

On Learnability under General Stochastic Processes

A. Philip Dawid[†], Ambuj Tewari^{‡,*}

[†] University of Cambridge

[‡] University of Michigan

ABSTRACT. Statistical learning theory under independent and identically distributed (iid) sampling and online learning theory for worst case individual sequences are two of the best developed branches of learning theory. Statistical learning under general non-iid stochastic processes is less mature. We provide two natural notions of learnability of a function class under a general stochastic process. We show that both notions are in fact equivalent to online learnability. Our results hold for both binary classification and regression.

Keywords: learnability, statistical learning, online learning, prequential principle, uniform martingale law of large numbers

MEDIA SUMMARY

This article is a contribution to “Learning Theory”, which has largely developed within the Computer Science community, but has much potential for cross-fertilisation with Theoretical Statistics. The basic framework is as follows. At successive time points we observe a signal, and wish to predict an associated outcome, using one from a given class of possible functions. Performance is measured by a specified loss, measuring the discrepancy between prediction and outcome. Learning theory aims to characterise conditions that allow us to make provably good predictions by analysing past data to determine a suitable prediction function. Two approaches have emerged, largely in isolation from each other. In the first, (iid) statistical learning theory, we assume all past and future signal-outcome pairs have the same probabilistic structure, independently. The second, online learning theory, makes no probabilistic assumptions, and aims to track the best predictor from those available, no matter what data are observed. In this work we introduce two natural ways of merging these approaches, assuming arbitrary probabilistic dependence across time. We show that both approaches are equivalent to online learnability. Our results hold for binary classification as well as regression.

* tewaria@umich.edu

1. INTRODUCTION

One of the most beautiful and best developed branches of machine learning theory is classical statistical learning theory (see the article by von Luxburg and Schölkopf (2011) for a non-technical overview and for more extensive references). However, it deals primarily with independent and identically distributed (iid) sampling of examples. There have been several attempts to deal with both dependence and non-stationarity: we discuss some of these extensions in Section 1.1. However in general the non-iid case is not as well developed as the classical iid case.

Another well developed branch of learning theory that has its own share of elegant mathematical ideas is online learning theory (the book by Cesa-Bianchi and Lugosi (2006) is an excellent if somewhat dated introduction). With roots in game theory and the area of information theory known as universal prediction of individual sequences, online learning theory, unlike statistical learning theory, does not use probabilistic foundations. It is therefore quite surprising that there are uncanny parallels between iid learning theory and online learning theory. The reader is invited to compare the statements of the fundamental theorems in these two areas (restated in this paper as Theorem 1 and Theorem 2).

Our main goal in this paper is to study learnability of a function class in the statistical setting under extremely general assumptions that do not require independence or stationarity. We first summarize the key theorems of iid and online learning in Section 3 and Section 4. Although this material is not new, we feel that the broader data science community might not be very familiar with results in online learning since it is a younger field compared to statistical learning theory. Also, presenting both iid learning and online learning results in a unified way allows us to draw parallels between the two theories and to motivate the need for novel theories that connect these two.

We propose a definition of learnability under general stochastic processes in Section 5. We show that learning under this general definition is equivalent to online learnability (Theorem 6). We give a prequential version of our main definition in Section 6. In the prequential version, as in online learning, the function output by the learning algorithm at any given time cannot peek into the future. We show that learnability under the prequential version of our general learning setting is also equivalent to online learnability (Theorem 9). We focus on the problem of binary classification for simplicity. But we also provide extensions of our equivalence results to the problem of real valued prediction (i.e., regression) in Section 7 (see Theorem 13 and Theorem 14).

1.1. Related Work. The iid assumption of statistical learning theory has been relaxed and replaced with various types of *mixing* assumptions, especially β -mixing (Mohri & Rostamizadeh, 2009; Vidyasagar, 2002). However, in this line of investigation, the stationary assumption is kept and the theory resembles the iid theory to a large extent since mixing implies approximate independence of random variables that are sufficiently separated in time. Mixing assumptions can be shown to hold for some interesting classes of processes, including some Markov and hidden Markov processes. Markov sampling has also been considered on its own as a generalization of iid sampling (Aldous & Vazirani, 1995; Gamarnik, 2003; Smale & Zhou, 2009).

There has been work on performance guarantees of specific algorithms like boosting (Lozano et al., 2006) and SVMs (Steinwart & Anghel, 2009; Steinwart et al., 2009) under non-iid assumptions. However, our focus here is not on any specific learning methodology. We would like to point out that, while we focus on *learnability* of functions in a fixed class, the question of *universal consistency* has also been studied in the context of general stochastic processes (Hanneke, 2017; Nobel, 1999).

There are a handful of papers that focus, as we do, on conditional risk given a sequence of observation drawn from a general non-iid stochastic processes (Pestov, 2010; Shalizi & Kontorovich, 2013; Zimin & Lampert, 2017). These papers focus on process decompositions: expressing a complex stochastic process as a mixture of simpler stochastic processes. For example, de Finetti’s theorem shows that exchangeable distributions are mixtures of iid distributions. The basic idea is to output a function with small expected loss one step beyond the observed sample where the expectation is also conditioned on the observed sample. While closely related, our performance measures are cumulative in nature and are inspired more by regret analysis in online learning than PAC bounds in computational learning theory.

The use of tools from online learning theory (e.g., sequential Rademacher complexity) for developing learning theory for dependent, non-stationary process was pioneered by Kuznetsov and Mohri (Kuznetsov & Mohri, 2015, 2017). However, their focus is on time series forecasting applications and therefore their performance measures always involve the expected loss of the function chosen by the learning algorithm some steps into the future (i.e., the part not seen by the learning algorithm) of the process. In contrast, our definition uses conditional distributions of the stochastic process *on the realized path* to define our performance measure. We also point out that there are earlier papers that apply learning theory tools to understand time series prediction (Alquier & Wintenberger, 2012; Meir, 2000; Modha & Masry, 1998). Some very recent work has also begun to extend some of the work on time series to processes with spatial structure and dependence such as those occurring on a network (Dagan et al., 2019).

A direct inspiration for this paper is the work of Skouras and Dawid (2000) on estimation in semi-parametric statistical models under misspecification. They highlighted that, under misspecification, M-estimators, including maximum likelihood estimators, may not converge to a deterministic limit even asymptotically. Instead, the limit can be stochastic. This is because, under misspecification, the “best” model can depend on the observed sequence of data. They gave examples showing that this can happen for non-ergodic processes or processes with long range dependencies that do not decay fast enough. Our work can be seen as a direct extension of their ideas to the learning theory setting where the focus is not on parameter estimation but on loss minimization over potentially massive function spaces.

2. PRELIMINARIES

We consider a supervised learning setting where we want to learn a mapping from an input space \mathcal{X} to an output space \mathcal{Y} . Two output spaces of interest to us in this paper are $\mathcal{Y} = \{-1, +1\}$ (binary classification) and $\mathcal{Y} = [-1, +1]$ (regression). Instead of talking about the difficulty of learning individual functions, we will define learnability for a *class* of functions that we will denote by $\mathcal{F} \subseteq \mathcal{Y}^{\mathcal{X}}$. Let $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ and let $\ell : \mathcal{Z} \times \mathcal{F} \rightarrow \mathbb{R}_+$ be a loss function mapping an input-output pair (x, y) and a function f to a non-negative loss. The set $\{1, \dots, n\}$ will be denoted by $[n]$ and we use $\mathbb{1}[C]$ to denote an indicator function that is 1 if the condition C is true and 0 otherwise. Two important loss functions are the 0-1 loss $\ell((x, y), f) = \mathbb{1}[y \neq f(x)]$ (in binary classification) and the absolute loss $\ell((x, y), f) = |y - f(x)|$ (in regression).

We often denote an input-output pair (x, y) by z . When the input-output pair is random, we will denote it by $Z = (X, Y)$, perhaps with additional time indices such as $Z_t = (X_t, Y_t)$. We will use the abbreviation $Z_{1:t}$ to denote the sequence Z_1, \dots, Z_t . A learning rule \hat{f}_n is a map from \mathcal{Z}^n to \mathcal{F} . We will abuse notation a bit and refer to the learning rule and the function output by the

learning rule both by \widehat{f}_n . An important learning rule is empirical risk minimization (ERM): given a sequence $z_{1:t}$ of input-output pairs, it outputs the function,

$$(2.1) \quad \widehat{f}_n^{\text{ERM}} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(z_t, f).$$

Note that, for infinite function classes, the minimum may not be achieved. In that case, one can work with functions achieving empirical risks that are arbitrarily close to the infimum of the empirical risk over the class \mathcal{F} .

Given a distribution P on \mathcal{Z} , the loss function can be extended as follows:

$$\ell(P, f) = \mathbb{E}_{z \sim P} [\ell(z, f)].$$

The function minimizing the expectation above is

$$f_P^* = \operatorname{argmin}_{f \in \mathcal{F}} \ell(P, f).$$

The P -regret of a function $f \in \mathcal{F}$ is defined as

$$\begin{aligned} \rho(P, f) &= \ell(P, f) - \inf_{f' \in \mathcal{F}} \ell(P, f') \\ &= \ell(P, f) - \ell(P, f_P^*). \end{aligned}$$

Note that the P -regret depends on the class \mathcal{F} but we hide this dependence when the function class is clear from the context.

3. LEARNABILITY IN THE IID SETTING

In this section we review some basic results of statistical learning theory under iid sampling. For more details the reader can consult standard texts in this area (Anthony & Bartlett, 1999; Shalev-Shwartz & Ben-David, 2014; Vidyasagar, 2002). In the standard formulation of statistical learning theory, we draw a sequence $Z_{1:n}$ of iid examples from a distribution P . That is, the joint distribution of $Z_{1:n}$ is a product distribution $\mathbf{P} = P \otimes P \otimes \dots \otimes P$. We adopt the minimax framework to define learnability of a class \mathcal{F} of functions with respect to a loss function ℓ . Define the worst case performance of a learning rule \widehat{f}_n by

$$V_n^{\text{iid}}(\widehat{f}_n, \mathcal{F}) = \sup_P \mathbb{E} [\rho(P, \widehat{f}_n)]$$

and the minimax value by

$$V_n^{\text{iid}}(\mathcal{F}) = \inf_{\widehat{f}_n} V_n^{\text{iid}}(\widehat{f}_n, \mathcal{F}).$$

For the sake of conciseness, the notation above hides the fact that \widehat{f}_n depends on the sequence $Z_{1:n}$. The expectation above is taken over the randomness in these samples.

Definition 1. We say that \mathcal{F} is learnable in the iid learning setting if

$$\limsup_{n \rightarrow \infty} V_n^{\text{iid}}(\mathcal{F}) = 0.$$

Furthermore, we say that \mathcal{F} is learnable via a sequence \widehat{f}_n of learning rules if

$$\limsup_{n \rightarrow \infty} V_n^{\text{iid}}(\widehat{f}_n, \mathcal{F}) = 0.$$

One of the major achievements of statistical learning theory was the determination of necessary and sufficient conditions for learnability of a class \mathcal{F} . Learnability in both binary classification with 0-1 loss and regression with absolute loss is known to be equivalent to a probabilistic condition, namely the uniform law of large numbers (ULLN) for the class \mathcal{F} :

$$(3.1) \quad \limsup_{n \rightarrow \infty} \sup_P \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{t=1}^n f(X_t) - Pf \right| \right] = 0.$$

Here $X_{1:n}$ are drawn iid from P and $Pf = \mathbb{E}_{X \sim P} [f(X)]$. Whether or not ULLN holds for a class \mathcal{F} depends on the finiteness of different combinatorial parameters, depending on whether we are in the binary classification or regression setting. We will discuss the binary classification case here, leaving the regression case to Section 7.

The VC dimension of \mathcal{F} , denoted by $\text{VCdim}(\mathcal{F})$, is the length n of the longest sequence $x_{1:n}$ shattered by \mathcal{F} . We say that a sequence $x_{1:n}$ is shattered by \mathcal{F} if

$$\forall \epsilon_{1:n} \in \{\pm 1\}^n, \exists f \in \mathcal{F}, \text{ s.t. } \forall t \in [n], f(x_t) = \epsilon_t.$$

Finally, we recall the definition of the (expected) Rademacher complexity of a function class with respect to a distribution P :

$$\mathfrak{R}_n(P, \mathcal{F}) = \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(X_t) \right]$$

Note that the expectation above is with respect to both $X_{1:n}$ and $\epsilon_{1:n}$. The former are drawn iid from P whereas the latter are iid $\{\pm 1\}$ -valued Rademacher (also called symmetric Bernoulli) random variables. The worst case, over P , Rademacher complexity is denoted by

$$\mathfrak{R}_n(\mathcal{F}) = \sup_P \mathfrak{R}_n(P, \mathcal{F}).$$

Theorem 1. *Consider binary classification with 0-1 loss in the iid setting. Then, the following are equivalent:*

- (1) \mathcal{F} is learnable.
- (2) \mathcal{F} is learnable via ERM.
- (3) The ULLN condition (3.1) holds for \mathcal{F} .
- (4) $\text{VCdim}(\mathcal{F}) < \infty$.
- (5) $\limsup_{n \rightarrow \infty} \mathfrak{R}_n(\mathcal{F}) = 0$.

A similar result holds for regression with absolute loss with the VC dimension condition (i.e., condition number 4 above) replaced with a similar one involving its scale-sensitive counterpart, called the fat shattering dimension (see Section 7.1 for details).

4. LEARNABILITY IN THE ONLINE SETTING

A second learning setting with a well-developed theory is the *online* learning setting, where no probabilistic assumptions are placed on the data-generating process. Compared to statistical learning theory under iid sampling, online learning theory is a younger field. The main combinatorial parameter in this area, the Littlestone dimension, was defined by Littlestone (1988). It was given the name ‘‘Littlestone dimension’’ by Ben-David et al. (2009), where it was also shown that it fully characterizes learnability in the binary classification setting. Scale-sensitive analogues of Littlestone

dimension for regression problems and the sequential version of Rademacher complexity were studied in Rakhlin et al. (2015a, 2015b).

The online learning setting takes an individual sequence approach, where results are sought that hold for every possible sequence $z_{1:n} \in \mathcal{Z}^n$ that might be encountered by the learning rule.

We consider a sequence $\widehat{f}_{0:n-1}$ of learning rules, where \widehat{f}_t takes in as input the sequence $z_{1:t}$ and outputs a (possibly random) function in \mathcal{F} . Define the expected (normalized) regret of $\widehat{f}_{0:n-1}$ on sequence $z_{1:n}$:

$$\rho^{\text{online}}(\widehat{f}_{0:n-1}, z_{1:n}) = \mathbb{E} \left[\frac{1}{n} \left(\sum_{t=1}^n \ell(z_t, \widehat{f}_{t-1}) - \inf_{f \in \mathcal{F}} \sum_{t=1}^n \ell(z_t, f) \right) \right].$$

This is similar in flavor to, but distinct from, the regret function ρ used in the iid setting. It obeys the *prequential principle* (Dawid, 1984): performance of \widehat{f}_{t-1} , which is learned using $z_{1:t-1}$, is judged using loss evaluated on z_t with no overlap between data used for learning and for performance evaluation. The expectation is needed because the learning rules $\widehat{f}_{0:n-1}$ may use internal randomization to achieve robustness to adversarial data. The regret nomenclature comes from the fact that $\widehat{f}_{0:n-1}$ cannot peek into the future to lower its loss but its cumulative performance is compared with lowest possible loss, in hindsight, over the entire sequence $z_{1:n}$. However, the comparator term has its own restriction: it uses the best fixed function f in hindsight, as opposed to the best sequence of functions.

The object of interest is now the following minimax value:

$$V_n^{\text{online}}(\mathcal{F}) = \inf_{\widehat{f}_{0:n-1}} V^{\text{online}}(\widehat{f}_{0:n-1}, \mathcal{F}),$$

where

$$V^{\text{online}}(\widehat{f}_{0:n-1}, \mathcal{F}) = \sup_{z_{1:n} \in \mathcal{Z}^n} \rho^{\text{online}}(\widehat{f}_{0:n-1}, z_{1:n})$$

is the worst-case performance of the sequence $\widehat{f}_{0:n-1}$ of learning rules, with \widehat{f}_t taking in as input the sequence $z_{1:t}$ and outputting a function in \mathcal{F} . The infimum is then taken over all such learning rule sequences.

Definition 2. We say that \mathcal{F} is learnable in the online learning setting if

$$\limsup_{n \rightarrow \infty} V_n^{\text{online}}(\mathcal{F}) = 0.$$

As in statistical learning, we have necessary and sufficient conditions for learnability that almost mirror those in Theorem 1. The ULLN condition gets replaced by the Uniform Martingale Law of Large Numbers (UMLLN). We say that UMLLN holds for \mathcal{F} if

$$(4.1) \quad \limsup_{n \rightarrow \infty} \sup_{\mathbf{P}, \mathcal{A}} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \left| \frac{1}{n} \sum_{t=1}^n (f(X_t) - \mathbb{E}[f(X_t) | \mathcal{A}_{t-1}]) \right| \right] = 0.$$

The crucial difference between the UMLLN condition and the ULLN condition is that here the supremum is taken over *all* joint distributions \mathbf{P} of $X_{1:n}$. In particular $X_{1:n}$ need not be iid. Also, to obtain a martingale structure, we use an arbitrary filtration $\mathcal{A} = (\mathcal{A}_t)_{t=0}^{n-1}$ such that X_t is \mathcal{A}_t -measurable. It is easy to see that UMLLN is a stronger condition than ULLN: simply restrict \mathbf{P} to be a product distribution and let \mathcal{A} be the natural filtration of X_t . Then the UMLLN condition reduces to the ULLN condition.

The VC dimension of \mathcal{F} is replaced by another combinatorial parameter, called the Littlestone dimension of \mathcal{F} , denoted by $\text{Ldim}(\mathcal{F})$. Before we present the definition of Littlestone dimension, we

need some notation to handle complete binary trees labeled with examples drawn from the input space \mathcal{X} . We think of a complete binary tree \mathbf{x} of depth n as defining a sequence $\mathbf{x}_t, 1 \leq t \leq n$, of maps. The map \mathbf{x}_t gives us the examples sitting at level t of the tree. For example, \mathbf{x}_1 is the root, $\mathbf{x}_2(-1)$ is the left child of the root, $\mathbf{x}_2(+1)$ is the right child of the root, and so on. In general $\mathbf{x}_t(\epsilon_{1:t-1})$ is the node at level t that we reach by following the path given by the sign sequence $\epsilon_{1:t-1} \in \{\pm 1\}^{t-1}$, where -1 means “go left” and $+1$ means “go right”. The Littlestone dimension of \mathcal{F} is the depth n of the largest complete binary tree \mathbf{x} shattered by \mathcal{F} . We say that a complete binary tree \mathbf{x} is shattered by \mathcal{F} if

$$\forall \epsilon_{1:n} \in \{\pm 1\}^n, \exists f \in \mathcal{F}, \text{ s.t. } \forall t \in [n], f(\mathbf{x}_t(\epsilon_{1:t-1})) = \epsilon_t.$$

Finally, Rademacher complexity gets replaced with its sequential analogue, called the sequential Rademacher complexity. We first define the sequential Rademacher complexity of \mathcal{F} given a tree \mathbf{x} of depth n as:

$$\mathfrak{R}_n^{\text{seq}}(\mathbf{x}, \mathcal{F}) = \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(\mathbf{x}_t(\epsilon_{1:t-1})) \right].$$

Note that the expectation above is only with respect to the Rademacher random variables $\epsilon_{1:t}$ as \mathbf{x} is a fixed tree. Taking the worst case over all complete binary trees \mathbf{x} of depth n gives us the sequential Rademacher complexity of \mathcal{F} :

$$\mathfrak{R}_n^{\text{seq}}(\mathcal{F}) = \sup_{\mathbf{x}} \mathfrak{R}_n^{\text{seq}}(\mathbf{x}, \mathcal{F}).$$

Theorem 2. *Consider binary classification with 0-1 loss in the online (individual sequence) setting. Then, the following are equivalent:*

- (1) \mathcal{F} is learnable.
- (2) The UMLLN condition (4.1) holds for \mathcal{F} .
- (3) $\text{Ldim}(\mathcal{F}) < \infty$.
- (4) $\limsup_{n \rightarrow \infty} \mathfrak{R}_n^{\text{seq}}(\mathcal{F}) = 0$.

As in the iid setting, a similar result holds for online regression with absolute loss, with the Littlestone dimension condition (i.e., condition number 3 above) replaced by a similar one involving its scale-sensitive counterpart, called the sequential fat shattering dimension (see Section 7.2 for details).

It is well known that online learnability is harder than iid learnability. That is, $\text{VCdim}(\mathcal{F}) \leq \text{Ldim}(\mathcal{F})$ for any \mathcal{F} , and the gap in this inequality can be arbitrarily large. For example, the set of threshold functions on \mathbb{R} :

$$(4.2) \quad \mathcal{F}_{\text{threshold}} = \{x \mapsto \mathbb{1}[x > \theta] : \theta \in \mathbb{R}\}$$

has $\text{VCdim}(\mathcal{F}_{\text{threshold}}) = 1$ but $\text{Ldim}(\mathcal{F}_{\text{threshold}}) = \infty$.

A conspicuous difference between Theorem 1 and Theorem 2 is the absence of the condition involving ERM. Indeed, ERM is not necessarily a good learning rule in the online setting: there exist classes learnable in the online setting that are not learnable via ERM. Unfortunately, the learning rules that learn a class \mathcal{F} in the online setting are quite complex (Ben-David et al., 2009). It is not known if there exists a rule as simple as ERM that will learn a class \mathcal{F} whenever \mathcal{F} is online learnable. In any case, ERM does not play as central a role in online learning as it does in learning in the iid setting.

5. LEARNABILITY UNDER GENERAL STOCHASTIC PROCESSES

In this section we move beyond the iid setting to cover *all* distributions, not just product distributions. For a general stochastic process \mathbf{P} , we still have an analogue of P at time t , namely

$$P_t(\cdot; z_{1:t-1}) = \mathbf{P}(\cdot | Z_{1:t-1} = z_{1:t-1}).$$

This is the conditional distribution of Z_t given $Z_{1:t-1}$. Just like P , this is unknown to the learning rule. However, unlike P in the iid case, P_t is *data-dependent*. Therefore the P_t -regret of a function $\rho(P_t, f)$ is data-dependent. We will often hide the dependence of P_t on past data $Z_{1:t-1}$. We can use the average of the P_t -regrets,

$$R_n(Z_{1:n}, f) = \frac{1}{n} \sum_{t=1}^n \rho(P_t, f)$$

as a performance measure. Note that the minimizer of this performance measure is data-dependent, unlike in the iid case. As in the iid setting a learning rule \hat{f}_n is a map from \mathcal{Z}^n to \mathcal{F} . To reduce clutter in our notation, we will continue to hide the dependence of \hat{f}_n on the realized sample $Z_{1:n}$. The value of a learning rule \hat{f}_n is now defined as

$$\begin{aligned} V_n^{\text{gen}}(\hat{f}_n, \mathcal{F}) &= \sup_{\mathbf{P}} \mathbb{E} \left[R_n(Z_{1:n}, \hat{f}_n) - \inf_{f \in \mathcal{F}} R_n(Z_{1:n}, f) \right] \\ &= \sup_{\mathbf{P}} \mathbb{E} \left[\frac{1}{n} \sum_{t=1}^n \ell(P_t, \hat{f}_n) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \right], \end{aligned}$$

where the supremum is now taken over all joint distributions \mathbf{P} over $Z_{1:n}$. This leads to consideration of the following minimax value to define learnability:

$$V_n^{\text{gen}}(\mathcal{F}) = \inf_{\hat{f}_n} V_n^{\text{gen}}(\hat{f}_n, \mathcal{F}).$$

Definition 3. We say that \mathcal{F} is process learnable if

$$\limsup_{n \rightarrow \infty} V_n^{\text{gen}}(\mathcal{F}) = 0.$$

Furthermore, we say that \mathcal{F} is process learnable via a sequence \hat{f}_n of learning rules if

$$\limsup_{n \rightarrow \infty} V_n^{\text{gen}}(\hat{f}_n, \mathcal{F}) = 0.$$

Note that in the iid case, when \mathbf{P} is a product distribution with marginal P , we have $P_t = P$ for all t and therefore, for any f ,

$$\frac{1}{n} \sum_{t=1}^n \ell(P_t, f) = \ell(P, f).$$

We have the following result as an immediate consequence.

Lemma 3. Fix any loss function ℓ and function class \mathcal{F} . For any learning rule \hat{f}_n , $V^{\text{gen}}(\hat{f}_n, \mathcal{F}) \geq V^{\text{iid}}(\hat{f}_n, \mathcal{F})$. This also means that $V_n^{\text{gen}}(\mathcal{F}) \geq V_n^{\text{iid}}(\mathcal{F})$.

The result above is not surprising: process learnability has to be harder than iid learnability. However, somewhat surprisingly, we can show that process learnability is at least as hard as online learnability.

Theorem 4. Consider binary classification with 0-1 loss in the general stochastic process setting. Suppose the class \mathcal{F} is not online learnable, i.e., $Ldim(\mathcal{F}) = \infty$. Then for any $n \geq 1$, $V_n^{gen}(\mathcal{F}) \geq 1/8$. Therefore, the class \mathcal{F} is not process learnable.

To complement the lower bound above, we will now give a performance guarantee for ERM in the general stochastic process setting. Given a loss ℓ and function class \mathcal{F} , define the loss class $\ell \circ \mathcal{F}$ as

$$\ell \circ \mathcal{F} = \{z \mapsto \ell(z, f) : f \in \mathcal{F}\}.$$

We define the sequential Rademacher complexity of a loss class $\ell \circ \mathcal{F}$ as

$$\begin{aligned} \mathfrak{R}_n^{\text{seq}}(\mathbf{z}, \ell \circ \mathcal{F}) &= \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \ell(\mathbf{z}_t(\epsilon_{1:t-1}), f) \right], \\ \mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) &= \sup_{\mathbf{z}} \mathfrak{R}_n^{\text{seq}}(\mathbf{z}, \ell \circ \mathcal{F}). \end{aligned}$$

Note that the supremum here is over \mathcal{Z} -valued trees that are labeled with input-output pairs. It is easy for us to connect the complexity to the loss class to the complexity of the underlying function class for a simple loss function like the 0-1 loss (see Appendix A for details.)

Theorem 5. Fix any loss function ℓ and function class \mathcal{F} . Let $\widehat{f}_n^{\text{ERM}}$ denote the ERM learning rule defined in (2.1). Then we have

$$V_n^{gen}(\mathcal{F}) \leq V_n^{gen}(\widehat{f}_n^{\text{ERM}}, \mathcal{F}) \leq 4\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}).$$

Proof. The first inequality is true by definition of $V_n^{gen}(\mathcal{F})$. So we just have to prove the second one.

Note, by definition of $\widehat{f}_n^{\text{ERM}}$,

$$\frac{1}{n} \sum_{t=1}^n \ell(Z_t, \widehat{f}_n^{\text{ERM}}) = \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f).$$

Therefore, we have

$$\begin{aligned} & R_n(Z_{1:n}, \widehat{f}_n^{\text{ERM}}) - \inf_{f \in \mathcal{F}} R_n(Z_{1:n}, f) \\ &= \frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_n^{\text{ERM}}) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \\ &= \frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_n^{\text{ERM}}) - \frac{1}{n} \sum_{t=1}^n \ell(Z_t, \widehat{f}_n^{\text{ERM}}) + \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \\ (5.1) \quad & \leq \sup_{f \in \mathcal{F}} \frac{1}{n} \left(\sum_{t=1}^n \ell(P_t, f) - \ell(Z_t, f) \right) + \sup_{f \in \mathcal{F}} \frac{1}{n} \left(\sum_{t=1}^n \ell(Z_t, f) - \ell(P_t, f) \right). \end{aligned}$$

The justification for the last inequality is as follows. First, we know that $\widehat{f}_n^{\text{ERM}} \in \mathcal{F}$. Second, when $\inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f)$ is achieved, at f^* say, we have,

$$\begin{aligned} & \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \\ & \leq \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f^*) - \frac{1}{n} \sum_{t=1}^n \ell(P_t, f^*) \\ & \leq \sup_f \frac{1}{n} \left(\sum_{t=1}^n \ell(Z_t, f) - \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f) \right). \end{aligned}$$

Taking expectations on both sides of (5.1) gives us

$$\begin{aligned} & \mathbb{E} \left[R_n(Z_{1:n}, \widehat{f}_n^{\text{ERM}}) - \inf_{f \in \mathcal{F}} R_n(Z_{1:n}, f) \right] \\ & \leq \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \left(\sum_{t=1}^n \ell(P_t, f) - \ell(Z_t, f) \right) \right] \\ & \quad + \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \left(\sum_{t=1}^n \ell(Z_t, f) - \ell(P_t, f) \right) \right] \\ & \leq 4\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}). \end{aligned}$$

Note that the last inequality follows from Theorem 2 of Rakhlin et al. (2015b). Since the last quantity above does not depend on \mathbf{P} , we can take supremum over \mathbf{P} on both sides to finish the proof. \square

We now have everything in place to be able to show the equivalence of process learnability and online learnability. A similar result can also be shown in the regression case (see Section 7.3).

Theorem 6. *Consider binary classification with 0-1 loss. Then all of the equivalent conditions in Theorem 2 are also equivalent to:*

- \mathcal{F} is process learnable.

Proof. Theorem 4 established that learnability in the general stochastic process setting implies online learnability. For the other direction, note that according to Theorem 5 we have

$$V_n^{\text{gen}}(\mathcal{F}) \leq 4\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) \leq 2\mathfrak{R}_n^{\text{seq}}(\mathcal{F}),$$

where the second inequality follows from Theorem 16 in Appendix A. Taking lim sup of both sides as n tends to infinity shows that online learnability implies process learnability. \square

Although online learnability turns out to be equivalent to process learnability, there is an important difference between the two settings which has to do with the importance of ERM. In the former ERM is not a good learning rule whereas in the latter a learnable class is learnable via ERM. Therefore ERM continues to play a special role in the general stochastic process setting just like the iid setting.

Also note that Theorem 6 is stated in terms of learnability which is an asymptotic concept. However, the proof clearly shows that the rate of convergence is determined by the sequential Rademacher complexity of \mathcal{F} which scales as $O\left(\sqrt{\text{Ldim}(\mathcal{F})/n}\right)$ (Alon et al., 2021)

5.1. **Examples.** We end this section with some examples showing that our definition of process learnability is natural, interesting and worth studying.

IID Sampling. Let us note once again that if $\mathbf{P} = P \otimes P \otimes \dots \otimes P$ is a product measure then $\ell(P_t, f)$ is just $\ell(P, f)$ and therefore not random. In this special but important case, our definition of learnability reduces to the standard definition of learnability under iid sampling.

Asymptotically Stationary Process. Suppose that \mathbf{P} is not a product measure but the process is *asymptotically stationary* in the sense that the random probability measure $\bar{P}_n = \frac{1}{n} \sum_{t=1}^n P_t$ converges to some fixed deterministic P^* in total variation $\|\cdot\|_{TV}$ as $n \rightarrow \infty$. For a class \mathcal{F} that is learnable in the general stochastic process setting and for loss function bounded by 1, we have

$$\begin{aligned} & \mathbb{E} \left[\ell(P^*, \hat{f}_n^{\text{ERM}}) \right] - \inf_{f \in \mathcal{F}} \ell(P^*, f) \\ &= \mathbb{E} \left[\ell(P^*, \hat{f}_n^{\text{ERM}}) - \ell(P^*, f_{P^*}) \right] \\ &\leq 2 \mathbb{E} \left[\sup_{f \in \mathcal{F}} |\ell(P^*, f) - \ell(\bar{P}_n, f)| \right] \\ &\quad + \mathbb{E} \left[\ell(\bar{P}_n, \hat{f}_n^{\text{ERM}}) - \ell(\bar{P}_n, f_{P^*}) \right] \\ &\leq 2 \mathbb{E} [\|P^* - \bar{P}_n\|_{TV}] \\ &\quad + \mathbb{E} \left[\ell(\bar{P}_n, \hat{f}_n^{\text{ERM}}) - \inf_{f \in \mathcal{F}} \ell(\bar{P}_n, f) \right]. \end{aligned}$$

By the stationarity assumption, the first term on the right in the last inequality goes to zero. Moreover, the rate of convergence can often be characterized in terms of the mixing coefficients of the stochastic process (Vidyasagar, 2002). By learnability of \mathcal{F} via ERM in the general stochastic process setting, the last term goes to zero. Note that $\ell(P, f)$ is linear in P and therefore $\ell(\bar{P}_n, f) = \frac{1}{n} \sum_{t=1}^n \ell(P_t, f)$. So, under stationarity, our learnability condition implies that ERM does well when its performance is measured under the (asymptotic) stationary distribution P^* .

Mixture of IID. Consider a simple mixture of product distributions

$$\mathbf{P} = \lambda P \otimes P \otimes \dots \otimes P + (1 - \lambda) Q \otimes Q \otimes \dots \otimes Q$$

where, for simplicity, assume that P and Q have disjoint supports. Then with probability λ we have $\forall t > 1 P_t = P$, and with probability $1 - \lambda$ we have $\forall t > 1 P_t = Q$. Therefore, the minimizer of

$$(5.2) \quad \frac{1}{n} \sum_{t=1}^n \ell(P_t, f)$$

is f_P^* with probability λ and f_Q^* with probability $1 - \lambda$ (assuming, again for simplicity, that the minimizers f_P^*, f_Q^* are unique). Here, unlike the iid and stationary examples, the “best” function, even with infinite data, is not deterministic but is random depending on which mixture component was selected. Still, if learnability in our general sense holds, then ERM will do well according to the performance measure (5.2). Note that this example can be easily generalized to a mixture of more than two iid processes. It can also be generalized, with additional technical conditions, to the case when $P \neq Q$. The main difference from the disjoint support case that we consider here will be that with probability λ , P_t will converge to P (in a suitable sense) and to Q otherwise. Similarly, the minimizer of (5.2) would not equal f_P^* or f_Q^* but it would converge (again, in some appropriate

sense determined by technical conditions) to one of them with probability λ and $1 - \lambda$ respectively.

Random Level. Fix the squared loss $\ell(z, f) = (y - f(x))^2$ and consider a class \mathcal{F} that is iid learnable and closed under translations by a constant, i.e., if $f \in \mathcal{F}$ then $f + c \in \mathcal{F}$ for any constant $c \in \mathbb{R}$. Let $X_{1:n}$ be iid drawn from some distribution P_X on $\mathcal{X} \subseteq \mathbb{R}^d$ that has a density with respect to Lebesgue measure on \mathbb{R}^d . Let $Y_t = f^*(X_t) + \xi_t + \xi_0$ for some $f^* \in \mathcal{F}$ and $1 \leq t \leq n$ where $(\xi_t)_{t=0}^n$ are iid standard normal. Note that the process $Z_t = (X_t, Y_t), 1 \leq t \leq n$ is *not* iid. It is not even mixing in any sense due to long range dependence in Y_t caused by ξ_0 . Now ERM over \mathcal{F} is given by:

$$\begin{aligned} \widehat{f}_n^{\text{ERM}}(Z_{1:n}) &= \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n (f^*(X_t) + \xi_0 + \xi_t - f(X_t))^2 \\ &= \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n (f^*(X_t) + \xi_t - (f(X_t) - \xi_0))^2 \\ &= \xi_0 + \operatorname{argmin}_{g \in \mathcal{F} - \xi_0} \frac{1}{n} \sum_{t=1}^n (f^*(X_t) + \xi_t - g(X_t))^2 \\ &= \xi_0 + \operatorname{argmin}_{g \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n (f^*(X_t) + \xi_t - g(X_t))^2, \end{aligned}$$

where the last equality holds because $\mathcal{F} - \xi_0 = \mathcal{F}$ and we have assumed that all empirical minimizers are unique with probability 1. Thus, we have shown that

$$\widehat{f}_n^{\text{ERM}}(Z_{1:n}) = \widehat{f}_n^{\text{ERM}}((X_t, f^*(X_t) + \xi_t)_{t=1}^n) + \xi_0.$$

Since \mathcal{F} is iid learnable, $\widehat{f}_n^{\text{ERM}}((X_t, f^*(X_t) + \xi_t)_{t=1}^n)$ converges (in $L_2(P_X)$ sense) to the function f^* which means the ERM on $Z_{1:n}$ converges to the *random* function $f^* + \xi_0$.

Next we compute $\ell(P_t, f)$, as follows. Let P'_t be the conditional distribution of Z_t given $X_{1:t-1}$ and $\xi_{0:t-1}$. Then we have

$$\begin{aligned} \ell(P'_t, f) &= \mathbb{E} [(Y_t - f(X_t))^2 | X_{1:t-1}, \xi_{0:t-1}] \\ &= \mathbb{E} [(f^*(X_t) + \xi_t + \xi_0 - f(X_t))^2 | X_{1:t-1}, \xi_{0:t-1}] \\ &= \mathbb{E} [(f^*(X_t) + \xi_t + \xi_0 - f(X_t))^2 | \xi_0] \\ &= 1 + \|f^* - f + \xi_0\|_{L_2(P_X)}^2 \end{aligned}$$

(with ξ_0 regarded as fixed). Then $\ell(P_t, f) = \mathbb{E}[\ell(P'_t) | Z_{1:t-1}] = 1 + \mathbb{E}[\|f^* - f + \xi_0\|_{L_2(P_X)}^2 | Z_{1:t-1}]$, where now ξ_0 (only) is regarded as random. It is easy to show that the distribution of ξ_0 , given $Z_{1:t-1}$, is normal with variance $1/t$ and mean

$$U_{t-1} = \frac{\sum_{i=1}^{t-1} (Y_i - f^*(X_i))}{t}.$$

Consequently

$$\ell(P_t, f) = 1 + \frac{1}{t} + \|f^* - f + U_{t-1}\|_{L_2(P_X)}^2.$$

In particular,

$$\inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \geq 1.$$

Now, with $\widehat{f}_n^{\text{ERM}} = \widehat{f}^{\text{ERM}}(Z_{1:n})$, consider $\ell(P_t, \widehat{f}_n^{\text{ERM}})$. We have shown $f^* - \widehat{f}_n^{\text{ERM}} \rightarrow -\xi_0$ in mean square. Also,

$$U_{t-1} = \xi_0 + \frac{1}{t} \left(\sum_{i=1}^{t-1} \xi_i - \xi_0 \right) \rightarrow \xi_0$$

in mean square. So $\frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_n^{\text{ERM}}) \rightarrow 1$, the smallest possible value. That is, asymptotically the minimiser of $\frac{1}{T} \sum_{t=1}^T \ell(P_t, f)$ over \mathcal{F} is $\widehat{f}_T^{\text{ERM}}$ (which converges, not to f^* , but to the random function $f^* + \xi_0$).

6. A PREQUENTIAL DEFINITION OF LEARNABILITY

The previous section generalized the statistical setting to include non-product distributions and extended the definition of learnability to a more general setting. In this section we will generalize the online learnability definition to obtain a prequential version of learnability, while still keeping the level of generality of the previous section. As in the online setting, consider a sequence of learning rules $\widehat{f}_{0:n-1}$, where \widehat{f}_t is a function only of $Z_{1:t-1}$, i.e. it cannot peek ahead to access $Z_{t:n}$. Unlike the online learning setting, $Z_{1:t}$ is a random sequence drawn from some general distribution \mathbf{P} over \mathcal{Z}^n . Now, define the minimax value

$$V_n^{\text{preq}}(\mathcal{F}) = \inf_{\widehat{f}_{0:n-1}} V_n^{\text{preq}}(\widehat{f}_{0:n-1}, \mathcal{F}),$$

where

$$V_n^{\text{preq}}(\widehat{f}_{0:n-1}, \mathcal{F}) = \sup_{\mathbf{P}} \mathbb{E} \left[\frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_{t-1}) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \right].$$

Note that the expectation above is with respect to both \mathbf{P} and any internal randomness used by the rules $\widehat{f}_{0:n-1}$. As before, the definition of the minimax value leads to the definition of learnability.

Definition 4. We say that \mathcal{F} is prequentially learnable if

$$\limsup_{n \rightarrow \infty} V_n^{\text{preq}}(\mathcal{F}) = 0.$$

The definition of $V_n^{\text{preq}}(\mathcal{F})$ can be obtained from the definition of $V_n^{\text{gen}}(\mathcal{F})$ by replacing \widehat{f}_n , which depends on the entire sequence $Z_{1:n}$, by \widehat{f}_{t-1} , which depends only on $Z_{1:t-1}$, in the loss term that involves P_t . It can also be thought as a generalization of V^{online} because of the following. When the distribution \mathbf{P} degenerates to a point mass at a specific sequence $z_{1:n}$ then P_t becomes a point mass at z_t and the difference of cumulative losses above reduces to the individual sequence regret of $\widehat{f}_{0:n-1}$ on $z_{1:n}$. This observation immediately gives us the following result.

Lemma 7. Fix any loss function ℓ and function class \mathcal{F} . Then we have $V_n^{\text{preq}}(\mathcal{F}) \geq V_n^{\text{online}}(\mathcal{F})$.

The lemma above says that prequential learnability is at least as hard as online learnability. Our next lemma provides a converse result.

Lemma 8. Fix any loss function and function class \mathcal{F} . Then for any sequence $\widehat{f}_{0:n-1}$ of learning rules we have

$$V_n^{\text{preq}}(\widehat{f}_{0:n-1}, \mathcal{F}) \leq V_n^{\text{online}}(\widehat{f}_{0:n-1}, \mathcal{F}) + 2\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}).$$

This also means that

$$V_n^{\text{preq}}(\mathcal{F}) \leq V_n^{\text{online}}(\mathcal{F}) + 2\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}).$$

Proof. Let \mathbf{P} be an arbitrary distribution. We have the following three term decomposition:

$$\begin{aligned}
& \frac{1}{n} \sum_{t=1}^n \ell(P_t, \hat{f}_{t-1}) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \\
&= \underbrace{\frac{1}{n} \sum_{t=1}^n \ell(P_t, \hat{f}_{t-1}) - \frac{1}{n} \sum_{t=1}^n \ell(Z_t, \hat{f}_{t-1})}_{(I)} \\
& \quad + \underbrace{\frac{1}{n} \sum_{t=1}^n \ell(Z_t, \hat{f}_{t-1}) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f)}_{(II)} \\
& \quad + \underbrace{\inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(Z_t, f) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f)}_{(III)}.
\end{aligned}$$

The term (I) involves a martingale difference sequence $\ell(P_t, \hat{f}_{t-1}) - \ell(Z_t, \hat{f}_{t-1})$ and hence has expectation zero under \mathbf{P} . Term (II) is the individual sequence regret of $\hat{f}_{0:n-1}$ on the sequence $Z_{1:n}$ and hence is bounded, in expectation, by $V_n^{\text{online}}(\hat{f}_{0:n-1}, \mathcal{F})$. Term (III), in expectation, is at most,

$$\mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n (\ell(Z_t, f) - \ell(P_t, f)) \right] \leq 2\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}),$$

where the inequality again follows from Theorem 2 of Rakhlin et al. (2015b).

The lemma now follows by taking expectations on both sides of the three term decomposition above and plugging in the upper bounds for each term's expected value. \square

We now have all the ingredients to characterize sequential learnability for binary classification.

Theorem 9. *Consider binary classification with 0-1 loss. Then all of the conditions in Theorem 2 are also equivalent to:*

- \mathcal{F} is sequentially learnable.

Proof. From Lemma 7, we know that if a class is sequentially learnable then it is online learnable. In the other direction, using Lemma 8, we have

$$\begin{aligned}
V_n^{\text{preq}}(\mathcal{F}) &\leq V_n^{\text{online}}(\mathcal{F}) + 2\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) \\
&\leq V_n^{\text{online}}(\mathcal{F}) + \mathfrak{R}_n^{\text{seq}}(\mathcal{F}),
\end{aligned}$$

where the second inequality follows from Theorem 16 in Appendix A. Under any of the equivalent conditions in Theorem 2, the lim sup of both of the quantities on the right goes to zero as n tends to infinity. \square

A similar result for the regression setting can be found in Section 7.4.

7. THE REGRESSION SETTING

In this section, we provide analogues of most of the binary classification results for the regression setting with absolute loss. Note that rates of convergence can depend on the loss function but *learnability* is quite robust to changes in the loss function. For example, we can also use squared loss $\ell((x, y), f) = (y - f(x))^2$. But we will keep our focus on the absolute loss in this section.

Our organization in this section is similar to the organization of results for binary classification. Section 7.1 and Section 7.2 review known results in iid and online learning, but give them a unified presentation. Section 7.3 and Section 7.4 present new results.

7.1. Statistical Learning. The fat shattering dimension of \mathcal{F} is a scale-sensitive parameter that takes a scale $\gamma > 0$ as an argument. The fat shattering dimension of \mathcal{F} at scale γ , denoted by $\text{fat}_\gamma(\mathcal{F})$, is the length n of the longest sequence $x_{1:n}$ that is γ -shattered by \mathcal{F} . We say that a sequence $x_{1:n}$ is γ -shattered by \mathcal{F} if there exists a witness sequence $s_{1:n}$ of real numbers such that

$$\forall \epsilon_{1:n} \in \{\pm 1\}^n, \exists f \in \mathcal{F}, \text{ s.t. } \forall t \in [n], \epsilon_t(f(x_t) - s_t) \geq \gamma.$$

Theorem 10. *Consider regression with absolute loss in the iid statistical setting. Then, the following are equivalent:*

- (1) \mathcal{F} is learnable.
- (2) \mathcal{F} is learnable via ERM.
- (3) The ULLN condition (3.1) holds for \mathcal{F} .
- (4) $\forall \gamma > 0, \text{fat}_\gamma(\mathcal{F}) < \infty$.
- (5) $\limsup_{n \rightarrow \infty} \mathfrak{R}_n(\mathcal{F}) = 0$.

The first four conditions are proved to be equivalent by Alon et al. (1997). For connections between fat shattering dimension and Rademacher complexity see the work of Mendelson (2002).

7.2. Online Setting. The fat shattering dimension of \mathcal{F} is replaced by its sequential analogue, just as VC dimension gets replaced by Littlestone dimension in the case of binary classification. The sequential fat shattering dimension of \mathcal{F} at scale γ , denoted by $\text{sfat}_\gamma(\mathcal{F})$, is the depth n of the deepest tree \mathbf{x} that is γ -shattered by \mathcal{F} . We say that a complete binary tree \mathbf{x} is γ -shattered by \mathcal{F} if there exists a complete binary real valued witness tree \mathbf{s} such that

$$\forall \epsilon_{1:n} \in \{\pm 1\}^n, \exists f \in \mathcal{F}, \text{ s.t. } \forall t \in [n], \epsilon_t(f(\mathbf{x}_t(\epsilon_{1:t-1}) - \mathbf{s}_t(\epsilon_{1:t-1}))) \geq \gamma.$$

Theorem 11. *Consider regression with absolute loss in the online (individual sequence) setting. Then, the following are equivalent:*

- (1) \mathcal{F} is learnable.
- (2) The UMLLN condition (4.1) holds for \mathcal{F} .
- (3) $\forall \gamma > 0, \text{sfat}_\gamma(\mathcal{F}) < \infty$.
- (4) $\limsup_{n \rightarrow \infty} \mathfrak{R}_n^{\text{seq}}(\mathcal{F}) = 0$.

The last three conditions are shown to be equivalent in Rakhlin et al., 2015b and the connection with learnability was established in Rakhlin et al., 2015a.

As in the binary classification setting, online learnability is harder than iid statistical learnability. That is, for any \mathcal{F} and any $\gamma > 0$, $\text{fat}_\gamma(\mathcal{F}) \leq \text{sfat}_\gamma(\mathcal{F})$ and the gap in this inequality can be arbitrarily large. For example, the set \mathcal{F}_{bv} of bounded variation functions from $[0, 1]$ to $[0, 1]$ with total variation at most V , has $\text{fat}_\gamma(\mathcal{F}_{\text{bv}}) < 1 + V/\gamma$ for all $\gamma > 0$ but $\text{sfat}_\gamma(\mathcal{F}_{\text{bv}}) = \infty$ for all $\gamma > 0$.

7.3. Learnability under General Stochastic Processes. We first state an analogue of Theorem 4 for the regression setting.

Theorem 12. *Consider regression with absolute loss in the general stochastic process setting. Suppose the class \mathcal{F} is not online learnable, i.e., there exists $\gamma > 0$ such that $\text{sfat}_\gamma(\mathcal{F}) = \infty$. Then for any $n \geq 1$, $V_n^{\text{gen}}(\mathcal{F}) \geq \gamma/500$. Therefore, the class \mathcal{F} is not process learnable.*

The result above allows us to extend Theorem 6 to the regression setting.

Theorem 13. *Consider regression with absolute loss with a class \mathcal{F} . Then all of the equivalent conditions in Theorem 11 are also equivalent to:*

- \mathcal{F} is process learnable.

Proof. Theorem 12 shows that process learnability implies online learnability. For the other direction, note that according to Theorem 5 we have

$$V_n^{\text{gen}}(\mathcal{F}) \leq 4 \mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}).$$

Taking lim sup of both sides as n tends to infinity and using Theorem 17 in Appendix A to bound the right hand side gives the desired implication. \square

7.4. Learnability under the Prequential Version.

Theorem 14. *Consider regression with absolute loss with a class \mathcal{F} . Then all of the equivalent condition in Theorem 11 are also equivalent to:*

- \mathcal{F} is prequentially learnable.

Proof. From Lemma 7, we know that if a class is prequentially learnable then it is online learnable. For the other direction, using Lemma 8, we have

$$V_n^{\text{preq}}(\mathcal{F}) \leq V_n^{\text{online}}(\mathcal{F}) + 2\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}).$$

From online learnability of \mathcal{F} and Theorem 17 in Appendix A, we know that the lim sup of both the quantities on the right is at most zero as n tends to infinity, giving the desired result. \square

8. CONCLUSION

In this paper we have proposed two new definitions of learnability of a class of functions under general non-iid stochastic processes. For the first definition, we showed that learnability is equivalent to online learnability. This equivalence also holds for the second definition, which is a prequential version of the first. We also showed how to extend our results from binary classification to the regression setting.

Our work poses several interesting questions for further investigation. First, we defined learnability using expectations. It will be good to derive high probability results. Second, we ignored the issue of convergence rates for simplicity. It should be possible to extend our analysis to extract information about rates of convergence. This is because the tools from iid and online learning that we use are powerful enough to give us information about rates. Third, instead of using a normalizing factor of n , the sample size, other data dependent normalizing factors could be of interest in applications. In this context, the theory of self-normalized processes comes to mind (Peña et al., 2008). Fourth, iid learning theory has been extended to deal with privacy constraints. Starting from the seminal work of Kasiviswanathan et al. (2011) these efforts have looked at a formalization of user

privacy known as *differential privacy*. Surprisingly, recent work (Alon et al., 2019; Bun et al., 2020) has shown that iid learnability under the additional constraint of approximate differential privacy is equivalent to online learnability! It will be interesting to study learnability under general stochastic processes with additional privacy constraints on the learning algorithm. Last, but certainly not least, there is a need to connect various strands of learning theory research on non-iid processes. It is unlikely that there is a single definition of learnability that is satisfactory for all purposes. We hope we have proposed two interesting and useful ones. Comparing and contrasting various existing definitions and approaches is an important goal for future work in this area.

Disclosure Statement. The authors have no conflicts of interest to declare.

Acknowledgments. Thanks to the organizers and attendees of the Fifth Bayesian, Fiducial, and Frequentist Conference (BFF5) held from May 6-9, 2018 in Ann Arbor, MI, USA. Their feedback on a preliminary version of this work was very useful. We also thank the two HDSR reviewers who made numerous helpful suggestions to improve the paper.

Contributions. APD and AT jointly conceived the research project, carried out the research, and wrote the paper.

APPENDIX A. RELATING THE COMPLEXITY OF THE LOSS CLASS TO THE FUNCTION CLASS

We first consider the 0-1 loss function and then give a result for the absolute loss.

A.1. Zero-One Loss. The result below is essentially already known. The main ideas are present in published work (Rakhlin et al., 2011). We just present the result in a form that is immediately useful to us. But before we do that, we need a useful lemma.

Lemma 15. *For any sequence $\epsilon_{1:n}$ of iid Rademacher random variables and any fixed $\{\pm 1\}$ -valued tree \mathbf{s} , the sequence $(\epsilon_t \mathbf{s}_t(\epsilon_{1:t-1}))_{t=1}^n$ is also iid Rademacher.*

Proof. It is easy to see that the sequence of random variables we have constructed takes values in $\{\pm 1\}$. We just need to check that the distribution of $\epsilon_t \mathbf{s}_t(\epsilon_{1:t-1})$ conditioned on $\epsilon_{1:t-1}$ is a fixed distribution independent of the past. This is readily verified since

$$\mathbb{E}[\epsilon_t \mathbf{s}_t(\epsilon_{1:t-1}) | \epsilon_{1:t-1}] = \mathbf{s}_t(\epsilon_{1:t-1}) \mathbb{E}[\epsilon_t | \epsilon_{1:t-1}] = 0.$$

Therefore, we have shown that the distribution of $\epsilon_t \mathbf{s}_t(\epsilon_{1:t-1})$ conditioned on $\epsilon_{1:t-1}$ is always Rademacher (symmetric Bernoulli). \square

Now we are ready to state and prove the main result of this subsection.

Theorem 16. *Let \mathcal{F} be a binary valued function class and let ℓ be the 0-1 loss function. Then we have,*

$$\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) = \frac{1}{2} \mathfrak{R}_n^{\text{seq}}(\mathcal{F}).$$

Proof. Instead of using \mathcal{Z} -valued trees to define $\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F})$, we will use a pair \mathbf{x}, \mathbf{y} of \mathcal{X} - and \mathcal{Y} -valued trees. The equality we are trying to prove can then be written as:

$$2 \sup_{\mathbf{x}, \mathbf{y}} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t \mathbb{1}[\mathbf{y}_t(\epsilon_{1:t-1}) \neq f(\mathbf{x}_t(\epsilon_{1:t-1}))] \right] = \sup_{\mathbf{x}} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(\mathbf{x}_t(\epsilon_{1:t-1})) \right].$$

For $y_1, y_2 \in \{\pm 1\}$, we can write $2\mathbb{1}[y_1 \neq y_2]$ as $1 - y_1 y_2$. Note that sequential Rademacher complexity is not affected if the entire function classes is shifted by a constant. Therefore, the left hand side is equal to

$$\sup_{\mathbf{x}, \mathbf{y}} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n -\epsilon_t \mathbf{y}_t(\epsilon_{1:t-1}) f(\mathbf{x}_t(\epsilon_{1:t-1})) \right].$$

Now consider the $\{\pm 1\}$ -valued tree $\mathbf{s} = -\mathbf{y}$. From Lemma 15, we know the above is equal to

$$\sup_{\mathbf{x}, \mathbf{s}} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(\mathbf{x}_t(\epsilon_{1:t-1} \cdot \mathbf{s}_{1:t-1}(\epsilon))) \right],$$

where $\epsilon_{1:t-1} \cdot \mathbf{s}_{1:t-1}(\epsilon)$ denotes the sequence

$$\epsilon_1 \cdot \mathbf{s}_1, \epsilon_2 \cdot \mathbf{s}_2(\epsilon_1), \dots, \epsilon_{t-1} \cdot \mathbf{s}_{t-1}(\epsilon_{1:t-2}).$$

Define the tree \mathbf{x}' as $\mathbf{x}'(\epsilon_{1:t}) = \mathbf{x}_t(\epsilon_{1:t-1} \cdot \mathbf{s}_{1:t-1}(\epsilon))$ and note that as \mathbf{x} ranges over all \mathcal{X} -valued trees and \mathbf{s} ranges over all $\{\pm 1\}$ -valued trees, \mathbf{x}' ranges over all \mathcal{X} -valued trees. Therefore, the supremum over the pair \mathbf{x}, \mathbf{s} above can simply be written as

$$\sup_{\mathbf{x}'} \mathbb{E} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \epsilon_t f(\mathbf{x}'_t(\epsilon_{1:t-1})) \right].$$

As we noted, \mathbf{x}' ranges over all \mathcal{X} -valued trees making the above quantity the same as $\mathfrak{R}_n^{\text{seq}}(\mathcal{F})$ which finishes the proof. \square

A.2. Absolute Loss. We now consider the absolute loss case. The only property of the absolute loss used in the proof below is that it is 1-Lipschitz (in either argument provided the other one is fixed).

Theorem 17. *Let \mathcal{F} be a bounded real valued function class such that $\text{sfat}_\gamma(\mathcal{F}) < \infty$ for all $\gamma > 0$. Let ℓ be the absolute loss. Then we have*

$$\limsup_{n \rightarrow \infty} \mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) \leq 0.$$

Proof. In this proof c, C will denote universal constants that can change from line to line. From Corollary 10 of Block et al. (2021), we have

$$(A.1) \quad \mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) \leq C \cdot \inf_{\alpha \geq 0} \left(\alpha + \frac{1}{\sqrt{n}} \int_{\alpha}^1 \sqrt{\log N'(\ell \circ \mathcal{F}, \delta)} d\delta \right).$$

Here $N'(\ell \circ \mathcal{F}, \delta)$ refers to the fractional covering number of the class $\ell \circ \mathcal{F}$ as defined in Definition 6 of Block et al. (2021). From the definition of the fractional covering number and the fact that the absolute loss is 1-Lipschitz, we have $N'(\ell \circ \mathcal{F}, \delta) \leq N'(\mathcal{F}, \delta)$. Moreover, by Theorem 13 of Block et al. (2021), we have

$$N'(\mathcal{F}, \delta) \leq \left(\frac{C}{\delta} \right)^{3 \text{sfat}_{c\delta}(\mathcal{F})}.$$

Fix some $\alpha > 0$. Plugging the bound above into (A.1) gives us,

$$\mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) \leq C \cdot \left(\alpha + \frac{1}{\sqrt{n}} \int_{\alpha}^1 \sqrt{\text{sfat}_{c\delta}(\mathcal{F}) \cdot \log(1/\delta)} d\delta \right).$$

Now since $\text{sfat}_{c\delta}(\mathcal{F}) < \infty$ for all $\delta > 0$, taking limits with respect to n on both sides gives us

$$\limsup_{n \rightarrow \infty} \mathfrak{R}_n^{\text{seq}}(\ell \circ \mathcal{F}) \leq C \cdot \alpha.$$

Since $\alpha > 0$ was arbitrary this proves the result. \square

APPENDIX B. PROOFS OF LOWER BOUNDS

Proof of Theorem 4. Since the Littlestone dimension of \mathcal{F} is infinite, by Theorem 3 of Alon et al. (2019) it contains $N = 2^{2^n+1}$ thresholds. This means that there are N functions $f_1, \dots, f_N \in \mathcal{F}$ and N examples $x_1, \dots, x_N \in \mathcal{X}$ such that for all $i, j \leq N$,

$$f_j(x_i) = 1 \text{ if and only if } i \leq j.$$

Without loss of generality identify these m examples with integers 1 through N written in binary notation (with enough zero padding to the left to make the binary encoding a bit vector of length exactly $2^n + 1$) and the functions with threshold functions

$$x \mapsto \mathbb{1}[x \leq \mathbf{b}]$$

for bit vectors \mathbf{b} of length $2^n + 1$.

We will now define a stochastic process $(X_t, Y_t), 1 \leq t \leq n$ indexed by bit vectors \mathbf{b} of length 2^n . The labels Y_t will be deterministic given X_t chosen as $Y_t = \mathbb{1}[X_t \leq \mathbf{b}\mathbf{1}]$ where $\mathbf{b}\mathbf{1}$ has a 1 added at the end and is therefore of length $2^n + 1$. So we only need to define a process X_1, \dots, X_n . This is defined as follows. Let $\epsilon_{1:n}$ be iid Rademacher random variables.

- $\ell_0 = 0$
- For $t = 1$ to n
 - If $\epsilon_t = +1$: $\ell_t = \ell_{t-1} + 2^{n-t}$ else: $\ell_t = \ell_{t-1}$
 - $\mathbf{b}_t = (2^n + 1)$ -length bit vector with same ℓ_t -length prefix as \mathbf{b} , $(\ell_t + 1)$ st bit equal to 1, and rest padded with zeros
 - Output $X_t = \mathbf{b}_t$

If all ϵ_t 's turn out to be +1 (an event with probability 2^{-n}), ℓ_n can become as large as

$$2^{n-1} + 2^{n-2} + \dots + 2 + 1 = 2^n - 1$$

which still leaves two bits to add a “10” at the end. So we do have enough bits available for all possible increases of resolution. Note that the true function is a threshold at an odd integer whereas the sampled X_t 's are always even integers. Finally, note that by construction $Y_t = \mathbb{1}[X_t \leq \mathbf{b}\mathbf{1}] = \mathbb{1}[\mathbf{b}_t \leq \mathbf{b}\mathbf{1}] = \mathbf{b}[\ell_t + 1]$ where $\mathbf{b}[\ell]$ is the ℓ th bit of \mathbf{b} .

Denote the stochastic process defined above by $\mathbf{P}_{\mathbf{b}}$. Our proof will follow the probabilistic method replacing the supremum over \mathbf{P} in the definition of $V_n^{\text{gen}}(\mathcal{F})$ with an expectation over \mathbf{P} with \mathbf{P} chosen to be $\mathbf{P}_{\mathbf{b}}$ with the bit vector \mathbf{b} chosen uniformly at random. That is,

$$\begin{aligned} V_n^{\text{gen}}(\widehat{f}_n, \mathcal{F}) &= \sup_{\mathbf{P}} \mathbb{E} \left[\frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_n) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \right] \\ &\geq \mathbb{E}_{\mathbf{b}} \left[\mathbb{E}_{\mathbf{P}_{\mathbf{b}}} \left[\frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_n) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(P_t, f) \right] \right]. \end{aligned}$$

Note that for every \mathbf{b} , $Y_t = f_{\mathbf{b}}(X_t)$ for some $f_{\mathbf{b}} \in \mathcal{F}$ which means that the infimum above is zero for every \mathbf{b} . Therefore, we have

$$V_n^{\text{gen}}(\widehat{f}_n, \mathcal{F}) \geq \mathbb{E}_{\mathbf{b}} \left[\mathbb{E}_{\mathbf{P}_{\mathbf{b}}} \left[\frac{1}{n} \sum_{t=1}^n \ell(P_t, \widehat{f}_n) \right] \right].$$

Now we want to argue that $\forall t \in \{1, \dots, n\}$,

$$(B.1) \quad \mathbb{E} \left[\ell(P_t, \widehat{f}_n) \right] \geq \frac{1}{8}.$$

Note the two sources of randomness in this expectation: the random bit vector \mathbf{b} and the Rademacher random variables $\epsilon_{1:n}$. Also, note that \widehat{f}_n does not have access to \mathbf{b} but only to the observed sample

$$(X_1, Y_1), \dots, (X_n, Y_n) = (\mathbf{b}_1, \mathbf{b}[\ell_1 + 1]), \dots, (\mathbf{b}_n, \mathbf{b}[\ell_n + 1])$$

Note that P_t puts probability 1/2 each on a “high resolution” value \mathbf{b}_t^+ (corresponding to $\epsilon_t = +1$) and a “low resolution” value \mathbf{b}_t^- (corresponding to $\epsilon_t = -1$). Therefore, we have,

$$\ell(P_t, \widehat{f}_n) = \frac{1}{2} \mathbb{1} \left[\widehat{f}_n(\mathbf{b}_t^-) \neq f_{\mathbf{b}}(\mathbf{b}_t^-) \right] + \frac{1}{2} \mathbb{1} \left[\widehat{f}_n(\mathbf{b}_t^+) \neq f_{\mathbf{b}}(\mathbf{b}_t^+) \right].$$

The expectation of the quantity above can be lower bounded as,

$$(B.2) \quad \begin{aligned} \mathbb{E} \left[\ell(P_t, \widehat{f}_n) \right] &\geq \frac{1}{2} \mathbb{E} \left[\mathbb{1} \left[\widehat{f}_n(\mathbf{b}_t^+) \neq f_{\mathbf{b}}(\mathbf{b}_t^+) \right] \right] \\ &\geq \frac{1}{2} \mathbb{P}(\epsilon_t = -1) \mathbb{E} \left[\mathbb{1} \left[\widehat{f}_n(\mathbf{b}_t^+) \neq f_{\mathbf{b}}(\mathbf{b}_t^+) \mid \epsilon_t = -1 \right] \right] \\ &= \frac{1}{4} \mathbb{E} \left[\mathbb{1} \left[\widehat{f}_n(\mathbf{b}_t^+) \neq f_{\mathbf{b}}(\mathbf{b}_t^+) \mid \epsilon_t = -1 \right] \right]. \end{aligned}$$

Now note that $f_{\mathbf{b}}(\mathbf{b}_t^+)$ is simply equal to $\mathbf{b}[\ell_{t-1} + 2^{n-t} + 1]$. Further note that when $\epsilon_t = -1$, the largest that ℓ_n can become is

$$\ell_{t-1} + 2^{n-t-1} + 2^{n-t-2} + \dots + 1 = \ell_{t-1} + 2^{n-t} - 1.$$

This means that, conditioned on $\epsilon_t = -1$, the entire sample is measurable w.r.t. $\mathbf{b}[1 : \ell_{t-1} + 2^{n-t}]$. Since \mathbf{b}_t^+ is measurable w.r.t. $\mathbf{b}[1 : \ell_{t-1} + 2^{n-t}]$, we can conclude that $f_{\mathbf{b}}(\mathbf{b}_t^+)$ is independent of the sample and \mathbf{b}_t^+ conditioned on $\epsilon_t = -1$. Since the unconditional distribution of $\mathbf{b}[\ell_{t-1} + 2^{n-t} + 1]$ is uniform on $\{0, 1\}$, this implies that

$$(B.3) \quad \mathbb{E} \left[\mathbb{1} \left[\widehat{f}_n(\mathbf{b}_t^+) \neq f_{\mathbf{b}}(\mathbf{b}_t^+) \mid \epsilon_t = -1 \right] \right] \geq \frac{1}{2}$$

which along with (B.2) gives (B.1). \square

Proof of Theorem 12. Since $\text{sfat}_{\gamma}(\mathcal{F}) = \infty$, by Theorem 8 of Jung et al. (2020), \mathcal{F} contains $N = 2^{2^n + 1}$ thresholds with margin $\gamma/5$. This means that there are N functions $\tilde{f}_1, \dots, \tilde{f}_N \in \mathcal{F}$, N examples $x_1, \dots, x_N \in \mathcal{X}$ and $u, u' \in [-1, +1]$ such that $|u - u'| \geq \gamma/5$, $|\tilde{f}_j(x_i) - u| \leq \gamma/100$ if $i \leq j$, and $|\tilde{f}_j(x_i) - u'| \leq \gamma/100$ if $i > j$.

Without loss of generality assume that $u > u'$. Suppose \mathcal{F} is learnable to accuracy ϵ . We want to show that by choosing ϵ to be sufficiently small, e.g., $\epsilon = \gamma/500$, we can get a contradiction with the lower bound established in proof of Theorem 4 above. Consider the stochastic process X_t, Y_t constructed in the proof above. We transform it into the regression setting by converting the binary labels Y_t into real values \tilde{Y}_t as follows:

$$\tilde{Y}_t = \begin{cases} u & \text{if } Y_t = 1, \\ u' & \text{if } Y_t = 0. \end{cases}$$

The transformed data consisting of pairs (X_t, \tilde{Y}_t) is fed into the ϵ -accurate learner \hat{f}_n for which we have the guarantee that

$$\mathbb{E} \left[\frac{1}{n} \sum_{t=1}^n \ell(\tilde{P}_t, \hat{f}_n) - \inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(\tilde{P}_t, f) \right] \leq \epsilon$$

where ℓ is the absolute loss and \tilde{P}_t is the conditional distribution defined with respect to the transformed process X_t, \tilde{Y}_t . Note that the labels Y_t were generated using a binary threshold function f_j which can be approximated to an error within $\gamma/100$ (in supremum norm) by some \tilde{f}_j . This means that

$$\inf_{f \in \mathcal{F}} \frac{1}{n} \sum_{t=1}^n \ell(\tilde{P}_t, f) \leq \gamma/100.$$

Therefore, for some t ,

$$\mathbb{E} \left[\ell(\tilde{P}_t, \hat{f}_n) \right] \leq \epsilon + \gamma/100$$

Now consider the *binary classifier* \hat{f}'_n obtained from the real valued function \hat{f}_n as follows:

$$\hat{f}'_n(x) = \begin{cases} 1 & \text{if } \hat{f}_n(x) > (u + u')/2, \\ 0 & \text{otherwise.} \end{cases}$$

We know from (B.2) and (B.3) that \hat{f}'_n will predict the label of \mathbf{b}_t^+ incorrectly with probability at least $1/4$. Because of the $\gamma/5$ gap between u and u' this means that \hat{f}'_n incurs an absolute loss of at least $\gamma/10$ on \mathbf{b}_t^+ and its real valued label (which is either u or u') with probability at least $1/4$. Recalling that \tilde{P}_t puts probability mass $1/2$ on \mathbf{b}_t^+ , we therefore have

$$\mathbb{E} \left[\ell(\tilde{P}_t, \hat{f}'_n) \right] \geq \frac{1}{2} \cdot \frac{1}{4} \cdot \frac{\gamma}{10} = \frac{\gamma}{80}$$

which means

$$\frac{\gamma}{80} \leq \epsilon + \frac{\gamma}{100}.$$

In order for this to give us a contradiction we just need to ensure that $\epsilon < \gamma/400$. The choice $\epsilon = \gamma/500$ does that. \square

REFERENCES

- Aldous, D., & Vazirani, U. (1995). A Markovian extension of Valiant’s learning model. *Information and Computation*, 117(2), 181–186.
- Alon, N., Ben-David, S., Cesa-Bianchi, N., & Haussler, D. (1997). Scale-sensitive dimensions, uniform convergence, and learnability. *Journal of the ACM*, 44(4), 615–631.
- Alon, N., Ben-Eliezer, O., Dagan, Y., Moran, S., Naor, M., & Yogev, E. (2021). Adversarial laws of large numbers and optimal regret in online classification. *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 447–455.
- Alon, N., Livni, R., Malliaris, M., & Moran, S. (2019). Private PAC learning implies finite Littlestone dimension. *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, 852–860.
- Alquier, P., & Wintenberger, O. (2012). Model selection for weakly dependent time series forecasting. *Bernoulli*, 18(3), 883–913.
- Anthony, M., & Bartlett, P. L. (1999). *Neural Network Learning: Theoretical Foundations*. Cambridge University Press.

- Ben-David, S., Pál, D., & Shalev-Shwartz, S. (2009). Agnostic online learning. *Proceedings of the 22nd Annual Conference on Learning Theory*. <https://www.cs.mcgill.ca/~colt2009/papers/032.pdf>
- Block, A., Dagan, Y., & Rakhlin, S. (2021). Majorizing measures, sequential complexities, and online learning. *arXiv preprint arXiv:2102.01729*.
- Bun, M., Livni, R., & Moran, S. (2020). An equivalence between private classification and online prediction. *Proceedings of the 61st Annual IEEE Symposium on Foundations of Computer Science*.
- Cesa-Bianchi, N., & Lugosi, G. (2006). *Prediction, Learning, and Games*. Cambridge University Press.
- Dagan, Y., Daskalakis, C., Dikkala, N., & Jayanti, S. (2019). Learning from weakly dependent data under Dobrushin’s condition. In A. Beygelzimer & D. Hsu (Eds.), *Proceedings of the Thirty-Second Conference on Learning Theory* (pp. 914–928). PMLR.
- Dawid, A. P. (1984). Present position and potential developments: Some personal views. statistical theory. the prequential approach (with discussion). *Journal of the Royal Statistical Society: Series A (General)*, 147(2), 278–290.
- Gamarnik, D. (2003). Extension of the PAC framework to finite and countable markov chains. *IEEE Transactions on Information Theory*, 49(1), 338–345.
- Hanneke, S. (2017). Learning whenever learning is possible: Universal learning under general stochastic processes. *arXiv preprint arXiv:1706.01418*.
- Jung, Y., Kim, B., & Tewari, A. (2020). On the equivalence between online and private learnability beyond binary classification. *Advances in Neural Information Processing Systems*, 33, 16701–16710.
- Kasiviswanathan, S. P., Lee, H. K., Nissim, K., Raskhodnikova, S., & Smith, A. (2011). What can we learn privately? *SIAM Journal on Computing*, 40(3), 793–826.
- Kuznetsov, V., & Mohri, M. (2015). Learning theory and algorithms for forecasting non-stationary time series. *Advances in Neural Information Processing Systems*, 541–549.
- Kuznetsov, V., & Mohri, M. (2017). Generalization bounds for non-stationary mixing processes. *Machine Learning*, 106(1), 93–117.
- Littlestone, N. (1988). Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4), 285–318.
- Lozano, A. C., Kulkarni, S. R., & Schapire, R. E. (2006). Convergence and consistency of regularized boosting algorithms with stationary β -mixing observations. *Advances in Neural Information Processing Systems*, 819–826.
- Meir, R. (2000). Nonparametric time series prediction through adaptive model selection. *Machine Learning*, 39(1), 5–34.
- Mendelson, S. (2002). Rademacher averages and phase transitions in Glivenko-Cantelli classes. *IEEE Transactions on Information Theory*, 48(1), 251–263.
- Modha, D. S., & Masry, E. (1998). Memory-universal prediction of stationary random processes. *IEEE Transactions on Information Theory*, 44(1), 117–133.
- Mohri, M., & Rostamizadeh, A. (2009). Rademacher complexity bounds for non-iid processes. *Advances in Neural Information Processing Systems*, 1097–1104.
- Nobel, A. B. (1999). Limits to classification and regression estimation from ergodic processes. *The Annals of Statistics*, 27(1), 262–273.

- Peña, V. H., Lai, T. L., & Shao, Q.-M. (2008). *Self-Normalized Processes: Limit Theory and Statistical Applications*. Springer Science & Business Media.
- Pestov, V. (2010). Predictive PAC learnability: A paradigm for learning from exchangeable input data. *2010 IEEE International Conference on Granular Computing*, 387–391.
- Rakhlin, A., Sridharan, K., & Tewari, A. (2011). Online learning: Stochastic, constrained, and smoothed adversaries. *Advances in Neural Information Processing Systems*, 1764–1772.
- Rakhlin, A., Sridharan, K., & Tewari, A. (2015a). Online learning via sequential complexities. *The Journal of Machine Learning Research*, 16(1), 155–186.
- Rakhlin, A., Sridharan, K., & Tewari, A. (2015b). Sequential complexities and uniform martingale laws of large numbers. *Probability Theory and Related Fields*, 161(1-2), 111–153.
- Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Shalizi, C., & Kontorovich, A. (2013). Predictive PAC learning and process decompositions. *Advances in Neural Information Processing Systems*, 1619–1627.
- Skouras, K., & Dawid, A. P. (2000). *Consistency in misspecified models* (tech. rep. No. 218). Department of Statistical Science, University College London. https://www.ucl.ac.uk/drupal/site_statistics/sites/statistics/files/rr218.pdf
- Smale, S., & Zhou, D.-X. (2009). Online learning with Markov sampling. *Analysis and Applications*, 7(01), 87–113.
- Steinwart, I., & Anghel, M. (2009). Consistency of support vector machines for forecasting the evolution of an unknown ergodic dynamical system from observations with unknown noise. *The Annals of Statistics*, 37(2), 841–875.
- Steinwart, I., Hush, D., & Scovel, C. (2009). Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1), 175–194.
- Vidyasagar, M. (2002). *A Theory of Learning and Generalization* (Second). Springer-Verlag.
- von Luxburg, U., & Schölkopf, B. (2011). Statistical learning theory: Models, concepts, and results. *Handbook of the History of Logic* (pp. 651–706). Elsevier.
- Zimin, A., & Lampert, C. (2017). Learning theory for conditional risk minimization. *Artificial Intelligence and Statistics*, 213–222.