

Open Problem: Is There an Online Learning Algorithm That Learns Whenever Online Learning Is Possible?

Steve Hanneke

STEVE.HANNEKE@GMAIL.COM

Toyota Technological Institute at Chicago

Editors: Mikhail Belkin and Samory Kpotufe

1. Background

One of the most classical topics in learning theory is *online learning*, wherein a learning algorithm observes a stream of data points X_t from a space \mathcal{X} , and for each time $t \in \mathbb{N}$ it makes a prediction $\hat{Y}_t \in \{0, 1\}$, after which it observes the *target* label $Y_t \in \{0, 1\}$. The prediction \hat{Y}_t may depend on the past observations $((X_s, Y_s), s < t)$ and X_t , but nothing else: that is, $\hat{Y}_t = f_t(X_{1:(t-1)}, Y_{1:(t-1)}, X_t)$ for a (possibly randomized) function f_t . The objective is to make few *mistakes*: that is, few times t for which $\hat{Y}_t \neq Y_t$. In particular, we are most interested in achieving a number of mistakes among $\{(X_t, Y_t)\}_{t \leq T}$ growing *sublinearly* in T : that is, $o(T)$ mistakes. In various versions of this problem, the sequences $\mathbb{X} := \{X_t\}_{t \in \mathbb{N}}$ and $\mathbb{Y} := \{Y_t\}_{t \in \mathbb{N}}$ can be either deterministic or random, and either oblivious or adaptive. Here we focus on the special case where $Y_t = f^*(X_t)$ for a fixed unknown *target concept* $f^* : \mathcal{X} \rightarrow \{0, 1\}$, and where the sequence \mathbb{X} may be random and is independent of any internal randomness in the learning algorithm f_t .

Since it is certainly not possible to guarantee few mistakes for *all* pairs (\mathbb{X}, f^*) , some restrictions are necessary, and in this respect theories of online learning may be grouped into three categories: (1) those which allow *arbitrary* sequences \mathbb{X} but restrict the allowed target concepts f^* (e.g., [Littlestone, 1988](#)), (2) those which restrict both \mathbb{X} and f^* (e.g., [Haussler, Littlestone, and Warmuth, 1994](#)), and (3) those which restrict the sequence \mathbb{X} of points but allow *arbitrary* target concepts f^* (e.g., [Stone, 1977](#)). In particular, a classic result in category (1) is that there exists a learning algorithm guaranteeing a *bounded* number of mistakes for every sequence \mathbb{X} *if and only if* the set of allowed target concepts f^* has a finite *Littlestone dimension* ([Littlestone, 1988](#)) (see [Ben-David, Pál, and Shalev-Shwartz, 2009](#), for the definition of Littlestone dimension). The works falling in category (2) vary widely in the types of restrictions they impose and the resulting guarantees that are possible. In particular, note that by introducing restrictions on the (possibly random) sequence \mathbb{X} , we can, to some extent, express the classic theory of *statistical learning*. For instance, there exist learning algorithms guaranteeing $O(\log(T))$ mistakes (in expectation) for every \mathbb{X} that is an *i.i.d.* process on \mathcal{X} when the set of allowed target concepts f^* has finite *VC dimension* ([Haussler, Littlestone, and Warmuth, 1994](#)). Other more-involved restrictions on the pair (\mathbb{X}, f^*) have also been considered (e.g., [Ryabko, 2006](#); [Uerner and Ben-David, 2013](#); [Bousquet, Hanneke, Moran, van Handel, and Yehudayoff, 2021](#)).

The subject of our present discussion is category (3): that is, *unrestricted* target concepts f^* , but with restrictions on the sequence \mathbb{X} . There has also been significant work in this category. As a simple example, there exist learning algorithms guaranteeing $o(T)$ mistakes (almost surely) for every target concept f^* and every \mathbb{X} that is an *i.i.d.* process on \mathcal{X} ([Stone, 1977](#); [Devroye, Györfi, and Lugosi, 1996](#); [Hanneke, 2021](#); [Hanneke, Kontorovich, Sabato, and Weiss, 2021](#)) (see Remark 4 for relevant technical conditions on \mathcal{X}). Indeed, for $\mathcal{X} = \mathbb{R}^d$, this even holds (in expectation) for the simple 1-nearest neighbor algorithm ([Cover and Hart, 1967](#); [Stone, 1977](#); [Devroye, Györfi, and](#)

Lugosi, 1996). There is also an extensive literature considering relaxations of the i.i.d. assumption, such as allowing \mathbb{X} to be *stationary ergodic* (Morvai, Yakowitz, and Györfi, 1996; Györfi, Lugosi, and Morvai, 1999; Györfi and Lugosi, 2002), or generally to satisfy a *law of large numbers* (Morvai, Kulkarni, and Nobel, 1999; Steinwart, Hush, and Scovel, 2009), while maintaining this guarantee of $o(T)$ mistakes (either in-expectation or almost surely).

The recent work of (Hanneke, 2021) unifies and weakens these various restrictions on \mathbb{X} , aiming (in part) to study the *fundamental limits* of this category. The following definition provides a formal criterion for online learning in this setting.

Definition 1 For a (possibly random) sequence $\mathbb{X} = \{X_t\}_{t \in \mathbb{N}}$, an online learning algorithm $f_t : \mathcal{X}^{t-1} \times \{0, 1\}^{t-1} \times \mathcal{X} \rightarrow \{0, 1\}$ (possibly randomized, independent of \mathbb{X}) is *weakly universally consistent under \mathbb{X}* if, for every (measurable) $f^* : \mathcal{X} \rightarrow \{0, 1\}$,

$$\mathbb{E} \left[\sum_{t=1}^T \mathbb{1}[f_t(X_{1:(t-1)}, f^*(X_{1:(t-1)}), X_t) \neq f^*(X_t)] \right] = o(T),$$

and is *strongly universally consistent under \mathbb{X}* if, for every (measurable) $f^* : \mathcal{X} \rightarrow \{0, 1\}$,

$$\sum_{t=1}^T \mathbb{1}[f_t(X_{1:(t-1)}, f^*(X_{1:(t-1)}), X_t) \neq f^*(X_t)] = o(T) \text{ (a.s.)}$$

Definition 2 For a (possibly random) sequence \mathbb{X} , we say (weak/strong) *universal online learning is possible under \mathbb{X}* if there exists an online learning algorithm that is (weakly/strongly) *universally consistent under \mathbb{X}* .

It is clear that not *every* \mathbb{X} admits universal online learning: for instance, for $\mathcal{X} = \mathbb{N}$, universal online learning is not possible under $\mathbb{X} = \{1, 2, 3, \dots\}$. Hence, to approach the goal of achieving the fundamental limits of online learning in this setting, (Hanneke, 2021) introduces a style of reasoning referred to as the **optimist’s decision theory**, described abstractly as follows.

The Optimist’s Decision Theory: Supposing we are tasked with achieving a given objective O in some scenario, then already we have implicitly committed to the assumption that achieving objective O is at least *possible* in that scenario: the *optimist’s assumption*. Since we must commit to this assumption to even begin designing a strategy for achieving objective O , we may *rely* on this assumption in our strategy for achieving the objective. We are then most interested in strategies guaranteed to achieve objective O *without any additional assumptions*. Such a strategy is *universal* in the most-general sense possible, since the optimist’s assumption is *necessary*. It will achieve the objective O in *all* scenarios where it is possible to do so. Moreover, such strategies have the satisfying property that, if ever they fail to achieve the objective, we may rest assured that no other strategy could have succeeded, so that nothing was lost.

Based on the above reasoning, an online learning algorithm that is universally consistent for every \mathbb{X} that admits universal online learning is called *optimistically universal*:

Definition 3 An online learning algorithm is *optimistically (weakly/strongly) universal* if it is (weakly/strongly) *universally consistent under every \mathbb{X} such that (weak/strong) universal online learning is possible under \mathbb{X}* .

In this context, we present two fundamental questions about online learning, originally posed in (Hanneke, 2021). The first asks whether there *exists* an optimistically universal online learning algorithm, while the second asks whether there is a basic property of \mathbb{X} that determines whether universal online learning is possible, proposing a particular candidate condition for concreteness. We now turn to these two questions in detail.

2. Open Problem 1: Optimistically Universal Online Learning

The first open problem concerns the existence of optimistically universal online learning algorithms.

Open Problem 1 *Does there exist an optimistically universal online learning algorithm? (in either the weak or strong sense)*

Prize: I am offering \$5,000 USD for a solution to this problem (be it positive or negative). If the weak/strong variants are solved in separate works, whichever is solved first will receive the prize.

Remark 4 (Remark on generality) *The general setting considered in (Hanneke, 2021) allows that \mathcal{X} is any nonempty space equipped with a separable metrizable topology \mathcal{T} , and the measurable sets are specified by the Borel σ -algebra generated by \mathcal{T} . However, I am willing to award the prizes for any solution general enough to address the case $\mathcal{X} = \mathbb{R}^d$ with the Euclidean topology.*

Notes: The answer to Open Problem 1 is known to be positive in the special case of *countable* \mathcal{X} , or for general \mathcal{X} but with the restriction to *deterministic* sequences \mathbb{X} (Hanneke, 2021). Indeed, as discussed below, in both cases a simple *memorization*-based algorithm suffices. However, for uncountable \mathcal{X} and general (random) sequences \mathbb{X} , there are simple cases where memorization fails: for instance, \mathbb{X} as any non-atomic i.i.d. process. Thus, the case that remains open concerns *uncountable* \mathcal{X} and general (random) sequences \mathbb{X} . Also note that it is conceivable that an optimistically strongly universal online learning algorithm is not necessarily also optimistically weakly universal, since the latter requires universal consistency under a strictly larger family of processes (see below). Nevertheless, I conjecture that the answers to the weak/strong variants will be the same.

3. Open Problem 2: When Is Universal Online Learning Possible?

The second open problem concerns characterizing the family of random sequences \mathbb{X} under which universal online learning is possible. In addition to being intrinsically interesting, this would likely also be an extremely helpful step toward resolving Open Problem 1. To make the problem concrete, (Hanneke, 2021) proposes the following two conditions.

Definition 5

- Let \mathcal{C}_w denote the family of all (possibly random) sequences $\mathbb{X} = \{X_t\}_{t \in \mathbb{N}}$ such that every disjoint sequence $\{A_i\}_{i \in \mathbb{N}}$ of measurable sets satisfies $\mathbb{E}[|\{i \in \mathbb{N} : X_{1:T} \cap A_i \neq \emptyset\}|] = o(T)$.
- Let \mathcal{C}_s denote the family of all (possibly random) sequences $\mathbb{X} = \{X_t\}_{t \in \mathbb{N}}$ such that every disjoint sequence $\{A_i\}_{i \in \mathbb{N}}$ of measurable sets satisfies $|\{i \in \mathbb{N} : X_{1:T} \cap A_i \neq \emptyset\}| = o(T)$ (a.s.).

We then have the following open problem.

Open Problem 2

- Is \mathcal{C}_w equal to the set of all \mathbb{X} such that weak universal online learning is possible under \mathbb{X} ?
- Is \mathcal{C}_s equal to the set of all \mathbb{X} such that strong universal online learning is possible under \mathbb{X} ?

Prize: I am offering \$1,000 USD for a solution to either of these questions (be it positive or negative). If the weak/strong variants are solved in separate works, whichever is solved first will receive the prize. Additionally, I note that Remark 4 also applies to this problem.

Notes: The work (Hanneke, 2021) establishes that $\mathbb{X} \in \mathcal{C}_w$ or $\mathbb{X} \in \mathcal{C}_s$ are *necessary* for weak or strong, respectively, universal online learning to be possible under \mathbb{X} . Moreover, in the special case of countable \mathcal{X} , or for general \mathcal{X} but with the restriction to deterministic sequences \mathbb{X} , (Hanneke, 2021) also shows that $\mathbb{X} \in \mathcal{C}_w$ or $\mathbb{X} \in \mathcal{C}_s$ are *sufficient* for weak or strong, respectively, universal online learning to be possible under \mathbb{X} . That is, both questions in Open Problem 2 have *positive* answers for countable \mathcal{X} , or for general \mathcal{X} with the restriction to deterministic \mathbb{X} . Indeed, it is an easy exercise to verify that the simple *memorization* algorithm is weakly or strongly universally consistent in these cases, when $\mathbb{X} \in \mathcal{C}_w$ or $\mathbb{X} \in \mathcal{C}_s$, respectively. Thus, the case that remains open in Open Problem 2 concerns whether $\mathbb{X} \in \mathcal{C}_w$ or $\mathbb{X} \in \mathcal{C}_s$ are sufficient conditions for weak or strong, respectively, universal online learning to be possible under \mathbb{X} for *uncountable* \mathcal{X} and general (random) sequences \mathbb{X} . The route to proving such a result (positively) would be to construct an online learning algorithm and show that it is weakly or strongly universally consistent under every \mathbb{X} in \mathcal{C}_w or \mathcal{C}_s , respectively. Note that, unlike Open Problem 1, an algorithm sufficient to positively resolve Open Problem 2 may even depend on the distribution of \mathbb{X} .

(Hanneke, 2021) also discusses relations between the sets involved in Open Problem 2. It is clear that $\mathcal{C}_s \subseteq \mathcal{C}_w$. Also, for any \mathbb{X} , if strong universal online learning is possible then weak universal online learning must also be possible. However, supposing \mathcal{X} is infinite, (Hanneke, 2021) gives an example \mathbb{X} in $\mathcal{C}_w \setminus \mathcal{C}_s$, so that the two sets are not equivalent. Moreover, (Hanneke, 2021) shows that weak universal online learning *is* possible under this \mathbb{X} , but *strong* universal online learning *is not* possible under this \mathbb{X} . Thus, the sets of (random) sequences \mathbb{X} under which universal online learning is possible in the weak and strong senses are *not equal*.

4. Connections to Related Settings

The work of (Hanneke, 2021) considers three learning settings: inductive, self-adaptive, and online. The inductive setting is most-familiar to the statistical learning literature, where a learning algorithm observes a finite *training* set and then produces a fixed hypothesis that is then used for all future predictions. The self-adaptive setting differs only in that it allows the learner to update its hypothesis based on the *unlabeled* data it has made predictions on so far. The interested reader is referred to that work for the precise definitions. (Hanneke, 2021) proves that there *do* exist optimistically universal *self-adaptive* learning algorithms, meaning that they are universally consistent for all \mathbb{X} such that universal *self-adaptive* learning is possible under \mathbb{X} (both weak and strong). On the other hand, (Hanneke, 2021) also proves that optimistically universal *inductive* learning is *impossible* (both weak and strong). Moreover, (Hanneke, 2021) provides a concise characterization of the family of all (possibly random) sequences \mathbb{X} such that (weak/strong) universal (inductive/self-adaptive) learning is possible under \mathbb{X} .

That work also makes connections between self-adaptive learning and online learning, providing a technique to convert any self-adaptive learning algorithm into an online learning algorithm, while preserving consistency. In particular, applying this conversion to the optimistically universal self-adaptive learning algorithm provides an online learning algorithm that is strongly universally consistent under *every* \mathbb{X} under which universal self-adaptive learning is possible. However, (Hanneke, 2021) also shows that for any infinite \mathcal{X} , there exist sequences \mathbb{X} (even deterministic) for which universal online learning is possible but universal self-adaptive learning is not possible. Thus, new techniques are needed to understand the sufficient conditions for universal online learning (Open Problem 2), and to approach the question of optimistically universal online learning (Open Problem 1).

References

- S. Ben-David, D. Pál, and S. Shalev-Shwartz. Agnostic online learning. In *Proceedings of the 22nd Conference on Learning Theory*, 2009.
- O. Bousquet, S. Hanneke, S. Moran, R. van Handel, and A. Yehudayoff. A theory of universal learning. In *Proceedings of the 53rd Annual ACM Symposium on Theory of Computing*, 2021.
- T. Cover and P. Hart. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, 13:21–27, 1967.
- L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer-Verlag New York, 1996.
- L. Györfi and G. Lugosi. Strategies for sequential prediction of stationary time series. In M. Dror, P. L’Ecuyer, and F. Szidarovszky, editors, *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, pages 225–248. Kluwer Academic Publishers, 2002.
- L. Györfi, G. Lugosi, and G. Morvai. A simple randomized algorithm for sequential prediction of ergodic time series. *IEEE Transactions on Information Theory*, 45(7):2642–2650, 1999.
- S. Hanneke. Learning whenever learning is possible: Universal learning under general stochastic processes. *Journal of Machine Learning Research*, 22:1–116, 2021.
- S. Hanneke, A. Kontorovich, S. Sabato, and R. Weiss. Universal Bayes consistency in metric spaces. *The Annals of Statistics*, To appear, 2021.
- D. Haussler, N. Littlestone, and M. Warmuth. Predicting $\{0, 1\}$ -functions on randomly drawn points. *Information and Computation*, 115(2):248–292, 1994.
- N. Littlestone. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning*, 2(4):285–318, 1988.
- G. Morvai, S. Yakowitz, and L. Györfi. Nonparametric inference for ergodic, stationary time series. *The Annals of Statistics*, 24(1):370–379, 1996.
- G. Morvai, S. R. Kulkarni, and A. B. Nobel. Regression estimation from an individual stable sequence. *Statistics*, 33:99–118, 1999.
- D. Ryabko. Pattern recognition for conditionally independent data. *Journal of Machine Learning Research*, 7(4):645–664, 2006.
- I. Steinwart, D. Hush, and C. Scovel. Learning from dependent observations. *Journal of Multivariate Analysis*, 100(1):175–194, 2009.
- C. J. Stone. Consistent nonparametric regression. *The Annals of Statistics*, 5(4):595–620, 1977.
- R. Urner and S. Ben-David. Probabilistic Lipschitzness A niceness assumption for deterministic labels. In *Learning Faster from Easy Data-Workshop @ NIPS*, 2013.