

Universal Online Learning with Unbounded Losses: Memory Is All You Need

Moïse Blanchard

Massachusetts Institute of Technology

MOISEB@MIT.EDU

Romain Cosson

Massachusetts Institute of Technology

COSSON@MIT.EDU

Steve Hanneke

Purdue University

STEVE.HANNEKE@GMAIL.COM

Abstract

We resolve an open problem of [13] on the subject of universally consistent online learning with non-i.i.d. processes and unbounded losses. The notion of an *optimistically universal learning rule* was defined by Hanneke [13] in an effort to study learning theory under minimal assumptions. A given learning rule is said to be optimistically universal if it achieves a low long-run average loss whenever the data generating process makes this goal achievable by some learning rule. [13] posed as an open problem whether, for every unbounded loss, the family of processes admitting universal learning are precisely those having a finite number of distinct values almost surely. In this paper, we completely resolve this problem, showing that this is indeed the case. As a consequence, this also offers a dramatically simpler formulation of an optimistically universal learning rule for any unbounded loss: namely, the simple *memorization* rule already suffices. Our proof relies on constructing random measurable partitions of the instance space and could be of independent interest for solving other open questions [14]. We extend the results to the non-realizable setting thereby providing an optimistically universal Bayes consistent learning rule.

Keywords: online learning, universal consistency, stochastic processes, measurable partitions, statistical learning theory, Borel measure

1. Introduction

Online learning. One of the main classical questions in statistical learning is *learnability*: whether it is possible to have low prediction loss on a prediction task given observations. In this paper, we study this question in the context of *online learning*. In this setting, there is a possibly random sequence of points $\mathbb{X} = (X_t)_{t \in \mathbb{N}}$ from a space \mathcal{X} of inputs and target values $\mathbb{Y} = (Y_t)_{t \in \mathbb{N}}$ from a space of outputs \mathcal{Y} . These two sequences are stochastic processes in general. Learning occurs sequentially, where at each time t the learner has observed $(X_s)_{s \leq t-1}$ and $(Y_s)_{s \leq t-1}$ and X_t , and makes a prediction $\hat{Y}_t \in \mathcal{Y}$ for the value of Y_t . Thus, we may consider $\hat{Y}_t = f_t((X_1, Y_1), \dots, (X_{t-1}, Y_{t-1}), X_t)$ for some function f_t (possibly randomized). We will assume that there is a measurable *target function* $f^* : \mathcal{X} \rightarrow \mathcal{Y}$ such that $Y_t = f^*(X_t)$: that is, the target values Y_t are deterministic in the input. However, we place no restrictions on this function. Given a loss function ℓ , the performance of the algorithm is quantified by the value $\ell(\hat{Y}_t, Y_t)$. In particular, we will say that an algorithm has learnt the prediction task if it guarantees a low long-run average loss, i.e. $\frac{1}{n} \sum_{t=1}^n \ell(\hat{Y}_t, Y_t) \xrightarrow[n \rightarrow \infty]{} 0$ (*a.s.*).

Universal online learning. In general, it is not possible to obtain such guarantees for all sequences \mathbb{X} and functions f^* . It is therefore necessary to constrain the set of considered pairs (\mathbb{X}, f^*) . There is a rich literature studying online learning with unrestricted sequences \mathbb{X} but with restrictions on the function f^* [17; 1], or with a mix of restrictions on \mathbb{X} and f^* [16; 21; 26; 3]. Here, we focus on *universal learning*, which imposes no assumptions on the set of target functions f^* , but restricts the input sequences \mathbb{X} . Thus, we are interested in algorithms which are *universally consistent* under a given family of stochastic processes \mathbb{X} . For instance, a classic result in this line of work states that in the Euclidian space $\mathcal{X} = \mathbb{R}^n$, even the inductive *nearest neighbor* predictor for classification has vanishing error rate for every i.i.d. input sequence \mathbb{X} and every target f^* [5; 24; 6], and thus can be shown to have vanishing long-run average loss in the online setting for this case as well.

In this work, we primarily focus on *unbounded* loss functions: i.e., the case $\sup_{y, y' \in \mathcal{Y}} \ell(y, y') = \infty$. In the context of i.i.d. processes, or various extensions thereof, there have been many works that consider unbounded losses, but with additional restrictions on the Y_t sequence (which effectively allow them to reduce back to the bounded case, in a certain sense). For instance, in the i.i.d. setting, [10] presents a variety of results on universal consistency for regression when Y_t has finite *variance*. Similarly, [11; 12; 2] develop several results on universal consistency for regression under *stationary ergodic* processes, under the restriction that Y_t has finite fourth moments. Extending these results to admit general families of processes (\mathbb{X}, \mathbb{Y}) is certainly an interesting and worthy direction of research. However, in the present work we focus on the case of *unrestricted* Y_t sequences, aside from the aforementioned assumption that $Y_t = f^*(X_t)$. In particular, in the regression setting, there are no variance restrictions on Y_t . As a consequence, the results we arrive at necessarily impose stricter requirements on the process \mathbb{X} of inputs. Indeed, one of the main contributions of the present work is illuminating just how restrictive the requirement on \mathbb{X} must necessarily be for universal consistency to be possible with unbounded losses.

Most of the literature on universal consistency has relied on conventional assumptions on stochastic processes, imported from the probability theory literature, such as the i.i.d. assumption or relaxations to stationary ergodic [18; 9; 8] or satisfying a law of large numbers [19; 22]. In contrast, in an effort to formulate a theory of learning under minimal assumptions, [13] recently introduced a framework that uses provably-minimal assumptions on \mathbb{X} . Specifically, the sole assumption on the input sequence is that there exists some algorithm which achieves universal consistency. This is known as the so-called *optimist's decision theory*: in order to achieve a given objective – in our case universal consistency – the minimal assumption is that this objective is at least possible. In that sense, the optimist's assumption is minimal, as it is a necessary assumption to prove any positive results on universal consistency. We are then particularly interested in designing online learning rules with guarantees that hold without further assumptions, which are referred to as *optimistically universal* learning rules. Such learning rules are consistent whenever learning is possible. Equivalently, if an optimistically universal learner fails for a specific instance, any other strategy would fail to be universally consistent under that process \mathbb{X} as well.

In the case of *unbounded* losses ℓ , the existence of an optimistically universal online learning rule was settled by [13].

This work also expresses a condition (condition FMV below) which characterizes the family of processes \mathbb{X} that admit the existence of universally consistent online learning rules for any (and all) unbounded losses. However, the definition of the optimistically universal learning rule given in that work, the proof that it satisfies this property, and also the proofs establishing that the proposed condition indeed characterizes the relevant family of processes, are actually quite complex. For

instance, the learning rule involves identifying a function contained in a certain countable function class $\tilde{\mathcal{F}}$, satisfying constraints on its losses relative to various other values, and the proof proceeds via arguing that there exists a choice of $\tilde{\mathcal{F}}$ that is *dense* in the set of all measurable functions, in a sense relevant to learning under every \mathbb{X} satisfying the condition. However, [13] also poses an interesting open problem (Open Problem 4 there) regarding a potential dramatic simplification of this theory. The essential question is the following:

Open Problem [13]: *For unbounded losses, is it true that there exist universally consistent online learning rules under \mathbb{X} if and only if \mathbb{X} almost surely has a finite number of distinct elements?*

Summary of contributions. We completely resolve the above question, proving that the original condition proposed by [13] (condition FMV below) is equivalent to the condition that \mathbb{X} almost surely contains only a finite number of distinct elements. Therefore, for unbounded losses, there exist universally consistent online learning rules under \mathbb{X} if and only if \mathbb{X} almost surely contains only a finite number of distinct elements. This result has immediate implications for drastically simplifying the theory of universally consistent learning with unbounded losses. Rather than the complicated learning strategy proposed by [13], it suffices to use the simple *memorization* algorithm, which simply remembers all past data points (X_s, Y_s) , $s < t$, and if the new X_t satisfies $X_t = X_s$ for some $s < t$, it predicts Y_s . If \mathbb{X} has only a finite number of distinct elements, then clearly this strategy has only finitely many non-zero losses, and hence would be universally consistent. Indeed, since this is the case for all such \mathbb{X} , for any unbounded loss this algorithm is also *optimistically* universal.

We refer to the above simple condition as the *finite support* (FS) condition:

Condition FS *Define FS as the set of all stochastic processes \mathbb{X} such that*

$$|\{x \in \mathcal{X} : \mathbb{X} \cap \{x\} \neq \emptyset\}| < \infty \quad (a.s.).$$

As in [13], we let SUOL denote the family of all processes \mathbb{X} under which there exist universally consistent online learning rules (defined more formally below). Our main result can then be stated as follows.

Theorem 1 *For \mathcal{X} any separable metric space and ℓ any unbounded loss, SUOL = FS.*

Although the existence of an optimistically universal learning algorithm was shown in [13] for unbounded losses, the above new characterization by FS drastically simplifies the definition of such a learning rule. As mentioned, as a consequence, we will prove that the simple “memorization rule” (described above) is optimistically universal.

Theorem 2 *If \mathcal{X} is a separable metric space and the loss is unbounded, the memorization rule is an optimistically universal online learning rule.*

Outline of paper. The rest of this paper is structured as follows. We introduce the formal setup and preliminaries in Section 2. We prove the main theorem on universal learning in Section 3, starting with the case $\mathcal{X} = [0, 1]$ and generalizing it to hold for any separable metric space in Section 4. We discuss consequences of the result for *inductive* and *self-adaptive* learning in Section 5. Finally, in Section 6, we consider a noisy setting and prove the existence of an optimistically universal Bayes consistent learning rule when the loss is unbounded. Remaining open problems on optimistically universal online learning will be recalled in the conclusion in Section 7.

2. Background and Preliminaries

2.1. Formal setup

Input space. We consider the general setup where (\mathcal{X}, ρ) is a separable metric space. We define the set \mathcal{B} of measurable subsets of \mathcal{X} to be the σ -algebra generated by the topology induced by the metric ρ . For more details on this setup, we refer to [20].

Value space and loss function. We fix the value space \mathcal{Y} , and an associated loss function $\ell : \mathcal{Y}^2 \rightarrow [0, \infty)$ that satisfies a relaxed triangle inequality $\forall y_1, y_2, y_3 \in \mathcal{Y}^3 : \ell(y_1, y_3) \leq c_\ell(\ell(y_2, y_1) + \ell(y_2, y_3))$, where c_ℓ is a finite constant, as well as that

$$\forall y \in \mathcal{Y}, \ell(y, y) = 0. \text{ For instance, the squared loss in regression satisfies this with } c_\ell = 2.$$

While there remain important open questions for universal online learning with *bounded* loss functions [14], in the present work we focus on *unbounded* loss functions: that is, we assume $\sup_{y_1, y_2 \in \mathcal{Y}} \ell(y_1, y_2) = \infty$.

Data generation process. The input data is generated by an infinite-horizon stochastic process $\mathbb{X} = \{X_t\}_{t=1}^\infty$, taking its values in \mathcal{X} . We impose no constraint on the nature of this process. We assume that the output data $\mathbb{Y} = \{Y_t\}_{t=1}^\infty$ is generated from \mathbb{X} by a deterministic measurable function $f^* : \mathcal{X} \rightarrow \mathcal{Y}$. In other words, we have $\{Y_t\}_{t=1}^\infty = \{f^*(X_t)\}_{t=1}^\infty$. Note that this framework substantially differs from other setups where the response function is noisy. When we look at a bounded time horizon $t \geq 1$, we will use the following notation: $\mathbb{X}_{\leq t} = \{X_1, \dots, X_t\}$ and $\mathbb{X}_{< t} = \{X_1, \dots, X_{t-1}\}$. We will also abuse notations and allow ourselves to use \mathbb{X} to denote the set of all the values taken by the stochastic process. For instance we will write $\#\mathbb{X} = \#\{x \in \mathcal{X} : \{x\} \cap \mathbb{X} \neq \emptyset\} \in \mathbb{N} \cup \{+\infty\}$ the number of different values taken by \mathbb{X} , we will also write $\#\mathbb{X}_{\leq t}$ the number of different values taken by \mathbb{X} before time t .

Online learning rule. An online learning rule is formally defined as a sequence $\{f_t\}_{t=1}^\infty$ of measurable functions $f_t : \mathcal{X}^{t-1} \times \mathcal{Y}^{t-1} \times \mathcal{X} \rightarrow \mathcal{Y}$. Given $t - 1$ training examples of the form $(X_i, f^*(X_i)) \in \mathcal{X} \times \mathcal{Y}$ and one new sample X_t , the learning rule f_t makes prediction $f_t(\mathbb{X}_{< t}, \mathbb{Y}_{< t}, X_t)$ for $f^*(X_t)$. As an important example, the *memorization learning rule* $\{f_t\}_{t=1}^\infty$ is defined as follows:

$$f_t((x_i)_{i<t}, (y_i)_{i<t}, x_t) = \begin{cases} y_i & \text{if } x_t = x_i, \\ y_0 & \text{if } x_t \notin \{x_i\}_{i<t}, \end{cases}$$

where $y_0 \in \mathcal{Y}$ is some arbitrary default response.

Learning task. The goal we pursue is to minimize the online loss defined as,

$$\mathcal{L}_{\mathbb{X}}(f, f^*, T) = \frac{1}{T} \sum_{t=1}^T \ell(f_t(\mathbb{X}_{< t}, \mathbb{Y}_{< t}, X_t), f^*(X_t)).$$

Specifically, letting

$$\mathcal{L}_{\mathbb{X}}(f, f^*) = \limsup_{T \rightarrow \infty} \mathcal{L}_{\mathbb{X}}(f, f^*; T),$$

we would like to guarantee $\mathcal{L}_{\mathbb{X}}(f, f^*) = 0$ (*a.s.*). This motivates the following definition of universal consistency.

Universal consistency and optimistically universal learning rules. We recall here the full statement of the definitions introduced by [13]. We say an online learning rule $\{f_t\}_{t=1}^\infty$ is *strongly universally consistent* under \mathbb{X} if for every measurable $f^* : \mathcal{X} \rightarrow \mathcal{Y}$, we have that $\mathcal{L}_{\mathbb{X}}(f, f^*) = 0$ (a.s.). We say the process \mathbb{X} admits *strong universal online learning* if there exists an online learning rule $\{f_t\}_{t=1}^\infty$ that is strongly universally consistent under \mathbb{X} . We denote by SUOL the set of all processes \mathbb{X} that admit strong universal online learning. An online learning rule $\{f_t\}_{t=1}^\infty$ is said to be *optimistically universal* if it is universally consistent under every \mathbb{X} in SUOL.

In this line of work, two of the main interests are: (1) providing concise conditions characterizing the family SUOL in terms of properties of the stochastic process \mathbb{X} , and (2) identifying particular learning rules that are optimistically universal: that is, learning rules that are universally consistent under every \mathbb{X} in SUOL. Establishing the existence of optimistically universal online learning rules remains an open problem for all *bounded* loss functions. In the case of *unbounded* losses this question has been settled. Specifically, [13] shows that, for any unbounded loss, there exists optimistically universal online learning rules. Moreover, [13] also expresses a condition which characterizes the family SUOL. The condition requires that, for every countable measurable partition of \mathcal{X} , the process \mathbb{X} visits a finite number of cells almost surely. This will be referred to as the “finite measurable visits” (FMV) condition:

Condition FMV [13] Define the set FMV as the set of all processes \mathbb{X} satisfying the condition that, for every disjoint sequence $\{A_k\}_{k=1}^\infty$ in \mathcal{B} with $\cup_{k=1}^\infty A_k = \mathcal{X}$ (i.e., every countable measurable partition),

$$\#\{k \in \mathbb{N} : A_k \cap \mathbb{X} \neq \emptyset\} < \infty \quad (a.s.).$$

It is worth noting, however, that the specification and analysis of the optimistically universal learning rule in [13], and the proof that Condition FMV indeed characterizes SUOL, are all quite complicated. For instance, the algorithm and its analysis directly rely on a construction of a countable dense subset of the set of measurable functions, under a metric appropriate to learning with unbounded losses. The learning algorithm then solves a sequence of constraint satisfaction problems specified in terms of this countable dense set of functions, to select one such function as its predictor. It is therefore desirable to *simplify* the theory, not only for practical reasons, but also to help us to intuitively understand the varieties of processes that admit universally consistent learners, and to clarify what kinds of learning rules can be optimistically universal. Toward this end, [13] poses an important open question regarding a potential simplification: is SUOL characterized by Condition FS? The main contribution of the present work is showing that indeed this is *true* (Theorem 1). This fact allows us to dramatically simplify the entire theory of universally consistent online learning with unbounded losses. Several simplifications are immediate from this:

1. This provides a new, stronger characterization of SUOL.
2. The proof establishing $\text{FS} = \text{SUOL}$ is significantly simpler than the original proof that $\text{FMV} = \text{SUOL}$.
3. The equivalence $\text{FS} = \text{SUOL}$ immediately implies that the simple *memorization* rule is optimistically universal. This contrasts with the complicated construction used in the optimistically universal learner of [13], which solves a sequence of constraint satisfaction problems in terms of a countable dense set of measurable functions.

A remark on dependency in the problem setup. At first glance, one might think that the class SUOL and its characterization should somehow depend on the specific setup given by (\mathcal{X}, ρ) , \mathcal{Y} and ℓ . However, as [13] showed (and as is implied by our Theorem 1), this dependency is very mild. In particular, the existence of an optimistically universal learning rule does not depend on the choice of (\mathcal{Y}, ℓ) as long as the loss is unbounded: $\sup_{y, y' \in \mathcal{Y}} \ell(y, y') = \infty$. Moreover, our results will hold for any separable metric space (\mathcal{X}, ρ) .

Outline of the proof. The essential strategy of the proof relies on the fact that $\text{FS} \subset \text{SUOL} \subset \text{FMV}$. The left inclusion is rather obvious. Indeed, if \mathbb{X} contains a finite number of distinct values (a.s.), even the simple *memorization* learning rule is universally consistent. For the sake of thoroughness, we include a brief proof of this observation in Section 2.2. The second inclusion, that $\text{SUOL} \subset \text{FMV}$, was shown by [13] as part of the proof that $\text{SUOL} = \text{FMV}$. We state this result formally in Section 2.3, and for the sake of being self-contained we include its proof in Appendix B. Given these inclusions $\text{FS} \subset \text{SUOL} \subset \text{FMV}$, what remains is establishing that $\text{FMV} = \text{FS}$. Establishing this equivalence is the main technical contribution of this work (Theorem 6). Its proof relies on constructing random measurable partitions of the space \mathcal{X} . We now turn to discussing the details of each of these components.

2.2. Sufficient condition for SUOL

We begin with the easiest of the claimed inclusions: namely, $\text{FS} \subset \text{SUOL}$. Recall that condition FS corresponds to having a finite number of values almost surely, i.e. $\#\mathbb{X} < \infty$ (a.s.). While it may be rather obvious that all such processes admit a strong universal learning rule (i.e., they belong to SUOL), for the sake of thoroughness we present a simple proof of this fact.

Proposition 3 $\text{FS} \subset \text{SUOL}$. *In particular, the memorization rule is universally consistent under every $\mathbb{X} \in \text{FS}$.*

Proof We will show that the memorization learning rule is universally consistent under every \mathbb{X} that takes a finite number of values almost surely. We can formally prove this result as follows. Let $\mathbb{X} \in \text{FS}$ be a given stochastic process and let $\{f_t\}_{t=1}^{\infty}$ be the memorization learning rule defined earlier. Observe that, for any measurable target function $f^* : \mathcal{X} \rightarrow \mathcal{Y}$, the (random) quantity $M = \max_{t \in \mathbb{N}} \ell(y_0, f^*(X_t))$ is always finite (a.s.), as it is a maximum over a finite set: $M < \infty$ (a.s.). Now observe that the memorization rule makes at most $\#\mathbb{X}$ errors, each of value at most M . Therefore, $0 \leq \hat{\mathcal{L}}_{\mathbb{X}}(f, f^*, T) \leq \frac{1}{T} M \cdot \#\mathbb{X} \xrightarrow{T \rightarrow \infty} 0$ (a.s.). ■

An important remark is that this condition is not *testable*: there does not exist a consistent hypothesis test for FS. In other terms it is not possible to decide from the stream of input data \mathbb{X} whether the process satisfies FS or not. Formally, a *hypothesis test* refers to a sequence of possibly random decision functions $\hat{t}_n : \mathcal{X} \rightarrow \{0, 1\}$. We then say that a test is *consistent* for a class of processes \mathcal{C} if for any process \mathbb{X} , $\hat{t}_n(\mathbb{X}_{\leq n}) \rightarrow \mathbf{1}_{\mathbb{X} \in \mathcal{C}}$ in probability.

Proposition 4 *If \mathcal{X} is infinite, there is no consistent hypothesis test for condition FS.*

Proof This is a consequence from Theorem 1 and the fact that there is no consistent hypothesis test for SUOL when \mathcal{X} is infinite, shown in [13] (Theorem 59). For completeness, we provide a direct and simplified proof in Appendix A. ■

This justifies the terminology “optimistic” in *optimistically universal learning rule* in the sense that belonging to the set of sequences for which universal learning is achievable, which we will prove is equal to FS, is a non-testable assumption.

2.3. Necessary condition for SUOL

We now recall a necessary condition for a stochastic process \mathbb{X} to admit strong universal online learning. It was shown that condition FMV, which requires that \mathbb{X} will only visit a finite number of zones for all given countable measurable partitions of \mathcal{X} , is necessary for universal learning.

Theorem 5 ([13]) $SUOL \subset FMV$.

We present the proof idea here, while for the purpose of being self-contained, a brief version of the full proof is included in Appendix B. Suppose the process \mathbb{X} does not satisfy FMV, then there exists a measurable partition $\{A_k\}_{k=1}^{\infty}$ such that with positive probability, \mathbb{X} visits an infinite number of the A_k . Now define a function f^* that is randomly piece-wise constant, i.e. constant on each A_k and taking random value $f^*(x \in A_k) \in \{y_{k,0}, y_{k,1}\} \in \mathcal{Y}$ where $\ell(y_{k,0}, y_{k,1})$ grows large at a sufficiently fast rate. Any online learning rule will fail to predict $Y_t = f^*(X_t)$ an infinite number of times, inducing non-vanishing loss. The reason is that the learning rule cannot leverage any of the training data $(\mathbb{X}_{<t}, \mathbb{Y}_{<t})$ when A_k is visited for the first time at time t .

2.4. Open Problem 4

In the previous sections, we recalled the inclusions $FS \subset SUOL \subset FMV$. This allows for a concise expression of the conjecture formulated by [13] and referred to as “Open Problem 4” (which, in light of the result of [13] that $SUOL = FMV$, is equivalent to the formulation of the problem stated earlier in Section 2).

Open Problem 4 ([13]) *Is it true that $FMV = FS$?*

We will prove that this equality holds for all separable metric spaces \mathcal{X} . This in turn implies that $SUOL = FS$ and therefore ensures that “memorization”, which we already saw is universally consistent for all processes in FS (Proposition 3), is an optimistically universal learning rule (thus establishing Theorem 2). The solution to the open problem will be detailed in Section 3 and generalised to all separable metric spaces in Section 4. We conclude this section by giving some additional inspiration for the proofs that will follow.

Remarks on Open Problem 4. In words, the question asked by Open Problem 4 is whether the set of countable measurable partitions is sufficiently large to separate all stochastic processes that take an infinite number of values.

It was already observed by [13] that when \mathcal{X} is countable or when \mathbb{X} is deterministic, FS and FMV are equal. However, both these setups come with a natural partition: if \mathcal{X} is countable, $\{\{x\} : x \in \mathcal{X}\}$ becomes a countable measurable partition of \mathcal{X} , and when \mathbb{X} is deterministic $\{\{x\} : \{x\} \cap \mathbb{X} \neq \emptyset\} \cup \{\mathcal{X} \setminus \mathbb{X}\}$ will also isolate all the different values taken by \mathbb{X} .

In the uncountable case, for instance when $\mathcal{X} = \mathbb{R}$, we aim to define a partition $\{A_k\}_{k=1}^{\infty}$ that scatters the space. We want to minimize the chance that two values taken by the process, say $X_t \neq X_{t'}$, fall in the same A_k . A classical and tempting way to build such a partition would be using the axiom of choice [28; 23]. Define the equivalence relation $x \sim_{\mathbb{Q}} y \iff x - y \in \mathbb{Q}$, where \mathbb{Q} can be enumerated $\mathbb{Q} = \{q_1, q_2, \dots\}$. Now for each of the equivalence classes of the form $\{x\} + \mathbb{Q}$, choose one representer. Denote by A the set of all representers and observe that $\{A + \{q\}\}_{q \in \mathbb{Q}}$ makes a countable partition of \mathbb{R} . Note that two different values of \mathbb{X} , say $X_t \neq X_{t'}$, fall in the same equivalence class only if $X_t - X_{t'}$ was chosen as a representer. This event could be made very rare if we were to shift all representers by a uniform random variable, or to choose the representer at random within their class of equivalence. The reason why this does not prove the result is that the corresponding partition $\{A_k\}_{k=1}^{\infty}$ is not measurable.

Another idea to create such a random partition $\{A_k\}_{k=1}^{\infty}$ would be to assign each $x \in \mathbb{R}$ to a set $A_{k(x)}$ where the index $k(x) \in \mathbb{N}$ is chosen independently at random following an exponential law $\mathcal{E}(\frac{1}{2})$: $\mathbb{P}(k(x) = k) = \frac{1}{2^k}$. The indices $\{k(X_1), k(X_2), \dots\}$ to which the elements of the sequence $\mathbb{X} = \{X_1, X_2, \dots\}$ are assigned, are almost surely unbounded when $\#\mathbb{X} = +\infty$, disproving condition FMV. We will refer to this construction as the partition \mathcal{P} as it will later be a useful inspiration. Unfortunately, as such, \mathcal{P} not define a proper partition because the sets A_k are not measurable in general. To solve this issue, instead of defining point-wise random sets, we will use countable union of small intervals. Depending on the scale of the process \mathbb{X} , these sets will give same behaviour as the parts of \mathcal{P} . We will make this idea more precise in the following paragraph.

We first recall a construction of dense open sets of \mathbb{R} with measure at most $\epsilon > 0$. Following a classical argument, one can consider the union of open intervals $\cup_{i \geq 1} (q_i - \frac{\epsilon}{2^i}, q_i + \frac{\epsilon}{2^i})$, where $\{q_1, q_2, \dots\}$ are i.i.d. sampled from some probability density of full support. If we denote the remainders $R_k = \cup_{i \geq k} (q_i - \frac{\epsilon}{2^i}, q_i + \frac{\epsilon}{2^i})$ and consider the partition $\{A_k\}_{k=0}^{\infty}$ defined by $A_k = R_k \setminus R_{k+1}$ where $R_0 = \mathbb{R}$, one could hope that any sequence \mathbb{X} taking infinite values will visit an infinite number of the A_k . In fact, this is true if the convergence rate of \mathbb{X} is not too fast but not in the general case. We will therefore use a decay rate adapted to the process \mathbb{X} through a parameter δ_k defined as follows,

$$\delta_k = \min \left\{ |x - y| \mid x, y \in \mathbb{X}_{\leq N}, \#\mathbb{X}_{\leq N} \geq 2^{2k+2} \right\}.$$

A key intuition is that the first k distinct points visited by \mathbb{X} have scale δ_k . Thus, a remainder R_i of smaller scale – such that the length of the intervals defining R_i is $\ll \delta_k$ – will appear uniformly random to the first k distinct inputs, similarly to the point-wise random sets from the partition \mathcal{P} introduced above.

3. Main Result

In this section, we state and prove the equivalence of FS and FMV therefore guaranteeing that memorization is an optimistically universal learning rule in the unbounded setup. The following result represents the main technical contribution of this work.

Theorem 6 (Main Result) *For any separable metric space (\mathcal{X}, ρ) , $\text{FMV} = \text{FS}$.*

Together with Proposition 3 and Theorem 5, this result implies Theorem 1 and Theorem 2. In this section, we prove the result for $\mathcal{X} = [0, 1]$, as it provides a direct and simple construction. The proof will then be generalised to all separable metric spaces in Section 4.

Proposition 7 *If $\mathcal{X} = [0, 1]$ with its usual topology, $\text{FMV} = \text{FS}$.*

Proof The inclusion $\text{FS} \subset \text{FMV}$ is a direct observation, therefore we focus on proving that $\text{FMV} \subset \text{FS}$. Let \mathbb{X} be a stochastic process which does not satisfy FS. The goal is to construct a countable measurable partition $\{A_k\}_{k=1}^\infty$ of \mathcal{X} which disproves condition FMV, i.e. such that $\{k \in \mathbb{N} : A_k \cap \mathbb{X} \neq \emptyset\}$ is infinite with nonzero probability. Denote by \mathcal{A} the event that \mathbb{X} takes an infinite number of values, i.e. $\mathcal{A} = \{\#\mathbb{X} = +\infty\}$. We have assumed that $\mathbb{P}(\mathcal{A}) > 0$ and will condition on \mathcal{A} for the rest of the proof. For $k \in \mathbb{N}$, define $N_k \in \mathbb{N}$ such that,

$$\mathbb{P}\left(\#\mathbb{X}_{\leq N_k} \geq \frac{1}{\mu_k^2} \mid \mathcal{A}\right) \geq 1 - \frac{1}{2^{k+1}}, \quad \text{where } \mu_k := \frac{1}{2^{k+1}}.$$

Note that N_k is a deterministic quantity that only depends on the process \mathbb{X} . It is well defined because $\mathbb{P}\left(\#\mathbb{X}_{\leq N} \geq \frac{1}{\mu_k^2} \mid \mathcal{A}\right) \rightarrow_{N \rightarrow \infty} 1$ since \mathbb{X} takes an infinite number of values in \mathcal{A} . Now also define $0 < \delta_k < \frac{1}{2^{k+1}}$ satisfying:

$$\mathbb{P}\left(\min_{1 \leq i, j \leq N_k, X_i \neq X_j} |X_i - X_j| > \delta_k \mid \mathcal{A}\right) \geq 1 - \frac{1}{2^{k+1}}.$$

Let \mathcal{E}_k be the intersection of the two events above, we have by union bound: $\mathbb{P}(\mathcal{E}_k \mid \mathcal{A}) \geq 1 - \frac{1}{2^k}$, where \mathcal{E}_k can be written as

$$\mathcal{E}_k : \quad \#\mathbb{X}_{\leq N_k} > \frac{1}{\mu_k^2} \quad \text{and} \quad \forall x \neq y \in \mathbb{X}_{\leq N_k}, |x - y| > \delta_k.$$

We are now ready to construct the partition. Let $\mathbf{q} = (q_i)_{i \geq 1}$ be an i.i.d. sequence of independent uniforms sampled from $\mathcal{U}([0, 1])$. Define $B_0 = [0, 1]$ and for $k \geq 1$,

$$B_k = \bigcup_{i_{k-1} < i \leq i_k} \left[q_i - \frac{\delta_k}{2}, q_i + \frac{\delta_k}{2} \right],$$

where $(i_k)_{k \geq 0}$ satisfies $i_0 = 0$ and $i_k = i_{k-1} + \lceil \frac{\mu_k}{\delta_k} \rceil$. Note that the Borel measure of B_k is roughly μ_k : $\mu(B_k) \leq \mu_k + \delta_k$. We use the remainders $R_k = [0, 1] \cap \bigcup_{l \geq k} B_l$ to define the partition $\{A_k\}_{k=-1}^\infty$ as follows:

$$A_k = R_k \setminus R_{k+1}, \quad k \geq 0$$

and $A_{-1} = \bigcap_{k \geq 0} R_k$. The sets $\{A_k\}_{k=-1}^\infty$ define a proper partition of \mathcal{X} since A_{-1} contains elements that appear infinitely often in $\{B_k\}_{k \geq 0}$ while for any $k \geq 0$, A_k contains elements that appear for the last time in the sequence $\{B_l\}_{l \geq 0}$ in B_k . This covers the whole space \mathcal{X} because by construction $B_0 = \mathcal{X}$. The interest of this (random) construction lies in the two following lemmas.

Lemma 8 *For any finite (deterministic) $S \subset [0, 1]$ with $\#S > \frac{1}{\mu_k^2}$ and $\forall x \neq y \in S, |x - y| > \delta_k$,*

$$\mathbb{P}(B_k \cap S = \emptyset) \leq e^{-2^{k+1}}.$$

Lemma 9 For any countable (deterministic) $S \subset [0, 1]$, $A_{-1} \cap S = \emptyset$, (a.s.)

We will now show that with probability $\mathbb{P}(\mathcal{A})$, the partition $\{A_k\}_{k=-1}^\infty$ disproves the condition FMV, in other terms that \mathbb{X} visits an infinite number of sets of the partition. Recall that the randomness is now both in terms of the stochastic process \mathbb{X} and the partition generated from \mathbf{q} . We have that,

$$\mathbb{P}(B_k \cap \mathbb{X} = \emptyset \mid \mathcal{A}) \leq (1 - \mathbb{P}(\mathcal{E}_k \mid \mathcal{A})) + \mathbb{P}(B_k \cap \mathbb{X}_{\leq N_k} = \emptyset \mid \mathcal{E}_k, \mathcal{A}) \leq \frac{1}{2^k} + e^{-2^{k+1}},$$

where in the last inequality we applied Lemma 8 to the set $\mathbb{X}_{\leq N_k}$ which has cardinality at least $\frac{1}{\mu_k^2}$ in \mathcal{E}_k . We can now apply the first Borel-Cantelli lemma to the sequence of events $\{B_k \cap \mathbb{X} = \emptyset\}$ conditionally on \mathcal{A} , which shows that almost surely only a finite number of these events are satisfied. Hence, conditionally on \mathcal{A} , there exists almost surely $\kappa \in \mathbb{N}$ such that for every $k \geq \kappa$, the sequence \mathbb{X} visits B_k . Further, by Lemma 9, with probability 1, \mathbb{X} does not visit A_{-1} . Therefore, conditionally on \mathcal{A} , the sequence almost surely visits an infinite number of sets of the partition $\{A_k\}_{k=-1}^\infty$. In summary,

$$\mathbb{P}_{\mathbf{q}, \mathbb{X}}(\#\{k \in \mathbb{N} : A_k \cap \mathbb{X} \neq \emptyset\} = +\infty) \geq \mathbb{P}(\mathcal{A}).$$

Thus, there exists a *deterministic* choice of \mathbf{q} yielding a partition $\{A_k\}_{k=-1}^\infty$ such that:

$$\mathbb{P}_{\mathbb{X}}(\#\{k \in \mathbb{N}, A_k \cap \mathbb{X} \neq \emptyset\} = +\infty) \geq \mathbb{P}(\mathcal{A}) > 0.$$

This shows the claim of the theorem. ■

Proof of Lemma 8 Note that the randomness of \mathbb{X} does not intervene in this lemma. The probability law \mathbb{P} only accounts for the randomness of the partition through the variables $\mathbf{q} = (q_i)_{i \geq 1}$. We enumerate $S = \{x_1, \dots, x_T\}$ where $T = \#S \geq \frac{1}{\mu_k^2}$.

$$\mathbb{P}(B_k \cap S = \emptyset) = \mathbb{P}(x_1 \notin B_k) \prod_{t=2}^T \mathbb{P}(x_t \notin B_k \mid x_1, \dots, x_{t-1} \notin B_k).$$

For the sake of simplicity we will use the notation $B(x, \delta) = [x - \delta, x + \delta]$. Note that events $\{x_i \notin B_k\}$ are negatively correlated. Indeed,

$$\begin{aligned} \mathbb{P}(x_t \notin B_k \mid x_1, \dots, x_{t-1} \notin B_k) &= \prod_{i=i_{k-1}+1}^{i_k} \mathbb{P} \left[q_i \notin B \left(x_t, \frac{\delta_k}{2} \right) \mid q_i \notin \bigcup_{1 \leq l \leq t-1} B \left(x_l, \frac{\delta_k}{2} \right) \right] \\ &= \prod_{i=i_{k-1}+1}^{i_k} \mathbb{P} \left[\tilde{q}_i \notin B \left(x_t, \frac{\delta_k}{2} \right) \right] \end{aligned}$$

where $\tilde{q}_i \sim \mathcal{U}(J)$ with $J := \mathcal{X} \setminus \bigcup_{1 \leq l \leq t-1} B(x_l, \frac{\delta_k}{2})$. Because $|x_t - x_l| > \delta_k$ for all $1 \leq l \leq t-1$, we have $B(x_t, \frac{\delta_k}{2}) \subset J$. Thus,

$$\mathbb{P}(x_t \notin B_k \mid x_1, \dots, x_{t-1} \notin B_k) = \prod_{i=i_{k-1}+1}^{i_k} \left(1 - \frac{\delta_k}{\mu(J)} \right) \leq \prod_{i=i_{k-1}+1}^{i_k} (1 - \delta_k) = \mathbb{P}(x_t \notin B_k).$$

Using the negative correlation, we have that

$$\mathbb{P}(B_k \cap S = \emptyset) \leq \prod_{t=1}^T \mathbb{P}(x_1 \notin B_k) = (1 - \delta_k)^{T(i_k - i_{k-1})} \leq (1 - \delta_k)^{\frac{1}{\mu_k \delta_k}} \leq e^{-\frac{1}{\mu_k}} = e^{-2^{k+1}}.$$

This ends the proof of the lemma. \blacksquare

Proof of Lemma 9 We start by proving that for a given $x \in \mathbb{R}$, $x \notin A_{-1}$ a.s. For $k \geq 1$ we have,

$$\mathbb{P}(x \in B_k) \leq \left\lceil \frac{\mu_k}{\delta_k} \right\rceil \delta_k \leq \left(\frac{\mu_k}{\delta_k} + 1 \right) \delta_k \leq \mu_k + \delta_k \leq \frac{1}{2^k}.$$

Therefore, $\mathbb{P}(x \in R_k) \leq \frac{1}{2^{k-1}}$. This shows that $\mathbb{P}(x \in A_{-1}) \leq \mathbb{P}(x \in \bigcap_k R_k) = 0$. Taking the union over all countable random variables in S , we have $\mathbb{P}(A_{-1} \cap S \neq \emptyset) = 0$. \blacksquare

Extension to all standard Borel spaces. Before we move on to proving the main theorem in the most general framework of separable metric spaces, observe that the proof for $\mathcal{X} = [0, 1]$ easily extends to all standard Borel space by Kuratowski's theorem, in particular for instance to $\mathcal{X} = \mathbb{R}^d$. Kuratowski's theorem states that if \mathcal{X} is an uncountable standard Borel space it is isomorphic to $[0, 1]$ with the Euclidean distance, meaning that there exists a measurable bijection $f : \mathcal{X} \rightarrow [0, 1]$. Let $\mathbb{X} = (X_i)_{i \geq 0}$ be a stochastic process on \mathcal{X} satisfying FMV. Then, because f is measurable, $\tilde{\mathbb{X}} := (f(X_i))_{i \geq 0}$ is a stochastic process on $[0, 1]$ which satisfies FMV. By Theorem 7, $\tilde{\mathbb{X}}$ satisfies FS. Thus, because f is bijective, \mathbb{X} also satisfies FS.

4. Extension to General Separable Metric Spaces

The original proof of SUOL = FMV by [13] holds for any separable metric space (\mathcal{X}, ρ) . In this section, we extend the proof above to hold in this more-general case as well, thus completely answering the question (Open Problem 4) posed by [13] in full generality, and completing the proof of Theorems 1 and 2. For the remainder of this section, we let (\mathcal{X}, ρ) denote a non-empty separable metric space, and we take as the set \mathcal{B} of measurable subsets of \mathcal{X} the Borel σ -algebra generated by the topology induced by ρ .

Theorem 6 (Restated) For any separable metric space (\mathcal{X}, ρ) , FMV = FS.

The main components of the proof are analogous to those for standard Borel spaces, with a few important changes: most importantly, the following lemma.

Lemma 10 For any \mathbb{X} satisfying condition FMV, for any $\delta, \varepsilon > 0$ and $m_0 \in \mathbb{N}$, there exists $M_{\varepsilon, \delta} \in \mathbb{N}$ with $M_{\varepsilon, \delta} \geq m_0$, and a sequence $\mathcal{G}^{\varepsilon, \delta} = \{G_1^{\varepsilon, \delta}, \dots, G_{M_{\varepsilon, \delta}}^{\varepsilon, \delta}\}$ in \mathcal{B} such that every distinct $i, j \in \{1, \dots, M_{\varepsilon, \delta}\}$ satisfy $G_i^{\varepsilon, \delta} \cap G_j^{\varepsilon, \delta} = \emptyset$, and every $i \in \{1, \dots, M_{\varepsilon, \delta}\}$ satisfies $\sup_{x, x' \in G_i^{\varepsilon, \delta}} \rho(x, x') \leq \delta$, and such that

$$\mathbb{P} \left(\mathbb{X} \cap \left(\mathcal{X} \setminus \bigcup_{i=1}^{M_{\varepsilon, \delta}} G_i^{\varepsilon, \delta} \right) \neq \emptyset \right) < \varepsilon.$$

In other words, $\mathcal{G}^{\varepsilon, \delta}$ is a sequence of disjoint measurable sets of diameter at most δ , which cover all of the points in \mathbb{X} with probability $1 - \varepsilon$.

Proof Let $\tilde{\mathcal{X}} \subseteq \mathcal{X}$ be a countable dense subset: that is, $\sup_{x \in \mathcal{X}} \inf_{\tilde{x} \in \tilde{\mathcal{X}}} \rho(\tilde{x}, x) = 0$. Enumerate $\tilde{\mathcal{X}}$ as $\{\tilde{x}_1, \tilde{x}_2, \dots\}$. Let $G_1^{\varepsilon, \delta} = \{x : \rho(x, \tilde{x}_1) \leq \delta/2\}$, and for integers $k \geq 2$ inductively define $G_k^{\varepsilon, \delta} = \{x : \rho(x, \tilde{x}_k) \leq \delta/2\} \setminus \bigcup_{k'=1}^{k-1} G_{k'}^{\varepsilon, \delta}$. In particular, this collection $\{G_k^{\varepsilon, \delta} : k \in \mathbb{N}\}$ forms a countable partition of \mathcal{X} into measurable subsets of diameter at most δ (by the triangle inequality). Now let \mathbb{X} be any process satisfying FMV. It remains only to show there exists a finite $M_{\varepsilon, \delta} \in \mathbb{N}$ satisfying the claim. Let $\hat{M} = \max\{k : \mathbb{X} \cap G_k^{\varepsilon, \delta} \neq \emptyset\}$, or $\hat{M} = \infty$ if there is no maximum. By hypothesis, $\mathbb{P}(\hat{M} < \infty) = 1$. Since the event $\{\hat{M} > M\}$ is non-increasing in M , $\lim_{M \rightarrow \infty} \mathbb{P}(\hat{M} > M) = \mathbb{P}(\hat{M} = \infty) = 0$. Thus, $\exists M_{\varepsilon, \delta} \in \mathbb{N}$ with $M_{\varepsilon, \delta} \geq m_0$ such that $\mathbb{P}(\hat{M} > M_{\varepsilon, \delta}) < \varepsilon$. In other words, $\mathbb{P}(\exists k > M_{\varepsilon, \delta} : \mathbb{X} \cap G_k^{\varepsilon, \delta} \neq \emptyset) < \varepsilon$. Since $\{G_k^{\varepsilon, \delta} : k \in \mathbb{N}\}$ is a partition of \mathcal{X} , this implies the claim in the lemma. \blacksquare

We are now ready for the main proof.

Proof of Theorem 6 Since condition FS clearly implies condition FMV, we focus on showing $\text{FMV} \subset \text{FS}$. Let \mathbb{X} be any process satisfying condition FMV, and for the sake of obtaining a contradiction, suppose that condition FS fails: that is, there is an event \mathcal{A} with $\mathbb{P}(\mathcal{A}) > 0$, on which $\#\{x \in \mathcal{X} : \mathbb{X} \cap \{x\} \neq \emptyset\} = \infty$.

For each $k \in \mathbb{N}$, let $N_k \in \mathbb{N}$ be such that

$$\mathbb{P}\left(\#\mathbb{X}_{\leq N_k} \geq 2^{2k+2} \mid \mathcal{A}\right) \geq 1 - \frac{1}{2^{k+2}},$$

and let $\delta_k > 0$ be such that

$$\mathbb{P}\left(\min_{i, j \leq N_k : X_i \neq X_j} \rho(X_i, X_j) > \delta_k \mid \mathcal{A}\right) \geq 1 - \frac{1}{2^{k+3}}.$$

Let $S_k = \{x \in \mathcal{X} : \mathbb{X}_{\leq N_k} \cap \{x\} \neq \emptyset\}$ and let $\varepsilon_k = \frac{1}{2^{k+3}}$. Let $\mathcal{G}^{\varepsilon_k, \delta_k}$ and $M_{\varepsilon_k, \delta_k}$ be as in Lemma 10, with $m_0 = 2^{k+2}$.

Let \mathcal{E}_k denote the event that $\#\mathbb{X}_{\leq N_k} \geq 2^{2k+2}$, $\min_{i, j \leq N_k : X_i \neq X_j} \rho(X_i, X_j) > \delta_k$, and $\mathbb{X} \cap (\mathcal{X} \setminus \bigcup \mathcal{G}^{\varepsilon_k, \delta_k}) = \emptyset$ all hold simultaneously. In particular, by the union bound, $\mathbb{P}(\mathcal{E}_k \mid \mathcal{A}) \geq 1 - 2^{-k-1}$.

For each $k \in \mathbb{N}$, let $b_k = \lceil 2^{-k-2} M_{\varepsilon_k, \delta_k} \rceil$, and let $Q_1^k, \dots, Q_{b_k}^k$ be independent uniform samples from $\mathcal{G}^{\varepsilon_k, \delta_k}$ (also independent across k and independent from \mathbb{X}). Then let $B_k = \bigcup_{i=1}^{b_k} Q_i^k$. For each $k \in \mathbb{N}$, let $R_k = \bigcup_{\ell \geq k} B_\ell$. Also let $A_{-1} = \bigcap_{k \in \mathbb{N}} R_k$ and for each $k \in \mathbb{N}$, let $A_k = R_k \setminus R_{k+1}$, and $A_0 = \mathcal{X} \setminus R_1$. We will show that (with non-zero probability) the countable measurable partition $\{A_k : k \in \mathbb{N} \cup \{-1, 0\}\}$ violates the condition FMV, thus obtaining a contradiction.

Now note that, on the event \mathcal{E}_k , every $x \in S_k$ is in a distinct set $G_i^{\varepsilon_k, \delta_k} \in \mathcal{G}^{\varepsilon_k, \delta_k}$: that is, by definition of \mathcal{E}_k , every $x \in S_k$ is in some $G_i^{\varepsilon_k, \delta_k} \in \mathcal{G}^{\varepsilon_k, \delta_k}$, and since each $G_i^{\varepsilon_k, \delta_k}$ has diameter at most δ_k , while every distinct $x, x' \in S_k$ are δ_k -separated (on event \mathcal{E}_k), no two elements of S_k can be in the same $G_i^{\varepsilon_k, \delta_k}$. Therefore, on the event \mathcal{E}_k we have that

$$\mathbb{P}(B_k \cap S_k = \emptyset \mid \mathbb{X}) = \mathbb{P}\left(Q_1^k \cap S_k = \emptyset \mid \mathbb{X}\right)^{b_k} = \left(1 - \frac{|S_k|}{M_{\varepsilon_k, \delta_k}}\right)^{b_k} \leq e^{-|S_k| b_k / M_{\varepsilon_k, \delta_k}} \leq e^{-2^k},$$

where the last inequality is based on the definition of b_k and the fact that $|S_k| \geq 2^{2k+2}$ on the event \mathcal{E}_k . Thus, on the event $\bigcap_{k \in \mathbb{N}} \mathcal{E}_k$,

$$\sum_{k=1}^{\infty} \mathbb{P}(B_k \cap S_k = \emptyset | \mathbb{X}) \leq \sum_{k=1}^{\infty} e^{-2^k} < \infty.$$

By the Borel-Cantelli lemma, this implies that there is an event \mathcal{E}' of probability one, such that on $\mathcal{E}' \cap \bigcap_{k \in \mathbb{N}} \mathcal{E}_k$, there exists $\kappa \in \mathbb{N}$ such that every $k \geq \kappa$ satisfies $B_k \cap S_k \neq \emptyset$, and hence also $\mathbb{X} \cap R_k \neq \emptyset$. Now, if $\mathbb{X} \cap A_{-1} = \emptyset$, this would further imply that $|\{k \in \mathbb{N} : \mathbb{X} \cap A_k \neq \emptyset\}| = \infty$.

We next turn to showing that $\mathbb{X} \cap A_{-1} = \emptyset$ (a.s.). For any $t, k \in \mathbb{N}$, by the union bound, $\mathbb{P}(X_t \in B_k) \leq \frac{b_k}{M_{\varepsilon_k, \delta_k}} \leq 2^{-k-1}$ (recalling that $M_{\varepsilon_k, \delta_k} \geq 2^{k+2}$, so that $b_k \leq 2^{-k-1} M_{\varepsilon_k, \delta_k}$). By the union bound, this further implies any $t, k \in \mathbb{N}$ satisfy $\mathbb{P}(X_t \in R_k) \leq \sum_{\ell \geq k} \mathbb{P}(X_t \in B_\ell) \leq \sum_{\ell \geq k} 2^{-\ell-1} = 2^{-k}$. Thus, $\mathbb{P}(X_t \in A_{-1}) = \mathbb{P}(X_t \in \bigcap_{k \in \mathbb{N}} R_k) \leq \lim_{k \rightarrow \infty} \mathbb{P}(X_t \in R_k) = 0$. By the union bound, $\mathbb{P}(\mathbb{X} \cap A_{-1} \neq \emptyset) = 0$. Thus, there is an event \mathcal{E}'' of probability one, on which $\mathbb{X} \cap A_{-1} = \emptyset$.

Altogether, we have that on the event $\mathcal{E}' \cap \mathcal{E}'' \cap \bigcap_{k \in \mathbb{N}} \mathcal{E}_k$, $|\{k \in \mathbb{N} : \mathbb{X} \cap A_k \neq \emptyset\}| = \infty$. Since $\mathbb{P}(\mathcal{E}') = \mathbb{P}(\mathcal{E}'') = 1$, and

$$\begin{aligned} \mathbb{P}\left(\bigcap_{k \in \mathbb{N}} \mathcal{E}_k\right) &\geq \mathbb{P}\left(\mathcal{A} \cap \bigcap_{k \in \mathbb{N}} \mathcal{E}_k\right) \geq \mathbb{P}(\mathcal{A}) - \sum_{k \in \mathbb{N}} \mathbb{P}(\mathcal{A}) (1 - \mathbb{P}(\mathcal{E}_k | \mathcal{A})) \\ &\geq \mathbb{P}(\mathcal{A}) - \sum_{k \in \mathbb{N}} \mathbb{P}(\mathcal{A}) 2^{-k-1} = \frac{1}{2} \mathbb{P}(\mathcal{A}), \end{aligned}$$

by the union bound we have $\mathbb{P}(\mathcal{E}' \cap \mathcal{E}'' \cap \bigcap_{k \in \mathbb{N}} \mathcal{E}_k) \geq \frac{1}{2} \mathbb{P}(\mathcal{A}) > 0$. In particular, this implies $\mathbb{P}(|\{k \in \mathbb{N} : \mathbb{X} \cap A_k \neq \emptyset\}| = \infty) > 0$. Moreover, by the law of total probability,

$$\mathbb{P}(|\{k \in \mathbb{N} : \mathbb{X} \cap A_k \neq \emptyset\}| = \infty) = \mathbb{E}\left[\mathbb{P}\left(|\{k \in \mathbb{N} : \mathbb{X} \cap A_k \neq \emptyset\}| = \infty \mid \{A_k : k \in \mathbb{N}\}\right)\right],$$

and hence (since \mathbb{X} is independent of the random partition $\{A_k : k \in \mathbb{N} \cup \{-1, 0\}\}$), there exists a *deterministic* choice of a partition $\{\hat{A}_k : k \in \mathbb{N} \cup \{-1, 0\}\}$ such that

$$\mathbb{P}\left(|\{k \in \mathbb{N} \cup \{-1, 0\} : \mathbb{X} \cap \hat{A}_k \neq \emptyset\}| = \infty\right) > 0,$$

contradicting condition FMV. This completes the proof. \blacksquare

5. Consequences on inductive and self-adaptive learning

Along with optimistically universal *online* learning, [13] identifies two other learning setups, namely *inductive learning* and *self-adaptive learning*.

Inductive learning. An inductive learning rule $\{f_t\}_{t=1}^{\infty}$ is a sequence of measurable functions $f_t : \mathcal{X}^{t-1} \times \mathcal{Y}^{t-1} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that given training data $(\mathbb{X}_{<t}, \mathbb{Y}_{<t})$ and input point $X_{t'}$ with $t' > t$ outputs prediction $f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_{t'})$. Its performance is measured in terms of,

$$\mathcal{L}_{\mathbb{X}}(f_t, f^*; t) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t'=t}^{t+T} \ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_{t'}), f^*(X_{t'})).$$

Let SUIL denote the set of all processes \mathbb{X} that admit strong universal inductive learning: i.e., for which there exists an inductive learning rule $\{f_t\}$ such that for every measurable $f^* : \mathcal{X} \rightarrow \mathcal{Y}$, $\mathcal{L}_{\mathbb{X}}(f_t, f^*; t) \rightarrow 0$ (a.s.). Note that the difference between an online learning rule and its inductive counterpart is that the latter will be fixed for an infinite horizon. It was therefore shown by [13] that $\text{SUIL} \subset \text{SUOL}$.

Self adaptive learning rule. A self-adaptive learning rule $\{f_{t_1, t_2}\}_{t_1 \leq t_2}^\infty$ is a sequence of measurable functions $f_{t_1, t_2} : \mathcal{X}^{t_2-1} \times \mathcal{Y}^{t_1-1} \times \mathcal{X} \rightarrow \mathcal{Y}$ such that given training data $(\mathbb{X}_{<t_2}, \mathbb{Y}_{<t_1})$ and input point X_{t_2} it performs prediction $f_{t_1, t_2}(\mathbb{X}_{<t_2}, \mathbb{Y}_{<t_1}, X_{t_2})$. Its performance is measured in terms of

$$\mathcal{L}_{\mathbb{X}}(f_{t_1, \cdot}, f^*; t_1) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t_2=t_1}^{t_1+T} \ell(f_{t_1, t_2}(\mathbb{X}_{<t_2}, \mathbb{Y}_{<t_1}, X_{t_2}), f^*(X_{t_2})).$$

Let SUAL denote the set of all processes \mathbb{X} that admit strong universal self-adaptive learning: i.e., for which there exists a self-adaptive learning rule $\{f_{t_1, t_2}\}$ such that for every measurable $f^* : \mathcal{X} \rightarrow \mathcal{Y}$, $\mathcal{L}_{\mathbb{X}}(f_{t_1, \cdot}, f^*; t_1) \rightarrow 0$ (a.s.). Note that self-adaptive learning rules are more expressive than inductive learning rules for they have access to additional unlabeled data, like in the semi-supervised learning setup studied in the literature [4], yet are still less powerful than online learning rules which would also have access to the respective labels. It was therefore shown by [13] that $\text{SUIL} \subset \text{SUAL} \subset \text{SUOL}$.

Consequence of the Main Theorem. For unbounded losses, [13] shows that $\text{SUIL} = \text{SUAL} = \text{SUOL}$. However, once again this proof relied on the aforementioned complicated arguments. But in light of our proof that $\text{SUOL} = \text{FS}$, it becomes immediately apparent that $\text{SUIL} = \text{SUAL} = \text{SUOL}$ and that these classes all admit memorization as an optimistically universal learning rule (merely noting that $\text{FS} \subset \text{SUIL}$, since for any $\mathbb{X} \in \text{FS}$, the inductive loss of the memorization rule f_t becomes zero once t exceeds the index of the last novel data point). This greatly simplifies the proof of these equivalences compared to the original proof of [13]. Note that while these three setups turn out to be equivalent when the loss is unbounded, interesting distinctions do exist in the bounded case for which [13] proved that there exists an optimistically universal self-adaptive learning rule (which surprisingly is necessarily different from nearest-neighbour), but no optimistically universal inductive learning rule.

6. Discussion on noise : an optimistically universal Bayes consistent learner

For simplicity we restricted the analysis to the *realizable* setting [17; 1] where there exists a measurable function f^* satisfying $\forall t \geq 1 : Y_t = f^*(X_t)$. A common variant allows for the function f^* to be *noisy*, i.e. to take the form of a conditional probability density $p_{Y|X}$. In this section, we will generalise the main results to this context. We start by recalling the adequate notion of consistency, e.g. [13; 15; 25]. We say that the learning rule f is strongly universally Bayes consistent under \mathbb{X} if for all conditional probability distribution $p_{Y|X}$ and any measurable function \bar{f} , almost surely

$$\mathcal{L}_{\mathbb{X}}(f, p_{Y|X}; T, \bar{f}) \leq 0$$

where $\hat{\mathcal{L}}_{\mathbb{X}}(f, p_{Y|X}; T, \bar{f})$ is the excess loss of the learning rule f against the constant predictor \bar{f} ,

$$\hat{\mathcal{L}}_{\mathbb{X}}(f, p_{Y|X}; T, \bar{f}) = \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T (\ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) - \ell(\bar{f}(X_t), Y_t)).$$

Theorem 11 *If (\mathcal{Y}, d) is a separable locally compact metric space with $\bar{d} = \infty$, there exists an optimistically universal Bayes consistent learning rule for the loss $\ell = d^p$, for any $p \geq 1$.*

Note that the Bayesian setting is more general than the *realizable* setting. Thus, any process \mathbb{X} that admits a universally Bayes consistent learning rule is in SUOL, and therefore it takes a finite number of values almost surely. To prove Theorem 11, it will suffice to define a learning rule that is universally Bayes consistent with any such process. For this purpose, we use a result from [7] which we slightly adapt to our setting. The proof of the following theorem can be found in Appendix C.

Theorem 12 ([7]) *Let $p \geq 1$ and (\mathcal{Y}, d) be a separable metric space such that any closed ball is compact. Let $(Y_i)_{i \geq 1}$ be an i.i.d. sequence of random variables in \mathcal{Y} of distribution Y satisfying $\mathbb{E}d^p(y_0, \mathcal{Y}) < \infty$ for some $y_0 \in \mathcal{Y}$. Denote \hat{y}_n a Fréchet sample mean of the samples Y_1, \dots, Y_n . Then,*

$$\frac{1}{n} \sum_{i=1}^n d^p(\hat{y}_n, Y_i) \rightarrow \min_{y \in \mathcal{Y}} \mathbb{E}d^p(y, Y) \quad (a.s.).$$

In the equation above, we use the notion of Fréchet sample mean that is defined as follows,

$$\hat{y}_n \in \operatorname{argmin}_{y \in \mathcal{Y}} \frac{1}{n} \sum_{i=1}^n d^p(y, Y_i).$$

Note that the minimum is well defined because the closed balls in the space (\mathcal{Y}, d) are compact and $d^p(y, y_0) \leq 2^{p-1} \frac{1}{n} \sum_{i=1}^n (d^p(y, Y_i) + d^p(y_0, Y_i))$, hence $\frac{1}{n} \sum_{i=1}^n d^p(y, Y_i) \geq 2^{1-p} d^p(y, y_0) - \frac{1}{n} \sum_{i=1}^n d^p(y_0, Y_i)$ for any $y_0 \in \mathcal{Y}$. Therefore the expression is minimized in the (compact) closed ball of radius at most $2 \left[\frac{1}{n} \sum_{i=1}^n d^p(y_0, Y_i) \right]^{1/p}$ around y_0 i.e. $d^p(\hat{y}_n, y_0) \leq \frac{2^p}{n} \sum_{i=1}^n d^p(y_0, Y_i)$. Similarly, the infimum $\inf_{y \in \mathcal{Y}} \mathbb{E}d^p(y, Y)$ is attained because for $y_0 \in \mathcal{Y}$ such that $\mathbb{E}d^p(y_0, Y) < \infty$, the quantity $\mathbb{E}d^p(y, Y)$ is minimized in the compact closed ball around y_0 of radius $2[\mathbb{E}d^p(y_0, Y)]^{1/p}$.

The discussion above allows to define the *Fréchet mean memorizer* learning rule, which is proved to be an optimistically universal Bayes consistent learner in Appendix D. Fix an arbitrary $y_0 \in \mathcal{Y}$,

$$f_t(\mathbf{x}_{<t}, \mathbf{y}_{<t}, x_t) = \begin{cases} y \in \operatorname{argmin}_{y \in \mathcal{Y}} \sum_{i=1}^{t-1} \mathbb{1}_{x_i=x_t} \ell(y, y_i), & \text{if } x_t \in \mathbf{x}_{<t}, \\ y_0 & \text{if } x_t \notin \mathbf{x}_{<t}. \end{cases}$$

Note that the discussion above covers the case of real regression with $\mathcal{Y} = \mathbb{R}$ and $\ell(y_1, y_2) = (y_1 - y_2)^2$. In this case, the Fréchet sample mean from Theorem 12 is the empirical average $\hat{y}_n = \frac{1}{n} \sum_{i=1}^n Y_i$ and the Bayes consistency of the corresponding learning rule can be obtained directly from the strong law of large numbers. In the setup where \mathbb{X} is an iid process, the question of Bayes consistency is addressed by [15; 25]. In particular, [25] uses the notion of medoid that adapts Fréchet means to compression-based algorithms.

7. Conclusion and Open Directions

In this paper, we showed that memorization is an optimistically universal learning rule when the loss is unbounded. This closes the study of unrestricted universal consistency with unbounded losses

for the *online*, the *inductive* and the *self-adaptive* setups. In a sense, our result may be viewed as a *negative* result, revealing that for unbounded losses, the processes in SUOL are all, to some extent, rather trivial. On the other hand, we know of many positive results for universal consistency with unbounded losses for i.i.d. or stationary ergodic processes, under *additional conditions* on the Y_t sequence, such as with *moment* conditions on Y_t in the regression setting [10; 11]. Thus, it would seem the next chapter in the study of universal consistency with unbounded losses and general non-i.i.d. families of processes should be to formulate broad sufficient conditions on the Y_t sequence (relative to the given X_t sequence) so that the family of processes \mathbb{X} admitting universal learning becomes rich, and in particular, includes within it all i.i.d. or stationary ergodic processes \mathbb{X} . It would be particularly interesting if there is a moment condition on the Y_t sequence (or more-generally, on the empirical moments of $\ell(y_0, Y_t)$ for some y_0), under which the set of all processes \mathbb{X} admitting strong universal learning are precisely the same as for the case of *bounded* losses: i.e., in the case of online learning, the set SUOL that would result from learning with a bounded loss (see [13; 14] for discussions regarding this set), or in the case of inductive or self-adaptive learning, the sets SUIL or SUAL, respectively, that would result from learning with a bounded loss (which have been characterized by [13]).

For concreteness, focusing on the setting of online learning, and letting SUOL_{01} denote the set of processes \mathbb{X} that admit strong universal online learning under the 0-1 loss for binary classification (i.e., $\mathcal{Y} = \{0, 1\}$ and $\ell(y, y') = \mathbb{1}[y \neq y']$), we ask the following question:

Open Problem: For every unbounded loss ℓ , does there exist an online learning rule $\{f_t\}_{t=1}^{\infty}$ with the property that, for every $\mathbb{X} \in \text{SUOL}_{01}$, we have $\mathcal{L}_{\mathbb{X}}(f, f^*) = 0$ (a.s.) for every measurable $f^* : \mathcal{X} \rightarrow \mathcal{Y}$ such that, for every $y_0 \in \mathcal{Y}$, $\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ell(y_0, f^*(X_t)) < \infty$ (a.s.)? In other words, f^* has empirically bounded losses in the long-run average.

If this is found to be true, it would generalize the known conditions on consistent regression (for deterministic functions) with the squared loss under i.i.d. processes with finite $\text{Var}(Y)$ [10]. Of course, even variations of the above problem using milder restrictions on (\mathbb{X}, f^*) would be interesting. As such, we may essentially pose a relaxed version of the question, which replaces the last condition with the mere requirement that, for every $y_0 \in \mathcal{Y}$, $\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \ell(y_0, f^*(X_t))^p < \infty$ (a.s.), where $p > 0$ is some constant: i.e., any limiting empirical *moment* condition.

In the case of *bounded* losses, both the question of the existence of optimistically universal online learning rules, and of concisely characterizing the set SUOL, remain open, as recently highlighted in a COLT open problem article [14]. In particular, [14] conjectures that for bounded losses, a process \mathbb{X} is in SUOL if and only if it has the property that, for any countable measurable partition $\{\mathcal{A}_k\}_{k=1}^{\infty}$ of \mathcal{X} , the number of visited sets grows sub-linearly with time: $\#\{k \in \mathbb{N} : \mathcal{A}_k \cap \mathbb{X}_{\leq T}\} = o(T)$ (a.s.). The success of memorization in the unbounded setup suggests that a simple rule such as nearest neighbour might possibly be optimistically universal for online learning with bounded losses. However, such intuitions can also be misleading. In the self-adaptive setup with bounded losses, [13] proved that nearest neighbour is not optimistically universal although another more-intricate learning rule is optimistically universal.

Acknowledgments

M. Blanchard was partly funded by ONR grant N00014-18-1-2122.

References

- [1] Shai Ben-David, Dávid Pál, and Shai Shalev-Shwartz. Agnostic online learning. In *COLT*, volume 3, page 1, 2009.
- [2] Gérard Biau, Kevin Bleakely, László Györfi, and György Ottucsák. Nonparametric sequential prediction of time series. *Journal of Nonparametric Statistics*, 22:297–317, 2010.
- [3] Olivier Bousquet, Steve Hanneke, Shay Moran, Ramon van Handel, and Amir Yehudayoff. A theory of universal learning. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 532–541, 2021.
- [4] Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. Semi-supervised learning. *IEEE Transactions on Neural Networks*, 20(3):542–542, 2009.
- [5] Thomas Cover and Peter Hart. Nearest neighbor pattern classification. *IEEE transactions on information theory*, 13(1):21–27, 1967.
- [6] Luc Devroye, László Györfi, and Gábor Lugosi. *A probabilistic theory of pattern recognition*, volume 31. Springer Science & Business Media, 2013.
- [7] Steven N Evans and Adam Q Jaffe. Strong laws of large numbers for Fréchet means. *arXiv preprint arXiv:2012.12859*, 2020.
- [8] László Györfi and Gábor Lugosi. Strategies for sequential prediction of stationary time series. In *Modeling uncertainty*, pages 225–248. Springer, 2002.
- [9] L Györfi, Gábor Lugosi, and Gusztáv Morvai. A simple randomized algorithm for sequential prediction of ergodic time series. *IEEE Transactions on Information Theory*, 45(7):2642–2650, 1999.
- [10] L. Györfi, M. Kohler, A. Krzyżak, and H. Walk. *A Distribution-Free Theory of Nonparametric Regression*. Springer-Verlag New York, 2002.
- [11] László Györfi and György Ottucsák. Sequential prediction of unbounded stationary time series. *IEEE Transactions on Information Theory*, 53(5):1866–1872, 2007.
- [12] László Györfi and György Ottucsák. Nonparametric sequential prediction of stationary time series. *Machine Learning For Financial Engineering*, pages 179–226, 2012.
- [13] Steve Hanneke. Learning whenever learning is possible: Universal learning under general stochastic processes. *Journal of Machine Learning Research*, 22(130):1–116, 2021.
- [14] Steve Hanneke. Open problem: Is there an online learning algorithm that learns whenever online learning is possible? In *Conference on Learning Theory*, pages 4642–4646. PMLR, 2021.

- [15] Steve Hanneke, Aryeh Kontorovich, Sivan Sabato, and Roi Weiss. Universal bayes consistency in metric spaces. In *2020 Information Theory and Applications Workshop (ITA)*, pages 1–33. IEEE, 2020.
- [16] David Haussler, Nick Littlestone, and Manfred K Warmuth. Predicting $\{0, 1\}$ -functions on randomly drawn points. *Information and Computation*, 115(2):248–292, 1994.
- [17] Nick Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine learning*, 2(4):285–318, 1988.
- [18] Gusztáv Morvai, Sidney Yakowitz, and László Györfi. Nonparametric inference for ergodic, stationary time series. *The Annals of Statistics*, 24(1):370–379, 1996.
- [19] Gusztáv Morvai, Sanjeev R Kulkarni, and Andrew B Nobel. Regression estimation from an individual stable sequence. *Statistics: A Journal of Theoretical and Applied Statistics*, 33(2): 99–118, 1999.
- [20] Kalyanapuram Rangachari Parthasarathy. *Probability measures on metric spaces*, volume 352. American Mathematical Soc., 2005.
- [21] Daniil Ryabko and Peter Bartlett. Pattern recognition for conditionally independent data. *Journal of Machine Learning Research*, 7(4), 2006.
- [22] Ingo Steinwart, Don Hush, and Clint Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1):175–194, 2009.
- [23] Jacques Stern. Le problème de la mesure. *Séminaire Bourbaki*, 1983:84, 1985.
- [24] Charles J Stone. Consistent nonparametric regression. *The annals of statistics*, pages 595–620, 1977.
- [25] Dan Tsir Cohen and Aryeh Kontorovich. Metric-valued regression. 2022.
- [26] Ruth Uner and Shai Ben-David. Probabilistic lipschitzness a niceness assumption for deterministic labels. In *Learning Faster from Easy Data-Workshop@ NIPS*, volume 2, page 1, 2013.
- [27] Veeravalli S Varadarajan. On the convergence of sample probability distributions. *Sankhyā: The Indian Journal of Statistics (1933-1960)*, 19(1/2):23–26, 1958.
- [28] Giuseppe Vitali. *Sul problema della misura dei Gruppi di punti di una retta: Nota*. Tip. Gamberini e Parmeggiani, 1905.

Appendix A. Proof of Proposition 4

Since \mathcal{X} infinite, let $\{x_i\}_{i \geq 0}$ a sequence of distinct points of \mathcal{X} . Let $\{\hat{t}_n\}_{n \geq 0}$ be a hypothesis test for condition FS. We suppose by contradiction that $\{\hat{t}_n\}$ is consistent and aim to construct a sequence \mathbb{X} on which it fails. Following a proof construction introduced in [13] (Theorem 47), we construct a (deterministic) process which fools the test by alternatively switching between two modes: constant

$X_t = x_0$ or visiting points of the distinct sequence $X_t = x_t$. Let $n_0 = 0$ and $X_0 = x_0$. We construct the sequences \mathbb{X} and $(n_i)_{i \geq 0}$ by induction. Suppose we have constructed n_t for $0 \leq t \leq k-1$ and X_t for $0 \leq t \leq n_{k-1}$.

- If k is even, consider the deterministic process \mathbb{Y} such that $Y_t = X_t$ for $t \leq n_{k-1}$ and $Y_t = x_0$ for $t > n_{k-1}$. Because $\{\hat{t}_n\}$ is consistent, we can define an index $n_k > n_{k-1}$ such that $\mathbb{P}(\hat{t}_{n_k}(\mathbb{Y}_{\leq n_k}) = 1) > \frac{3}{4}$.
- If k odd, consider the deterministic process \mathbb{Y} such that $Y_t = X_t$ for $t \leq n_{k-1}$ and $Y_t = x_t$ for $t > n_{k-1}$. Similarly, let $n_k > n_{k-1}$ such that $\mathbb{P}(\hat{t}_{n_k}(\mathbb{Y}_{\leq n_k}) = 0) > \frac{3}{4}$.

We then set $X_t = Y_t$ for $n_{k-1} < n_k$. Note that for all $k \geq 0$, $\mathbb{P}(\hat{t}_{n_{2k}}(\mathbb{X}_{\leq n_{2k}}) = 1) > \frac{3}{4}$ and $\mathbb{P}(\hat{t}_{n_{2k+1}}(\mathbb{X}_{\leq n_{2k+1}}) = 1) < \frac{1}{4}$. Then, $\hat{t}_n(\mathbb{X}_{\leq n})$ does not converge in probability and the hypothesis test $\{\hat{t}_n\}$ is not consistent. This ends the proof of the proposition.

Appendix B. Proof of Theorem 5

Let \mathbb{X} be a stochastic process that does not satisfy FMV and f_n be a learning rule, we aim to show that this learning rule cannot be universally consistent. By hypothesis, there exists a finite measurable partition $\{A_k\}_{k=1}^\infty$ in \mathcal{B} such that \mathbb{X} visits an infinite number of the A_k with probability $p > 0$. We denote by \mathcal{A} this event. We call \mathcal{F} the class of measurable functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ that takes constant values on each of the A_k . We will show that some objective function $f^* \in \mathcal{F}$ cannot be learnt by f_n .

First, let us define τ_k the first instant at which \mathbb{X} attains A_k .

$$\tau_k = \begin{cases} \min\{t \in \mathbb{N} : X_t \in A_k\} & \text{if } A_k \cap \mathbb{X} \neq \emptyset \\ 0 & \text{otherwise.} \end{cases}$$

We also define a deterministic quantity $T_k \in \mathbb{N}$ that upper bounds the τ_k with high probability, i.e.

$$\mathbb{P}(\tau_k \leq T_k) > 1 - 2^{-k}, \quad \forall k \geq 1.$$

By the Borel Cantelli lemma, since $\sum_k \mathbb{P}(\tau_k > T_k) < \infty$, almost surely there exists $\kappa \in \mathbb{N}$ such that $\tau_k \leq T_k$ for $k \geq \kappa$. We will denote \mathcal{E} this event. We now sample f^* randomly from \mathcal{F} as follows:

$$f^*(x \in A_k) = \begin{cases} y_{k,0} & \text{with proba } 1/2, \\ y_{k,1} & \text{with proba } 1/2, \end{cases}$$

where $y_{k,1}$ and $y_{k,0}$ are selected such that $\ell(y_{k,1}, y_{k,0}) \geq 2c_\ell T_k$ (recalling that c_ℓ denotes the constant from the relaxed triangle inequality satisfied by ℓ). Note that taking the expectation over the randomness in f^* allows to write:

$$\sup_{g \in \mathcal{F}} \mathbb{E}_{\mathbb{X}}(\mathcal{L}_{\mathbb{X}}(f, g)) \geq \mathbb{E}_{f^*, \mathbb{X}}(\mathcal{L}_{\mathbb{X}}(f, f^*)). \quad (1)$$

We first prove a lower bound on the right term. Conditionally on $\mathcal{A} \cap \mathcal{E}$, observe that for any $k \geq \kappa$, $(\mathbb{X}_{< \tau_k}, \mathbb{Y}_{< \tau_k})$ provides no information on $f^*(X_{\tau_k})$. Then, the average of the corresponding

prediction error satisfies $\mathbb{E}(\ell(f_{\tau_k}(\mathbb{X}_{<\tau_k}, \mathbb{Y}_{<\tau_k}, X_{\tau_k}), f^*(X_{\tau_k})) \geq T_k \geq \tau_k$, where we used the fact that ℓ satisfies the relaxed triangle inequality. Thus, in $\mathcal{A} \cap \mathcal{E}$, $\mathbb{E}_{f^*}(\mathcal{L}_{\mathbb{X}}(f, f^*, \tau_k)) \geq 1$ for any $k \geq \kappa$, hence by Fatou's lemma

$$\mathbb{E}_{f^*}(\mathcal{L}_{\mathbb{X}}(f, f^*)) \geq \limsup_{t \in \mathbb{N}} \mathbb{E}_{f^*}(\mathcal{L}_{\mathbb{X}}(f, f^*, \tau_k)) \geq 1.$$

Therefore, $\mathbb{P}_{\mathbb{X}}(\mathbb{E}_{f^*}(\mathcal{L}_{\mathbb{X}}(f, f^*)) \geq 1) \geq \mathbb{P}(\mathcal{A} \cap \mathcal{E}) = p$, which yields $\mathbb{E}_{f^*, \mathbb{X}}(\mathcal{L}_{\mathbb{X}}(f, f^*)) \geq p$. Equation 1 then shows that there exists $g \in \mathcal{F}$ such that $\mathbb{E}_{\mathbb{X}}(\mathcal{L}_{\mathbb{X}}(f, g)) > 0$, hence $\mathcal{L}_{\mathbb{X}}(f, g) > 0$ with nonzero probability. This ends the proof of the result.

Appendix C. Proof of Theorem 12

Because \mathcal{Y} is separable, the sequence of empirical measures converges weakly to the true measure of Y almost surely [27]. Hence, by Lemma 2.1 of [7], almost surely we have for all $y \in \mathcal{Y}$, $\frac{1}{n} \sum_{i=1}^n d^p(y, Y_i) \rightarrow \mathbb{E}d^p(y, Y)$. In the rest of this proof, we suppose that this event is met. We now denote by $R^* := \min_{y \in \mathcal{Y}} \mathbb{E}d^p(y, Y)$ the minimal risk and $y_0 \in \mathcal{Y}$ such that $\mathbb{E}d^p(y_0, Y) < \infty$. We first note that the sequence $(\hat{y}_n)_n$ is bounded almost surely since $d^p(\hat{y}_n, y_0) \leq \frac{2^p}{n} \sum_{i=1}^n d^p(y_0, Y_i) \rightarrow 2^p \mathbb{E}d^p(y_0, Y)$. Therefore, almost surely the sequence lies in a compact. We will suppose that this condition is also met in the rest of the proof. Now suppose by contradiction that the convergence does not hold. Let $\epsilon > 0$ and consider a subsequence ϕ such that $|\frac{1}{\phi(k)} \sum_{i=1}^{\phi(k)} d^p(\hat{y}_{\phi(k)}, Y_i) - R^*| \geq \epsilon$. Because the sequence \hat{y}_n is confined to a compact closed ball, there exists a subsequence ψ and $y^* \in \mathcal{Y}$ such that $\hat{y}_{\phi(\psi(l))} \rightarrow y^*$. For simplicity, we omit the subsequences and simply write $\hat{y}_l \rightarrow y^*$. Fix $\epsilon > 0$. Using the constant c_ϵ dependent on p from Lemma 2.3 of [7] such that for any $a, b \geq 0$, $(a + b)^p \leq (1 + \epsilon)a^p + c_\epsilon b^p$, we can write for any $y \in \mathcal{Y}$,

$$\begin{aligned} \mathbb{E}d^p(y^*, Y) &= \lim \frac{1}{l} \sum_{i=1}^l d^p(y^*, Y_i) \leq \liminf c_\epsilon d^p(y^*, \hat{y}_l) + \frac{1 + \epsilon}{l} \sum_{i=1}^l d^p(\hat{y}_l, Y_i) \\ &\leq \liminf c_\epsilon d^p(y^*, \hat{y}_l) + \frac{1 + \epsilon}{l} \sum_{i=1}^l d^p(y, Y_i) \\ &= (1 + \epsilon) \mathbb{E}d^p(y, Y). \end{aligned}$$

Since this holds for any $\epsilon > 0$ and $y \in \mathcal{Y}$, this shows that $\mathbb{E}d^p(y^*, Y) = R^*$. We now observe that

$$\begin{aligned} \limsup \frac{1}{l} \sum_{i=1}^l d^p(\hat{y}_l, Y_i) &\leq \limsup c_\epsilon d^p(\hat{y}_l, y^*) + \frac{1 + \epsilon}{l} \sum_{i=1}^l d^p(y^*, Y_i) = (1 + \epsilon) \mathbb{E}d^p(y^*, Y), \\ \liminf \frac{1}{l} \sum_{i=1}^l d^p(\hat{y}_l, Y_i) &\geq \limsup \frac{1}{(1 + \epsilon)l} \sum_{i=1}^l d^p(y^*, Y_i) - \frac{c_\epsilon}{1 + \epsilon} d(\hat{y}_l, y^*) = \frac{\mathbb{E}d^p(y^*, Y)}{1 + \epsilon}. \end{aligned}$$

Therefore, $\frac{1}{l} \sum_{i=1}^l d^p(\hat{y}_l, Y_i) \rightarrow R^*$ which contradicts the hypothesis and ends the proof of the theorem.

Appendix D. Proof of Theorem 11

We fix a process \mathbb{X} that admits a strong universal Bayes consistent learning rule and define $S = \{x \in \mathcal{X}, \{x\} \cap \mathbb{X} \neq \emptyset\}$ the random support of \mathbb{X} . This set is almost surely finite because $\mathbb{X} \in \text{SUOL}$. We fix a measurable function \bar{f} and we write,

$$\begin{aligned} \mathcal{L}_{\mathbb{X}}(f, p_{Y|X}; T, \bar{f}) &= \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \sum_{x \in S, X_t=x} (\ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) - \ell(\bar{f}(X_t), Y_t)), \\ &\leq \sum_{x \in S} \limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{X_t=x} (\ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) - \ell(\bar{f}(X_t), Y_t)). \end{aligned}$$

Now fix $x \in \mathcal{X}$, And denote $\mathcal{L}_x(T) := \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{X_t=x} (\ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) - \ell(\bar{f}(X_t), Y_t))$. If $\sum_{t \geq 1} \mathbb{1}_{X_t=x} < \infty$, then we have directly $\limsup_T \mathcal{L}_x(T) = 0$. We now turn to the case $\sum_{t \geq 1} \mathbb{1}_{X_t=x} = \infty$. Note that the sequence $(Y_u)_{u, X_u=x}$ is an i.i.d. sequence of variables following the distribution $p_{Y|X=x}$. First suppose that there exists $y_0 \in \mathcal{Y}$ such that $\mathbb{E}_{Y|X=x} \ell(y_0, Y) < \infty$. Then, by Theorem 12,

$$\frac{1}{\sum_{t=1}^T \mathbb{1}_{X_t=x}} \sum_{t=1}^T \mathbb{1}_{X_t=x} \ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) \rightarrow \min_{y \in \mathcal{Y}} \mathbb{E}_{Y|X=x} \ell(y, Y).$$

Otherwise, we have observe directly

$$\frac{1}{\sum_{t=1}^T \mathbb{1}_{X_t=x}} \sum_{t=1}^T \mathbb{1}_{X_t=x} \ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) \leq \infty = \min_{y \in \mathcal{Y}} \mathbb{E}_{Y|X=x} \ell(y, Y).$$

Further note that

$$\frac{1}{\sum_{t=1}^T \mathbb{1}_{X_t=x}} \sum_{t=1}^T \mathbb{1}_{X_t=x} \ell(\bar{f}(X_t), Y_t) \rightarrow \mathbb{E}_{Y|X=x} \ell(\bar{f}(x), Y) \geq \min_{y \in \mathcal{Y}} \mathbb{E}_{Y|X=x} \ell(y, Y).$$

Therefore, we obtain in all cases

$$\limsup_{T \rightarrow \infty} \frac{1}{\sum_{t=1}^T \mathbb{1}_{X_t=x}} \sum_{t=1}^T \mathbb{1}_{X_t=x} (\ell(f_t(\mathbb{X}_{<t}, \mathbb{Y}_{<t}, X_t), Y_t) - \ell(\bar{f}(X_t), Y_t)) \leq 0.$$

Because we have $\sum_{t=1}^T \mathbb{1}_{X_t=x} \leq T$, this yields $\limsup_T \mathcal{L}_x(T) \leq 0$. Now recall that because $\mathbb{X} \in \text{SUOL}$, the set S is finite almost surely, hence $\mathcal{L}_{\mathbb{X}}(f, p_{Y|X}; T, \bar{f}) \leq 0$ (a.s.).