Misspecification in Prediction Problems and Robustness via Improper Learning

Annie Marsden
Stanford University
marsden@stanford.edu

John Duchi
Stanford University
jduchi@stanford.edu

Gregory Valiant
Stanford University
gvaliant@stanford.edu

Abstract

We study probabilistic prediction games when the underlying model is misspecified, investigating the consequences of predicting using an incorrect parametric model. We show that for a broad class of loss functions and parametric families of distributions, the regret of playing a "proper" predictor—one from the putative model class—relative to the best predictor in the same model class has lower bound scaling at least as $\sqrt{\gamma n}$, where γ is a measure of the model misspecification to the true distribution in terms of total variation distance. In contrast, using an aggregation-based (improper) learner, one can obtain regret $d \log n$ for any underlying generating distribution, where d is the dimension of the parameter; we exhibit instances in which this is unimprovable even over the family of all learners that may play distributions in the convex hull of the parametric family. These results suggest that simple strategies for aggregating multiple learners together should be more robust, and several experiments conform to this hypothesis.

1 Introduction

Suppose we would like to find a probability distribution that models outcomes y given data x. Typically one chooses a parametric family of probability distributions and aims to find the "best" member of this family according to a given loss. It is rarely realistic to assume that the parametric family is well-specified—that it includes the true distribution of the outcomes y and data x— and thus it is important to understand the consequences of misspecification and how to circumvent these downsides. To address these challenges, in this paper we derive a new measure of a problem's robustness to misspecification that relies on the curvature of the loss at hand and putative parametric family, proving that this measure lower bounds convergence rates for prediction error and certifies the *failure* of a parametric family and loss to be robust (or achieve optimal convergence rates for prediction). To complement this new family of lower bounds for probabilistic prediction problems, we build out of earlier work on *improper learning* (Vovk, 1998; Foster et al., 2018)—when we may choose predictions $p(y \mid x)$ outside the given model family—to show how it is possible to be robust to such misspecification, and moreover, we give new optimality guarantees for such improper procedures.

Formalizing our setting, we consider the following probabilistic game: a player receives a covariate vector $x \in \mathcal{X}$, plays a distribution $p(\cdot \mid x)$ on a target set $\mathcal{Y}$, then receives $y \in \mathcal{Y}$ and suffers loss

$$L(p(\cdot \mid x), y).$$

We study both a sequential and a stochastic variant of this problem. In the former, for a sequence of examples $\{(x_i, y_i)\}_{i=1}^n$, a player chooses a distribution p_k depending on the past examples $\{(x_i, y_i)\}_{i=1}^{k-1}$, and then for a fixed conditional distribution p on $Y \mid X$, suffers regret

$$\mathsf{Reg}_n(p) := \sum_{i=1}^n L(p_i(\cdot \mid x_i), y_i) - \sum_{i=1}^n L(p(\cdot \mid x_i), y_i).$$

In the stochastic variant, the examples (x_i, y_i) are i.i.d. from an unknown distribution P, and we consider the risk of the conditional p.m.f. p,

$$\mathsf{Risk}_P(p) := \mathbb{E}[L(p(Y \mid X))],$$

where the expectation is taken over $(X, Y) \sim P$. The goal is to play p_i or p above to make the regret and risk as small as possible.

This regret and risk minimization framework is familiar from the universal prediction and probabilistic forecasting literature (Merhav and Feder, 1998; Grünwald, 2007;

Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS) 2021, San Diego, California, USA. PMLR: Volume 130.

Gneiting and Raftery, 2007; Cover and Thomas, 2006; Cesa-Bianchi and Lugosi, 2006), which considers best possible estimators and online learners for the regret Reg_n over various losses L and families $\mathcal{P}$ of possible predictive distributions. In this paper, we study these regret and risk minimization formulations over parametric families of distributions $\{p_\theta(\cdot \mid x)\}_{\theta \in \Theta}$, where $\Theta \subset \mathbb{R}^d$ is a convex set. We shall either consider the regret

$$\mathsf{Reg}_n^\Theta := \sup_{\theta^\star \in \Theta} \mathsf{Reg}_n(p_{\theta^\star}) \tag{1a}$$

or—in the stochastic version of the problem—the excess risk relative to this family,

$$\mathsf{Risk}_P^\Theta(p) := \mathsf{Risk}_P(p) - \inf_{\theta^\star \in \Theta} \mathsf{Risk}_P(p_{\theta^\star}). \tag{1b}$$

When the conditional distribution of $Y \mid X$ belongs to the parametric family $\{p_\theta\}_{\theta \in \Theta}$ where $\Theta \subset \mathbb{R}^d$, maximum likelihood estimators enjoy rates of convergence of $O(d/n)$ for the excess risk (1b) as n grows (van der Vaart, 1998). In typical practice, however, the data generating distribution is misspecified, so it is important to understand how this misspecification impacts possible convergence rates and optimal estimators.

We thus consider three intertwined objects: the parametric family $\{p_\theta\}_{\theta \in \Theta}$ against which we compare the performance of our prediction p, a family Γ of distributions on Y given X that we may return (i.e. predict from), and the family $\mathcal{P}$ of data generating distributions that nature may choose. We study the interaction between these three and the impact of allowing the family $\mathcal{P}$ to differ from the parametric model $\{p_\theta\}_{\theta \in \Theta}$. The traditional approach considers the minimax excess risk over the family Θ,

$$\inf_{\widehat{p}_n} \sup_{\theta \in \Theta} \mathbb{E}_{P_\theta^n} \left[\mathsf{Risk}_P^\Theta(\widehat{p}_n) \right], \tag{2}$$

where the infimum is over all estimators $\widehat{p}_n$ that use the n points $\{(X_i, Y_i)\}_{i=1}^n$ to output a distribution $p(Y \mid X)$, and the expectation is taken over $\{(X_i, Y_i)\}_{i=1}^n \stackrel{\text{iid}}{\sim} P_\theta$, where we have abused notation to use P_θ to denote the joint over (X, Y) when $Y \mid X = x$ follows $p_\theta(\cdot \mid x)$. We elaborate this setting slightly. First, we restrict the estimator $\widehat{p}_n$ to take values in a set Γ of distributions (for example, we might take $\Gamma = \{p_\theta\}_{\theta \in \Theta}$, the parametric family, or its convex hull), which we write as $\widehat{p}_n \in \Gamma$. Second, we expand the supremum (2) to also include distributions P *near* the model P_θ: recalling the definition of the total-variation distance $\|P - Q\|_{\text{TV}} := \sup_A |P(A) - Q(A)|$, we consider distributions P for which there is some $\theta \in \Theta$ such that $\|P - P_\theta\|_{\text{TV}} \le \gamma$. This gives us our misspecified minimax risk.

Definition 1.1. *Let $\Theta \subset \mathbb{R}^d$, $\gamma \ge 0$, and Γ be a set of allowable distributions $p(Y \mid X)$. The* minimax risk *at variation distance γ is*

$$\mathfrak{M}_n(\Theta, \Gamma, \gamma) := \inf_{\widehat{p}_n \in \Gamma} \sup_{\substack{\theta \in \Theta \\ P : \|P - P_\theta\|_{\text{TV}} \le \gamma}} \mathbb{E}_{P^n}[\mathsf{Risk}_P^\Theta(\widehat{p}_n)]. \tag{3}$$

The quantity (3) is somewhat complex. The idea is to quantify—via the parameter γ—the impact of allowing the data generating distribution P to depart slightly from the parametric family $\{p_\theta\}_{\theta \in \Theta}$ while constraining ourselves to play a prediction from the family Γ.

The typical setting in online convex optimization and learning (Cesa-Bianchi and Lugosi, 2006) is to take the family of "playable" distributions to be the parametric family $\Gamma = \{p_\theta\}_{\theta \in \Theta}$. In this case, standard minimax risk bounds show that in the well-specified setting that the data comes from the parametric family (i.e. $\gamma = 0$ in Def. 1.1) and the loss L is smooth, then we expect the risk to scale as d/n (cf. van der Vaart (1998); Wainwright (2019); Bach (2014)). Yet as we show in the first part of this work, such results need not be stable to perturbations away from the parametric model. We show that the curvature of the loss relative to predictions and the parameter space Θ essentially governs convergence rates: when losses are appropriately "flat," there is little information and rates are necessarilty slow and misspecification carries a potentially heavy penalty; conversely, when there is substantial curvature, rates exhibit less antagonistic behavior. Accordingly, we introduce what we term the *linearity constant* Lin of the loss L, family $\{p_\theta\}_{\theta \in \Theta}$, and misspecification γ, showing a lower bound of roughly $\min\{1/\sqrt{n}, \text{Lin}/n\}$ on the minimax risk (3). In some cases we delineate, Lin may be exponentially large in problem parameters, so convergence rates slow to the worst-case rates for general online convex optimization (Zinkevich, 2003; Shalev-Shwartz et al., 2009; Agarwal et al., 2012), and we consider the family sensitive to misspecification.

To complement these negative results, we highlight a solution to this instability by considering the convex hull of the parameteric family, that is the set of mixtures, aggregations, or ensembles $\text{Conv}\{p_\theta\}_{\theta \in \Theta} := \{\int_{\theta \in \Theta} p_\theta d\mu(\theta) \text{ s.t. } \int_{\theta \in \Theta} d\mu(\theta) = 1 \text{ and } d\mu \ge 0\}$. The idea to combine probabilistic forecasts is classical (Granger and Ramanathan, 1984; Genest and Zidek, 1986; Jacobs, 1995; Clemen and Winkler, 1999; Hall and Mitchell, 2007). When the loss function is *mixable* (which we define later), Vovk's Aggregating Algorithm and its variants, e.g. exponential weights, Exp3, and Bayesian universal prediction (Vovk, 1990, 1992, 1998; Grünwald, 2007; Cesa-Bianchi and Lugosi, 2006; Auer et al., 2002), provide stability and achieve minimax regret $O(d \log n)$ for any γ in Definition 1.1. By a standard online-to-batch conversion (Jensen's inequality) (Cesa-Bianchi et al., 2002), this guarantees a

minimax excess risk (3) of at most $O(d \log n/n)$. We also show that for generalized linear models, Vovk's Aggregating Algorithm is optimal up to log-factors with respect to the misspecification parameter γ. That is, there is no better algorithm to use if you are guaranteed certain values of γ.

It is perhaps unfair to consider the entire convex hull of $\{p_\theta\}_{\theta\in\Theta}$ for the class Γ, as this could potentially yield much smaller risk than the parametric family in the risk (1b). Indeed, we give an example in which the best parameteric predictor has no predictive power, while returning a mixture of two parameterized distributions achieves zero loss (though we also give examples in parametric families where considering the convex hull provides no benefit). To justify the aggregation strategy we give a Bernstein von-Mises theorem under misspecification, which shows that the strategy converges to a Gaussian centered at the risk-minimizing parameter estimate $\widehat{\theta}_n$ with covariance shrinking at rate $O(1/n)$; a corollary of this is that Vovk's Aggregating Algorithm returns a distribution which converges in total variation distance to $p_{\widehat{\theta}_n} \in \{p_\theta\}_{\theta\in\Theta}$. Thus, aggregation (or exponential weights) stabilizes predictions while asymptotically enjoying identical convergence to standard risk-minimization procedures.

1.1 Related Work

Our results broadly fall under probabilistic universal prediction in which the data (x_t, y_t) can be any arbitrary sequence (Rissanen, 1984; Merhav and Feder, 1998; Cesa-Bianchi and Lugosi, 1999, 2006; Grünwald, 2007; Cover, 1991). That Vovk's Aggregating Algorithm provides minimax rate stability is known Foster et al. (2018), and this is similar to the minimax guarantees of Bayesian models in universal prediction (Grünwald, 2007); a long line of work gives the same logarithmic minimax rates (Merhav and Feder, 1998; DeSantis et al., 1988; Haussler and Barron, 1992; Yamanishi, 1995). Early work in these prediction problems focuses on the logarithmic loss $L_{\log}(p(\cdot), y) = -\log p(y)$, while more recent work extends these bounds to exp-concave and so-called "mixable" losses (Hazan et al., 2007; Cesa-Bianchi and Lugosi, 2006). Our results on minimax lower bounds, distinguishing carefully between well-specified and misspecified models and proper and improper predictions, are novel.

While our results are general, applying to exponential families and beyond, related results are available for logistic regression. In this case, for $B, R > 0$ we consider $\Theta = \{\theta : \|\theta\| \le B\}$, $\mathcal{X} \subset \{x : \|x\| \le R\}$, and let $\{p_\theta\}_{\theta\in\Theta}$ be the family of binary logistic distributions, $p_\theta(y \mid x) = (1 + e^{-y\theta^T x})^{-1}$, with log loss. Hazan et al. (2014) show that any algorithm returning some p_θ suffers minimax risk (recall (3)) $\Omega(\sqrt{B/n})$ in the regime where $n = O(\exp cB)$ for some positive constant $c > 0$, $R = 1$, and the allowable perturbation $\gamma = 1$. Foster et al. (2018) show that Vovk's Aggregating Algorithm (Vovk, 1998) guarantees minimax risk $O(d \log(Bn)/n)$, allowing one to sidestep this lower bound via improper learning, which we also leverage. In the special case of logistic regression—see Example 2 to come in Section 2.2—a simplification of our results gives lower bound $\Omega(1)\sqrt{\gamma BR/n}$ if $n \le \exp(RB/2)$ and $\Omega(1)\exp(2BR/5)/n$ otherwise. We thus show that even when the perturbations away from the parametric family are small, the minimax risk when the set of playable distributions is $\Gamma = \{p_\theta\}_{\theta\in\Theta}$ may grow substantially; this generalizes Hazan et al. (2014), where $R = 1$ and $\gamma = 1$, and gives somewhat sharper constants.

2 Parametric Model Instability

Our first step towards understanding sensitivity to misspecification is to provide optimality guarantees for the minimax risk in Definition 1.1 when the player can play only elements of the parametric family of interest, that is, when $\Gamma = \{p_\theta\}_{\theta\in\Theta}$. We focus on losses that depend specifically on $\theta^T x$, where we have

$$L(p_\theta(\cdot \mid x), y) = \ell(\theta^T x, y) \tag{4}$$

for some twice differentiable and convex $\ell : \mathbb{R}\times\mathcal{Y} \to \mathbb{R}_+$. A broad range of models and losses take this form, including all generalized linear models (Hastie et al., 2009).

2.1 Main lower bound

Our key contribution is to lower bound the minimax risk via a quantity we term the *linearity constant* of the induced loss ℓ, which measures the sensitivity of ℓ around various points in its domain. The first component is (roughly) a measure of the signal contained in ℓ for different targets y, where for $t \in \mathbb{R}$, $y \in \mathcal{Y}$, and $\ell'(t_0, y_0)$ as shorthand for $\frac{\partial}{\partial t}\ell(t, y)|_{(t,y)=(t_0,y_0)}$ we define

$$q^\star_\ell(t, y) := \sup_{\substack{\tilde{y}\in\mathcal{Y}, \alpha\in[-1,1] \\ (\tilde{y},\alpha)\in \mathsf{C}_y(\ell,t)}} \left\{ \frac{\alpha\ell'(\alpha t, \tilde{y})}{\alpha\ell'(\alpha t, \tilde{y}) - \ell'(t, y)} \right\}, \tag{5}$$

where $\mathsf{C}_y(\ell, t) := \{(\tilde{y}, \alpha) \mid \text{sign}(\alpha\ell'(\alpha t, \tilde{y})) \ne \text{sign}(\ell'(t, y))\}$.
We always have $q^\star_\ell(t, y) \in [0, 1]$. For many cases, this quantity is a positive numerical constant (e.g. for the squared error $\ell(t, y) = \frac{1}{2}(t - y)^2$ with $\mathcal{Y} = \mathbb{R}$, we have $q^\star_\ell(t, y) = 1$). For a given misspecification size $\gamma \in [0, 1]$, sample size n, radii R and B, we define our measure

of sensitivity to misspecification,

$$\mathsf{Lin}(\ell, \mathcal{Y}, R, B, n, \gamma) := \sup_{y \in \mathcal{Y}} \sup_{t^2 + \delta^2 \le \frac{R^2 B^2}{2}} \left\{ |\ell'(t, y)| \right. \\ \left. \times \min \left\{ \delta \sqrt{\gamma q_\ell^\star(t, y)}, \frac{|\ell'(t, y)|}{\sup_{|\Delta| \le \gamma} \ell''(t + \Delta, y)} \frac{q_\ell^\star(t, y)}{\sqrt{2n}} \right\} \right\}. \tag{6}$$

This linearity constant roughly measures the extent to which the loss grows quickly without substantial curvature, that is, $\ell'(t, y)$ is large while $\ell''(t, y)$ is small. A heuristic simplification may help with intuition: by ignoring the $q_\ell^\star$ term and perturbation by Δ in ℓ'', we roughly have

$$\mathsf{Lin} \approx O(1) \sup_{\substack{y \in \mathcal{Y} \\ |t| \le RB}} |\ell'(t, y)| \min \left\{ RB\sqrt{\gamma}, \frac{|\ell'(t, y)|}{\ell''(t, y)} \frac{1}{\sqrt{n}} \right\}, \tag{7}$$

which makes clearer the various relationships. When the ratio of $\ell'(t, y)$ to $\ell''(t, y)$ is large, estimation and optimization are intuitively hard: there is little curvature to help identify optimal parameters, while small changes in the parameter induce large changes in the loss (as $\ell'(t, y)$ is large relative to ℓ''). The allowable misspecification of the model—via the parameter γ—means that in the lower bound, an adversary may essentially put positive mass on those points for which the ratio ℓ'/ℓ'' is large, so that one must pay this worst-case cost.

We then have the following theorem, whose proof we provide in Appendix A.

Theorem 1. *Let the loss L and family $\{p_\theta\}$ satisfy Eq. (4), where $\mathcal{X} = \{x : \|x\| \le R\}$. Consider $\Gamma = \{p_\theta\}_{\theta \in \Theta}$, where $\Theta = \{\theta : \|\theta\| \le B\}$. Then*

$$\mathfrak{M}_n(\Theta, \Gamma, \gamma) \ge \frac{1}{4\sqrt{n}} \mathsf{Lin}(\ell, \mathcal{Y}, R, B, n, \gamma).$$

Using the heuristic display (7) above can provide some intuition. When $|\ell'(t, y)|/\ell''(t, y) \gtrsim \sqrt{n}$, so that the problem has little curvature, the (heuristic) linearity constant (7) scales as $\sup_{t,y} |\ell'(t, y)| RB$, which gives the lower bound $\sup_{t,y} \frac{|\ell'(t,y)|RB}{\sqrt{n}}$; this is the familiar worst-case minimax bound for stochastic convex optimization with Lipschitz objective on a compact domain (Agarwal et al., 2012). As the worst-case constructions look very little like standard prediction problems, one might hope (at least in the absence of misspecification) to achieve better rates; Theorem 1 helps to delineate problems where this may be impossible.

2.2 Examples with the logarithmic loss

Here we consider a few examples to build intuition for Theorem 1 beyond the heuristic (7). As we shall see, its generality allows exploration of many losses, including various scoring rules (Gneiting and Raftery, 2007); for this section, we focus on the common logarithmic loss for three well-known exponential family models. We will use the following notation: For a set Ω such that $f, g : \Omega \to \mathbb{R}$ we write $f \gtrsim g$ if there exists a finite numerical constant C such that for any $\omega \in \Omega$, $f(\omega) \ge Cg(\omega)$ and we write $f \asymp g$ if $f \gtrsim g \gtrsim f$.

Example 1 (Linear regression)**:** For linear regression we consider the model $y \mid x \sim \mathsf{N}(\theta^T x, 1)$ with log loss, so that the loss becomes $L_{\log}(p_\theta(\cdot \mid x), y) = \frac{1}{2}(\theta^T x - y)^2$ and $\ell(t, y) = \frac{1}{2}(t - y)^2$. In this case, we may take $q_\ell^\star \gtrsim 1$ in Eq. (5). Letting $\mathrm{d}(\mathcal{Y})$ be shorthand for $\mathrm{diam}(\mathcal{Y})$, the linearity constant (6) becomes

$$\begin{aligned} &\mathsf{Lin}(\ell, \mathcal{Y}, R, B, n, \gamma) \\ &\asymp \sup_{y \in \mathcal{Y}} \sup_{t^2 + \delta^2 \le \frac{R^2 B^2}{2}} \left\{ (t - y) \min \left\{ \delta\sqrt{\gamma}, \frac{t - y}{\sqrt{n}} \right\} \right\} \\ &\asymp \min \left\{ RB \max\{RB, \mathrm{d}(\mathcal{Y})\} \sqrt{\gamma}, \frac{\max\{(RB)^2, \mathrm{d}(\mathcal{Y})^2\}}{\sqrt{n}} \right\}, \end{aligned}$$

yielding minimax lower bound,

$$\mathfrak{M}_n(\Theta, \Gamma, \gamma) \ge c \min \left\{ \frac{RB \max\{RB, \mathrm{d}(\mathcal{Y})\} \sqrt{\gamma}}{\sqrt{n}}, \right. \\ \left. \frac{\max\{(RB)^2, \mathrm{d}(\mathcal{Y})^2\}}{n} \right\}.$$

This is sharp: the stochastic gradient method achieves the bound (Hazan and Kale, 2011), as ℓ is strongly convex and has Lipschitz constant $\max\{BR, \mathrm{diam}(\mathcal{Y})\}$.

With that said, this behavior—which depends on $\mathrm{diam}(\mathcal{Y})$—is worse than what one achieves in well-specified or stochastic settings, where the stochasticity means that rates of $1/n$ are achievable. $\Diamond$

Example 2 (Logistic regression)**:** For our second example, we consider logistic regression, showing that if we must play proper predictions p_θ, parametric $1/n$ rates are impossible until n is very large or if the radii R and B are small. For logistic regression with logarithmic loss, we have $\mathcal{Y} = \{-1, 1\}$, $p_\theta(y \mid x) = \frac{1}{1 + \exp(-y\theta^T x)}$, and $\ell(t, y) = \log(1 + e^{-ty})$. Appendix B.1 works through the derivation of $q_\ell^\star(t, y)$ and $\mathsf{Lin}(\ell, \mathcal{Y}, R, B, n, \gamma)$ which results in minimax risk lower bound,

$$\mathfrak{M}_n(\Theta, \Gamma, \gamma) \ge c \min \left\{ \frac{\sqrt{\gamma RB}}{\sqrt{n}}, \frac{\exp(2RB/5)}{n} \right\}, \tag{8}$$

when $RB \ge 1$. Notice that the expression allows for the regime where $RB \gg 1$—for example, in the natural case that the data and parameter radii scale with the dimension of the problem. We may contrast this lower bound with previous results. In the regime where $\gamma = 1$,

Hazan et al. (2014) show that for $R = 1$ and numerical constants $c_0, c_1 > 0$, any algorithm playing parametric predictors p_θ necessarily suffers minimax risk $\Omega(\sqrt{B/n})$ whenever $n \le c_0 \exp(c_1 B)$. The result (16) recovers this lower bound while applying whenever $\gamma > 0$. $\diamondsuit$

Example 3 (Geometric distributions)**:** We say $Y \sim \mathsf{Geo}(\lambda)$ for some $\lambda \in (0,1)$ if Y has support $\{0, 1, 2, \ldots, \}$ and $P(Y = y) = \lambda(1-\lambda)^y$. We model this via $Y \mid x \sim \mathsf{Geo}(e^{\theta^T x}/(1 + e^{\theta^T x}))$, giving losses $L_{\log}(p_\theta(\cdot \mid x), y) = (y+1)\log(1+\exp(\theta^T x)) - \theta^T x$ and $\ell(t, y) = (y+1)\log(1+e^t) - t$. In Appendix B.2 we obtain the analogue of inequality (16), that is,

$$\mathfrak{M}_n(\Theta, \Gamma, \gamma) \ge c_0 \mathrm{diam}(\mathcal{Y}) \min\left\{\frac{\sqrt{\gamma RB}}{\sqrt{n}}, \frac{e^{c_1 RB}}{n}\right\}.$$

Again, we see that until $n \gtrsim \exp(cRB)$, any method playing the models p_θ for points $\theta \in \Theta$ necessarily cannot converge faster than $\mathrm{diam}(\mathcal{Y})/\sqrt{n}$. $\diamondsuit$

2.3 Scoring rules and general losses

In probabilistic prediction and forecasting, one more generally may consider *scoring rules* (Gneiting and Raftery, 2007), which are losses designed to engender various behaviors: honesty in eliciting predictions, calibration of forecasts, robustness, or other reasons. Typically, these induced losses are exp-concave (as we discuss in the next section, which will allow us to describe an efficient algorithm for them). For example, to achieve robustness (Mei et al., 2018; Huber and Ronchetti, 2009) (there are no unbounded losses) one might consider the squared error or Hellinger-type losses

$$L_{\mathrm{sq}}(p(\cdot \mid x), y) := \frac{1}{2}(p(y \mid x) - 1)^2 \tag{9}$$

and

$$L_{\mathrm{hel}}(p(\cdot \mid x), y) = (\sqrt{p(y \mid x)} - 1)^2, \tag{10}$$

neither of which is proper (so that the true distribution may not minimize the loss). Alternatively, proper scoring rules (Gneiting and Raftery, 2007) are minimized by the true predictive distribution, and include the logarithmic loss and the quadratic scoring rule with loss

$$L_{\mathrm{quad}}(p(\cdot \mid x), y) := \frac{1}{2}\sum_{k \in \mathcal{Y}} (p(k \mid x) - \mathbf{1}\{k = y\})^2. \tag{11}$$

As these scoring rules are differentiable and at least $\mathcal{C}^2$ on $p \in (0,1)$, an argument by the delta method (van der Vaart, 1998) shows that in well-specified cases, one expects to achieve convergence at rate $1/n$. Moreover, as we will see in the next section, aggregating algorithms can achieve regret scaling as $\log n$ for each of these loss measures.

Yet the minimax bound in Theorem 1 shows that this is unachievable with misspecification. Here, because of the complexity of the losses and resulting calculations, we take a completely asymptotic perspective, saying that we have an *asymptotic rate* $\mathsf{r}(n)$ *minimax lower bound* if $\mathsf{r}(n) \to 0$ as $n \to \infty$, while

$$\liminf_{n \to \infty} \frac{\mathfrak{M}_n(\Theta, \Gamma, \gamma)}{\mathsf{r}(n)} > 0$$

for all $\gamma > 0$. As usual we consider losses taking the form $L(p_\theta(\cdot \mid x), y) = \ell(\theta^T x, y)$ for some scalar induced loss $\ell : \mathbb{R} \times \mathcal{Y} \to \mathbb{R}$. We restrict our focus to losses for which no *universally perfect* prediction exists, meaning that if $t \in \mathbb{R}$ and $y \in \mathcal{Y}$ satisfy $\ell'(t, y) \ne 0$, there exists $y_0 \in \mathcal{Y}$ such that $\ell'(t, y)\ell'(t, y_0) < 0$. We have the following result, whose proof we provide in Appendix A.4.

Proposition 2. *Assume that the scalar loss allows no universally perfect prediction. If there exists $|t| \le RB/2$ and $y \in \mathcal{Y}$ such that $\ell'(t, y) \ne 0$ while $\ell''(t, y) = 0$ and $\ell(\cdot, y)$ is $\mathcal{C}^3$ near t, then the prediction family $\Gamma = \{p_\theta\}_{\theta \in \Theta}$ has asymptotic rate $\mathsf{r}(n) = n^{-3/4}$ minimax lower bound.*

Roughly, the result in the proposition is simple: if the induced scalar loss ℓ is not convex in t, then proper predictions cannot be rate-optimal (as rates scaling as $1/n$ are achievable here).

While it is possible in some cases to achieve a non-asymptotic result, in general it is somewhat tedious. Nonetheless, we have the following result, whose tediousness in verification precludes our including a formal proof, but essentially we need simply note that the induced losses $\ell(t, y)$ for each problem are non-convex in t but are smooth.

Corollary 3. *Let $\Gamma = \{P_\theta\}$ be any of the logistic-, geometric-, poisson-, or linear-regression families of predictive densities. Then for any of the squared L_{sq}, Hellinger L_{hel}, or quadratic L_{quad} losses (Eqs. (9) and (11)), $\mathfrak{M}_n(\Theta, \Gamma, \gamma)$ has asymptotic rate lower bound $\mathsf{r}(n) = \frac{1}{n^{3/4}}$ for any $\gamma > 0$.*

Theorem 1 and its consequences via Proposition 2 assert that *any* algorithm returning elements of the parametric family $\{p_\theta\}_{\theta \in \Theta}$ must suffer when misspecification is possible. These results are information-theoretic, and as such, we see that in situations where misspecification is possible, to perform better it is essential that the family Γ of allowable distributions be improper.

3 Robustness via Improper Learning

While proper algorithms evidently must suffer losses when they are misspecified, we now show that simply by considering $\Gamma = \mathrm{Conv}\{p_\theta\}_{\theta \in \Theta}$ we can sidestep the lower bound of Theorem 1. Specifically, we consider Vovk's

Aggregating Algorithm and show that this provides stability to misspecification. We present the algorithm in an online setting, though standard online-to-batch conversion techniques Cesa-Bianchi and Lugosi (2006) extend the result to the stochastic optimization setting in which we have proved each of our lower bounds.

Algorithm 1 Vovk's Aggregating Algorithm (Online Setting)

Let $\Gamma = \mathrm{Conv}\{p_\theta\}_{\theta\in\Theta}$ and $\eta > 0$ be some fixed value.
For $t = 0, \ldots, n$
  Nature reveals x_t.
  Define $d\widehat{\mu}^{\mathsf{Vovk}}_{t,\eta}(\theta) \propto e^{-\eta \sum_{s=0}^{t-1} L(P_\theta(y_s|x_s))}$.
  Decision Maker plays
    $\widehat{P}^{\mathsf{Vovk}}_{t,\eta} := \int_{\theta\in\Theta} P_\theta(\cdot \mid x_t) d\widehat{\mu}^{\mathsf{Vovk}}_{t,\eta}(\theta)$.
  Nature reveals y_t and Decision Maker suffers loss $L(\widehat{P}^{\mathsf{Vovk}}_{t,\eta}(y_t \mid x_t))$.

	$L_{\log}$	L_{sq}	L_{hel}	L_{quad}
Mixability Constant η	1	1	3	1/2

Table 1. Mixability constants for common losses in Eqs. (11) and Eqs. (9), where $L_{\log}(p, y) = -\log p(y)$, $L_{\mathrm{sq}}(p, y) = \frac{1}{2}(p(y) - 1)^2$, $L_{\mathrm{hel}}(p, y) = (\sqrt{p(y)} - 1)^2$, and $L_{\mathrm{quad}}(p, y) = \frac{1}{2}\|p(\cdot) - e_y\|_2^2$. See Appendix B.3.

To give a regret bound, we continue our usual focus by considering losses L and families $\{p_\theta\}_{\theta\in\Theta}$ for which we can write $L(p_\theta(\cdot \mid x), y) = \ell(\theta^T x, y)$. We restrict ourselves to considering *mixable* losses, where for some $\eta > 0$, the function $p \mapsto \exp(-\eta L(p, y))$ is concave over the collection $\mathcal{P}(\mathcal{Y})$ of distributions on $\mathcal{Y}$. This constant η guaranteeing the exp-concavity of L bounds the *mixability* constant, which allows one to obtain "fast rates" via exponentially-weighted averaging in many online learning problems Vovk (1998); Cesa-Bianchi and Lugosi (2006); Van Erven et al. (2015). In Table 1, we record the mixability constants for several example losses—each of which we touch on in Sec. 2.3—showing that Vovk's aggregating algorithm achieves logarithmic regret for any of them once we apply the coming convergence result.

Corollary 4 (Foster et al. (2018), Theorem 1)**.** *Let $\widehat{P}^{\mathsf{Vovk}}_{t,\eta}$ be as defined above. Let $\mathrm{Lip}_\ell(T)$ be the Lipschitz constant of ℓ restricted to $[-T, T] \times \mathcal{Y}$. Then for any sequence $(x_i, y_i)_{i=1}^n$ and $\theta^\star \in \Theta$,*

$$\mathsf{Reg}_n(\widehat{P}^{\mathsf{Vovk}}_{n,\eta}, p_{\theta^\star}) \le 5\frac{d}{\eta}\log\left(\frac{\mathrm{Lip}_\ell(RB)n}{d} + e\right).$$

Using Corollary 4 and Theorem 1, we see that the aggregating algorithm typically provides a stronger convergence guarantee than algorithms constrained to $\{p_\theta\}_{\theta\in\Theta}$ can attain. In each of Examples 1–2, the lower bounds necessarily suffer exponential dependence $\exp(\Omega(1)RB)$ as $n \to \infty$, and so as long as the Lipschitz constant $\mathrm{Lip}_\ell(RB)$ is not super-exponential in RB, Corollary 4 guarantees better convergence. Even more, in some cases—for example, when using the general (potentially non-convex in θ) losses as in Sec. 2.3—even for fixed radii R, B we have $\sqrt{n}\mathsf{Lin}(\ell, \mathcal{Y}, R, B, n, \gamma) \to \infty$ as $n \to \infty$. In this case, Corollary 4 even guarantees a better asymptotic rate in n.

3.1 Lower bounds for arbitrary improper algorithms

An important question is whether the Aggregating Algorithm 1 achieves optimal rates when the misspecification parameter γ changes. As the coming Theorem 5 shows, Corollary 4 is tight to within logarithmic factors, as we can show an (asymptotic) lower bound of d/n even for well-specified families of generalized linear models. The theorem also shows that while playing in the convex hull of a parameterized family $\{p_\theta\}_{\theta\in\Theta}$ allows more powerful mechanisms, in natural (well-specified) scenarios, this extra power is no panacea: the upper bound on the risk of the Aggregating Algorithm 1 is tight to a logarithmic factor over all algorithms that may play arbitrary elements in the convex hull (or even algorithms playing any probability distribution). Indeed, let Γ_{all} be the collection of *all* probability distributions on $Y \mid X$, so that an algorithm may play an arbitrary distribution (which is of course larger than $\mathrm{Conv}\{p_\theta\}_{\theta\in\Theta}$). We then take the risk to be the expected logarithmic loss, where for a conditional distribution $p(y \mid x)$ we define

$$\mathsf{Risk}_P(p) = -\mathbb{E}_P[\log p(Y \mid X)],$$

and we let $\mathsf{Risk}^\star_P = \inf_p \mathsf{Risk}_P(p)$ be the smallest risk across all predictive distributions $p(y \mid x)$. When the models p_θ are not misspecified, i.e., $\gamma = 0$, we have $\mathsf{Risk}^\star_{P_\theta} = \mathsf{Risk}_{P_\theta}(p_\theta)$, and we have the following lower bound.

Theorem 5. *Let $\{p_\theta\}_{\theta\in\Theta}$ be a generalized linear model of the form*

$$dP_\theta(y \mid x) = \exp(y\theta^T x - A(\theta^T x))d\nu(y),$$

where ν is a base measure and $X \sim \mathsf{Uni}(\{-1, 1\}^d)$, 0 is in the interior of $\mathrm{dom}\, A$, and let Θ contain $\mathbf{0}$ in its interior. Then for the log loss $L_{\log}$, there exists a numerical constant $c > 0$ such that for all large enough n

$$\mathfrak{M}_n(\Theta, \Gamma_{\mathrm{all}}, 0) = \inf_{\widehat{p}_n}\sup_{\theta\in\Theta}\mathbb{E}^n_{P_\theta}\left[\mathsf{Risk}_{P_\theta}(\widehat{p}_n) - \mathsf{Risk}^\star_{P_\theta}\right] \ge c\frac{d}{n}.$$

See Appendix B.4 for the proof of Theorem 5. As we essentially only care about the conditional distribution $p_\theta(y \mid x)$, here we chose the marginal over X to be

uniform for convenience; other choices suffice as well.

Comparing the lower bound Theorem 5 provided with the regret bound in Corollary 4, we see that holding RB constant (and $\mathrm{Lip}_\ell(RB)$ constant) then for risk functional $\mathsf{Risk}_P(p) := \mathbb{E}_P[L_{\log}(p(\cdot \mid X), Y)]$, a standard online-to-batch conversion (or Jensen's inequality) implies

$$\mathbb{E}\left[\mathsf{Risk}_P(\widehat{p}_n^{\mathsf{Vovk}}) - \inf_{\theta^\star \in \Theta} \mathsf{Risk}_P(p_\theta^\star)\right] \le (5 + o(1)) \frac{d \log n}{n}$$

as $n \to \infty$, where $o(1) \to 0$ hides problem-dependent constants. To within a factor $O(1) \log n$, then, Vovk's aggregating algorithm is generally unimprovable.

3.2 Asymptotically aggregating to point predictions

While in general the ability to play elements in $\mathrm{Conv}\{p_\theta\}_{\theta \in \Theta}$ could (in principle) yield much better performance than any individual element p_θ for $\theta \in \Theta$, in a sense, the aggregating algorithm is only performing a small amount of averaging to substantially increase its robustness. Indeed, when the risk minimization problem at hand is classical—the loss has continuous derivatives and the population risk Risk_P is strongly convex in a neighborhood of its minimizer $\theta^\star$—then we can show that Vovk's aggregating algorithm asymptotically plays points very close to $\{p_\theta\}_{\theta \in \Theta}$. That is, in "nice" cases, the aggregating algorithm more or less behaves as the empirical risk minimizer, which is asymptotically optimal (Duchi and Ruan, 2020). In these cases, stochastic gradient methods (which are necessarily proper, as they optimizer over θ) similarly achieve optimal asymptotic rates (Duchi and Ruan, 2020), and sometimes similarly strong finite sample rates (Bach, 2014).

We make this more formal via a generalized Bernstein von-Mises Theorem, which shows that when a unique minimizer exists, the density $d\widehat{\mu}_{n,\eta}^{\mathsf{Vovk}}$ in Alg. 1 converges to a normal density centered at the empirical minimizer $\widehat{\theta}_n \in \Theta$ with covariance operator shrinking at the rate $1/n$. Such posterior limiting normality results are relatively well-known: (see van der Vaart (1998); Kleijn and Van der Vaart (2012)). In our setting, rather than considering just the posterior distribution, we consider distributions over Θ of the form of $\widehat{\mu}_{n,\eta}^{\mathsf{Vovk}}$; when L is the log-loss, this is the usual posterior. We first define the class of families and losses for which Theorem 6 holds.

Definition 3.1. *The family and loss pair $(\{p_\theta\}_{\theta \in \Theta}, L)$ is* Bernstein von-Mises generalizable *if there exist $\epsilon_1, \epsilon_2 > 0$ such that the risk $\mathsf{Risk}_P(\theta) := \mathbb{E}_P[L(P_\theta(Y \mid X))]$ satisfies the following conditions:*

(i) The minimizer $\theta^\star = \mathrm{argmin}_{\theta \in \Theta} \mathsf{Risk}_P(\theta)$ is unique and has positive definite Hessian $\nabla^2 \mathsf{Risk}_P(\theta^\star) \succ 0$.

(ii) On the ϵ_1-ball around $\theta^\star$, $\theta^\star + \epsilon_1 \mathbb{B}_2^d$, the loss $\theta \mapsto L(P_\theta(\cdot \mid x), y)$ is $M_{\mathrm{Lip},0}(x,y)$ Lipschitz and has $M_{\mathrm{Lip},2}(x,y)$-Lipschitz Hessian, where $\mathbb{E}[M_{\mathrm{Lip},0}(X,Y)^2] < \infty$ and $\mathbb{E}[M_{\mathrm{Lip},2}(X,Y)] < \infty$.

(iii) For all $\theta \in \Theta \setminus \{\theta^\star + \epsilon_1 \mathbb{B}_2^d\}$, we have $\mathsf{Risk}_P(\theta) \ge \mathsf{Risk}_P(\theta^\star) + \epsilon_2$.

When $\theta \mapsto L(P_\theta(\cdot \mid x), y)$ is convex, condition (iii) is redundant given the others, and the other conditions of Definition 3.1 hold for generalized linear models. Under the conditions Definition 3.1 specifies, we then obtain the following convergence guarantee, whose proof we provide in Appendix B.5.

Theorem 6 (Generalized Bernstein von-Mises)**.** *Let the pair $(\{p_\theta\}_{\theta \in \Theta}, L)$ be Bernstein von-Mises generalizable (Definition 3.1), and for $(X_i, Y_i) \overset{\text{iid}}{\sim} P$ define $\mathsf{Risk}_n(\theta) = \frac{1}{n}\sum_{i=1}^n L(P_\theta(\cdot \mid X_i), Y_i)$. Assume Θ is compact and $\theta^\star \in \operatorname{int} \Theta$. Let $\widehat{\theta}_n := \mathrm{argmin}_{\theta \in \Theta} \mathsf{Risk}_n(\theta)$ and $\widehat{\mu}_{n,\eta}^{\mathsf{Vovk}}$ be as defined in Vovk's Aggregating Algorithm. Then*

$$\left\| \widehat{\mu}_{n,\eta}^{\mathsf{Vovk}} - \mathsf{N}\left(\widehat{\theta}_n, \frac{1}{n} \nabla^2 \mathsf{Risk}_n(\widehat{\theta}_n)^{-1}\right) \right\|_{\mathrm{TV}} \xrightarrow[P]{a.s.} 0.$$

Using the theorem and its proof, we can also (under a minor continuity condition) establish a convergence guarantee showing roughly that the aggregating algorithm asymptotically plays essentially the empirical point estimator. We consider the following assumption.

Assumption 1. *There exists a neighborhood B of $\theta^\star$ such that the log-likelihood $\theta \mapsto \log p_\theta(y \mid x)$ is $\mathrm{Lip}_p(x,y)$-Lipschitz on B, and $\mathrm{Lip}_p(x) := \sup_{\theta \in B} \int_{\mathcal{Y}} \mathrm{Lip}_p(x,y) dP_\theta(y \mid x) < \infty$ for each x.*

We then have the following corollary, whose proof we provide in Appendix B.6.

Corollary 7. *In addition to the conditions of Theorem 6, let Assumption 1 hold. Then for each $x \in \mathcal{X}$,*

$$\left\| \widehat{P}_{n,\eta}^{\mathsf{Vovk}}(\cdot \mid x) - P_{\widehat{\theta}_n}(\cdot \mid x) \right\|_{\mathrm{TV}} \xrightarrow{a.s.} 0.$$

Roughly, Theorem 6 and Corollary 7 show that the aggregating algorithm 1 is asymptotically constrained to making predictions in $\{p_\theta\}_{\theta \in \Theta}$, at least in non-adversarial cases. In a sense, then, the aggregating algorithm 1 is not taking full advantage of its improperness: while it can return any distribution in $\mathrm{Conv}\{p_\theta\}_{\theta \in \Theta}$, it (eventually) is nearly playing elements of $\{p_\theta\}_{\theta \in \Theta}$. While this is optimal in some cases (Theorem 5), the question of how to efficiently and optimally return predictions in $\mathrm{Conv}\{p_\theta\}_{\theta \in \Theta}$ remains open and a natural direction for future work.

4 Experiments and Implementation Details

Before discussing our experiments, we make a few remarks on the computability of Vovk's Aggregating Algorithm. Whenever $L(P_\theta(\cdot \mid x), y)$ is convex in θ and β-smooth, there exists an algorithm Foster et al. (2018) approximating $\widehat{P}^{\mathsf{Vovk}}_{n,\eta}$ that achieves the regret bound in Corollary 4 to within an additive factor $1/n$, and the algorithm is polynomial in $(RB, d, \mathrm{Lip}_\ell(RB), n)$. Yet these algorithms are still computationally intensive; assuming our theoretical results are predictive of actual performance, one might expect that aggregating-type strategies could still yield improvements over standard empirical risk minimization. Indeed, Jézéquel et al. (2020) take the computational difficulty of the approximating algorithms in the paper Foster et al. (2018) as motivation to develop an efficient improper learning algorithm for the special case of logistic regression, which (roughly) hedges its predictions by pretending to receive both positive and negative examples in future time steps, constructing a loss that depends explicitly on the new data x_t; Jézéquel et al. show that it achieves a regret bound with a multiplicative RB factor of the logarithmic regret in Corollary 4. It is unclear how to extend this approach to situations in which the cardinality $|\mathcal{Y}|$ of $\mathcal{Y}$ is much larger than 1, though this is an interesting question for future work. In our experiments, we take a heuristic approach, focusing on the risk minimization setting, and perform aggregation of subsampled maximum likelihood estimators; this approach is reminiscent of the subsampled and bootstrapped estimators Zhang et al. (2013); Kleiner et al. (2012), but we use aggregation as in Alg. 1 to weight predictions. We call the procedure AHA (**A** **H**euristic **A**ggregation) for short.

Algorithm 2 AHA (**A** **H**euristic **A**ggregation)

Input: $\{(x_i, y_i)\}_{i=1}^n$ and parameter radius B
Output: $P^{\mathsf{mix}}_{n,B}$
For $k = 1, \cdots, K$
$S_k \leftarrow$ random subset of the data of size $|S_k| = 2n/3$
$\widehat{\theta}^k_n \leftarrow \mathrm{argmin}_{\|\theta\| \le B} \sum_{(x,y) \in S_k} L(p_\theta(y \mid x))$
$\mu^k_n \leftarrow \exp(\sum_{i=1}^n -L(p_{\widehat{\theta}^k_n}(y_i \mid x_i)))$
$P^{\mathsf{mix}}_{n,B} \leftarrow \left(\sum_{k=1}^K \mu^k_n p_{\widehat{\theta}^k_n}\right) / \left(\sum_{k=1}^K \mu^k_n\right)$

Our results suggest that when performing a probabilistic forecasting task with parameterized model $\{p_\theta\}_{\theta \in \Theta}$, returning the mixture distribution from Vovk's Aggregating Algorithm should be more robust to misspecification than an algorithm which returns $P_{\widehat{\theta}}$; we thus expect Algorithm 2 should exhibit more robustness as well. To that end, we consider two experiments: the first a synthetic experiment with linear regression, where we may explicitly control the degree of misspecification, and the second a logistic regression problem on real digit recognition data, where we mix two populations and we expect (roughly) that a model should do well on one, but may be missing important aspects of the other.

Improper Linear Regression For the synthetic data, we let $X \in \mathbb{R}^d$ be an observed covariate and $H \in \mathbb{R}$ a hidden variable, and for $\tau \in \mathbb{R}_+$ we let y have density

$$p_\tau(y \mid X = x, H = h) = \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}(y - (x^T\theta^\star + \tau h)^2}.$$

We fix the dimension $d = 10$ and let $\theta^\star \in \mathbb{R}^d$ be uniform on $\mathbb{S}^{d-1}$; we generate data by drawing $(X, H) \sim \mathsf{N}(0, I_d) \times \mathsf{N}(0, 1)$. We then use the parametric model $\{p_\theta\}_{\theta \in \Theta}$ to model $Y \mid X = x \sim \mathsf{N}(\theta^T x, 1)$, which is misspecified when $\tau > 0$. As τ grows larger—increasing misspecification—we expect greater differences between the M.L.E. $\widehat{\theta}_n$ and the AHA Algorithm 2.

Logistic Regression We consider the MNIST handwritten digits LeCun et al. (1995), where we mix in typed digits; as our base featurization, we use a standard 7-layer convolutional neural network trained on the MNIST data, so that the typewritten digits (roughly) represent a misspecified sub-population, and as the proportion τ of typewritten digits increases, we expect increasing misspecification. We consider a simplified binary version of this problem, where we seek to distinguish digits 3 and 8, and we use a logistic regression model $p_\theta(y \mid x) = (1 + e^{-yx^T\theta})^{-1}$ with log loss.

Experiment For both linear regression and logistic regression, we conduct the following experiment: For a training sample size n and parameter radius B, we compute the constrained MLE

$$\widehat{\theta}_n := \mathop{\mathrm{argmin}}_{\|\theta\| \le B} \sum_{i=1}^n L(P_\theta(\cdot \mid x_i), y_i)$$

and return $P_{\widehat{\theta}_n}$, and also compute $P^{\mathsf{mix}}_{n,B}$ as the output of Alg. 2 with resampling size $K = 20$ for the improper linear regression experiment and $K = 10$ for the MNIST experiment. We use a held-out test set of size $N = 5000$ to approximate the risk $\mathsf{Risk}(p)$ of the returned conditional probability p. We plot how this approximated risk decays as we increase training sample size n up to 1000 for improper linear regression and up to 200 for logistic regression.

Within each experiment, we implement several regularization schedules. We test $B = c$, $B = c \log n$, $B = c\sqrt{n}$, and $B = cn$ for $c \in \{0.1, 0.2, 1\}$. In Figures

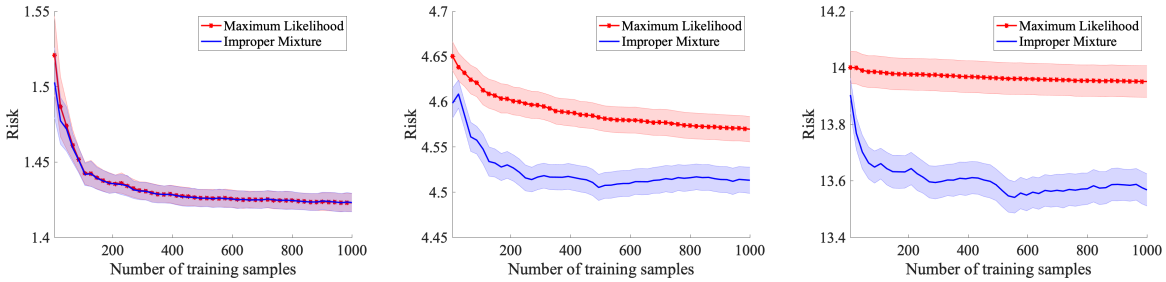

(a) $\tau = 0$ (b) $\tau = 2.5$ (c) $\tau = 5$

Figure 1. Linear Regression, Synthetic Data. As misspecification τ increases, the improper learning algorithm AHA (Alg. 2) outperforms the best constrained MLE.

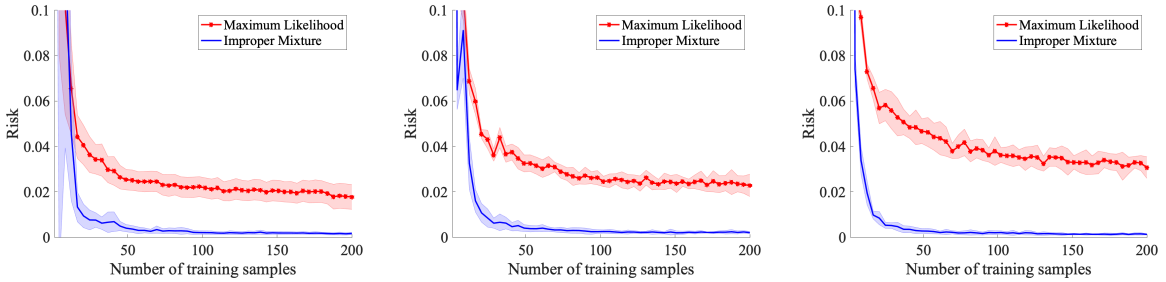

(a) $\tau = 0\%$ (b) $\tau = 5\%$ (c) $\tau = 20\%$

Figure 2. Logistic Regression, MNIST Data mixed with typed data. As misspecification τ increases, the improper learning algorithm AHA's performance remains stable, while the best regularized MLE's performance worsens.

1 and 2 *we only show the results for the best choice of B according to the performance of the Maximum Likelihood Estimator.* We repeat the experiment 100 times on the synthetic data and 10 times on the real dataset and average the results. We run the experiment for $\tau = 0, 2.5, 5$ on the synthetic dataset and $\tau = 0\%, 5\%, 20\%$ for the real dataset.

The results of Figure 1 and Figure 2 are consistent with our expectations: as the magnitude of misspecification (as measured by $\tau \geq 0$) increases, the gap in performance between the maximum likelihood estimator and the aggregated solution increases. Even more, if we may be so bold, the results suggest that using a subsampling and aggregation strategy as in Alg. 2 may be a useful primitive for other learning problems; we leave this as a possibility for future work.

5 Discussion

This work takes steps toward addressing the fundamental and practically important challenge of the cost of inaccurate modeling. While modeling assumptions are ubiquitous throughout statistics, machine learning, and data science—allowing analyses that demonstrate fast convergence rates, efficient algorithms, and interpretable conclusions—most such assumptions are (at least) slightly flawed. This misspecification can have downsides: in addition to perhaps faulty conclusions from a faulty model, even convergence rates of estimators may degrade. This adds a wrinkle to data-modeling tasks: not only must we choose a model that closely fits the data, but we must be mindful of the cost of model misspecification, as this cost is not uniform across models. Our development of the linearity constant $\mathsf{Lin}(\ell, \mathcal{Y}, R, B, n, \gamma)$ in Eq. (6) of the model family gives a reasonably concise description of potential sensitivity to misspecification for many model families.

Yet as we additionally consider, for probabilistic prediction problems aggregation strategies can at least ameliorate these challenges. While aggregation approaches are familiar throughout statistical learning Tsybakov (2004); Dalalyan and Tsybakov (2008); Van Erven et al. (2015), we believe their potential for improvement *beyond* "optimal" point estimators remains unexplored; our results provide one lens for viewing this problem.

6 Acknowledgments

ACM and JD were supported by NSF CAREER Award CCF-1553086, ONR YIP N00014-19-2288, and the DAWN Consortium. AM and GV were supported by DOE award DE-SC0019205, NSF awards 1813049 and 1704417, and ONR YIP N00014-18-1-2295. AM

is additionally supported by an ACM SIGHPC/Intel Fellowship.

References

Alekh Agarwal, Peter L. Bartlett, Pradeep Ravikumar, and Martin J. Wainwright. Information-theoretic lower bounds on the oracle complexity of convex optimization. *IEEE Transactions on Information Theory*, 58(5):3235–3249, 2012.

Peter Auer, Nicol'o Cesa-Bianchi, Yoav Freund, and Robert E. Schapire. The nonstochastic multiarmed bandit problem. *SIAM Journal on Computing*, 32 (1):48–77, 2002.

Francis Bach. Adaptivity of averaged stochastic gradient descent to local strong convexity for logistic regression. *Journal of Machine Learning Research*, 15(1):595–627, 2014.

N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. In *Advances in Neural Information Processing Systems 14*, pages 359–366, 2002.

Nicoló Cesa-Bianchi and Gábor Lugosi. On prediction of individual sequences. *Annals of Statistics*, 27(6): 1865–1895, 1999.

Robert T Clemen and Robert L Winkler. Combining probability distributions from experts in risk analysis. *Risk Analysis*, 19(2):187–203, 1999.

Thomas M. Cover. Universal portfolios. *Mathematical Finance*, 1(1):1–29, January 1991.

Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory, Second Edition*. Wiley, 2006.

Arnak Dalalyan and Alexandre B. Tsybakov. Aggregation by exponential weighting, sharp oracle inequalities and sparsity. *Machine Learning*, 72:39–61, 2008.

A. DeSantis, G. Markowsky, and M.N. Wegman. Learning probabilistic prediction functions. In *COLT*, pages 312–328, 1988.

John C. Duchi. Introductory lectures on stochastic convex optimization. In *The Mathematics of Data*, IAS/Park City Mathematics Series. American Mathematical Society, 2018.

John C. Duchi and Feng Ruan. Asymptotic optimality in stochastic optimization. *Annals of Statistics*, To Appear, 2020.

Dylan J Foster, Satyen Kale, Haipeng Luo, Mehryar Mohri, and Karthik Sridharan. Logistic regression: The importance of being improper. In *Proceedings of the Thirty First Annual Conference on Computational Learning Theory*, pages 167–208, 2018.

Christian Genest and James V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–135, 1986.

Tilmann Gneiting and Adrian Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477): 359–378, 2007.

Clive WJ Granger and Ramu Ramanathan. Improved methods of combining forecasts. *Journal of Forecasting*, 3(2):197–204, 1984.

Peter Grünwald. *The Minimum Description Length Principle*. MIT Press, 2007.

Stephen G Hall and James Mitchell. Combining density forecasts. *International Journal of Forecasting*, 23 (1):1–13, 2007.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer, second edition, 2009.

David Haussler and Andrew Barron. How well do bayes methods work for on-line prediction of +1;-1 values? In *Proceedings of the Third NEC Symposium on Computation and Cognition. SIAM*, 1992.

E. Hazan, A. Agarwal, and S. Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

Elad Hazan and Satyen Kale. An optimal algorithm for stochastic strongly convex optimization. In *Proceedings of the Twenty Fourth Annual Conference on Computational Learning Theory*, 2011. URL http://arxiv.org/abs/1006.2425.

Elad Hazan, Tomer Koren, and Kfir Y Levy. Logistic regression: Tight bounds for stochastic and online optimization. In *Proceedings of the Twenty Seventh Annual Conference on Computational Learning Theory*, pages 197–209, 2014.

Peter J. Huber and Elvezio M. Ronchetti. *Robust Statistics*. John Wiley and Sons, second edition, 2009.

Robert A Jacobs. Methods for combining experts' probability assessments. *Neural computation*, 7(5): 867–888, 1995.

Rémi Jézéquel, Pierre Gaillard, and Alessandro Rudi. Efficient improper learning for online logistic regression. *arXiv:2003.08109*, 2020.

B.J.K. Kleijn and A.W. Van der Vaart. The bernstein-von-mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381, 2012.

Ariel Kleiner, Ameet Talwalkar, Purnamrita Sarkar, and Michael Jordan. Bootstrapping big data. In *Proceedings of the 29th International Conference on Machine Learning*, 2012.

Y. LeCun, L. D. Jackel, L. Bottou, A. Brunot, C. Cortes, J. S. Denker, H. Drucker, I. Guyon, U. A. Muller, E. Sackinger, P. Simard, and V Vapnik. Comparison of learning algorithms for handwritten digit recognition. In *International Conference on Artificial Neural Networks*, pages 53–60, 1995.

Song Mei, Yu Bai, and Andrea Montanari. The landscape of empirical risk for non-convex losses. *Annals of Statistics*, 46(6):2747–2774, 2018.

Neri Merhav and Meir Feder. Universal prediction. *IEEE Transactions on Information Theory*, 44(6): 2124–2147, 1998.

Jorma Rissanen. Universal coding, information, prediction, and estimation. *IEEE Transactions on Information Theory*, 30:629–636, 1984.

Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *Proceedings of the Twenty Second Annual Conference on Computational Learning Theory*, 2009.

Alexandre B Tsybakov. Optimal aggregation of classifiers in statistical learning. *Annals of Statistics*, 32 (1):135–166, 2004.

Aad W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998.

Tim Van Erven, Peter Grunwald, Nishant A Mehta, Mark Reid, Robert Williamson, et al. Fast rates in statistical and online learning. 2015.

V. G. Vovk. Universal forcasting algorithms. *Information and Computation*, (96):245–277, 1992.

V. G. Vovk. A game of prediction with expert advice. *Journal of Computer and System Sciences*, 56(2): 153–173, April 1998.

Volodimir G. Vovk. Aggregating strategies. In *Proceedings of the Third Annual Workshop on Computational Learning Theory*, pages 371–383, 1990.

Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.

Kenji Yamanishi. A loss bound model for on-line stochastic prediction algorithms. *Information and Computation*, 119(1):39–54, 1995.

Yuchen Zhang, John C. Duchi, and Martin J. Wainwright. Communication-efficient algorithms for statistical optimization. *Journal of Machine Learning Research*, 14:3321–3363, 2013.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the Twentieth International Conference on Machine Learning*, 2003.