

On Learnability, Complexity and Stability

Silvia Villa, Lorenzo Rosasco and Tomaso Poggio

1 Introduction

A key question in statistical learning is which hypotheses (function) spaces are learnable. Roughly speaking, a hypotheses space is learnable if there is a consistent learning algorithm, i.e. one returning an optimal solution as the number of sample goes to infinity. Classic results for supervised learning characterize learnability of a function class in terms of its complexity (combinatorial dimension) [17, 16, 1, 2, 9, 3]. Indeed, minimization of the empirical risk on a function class having finite complexity can be shown to be consistent. A key aspect in this approach is the connection with empirical process theory results showing that finite combinatorial dimensions characterize function classes for which a uniform law of large numbers holds, namely uniform Glivenko-Cantelli classes [7].

More recently, the concept of stability has emerged as an alternative and effective method to design consistent learning algorithms [4]. Stability refers broadly to continuity properties of learning algorithm to its input and it is known to play a crucial role in regularization theory [8]. Surprisingly, for certain classes of loss functions, a suitable notion of stability of ERM can be shown to characterize learnability of a function class [10, 12, 11].

In this paper, after recalling some basic concepts (Section 2), we review results characterizing learnability in terms of complexity and stability in supervised learning (Section 3) and in the so called general learning (Section 4). We conclude with some remarks and open questions.

Silvia Villa
LCSL, MIT and IIT, USA e-mail: silvia.villa@iit.it

Lorenzo Rosasco
LCSL, MIT and IIT, USA and DIBRIS, Università di Genova, Italy, e-mail: lrosasco@mit.edu

Tomaso Poggio
CBCL, MIT, USA, e-mail: tp@ai.mit.edu

2 Supervised Learning, Consistency and Learnability

In this section, we introduce basic concepts in Statistical Learning Theory (SLT). First, we describe the supervised learning setting, and then, define the notions of consistency of a learning algorithm and of learnability of a hypotheses class.

Consider a probability space (\mathcal{Z}, ρ) , where $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, with \mathcal{X} a measurable space and \mathcal{Y} a closed subset of \mathbb{R} . A loss function is a measurable map $\ell : \mathbb{R} \times \mathcal{Y} \rightarrow [0, +\infty)$. We are interested in the problem of minimizing the expected risk,

$$\inf_{\mathcal{F}} \mathcal{E}_{\rho}, \quad \mathcal{E}_{\rho}(f) = \int_{\mathcal{X} \times \mathcal{Y}} \ell(f(x), y) d\rho(x, y), \quad (1)$$

where $\mathcal{F} \subset \mathcal{Y}^{\mathcal{X}}$ is the set of measurable functions from \mathcal{X} to \mathcal{Y} (endowed with the product topology and the corresponding Borel σ -algebra). The probability distribution ρ is assumed to be fixed but known only through a training set, i.e. a set of pairs $\mathbf{z}_n = ((x_1, y_1), \dots, (x_n, y_n)) \in \mathcal{Z}^n$ sampled identically and independently according to ρ . Roughly speaking, the problem of supervised learning is that of approximately solving Problem (1) given a training set \mathbf{z}_n .

Example 1 (Regression and Classification). In (bounded) regression \mathcal{Y} is a bounded interval in \mathbb{R} , while in binary classification $\mathcal{Y} = \{0, 1\}$. Examples of loss functions are the square loss $\ell(t, y) = (t - y)^2$ in regression and the misclassification loss $\ell(t, y) = \mathbb{1}_{\{t \neq y\}}$ in classification. See [16] for a more exhaustive list of loss functions.

In the next section, the notion of approximation considered in SLT is defined rigorously. We first introduce the concepts of hypotheses space and learning algorithm.

Definition 1. A *hypotheses space* is a set of functions $\mathcal{H} \subseteq \mathcal{F}$. We say that \mathcal{H} is *universal* if $\inf_{\mathcal{F}} \mathcal{E}_{\rho} = \inf_{\mathcal{H}} \mathcal{E}_{\rho}$, for all distributions ρ on \mathcal{Z} .

Definition 2. A *learning algorithm* A on \mathcal{H} is a map,

$$A : \bigcup_{n \in \mathbb{N}} \mathcal{Z}^n \rightarrow \mathcal{H}, \quad \mathbf{z}_n \mapsto A_{\mathbf{z}_n} = A(\mathbf{z}_n),$$

such that, for all $n \geq 1$, $A|_{\mathcal{Z}^n}$ is measurable with respect to the completion of the product σ -algebra on \mathcal{Z}^n .

Empirical Risk Minimization (ERM) is arguably the most popular example of learning algorithm in SLT.

Example 2. Given a training set \mathbf{z}_n the empirical risk $\mathcal{E}_{\mathbf{z}_n} : \mathcal{F} \rightarrow \mathbb{R}$ is defined as

$$\mathcal{E}_{\mathbf{z}_n}(f) = \frac{1}{n} \sum_{i=1}^n \ell(f(x_i), y_i).$$

Given a hypotheses space \mathcal{H} , ERM on \mathcal{H} is defined by minimization of the empirical risk on \mathcal{H} .

We add one remark.

Remark 1 (ERM and Asymptotic ERM). In general some care is needed while defining ERM since a (measurable) minimizer might not be ensured to exist. When $\mathcal{Y} = \{0, 1\}$ and ℓ is the misclassification loss function, it is easy to see that a minimizer exists (possibly non unique). In this case measurability is studied for example in Lemma 6.17 in [15]. When considering more general loss functions or regression problems one might need to consider learning algorithms defined by suitable (measurable) almost minimizers of the empirical risk (see e.g. Definition 10).

2.1 Consistency and Learnability

Aside from computational considerations, the following definition formalizes in which sense a learning algorithm approximatively solves Problem (1).

Definition 3. We say that a learning algorithm A on \mathcal{H} is *uniformly consistent*¹ if

$$\forall \varepsilon > 0, \quad \lim_{n \rightarrow +\infty} \sup_{\rho} \rho^n(\{\mathbf{z}_n : \mathcal{E}_{\rho}(A_{\mathbf{z}_n}) - \inf_{\mathcal{H}} \mathcal{E}_{\rho} > \varepsilon\}) = 0,$$

and *universally uniformly consistent* if \mathcal{H} is universal.

The next definition shifts the focus from a learning algorithm on \mathcal{H} , to \mathcal{H} itself.

Definition 4. We say that a space \mathcal{H} is *uniformly learnable* if there exists a uniformly consistent learning algorithm on \mathcal{H} . If \mathcal{H} is also universal we say that it is *universally uniformly learnable*.

Note that, in the above definitions, the term “uniform” refers to the distribution for which consistency holds, whereas “universal” refers to the possibility of solving Problem (1) without a bias due to the choice of \mathcal{H} . The requirement of uniform learnability implies the existence of a learning rate for A [15] or equivalently a bound on the sample complexity [2]. The following classical result, sometimes called the “no free lunch” theorem, shows that uniform universal learnability of a hypotheses space is too much to hope for.

Theorem 1. *Let $\mathcal{Y} = \{0, 1\}$, and \mathcal{X} such that there exists a measure μ on \mathcal{X} having an atom-free distribution. Let ℓ be the misclassification loss. If \mathcal{H} is universal, then \mathcal{H} is not uniformly learnable.*

The proof of the above result is based on Theorem 7.1 in [6], which shows that for each learning algorithm A on \mathcal{H} and any fixed n , there exists a measure ρ on $\mathcal{X} \times \mathcal{Y}$ such that the expected value of $\mathcal{E}_{\rho}(A_{\mathbf{z}_n}) - \inf_{\mathcal{H}} \mathcal{E}_{\rho}$ is greater than $1/4$. A general form of the no free lunch theorem, beyond classification, is given in [15] (see

¹ Consistency can be defined with respect to other convergence notions for random variables. If the loss function is bounded, convergence in probability is equivalent to convergence in expectation.

Corollary 6.8). In particular, this result shows that the no free lunch theorem holds for convex loss functions, as soon as there are two probability distributions ρ_1, ρ_2 such that $\inf_{\mathcal{H}} \mathcal{E}_{\rho_1} \neq \inf_{\mathcal{H}} \mathcal{E}_{\rho_2}$ (assuming that minimizers exist). Roughly speaking, if there exist two learning problems with distinct solutions, then \mathcal{H} cannot be universal uniformly learnable (this latter condition becomes more involved when the loss is not convex).

The no free lunch theorem shows that universal uniform consistency is too strong of a requirement. Restrictions on either the class of considered distributions ρ or the hypotheses spaces/algorithms are needed to define a meaningful problem. In the following, we will follow the latter approach where assumptions on \mathcal{H} (or A), but not on the class distributions ρ , are made.

3 Learnability of a Hypotheses space

In this section we study uniform learnability by putting appropriate restrictions on the hypotheses space \mathcal{H} . We are interested in conditions which are not only sufficient but also necessary. We discuss two series of results. The first is classical and characterizes learnability of a hypotheses space in terms of suitable complexity measures. The second, more recent, is based on the stability (in a suitable sense) of ERM on \mathcal{H} .

3.1 Complexity and Learnability

Classically assumptions on \mathcal{H} are imposed in the form of restrictions on its "size" defined in terms of suitable notions of combinatorial dimensions (complexity). The following definition of complexity for a class of binary valued functions has been introduced in [17].

Definition 5. Assume $\mathcal{Y} = \{0, 1\}$. We say that \mathcal{H} *shatters* $S \subseteq \mathcal{X}$ if for each $E \subseteq S$ there exists $f_E \in \mathcal{H}$ such that $f_E(x) = 0$, if $x \in E$, and $f_E(x) = 1$ if $x \in S \setminus E$. The *VC-dimension* of \mathcal{H} is defined as

$$\text{VC}(\mathcal{H}) = \max\{d \in \mathbb{N} : \exists S = \{x_1, \dots, x_d\} \text{ shattered by } \mathcal{H}\}$$

The VC-dimension turns out to be related to a special class of functions, called uniform Glivenko-Cantelli, for which a uniform form of the law of large numbers holds [7].

Definition 6. We say that \mathcal{H} is a *uniform Glivenko-Cantelli (uGC) class* if it has the following property

$$\forall \varepsilon > 0, \quad \lim_{n \rightarrow +\infty} \sup_{\rho} \rho^n \left(\left\{ \mathbf{z}_n : \sup_{f \in \mathcal{H}} |\mathcal{E}_{\rho}(f) - \mathcal{E}_{\mathbf{z}_n}(f)| > \varepsilon \right\} \right) = 0.$$

The following theorem completely characterizes learnability in classification.

Theorem 2. *Let $\mathcal{Y} = \{0, 1\}$ and ℓ be the misclassification loss. Then the following conditions are equivalent:*

1. \mathcal{H} is uniformly learnable,
2. ERM on \mathcal{H} is uniformly consistent,
3. \mathcal{H} is a uGC-class,
4. the VC-dimension of \mathcal{H} is finite.

The proof of the above result can be found for example in [2] (see Theorems 4.9, 4.10 and 5.2). The characterization of uGC classes in terms of combinatorial dimensions is a central theme in empirical process theory [7]. The results on binary valued functions are essentially due to Vapnik and Chervonenkis [17]. The proof that uGC of \mathcal{H} implies its learnability is straightforward. The key step in the above proof is showing that learnability is sufficient for finite VC-dimension, i.e. $\text{VC}(\mathcal{H}) < \infty$. The proof of this last step crucially depends on the considered loss function.

A similar result holds for bounded regression with the square [1, 2] and absolute loss functions [9, 3]. In this case, a new notion of complexity needs to be defined since the VC-dimension of real valued function classes is not defined. Here, we recall the definition of γ -fat shattering dimension of a class of functions \mathcal{H} originally introduced in [9].

Definition 7. Let \mathcal{H} be a set of functions from \mathcal{X} to \mathbb{R} and $\gamma > 0$. Consider $S = \{x_1, \dots, x_d\} \subset \mathcal{X}$. Then S is γ -shattered by \mathcal{H} if there are real numbers r_1, \dots, r_d such that for each $E \subseteq S$ there is a function $f_E \in \mathcal{H}$ satisfying

$$\begin{cases} f_E(x) \leq r_i - \gamma & \forall x \in S \setminus E \\ f_E(x) \geq r_i + \gamma & \forall x \in E. \end{cases}$$

We say that (r_1, \dots, r_d) witnesses the shattering. The γ -fat shattering dimension of \mathcal{H} is

$$\text{fat}_{\mathcal{H}}(\gamma) = \max\{d : \exists S = \{x_1, \dots, x_d\} \subseteq \mathcal{X} \text{ s.t. } S \text{ is } \gamma\text{-shattered by } \mathcal{H}\}.$$

As mentioned above, an analogous of Theorem 2 can be proved for bounded regression with the square and absolute losses, if condition 4) is replaced by $\text{fat}_{\mathcal{H}}(\gamma) < +\infty$ for all $\gamma > 0$. We end noting that is an open question proving that the above results holds for loss function other than the square and absolute loss.

3.2 Stability and Learnability

In this section we show that learnability of a hypotheses space \mathcal{H} is equivalent to the stability (in a suitable sense) of ERM on \mathcal{H} . It is useful to introduce the following notation. For a given loss function ℓ , let $L : \mathcal{F} \times Z \rightarrow [0, \infty)$ be defined as

$L(f, z) = \ell(f(x), y)$, for $f \in \mathcal{F}$ and $z = (x, y) \in \mathcal{Z}$. Moreover, let \mathbf{z}_n^i be the training \mathbf{z}_n with the i -th point removed. With the above notation, the relevant notion of stability is given by the following definition.

Definition 8. A learning algorithm A on \mathcal{H} is *uniformly CV_{loo} stable* if there exist sequences $(\beta_n, \delta_n)_{n \in \mathbb{N}}$ such that $\beta_n \rightarrow 0$, $\delta_n \rightarrow 0$ and

$$\sup_{\rho} \rho^n \{ |L(A_{\mathbf{z}_n^i}, z_i) - L(A_{\mathbf{z}_n}, z_i)| \leq \beta_n \} \geq 1 - \delta_n, \quad (2)$$

for all $i \in \{1, \dots, n\}$.

Before illustrating the implications of the above definition to learnability we first add a few comments and historical remarks. We note that, in a broad sense, stability refers to a quantification of the continuity of a map with respect to its input. The key role of stability in learning has long been advocated on the basis of the interpretation of supervised learning as an ill-posed inverse problems [11]. Indeed, the concept of stability is central in the theory of regularization of ill-posed problem [8]. A first quantitative connection between the performance of a symmetric learning algorithm² and a notion of stability is derived in the seminal paper [4]. Here a notion of stability, called uniform stability, is shown to be sufficient for consistency. If we let $\mathbf{z}_n^{i,u}$ be the training \mathbf{z}_n with the i -th point replaced by u , uniform stability is defined as,

$$|L(A_{\mathbf{z}_n^{i,u}}, z) - L(A_{\mathbf{z}_n}, z)| \leq \beta_n, \quad (3)$$

for all $\mathbf{z}_n \in \mathcal{Z}^n$, $u, z \in \mathcal{Z}^n$ and $i \in \{1, \dots, n\}$. A thorough investigation of weaker notions of stability is given in [10]. Here, many different notions of stability are shown to be sufficient for consistency (and learnability) and the question is raised of whether stability (of ERM on \mathcal{H}) can be shown to be necessary for learnability of \mathcal{H} . In particular a definition of CV stability for ERM is shown to be necessary and sufficient for learnability in a Probably Approximate Correct (PAC) setting, that is when $\mathcal{Y} = \{0, 1\}$ and for some $h^* \in \mathcal{H}$, $y = h^*(x)$, for all $x \in \mathcal{X}$. Finally, Definition ?? of CV_{loo} stability is given and studied in [11]. When compared to uniform stability, we see that: 1) the ‘‘replaced one’’ training set $\mathbf{z}_n^{i,u}$ is considered instead of the ‘‘leave one out’’ training set \mathbf{z}_n^i ; 2) the error is evaluated on the point z_i which is left out, rather than any possible $z \in \mathcal{Z}$; finally 3) the condition is assumed to hold for a fraction $1 - \delta_n$ of training sets (which becomes increasingly larger as n increases) rather than uniformly for any training set $\mathbf{z}_n \in \mathcal{Z}^n$.

The importance of CV_{loo} stability is made clear by the following result.

Theorem 3. Let $\mathcal{Y} = \{0, 1\}$ and ℓ be the misclassification loss function. Then the following conditions are equivalent,

1. \mathcal{H} is uniformly learnable,
2. ERM on \mathcal{H} is CV_{loo} stable

² We say that a learning algorithm A is symmetric if it does not depend on the order of the points in \mathbf{z}_n .

The proof of the above result is given in [11] and is based on essentially two steps. The first is proving that CV_{100} stability of ERM on \mathcal{H} implies that ERM is uniformly consistent. The second is showing that if \mathcal{H} is a uGC class then ERM on \mathcal{H} is CV_{100} stable. Theorem 3 then follows from Theorem 2 (since uniform consistency of ERM on \mathcal{H} and \mathcal{H} being uGC are equivalent).

Both steps in the above proof can be generalized to regression as long as the loss function is assumed to be bounded. The latter assumption holds for example if the loss function satisfies a suitable Lipschitz condition and \mathcal{Y} is compact (so that \mathcal{H} is a set of uniformly bounded functions). However, generalizing Theorem 3 beyond classification requires the generalization of Theorem 2. For the the square and absolute loss functions and \mathcal{Y} compact, the characterization of learnability in terms of γ -fat shattering dimension can be used. It is an open question whether there is a more direct way to show that learnability is sufficient for stability, independently to Theorem 2 and to extend the above results to more general classes of loss functions. We will see a partial answer to this question in Section 4.

4 Learnability in the General Learning Setting

In the previous sections we focused our attention on supervised learning. Here we ask whether the results we discussed extend to the so called general learning [16].

Let (\mathcal{Z}, ρ) be a probability space and \mathcal{F} a measurable space. A loss function is a map $L: \mathcal{F} \times \mathcal{Z} \rightarrow [0, \infty)$, such that $L(f, \cdot)$ is measurable for all $f \in \mathcal{F}$. We are interested in the problem of minimizing the expected risk,

$$\inf_{\mathcal{H}} \mathcal{E}_{\rho}, \quad \mathcal{E}_{\rho}(f) = \int_{\mathcal{Z}} L(f, z) d\rho(z), \quad (4)$$

when ρ is fixed but known only through a training set, $\mathbf{z}_n = (z_1, \dots, z_n) \in \mathcal{Z}^n$ sampled identically and independently according to ρ . Definition 2 of a learning algorithm on \mathcal{H} applies as is to this setting and ERM on \mathcal{H} is defined by the minimization of the empirical risk

$$\mathcal{E}_{\mathbf{z}_n}(f) = \frac{1}{n} \sum_{i=1}^n L(f, z_i).$$

While general learning is close to supervised learning, there are important differences. The data space \mathcal{Z} has no natural decomposition, \mathcal{F} needs not to be a space of functions. Indeed, \mathcal{F} and \mathcal{Z} are related only via the loss function L . For our discussion it is important to note that the distinction between \mathcal{F} and the hypotheses space \mathcal{H} becomes blurred. In supervised learning \mathcal{F} is the largest set of functions for which Problem (1) is well defined (measurable functions in $\mathcal{Y}^{\mathcal{X}}$). The choice of a hypotheses corresponds intuitively to a more "manageable" function space. In general learning the choice of \mathcal{F} is more arbitrary as a consequence the the definition of universal hypotheses space is less clear. The setting is too general for an analogue of the no free lunch theorem to hold. Given these premises, in what follows

we will simply identify $\mathcal{F} = \mathcal{H}$ and consider the question of learnability, noting that the definition of uniform learnability extends naturally to general learning. We present two sets of ideas. The first, due to Vapnik, focuses on a more restrictive notion of consistency of ERM. The second, investigates the characterization of uniform learnability in terms of stability.

4.1 Vapnik's Approach and Non Trivial Consistency

The extension of the classical results characterizing learnability in terms of complexity measure is tricky. Since \mathcal{H} is not a function space the definitions of VC or V_γ dimensions do not make sense. A possibility is to consider the class $L \circ \mathcal{H} := \{z \in \mathcal{Z} \mapsto L(f, z) \text{ for some } f \in \mathcal{H}\}$ and the corresponding VC dimension (if L is binary valued) or V_γ dimension (if L is real valued). Classic results about the equivalence between the uGC property and finite complexity apply to the class $L \circ \mathcal{H}$. Moreover, uniform learnability can be easily proved if $L \circ \mathcal{H}$ is a uGC class. On the contrary, the reverse implication does not hold in the general learning setting. A counterexample is given in [16] (Sec. 3.1) showing that it is possible to design hypotheses classes with infinite VC (or V_γ) dimension, which are uniformly learnable with ERM. The construction is as follows. Consider an arbitrary set \mathcal{H} and a loss L for which the class $L \circ \mathcal{H}$ has infinite VC (or V_γ) dimension. Define a new space $\widetilde{\mathcal{H}} := \mathcal{H} \cup \{\tilde{h}\}$ by adding to \mathcal{H} an element \tilde{h} such that $L(\tilde{h}, z) \leq L(h, z)$ for all $z \in \mathcal{Z}$ and $h \in \mathcal{H}$ ³. The space $L \circ \widetilde{\mathcal{H}}$ has infinite VC, or V_γ , dimension and is trivially learnable by ERM, which is constant and coincides with \tilde{h} for each probability measure ρ . The previous counterexample proves that learnability, and in particular learnability via ERM, does not imply finite VC or V_γ dimension. To avoid these cases of “trivial consistency” and to restore the equivalence between learnability and finite dimension, the following stronger notion of consistency for ERM has been introduced by Vapnik [16].

Definition 9. ERM on \mathcal{H} is *strictly uniformly consistent* if and only if

$$\forall \varepsilon > 0, \quad \limsup_{n \rightarrow \infty} \rho^n \left(\inf_{\mathcal{H}_c} \mathcal{E}_{\mathbf{z}_n}(f) - \inf_{\mathcal{H}_c} \mathcal{E}_\rho(f) > \varepsilon \right) = 0,$$

where $\mathcal{H}_c = \{f \in \mathcal{H} : \mathcal{E}_\rho(f) \geq c\}$.

The following result characterizes strictly uniform consistency in terms of uGC property of the class $L \circ \mathcal{H}$ (see Theorem 3.1 and its Corollary in [16])

Theorem 4. Let $B > 0$ and assume $L(f, z) \leq B$ for all $f \in \mathcal{H}$ and $z \in \mathcal{Z}$. Then the following conditions are equivalent,

1. ERM on \mathcal{H} is strictly consistent,
2. $L \circ \mathcal{H}$ is a uniform one-sided Glivenko-Cantelli class.

³ Note that this construction is not possible in classification or in regression with the square loss.

The definition of one-sided Glivenko-Cantelli class simply corresponds to omitting the absolute value in Definition 6.

4.2 Stability and Learnability for General Learning

In this section we discuss ideas from [14] extending the stability approach to general learning. The following definitions are relevant.

Definition 10. A *uniform Asymptotic ERM (AERM) algorithm* A on \mathcal{H} is a learning algorithm such that

$$\forall \varepsilon > 0, \quad \lim_{n \rightarrow \infty} \sup_p \rho^n(\{\mathbf{z}_n : \mathcal{E}_{\mathbf{z}_n}(A_{\mathbf{z}_n}) - \inf_{\mathcal{H}} \mathcal{E}_{\mathbf{z}_n} > \varepsilon\}) = 0.$$

Definition 11. A learning algorithm A on \mathcal{H} is *uniformly replace one (RO) stable* if there exists a sequence $\beta_n \rightarrow 0$ such that

$$\frac{1}{n} \sum_{i=1}^n |L(A_{\mathbf{z}_n^{i,u}}, z) - L(A_{\mathbf{z}_n}, z)| \leq \beta_n.$$

for all $\mathbf{z}_n \in \mathcal{Z}^n$, $u, z \in \mathcal{Z}^n$ and $i \in \{1, \dots, n\}$.

Note that the above definition is close to that of uniform stability (3), although the latter turns out to be a stronger condition. The importance of the above definitions is made clear by the following result.

Theorem 5. Let $B > 0$ and assume $L(f, z) \leq B$ for all $f \in \mathcal{H}$ and $z \in \mathcal{Z}$. Then the following conditions are equivalent,

1. \mathcal{H} is uniformly learnable,
2. there exists an AERM algorithm on \mathcal{H} which is RO stable.

As mentioned in Remark 1, Theorem 3 holds not only for exact minimizers of the empirical risk, but also for AERM. In this view, there is a subtle difference between Theorem 3 and Theorem 5. In supervised learning, Theorem 3 shows that uniform learnability implies that *every* ERM (AERM) is stable, while in general learning, Theorem 5 shows that uniform learnability implies the *existence* of a stable AERM (whose construction is not explicit).

The proof of the above result is given in Theorem 7 in [14]. The hard part of the proof is showing that learnability implies existence of a RO stable AERM. This part of the proof is split in two steps showing that: 1) if there is a uniformly consistent algorithm A , then there exists a uniformly consistent AERM A' (Lemma 20 and Theorem 10); 2) every uniformly consistent AERM is also RO stable (Theorem 9). Note that the results in [14] are given in expectation and with some quantification of how different convergence rates are related. Here we give results in probability to be uniform with the rest of the paper and state only asymptotic results to simplify the presentation.

5 Discussion

In this paper we reviewed several results concerning the learnability of a hypotheses space. Extensions of these ideas can be found in [5] (and references therein) for multi-category classification, and in [13] for sequential prediction. It would be interesting to devise constructive proofs in general learning suggesting how stable learning algorithms can be designed. Moreover, it would be interesting to study universal consistency and learnability in the case of samples from non stationary processes.

References

1. Alon, N., Ben-David, S., Cesa-Bianchi, N., Haussler, D.: Scale-sensitive dimensions, uniform convergence, and learnability. *J. ACM* **44**(4), 615–631 (1997)
2. Anthony, M., Bartlett, P.L.: *Neural network learning: theoretical foundations*. Cambridge University Press, Cambridge (1999)
3. Bartlett, P., Long, P., Williamson, R.: Fat-shattering and the learnability of real-valued functions. *Journal of Computer and System Sciences* **52**, 434–452 (1996)
4. Bousquet, O., Elisseeff, A.: Stability and generalization. *Journal of Machine Learning Research* **2**, 499–526 (2002)
5. Daniely, A., Sabato, S., Ben-David, S., Shalev-Shwartz, S.: Multiclass learnability and the erm principle. *Journal of Machine Learning Research - Proceedings Track* **19**, 207–232 (2011)
6. Devroye, L., Györfi, L., Lugosi, G.: *A Probabilistic Theory of Pattern Recognition*. No. 31 in *Applications of mathematics*. Springer, New York (1996)
7. Dudley, R., Giné, E., Zinn, J.: Uniform and universal Glivenko-Cantelli classes. *J. Theoret. Prob.* **4**, 485–510 (1991)
8. Engl, H.W., Hanke, M., Neubauer, A.: *Regularization of inverse problems, Mathematics and its Applications*, vol. 375. Kluwer Academic Publishers Group, Dordrecht (1996)
9. Kearns, M.J., Schapire, R.E.: Efficient distribution-free learning of probabilistic concepts. In: *Computational learning theory and natural learning systems*, Vol. I, Bradford Book, pp. 289–329. MIT Press, Cambridge, MA (1994)
10. Kutin, S., Niyogi, P.: Almost-everywhere algorithmic stability and generalization error. Tech. Rep. TR-2002-03, Department of Computer Science, The University of Chicago (2002)
11. Mukherjee, S., Niyogi, P., Poggio, T., Rifkin, R.: Learning theory: stability is sufficient for generalization and necessary and sufficient for consistency of empirical risk minimization. *Adv. Comput. Math.* **25**(1-3), 161–193 (2006)
12. Poggio, T., Rifkin, R., Mukherjee, S., Niyogi, P.: General conditions for predictivity in learning theory. *Nature* **428**, 419–422 (2004)
13. Rakhlin, A., Sridharan, K., Tewari, A.: Online learning: Beyond regret. *Journal of Machine Learning Research - Proceedings Track* **19**, 559–594 (2011)
14. Shalev-Shwartz, S., Shamir, O., Srebro, N., Sridharan, K.: Learnability, stability and uniform convergence. *J. Mach. Learn. Res.* **11**, 2635–2670 (2010)
15. Steinwart, I., Christmann, A.: *Support vector machines*. Information Science and Statistics. Springer, New York (2008)
16. Vapnik, V.: *The Nature of Statistical Learning Theory*. Springer Verlag, New York (1995)
17. Vapnik, V.N., Chervonenkis, A.Y.: Theory of uniform convergence of frequencies of events to their probabilities and problems of search for an optimal solution from empirical data. *Avtomat. i Telemek. (2)*, 42–53 (1971)