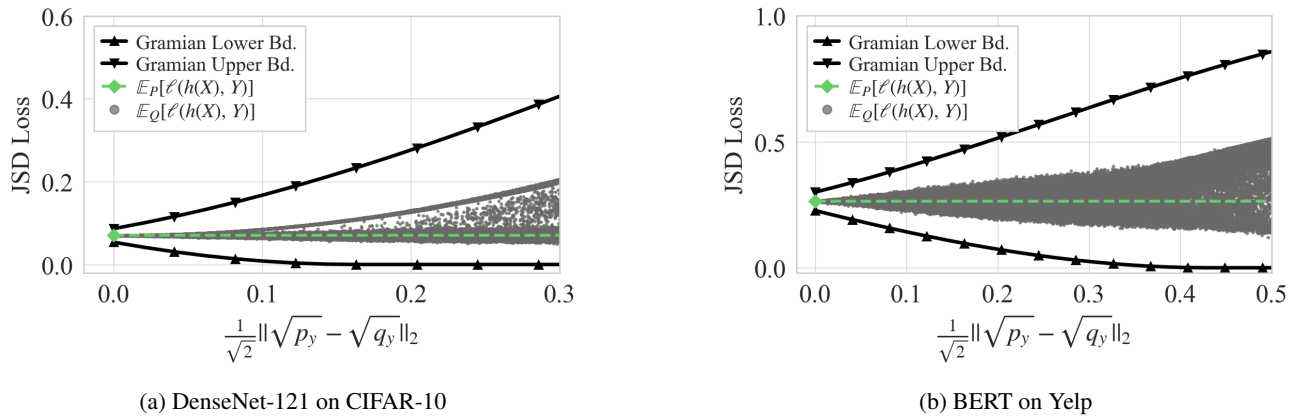




Figure 8. Certified JSD Loss with label distribution shifts.




(a) EfficientNet-B7 on ImageNet-1k




(b) DeBERTa on SNLI

Figure 9. Certified JSD Loss with label distribution shifts.



(a) EfficientNet-B7 on ImageNet-1k



(b) DeBERTa on SNLI

Figure 10. Certified classification error with label distribution shifts.




Figure 11. Certified Jensen-Shannon divergence loss for the colored MNIST dataset.