

A Tight Excess Risk Bound via a Unified PAC-Bayesian–Rademacher–Shtarkov–MDL Complexity

Peter D. Grünwald

Centrum Wiskunde & Informatica and Leiden University

PDG@CWI.NL

Nishant A. Mehta

University of Victoria

NMEHTA@UVIC.CA

Abstract

We present a novel notion of complexity that interpolates between and generalizes some classic existing complexity notions in learning theory: for estimators like empirical risk minimization (ERM) with arbitrary bounded losses, it is upper bounded in terms of data-independent Rademacher complexity; for generalized Bayesian estimators, it is upper bounded by the data-dependent information complexity (also known as stochastic or PAC-Bayesian, $\text{KL}(\text{posterior} \parallel \text{prior})$ complexity. For (penalized) ERM, the new complexity reduces to (generalized) normalized maximum likelihood (NML) complexity, i.e. a minimax log-loss individual-sequence regret. Our first main result bounds excess risk in terms of the new complexity. Our second main result links the new complexity via Rademacher complexity to $L_2(P)$ entropy, thereby generalizing earlier results of Opper, Haussler, Lugosi, and Cesa-Bianchi who did the log-loss case with L_∞ . Together, these results recover optimal bounds for VC- and large (polynomial entropy) classes, replacing localized Rademacher complexity by a simpler analysis which almost completely separates the two aspects that determine the achievable rates: ‘easiness’ (Bernstein) conditions and model complexity.

1. Introduction

We simultaneously address four questions of learning theory:

- (A) We establish a precise relation between Rademacher complexities for arbitrary bounded losses and the minimax cumulative log-loss regret, also known as the *Shtarkov integral* and *normalized maximum likelihood (NML) complexity*.
- (B) We bound this minimax regret in terms of L_2 entropy. Past results were based on L_∞ entropy.
- (C) We introduce a new type of complexity that enables a unification of data-dependent PAC-Bayesian and empirical-process-type excess risk bounds into a single clean bound; this bound recovers minimax optimal rates for large classes under Bernstein ‘easiness’ conditions.
- (D) We extend the link between excess risk bounds for arbitrary losses and codelengths of Bayesian codes to general codes.

All four results are part of the chain of bounds in Figure 1. The \leftarrow arrow stands for ‘bounded in terms of’; the precise bounds (which may hold in probability and expectation or may even be an equality) are given in the respective results in the paper. Red arrows indicate results that are new. We start with a family of predictors \mathcal{F} for an arbitrary *loss function* ℓ , which, for example, may be log-loss, squared error loss or 0/1-loss, and an estimator $\hat{\Pi}$ which on each sample $Z^n = Z_1, \dots, Z_n$ outputs a distribution $\hat{\Pi} \mid Z^n$ on \mathcal{F} ; classic deterministic estimators \hat{f} such as ERM are represented

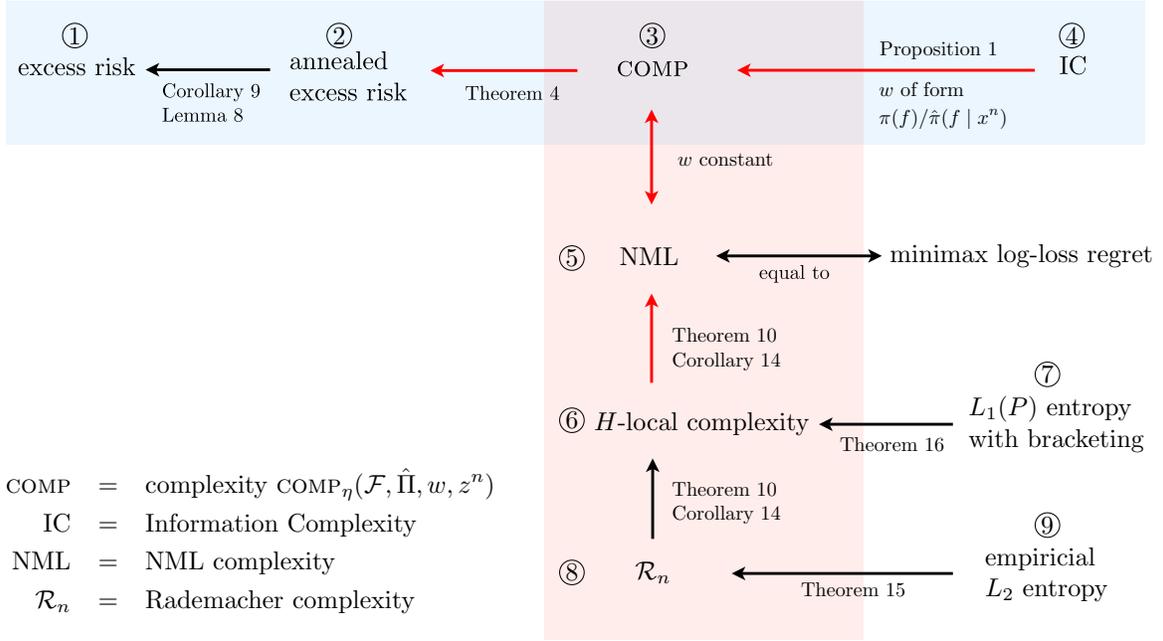


Figure 1: The tree of bounds we provide; red arrows indicate new results.

by taking a $\hat{\Pi}$ that outputs the Dirac measure on \hat{f} . The main bound ② \leftarrow ③, Theorem 4, bounds the *annealed excess risk* of a fixed but arbitrary estimator $\hat{\Pi}$ in terms of its empirical risk on the training data Z^n plus a novel notion of complexity, $\text{COMP}_\eta(\mathcal{F}, \hat{\Pi}, w, Z^n)$ (formulas for COMP_η and all other concepts in the paper are summarized in the Glossary on page 27). The annealed excess risk is a proxy (and lower bound) of the actual excess risk, the expected loss difference between predicting with $\hat{\Pi} | Z^n$ and predicting with the actual risk minimizer f^* within \mathcal{F} . The bound ① \leftarrow ② (Corollary 9, based on Lemma 8, itself from Grünwald (2012)) bounds the actual excess risk in terms of the annealed excess risk, so that we get a true excess risk bound for $\hat{\Pi}$. The complexity $\text{COMP}_\eta(\mathcal{F}, \hat{\Pi}, w, Z^n)$ is dependent on a *luckiness function* $w : \mathcal{Z}^n \times \mathcal{F} \rightarrow \mathbb{R}_0^+$; w can be chosen freely; different choices lead to different complexities and excess risk bounds. For nonconstant w , the complexity becomes data-dependent; in particular, for w of the form $\pi(f)/\hat{\pi}(f | z^n)$, where π is the density of a ‘prior’ distribution Π on \mathcal{F} , the complexity becomes, by ③ \leftarrow ④ (Proposition 1) (strictly) upper bounded by the *information complexity* of Zhang (2006a,b), involving a Kullback-Leibler (KL) divergence term $\text{KL}(\hat{\Pi} | Z^n \| \Pi)$. Information complexity generalizes earlier complexity notions and accompanying bounds from the information theory literature such as (extended) *stochastic complexity* (Rissanen, 1989; Yamanishi, 1998), *resolvability* (Barron and Cover, 1991; Barron et al., 1998), and also excess risk bounds from the PAC-Bayesian literature (Audibert, 2004; Catoni, 2007). Together, ① \leftarrow ② \leftarrow ③ \leftarrow ④ recover and strengthen Zhang’s bounds.

For constant w , the complexity is independent of the data and turns out (③ = ⑤), Section 2.2 to be equal to the *minimax cumulative individual sequence regret* for sequential prediction with log-loss relative to a family $\mathcal{Q}_{\mathcal{F}}$ of probability measures defined in terms of \mathcal{F} , also known as the *log-Shtarkov integral* or *NML (Normalized Maximum Likelihood) complexity*. NML complexity has been much studied in the MDL (minimum description length) literature (Rissanen, 1996; Grünwald, 2007).

Problem A: NML and Rademacher NML complexity can itself be bounded in terms of a new complexity we introduce, H -local complexity, which is further bounded in terms of Rademacher complexity \mathcal{R}_n (Theorem 10 and Corollary 14, ⑤ ← ⑥ ← ⑧). Both Rademacher and NML complexities are used as penalties in model selection (albeit with different motivations), and the close conceptual similarity between NML and Rademacher complexity has been noted by several authors (e.g. Grünwald (2007); Zhu et al. (2009); Roos (2016)). For example, as shown by Grünwald (2007, Open Problem 19, page 583) in classification problems, both the empirical Rademacher complexity for 0-1 loss and the NML complexity of a family of conditional distributions can be simply expressed in terms of a (possibly transformed) minimized cumulative loss that is uniformly averaged over all possible values of the data to be predicted, thereby measuring how well the model \mathcal{F} can fit random noise. Theorem 10 and Corollary 14 establish, for the first time, a precise and tight link between NML and Rademacher complexity. The proofs extend a technique due to Opper and Haussler (1999), who bound NML complexity in terms of L_∞ entropy using an empirical process result of Yurinskii (1976). By using Talagrand’s inequality instead, we get a bound in terms of Rademacher complexity.

Problem B: Bounding NML Complexity with $L_2(P)$ entropy and empirical L_2 Entropy If \mathcal{F} is of VC-type or a class of polynomial empirical L_2 entropy, the Rademacher complexity can be further bounded, (Theorem 15, ⑧ ← ⑨), in terms of the empirical L_2 entropy; if \mathcal{F} admits polynomial $L_1(P)$ entropy with bracketing, then T_n is further bounded, (Theorem 16, ⑥ ← ⑦), in terms of this $L_1(P)$ entropy with bracketing. These latter two results are well-known, due to Koltchinskii (2011) and Massart and Nédélec (2006) respectively, but *in conjunction* with ⑤ ← ⑥ they become of significant interest for log-loss individual sequence prediction. Whereas previous bounds on minimax log-loss regret were invariably in terms of L_∞ entropy (Opper and Haussler, 1999; Cesa-Bianchi and Lugosi, 2001; Rakhlin and Sridharan, 2015), the aforementioned two results allow us to obtain bounds in terms of $L_1(P)$ entropy and empirical L_2 entropy, where P can be any member of the class \mathcal{Q}_f . Unlike the latter two works, however, our results are restricted to static experts that treat the data as i.i.d.

Problem C: Unifying data-dependent and empirical process-type excess risk bounds As lamented by Audibert (2004; 2009), despite their considerable appeal, standard PAC-Bayesian and KL excess risk bounds do not lead to the right rates for large classes, i.e. with polynomial $L_2(P)$ entropy. On the other hand, standard Rademacher complexity generalization and excess risk bound analyses are not easily extendable to either penalized estimators or generalized Bayesian estimators that are based on updating a prior distribution; also handling logarithmic loss appears difficult. Yet ① ← ② ← ③ shows that there does exist a single bound capturing all these applications — by varying the function w one can get both (a strict strengthening of) the KL bounds and a Rademacher complexity-type excess risk bound. In this way, via the chain of bounds ① ← ... ← ⑦/⑨, we recover rates for empirical risk minimization (ERM) that either are minimax optimal (for classification) or the best known rates for ERM (for other losses), even for VC and polynomial entropy classes; the rates depend in the right way on the ‘easiness’ of the problem as measured by the *central condition* (Van Erven et al., 2015), which generalizes Tsybakov’s (2004) margin condition and Bernstein conditions (Bartlett et al., 2005).

Problem D: Excess Risk Bounds and Data Compression While Zhang’s bound holds for arbitrary ‘posteriors’ $\hat{\Pi}$, one gets the best excess risk bounds if one takes $\hat{\Pi}$ to be a *generalized Bayesian estimator*. With such a $\hat{\Pi}$, the information complexity can be expressed in terms of a (generalization

of) the cumulative log-loss of a Bayesian sequential prediction strategy (Zhang, 2006a; Grünwald, 2012) defined relative to the constructed probability model $\mathcal{Q}_{\mathcal{F}}$. By the correspondence between codelengths and cumulative log-loss (reviewed in Section 2), we may say that we bound an *excess risk in terms of a codelength*. ③ = ⑤ shows that we also get a useful excess risk bound in terms of the codelengths of the minimax (NML) code — interestingly, Bayes and NML codes are the two central codes in the *universal coding* literature (Barron et al., 1998; Grünwald, 2007). In fact, our work shows that the correspondence between excess risk bounds and codelengths is a quite general phenomenon, not particular to Bayes and NML codes: in Section 2 we show that there is a 1-to-1 relation between luckiness functions w and codes for data based on \mathcal{Q}_f : each luckiness function w defines, up to scaling, a different code giving rise to a different complexity $\text{COMP}_{\eta}(\mathcal{F}, \hat{\Pi}, w, z^n)$ and hence a different excess risk bound, and vice versa. Because of its relation to data compression, it becomes easy to extend the approach to model selection by using ‘two-part codes’ (Section 5.1).

For yet other choices of w , we obtain bounds for penalized ERM with arbitrary bounded penalization functions (Section 5.2); if we specialize to log-loss, the complexity bound becomes a minimax regret-with-luckiness-term as considered in recent papers on sequential log-loss prediction such as (Kakade et al., 2006; Bartlett et al., 2013). Many other choices of w are possible and remain subject for future investigation.

Additional Features and Limitations The full story above can only be told for bounded losses, although the bounds ② \leftarrow ③, ③ \leftarrow ④ also hold for unbounded losses, and ① \leftarrow ② was recently extended to unbounded losses under a mild additional condition (Grünwald and Mehta, 2016). Remarkably, ② \leftarrow ③ is as tight as can be: when viewed in terms of exponential moments, it is really an equality rather than a bound; this suggests that, no matter the choice of w , the resulting bound is essentially unimprovable. The complexity COMP_{η} depends on a *learning rate parameter* η , and all bounds in the figure become different depending on the choice of η . The optimal η depends on the easiness of the problem at hand, as measured by central/Bernstein/Tsybakov’s conditions (see above). ERM can be applied (and optimal rates can be obtained) without knowledge of the optimal η ; to get the right rates for Bayesian and penalized ERM algorithms however, these algorithms should be made dependent on η (η is akin to $1/\lambda$ in the lasso and ridge regression); in practice, one can learn it from the data using an algorithm such as the ‘safe Bayes’ algorithm of Grünwald (2012) (Section 5.1).

Contents In Section 2.2, we introduce the simple data-independent version of our complexity, $\text{COMP}(\mathcal{F}, \hat{f})$, which is really the NML complexity. In Section 2.4 we extend our notion of complexity to the generalized data-dependent form $\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n)$. Section 3 contains our first main result, Theorem 4. In Section 4, we derive our second main result, Theorem 10 and its Corollary 14, a bound on $\text{COMP}(\mathcal{F}, \hat{f})$ in terms of Rademacher complexity; we also present a concrete application of this result, Theorem 20, which provides the best known rates for ERM under Bernstein conditions for bounded loss functions in a number of situations. Section 5 gives various applications of our result. Finally, Section 6 closes with a discussion of our work in the context of other recent works. All long proofs can be found in the appendix. Mathematical definitions and notations are summarized in the Glossary on page 27.

2. The Novel Complexity Notion

2.1 Preliminaries

In the statistical learning problem (Vapnik, 1998), a labeled sample $Z^n = Z_1, \dots, Z_n$ is drawn independently from probability distribution P over $\mathcal{Z} = (\mathcal{X} \times \mathcal{Y})$, where, for each $j \in [n]$, we have $Z_j = (X_j, Y_j)$. We are given an action space or *model* \mathcal{F} and a *loss function* $\ell : \mathcal{F} \times \mathcal{Z} \rightarrow \mathbb{R}$, where we denote the loss that action or predictor f makes on z as $\ell_f(z)$. Loss functions such as 0-1, squared error, and log-loss (for joint densities on $z = (x, y)$) can all be expressed this way: in the former two cases \mathcal{F} consists of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$, with $\ell_f(x, y) = |y - f(x)|$ and $(y - f(x))^2$ respectively, and in the latter case \mathcal{F} is a set of probability densities on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ relative to some underlying measure ν and $\ell_f(x, y) = -\log f(x, y)$, with log denoting natural logarithm in this paper. An *estimator* or *learning algorithm* $\hat{\Pi}$ is a function from \mathcal{Z}^n to distributions over \mathcal{F} . Here and in the sequel we simply assume that \mathcal{F} is endowed with a suitable sigma-algebra Σ so that (\mathcal{F}, Σ) is a measurable space and all functions we refer to are measurable. We will write $\hat{\Pi} | z^n$ to denote the distribution chosen for data z^n . In practice $\hat{\Pi}$ is often supported entirely on a single function $\hat{f} \in \mathcal{F}$; in that case we simply write the estimator as \hat{f} and the f chosen for given data z^n as $\hat{f}_{|z^n}$. An example of such a *deterministic estimator* is ERM, the empirical risk minimizer. An example of a *randomized estimator* is obtained by setting $\hat{\Pi} | z^n$ to be the generalized η -Bayesian posterior (Zhang, 2006b; Catoni, 2007), which we explicitly define in Section 2.4. Henceforth, we simply call $\hat{\Pi}$ an ‘estimator’ irrespective of whether it is deterministic or randomized.

We aim to learn distributions $\hat{\Pi}$ that obtain low expected *risk* $\mathbb{E}_{f \sim \hat{\Pi}}[\mathbb{E}_{Z \sim P}[\ell_f(Z)]]$. The risk of an action f is the expected value of the loss $\ell_f(Z)$ suffered when playing action f and the actual outcome is Z . A natural way to measure the quality of $\hat{\Pi}$ on data z^n is therefore the *excess risk* $\mathbb{E}_{f \sim \hat{\Pi} | z^n}[\mathbb{E}_{Z \sim P}[\ell_f(Z) - \ell_{f^*}(Z)]]$, where f^* is a minimizer of the risk over \mathcal{F} ; like many other authors (e.g Bartlett et al. (2005)) we assume throughout this work that such a minimizer exists. We use the notation $R_f(z) = \ell_f(z) - \ell_{f^*}(z)$, extended to samples $z^n = (z_1, \dots, z_n) \in \mathcal{Z}^n$ as $R_f(z^n) = \sum_{i=1}^n (\ell_f(z_i) - \ell_{f^*}(z_i))$. Note that, when $Z^n \sim P$, $\hat{\Pi} | Z^n$ and $\hat{f}_{|Z^n}$ can be thought of as random variables so we simply write them without Z^n whenever this cannot cause any confusion.

2.2 The Novel Complexity Measure, Simple Case

To prepare for the definition of our complexity measure COMP, we first need to associate each $f \in \mathcal{F}$ with an associated probability distribution Q_f . We may assume without loss of generality that the underlying distribution P on Z has a density p with respect to some base measure ν (we could for example take $\nu = P$ but the formulas below are easier to parse for general ν). Now for each $f \in \mathcal{F}$, we define Q_f to be the distribution over \mathcal{Z} with density (with respect to the same base measure ν)

$$q_f(z) := \frac{p(z) \cdot e^{-\eta R_f(z)}}{\mathbb{E}_{Z \sim P} [e^{-\eta R_f(Z)}]}. \quad (1)$$

We extend the definition to n outcomes by taking the product densities, $q_f(z^n) := \prod_{i=1}^n q_f(z_i)$. In this way the model \mathcal{F} is itself mapped to a set $\mathcal{Q}_{\mathcal{F}} = \{q_f : f \in \mathcal{F}\}$ of probability densities, the mapping depending on the loss function ℓ of interest, but also (suppressed in notation) on η , f^* , and on the ‘true’ P ; this is an instance of the ‘entropification procedure’ suggested by Grünwald (1999).

We are now ready to define our new complexity measure. For simplicity of we first present the very special case of deterministic estimators \hat{f} without data dependence; this case is sufficient to

make the connection to minimax regret and Rademacher complexities in Section 4. For this setting we define the *Shtarkov integral* (name to be explained below) as

$$\begin{aligned} S(\mathcal{F}; \hat{f}) &:= \mathbb{E}_{Z^n \sim P} \left[\frac{e^{-\eta R_{\hat{f}|Z^n}(Z^n)}}{C(\hat{f}|Z^n)} \right] = \int_{Z^n} q_{\hat{f}|z^n}(z^n) d\nu(z^n) \\ &\quad \dots \left(\text{(if } \ell \text{ log-loss, } \eta = 1, \text{ model correct)} \int_{Z^n} p_{\hat{f}|z^n}(z^n) d\nu(z^n) \right), \end{aligned} \quad (2)$$

where, for any $f \in \mathcal{F}$, $C(f) := \mathbb{E}_{Z^n \sim P} [e^{-\eta R_f(Z^n)}]$ is the normalization constant. Whenever $S(\mathcal{F}, \hat{f})$ is finite (as will automatically be the case with bounded loss), the corresponding *complexity of model \mathcal{F} equipped with \hat{f}* is defined as

$$\text{COMP}(\mathcal{F}, \hat{f}) := \eta^{-1} \log S(\mathcal{F}, \hat{f}). \quad (3)$$

COMP, S , q_f , and normalizer C all depend on η , but this is suppressed in notation unless needed for clarity. The final equality in (2) holds in the very special case that the original loss function is log-loss, $\eta = 1$, and \mathcal{F} contains the density p of P ('the model is correct'). In that case $f^* = p$ (since log-loss is a *proper loss*, see e.g. (Gneiting and Raftery, 2007)) $C(f)$ evaluates to 1 for all $f \in \mathcal{F}$, $\mathcal{Q}_{\mathcal{F}, \ell}$ is equal to \mathcal{F} , and $R_f(z) = -\log f(z) + \log p(z)$; thus, (1) reduces to $q_f(z) = f(z)$, and the final equation in (2) follows. We further define the *maximal complexity* $\text{COMP}(\mathcal{F})$ as

$$S(\mathcal{F}) := \int_{Z^n} \sup_{f \in \mathcal{F}} q_f(z^n) d\nu(z^n) \quad ; \quad \text{COMP}(\mathcal{F}) := \eta^{-1} \log S(\mathcal{F}) = \sup_{\hat{f}} \text{COMP}(\mathcal{F}, \hat{f}), \quad (4)$$

where the final equality is a trivial consequence of the definition, the sup ranging over all deterministic estimators that can be defined on \mathcal{F} .

We often use the following observation due to (e.g.) Opper and Haussler (1999): Let \mathcal{K} be a finite set and let $\{\mathcal{F}_k : k \in \mathcal{K}\}$ be a partition of \mathcal{F} . Then for every deterministic estimator,

$$\text{COMP}(\mathcal{F}, \hat{f}) \leq \frac{\log |\mathcal{K}|}{\eta} + \max_{k \in \mathcal{K}} \text{COMP}(\mathcal{F}_k). \quad (5)$$

This result follows as a special case of Proposition 22 in Section 5, but its proof is simple enough to state in just a few lines:

$$\begin{aligned} \text{COMP}(\mathcal{F}, \hat{f}) &\leq \eta^{-1} \log \int_{Z^n} \max_{k \in \mathcal{K}} \sup_{f \in \mathcal{F}_k} q_f(z^n) d\nu(z^n) \\ &\leq \eta^{-1} \log \int_{Z^n} \sum_{k \in \mathcal{K}} \sup_{f \in \mathcal{F}_k} q_f(z^n) d\nu(z^n) \\ &\leq \eta^{-1} \log |\mathcal{K}| + \eta^{-1} \max_{k \in \mathcal{K}} \log \int_{Z^n} \sup_{f \in \mathcal{F}_k} q_f(z^n) d\nu(z^n). \end{aligned}$$

Using (5), we can link COMP to Rademacher complexity, which we will do in Section 4.1 and 4.2. Below, we first link COMP to log-loss prediction, extend it to encompass data-dependent and PAC-Bayesian complexities and present our excess risk bound for the general complexities (Section 4.1 and 4.2) can be read without this material).

Minimax Cumulative Log-Loss Interpretation of COMP For every given estimator \hat{f} , we can define a density r on \mathcal{Z}^n relative to ν by setting

$$r(z^n) := \frac{q_{\hat{f}}(z^n)}{S(\mathcal{F}, \hat{f})}, \quad (6)$$

which evidently integrates to 1 and hence is a probability density (different choice of estimator \hat{f} leads to different r ; this is suppressed in the notation). We can use density r to sequentially predict Z_1, Z_2, \dots, Z_n by predicting Z_i with the corresponding conditional density $r(Z_i | Z^{i-1})$. The cumulative log-loss obtained this way is given by

$$\sum_{i=1}^n -\log r(Z_i | Z^{i-1}) = -\log r(Z^n),$$

the latter equality following by definition of conditional probability and telescoping. Because of the correspondence, via Kraft’s inequality, of log-loss prediction and data compression, we can also think of this quantity as a codelength. Similarly, $\min_{f \in \mathcal{F}} -\log q_f(Z^n)$ is the minimum cumulative loss one could have obtained *with hindsight*, i.e. if one had sequentially predicted the Z_i by the q_f that turned out to minimize $-\log q_f$ on Z^n . Assuming this minimum is well-defined it is of course achieved by \hat{f}^{ML} , the maximum likelihood estimator relative to \mathcal{Q} , for which evidently also $\text{COMP}(\mathcal{F}) = \text{COMP}(\mathcal{F}, \hat{f}^{\text{ML}})$. Thus we get that for all $z^n \in \mathcal{Z}^n$,

$$\begin{aligned} \eta^{-1} \cdot \text{COMP}(\mathcal{F}, \hat{f}) &= \log S(\mathcal{F}, \hat{f}) = -\log r(z^n) - \left(-\log q_{\hat{f}}(z^n)\right) \\ \text{if } \hat{f} &\stackrel{\text{ML}}{=} \hat{f}^{\text{ML}} = \eta^{-1} \cdot \text{COMP}(\mathcal{F}) = -\log r(z^n) - \min_{f \in \mathcal{F}} \left(-\log q_f(z^n)\right), \end{aligned} \quad (7)$$

the first equation holding for general \hat{f} and the second for \hat{f}^{ML} . The final expression is just the (cumulative log-loss) *regret* of r on data z^n , which, by (7), is constant on z^n . As first noted by [Shtarkov \(1987\)](#), this implies that (7) is also the *minimax individual-sequence regret* relative to the model \mathcal{Q} when sequentially predicting outcomes Z_1, \dots, Z_n with the log-loss; the corresponding optimal sequential prediction strategy r is usually called the normalized maximum likelihood (NML) or Shtarkov density; see [Rissanen \(1996\)](#); [Grünwald \(2007\)](#); [Opper and Haussler \(1999\)](#); [Cesa-Bianchi and Lugosi \(2001\)](#) for details.

2.3 Allowing Data-Dependency

We now generalize the complexity definition above for arbitrary deterministic \hat{f} so that it becomes data dependent; further extension to randomized estimators follows in Section 2.4. The central concept we need is that of a *luckiness function* $w : \mathcal{Z}^n \rightarrow \mathbb{R}_0^+$; every combination of estimator and luckiness function will, up to scaling, define a unique version of complexity; and every such complexity will induce a different data-dependent bound on excess risk. We call w ‘luckiness function’ since it will influence our excess risk bounds so that they become better iff we are ‘lucky’ in the sense that P is such that $w(X^n)$ will be large with high probability).

The *generalized Shtarkov integral* for estimator \hat{f} relative to luckiness function w is defined as

$$S(\mathcal{F}, \hat{f}, w) := \mathbb{E}_{Z^n \sim P} \left[\frac{e^{-\eta R_{\hat{f}|Z^n}(Z^n)}}{C(\hat{f}|Z^n)} \cdot w(Z^n) \right] = \int_{\mathcal{Z}^n} q_{\hat{f}|z^n}(z^n) w(z^n) d\nu(z^n), \quad (8)$$

and, whenever $S(\mathcal{F}, \hat{f}, w) < \infty$, we define the corresponding data-dependent complexity as

$$\text{COMP}(\mathcal{F}, \hat{f}, w, z^n) := \frac{1}{\eta} \left(-\log w(z^n) + \log S(\mathcal{F}, \hat{f}, w) \right). \quad (9)$$

Both expressions evidently reduce to (2) and (3) if we take w constant over Z^n .

Cumulative Log-Loss Interpretation Fix an arbitrary estimator \hat{f} . Then for any luckiness function w with $S(\mathcal{F}, \hat{f}, w) < \infty$, we can define the probability density

$$r_w(z^n) := \frac{q_{\hat{f}|z^n}(z^n) \cdot w(z^n)}{S(\mathcal{F}, \hat{f}, w)}, \quad (10)$$

with (6) being the special case with $w \equiv 1$. Just as with r_1 , for general such w , r_w can be thought of as a sequential prediction strategy, and $\eta \cdot \text{COMP}(\mathcal{F}, \hat{f}, w, z^n) - \log q_{\hat{f}|z^n}(z^n) = -\log r_w(z^n)$ is the cumulative log-loss achieved by r_w . Different (up to scaling) w generate different log-loss prediction strategies (codes) and corresponding complexities. Conversely, for every probability density r' relative to ν on Z^n , we can set a luckiness measure $w(z^n)$ proportional to $r'(z^n)/q_{\hat{f}|z^n}(z^n)$; with the appropriately scaled choice of w , r_w will coincide r' ; we thus have a 1-to-1-correspondence between luckiness functions w with $S(\mathcal{F}, \hat{f}, w) < \infty$, codes and complexities.

2.4 The Novel Complexity Measure, General Case

Here we further generalize the complexity definition so that it can output distributions $\hat{\Pi} | Z^n$ on \mathcal{F} . For this we need to extend the domain of the luckiness function to encompass \mathcal{F} , i.e. we now take arbitrary functions of the form $w : Z^n \times \mathcal{F} \rightarrow \mathbb{R}_0^+$.

The generalized Shtarkov integral for estimator $\hat{\Pi}$ relative to luckiness function w is defined as

$$S(\mathcal{F}, \hat{\Pi}, w) := \mathbb{E}_{Z^n \sim P} \left[\exp \left(-\mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} \left[\eta R_{\underline{f}}(Z^n) + \log C(\underline{f}) - \log w(Z^n, \underline{f}) \right] \right) \right], \quad (11)$$

and the generalized (data-dependent) model complexity corresponding to (11) is now defined as

$$\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) := \frac{1}{\eta} \cdot \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi} | z^n} \left[-\log w(z^n, \underline{f}) \right] + \log S(\mathcal{F}, \hat{\Pi}, w) \right). \quad (12)$$

Both expressions are readily seen to generalize (8) and (9) respectively: if, for a given deterministic estimator \hat{f} , we take $\hat{\Pi}(\cdot | Z^n)$ to be $\delta_{\hat{f}}$ (the Dirac measure on $\hat{f}|_{Z^n}$) and we take a function $w(z^n, f) \equiv w(z^n)$ that does not depend on f , then the expressions above simplify trivially to (8) and (9) respectively; thus $\text{COMP}(\mathcal{F}, \delta_{\hat{f}}, w, z^n) = \text{COMP}(\mathcal{F}, \hat{f}, w, z^n)$. Finally, we define

$$\text{COMP}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, z^n) := \text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) + \mathbb{E}_{\underline{f} \sim \hat{\Pi} | z^n} [R_{\underline{f}}(z^n)] \quad (13)$$

as the sum of the complexity and the expected excess loss that a random draw from $\hat{\Pi}$ achieves on the data.

COMP generalizes information complexities To explain how PAC-Bayesian type complexity arise a special case of COMP, we consider luckiness measures w that are defined in terms of probability distributions Π on \mathcal{F} that do not depend on the data; we call these ‘priors’. For notational convenience it is useful to assume (without loss of generality) that Π has a density π relative to some underlying measure ρ on \mathcal{F} and that, for all $z^n \in \mathcal{Z}^n$, $\hat{\Pi} | z^n$ also has a density $\hat{\pi} | z^n$ relative to ρ .

Proposition 1 Consider arbitrary Π and $\hat{\Pi}$ as above with densities π and $\hat{\pi} | z^n$ relative to some ρ . Set $w(z^n, f) := \pi(f)/\hat{\pi}(f | z^n)$. Then we have

$$S(\mathcal{F}, \hat{\Pi}, w) \leq 1. \quad (14)$$

Consequently,

$$\text{COMP}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, z^n) \leq \mathbb{E}_{\underline{f} \sim \hat{\Pi} | z^n} [R_{\underline{f}}(z^n)] + \eta^{-1} \cdot \text{KL}(\hat{\Pi} | z^n \| \Pi), \quad (15)$$

where $\text{KL}(\hat{\Pi} | z^n \| \Pi) = \mathbb{E}_{\underline{f} \sim \hat{\Pi} | z^n} [\log \hat{\pi}(f | z^n)/\pi(f)]$ is KL divergence.

Proof By Jensen’s inequality applied to (11), we have, using the definition of w and Fubini’s theorem,

$$\begin{aligned} S(\mathcal{F}, \hat{\Pi}, w) &\leq \mathbb{E}_{Z^n \sim P} \mathbb{E}_{\underline{f} \sim \hat{\Pi} | z^n} \left[\frac{e^{-\eta R_{\underline{f}}(z^n)}}{C(\underline{f})} \cdot w(z^n, \underline{f}) \right] = \mathbb{E}_{Z^n \sim P} \mathbb{E}_{\underline{f} \sim \Pi} \left[\frac{e^{-\eta R_{\underline{f}}(z^n)}}{C(\underline{f})} \right] \\ &= \mathbb{E}_{\underline{f} \sim \Pi} \left[\int_{z^n} q_{\underline{f}}(z^n) d\nu(z^n) \right] = 1, \end{aligned}$$

which gives (14); (15) follows by plugging in our choice of w into the definition of COMP. \blacksquare

We thus see that $\text{COMP}^{\text{FULL}}$ is upper bounded by *information complexity* defined relative to prior Π (Zhang, 2006a,b), which is just (15) normalized (divided by n). The notion of information complexity is also used to bound excess risks in the PAC-Bayesian approach of Catoni (2007) and Audibert (2004). As noted by Zhang (2006b), the right-hand side of (15) is minimized if we take as our estimator $\hat{\Pi}$ the η -generalized Bayesian estimator,

$$\hat{\pi}(f | z^n) := \frac{\exp(-\eta \sum_{i=1}^n \ell_f(z_i)) \cdot \pi(f)}{\int \exp(-\eta \sum_{i=1}^n \ell_f(z_i)) \cdot \pi(f) d\rho(f)}, \quad (16)$$

and in that case is equal to the *generalized marginal likelihood*, known in the MDL literature as the *extended stochastic complexity* (Yamanishi, 1998)

$$-\eta^{-1} \log \mathbb{E}_{\underline{f} \sim W} [\exp(-\eta R_{\underline{f}}(Z^n))], \quad (17)$$

which, for $\eta = 1$ and ℓ the log-loss, coincides with the standard log Bayesian marginal likelihood.

We provide some further simple properties of $\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n)$ for general $\hat{\Pi}$ and w in Section 5.

Cumulative Log-Loss Interpretation Just as for deterministic estimators, we note that every randomized estimator $\hat{\Pi}$ and luckiness function w defines a probability density/prediction strategy on Z^n by setting

$$r_w(z^n) := \frac{e^{\left(\mathbb{E}_{\underline{f} \sim \hat{\Pi}|z^n} [\log q_{\underline{f}}(z^n) \cdot w(z^n, \underline{f})]\right)}}{S(\mathcal{F}, \hat{\Pi}, w)},$$

and just as before, COMP can be interpreted in terms of the ‘code’ r_w .

3. First Main Result: Bounding Excess Risk in Terms of New Complexity

In this section and Section 4, we restrict to the bounded loss setting; given this restriction, it is without loss of generality that we assume that

$$\sup_{f, g \in \mathcal{F}} \text{ess sup} |\ell_f(Z) - \ell_g(Z)| \leq \frac{1}{2}, \quad (\text{A1})$$

as this always can be accomplished by an appropriate scaling of the loss.

Before presenting our first main result, it will be useful to introduce a variant of an ordinary expectation as well as some notation. For $\eta > 0$ and general random variables U , we define the *annealed expectation* (see Grünwald and Mehta (2016) for the origin of this terminology) as

$$\mathbb{E}^{\text{ANN}, \eta} [U] = -\frac{1}{\eta} \log \mathbb{E} \left[e^{-\eta U} \right]. \quad (\text{18})$$

Below we will first bound the annealed excess risk rather than the standard excess risk and then continue to bound the latter in terms of the former. Our first main result below may be expressed succinctly via the notion of *exponential stochastic inequality*,

Definition 2 (Exponential Stochastic Inequality (ESI)) Let $\eta > 0$ and let U, U' be random variables on some probability space with probability measure P . We define

$$U \triangleleft_{\eta} U' \Leftrightarrow \mathbb{E}_{U, U' \sim P} \left[e^{\eta(U-U')} \right] \leq 1, \quad (\text{19})$$

and we write $U \triangleleft_{\eta}^* U'$ iff the right hand of (19) holds with equality.

Clearly $U \triangleleft_{\eta}^* U' \Rightarrow U \triangleleft_{\eta} U'$. An ESI simultaneously captures high probability and in-expectation results:

Proposition 3 (ESI Implications) For all $\eta > 0$, if $U \triangleleft_{\eta} U'$ then, (i), $\mathbb{E}[U] \leq \mathbb{E}[U']$; and, (ii), for all $K > 0$, with P -probability at least $1 - e^{-K}$, $U \leq U' + K/\eta$.

Proof Jensen’s inequality yields (i). Apply Markov’s inequality to $e^{-\eta(U-U')}$ for (ii). ■

We now present our first main result, a new bound that interpolates between the Zhang bound and standard empirical process theory bounds for handling large classes and that is sharp in the sense that it really is an equality of exponential moments.

Theorem 4 For every randomized estimator $\hat{\Pi}$ and every luckiness function $w : \mathcal{Z}^N \times \mathcal{F} \rightarrow \mathbb{R}_0^+$, we have

$$\mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} \left[\mathbb{E}_{\bar{Z} \sim P}^{\text{ANN}, \eta} \left[R_f(\bar{Z}) \right] \right] \leq_{n\eta}^* \frac{1}{n} \cdot \text{COMP}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, Z^n). \quad (20)$$

The proof is in the appendix; it is merely a sequence of straightforward rewritings, where the key observation is that for every $f \in \mathcal{F}$, the annealed risk $\mathbb{E}_{\bar{Z} \sim P}^{\text{ANN}, \eta} [R_f(\bar{Z})]$ is related to the normalization factor appearing in the definition (1) of the probability density q_f and its n -fold product $C(f)$ appearing in (2) via the following equality, as follows immediately from the definitions:

$$\mathbb{E}_{\bar{Z} \sim P}^{\text{ANN}, \eta} [R_f(\bar{Z})] = \frac{1}{n} \cdot \frac{-\log C(f)}{\eta}. \quad (21)$$

We note that, by taking w as in Proposition 1, via (15), this theorem strictly generalizes Theorem 3.1 of Zhang (2006b), the left-hand side of Zhang’s inequality being equal to the annealed excess risk and the right-hand side to the information complexity, i.e. the right hand of (15). However, by taking different w , we get different bounds which are not covered by Zhang’s results and which, as we’ll see, can be used to recover minimax excess risk bounds for certain large classes of polynomial entropy.

The above ESI’s have annealed expectations on their left-hand sides and thus still fall short of providing such excess risk bounds.

This gap can be resolved under the v -central condition.

Definition 5 (v-central condition) Let $v: x \mapsto \eta_0 \cdot x^\alpha$ be a function with domain $[0, \infty)$ for some $\eta_0 > 0$ and $\alpha \geq 0$. We say that (P, ℓ, \mathcal{F}) satisfies the v -central condition if, for all $\gamma > 0$,

$$\mathbb{E}[e^{-v(\gamma)R_f(Z)}] \leq e^{v(\gamma)\gamma}.$$

In the special case of $\alpha = 0$, we say that the η -central condition holds (for $\eta = \eta_0$).

If the loss is η -exp-concave and the class \mathcal{F} is convex, it is known that the η -central condition holds (see the figure on page 1798 of Van Erven et al. (2015), or Lemma 1 of Mehta (2017) for an explicit proof). More generally, in the case of bounded losses the v -central condition is in fact equivalent to the well-known *Bernstein condition*.

Definition 6 (Bernstein condition) Let $\beta \in [0, 1]$. We say that (P, ℓ, \mathcal{F}) satisfies the β -Bernstein condition if, for some constant $B < \infty$

$$\mathbb{E} \left[R_f(Z)^2 \right] \leq B \mathbb{E} \left[R_f(Z) \right]^\beta \quad \text{for all } f \in \mathcal{F}.$$

We only recall one direction of the equivalence of the v -central condition to the Bernstein condition here; the full equivalence is due to (Van Erven et al., 2015, Theorem 5.4).

Lemma 7 (Bernstein implies v-central) Assume for all $f \in \mathcal{F}$ that $R_f(Z) \in [-1/2, 1/2]$ a.s. If the β -Bernstein condition holds for some $\beta \in [0, 1]$ and some constant B , then the v -central condition holds for

$$v(\gamma) = \min \left\{ \frac{\gamma^{1-\beta}}{B}, 1 \right\}.$$

Proof For clarity, let \bar{a} and \bar{b} refer to the constants a and b from part 1(a) of Theorem 5.4 of [Van Erven et al. \(2015\)](#). Apply that result with $\bar{b} = \frac{1}{2\bar{a}}$, $\bar{a} = 1/2$, and the u function there set to $x \mapsto Bx^\beta$. Note that although the statement of Theorem 5.4 actually imposes the stronger condition that the loss ℓ be $[0, 1/2]$ -valued, the proof thereof only requires that $R_f \in [-1/2, 1/2]$ a.s. for all $f \in \mathcal{F}$. ■

Note that for such bounded loss functions, the weakest Bernstein condition with $\beta = 0$ holds automatically and so does the v -central condition with $v(\gamma) \propto \gamma$.

The following lemma is a translation of Lemma 2 of [Grünwald \(2012\)](#) which addresses the aforementioned gap between the annealed and actual expectations.

Lemma 8 *Suppose that the v -central condition holds for some $\eta_0 > 0$ and $\alpha \geq 0$. If $R_f(Z) \in [-1/2, 1/2]$ a.s., then for all $\gamma > 0$, for all $\eta \leq \frac{v(\gamma)}{2}$*

$$\mathbb{E}_{Z \sim P} [R_f(Z)] \leq C_\eta \cdot \mathbb{E}_{Z \sim P}^{\text{ANN}, \eta} [R_f(Z)] + \frac{C_\eta - 1}{\eta} v(\gamma) \cdot \gamma,$$

with $C_\eta = 2 + 2\eta$. In particular, taking $\eta = v(\gamma)/2$, we have

$$\mathbb{E}_{Z \sim P} [R_f(Z)] \leq C_{v(\gamma)/2} \cdot \mathbb{E}_{Z \sim P}^{\text{ANN}, v(\gamma)/2} [R_f(Z)] + 2(C_{v(\gamma)/2} - 1)\gamma.$$

We note that a version of the above result also holds for general bounded losses. In fact, [Grünwald and Mehta \(2016\)](#) (still under review) provide a refined version of this lemma that works even for unbounded losses and for any $\eta < v(\gamma)$, with sharper bounds for $\eta < v(\gamma)/2$.

With all the pieces in place, the next two excess risk bounds in terms of COMP are nearly immediate.

Corollary 9 *Take the same setup as Lemma 8. If $\hat{\Pi}$ is an arbitrary randomized estimator, and w is an arbitrary luckiness function, then*

$$\mathbb{E}_{\underline{f} \sim \hat{\Pi}_{|Z^n}} \left[\mathbb{E}_{Z \sim P} [R_{\underline{f}}(Z)] \right] \leq_{v(\gamma) \cdot n/6} \frac{3 \text{COMP}_{v(\gamma)/2}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, Z^n)}{n} + 4\gamma. \quad (22)$$

If \hat{f} is ERM, then

$$\mathbb{E}_{Z \sim P} [R_{\hat{f}}(Z)] \leq_{v(\gamma) \cdot n/6} \frac{3 \text{COMP}_{v(\gamma)/2}(\mathcal{F}, \hat{f})}{n} + 4\gamma. \quad (23)$$

Proof For (22), start with Lemma 8 with $\eta = v(\gamma)/2$ for the desired $\gamma > 0$, and then apply Theorem 4 to (stochastically) upper bound the annealed excess risk term. Observe that since $v(\gamma) \leq 1$ by assumption, we have $C_{v(\gamma)/2} \leq 3$.

For (23), start with (22), and take $w \equiv 1$ and $\hat{\Pi}(\cdot | Z^n)$ equal to the probability measure that places mass 1 on $\hat{f}_{|Z^n}$ and 0 elsewhere. From these settings and the optimality of ERM for the empirical risk, $\text{COMP}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, Z^n)$ reduces to the simpler form $\text{COMP}(\mathcal{F}, \hat{f})$. ■

To aid in the interpretation of the corollary, let us remark on two special cases of (22). In both cases, we will suppose that, as in Proposition 1, $w(z^n, f) := \pi(f)/\hat{\pi}(f | z^n)$ where π is the density of a fixed probability measure on \mathcal{F} independent of the sample, so that COMP is bounded by

information complexity. First, if the η -central condition holds, then, setting $\eta' = \eta/2$ and using (15), it further follows that

$$\mathbb{E}_{\underline{f} \sim \hat{\Pi}_{|Z^n}} \left[\mathbb{E}_{Z \sim P} \left[R_{\underline{f}}(Z) \right] \right] \leq \frac{\eta}{6} \frac{3}{n} \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi}_{|Z^n}} [R_{\underline{f}}(Z^n)] + \frac{2 \text{KL}(\hat{\Pi} | Z^n \| \Pi)}{\eta} \right).$$

In the second case, we take $\hat{\Pi}$ to be any posterior whose $\hat{\Pi}$ -expected empirical risk is at most the empirical risk of f^* (for example, $\hat{\Pi}$ could be Dirac measure on, hence coincide with, ERM), and for simplicity we further assume that a β -Bernstein condition holds for some $B \geq 2$ (if it holds for a smaller B , we will simply weaken the condition). Thus, from the bounded loss assumption the v -central condition holds for $v(\gamma) = \frac{\gamma^{1-\beta}}{B}$ (provided that we only consider $\gamma \leq B^{1/(1-\beta)}$), and tuning γ yields $\gamma = A_1 \cdot n^{-1/(2-\beta)} \text{KL}(\hat{\Pi} | Z^n \| \Pi)^{1/(2-\beta)}$ for a constant A_1 depending only on β and B , so that

$$\mathbb{E}_{\underline{f} \sim \hat{\Pi}_{|Z^n}} \left[\mathbb{E}_{Z \sim P} \left[R_{\underline{f}}(Z) \right] \right] \leq_{n \cdot a_n} A_1 \cdot \left(\frac{\text{KL}(\hat{\Pi} | Z^n \| \Pi)}{n} \right)^{1/(2-\beta)}, \quad (24)$$

where $a_n = A_2 (\text{KL}(\hat{\Pi} | Z^n \| \Pi)/n)^{(1-\beta)/(2-\beta)}$ for a constant A_2 depending only on β and B . Lastly, as usual, in both cases when the class is finite and the prior Π is uniform, the KL-divergence term reduces to $\log |\mathcal{F}|$. We thus retrieve the familiar bounds $O(n^{-1/2})$ in the worst-case ($\beta = 0$, for which the Bernstein condition holds vacuously for bounded losses) and $O(n^{-1})$ for the best case, $\beta = 1$.

In the next section, we derive excess risk bounds for large classes by suitably controlling COMP and then applying Corollary 9.

4. Bounds on Maximal Complexity $\text{COMP}(\mathcal{F})$ and the excess risk bounds they imply

The results of the previous section do not yet yield explicit excess risk bounds as they still involve a COMP term. In this section, we leverage and extend ideas from [Opper and Haussler \(1999\)](#) as well as results from empirical process theory to provide explicit bounds on COMP for several important types of large classes: classes of VC-type, classes whose empirical entropy grows polynomially, and sets of classifiers whose entropy with bracketing grows polynomially. Along the way, we form vital connections to expected suprema of certain empirical processes, including Rademacher complexity. At the end of this section we present explicit excess risk bounds; these bounds are simple consequences of the bounds we developed on COMP.

4.1 Preliminaries

Losses and Lipschitzness. To properly capture losses like log-loss and supervised losses like 0-1 loss and squared loss, we introduce two different parameterizations of the loss function:

1. the *direct* parameterization: $\ell_f(z) = f(x, y)$;
2. and the *supervised loss* parameterization: $\ell_f(z) = \ell(y, f(x))$.

For example, in the case of conditional density estimation with log loss, we then have $\ell_f(z) = f(x, y) = -\log p_f(y | x)$. Thus, each function $f \in \mathcal{F}$ has domain \mathcal{Z} , and the equivalence $\mathcal{F} = \{\ell_f :$

$f \in \mathcal{F}$ holds. On the other hand, for supervised losses, each function $f \in \mathcal{F}$ has domain \mathcal{X} while each loss-composed function ℓ_f has domain \mathcal{Z} .

Unlike previous sections, in this section we require an additional assumption in the case of the supervised loss parameterization: we assume that, for each outcome $(x, y) = z \in \mathcal{Z}$, the loss $\ell_f(z) = \ell(y, f(x))$ is L -Lipschitz in its second argument, i.e. for all $f, g \in \mathcal{F}$,

$$|\ell(y, f(x)) - \ell(y, g(x))| \leq L |f(x) - g(x)|. \quad (\text{A2})$$

In the case of classification with 0-1 loss, \mathcal{F} is the set of classifiers taking values in $\{0, 1\}$ and $\mathcal{Y} = \{0, 1\}$, and so (A2) will hold with $L = 1$ (and is in fact an equality). For convenience in the analysis, in the case of the direct parameterization we may always take $L = 1$.

Standard complexity measures. It will be useful to review some of the standard notions of complexity before presenting our bounds. In the below, let \mathcal{H} be a class of functions mapping from some space \mathcal{S} to \mathbb{R} ; we typically will take \mathcal{S} equal to either \mathcal{X} or \mathcal{Z} .

For a pseudonorm $\|\cdot\|$, the ε -covering number $\mathcal{N}(\mathcal{H}, \|\cdot\|, \varepsilon)$ is the minimum number of radius- ε balls in the pseudonorm $\|\cdot\|$ whose union contains \mathcal{H} . We will work with the $L_2(Q)$ (or $L_1(Q)$) pseudonorms for some probability measure Q . A case that will occur frequently is when $Q = P_n$ is the empirical measure $\frac{1}{n} \sum_{j=1}^n \delta_{S_j}$ based on a sample S_1, \dots, S_n ; here, δ_s (for $s \in \mathcal{S}$) is a Dirac measure, and the sample will always be clear from the context.

For two functions $h^{(l)}$ and $h^{(u)}$, the *bracket* $[h^{(l)}, h^{(u)}]$ is the set of all functions f that satisfy $h^{(l)} \leq f \leq h^{(u)}$. An ε -bracket (in some pseudonorm $\|\cdot\|$) is a bracket $[h^{(l)}, h^{(u)}]$ satisfying $\|h^{(l)} - h^{(u)}\| \leq \varepsilon$. The ε -bracketing number $\mathcal{N}_{[\cdot]}(\mathcal{H}, \|\cdot\|, \varepsilon)$ is the minimum number of ε -brackets that cover \mathcal{H} ; the logarithm of the ε -bracketing number is called the ε -entropy with bracketing.

Let $\epsilon_1, \dots, \epsilon_n$ be independent Rademacher random variables (distributed uniformly on $\{-1, 1\}$). The *empirical Rademacher complexity* of \mathcal{H} and the *Rademacher complexity* of \mathcal{H} respectively are

$$\mathcal{R}_n(\mathcal{H} \mid S_1, \dots, S_n) := \mathbb{E}_{\epsilon_1, \dots, \epsilon_n} \left[\sup_{h \in \mathcal{H}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i h(S_i) \right| \right] \quad \mathcal{R}_n(\mathcal{H}) := \mathbb{E} \left[\sup_{h \in \mathcal{H}} \left| \frac{1}{n} \sum_{i=1}^n \epsilon_i h(S_i) \right| \right],$$

where the first expectation is conditional on S_1, \dots, S_n .

4.2 H -local complexity and Rademacher complexity bounds on the NML complexity

We first show that the simple form of the complexity $\text{COMP}(\mathcal{F}, \hat{f}) \leq \text{COMP}(\mathcal{F})$ can be directly upper bounded in terms of two other complexity notions, the H -local complexity (defined below) and Rademacher complexity, up to a constant depending on the $L_2(P)$ diameter of \mathcal{F} .

Theorem 10 Fix $\varepsilon > 0$ and let \mathcal{F} have diameter ε in the $L_2(P)$ pseudometric. Define $\sigma := eL\varepsilon$, fix arbitrary $f_0 \in \mathcal{F}$, and define the loss class $\mathcal{G} := \{\ell_{f_0} - \ell_f : f \in \mathcal{F}\}$. Define

$$T_n := \sup_{f \in \mathcal{F}} \left\{ \sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) - \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) \right] \right\}.$$

Then

$$\text{COMP}_\eta(\mathcal{F}) \leq 3 \mathbb{E}_{Z^n \sim Q_{f_0}} [T_n] + n\eta\sigma^2 \quad (25)$$

$$\leq 6n \mathbb{E}_{Z^n \sim Q_{f_0}} [\mathcal{R}_n(\mathcal{G})] + n\eta\sigma^2. \quad (26)$$

Two remarks are in order. First, we refer to the quantity $\mathbb{E}_{Z^n \sim Q_{f_0}} [T_n]$ as an *entropified local complexity*, or *H-local complexity* for short. The “local” part of the name stems from how, in the empirical process inside the supremum defining T_n , the losses are centered/localized around ℓ_{f_0} ; the “entropified” part of the name is due to the fact that the sample is distributed according to Q_{f_0} , itself defined via entropification. Second, the attentive reader may have noticed that in the above theorem, the expectation in the Rademacher complexity is relative to the distribution Q_{f_0} for arbitrary $f_0 \in \mathcal{F}$, rather than the distribution P generating the data. Moreover, the appearance of Q_{f_0} appears to dampen the utility of \mathcal{F} having small $L_2(P)$ diameter. This apparent mismatch will be of no concern due to a technical lemma (Lemma 23 in Appendix A.2), which relates the $L_2(Q_{f_0})$ and $L_2(P)$ pseudometrics.

We now sketch a proof of this theorem in three steps; the first part, (25), is an immediate consequence of Lemmas 11 and 12 below. The proofs of these results can be found in Appendix A.2.

First step: Relating COMP to exponential moment of T_n . The following result follows from a straightforward generalization of an argument of Opper and Haussler (1999):

Lemma 11 *Take arbitrary \mathcal{F} and fix arbitrary $f_0 \in \mathcal{F}$. Then:*

$$\text{COMP}_\eta(\mathcal{F}) \leq \frac{1}{\eta} \log \mathbb{E}_{Z^n \sim Q_{f_0}} [e^{\eta T_n}]. \quad (27)$$

Second step: Bounding exponential moment of T_n . It remains to bound $\mathbb{E}[e^{\eta T_n}]$. The next lemma does this by leveraging Talagrand’s inequality.

Lemma 12 *Suppose that \mathcal{F} has $L_2(P)$ diameter at most ϵ . Recall that $\sigma = e L \epsilon$ and $\mathcal{G} := \{\ell_{f_0} - \ell_f : f \in \mathcal{F}\}$. Then*

$$\mathbb{E}_{Z^n \sim Q_{f_0}} [e^{\eta T_n}] \leq \exp\left(3\eta \mathbb{E}_{Z^n \sim Q_{f_0}} [T_n] + n\eta^2 \sigma^2\right) \quad (28)$$

We note that Opper and Haussler (1999) obtained a result similar in form to (28) but under the considerably stronger assumption that the original class has finite sup-norm entropy and, consequently, that the class \mathcal{G}_k has sup-norm radius at most $O(\epsilon)$.

Inequality (25) of Theorem 10 now follows. Proving the second part, inequality (26), is based on the following standard result from empirical process theory.

Lemma 13

$$\mathbb{E}_{Z^n \sim Q_{f_0}} [T_n] \leq 2n \mathbb{E}_{Z^n \sim Q_{f_0}} [\mathcal{R}_n(\mathcal{G})].$$

This lemma is proved in Appendix A.2 for completeness.

To make the bounds in Theorem 10 useful for general \mathcal{F} with possibly large $L_2(P)$ diameter, we first decompose $\text{COMP}_\eta(\mathcal{F})$ in terms of the $L_2(P)$ covering numbers at some small, optimally-tuned resolution ϵ and the maximal complexity among all Voronoi cells induced by the cover, as in (4). We then use existing bounds on H -local complexity and Rademacher complexity to get sharp bounds on $\text{COMP}(\mathcal{F})$ in terms of covering numbers.

To this end, let \mathcal{F} be arbitrary and let $\{f_1, f_2, \dots, f_{N_\epsilon}\}$ form an $(\epsilon/2)$ -net for \mathcal{F} in the $L_2(P)$ pseudometric, with $N_\epsilon := \mathcal{N}(\mathcal{F}, L_2(P), \epsilon/2)$, and let $\mathcal{F}_{\epsilon,1}, \dots, \mathcal{F}_{\epsilon,N_\epsilon}$ be the corresponding partition of \mathcal{F} into Voronoi cells according to the $L_2(P)$ pseudometric. That is, for each $k \in [N_\epsilon]$, the Voronoi

cell $\mathcal{F}_{\varepsilon,k}$ is defined as $\{f \in \mathcal{F} : k = \arg \min_{i \in [N_\varepsilon]} \|f - f_i\|_{L_2(P)}\}$.¹ Clearly, each cell $\mathcal{F}_{\varepsilon,k}$ has $L_2(P)$ diameter at most ε . For each k , fix an arbitrary $f_k \in \mathcal{F}_{\varepsilon,k}$ and let $T_n^{(k)}$ be defined as T_n above with f_k in the role of f_0 . Inequality (5) immediately gives the following corollary of Theorem 10.

Corollary 14 *Let $\sigma := eL\varepsilon$ and, for each $k \in [N_\varepsilon]$, define the loss class $\mathcal{G}_k := \{\ell_{f_k} - \ell_f : f \in \mathcal{F}_{\varepsilon,k}\}$. Then*

$$\text{COMP}_\eta(\mathcal{F}) \leq \frac{\log N_\varepsilon}{\eta} + \max_{k \in [N_\varepsilon]} \left(3 \mathbb{E}_{Z^n \sim Q_{f_k}} [T_n^{(k)}] + \eta n \sigma^2 \right) \quad (29)$$

$$\leq \frac{\log N_\varepsilon}{\eta} + \max_{k \in [N_\varepsilon]} \left(6n \mathbb{E}_{Z^n \sim Q_{f_k}} [\mathcal{R}_n(\mathcal{G}_k)] + \eta n \sigma^2 \right). \quad (30)$$

4.3 From H -local complexity and Rademacher complexity to excess risk bounds

We now show some concrete implications of our link between COMP , $\mathbb{E}[T_n]$ and \mathcal{R}_n for three types of classes: classes of VC-type, classes with polynomial empirical entropy, and sets of classifiers of polynomial $L_2(P)$ entropy with bracketing. Each of these types of classes will be defined in sequence. Let \mathcal{H} be a class of functions over a space \mathcal{S} . The class \mathcal{H} is said to be of *VC-type* if, for some $A \in (0, \infty)$ and $V > 0$, for all $\varepsilon > 0$, the empirical covering numbers of \mathcal{H} satisfy

$$\sup_{s_1, \dots, s_n \in \mathcal{S}} \mathcal{N}(\mathcal{H}, L_2(P_n), \varepsilon) \leq (A/\varepsilon)^V. \quad (31)$$

Such classes often are called parametric classes.

The class \mathcal{H} is said to have *polynomial empirical entropy* if, for some $A \in (0, \infty)$, $\rho \in (0, 1)$, for all $\varepsilon > 0$, the empirical entropy of \mathcal{H} satisfies

$$\sup_{s_1, \dots, s_n \in \mathcal{S}} \log \mathcal{N}(\mathcal{H}, L_2(P_n), \varepsilon) \leq (A/\varepsilon)^{2\rho}. \quad (32)$$

These classes are nonparametric.

We say the class \mathcal{H} has *polynomial $L_1(P)$ entropy with bracketing* if, for some $A \in (0, \infty)$, $\rho \in (0, 1)$, for all $\varepsilon > 0$, the $L_1(P)$ entropy with bracketing of \mathcal{H} satisfies

$$\log \mathcal{N}_{[\cdot, \cdot]}(\mathcal{H}, L_1(P), \varepsilon) \leq (A^2/\varepsilon)^\rho. \quad (33)$$

To obtain explicit bounds from Corollary 14, we require suitable upper bounds on either the Rademacher complexity $\mathbb{E}_{Z^n \sim Q_{f_k}} [\mathcal{R}_n(\mathcal{G}_k)]$ or directly on the H -local complexity $\mathbb{E}_{Q_{f_k}} [T_n^{(k)}]$ itself in the three cases of interest: VC-type classes, classes of polynomial empirical entropy, and sets of classifiers of polynomial entropy with bracketing.

It is simple to obtain such bounds using Dudley's entropy integral, itself a product of the well-known chaining method from empirical process theory. However, the trick here is that we would like to leverage the fact that \mathcal{G}_k has small $L_2(P)$ diameter. By making use of Talagrand's generic chaining complexity, Koltchinskii (2011) obtained bounds which improve with reductions in the $L_2(P)$ diameter. We restate simplified versions of these bounds here (see equations (3.17) and (3.19) of Koltchinskii (2011)):

In the following and all subsequent results, \lesssim indicates inequality up to multiplication by a universal constant.

1. Ties are broken arbitrarily.

Theorem 15 (Rademacher complexity bounds (Koltchinskii 2011)) *Let \mathcal{H} be a class of functions over \mathcal{Z} , and let $Q \in \Delta(\mathcal{Z})$. Let $\sup_{h \in \mathcal{H}} \mathbb{E}_{Z \sim Q}[h(Z)^2] \leq \sigma^2$ and $U := \sup_{h \in \mathcal{H}} \|h\|_\infty$. Assume that \mathcal{H} is of VC-type as in (31) with exponent V . Then, for $\sigma^2 \geq \frac{c}{n}$ (for some constant c)*

$$\mathbb{E}_{Z^n \sim Q} [\mathcal{R}_n(\mathcal{H})] \lesssim \max \left\{ \sqrt{\frac{V}{n}} \sigma \sqrt{\log \frac{A}{\sigma}}, \frac{VU}{n} \log \frac{A}{\sigma} \right\}. \quad (34)$$

Take \mathcal{H} , Q , σ , and U as before, but now assume that \mathcal{G} is of polynomial empirical entropy as in (32) with exponent ρ . Then

$$\mathbb{E}_{Z^n \sim Q} [\mathcal{R}_n(\mathcal{H})] \lesssim \max \left\{ \frac{A^\rho}{\sqrt{n}} \sigma^{1-\rho}, \frac{A^{2\rho/(\rho+1)} U^{(1-\rho)/(1+\rho)}}{n^{1/(1+\rho)}} \right\}. \quad (35)$$

For the case of classes of polynomial entropy with bracketing, we appeal to upper bounds on $\mathbb{E}_{Q_{f_k}} [T_n^{(k)}]$. If the class \mathcal{G}_k has small $L_1(Q_{f_k})$ diameter and, moreover, if it also has polynomial $L_1(Q_{f_k})$ entropy with bracketing, then Lemma A.4 of Massart and Nédélec (2006) provides precisely such a bound. Below, we present a straightforward consequence thereof.

Theorem 16 (Expected supremum bounds (Massart and Nédélec 2006)) *Let \mathcal{H} be a class of functions over \mathcal{Z} , and let $Q \in \Delta(\mathcal{Z})$. Let $\sup_{h \in \mathcal{H}} \mathbb{E}_{Z \sim Q}[|h(Z)|] \leq \sigma^2$ and $\sup_{h \in \mathcal{H}} \|h\|_\infty \leq 1$. Assume that \mathcal{H} has polynomial entropy with bracketing as in (33) with exponent ρ . Then*

$$\mathbb{E}_{Z^n \sim Q} \left[\sup_{h \in \mathcal{H}} \left\{ \frac{1}{n} \sum_{j=1}^n h(Z_j) - \mathbb{E}[h(Z)] \right\} \right] \lesssim \max \left\{ \frac{A^\rho}{\sqrt{n}} \sigma^{1-\rho}, \frac{A^{2\rho/(\rho+1)}}{n^{1/(1+\rho)}} \right\}. \quad (36)$$

The following theorem builds on Corollary 14 and nearly follows by plugging in either (34) or (35) into (30) and tuning ε in terms of n and η (which gives the VC case, (37)), and plugging in (36) into (29) and then tuning (which gives the polynomial entropy case, (38)). The remaining work is to resolve a minor discrepancy between $L_2(P)$ pseudonorms and $L_2(Q_{f_k})$ pseudonorms (or the L_1 versions thereof). This theorem will allow us to show optimal rates under Bernstein conditions.

Theorem 17 *If \mathcal{F} is of VC-type as in (31) with exponent V , then for all $\eta \in (0, 1]$,*

$$\frac{\text{COMP}_\eta(\mathcal{F})}{n} \lesssim V \log \frac{ALn}{V} \cdot n^{-1} \cdot \eta^{-1}. \quad (37)$$

If \mathcal{F} has polynomial empirical entropy as in (32) or is a set of classifiers of polynomial entropy with bracketing as in (33) with exponent ρ , then, for all $0 < \eta < 1$,

$$\frac{\text{COMP}_\eta(\mathcal{F})}{n} \lesssim (AL)^{\frac{2\rho}{1+\rho}} \cdot n^{-\frac{1}{1+\rho}} \cdot \eta^{-\frac{1-\rho}{1+\rho}}. \quad (38)$$

The proof of Theorem 17 can be found in Appendix A.3. We will now prepare for our results on the rates of ERM on a class \mathcal{F} . We considerably generalize these results in the next section, using the concept of ‘ERM-like’ estimators. For now, the reader may skip the following definition and simply equate ERM-like with ‘ERM’.

Definition 18 Consider two models $\bar{\mathcal{F}}$ and $\mathcal{F} \subseteq \bar{\mathcal{F}}$ and let \hat{f} be any deterministic estimator on the larger model $\bar{\mathcal{F}}$ and $w : \mathcal{Z}^n \rightarrow \mathbb{R}_0^+$ be a luckiness function. We say that \hat{f} is ERM-like relative to \mathcal{F} and w if for some $\tau \geq 0$, for all $\eta > 0$, all $z^n \in \mathcal{Z}^n$,

$$\text{COMP}_\eta^{\text{FULL}}(\bar{\mathcal{F}}, \hat{f}, w, z^n) \lesssim \text{COMP}_\eta(\mathcal{F}) + \frac{\tau}{\eta}. \quad (39)$$

Note that if we take $\bar{\mathcal{F}} = \mathcal{F}$ and set \hat{f} to be ERM on \mathcal{F} , then \hat{f} is indeed ERM-like with $\tau = 0$, since then $R_{\hat{f}}(z^n) \leq 0$.

In the following corollary, note that in both cases, the occurrence of the Bernstein exponent β (or κ^{-1}) is consistent with its occurrence in the simple finite \mathcal{F} setting of (24).

Corollary 19 Assume that a β -Bernstein condition holds for \mathcal{F} as in Definition 6 for some β and B , and impose assumption (A1). Define $\kappa := \beta^{-1}$. Let \hat{f} be a deterministic estimator on a model $\bar{\mathcal{F}} \supseteq \mathcal{F}$ that is ERM-like relative to \mathcal{F} and some luckiness function w .

Then, (37) further implies, taking $\gamma = \left(B \left(\frac{V \log \frac{ALn}{V}}{n} + \tau \right) \right)^{\kappa/(2\kappa-1)}$ and taking n large enough so that $\frac{V}{n} \log \frac{ALn}{V} + \tau \leq B^{1/(1-\beta)}$, that for all $z^n \in \mathcal{Z}^n$,

$$\frac{\text{COMP}_{v(\gamma)}^{\text{FULL}}(\bar{\mathcal{F}}, \hat{f}, W, z^n)}{n} + \gamma \lesssim \left(B \left(V \log \frac{ALn}{V} + \tau \right) \right)^{\frac{\kappa}{2\kappa-1}} \cdot n^{-\frac{\kappa}{2\kappa-1}}. \quad (40)$$

Analogously, under such a Bernstein condition, (38) further implies, taking $\gamma = n^{-\frac{\kappa}{2\kappa-1+\rho}}$ and assuming that $n > (2B)^{-\frac{2\kappa-1+\rho}{\kappa-1}}$, that for all $z^n \in \mathcal{Z}^n$,

$$\frac{\text{COMP}_{v(\gamma)}^{\text{FULL}}(\bar{\mathcal{F}}, \hat{f}, W, z^n)}{n} + \gamma \lesssim \left((AL)^{\frac{2\rho}{\rho+1}} + 1 + \tau \right) \cdot B^{\frac{1-\rho}{1+\rho}} n^{-\frac{\kappa}{2\kappa-1+\rho}}. \quad (41)$$

We used the notation $\kappa = \beta^{-1}$ here to make the results easier comparable to Tsybakov (2004) and Audibert (2004); the result still holds for the case $\beta = 0$ though, if we simply replace κ^{-1} by β in all exponents above; e.g. $\kappa/(2\kappa-1)$ in (40) becomes $1/(2-\kappa^{-1}) = 1/(2-\beta)$.

Proof To see (40), we begin by upper bounding $\frac{\text{COMP}_{v(\gamma)}(\mathcal{F})}{n}$ using (37) with $\eta = v(\gamma) = \min \left\{ \frac{\gamma^{1-\beta}}{B}, 1 \right\}$ (from Lemma 7). Tentatively suppose that $B\gamma^{-(1-\beta)} \geq 1$; then $v(\gamma)^{-1} \lesssim B\gamma^{-(1-\beta)}$, and hence

$$\frac{\text{COMP}_{v(\gamma)}^{\text{FULL}}(\bar{\mathcal{F}}, \hat{f}, W, z^n)}{n} + \gamma \lesssim \frac{B}{n} \left(V \log \frac{ALn}{V} + \tau \right) \gamma^{-(1-\beta)} + \gamma.$$

Tuning γ such that it is equal to the first term on the RHS above yields (40); it is simple to verify that the supposition $B\gamma^{-(1-\beta)} \geq 1$ is ensured by the constraint on n stated in the corollary.

We now prove (41). We first prove the case for \hat{f} ERM on \mathcal{F} , i.e. with $\text{COMP}(\mathcal{F})$ rather than $\text{COMP}^{\text{FULL}}(\bar{\mathcal{F}})$ on the left. Let $\gamma_n := n^{-\frac{\kappa}{2\kappa-1+\rho}}$ be the value of γ used at sample size n . We get from the definition of the Bernstein condition and Lemma 7 that

$$\eta_n := v(\gamma_n) = B^{-1} \left(n^{-\frac{\kappa}{2\kappa-1+\rho}} \right)^{(\kappa-1)/\kappa} = B^{-1} n^{-\frac{\kappa-1}{2\kappa-1+\rho}}$$

for all n for which the RHS above is at most 1. This will hold whenever $n \geq (1/B)^{\frac{2\kappa-1+\rho}{\kappa-1}}$. For such n , we will also have $\eta_n \leq 1$ and thus can apply (38), plugging in $\eta = \eta_n = v(\gamma_n)$. The result follows by simple algebra for all n larger than the given bound. To extend the result to ERM-like estimators, we note that from (39) and (38), using $\rho \geq 0$, we get

$$\frac{\text{COMP}_\eta^{\text{FULL}}(\mathcal{F}, \hat{f}, W, z^n)}{n} \lesssim \left((AL)^{\frac{2\rho}{1+\rho}} \cdot n^{-\frac{1}{1+\rho}} + \tau \right) \cdot \eta^{-\frac{1-\rho}{1+\rho}}.$$

We now proceed as above plugging in η_n into the above expression rather than (38). \blacksquare

4.4 Recovering Bounds under Bernstein Conditions for Large Classes

Let us now see how Lemma 8 can recover the known optimal rates under the Tsybakov margin condition and also the best known rates for empirical risk minimization under Bernstein-type conditions. The theorem below works for general deterministic estimators, not just ERM; the (simple) rate implications for ERM are discussed underneath the theorem. While there are other techniques that can achieve the same rates for ERM, we feel that our approach embodies a simpler analysis for the polynomial entropy case; it also leads to some new results for other estimators, which we detail in the next section.

Theorem 20 *Assume that the β -Bernstein condition holds for \mathcal{F} as in Corollary 19 and define $\kappa := \beta^{-1}$. Let \hat{f} be a deterministic estimator on $\bar{\mathcal{F}} \supseteq \mathcal{F}$ that is ERM-like relative to \mathcal{F} (with τ as in Definition 18) and some luckiness function w . First suppose \mathcal{F} is of VC-type as in (31) with exponent V . Then there is a universal constant C_1 such that for all n large enough so that $\frac{V}{n} \log \frac{ALn}{V} + \tau \leq B^{1/(1-\beta)}$, we have*

$$\mathbb{E}_{Z \sim P} \left[R_{\hat{f}}(Z) \right] \leq_{\psi_1(n)} C_1 \left(B \left(V \log \frac{ALn}{V} + \tau \right) \right)^{\frac{\kappa}{2\kappa-1}} \cdot n^{-\frac{\kappa}{2\kappa-1}}, \quad (42)$$

where $\psi_1(n) = \frac{1}{6B} \cdot \left(B \left(V \log \frac{ALn}{V} + \tau \right) \right)^{(\kappa-1)/(2\kappa-1)} n^{\kappa/(2\kappa-1)} \asymp (\log n)^{(\kappa-1)/(2\kappa-1)} \cdot n^{\kappa/(2\kappa-1)}$.

Analogously, suppose that $\mathcal{F} \subseteq \bar{\mathcal{F}}$ has polynomial empirical entropy as in (32) or is a set of classifiers of polynomial entropy with bracketing as in (33) with exponent ρ . Then there is a C_2 such that for all n large enough so that $n > (2B)^{-\frac{2\kappa-1+\rho}{\kappa-1}}$, we have

$$\mathbb{E}_{Z \sim P} \left[R_{\hat{f}}(Z) \right] \leq_{\psi_2(n)} C_2 \left[\left((AL)^{\frac{2\rho}{\rho+1}} + 1 + \tau \right) \cdot B^{\frac{1-\rho}{1+\rho}} \cdot n^{-\frac{\kappa}{2\kappa-1+\rho}} \right], \quad (43)$$

where $\psi_2(n) = \frac{1}{6} \cdot n^{\frac{\kappa+\rho}{2\kappa-1+\rho}}$.

Proof of Theorem 20 For both results, observe from Corollary 9 and the definition of $\text{COMP}^{\text{FULL}}$ that, for all $\gamma > 0$,

$$\mathbb{E}_{Z \sim P} \left[R_{\hat{f}}(Z) \right] \leq_{v(\gamma)n/6} \frac{3 \left(\text{COMP}_{v(\gamma)/2}^{\text{FULL}}(\mathcal{F}, \hat{f}) \right)}{n} + 4\gamma. \quad (44)$$

The result now follows by plugging in (40) and (41). \blacksquare

Theorem 20 combined with part (ii) of Proposition 3 implies that, with probability at least $1 - \delta$, ERM obtains the rate

$$n^{-\frac{\kappa}{2\kappa-1+\rho}} + n^{-\frac{\kappa+\rho}{2\kappa-1+\rho}} \cdot \log \frac{1}{\delta}.$$

For sets of classifiers of polynomial entropy with bracketing, the rate $n^{-\kappa/(2\kappa-1+\rho)}$ is known to be optimal and, in particular, matches the results of Tsybakov (2004) (see Theorem 1), Audibert (2004) (see the discussion after Theorem 3.3), and Koltchinskii (2006, p. 36). Outside the realm of classification, for classes with polynomial empirical entropy, the rate we obtain is to our knowledge the best known for ERM. In particular, if these nonparametric classes are convex and the loss is exp-concave, then $\kappa = \beta = 1$, and the rates we obtain for ERM are known to be minimax optimal (Rakhlin et al., 2017, Theorem 7). We note, however, that there are cases where $\beta < 1$ and yet, by using an aggregation scheme, one can obtain a rate as if $\beta = 1$; one such example is in the case of squared loss with a non-convex class (Rakhlin et al., 2017; Liang et al., 2015).

5. Properties and Applications of COMP

Here we provide two applications of the developments in this paper, emphasizing that we do not just recover existing results but also generate new ones. We first provide some general properties of $\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n)$ that will be used in the applications.

The first property concerns data-independent complexities, i.e. based on constant luckiness functions. Then for any randomized estimator, just as in the deterministic case, $\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n)$ is upper bounded by the maximal complexity $\text{COMP}(\mathcal{F})$:

Proposition 21 *Let $\hat{\Pi}$ be an arbitrary randomized estimator and let $w(f, z^n) = 1$ for all $f \in \mathcal{F}$, $z^n \in \mathcal{Z}^n$. Then:*

$$\sup_{z^n \in \mathcal{Z}^n} \text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) \leq \text{COMP}(\mathcal{F}). \quad (45)$$

The second property is about a decomposition of COMP that vastly generalizes (5). Let \mathcal{K} be a countable set and consider a partition $\{\mathcal{F}_k \mid k \in \mathcal{K}\}$ of the model \mathcal{F} . Suppose we have an estimator $\hat{\Pi}$ for \mathcal{F} . This induces conditional estimators $\hat{\Pi}_{|k}$ for each \mathcal{F}_k in the following way: for each k and $z^n \in \mathcal{Z}^n$, $\hat{\Pi}_{|k, z^n}$ is a distribution on \mathcal{F}_k ; if $\hat{\Pi}(f \in \mathcal{F}_k \mid z^n) > 0$, then $\hat{\Pi}_{|k, z^n}$ is the distribution of f according to $\hat{\Pi}$ conditioned on both data z^n and $f \in \mathcal{F}_k$; for z^n with $\hat{\Pi}(f \in \mathcal{F}_k \mid z^n) = 0$, $\hat{\Pi}_{|k, z^n}$ can be set to an arbitrary distribution on \mathcal{F}_k (the choice will not affect the results). Thus, $\hat{\Pi}_{|k}$ is now well-defined as an estimator that maps \mathcal{Z}^n into the set of distributions on \mathcal{F}_k . For all $k \in \mathcal{K}$, let w_k be a luckiness function on $\mathcal{Z}^n \times \mathcal{F}_k$. The following proposition shows that, for arbitrary such luckiness functions w_k and sub-models \mathcal{F}_k , we can construct an overall luckiness function w such that the complexity $\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n)$ may be decomposed into the sub-complexities $\text{COMP}(\mathcal{F}_k, \hat{\Pi}_{|k}, w_k, z^n)$.

Formally, for each $z^n \in \mathcal{Z}^n$, let $\hat{\Pi}_{\mathcal{K}} \mid z^n$ be the marginal distribution on \mathcal{K} with probability mass function $\hat{\pi}_{\mathcal{K}} \mid z^n$, induced by $\hat{\Pi}$, i.e. $\hat{\pi}_{\mathcal{K}}(k \mid z^n) = \hat{\Pi}(f \in \mathcal{F}_k \mid z^n)$. Let $\Pi_{\mathcal{K}}$ be a prior probability measure on \mathcal{K} with probability mass function $\pi_{\mathcal{K}}$.

Proposition 22 *With the definitions above, for $k \in \mathcal{K}$, $f \in \mathcal{F}_k$, $z^n \in \mathcal{Z}^n$, set the global luckiness function w to:*

$$w(z^n, f) := w_k(z^n, f) \cdot \frac{\pi_{\mathcal{K}}(k)}{\hat{\pi}_{\mathcal{K}}(k \mid z^n)}. \quad (46)$$

Then for each $z^n \in \mathcal{Z}^n$:

$$\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) \leq \frac{\text{KL}(\hat{\Pi}_{\mathcal{K}} | z^n \| \Pi_{\mathcal{K}})}{\eta} + \mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\text{COMP}(\mathcal{F}_{\underline{k}}, \hat{\Pi}_{|\underline{k}}, w_{\underline{k}}, z^n) \right]. \quad (47)$$

We note that this result generalizes two of our previous observations. First, it generalizes (15) for the case of countable \mathcal{F} : the second term of (15) can be retrieved by taking all the \mathcal{F}_k to be singletons and $w_k \equiv 1$; note that the second term of (47) then vanishes. Second, it generalizes (5), which can be retrieved by taking \mathcal{K} finite, taking $\hat{\Pi}$ to be an arbitrary deterministic estimator, setting the w_k to be equal to 1 and then applying Proposition 21.

5.1 Two-Part MDL Estimator Achieving Optimal Rate

Consider a given (finite or countable) partition $\{\mathcal{F}_k : k \in \mathcal{K}\}$ of our model \mathcal{F} and let \hat{f}_k represent ERM within \mathcal{F}_k . Let, for each k , $\overline{\text{COMP}}(\mathcal{F}_k, \hat{f}_k)$ be any number larger than or equal to $\text{COMP}(\mathcal{F}_k, \hat{f}_k)$.

Fix a ‘prior’ distribution $\pi_{\mathcal{K}}$ on \mathcal{K} . Suppose that there exists a $k^* \in \mathcal{K}$ that achieves excess risk $\inf_{k \in \mathcal{K}} \inf_{f \in \mathcal{F}_k} \mathbb{E}[R_f(Z)]$; if there’s more than one \mathcal{F}_k achieving the minimum, take k^* to be the one with the largest mass $\pi_{\mathcal{K}}(k)$; further ties can be resolved arbitrarily. f^* denotes the risk minimizer within \mathcal{F}_{k^*} which, consistently with earlier notation, is then also the risk minimizer within \mathcal{F} .

Consider a two-stage deterministic estimator \check{f} that proceeds by first selecting $\check{k} \equiv \check{k}_{|z^n} \in \mathcal{K}$ based on data Z^n and then uses the ERM $\hat{f}_{\check{k}}$ within $\mathcal{F}_{\check{k}}$. From Proposition 22, we see that, by choosing the luckiness functions w and w_k appropriately (all w_k are set to 1 for all $f \in \mathcal{F}_k, z^n \in \mathcal{Z}^n$), we get, with $\check{f}_{|z^n} := \hat{f}_{\check{k}_{|z^n}}$ denoting the $f \in \mathcal{F}_{\check{k}_{|z^n}}$ selected by the estimator \check{f} upon observing Z^n , that

$$\begin{aligned} \text{COMP}_{\eta}^{\text{FULL}}(\mathcal{F}, \check{f}, w, z^n) &\leq - \sum_{i=1}^n \ell_{f^*}(z_i) + \\ &\sum_{i=1}^n \ell_{\hat{f}_{\check{k}_{|z^n}}}(z_i) + \frac{-\log \pi_{\mathcal{K}}(\check{k}_{|z^n})}{\eta} + \overline{\text{COMP}}_{\eta}(\mathcal{F}_{\check{k}}, \hat{f}_{\check{k}}). \end{aligned} \quad (48)$$

For reasons to become clear, we may call the particular choice of estimator \check{k} that is defined as minimizing the right-hand side of the second line of (48) the η -generalized MDL estimator. For this estimator we must then further have, for all $\eta > 0$, using that $\sum_{i=1}^n \ell_{\hat{f}_{\check{k}_{|z^n}}}(z_i) + \frac{-\log \pi_{\mathcal{K}}(\check{k}_{|z^n})}{\eta} \leq \sum_{i=1}^n \ell_{\hat{f}_{k^*|z^n}}(z_i) + \frac{-\log \pi_{\mathcal{K}}(k^*)}{\eta}$, that

$$\text{COMP}_{\eta}^{\text{FULL}}(\mathcal{F}, \check{f}, W, z^n) \leq \frac{-\log \pi_{\mathcal{K}}(k^*)}{\eta} + \overline{\text{COMP}}_{\eta}(\mathcal{F}_{k^*}, \hat{f}_{k^*}). \quad (49)$$

Let us first consider the case that $\overline{\text{COMP}}_{\eta}(\mathcal{F}_{k^*}, \hat{f}_{k^*}) = \text{COMP}_{\eta}(\mathcal{F}_{k^*}, \hat{f}_{k^*})$. In that case, for each $\eta > 0$, the η -generalized MDL estimator for \mathcal{F} has essentially the same complexity bound as does ERM within the optimal submodel \mathcal{F}_{k^*} ; indeed it is ERM-like according to Definition 18. Thus, by Theorem 20, if a β -Bernstein condition holds for \mathcal{F}_{k^*} , the two-part MDL estimator achieves the same rate as ERM within the optimal subclass \mathcal{F}_{k^*} if we choose $\eta = v(\gamma)$ with γ set to the same value as we would for ERM relative to the submodel \mathcal{F}_{k^*} . Two-part $v(\gamma)$ -MDL thus serves as an optimal model selection criterion, and this holds even if the number of alternatives \mathcal{F}_k considered is infinite. However, even setting aside computational issues, there are two obstacles to applying such

an MDL principle in practice: first, for each fixed η , $\text{COMP}_\eta(\mathcal{F}_{k^*}, \hat{f}_{k^*})$ depends on the unknown distribution P , and second, the desired η , i.e. $\eta = v(\gamma)$ cannot be calculated since it depends on the unknown β in the Bernstein condition (and hence also on P).

The first obstacle is overcome if, for each fixed η , we base the two-part estimate \hat{f} on upper bounds of $\overline{\text{COMP}}_\eta(\mathcal{F}_k, \hat{f}_k)$ that can be calculated for each k without knowing P . If, for example, in the polynomial entropy case, for each k we plug in the upper bound (38) with \mathcal{F} set to \mathcal{F}_k (which can in principle be calculated), then 2-part MDL will still achieve the optimal rate once we use the right η ; similarly if \mathcal{F} is a VC-class and we plug in, for each k , (37).

As for the second obstacle, the optimal values of η are induced by the best β for which a β -Bernstein condition holds; in practice, one can learn it from the data using an algorithm such as the ‘safe Bayes’ algorithm of Grünwald (2012).

Remarkably, this model selection estimator has, for each fixed η , an interpretation as minimizing a 2-part codelength of the data: in the first part, one encodes a model index k (using the code with lengths $-\log \pi_{\mathcal{K}}(k)$; each prior induces such a code by Kraft’s inequality) and in the second part, one encodes the data using the NML code, i.e. the optimal universal code relative to \mathcal{F}_k , and one picks the k minimizing the total codelength. In fact, exactly this minimization, for the case of $\eta = 1$ and log-loss, was suggested by Rissanen (1996) in the context of his MDL Principle, and has been much applied since under the name ‘refined MDL’ (Grünwald, 2007). Rissanen suggested this method simply because, viewed as a coding strategy, it led to small codelengths (cumulative log loss) of the data, and gave no frequentist justification in terms of convergence rates; we have just shown that, with a correctly set η , optimal rates for ERM within the optimal subclass can be recovered.

For the log-loss case with $\eta = 1$, we get $\text{COMP}(\mathcal{F}_k, \hat{f}_k) = \log S(\mathcal{F}_k, \hat{f}_k)$ and (2), so the refined MDL estimator will pick the k minimizing

$$\sum_{i=1}^n \ell_{\hat{f}_{k|Z^n}}(Z_i) - \log \pi_{\mathcal{K}}(k) + \log \int_{\mathcal{Z}^n} \hat{f}_{|z^n}(z^n) d\nu(z^n), \quad (50)$$

and thus avoids the problem of COMP being uncomputable without knowledge of P .

5.2 Penalized ERM Bounds: Lasso, Ridge and Luckiness NML

Consider a *penalization function* $\Gamma : \mathcal{F} \rightarrow \mathbb{R}$. Let \hat{f} be the penalized empirical risk minimizer defined as

$$\hat{f}_{|z^n} := \arg \min_{f \in \mathcal{F}} \sum_{i=1}^n \ell_f(z_i) + \eta^{-1} \cdot \Gamma(f) = \arg \min_{f \in \mathcal{F}} R_f(z^n) + \eta^{-1} \cdot \Gamma(f), \quad (51)$$

with ties resolved arbitrarily; we assume that the minimum is always achieved. Obviously, successful estimation procedures such as the lasso (with multiplier $\lambda := \eta^{-1}$) and ridge regression (Hastie et al., 2001) can be expressed in this form. In this subsection we show how our results give tight annealed excess risk bounds for such estimators for arbitrary penalization functions Γ ; these can then be turned into real excess risk bounds using Corollary 9. The main interest of this fact is that neither of the two already pre-existing specializations of COMP, i.e. PAC-Bayesian information complexity and Rademacher-type complexity, can easily handle penalized ERM-type methods. In contrast, we can

simply define the luckiness function $w(z^n) = \exp(-\Gamma(\hat{f}_{|z^n}))$. We can then write $\text{COMP}_\eta^{\text{FULL}}$ as

$$\begin{aligned} \text{COMP}_\eta^{\text{FULL}}(\mathcal{F}, \hat{f}, W, z^n) &= R_{\hat{f}_{|z^n}}(z^n) + \frac{1}{\eta} \cdot \left(\Gamma(\hat{f}_{|z^n}) + \log S_\eta(\mathcal{F}, \hat{f}, w) \right) \\ &= \min_{f \in \mathcal{F}} \left(R_f(z^n) + \frac{\Gamma(f)}{\eta} \right) + \frac{1}{\eta} \log S_\eta(\mathcal{F}, \hat{f}, w). \end{aligned} \quad (52)$$

with

$$S_\eta(\mathcal{F}, \hat{f}, w) = \mathbb{E}_{Z^n \sim P} \left[\frac{e^{-(\eta R_{\hat{f}_{|Z^n}}(Z^n) + \Gamma(\hat{f}_{|Z^n}))}}{C(\hat{f}_{|Z^n})} \right] = \int q_{\hat{f}_{|z^n}}(z^n) w(\hat{f}_{|z^n}) d\nu(z^n),$$

We can now use Corollary 9 of Theorem 4 once again to get actual excess risk bounds under the v -central condition (recall that the β -Bernstein condition implies v -central with $v(\gamma) \asymp \gamma^{1-\beta}$), plugging in the above expression (52). We can expect this risk bound to be tight since (52) is really an equality, and Corollary 9, the link between annealed and actual excess risk for a given v -central condition, is also tight up to constant factors. Of course, to make such a risk bound insightful we would have to further bound $\log S_\eta$, in a manner similar as was done for ERM-like estimators in Theorem 17. Penalized empirical risk methods such as the lasso have been thoroughly studied over the last fifteen years, and we do not yet know whether the approach we just sketched will lead to new results; our goal here is mainly to show that penalized ERM and generalized Bayesian (randomized) estimators can both be analyzed using the same technique, which bounds annealed risk in terms of cumulative log-loss differences.

Cumulative Log-Loss Bounds — Luckiness Regret If the original loss function ℓ is log-loss and we take $\eta = 1$, then we can interpret the penalized estimator (51) in terms of ‘minimax luckiness regret’, which features prominently in recent papers on sequential individual sequence prediction with log-loss such as (Kakade et al., 2006; Bartlett et al., 2013), with the ‘luckiness’ terminology introduced by Grünwald (2007): for arbitrary probability densities (sequential log-loss prediction strategies) r on \mathcal{Z}^n , we define the *luckiness regret* of r on z^n with *slack function* Γ relative to set of densities $\{p_f : f \in \mathcal{F}\}$ as

$$-\log r(z^n) - \min_{f \in \mathcal{F}} \left(-\log p_f(z^n) + \Gamma(f) \right), \quad (53)$$

i.e. the difference between the log-loss of r and the log-loss achieved by the Γ -penalized predictor \hat{f} which minimizes the penalized loss in hindsight. Now, if we take luckiness function $w(z^n) := \exp(-\Gamma(\hat{f}(z^n)))$ and we take as r in (53) the density r_w for the penalized estimator \hat{f} as in definition (10) (note that $p_f = q_f$, since we work with log-loss), then from (53) and that definition we get that for each z^n , the luckiness regret of r_w on z^n is given by

$$\begin{aligned} &-\log r_w(z^n) - \min_{f \in \mathcal{F}} \left(-\log p_f(z^n) + \Gamma(f) \right) = \\ &-\log p_{\hat{f}_{|z^n}}(z^n) + \Gamma(\hat{f}_{|z^n}) + \log S(\mathcal{F}, \hat{f}, w) - \min_{f \in \mathcal{F}} \left(-\log p_f(z^n) + \Gamma(f) \right) = \log S(\mathcal{F}, \hat{f}, w), \end{aligned}$$

so that the luckiness regret of r_w is constant over z^n . r_w is thus an equalizer strategy, and, as explained by Grünwald (2007), this implies that r_w minimizes, over all probability densities r , the maximum, over all $z^n \in \mathcal{Z}^n$, of (53), thus achieving the *minimax luckiness regret*. This minimax luckiness regret is then also equal to $\log S(\mathcal{F}, \hat{f}, w)$.

Reconsidering Chatterjee and Barron (2014) Our first main result, the annealed risk convergence bound of Theorem 4, when specialized to log-loss and $\eta = 1/2$, implies a classic result of Barron and Cover (1991) that gives nonasymptotic Hellinger convergence rates for two-part MDL estimators for well-specified models, implying that (two-part) data compression implies learning. Such two-part MDL estimators invariably work with a countable discretization of the parameter space. Chatterjee and Barron (2014) sought to use those bounds to prove convergence at the right rate of Lasso-type estimators in a Gaussian regression setting, showing that the ℓ_1 -penalization can be linked to a minimization over a discretized grid of parameter values that allows it to be related to two-part MDL so that the Barron-Cover result can be used to prove rates of convergence. The present development suggests that this can perhaps be done much more generally — there is no need to consider only two-part codes or a probabilistic setting: every Γ -penalized estimator \hat{f} for every bounded loss function defines a corresponding density r_w with $w(x^n) = \exp(-\Gamma(\hat{f}))$, and hence a code with lengths $-\log r_w(z^n)$ in terms of which one can prove an excess risk bound via Theorem 4 and Corollary 9.

6. Discussion and Future Work

Our strategy for controlling $\text{COMP}(\mathcal{F})$ owes much to an ingenious argument of Opper and Haussler (1999). They analyzed the minimax regret in the individual sequence prediction setting with log loss, where the class of comparators is the set of static experts (i.e. experts that predict according to the same distribution in each round). Cesa-Bianchi and Lugosi (2001) obtain bounds in the more general setting where the comparator class consists of *arbitrary* experts that can predict conditionally on the past (for a further considerable extension within the realm of log loss, see Rakhlin and Sridharan (2015) who use sequential complexities). Whereas the works of Opper and Haussler (1999) and Cesa-Bianchi and Lugosi (2001) both operate under some kind of bounded L_∞ metric entropy (the metric entropy in the latter work differs due to the experts' sequential nature), the present paper operates under the much weaker assumption of bounded L_2 metric entropy. We note, however, that unlike the non-i.i.d. setting of Cesa-Bianchi and Lugosi (2001), the present paper is restricted to the i.i.d./static experts setting. Yet, the extension to general losses we introduce appears to be completely new.

Theorem 20 offers a distribution-dependent bound whose derivation we view as simpler than similar bounds based on local Rademacher complexities. In particular, the strategy adopted in the present paper completely avoids complicated (at least in the view of the authors) fixed point equations that have been used to obtain good excess risk bounds in other works (such as Koltchinskii and Panchenko (1999); Bartlett et al. (2005); Koltchinskii (2006)). In the case of classes of VC-type, one can obtain optimal rates by decoupling the optimization of the parameters ε and γ ; thus, one can obtain a suitable bound on COMP without considering γ , leading to a rather easy tuning problem. In the case of larger classes of polynomial empirical entropy or sets of classifiers of polynomial entropy with bracketing, while γ and ε must be tuned jointly to obtain optimal rates, we have shown that an optimal tuning can be obtained without great effort.

We note, however, that the bounds in the present paper lack the kind of data-dependence exhibited by previous works leveraging local Rademacher complexities. Indeed, the bound in Theorem 20 is an exact oracle inequality which is distribution-dependent and, consequently, is not computable by a practitioner who does not know the β for which a Bernstein condition holds. In contrast, bounds obtained via local Rademacher complexities can be computed without distributional knowledge and

have been shown to behave like the correct (but unknown to the practitioner) distribution-dependent bounds asymptotically (see Theorem 4.2 of [Bartlett et al. \(2005\)](#)).

Yet, the present work gives rise to results which allow a different kind of data-dependence: a PAC-Bayesian improvement for situations when the posterior distribution is close to a prior distribution. This improvement (which is also algorithm-dependent) is already apparent from the simplified setting of Proposition 22 in which one places a prior over submodels, and we expect that much more can be accomplished by using Theorem 4 as a starting point.

Theorem 4 is also related to the main results of [Audibert and Bousquet \(2007\)](#), who provide bounds on the excess risk for bounded loss functions that can involve the generic chaining technique of Fernique and [Talagrand \(2014\)](#); this technique generalizes the standard chaining technique of Dudley and can lead to smaller complexities in some cases. To discuss the connection, first note that, as far as we know, the standard chaining technique is used at some point in *all* approaches that achieve optimal rates for polynomial entropy classes under Tsybakov or Bernstein conditions, although this sometimes remains hidden². In our approach, chaining remains completely under the hood, but as mentioned earlier it is present in the proof of Koltchinskii’s (2011) result linking Rademacher complexities to empirical entropy. Like we do, [Audibert and Bousquet \(2007\)](#) provides bounds on excess risk that allow for the use of priors, that can exploit Bernstein conditions, and that lead to optimal rates for large classes. However, whereas in our work chaining remains under the hood, their analogue of our ‘complexity’ (the right-hand side of their deviation bound) involves chaining explicitly, replacing the KL-term by an infinite sum over (roots of) KL terms. This makes it possible to design partitions of \mathcal{F} and priors thereon that allow one to use generic chaining. On the other hand, they directly bound the excess risk — there is no ‘annealed’ step in between and hence no direct analogue of Theorem 4 either — so that it is not clear whether their approach lends itself to the relatively easy fixed-point-free tuning that is possible using our approach; also, the Shtarkov integral and hence the connection to minimax log-loss regret does not appear in their work, making the two approaches somewhat orthogonal.

Thus, the aforementioned works go beyond our work in that they either allow data-dependent analogues of Rademacher complexities (turning oracle bounds into empirical bounds) or allow one to use generic chaining; it is at this point unclear (and an interesting open problem) whether our approach can be extended in these directions. We stress, however, that these papers make no connection between excess risk and NML complexity nor between NML complexity and Rademacher complexities; these connections are, as far as we know, completely new.

The recently developed notion of offset Rademacher complexity provides a powerful alternative to analyses based on local Rademacher complexities. [Liang et al. \(2015\)](#) introduced offset Rademacher complexities for the i.i.d. statistical learning setting to obtain faster rates under squared loss with unbounded noise (and hence unbounded loss); their bounds hold for Audibert’s star estimator ([Audibert, 2008](#)) — an aggregation method — and obtain faster rates even in non-convex situations. The techniques of the present paper, while for general loss functions, notably do not currently handle unbounded losses nor do they leverage aggregation; in light of this latter trait, the rates obtained by Theorem 20 in the case of squared loss with non-convex classes are not minimax optimal as ERM itself fails to be an optimal procedure ([Juditsky et al., 2008](#)). On the other hand, the rate provided by Theorem 20 is known to tightly characterize the performance of ERM in a number of situations,

2. For example, the proofs of [Tsybakov \(2004\)](#) are based on various results of [van de Geer \(2000, Chapter 5\)](#) which are in turn based on chaining. We also note that if one strengthens the Bernstein condition to a two-sided version then, with 0/1-loss, one can avoid chaining, see [Audibert \(2004\)](#).

and it is unclear (to the authors) how to recover such results for ERM from the offset Rademacher complexity-based analysis of [Liang et al. \(2015\)](#).

[Zhivotovskiy and Hanneke \(2016\)](#) use a combination of offset Rademacher complexities with a shifted empirical process to obtain tight bounds for ERM for the case of classification with VC classes under Massart's noise condition. While in this setting our bounds are not as tight as those of [Zhivotovskiy and Hanneke \(2016\)](#), our analysis applies to the case of general noise, general losses, and large classes. We note that in the case of classification and bounded noise, existing lower bounds imply that classes of infinite VC dimension fail to be learnable.

Glossary

Notation	Description	Page
Losses		
$R_f(z)$	Excess loss, $\ell_f(z) - \ell_{f^*}(z)$	5
$R_f(z^n)$	Cumulative excess loss, $\ell_f(z^n) - \ell_{f^*}(z^n)$	5
$q_f(z)$	Entropified version of f , $\frac{p(z) \cdot e^{-\eta R_f(z)}}{\mathbb{E}_{Z \sim P} [e^{-\eta R_f(Z)}]}$	5
Notation		
$U \triangleq_{\eta} U'$	Exponential stochastic inequality, $\mathbb{E}_{U, U' \sim P} [e^{\eta(U-U')}] \leq 1$	10
$\mathbb{E}^{\text{ANN}, \eta} [U]$	Annealed expectation, $-\frac{1}{\eta} \log \mathbb{E} [e^{-\eta U}]$	10
Complexities		
$C(f)$	Normalization constant for Shtarkov integral, $\mathbb{E}_{Z^n \sim P} [e^{-\eta R_f(Z^n)}]$	6
$S(\mathcal{F}; \hat{f})$	Shtarkov integral (deterministic version), $\mathbb{E}_{Z^n \sim P} \left[\frac{e^{-\eta R_{\hat{f} Z^n}(Z^n)}}{C(\hat{f} Z^n)} \right]$	6
$S(\mathcal{F})$	Maximal Shtarkov integral (deterministic version), $\int_{Z^n} \sup_{f \in \mathcal{F}} q_f(z^n) d\nu(z^n)$	6
$S(\mathcal{F}; \hat{f}, w)$	Generalized Shtarkov integral (deterministic version) $\mathbb{E}_{Z^n \sim P} \left[\frac{e^{-\eta R_{\hat{f} Z^n}(Z^n)}}{C(\hat{f} Z^n)} \cdot w(Z^n) \right] = \int_{Z^n} q_{\hat{f} z^n}(z^n) w(z^n) d\nu(z^n)$	7
$S(\mathcal{F}, \hat{\Pi}, w)$	Generalized Shtarkov integral, $\mathbb{E}_{Z^n \sim P} \left[\exp \left(-\mathbb{E}_{\underline{f} \sim \hat{\Pi} Z^n} [\eta R_{\underline{f}}(Z^n) + \log C(\underline{f}) - \log w(Z^n, \underline{f})] \right) \right]$	8
$\text{COMP}(\mathcal{F}, \hat{f})$	Complexity, $\eta^{-1} \log S(\mathcal{F}, \hat{f})$	6
$\text{COMP}(\mathcal{F}, \hat{f}, w, z^n)$	Generalized complexity (deterministic version), $\frac{1}{\eta} \left(-\log w(z^n) + \log S(\mathcal{F}, \hat{f}, w) \right)$	8
$\text{COMP}(\mathcal{F})$	Maximal complexity, $\eta^{-1} \log S(\mathcal{F}) = \sup_{\hat{f}} \text{COMP}(\mathcal{F}, \hat{f})$	6
$\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n)$	Generalized complexity, $\frac{1}{\eta} \cdot \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi} z^n} \left[-\log w(z^n, \underline{f}) \right] + \log S(\mathcal{F}, \hat{\Pi}, w) \right)$	8
$\text{COMP}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, z^n)$	Full generalized complexity, $\text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) + \mathbb{E}_{\underline{f} \sim \hat{\Pi} z^n} [R_{\underline{f}}(z^n)]$	8
T_n	$\sup_{f \in \mathcal{F}} \left\{ \sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) - \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) \right] \right\}$	14
$\mathbb{E}_{Z^n \sim Q_{f_0}} [T_n]$	H -local complexity	15
$\mathcal{N}(\mathcal{H}, \ \cdot\ , \varepsilon)$	ε -covering number for \mathcal{H} in the norm $\ \cdot\ $	14
$\mathcal{N}_{[\cdot]}(\mathcal{H}, \ \cdot\ , \varepsilon)$	ε -bracketing number for \mathcal{H} in the norm $\ \cdot\ $	14
$\mathcal{R}_n(\mathcal{H} S_1, \dots, S_n)$	Empirical Rademacher complexity, $\mathbb{E}_{\varepsilon_1, \dots, \varepsilon_n} \left[\sup_{h \in \mathcal{H}} \left \frac{1}{n} \sum_{i=1}^n \varepsilon_i h(S_i) \right \right]$	14
$\mathcal{R}_n(\mathcal{H})$	Rademacher complexity, $\mathbb{E} \left[\sup_{h \in \mathcal{H}} \left \frac{1}{n} \sum_{i=1}^n \varepsilon_i h(S_i) \right \right]$	14

References

- Jean-Yves Audibert. *PAC-Bayesian statistical learning theory*. PhD thesis, Université Paris VI, 2004.
- Jean-Yves Audibert. Progressive mixture rules are deviation suboptimal. In *Advances in Neural Information Processing Systems*, pages 41–48, 2008.
- Jean-Yves Audibert. Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, 37(4):1591–1646, 2009.
- Jean-Yves Audibert and Olivier Bousquet. Combining PAC-Bayesian and generic chaining bounds. *The Journal of Machine Learning Research*, 8:863–889, 2007.
- A. Barron, J. Rissanen, and B. Yu. The minimum description length principle in coding and modeling. *IEEE Trans. Inf. Theory*, 44(6):2743–2760, 1998. Special Commemorative Issue: Information Theory: 1948-1998.
- Andrew R. Barron and Thomas M. Cover. Minimum complexity density estimation. *IEEE Transactions on Information Theory*, 37(4):1034–1054, 1991.
- Peter Bartlett, Peter Grünwald, Peter Harremoës, Fares Hedayati, and Wojciech Kotłowski. Horizon-independent optimal prediction with log-loss in exponential families. In *Conference on Learning Theory*, pages 639–661, 2013.
- Peter L. Bartlett, Olivier Bousquet, and Shahar Mendelson. Local Rademacher complexities. *The Annals of Statistics*, 33(4):1497–1537, 2005.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Olivier Bousquet. A Bennett concentration inequality and its application to suprema of empirical processes. *Comptes Rendus Mathématique*, 334(6):495–500, 2002.
- Olivier Catoni. *PAC-Bayesian Supervised Classification*. Lecture Notes-Monograph Series. IMS, 2007.
- Nicolò Cesa-Bianchi and Gábor Lugosi. Worst-case bounds for the logarithmic loss of predictors. *Machine Learning*, 43(3):247–264, 2001.
- Sabyasachi Chatterjee and Andrew Barron. Information theoretic validity of penalized likelihood. In *Information Theory (ISIT), 2014 IEEE International Symposium on*, pages 3027–3031. IEEE, 2014.
- Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Peter Grünwald. The safe Bayesian. In *International Conference on Algorithmic Learning Theory*, pages 169–183. Springer, 2012.

- Peter D. Grünwald. Viewing all models as “probabilistic”. In *Proceedings of the Twelfth ACM Conference on Computational Learning Theory (COLT’ 99)*, pages 171–182, 1999.
- Peter D. Grünwald. *The Minimum Description Length Principle*. MIT Press, Cambridge, MA, 2007.
- Peter D. Grünwald and Nishant A. Mehta. Fast rates with unbounded losses. *arXiv preprint 1605.00252*, 2016.
- Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer Verlag, 2001.
- David Haussler. Sphere packing numbers for subsets of the boolean n-cube with bounded Vapnik-Chervonenkis dimension. *Journal of Combinatorial Theory, Series A*, 69(2):217–232, 1995.
- Anatoli Juditsky, Philippe Rigollet, and Alexandre B Tsybakov. Learning by mirror averaging. *The Annals of Statistics*, 36(5):2183–2206, 2008.
- Sham M. Kakade, Matthias W. Seeger, and Dean P. Foster. Worst-case bounds for Gaussian process models. In *Proceedings of the 2005 Neural Information Processing Systems Conference (NIPS 2005)*, 2006.
- Vladimir Koltchinskii. Local Rademacher complexities and oracle inequalities in risk minimization. *The Annals of Statistics*, 34(6):2593–2656, 2006.
- Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: École D’Été de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.
- Vladimir Koltchinskii and Dmitry Panchenko. Rademacher processes and bounding the risk of function learning. In *High Dimensional Probability II*, 1999.
- Tengyuan Liang, Alexander Rakhlin, and Karthik Sridharan. Learning with square loss: Localization through offset Rademacher complexity. In *Proceedings of The 27th Conference on Learning Theory (COLT 2015)*, pages 1260–1285, 2015.
- Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.
- Nishant A. Mehta. Fast rates with high probability in exp-concave statistical learning. In *Artificial Intelligence and Statistics*, pages 1085–1093, 2017.
- Manfred Opper and David Haussler. Worst case prediction over sequences under log loss. In *The Mathematics of Information Coding, Extraction and Distribution*, pages 81–90. Springer, 1999.
- Alexander Rakhlin and Karthik Sridharan. Sequential probability assignment with binary alphabets and large classes of experts. *arXiv preprint arXiv:1501.07340*, 2015.
- Alexander Rakhlin, Karthik Sridharan, and Alexandre B Tsybakov. Empirical entropy, minimax regret and minimax risk. *Bernoulli*, 23(2):789–824, 2017.

- Jorma Rissanen. *Stochastic Complexity in Statistical Inquiry*. World Scientific, Hackensack, NJ, 1989.
- Jorma Rissanen. Fisher information and stochastic complexity. *IEEE Trans. Inf. Theory*, 42(1): 40–47, 1996.
- Teemu Roos. Informal remark, 2016. Remarks made during the discussion of J. Shawe-Taylor’s invited talk at the WITMSE 2016 conference.
- Yu. M. Shtarkov. Universal sequential coding of single messages. *Problems of Information Transmission*, 23(3):3–17, 1987.
- Michel Talagrand. *Upper and lower bounds for stochastic processes: Modern methods and classical problems*, volume 60. Springer Science & Business Media, 2014.
- Alexander B Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.
- Sara van de Geer. *Empirical Processes in M-Estimation (Cambridge Series in Statistical and Probabilistic Mathematics)*. Cambridge University Press Cambridge, 2000.
- Tim van Erven, Peter D. Grünwald, Nishant A. Mehta, Mark D. Reid, and Robert C. Williamson. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16:1793–1861, 2015. Special issue in Memory of Alexey Chervonenkis.
- Vladimir N. Vapnik. *Statistical Learning Theory*. Wiley, New York, 1998.
- Mathukumalli Vidyasagar. *Learning and Generalization with Applications to Neural Networks*. Springer, 2002.
- Kenji Yamanishi. A decision-theoretic extension of stochastic complexity and its applications to learning. *Information Theory, IEEE Transactions on*, 44(4):1424–1439, 1998.
- V.V Yurinskii. Exponential inequalities for sums of random vectors. *Journal of multivariate analysis*, 6(4):473–499, 1976.
- Tong Zhang. From ε -entropy to KL-entropy: Analysis of minimum information complexity density estimation. *The Annals of Statistics*, 34(5):2180–2210, 2006a.
- Tong Zhang. Information-theoretic upper and lower bounds for statistical estimation. *Information Theory, IEEE Transactions on*, 52(4):1307–1321, 2006b.
- Nikita Zhivotovskiy and Steve Hanneke. Localization of VC classes: Beyond local Rademacher complexities. In *International Conference on Algorithmic Learning Theory*, pages 18–33. Springer, 2016.
- Xiaojin Zhu, Bryan R Gibson, and Timothy T Rogers. Human Rademacher complexity. In *Advances in Neural Information Processing Systems*, pages 2322–2330, 2009.

Appendix A.

This section contains proofs omitted from the main text.

A.1 Theorem 4

Proof Let us abbreviate $\text{ANN}(f) = n \mathbb{E}_{\bar{Z} \sim P}^{\text{ANN}, \eta} [R_f(\bar{Z})]$. By the definition of ESI (19) we see that the statement in the theorem is equivalent to

$$\mathbb{E}_{Z^n \sim P} \left[\exp \left(\eta \cdot \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} [\text{ANN}(\underline{f})] - \text{COMP}^{\text{FULL}}(\mathcal{F}, \hat{\Pi}, w, Z^n) \right) \right) \right] = 1. \quad (54)$$

Plugging in the definition of $\text{COMP}^{\text{FULL}}$ and then COMP , the left side can be rewritten as

$$\begin{aligned} & \mathbb{E} \left[\exp \left(\eta \cdot \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} [\text{ANN}(\underline{f}) - R_{\underline{f}}(Z^n)] - \frac{1}{\eta} \cdot \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} [-\log w(\underline{f}, Z^n)] + \log S(\mathcal{F}, \hat{\Pi}, w) \right) \right) \right) \right] = \\ & \frac{\mathbb{E} \left[\exp \left(\eta \cdot \left(\mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} [\text{ANN}(\underline{f}) - R_{\underline{f}}(Z^n)] - \frac{1}{\eta} \cdot \mathbb{E}_{\underline{f} \sim \hat{\Pi} | Z^n} [-\log w(\underline{f}, Z^n)] \right) \right) \right]}{\mathbb{E}_{Z^n \sim P} \left[\exp \left(-\mathbb{E}_{\underline{f} \sim \hat{\Pi} | z^n} [\eta R_{\underline{f}}(z^n) + \log C(\underline{f}) - \log w(z^n, \underline{f})] \right) \right]}, \end{aligned}$$

where the denominator is just the definition of S . It is thus sufficient to prove that this expression is equal to 1. But this is immediate from the definition of $C(f)$ and $\text{ANN}(\cdot)$. \blacksquare

A.2 Proof of second main result, Theorem 10

We first prove the results that imply (25) and then prove the result that implies (26).

A.2.1 PROOF OF (25)

Inequality (25) from Theorem 10 is a consequence of Lemmas 11 and 12, which we prove in turn.

Proof of Lemma 11

$$\begin{aligned} e^{\eta \cdot \text{COMP}_\eta(\mathcal{F})} &= S(\mathcal{F}) = \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\sup_{f \in \mathcal{F}} \frac{q_f(Z^n)}{q_{f_0}(Z^n)} \right] \\ &= \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\exp \left(\sup_{f \in \mathcal{F}} \log \frac{q_f(Z^n)}{q_{f_0}(Z^n)} \right) \right] \\ &\leq \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\exp \left(\sup_{f \in \mathcal{F}} \left\{ \log \frac{q_f(Z^n)}{q_{f_0}(Z^n)} - \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\log \frac{q_f(Z^n)}{q_{f_0}(Z^n)} \right] \right\} \right) \right], \end{aligned}$$

where the inequality follows because the second term inside the supremum is a negative KL-divergence. Now, using the definition of Q_f and Q_{f_0} , the above is equal to

$$\mathbb{E}_{Z^n \sim Q_{f_0}} \left[\exp \left(\underbrace{\eta \sup_{f \in \mathcal{F}} \left\{ \sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) - \mathbb{E}_{Z^n \sim Q_{f_0}} \left[\sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) \right] \right\}}_{T_n} \right) \right].$$

■

It remains to prove Lemma 12.

Proof of Lemma 12 First, from our assumption on the loss and $\eta \leq 1$ together imply that

$$\sup_{f, g \in \mathcal{F}} \text{ess sup} \left\{ \eta (\ell_f(Z) - \ell_g(Z) - \mathbb{E}[\ell_f(Z) - \ell_g(Z)]) \right\} \leq 1.$$

Our goal now is to be able to apply Talagrand's inequality. To this end, observe that

$$\sup_{f, g \in \mathcal{F}} \text{Var}[\eta (\ell_f(Z) - \ell_g(Z) - \mathbb{E}[\ell_f(Z) - \ell_g(Z)])] \leq \eta^2 \sup_{f, g \in \mathcal{F}} \|\ell_f - \ell_g\|_{L_2(Q_{f_0})}^2.$$

Now, if \mathcal{F} had a small $L_2(Q_{f_0})$ diameter, then the Lipschitzness of the loss would imply that the above term is also small. However, by assumption, the class \mathcal{F} is only known to have small $L_2(P)$ diameter (of at most ε). Lemma 23 (stated after this proof) effectively bridges the gap between these two pseudonorms, showing that

$$\sup_{f, g \in \mathcal{F}} \|\ell_f - \ell_g\|_{L_2(Q_{f_0})} \leq e L \sup_{f, g \in \mathcal{F}} \|f - g\|_{L_2(P)}, \quad (55)$$

which is then at most $e L \varepsilon = \sigma$.

Bousquet's version of Talagrand's inequality (see Theorem 2.3 of Bousquet (2002) or, for a more direct presentation, Theorem 12.5 of Boucheron et al. (2013)) now yields

$$\mathbb{E}_{Q_{f_0}} [e^{\lambda \eta T_n^{(k)}}] \leq \exp \left(\mathbb{E}_{Q_{f_0}} [\eta T_n^{(k)}] + (e^\lambda - (\lambda + 1))(n\eta^2 \sigma^2 + 2 \mathbb{E}_{Q_{f_0}} [\eta T_n^{(k)}]) \right).$$

Inequality (28) now follows by taking $\lambda = 1$. ■

The following lemma was used to control the complexity of the class \mathcal{F} .

Lemma 23 *For the supervised loss parameterization,*

$$\|\ell_f - \ell_g\|_{L_2(Q_{f_0})} \leq e \cdot L \|f - g\|_{L_2(P)}. \quad (56)$$

For the direct parameterization,

$$\|\ell_f - \ell_g\|_{L_2(Q_{f_0})} \leq e \|f - g\|_{L_2(P)}. \quad (57)$$

Proof of Lemma 23 We first prove (56), the supervised loss parameterization result. The Lipschitz assumption on the loss implies that

$$\mathbb{E}_{(X,Y) \sim Q_{f_0}} \left[\left(\ell_f(X,Y) - \ell_g(X,Y) \right)^2 \right] \leq L^2 \mathbb{E}_{X \sim Q_{f_0}} \left[\left(f(X) - g(X) \right)^2 \right].$$

Next, observe that for $\Delta(x) = \frac{q_{f_0}(x)}{p(x)}$

$$\mathbb{E}_{X \sim Q_{f_0}} \left[\left(f(X) - g(X) \right)^2 \right] = \mathbb{E}_{X \sim P} \left[\Delta(x) \left(f(X) - g(X) \right)^2 \right].$$

Since the inside of the expectation is nonnegative, it remains to upper bound $\Delta(x)$. By definition,

$$\Delta(x) = \frac{p(x) \int p(y | x) e^{-\eta R_f(x,y)} dy}{p(x) \mathbb{E}_{(\bar{X}, \bar{Y}) \sim P} \left[e^{-\eta R_f(\bar{X}, \bar{Y})} \right]} = \frac{\mathbb{E}_{Y \sim P | X=x} \left[e^{-\eta R_f(x,Y)} \right]}{\mathbb{E}_{(\bar{X}, \bar{Y}) \sim P} \left[e^{-\eta R_f(\bar{X}, \bar{Y})} \right]} \leq e^\eta \leq e,$$

since $\eta \leq 1$ and the excess loss random variable takes values in $[-1/2, 1/2]$.

We now prove the direct parameterization result (57). Observe that for $\Delta(z) = \frac{q_{f_0}(z)}{p(z)}$

$$\mathbb{E}_{Z \sim Q_{f_0}} \left[\left(\ell_f(Z) - \ell_g(Z) \right)^2 \right] = \mathbb{E}_{Z \sim P} \left[\Delta(Z) \left(f(Z) - g(Z) \right)^2 \right],$$

where we use the fact that $\ell_f = f$ for all $f \in \mathcal{F}$ in the direct parameterization. As above, it remains to upper bound $\Delta(z)$. By definition,

$$\Delta(z) = \frac{p(z) e^{-\eta R_f(z)}}{p(z) \mathbb{E}_{\bar{Z} \sim P} \left[e^{-\eta R_f(\bar{Z})} \right]} \leq e^\eta \leq e.$$

■

A.2.2 PROOF OF (26)

Inequality (26) is a consequence of (25) and a standard empirical process theory result, Lemma 13. For completeness, we provide a proof of this result below.

Proof of Lemma 13 Recall that $\mathcal{G} = \{\ell_{f_0} - \ell_f : f \in \mathcal{F}\}$, and let $\epsilon_1, \dots, \epsilon_n$ be independent Rademacher random variables. In the below, both Z^n and \bar{Z}^n are drawn from Q_{f_0} .

The following sequence of inequalities is a standard use of symmetrization from empirical process theory:

$$\begin{aligned}
& \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left\{ \sum_{j=1}^n (\ell_{f_0}(Z_j) - \ell_f(Z_j)) - \mathbb{E} \left[\sum_{j=1}^n (\ell_{f_0}(\bar{Z}_j) - \ell_f(\bar{Z}_j)) \right] \right\} \right] \\
&= \mathbb{E} \left[\sup_{g \in \mathcal{G}} \left\{ \sum_{j=1}^n g(Z_j) - \mathbb{E} \left[\sum_{j=1}^n g(\bar{Z}_j) \right] \right\} \right] \\
&\leq \mathbb{E} \left[\sup_{g \in \mathcal{G}} \sum_{j=1}^n (g(Z_j) - g(\bar{Z}_j)) \right] \\
&= \mathbb{E} \left[\sup_{g \in \mathcal{G}} \sum_{j=1}^n \epsilon_j (g(Z_j) - g(\bar{Z}_j)) \right] \\
&\leq 2 \mathbb{E} \left[\sup_{g \in \mathcal{G}} \sum_{j=1}^n \epsilon_j g(Z_j) \right] \\
&\leq 2 \mathbb{E} \left[\sup_{g \in \mathcal{G}} \left| \sum_{j=1}^n \epsilon_j g(Z_j) \right| \right].
\end{aligned}$$

■

A.3 Proof of Theorem 17

Proof of Theorem 17 Taking the results of Corollary 14 and dividing by n gives the two inequalities

$$\frac{\text{COMP}_\eta(\mathcal{F})}{n} \leq \frac{\log \mathcal{N}(\mathcal{F}, L_2(P), \varepsilon/2)}{n\eta} + \frac{3}{n} \max_{k \in [N_\varepsilon]} \mathbb{E}_{Z^n \sim Q_{f_k}} [T_n^{(k)}] + \eta\sigma^2 \quad (58)$$

and

$$\frac{\text{COMP}_\eta(\mathcal{F})}{n} \leq \frac{\log \mathcal{N}(\mathcal{F}, L_2(P), \varepsilon/2)}{n\eta} + 6 \max_{k \in [N_\varepsilon]} \mathbb{E}_{Z^n \sim Q_{f_k}} [\mathcal{R}_n(\mathcal{G}_k)] + \eta\sigma^2, \quad (59)$$

where we remind the reader that $N_\varepsilon = \mathcal{N}(\mathcal{F}, L_2(P), \varepsilon/2)$.

In the below applications of Theorems 15 and 16, we make use of the following two observations. First, from Lemma 23 (which we previously applied to yield (55)), it follows that the $L_2(Q_{f_k})$ diameter of \mathcal{G}_k is at most σ . Second, for any distribution $Q \in \Delta(\mathcal{Z})$, for all $u > 0$,

$$\mathcal{N}(\mathcal{G}_k, L_2(Q), u) = \mathcal{N}(\{\ell_f : f \in \mathcal{F}_{\varepsilon,k}\}, L_2(Q), u) \leq \mathcal{N}(\mathcal{F}_{\varepsilon,k}, L_2(Q), u/L) \quad (60)$$

and (in the case of sets of classifiers)

$$\begin{aligned}
N_{[\cdot]}(\mathcal{G}_k, L_2(Q_{f_k}), u) &= \mathcal{N}_{[\cdot]}(\{\ell_f : f \in \mathcal{F}_{\varepsilon,k}\}, L_2(Q_{f_k}), u) = \mathcal{N}_{[\cdot]}(\mathcal{F}_{\varepsilon,k}, L_2(Q_{f_k}), u) \\
&\leq \mathcal{N}_{[\cdot]}(\mathcal{F}_{\varepsilon,k}, L_2(P), u/e);
\end{aligned} \quad (61)$$

in both (60) and (61), the first equality holds because \mathcal{G}_k is a shifted version of $\{\ell_f : f \in \mathcal{F}\}$. In the case of the supervised loss parameterization, the inequality in (60) holds from the Lipschitzness of

the loss, and, in the case of the direct parameterization, the inequality is actually equality (recall that $L = 1$ in this case). The second equality of (61) holds because we only consider sets of classifiers with 0-1 loss. Lastly, the inequality in (61) is due to the 1-Lipschitzness of 0-1 loss for sets of classifiers and Lemma 23. From (60), if \mathcal{F} is a VC-type class (and hence so is $\mathcal{F}_{\varepsilon,k}$), then \mathcal{G}_k also is a VC-type class. Analogously, if \mathcal{F} has polynomial empirical entropy, the same property extends to \mathcal{G}_k . From (61), if \mathcal{F} is a class whose $L_2(P)$ entropy with bracketing is polynomial (and hence so is $\mathcal{F}_{\varepsilon,k}$), then \mathcal{G}_k is a class whose $L_2(Q_{f_k})$ entropy with bracketing is polynomial with the same exponent.

VC-type classes. First, Theorem 24 (stated after this proof) implies that, for all $u > 0$,

$$\mathcal{N}(\mathcal{F}, L_2(P), u) \leq \left(\frac{2A}{u} \right)^V.$$

Starting from (59), inequality (34) from Theorem 15 combined with (60) then implies that (coarsely using $\eta \leq 1$)

$$\begin{aligned} \frac{\text{COMP}_\eta(\mathcal{F})}{n} &\lesssim \frac{V \log \frac{4A}{\varepsilon}}{n\eta} + \max \left\{ \sqrt{\frac{V}{n}} \sigma \sqrt{\log \frac{AL}{\sigma}}, \frac{VU}{n} \log \frac{AL}{\sigma} \right\} + \eta \sigma^2 \\ &\lesssim \frac{V \log \frac{4A}{\varepsilon}}{n\eta} + \max \left\{ \sqrt{\frac{V}{n}} L \varepsilon \sqrt{\log \frac{A}{\varepsilon}}, \frac{VU}{n} \log \frac{A}{\varepsilon} \right\} + (L\varepsilon)^2. \end{aligned}$$

Finally, setting $\varepsilon = \frac{4}{L} \sqrt{\frac{V}{n}}$ yields (up to a universal multiplicative constant) the bound

$$\frac{V \log \frac{ALn}{V}}{n\eta} + \max \left\{ \frac{V}{n} \sqrt{\log \frac{ALn}{V}}, \frac{V}{n} \log \frac{ALn}{V} \right\} + \frac{V}{n} \lesssim \frac{V \log \frac{ALn}{V}}{n\eta},$$

where we used the assumption that $\eta \leq 1$. This proves (37).

Classes of polynomial empirical entropy or polynomial entropy with bracketing. The first order of business is to control $N_\varepsilon = \log \mathcal{N}(\mathcal{F}, L_2(P), \varepsilon/2)$. In the case of classes of polynomial empirical entropy, we again invoke Theorem 24 to conclude that, for all $u > 0$,

$$\log \mathcal{N}(\mathcal{F}, L_2(P), u) \leq \left(\frac{2A}{u} \right)^{2\rho}.$$

In the case of sets of classifiers of polynomial entropy with bracketing, the $L_2(P)$ entropy can be controlled by the relationship

$$\log \mathcal{N}(\mathcal{F}, L_2(P), u) \leq \log \mathcal{N}_{[\cdot]}(\mathcal{F}, L_2(P), u) = \log \mathcal{N}_{[\cdot]}(\mathcal{F}, L_1(P), u^2) \leq \left(\frac{A}{u} \right)^{2\rho}.$$

Next, for (i) classes of polynomial empirical entropy, we start from (59) and apply inequality (35) from Theorem 15 combined with (60); or (ii) for classes of polynomial entropy with bracketing, we

start from (58) and apply³ Theorem 16 combined with (61); both cases imply that, for $0 < \eta \leq 1$, using $\rho < 1$,

$$\begin{aligned} \frac{\text{COMP}_\eta(\mathcal{F})}{n} &\lesssim \frac{1}{n\eta} \left(\frac{2A}{\varepsilon} \right)^{2\rho} + \max \left\{ \frac{(AL)^\rho}{\sqrt{n}} \sigma^{1-\rho}, \frac{(AL)^{2\rho/(\rho+1)} U^{(1-\rho)/(1+\rho)}}{n^{1/(1+\rho)}} \right\} + \eta \sigma^2 \\ &\lesssim \frac{1}{n\eta} \left(\frac{A}{\varepsilon} \right)^{2\rho} + \frac{A^\rho L}{\sqrt{n}} \varepsilon^{1-\rho} + \eta^{\frac{\rho-1}{\rho+1}} \cdot \frac{(AL)^{2\rho/(\rho+1)}}{n^{1/(1+\rho)}} + \eta \cdot (L\varepsilon)^2. \end{aligned} \quad (62)$$

(the enlargement of the third term will not affect the rates, as will now become clear). We now set $\varepsilon := C_0 n^{-\frac{1}{2(1+\rho)}} \cdot \eta^{-\frac{1}{1+\rho}}$ for a constant $C_0 > 0$ to be determined later (this choice for ε was obtained by minimizing the sum of the first and second terms in the last line of (62) by setting the derivative to 0). With this choice, we get, as a very simple yet tedious calculation shows:

$$\begin{aligned} n^{-1} \eta^{-1} \varepsilon^{-2\rho} &= C_0^{-2\rho} \cdot n^{-\frac{1}{1+\rho}} \cdot \eta^{\frac{\rho-1}{\rho+1}} \\ n^{-1/2} \varepsilon^{1-\rho} &= C_0^{1-\rho} \cdot n^{-\frac{1}{1+\rho}} \cdot \eta^{\frac{\rho-1}{\rho+1}} \\ \eta \varepsilon^2 &= C_0^2 \cdot n^{-\frac{1}{1+\rho}} \cdot \eta^{\frac{\rho-1}{\rho+1}} \end{aligned}$$

so that (62) becomes

$$\frac{\text{COMP}_\eta(\mathcal{F})}{n} \lesssim C_{A,C_0,L} \cdot n^{-\frac{1}{1+\rho}} \cdot \eta^{\frac{\rho-1}{\rho+1}} \quad (63)$$

where

$$C_{A,C_0,L} = \left(\frac{2A}{C_0} \right)^{2\rho} + A^\rho L^\rho (C_0 e)^{1-\rho} + (AL)^{2\rho/(\rho+1)} + (e L C_0)^2. \quad (64)$$

Plugging in $C_0 = A^{\rho/(\rho+1)} L^{-1/(\rho+1)}$, the four terms become of the same order:

$$\begin{aligned} C_{A,C_0,L} &\lesssim \left(L^{1/(\rho+1)} A^{1-\frac{\rho}{\rho+1}} \right)^{2\rho} + L^{1-\frac{1-\rho}{1+\rho}} A^{\rho+\frac{\rho(1-\rho)}{\rho+1}} + (AL)^{2\rho/(\rho+1)} + \left(L^{1-\frac{1}{1+\rho}} A^{\frac{\rho}{\rho+1}} \right)^2 \\ &\lesssim (AL)^{2\rho/(\rho+1)}, \end{aligned}$$

and (38) follows. ■

The above proof made use of the universal $L_2(P)$ metric entropy being essentially equivalent to the universal $L_2(P_n)$ metric entropy. This result extends an analogous result of Haussler (1995) for VC classes (see Corollary 1 therein).

Theorem 24 (Extended Haussler) *Let \mathcal{F} be a class of functions over a space \mathcal{S} . Suppose that, for all $\varepsilon > 0$ and all $n \in \mathbb{N}$, there is some function $\psi: \mathbb{R}_+ \rightarrow \mathbb{N}$ such that*

$$\sup_{s_1, \dots, s_n \in \mathcal{S}} \mathcal{N}(\mathcal{F}, L_2(P_n), \varepsilon) \leq \psi(\varepsilon).$$

Then, for any probability measure $P \in \Delta(\mathcal{S})$ and any $\varepsilon > 0$,

$$\mathcal{N}(\mathcal{F}, L_2(P), \varepsilon) \leq \psi(\varepsilon/2).$$

3. Note that in classification, for any Q , the $L_1(P)$ diameter is equal to the square of the $L_2(P)$ diameter.

The proof is essentially due to Haussler with little change to the argument for the more general result.

Proof of Theorem 24 Let d be some pseudometric on \mathcal{F} . We say that $U \subset \mathcal{F}$ is ε separated if, for all $f, g \in U$, it holds that $d(f, g) > \varepsilon$. Let the ε -packing number $\mathcal{M}(\mathcal{F}, d, \varepsilon)$ be the maximal size of an ε -separated set in \mathcal{F} .

The packing numbers and covering numbers satisfy the following relationship (Vidyasagar, 2002, Lemma 2.2)

$$\mathcal{M}(\mathcal{F}, d, \varepsilon) \leq \mathcal{N}(\mathcal{F}, d, \varepsilon/2).$$

Thus, it is sufficient to bound $\mathcal{M}(\mathcal{F}, L_2(P), \varepsilon)$.

Suppose that $\mathcal{M}(\mathcal{F}, L_2(P), \varepsilon) > \mathcal{M}(\mathcal{F}, L_2(P_n), \varepsilon)$, and take U to be some ε -separated subset of \mathcal{F} in the $L_2(P)$ pseudometric of cardinality $|U| > \mathcal{M}(\mathcal{F}, L_2(P_n), \varepsilon)$.

Next, draw s_1, \dots, s_n i.i.d. from P . Since U is finite, by taking n large enough we can ensure that the event $A_{f,g}$, defined as,

$$\|f - g\|_{L_2(P_n)} = \left(\frac{1}{n} \sum_{j=1}^n (f(s_j) - g(s_j))^2 \right)^{1/2} < \varepsilon,$$

occurs with probability at most $\frac{1}{|U|^2}$. Since $\binom{|U|}{2} < |U|^2$, it follows that the probability that no event $A_{f,g}$ occurs among all $f, g \in U$ is positive. Hence, there exists a set of points s_1, \dots, s_n for which U is an ε -packing in the $L_2(P_n)$ pseudometric. But then it must be the case that $\mathcal{M}(\mathcal{F}, L_2(P_n), \varepsilon) \geq |U|$, contradicting our assumption that $|U| > \mathcal{M}(\mathcal{F}, L_2(P_n), \varepsilon)$. ■

A.4 Proofs for Section 5

A.4.1 PROOF OF PROPOSITION 21

Using $w(f, z^n) \equiv 1$, we can write:

$$\begin{aligned} \text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) &= \frac{1}{\eta} \log \mathbf{E}_{Z^n \sim P} \left[\exp \left(- \mathbf{E}_{\underline{f} \sim \hat{\Pi} | Z^n} \left[\eta R_{\underline{f}}(Z^n) + \log C(\underline{f}) \right] \right) \right] \\ &\leq \frac{1}{\eta} \log \mathbf{E}_{Z^n \sim P} \left[\sup_{f \in \mathcal{F}} \frac{e^{-\eta R_f(Z^n)}}{C(f)} \right], \end{aligned}$$

which is just $\text{COMP}(\mathcal{F})$.

A.4.2 PROOF OF PROPOSITION 22

Plugging the definition w into the definition of COMP, a sequence of straightforward rewritings gives:

$$\begin{aligned}
& \text{COMP}(\mathcal{F}, \hat{\Pi}, w, z^n) \\
&= \frac{1}{\eta} \cdot \left(\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \mathbb{E}_{\underline{f} \sim \hat{\Pi} | \underline{k}, z^n} \left[\log \frac{\hat{\pi}_{\mathcal{K}}(\underline{k} | z^n)}{\pi_{\mathcal{K}}(\underline{k}) \cdot w_{\underline{k}}(z^n, \underline{f})} \right] + \log \text{S}(\mathcal{F}, \hat{\Pi}, w) \right) \\
&= \frac{1}{\eta} \cdot \left(\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\log \frac{\hat{\pi}_{\mathcal{K}}(\underline{k} | z^n)}{\pi_{\mathcal{K}}(\underline{k})} + \mathbb{E}_{\underline{f} \sim \hat{\Pi} | \underline{k}, z^n} \left[-\log w_{\underline{k}}(z^n, \underline{f}) \right] \right] + \log \text{S}(\mathcal{F}, \hat{\Pi}, w) \right) \\
&= \frac{1}{\eta} \cdot \text{KL}(\hat{\Pi}_{\mathcal{K}} | z^n \| \Pi_{\mathcal{K}}) + \mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \mathbb{E}_{\underline{f} \sim \hat{\Pi} | \underline{k}, z^n} \left[-\log w_{\underline{k}}(z^n, \underline{f}) \right] + \frac{1}{\eta} \log \text{S}(\mathcal{F}, \hat{\Pi}, w). \quad (65)
\end{aligned}$$

If we can further show that

$$\log \text{S}(\mathcal{F}, \hat{\Pi}, w) \leq \mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\log \text{S}(\mathcal{F}_{\underline{k}}, \hat{\Pi}_{\underline{k}}, w_{\underline{k}}) \right] \quad (66)$$

then the result follows by plugging this into the last line of (65). We thus proceed to show (66). Setting

$$g(\underline{k}, z^m) = \mathbb{E}_{\underline{f} \sim \hat{\Pi} | \underline{k}, z^m} \left[\eta R_{\underline{f}} + \log C(\underline{f}) - \log w_{\underline{k}}(z^m, \underline{f}) \right],$$

we can write:

$$\begin{aligned}
\log \text{S}(\mathcal{F}, \hat{\Pi}, w) &= \log \mathbb{E}_{Z^n \sim P} \left[\exp \left(\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\log \frac{\pi_{\mathcal{K}}(\underline{k})}{\hat{\pi}_{\mathcal{K}}(\underline{k} | z^n)} \right] \right) \cdot \exp \left(-\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} [g(\underline{k}, Z^n)] \right) \right] \\
&\leq \log \mathbb{E}_{Z^n \sim P} \left[\left(\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\frac{\pi_{\mathcal{K}}(\underline{k})}{\hat{\pi}_{\mathcal{K}}(\underline{k} | z^n)} \right] \right) \cdot \exp \left(-\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} [g(\underline{k}, Z^n)] \right) \right] \\
&= \log \mathbb{E}_{Z^n \sim P} \left[\exp \left(-\mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} [g(\underline{k}, Z^n)] \right) \right] \\
&\leq \mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\log \mathbb{E}_{Z^n \sim P} \left[\exp \left(-g(\underline{k}, Z^n) \right) \right] \right] = \mathbb{E}_{\underline{k} \sim \hat{\Pi}_{\mathcal{K}} | z^n} \left[\log \text{S}(\mathcal{F}, \hat{\Pi}_{\underline{k}}, w_{\underline{k}}) \right]
\end{aligned}$$

where the first and last equalities are just definition chasing, the first inequality is Jensen's and the last inequality is Lemma 3.2. from Audibert (2009); the result follows.