

PROSPECTIVE LEARNING: PRINCIPLED EXTRAPOLATION TO THE FUTURE

The Future Learning Collective

ABSTRACT

Learning is a process which can update decision rules, based on past experience, such that future performance improves. Traditionally, machine learning is often evaluated under the assumption that the future will be identical to the past in distribution or change adversarially. But these assumptions can be either too optimistic or pessimistic for many problems in the real world. Real world scenarios evolve over multiple spatiotemporal scales with partially predictable dynamics. Here we reformulate the learning problem to one that centers around this idea of dynamic futures that are partially learnable. We conjecture that certain sequences of tasks are not retrospectively learnable (in which the data distribution is fixed), but are prospectively learnable (in which distributions may be dynamic), suggesting that prospective learning is more difficult in kind than retrospective learning. We argue that prospective learning more accurately characterizes many real world problems that (1) currently stymie existing artificial intelligence solutions and/or (2) lack adequate explanations for how natural intelligences solve them. Thus, studying prospective learning will lead to deeper insights and solutions to currently vexing challenges in both natural and artificial intelligences.

1 INTRODUCTION

All learning is for the future. Learning involves updating decision rules or policies, based on past experiences, to improve *future* performance. The dogma of much of machine learning is that future scenarios faced by the learner are similar to past ones, e.g., all data, including what the learner learns from and what the learner gets tested on, are independent and sampled from the same probability distribution (Glivenko, 1933; Cantelli, 1933; Vapnik & Chervonenkis, 1971; Vapnik, 1991; Valiant, 1984). This is not specific to supervised learning: algorithms for sub-fields ranging from robust learning (White, 1982; Huber & Ronchetti, 2009) (test distribution can be slightly different from the training distribution), online learning (Cesa-Bianchi & Lugosi, 2006; Orabona, 2019) (data distribution can arbitrarily drift over time) to adversarial learning (Madry et al., 2017) (test distribution can be adversarial) and reinforcement learning (Bellman, 1954; Recht, 2019) (data distribution changes as the learner procures more data), all use variants of this assumption. We refer to these algorithms as *retrospective learners* because they learn representations that are optimal for past data, not future data. Retrospective learning is certainly useful: it is the basis for, arguably all, real-world learning systems today. Modifications to the core retrospective assumption, e.g., covariate shift detection, domain adaptation, meta-learning and continual learning, help bridge the gap that arises from using representations that are optimal on past data to make inferences on future data. Each of these ideas makes a specific assumption about future data, e.g., covariate shift assumes that only the marginal on the inputs changes, meta-learning assumes that new tasks are similar to past tasks. Learning systems can fail catastrophically when they are faced with scenarios outside these assumptions (Giurfa et al., 2001; Vermaercke et al., 2014; Fleuret et al., 2011; Kim et al., 2018; Yao et al., 2022; Lazaridou et al., 2021).

This paper attempts to rethink the fundamental assumptions of learning to address the above gap. We propose the concept of **prospective learning**, a generalization of retrospective learning, that relaxes the assumption that past (training) and future (test) distributions are identical and stationary. The goal of prospective learning explicitly emphasizes the ability to predict the future, or *prospection*. In comparison, adversarial learning or online learning make no attempt at predicting the future. While the idea of prospection is by no means new (Seligman et al., 2013), its role has largely been neglected in machine learning. Prospective learning embodies the eventual goal of machine learning which is to make predictions for the foreseeable future.

2 A FAILURE MODE OF RETROSPECTIVE LEARNING

Consider an infinite sequence of distributions denoted by $P = \{P_t\}_{t \in \mathbb{N}}$ where time t is a natural number. We see data from the first t' distributions and we seek a hypothesis that generalizes to distributions that we will see in the future.

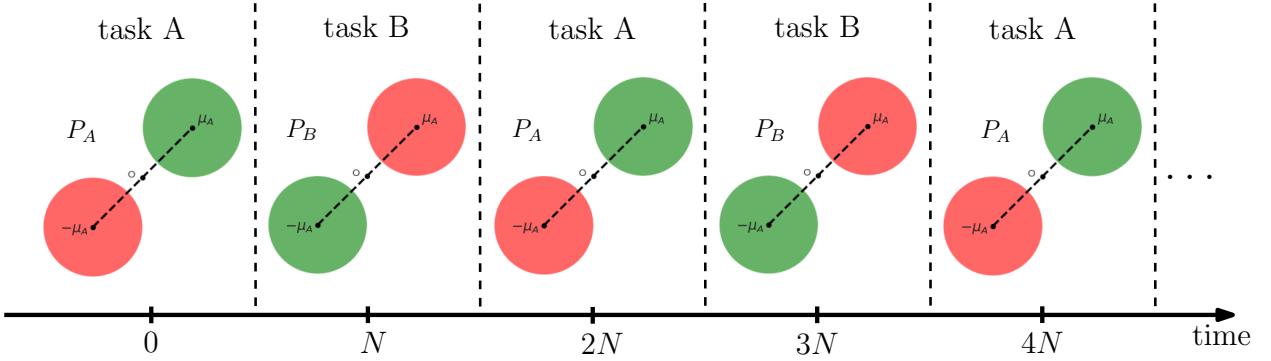


Figure 1: This figure depicts an infinite sequence of data distributions formed by alternating between two distributions, P_A and P_B . Both distributions define binary classification problems, namely task A and task B, with 2-dimensional inputs where the classes are separable using a single line. A single hypothesis, that doesn't change with time, cannot generalize arbitrarily well to this sequence of distributions. However, the sequence of distributions evolve according to a simple alternating pattern which can be exploited to generalize to tasks in the future. A prospective learner should ideally leverage this pattern by modeling the dynamics under which the tasks evolve with time.

We illustrate the insufficiency of retrospective learning using a simple example. Consider two distributions P_A and P_B ; both P_A and P_B define binary classification tasks with 2-dimensional inputs. The classes in both distributions can be perfectly separated by a linear classifier. Consider the sequence of distributions P , formed by alternating between P_A and P_B at regular intervals. We would like to design a learner L that learns using a finite amount of data and generalizes to the infinite sequence of distributions depicted in Figure 1.

No single hypothesis can achieve an arbitrarily small error on both P_A and P_B simultaneously. As a result, if a retrospective learner picks a hypothesis that is unchanged over time, then it cannot achieve an arbitrarily small error on the sequence of distributions even if the learner has access to a large number of samples.

We can also use methods from continual or meta-learning to tackle a sequence of distributions. At every time step, the learner encounters an unknown distribution and uses a few samples from this distribution in order to build a hypothesis for it. Even if the learner has seen a large number of distributions, it still requires at least 1 sample at every time step. Such a learner lowers the sample complexity of learning the unknown task with the help of a learnt inductive bias (Baxter, 2000). However, it still fails to model one key aspect of the problem.

Many standard methods do not predict the future sequence of distributions; this is helpful if the sequence of distributions follows some pattern like in the example in Figure 1. For this example, if we know that the tasks alternate at every N time steps, then we can maintain two different hypotheses—one for P_A and another for P_B —and alternate between them. This helps us generalize better but more importantly allows us to do so using fewer samples. We do not need additional samples when the task switches to adapt to an unknown distribution, since we can perfectly predict the future sequence of distributions. In general, data distributions can change with time, and approximating the underlying dynamics of distribution shift can help us build a hypothesis that predicts the next distribution.

3 THE THEORETICAL FOUNDATIONS FOR PROSPECTIVE LEARNING

3.1 RETROSPECTIVE LEARNING

To formalize prospective learning, we contrast it with traditional (retrospective) learning (Mohri et al., 2018). Consider a distribution P over inputs x and outputs y , in which case we have a supervised learning problem. A learning task is generally associated with a distribution P , a hypothesis class \mathcal{H} , and a risk function $R(\cdot)$. For simplicity, in this work, we represent a learning task by its associated distribution P , assuming that \mathcal{H} and $R(\cdot)$ are fixed. Therefore, we use the terms “task” and “distribution” interchangeably. Using the dataset D , drawn independently and identically from $P \in \mathcal{P}$, we seek a hypothesis h that minimizes the risk $R(\cdot)$. The hypothesis h is selected from the hypothesis class \mathcal{H} . The risk of a hypothesis h can be defined as $R(h) = \mathbb{E}_{(x,y) \sim P} [\ell((x,y), h)]$ where ℓ is the loss function. For a classification problem, we often consider the 0-1 loss, $\ell((x,y), h) = 1\{h(x) \neq y\}$, and for a regression problem, we often consider the mean-squared error $\ell((x,y), h) = (h(x) - y)^2$.

Classical machine learning posits that a learning task is unchanged with respect to time t , i.e., $P_t \equiv P; \forall t$ (Glivenko, 1933; Cantelli, 1933). One may formalize learning for a class of problems using the probably approximately correct (PAC) model (Valiant, 1984; Vapnik & Chervonenkis, 1971). Under this model, we seek a learning algorithm that achieves an arbitrarily small risk on any given task P using samples from that task.

Definition 1 (PAC-learnability) (From Valiant (1984); Kearns & Vazirani (1994)) *A family of tasks \mathcal{P} is PAC-learnable using hypothesis class \mathcal{H} if there exists a learner L with the following property: For all $\epsilon, \delta > 0$ and tasks $P \in \mathcal{P}$, if L has access to dataset D of at least n independent and identically distributed samples from P , then it outputs a hypothesis $\hat{h} \in \mathcal{H}$ which satisfies*

$$\mathbb{P}_D[|R(\hat{h}) - R(h^*)| < \epsilon] \geq 1 - \delta, \quad (1)$$

where \mathbb{P}_D is the probability over draws of the dataset D and h^* is the optimal hypothesis from the hypothesis class \mathcal{H} . The learner L is said to be efficiently PAC-learnable if the sample complexity n is a polynomial function of ϵ^{-1} and δ^{-1} .

While Inequality 1 is relatively simple, it can be used to frame the foundation of the most recent AI revolution, spanning computer vision (Krizhevsky et al., 2012), natural language processing (Devlin et al., 2019), diagnostics (McKinney et al., 2020) and protein folding (Senior et al., 2020). Yet, its limitations have been widely acknowledged for decades in ML and statistics (see Bozinovski & Fulgosi (1976); Hand (2006)). Specifically, real-world data rarely exhibit such simple dynamics. Specifically, the assumption that training (past) data and test/evaluation (future) data are sampled from the exact same fixed distribution is woefully inadequate. Empirically, modern deep learning approaches continue to fail in scenarios that are trivially easy for humans (Madry et al., 2017; Alcorn et al., 2019), and which we believe would benefit from prospective learning capabilities. Moreover, existing retrospective learning theory is inadequate to characterize most of the above generalized learning scenarios (Thrun, 1996) (though see Arjovsky (2021); Baxter (2000) for some compelling results).

3.2 PROSPECTIVE LEARNING

Sub-fields like online (Cesa-Bianchi & Lugosi, 2006), transfer (Bozinovski & Fulgosi, 1976), multitask (Bengio et al., 1992; Caruana, 1997) meta- (Baxter, 2000; Wang, 2020), continual (Ring, 1994), and lifelong learning (Thrun & Mitchell, 1995) present different ways to generalize the retrospective learner defined in Inequality 1. Most generalizations focus on *adapting* to new tasks; given data from a set of tasks, the goal is to adapt to a new task using a small number of samples from it. On the other hand, a prospective learner attempts to model how the tasks evolve with time; the emphasis is on *predicting* what new tasks will look like prospectively, as opposed to tackling them retrospectively.

We propose a working formal definition of prospective learning which generalizes the classical retrospective learning scenario described by Inequality 1. Instead of learning a task where the distribution, risk, and hypothesis are fixed across time, we assume that they can evolve over time. For simplicity, hereafter we keep the risk and hypothesis class fixed, and assume that only the data distribution changes over time. To formalize this idea, we consider a fixed sequence of distributions $P := \{P_t\}_{t \geq 0}$. Here, P_t is the distribution at time t . Time t can take values from $[0, \infty)$ or be limited to the set of natural numbers, among other choices.

Let $D_{t'}$ denote a dataset drawn from P before time t' . Using a prospective learner L that takes $D_{t'}$ as input, we wish to estimate a time-indexed sequence of decision functions or hypotheses that generalizes to future tasks encountered after time t' . We denote this estimated hypothesis sequence by $\hat{h} = \{\hat{h}_t(\cdot)\}_{t \geq 0}$ where \hat{h}_t is the hypothesis at time t . We provide a visual illustration of the objective of prospective learning in Figure 2.

Finally, let \mathcal{H} denote a hypothesis class that contains sequences of hypotheses. For a hypothesis sequence $h \in \mathcal{H}$, we can define the risk at time t , assuming we have supervised learning problems, as $R_t(h_t) = \mathbb{E}_{(x,y) \sim P_t} [\ell((x,y), h_t)]$ where ℓ is some choice of a loss function. We can now provide a formal definition of prospective-learnability of the sequence P using the hypothesis class \mathcal{H} as follows.

Definition 2 (Prospective-Learnability) Consider a sequence of distributions $P := \{P_t\}_{t \geq 0}$ and a hypothesis class \mathcal{H} . We say that \mathcal{H} is prospectively-learnable with respect to P and a chosen reference hypothesis sequence h , with precision $\epsilon > 0$ and confidence $\delta > 0$ if there exists a learner L and $\bar{t} \geq 0$ that satisfies the following property: for any $t' > \bar{t}$, if L has access to a dataset $D_{t'}$ with at least $n_{t'}$ samples drawn from P up to time t' , then it outputs a hypothesis sequence $\hat{h} \in \mathcal{H}$ such that it satisfies

$$\lim_{T \rightarrow \infty} \frac{1}{T - t'} \int_{t'}^T dt \mathbb{P}_{D_{t'}} [|R_t(\hat{h}) - R_t(h)| < \epsilon] \geq 1 - \delta. \quad (2)$$

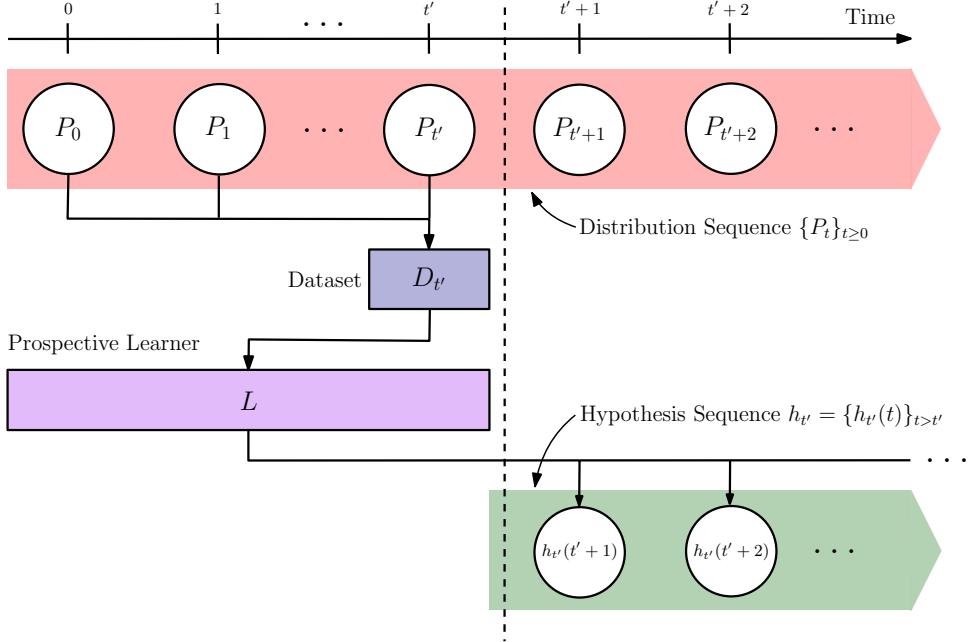


Figure 2: A schematic diagram illustrating the objective of prospective learning and the setup leading up to the Definition 2. $P = \{P_t\}_{t \geq 0}$ is the sequence of distributions and $D_{t'}$ denotes a dataset drawn from P , over the finite time interval $0 \leq t \leq t'$. Using the prospective learner L that takes $D_{t'}$ as input, we wish to estimate a time-indexed sequence of decision functions or hypotheses $\hat{h}_{t'}(t)$ that generalizes to future tasks encountered after time t' . Therefore, the focus of prospective learning is to learn a sequence of hypotheses from data received up to t' such that it generalizes to future tasks after time t' .

Here, $n_{t'}$ and \bar{t} would depend on the sequence P , the complexity of the hypothesis class \mathcal{H} , the sequence of risks $\{R_t\}_{t \geq 0}$, and the choice of ϵ and δ .

In words, the sequence of hypotheses \hat{h} , learned by L from all the data points $D_{t'}$ received up to time $t' > \bar{t}$, makes predictions, with precision no worse than ϵ and with a probability of at least $1 - \delta$ on future tasks beyond time t' , with respect to a reference hypothesis sequence h .

Definition 2 offers some flexibility when choosing the reference hypothesis sequence h ; the performance of the prospective learner L is measured against this reference. Some reasonable choices for h are:

1. Strong reference h^* , the optimal sequence of hypotheses given the entire distribution sequence P where h_t^* is the optimal hypothesis for P_t ;
2. Weak reference h^0 , the sequence of hypotheses that performs at random chance where h_t^0 is the hypothesis that performs at random chance with respect to P_t .

Depending on the “strictness” of these choices, the class of sequences P that are prospectively-learnable according to Definition 2 will vary.

In contrast to a retrospective learner, a prospective learner should respect Inequality 2 for all $t' > \bar{t}$ where \bar{t} is finite. This is a natural condition to impose because the prospective learner needs to see a few tasks before it can build an accurate model of how the tasks evolve in the future. Here, \bar{t} is the minimum amount of time that we need to observe data from the sequence P in order for P to be prospectively learnable. Therefore, \bar{t} is a function of P , the hypothesis class \mathcal{H} , the risk, R , in addition to ϵ, δ .

Equipped with this definition of prospective learnability, we can attempt to understand the properties of prospectively learnable sequences of distributions. Once we have defined the sequence of distributions that governs a given prospective learning setting (that is, those that govern changes in distributions, risks, and constraints), defining quantities that we use in retrospective learning theory, such as VC-dimension (Vapnik & Chervonenkis, 2015), is analogous. In this way, the primary research question for prospective learning is determining which distribution sequences are intrinsi-

cally low-dimensional (e.g., have a dependence structure), because otherwise learning cannot be achieved regardless of sample complexity.

When designing and implementing algorithms to perform prospective learning, it is important to determine the desirable properties of these algorithms and their output hypotheses. We state the following conjecture as an entry point to the discussion about the hypothesis sequences learned by prospective-learners.

Conjecture 1 (Continuum Hypothesis of Learning Hypotheses) Suppose $P = \{P_t\}_{t \geq 0}$ is a sequence of distributions. We conjecture that there are three distinct complexity classes for learning.

1. P is PAC-Learnable: With the introduction of potentially time-varying P_t , R_t and \mathcal{H}_t , PAC-learnability takes a different form. Specifically, whenever any of the above change in such a way that the relative risk, $|R_t(\hat{h}) - R_t(h)|$ remains bounded above by ϵ with probability $1 - \delta$, then it remains PAC-learnable. A simple example of such a scenario is in a classification problem, where the posteriors can change dramatically, but the discriminant boundaries remain fixed, so the relative risk (assuming risk is expected 0 – 1 loss) remains constant.
2. P is prospectively learnable: This is where the task evolves such that some complexity measure of the sequence P is low, in which case, we conjecture that a prospective learner can learn a sequence of hypotheses $\hat{h} \in \mathcal{H}$ after observing $n_{t'}$ amount of data drawn from P for a duration $t' > \bar{t}$ such that \hat{h} can make accurate predictions in the future. Note that \mathcal{H} , t and $n_{t'}$ would depend on the complexity of P . For instance, when the complexity of P is high, we may need to search for \hat{h} from a more complex class of hypotheses and observe a larger amount of data from P for a longer time.
3. P is not prospectively learnable: Learning such a sequence of distributions would be impossible and it requires us to observe the sequence P for an infinite period of time or collect an infinite amount of data from it. We deem such sequences too complicated to learn prospectively.

While we do not propose a complexity measure for a distribution sequence in this work, we note that quantities such as predictive information (Bialek et al., 2001), which compute the mutual information between the past and the future of a time series, can inspire the construction of an appropriate measure. We leave this and the discovery of further desirable properties of hypothesis sequences suitable for prospective learning as future work.

Conjecture 1 motivates the development of prospective learning algorithms that can output dynamic sequences of hypotheses. Such learners must update hypotheses as time evolves. Note that prospective learning as defined here is equivalent to retrospective learning if the distribution sequence is a constant. Thus traditional retrospective learning can be seen as a special case of prospective learning. But this formulation of prospective learning is also quite general, and it has relations with continual-, lifelong-, transfer-, meta-, online-, and OOD-learning. However, we leave the task of verifying whether the learnability of these paradigms is encapsulated by Definition 2 as future work.

4 RELATED WORK

Online Learning Both online learning (Cesa-Bianchi & Lugosi, 2006; Hoi et al., 2021; Ying & Pontil, 2008) and prospective learning deal with task sequences. The goal of the online learner is to iteratively estimate the hypothesis such that it minimizes regret, which is the difference between the learner’s risk and the risk of the best model in hindsight. It is usually studied using a worst-case analysis and we assume that the sequence of tasks can even be generated by an adversary. Many algorithms use a follow-the-leader (FTL) strategy which selects the best model averaged over the past, since it achieves an asymptotically optimal regret. As a result, an online learner that uses FTL doesn’t consider the temporal dynamics of the sequence and has no prospective abilities. The same is true for other popular algorithms like online gradient descent (OGD) and its variants.

In Figure 3, we evaluate FTL and OGD on two different alternating sequence of binary classification tasks in the online setting. In Figure 3a, we consider two tasks with their labels flipped and observe that both OGD and FTL usually have high test error when the task switches. Hence both learners fail to exploit the dynamics of the task sequence. In Figure 3b, we consider an alternating sequence of tasks where the first and second task have different class-conditional means. We observe that this sequence is weakly prospective learnable by FTL since it converges to some fixed hypothesis that works better than random on both tasks. We expect that a *bona fide* prospective learner

would successfully and optimally tackle both of these problems as it would model the temporal dynamics of the task sequence.

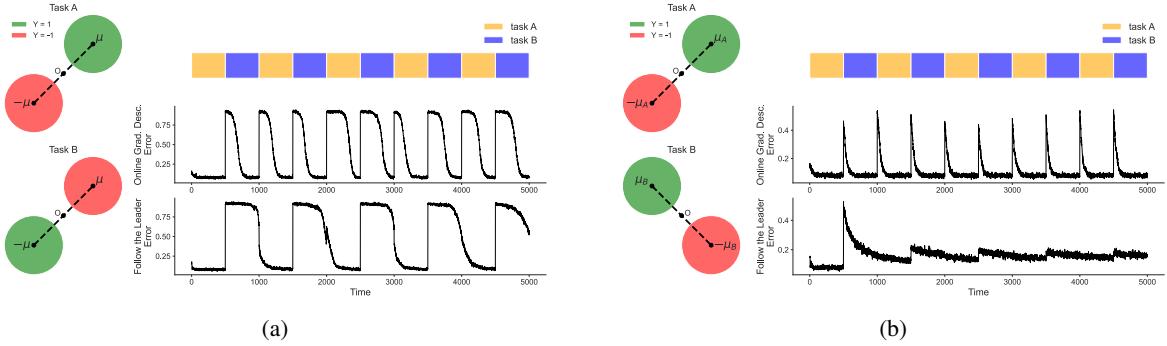


Figure 3: We consider two different alternating sequences of tasks where the learning task shifts between task A and task B periodically every 500 time-steps. In this online setting, we see one new sample every time-step. We plot the test error of the current task against the number of time-steps for FTL and OGD.

(a) For task A, the samples from one class belong to a Gaussian with mean $[1, 1]^\top$ and samples from the other class belong to a Gaussian with mean $[-1, -1]^\top$. Both Gaussians have isotropic covariance I . Task B is similar to Task A, except that the distributions for the two labels are interchanged. The OGD algorithm’s error peaks at each task switch and slowly reduces afterwards as it learns the new task by unlearning the one before. On the other hand, the FTL algorithm seems to learn task A successfully, but completely fails at learning task B. We speculate that for FTL, the model weights are biased to task A.

(b) For task A, the samples from one class belong to a Gaussian with mean $[1, 1]^\top$ and samples from the other class belong to a Gaussian with mean $[-1, -1]^\top$. For task B, the samples from one class belong to a Gaussian with mean $[1, -1]^\top$ and samples from the other class belong to a Gaussian with mean $[-1, 1]^\top$. All Gaussians have isotropic covariance I . Unlike (a), the class conditional means are completely different making the task sequence slightly easier. OGD learns each task much faster, but test error of the current task peaks at each task switch. Interestingly, FTL is able to progressively reduce its error as it observes more data from the sequence. It manages to converge to a fixed hypothesis that achieves reasonably low error on both task A and task B. However, FTL does not converge to the Bayes optimal hypothesis for either task since that requires it to be prospective.

In conclusion, neither OGD nor FTL attempt to model the temporal dynamics of the task sequence. We expect a (reasonable) prospective learner would achieve smaller error on both problems.

Online Meta-Learning Online meta-learning (Finn et al., 2019; Denevi et al., 2019; Yao et al., 2020) refers to a class of algorithms that implement meta-learning updates in an online manner. These methods seek to minimize a variant of regret which inspires the use of the follow-the-meta-leader strategy (Finn et al., 2019). In practice, this translates to implementing model-agnostic meta-learning (MAML) (Finn et al., 2017) but with the tasks sampled uniformly at random from all previously seen tasks. While this minimizes the worst-case regret, this choice is sub-optimal if we know that the task sequence evolves with time according to a predictable pattern. Hence, methods like online-MAML will fail to be strongly prospective learnable for some task sequences.

Continual Learning Unlike PAC Learning and online learning, continual learning (Van de Ven & Tolias, 2019; Zenke et al., 2017; Ramesh & Chaudhari, 2021; Vogelstein et al., 2020; Geisa et al., 2021) lacks a universally agreed-upon definition. In the current literature, it often refers to a wide gamut of scenarios that vary in the goal (e.g., avoiding catastrophic forgetting, positive forward/backward transfer), or storage and computational constraints. A lack of a universally agreed-upon definition limits our ability to formally delineate differences. That said, prospective learning resembles continual learning, albeit with some key differences. Suppose we would like to tackle a sequence of tasks where the task switches from A to B, and back, at regular intervals. Many existing continual learning methods will eventually learn both task A and B, and will therefore be able to switch back and forth with high accuracy—but shortly after the task switches. In other words, current continual algorithms will **fail to predict when the task will switch**. In contrast, a prospective learner, that satisfies Definition 2, will learn how the sequence of tasks evolves, and therefore predict when the task switches. Formally, this learner should perform at least as well as the continual learner that does not model the dynamics of the tasks. In simple terms, a continual learner would iteratively “adapt” to future tasks while a prospective learner would “anticipate” them *a priori*. We certainly expect continual learners to benefit

from prospective abilities and believe learners should ideally incorporate both adaptation and anticipation to actively learn from and make inferences about the future.

Streaming Learning The setting of streaming learning is also closely related to prospective learning but the objective can often vary. For example, the goal of [Gama et al. \(2013\)](#) is to minimize the error on samples from an underlying stationary process. More recent works on streaming learning such as [Hayes et al. \(2020\)](#) minimize the error on a fixed held-out dataset or on all the past data—neither of which emphasizes prospectiveness.

Out-of-distribution (OOD) Generalization The OOD generalization problem ([Shen et al., 2021](#); [Geisa et al., 2021](#); [Arjovsky, 2020](#); [Ye et al., 2022](#)) seeks to minimize the generalization error on an OOD task where the hypothesis is selected using data from a potentially different set of domains. OOD is typically defined for batch data, not sequential data. That said, we recover the OOD generalization setting if we assume that the distribution shifts at time t' —which is the point up to which we see data—and then never changes again.

Few-shot Learning Few-shot learning ([Triantafillou et al., 2019](#); [Wang et al., 2020](#)) refers to a broad range of settings where the model adapts to a new task using limited amounts of labeled data. Usually, the adaptation step has a limited computational budget but this is difficult to write down formally. Prospective learners are also expected to learn using very few samples with one key difference: they generalize to future tasks using ‘time’ as another context variable.

It is evident that when the stochastic process is a constant process, prospective learning becomes equivalent to PAC learning. Thus PAC learning can be seen as a special case of prospective learning. Furthermore, as discussed above, prospective learning is quite general and it has relations with various learning paradigms such as continual-, lifelong-, transfer-, meta-, online-, and OOD-learning. However, learning paradigms besides PAC learning and online learning lack formal and universally agreed-upon definitions of learnability. Therefore, concrete comparison of these paradigms with prospective learning is limited.

5 DISCUSSION

Classical machine learning limits the ability of learning systems to effectively learn in environments that (somewhat) predictably change. Here we have proposed an alternative to the classical learning problem that works around the assumption that training (past) and test (future) distributions are not identical, but predictably related to some degree. This reframing of the learning problem can produce flexible agents that effectively adapt in novel contexts with invariance along key dimensions relevant to learning. As a result, algorithms that follow a prospective learning framework can effectively learn for the future, overcoming a fundamental limitation of most modern AI.

We have provided an answer to the question “What is prospective learning?”, yet we have largely ignored the broader question of “how is prospective learning implemented?”. Assuming that natural intelligences (NI; including, but not necessarily limited to, humans) implement prospective learning [Seligman et al. \(2013\)](#), motivates synergistic research between NI and AI. Carefully designed, ecologically appropriate experiments across species, phyla, and substrates, will enable (i) quantifying the limitations and capabilities of existing NIs with respect to prospective learning criteria, and (ii) refining the mechanistic hypotheses generated by the AI theory.

Discovering *how* prospective learning is implemented in NIs may accelerate incorporating prospective learning capabilities in AI. For example, experiments in NI across taxa can establish milestones for experiments in AI (e.g., [\(Turing, 1950\)](#)). Implementing and deploying AIs and observing NIs exhibiting prospective learning ‘in the wild’ will provide real-world evidence to deepen our understanding of learning for the future. These AI implementations could be purely software, leverage specialized (neuromorphic) hardware ([Vogelstein et al., 2007](#); [Davies, 2019](#)), or even include wetware and hybrids ([Silva et al., 2020](#)), as prospective learning is, by definition, substrate independent.

In closing, there are many ways in which how we frame intelligence in machine learning differs from the kinds of problems that we really want to (and evolved to) solve in the real world, including both objective functions and prior knowledge. Other discrepancies are hard to even formalize ([Mitchell, 2019](#)). Here we have argued that the problem of shifts in the world can be formulated considerably better than to just appeal to out of distribution learning. We highlight how many current developments in machine learning can be seen as steps towards prospective learning. We believe that formulating the learning problem as a prospective one will produce important synergies between multiple branches of machine learning and cognitive sciences, providing a more thorough foundation and new ways of providing logical precision.

ACKNOWLEDGEMENTS

This paper was supported by an NSF AI Institute Planning award (# 2020312), as well as support from Microsoft Research, and DARPA. RR and PC were supported by grants from the National Science Foundation (IIS-2145164, CCF-2212519), the Office of Naval Research (N00014-22-1-2255), and cloud computing credits from Amazon Web Services. The authors would like to especially thank Kathryn Vogelstein for putting up with endless meetings at the Vogelstein residence in order to make these ideas come to life.

THE FUTURE LEARNING COLLECTIVE

Ashwin De Silva^{1*}; Rahul Ramesh^{2*}; Lyle Ungar²; Marshall Hussain Shuler¹; Noah J. Cowan¹; Michael Platt²; Chen Li¹; Leyla Isik¹; Seung-Eon Roh¹; Adam Charles¹; Archana Venkataraman¹; Brian Caffo¹; Javier J. How¹; Justus M Kebschull¹; John W. Krakauer¹; Maxim Bichuch¹; Kaleab Alemayehu Kinfu¹; Eva Yezerets¹; Dinesh Jayaraman²; Jong M. Shin¹; Soledad Villar¹; Ian Phillips¹; Carey E. Priebe¹; Thomas Hartung¹; Michael I. Miller¹; Jayanta Dey¹; Ningyuan (Teresa) Huang¹; Eric Eaton²; Ralph Etienne-Cummings¹; Elizabeth L. Ogburn¹; Randal Burns¹; Onyema Osuagwu⁴; Brett Mensh⁵; Alysson R. Muotri⁶; Julia Brown⁷; Chris White⁸; Weiwei Yang⁸; Andrei A. Rusu⁹ Timothy Verstynen³; Konrad P. Kording²; Pratik Chaudhari²✉; Joshua T. Vogelstein¹✉;

¹ Johns Hopkins University; ² University of Pennsylvania; ³ Carnegie Mellon University; ⁴ Morgan State University ⁵ Howard Hughes Medical Institute; ⁶ University of California, San Diego ⁷ MindX; ⁸ Microsoft Research ⁹ DeepMind

* Equal contribution

✉ Corresponding authors: pratikac@seas.upenn.edu, jovo@progl.ai

REFERENCES

Michael A Alcorn, Qi Li, Zhitao Gong, Chengfei Wang, Long Mai, Wei-Shinn Ku, and Anh Nguyen. Strike (with) a pose: Neural networks are easily fooled by strange poses of familiar objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4845–4854. openaccess.thecvf.com, 2019.

Martin Arjovsky. *Out of distribution generalization in machine learning*. PhD thesis, New York University, 2020.

Martin Arjovsky. Out of distribution generalization in machine learning, March 2021.

Jonathan Baxter. A model of inductive bias learning. *J. Artif. Intell. Res.*, 12(1):149–198, March 2000.

Richard Bellman. The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6):503–515, 1954.

Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, volume 2. researchgate.net, 1992.

William Bialek, Ilya Nemenman, and Naftali Tishby. Predictability, complexity, and learning. *Neural computation*, 13(11):2409–2463, 2001.

Stevo Bozinovski and Ante Fulgosi. The influence of pattern similarity and transfer learning upon the training of a base perceptron B2. *Proceedings of Symposium Informatica*, pp. 3–121–5, 1976.

Francesco Paolo Cantelli. Sulla determinazione empirica delle leggi di probabilita. *Giorn. Ist. Ital. Attuari*, 4(421-424), 1933.

Rich Caruana. Multitask learning. *Mach. Learn.*, 28(1):41–75, July 1997.

Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, March 2006.

M Davies. Progress in Neuromorphic Computing : Drawing Inspiration from Nature for Gains in AI and Computing. In *2019 International Symposium on VLSI Design, Automation and Test (VLSI-DAT)*, pp. 1–1, April 2019.

Giulia Denevi, Dimitris Stamos, Carlo Ciliberto, and Massimiliano Pontil. Online-within-online meta-learning. *Advances in Neural Information Processing Systems*, 32, 2019.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*, January 2019.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, pp. 1126–1135. PMLR, 2017.

Chelsea Finn, Aravind Rajeswaran, Sham Kakade, and Sergey Levine. Online meta-learning. In *International Conference on Machine Learning*, pp. 1920–1930. PMLR, 2019.

François Fleuret, Ting Li, Charles Dubout, Emma K Wampler, Steven Yantis, and Donald Geman. Comparing machines and humans on a visual categorization test. *Proc. Natl. Acad. Sci. U. S. A.*, 108(43):17621–17625, October 2011.

Joao Gama, Raquel Sebastiao, and Pedro Pereira Rodrigues. On evaluating stream learning algorithms. *Machine learning*, 90:317–346, 2013.

Ali Geisa, Ronak Mehta, Hayden S Helm, Jayanta Dey, Eric Eaton, Jeffery Dick, Carey E Priebe, and Joshua T Vogelstein. Towards a theory of out-of-distribution learning. *arXiv preprint arXiv:2109.14501*, 2021.

M Giurfa, S Zhang, A Jenett, R Menzel, and M V Srinivasan. The concepts of ‘sameness’ and ‘difference’ in an insect. *Nature*, 410(6831):930–933, April 2001.

V Glivenko. Sulla determinazione empirica delle leggi di probabilita. *Gion. Ist. Ital. Attauri.*, 4:92–99, 1933.

David J Hand. Classifier technology and the illusion of progress. *Statistical Science*, 21(1):1–14, February 2006.

Tyler L Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII* 16, pp. 466–483. Springer, 2020.

Steven CH Hoi, Doyen Sahoo, Jing Lu, and Peilin Zhao. Online learning: A comprehensive survey. *Neurocomputing*, 459:249–289, 2021.

Peter J Huber and Elvezio M Ronchetti. *Robust Statistics*. Wiley, 2 edition edition, February 2009.

Michael J Kearns and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994.

Junkyung Kim, Matthew Ricci, and Thomas Serre. Not-So-CLEVR: learning same-different relations strains feedforward neural networks. *Interface Focus*, 8(4):20180011, August 2018.

A Krizhevsky, I Sutskever, and G E Hinton. Imagenet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.*, 2012.

Angeliki Lazaridou, Adhi Kuncoro, Elena Gribovskaya, Devang Agrawal, Adam Liska, Tayfun Terzi, Mai Gimenez, Cyprien de Masson d’Autume, Tomas Kociský, Sebastian Ruder, et al. Mind the gap: Assessing temporal generalization in neural language models. *Advances in Neural Information Processing Systems*, 34:29348–29363, 2021.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

Scott Mayer McKinney, Marcin Sieniek, Varun Godbole, Jonathan Godwin, Natasha Antropova, Hutan Ashrafian, Trevor Back, Mary Chesus, Greg S Corrado, Ara Darzi, et al. International evaluation of an ai system for breast cancer screening. *Nature*, 577(7788):89–94, 2020.

M Mitchell. Artificial intelligence hits the barrier of meaning. *Information*, 2019.

Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of Machine Learning*. MIT Press, November 2018.

Francesco Orabona. *A Modern Introduction to Online Learning*. December 2019.

Rahul Ramesh and Pratik Chaudhari. Model zoo: A “growing” brain” that learns continually. *arXiv preprint arXiv:2106.03027*, 2021.

Benjamin Recht. A Tour of Reinforcement Learning: The View from Continuous Control. *Annual Review of Control, Robotics, and Autonomous Systems*, May 2019.

M B Ring. *Continual learning in reinforcement environments*. PhD thesis, University of Texas at Austin, 1994.

Martin EP Seligman, Peter Railton, Roy F Baumeister, and Chandra Sripada. Navigating into the future or driven by the past. *Perspectives on psychological science*, 8(2):119–141, 2013.

Andrew W Senior, Richard Evans, John Jumper, James Kirkpatrick, Laurent Sifre, Tim Green, Chongli Qin, Augustin Žídek, Alexander WR Nelson, Alex Bridgland, et al. Improved protein structure prediction using potentials from deep learning. *Nature*, 577(7792):706–710, 2020.

Zheyuan Shen, Jiashuo Liu, Yue He, Xingxuan Zhang, Renzhe Xu, Han Yu, and Peng Cui. Towards out-of-distribution generalization: A survey. *arXiv preprint arXiv:2108.13624*, 2021.

G A Silva, A R Muotri, and C White. Understanding the human brain using brain organoids and a Structure-Function theory, 2020.

Sebastian Thrun. Is learning the n-th thing any easier than learning the first? In D S Touretzky, M C Mozer, and M E Hasselmo (eds.), *Advances in Neural Information Processing Systems* 8, pp. 640–646. MIT Press, 1996.

Sebastian Thrun and Tom M Mitchell. Lifelong robot learning. *Rob. Auton. Syst.*, 15(1):25–46, July 1995.

Eleni Triantafillou, Tyler Zhu, Vincent Dumoulin, Pascal Lamblin, Utku Evci, Kelvin Xu, Ross Goroshin, Carles Gelada, Kevin Swersky, Pierre-Antoine Manzagol, et al. Meta-dataset: A dataset of datasets for learning to learn from few examples. *arXiv preprint arXiv:1903.03096*, 2019.

A M Turing. COMPUTING MACHINERY AND INTELLIGENCE. *Mind*, LIX(236):433–460, October 1950.

L G Valiant. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, November 1984.

Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.

V Vapnik and A Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. *Theory Probab. Appl.*, 16(2):264–280, January 1971.

Vladimir Vapnik. Principles of risk minimization for learning theory. *Advances in neural information processing systems*, 4, 1991.

Vladimir N Vapnik and A Ya Chervonenkis. On the uniform convergence of relative frequencies of events to their probabilities. In *Measures of complexity*, pp. 11–30. Springer, 2015.

Ben Vermaercke, Elys Cop, Sam Willems, Rudi D’Hooge, and Hans P Op de Beeck. More complex brains are not always better: rats outperform humans in implicit category-based generalization by implementing a similarity-based strategy. *Psychon. Bull. Rev.*, 21(4):1080–1086, August 2014.

Joshua T Vogelstein, Jayanta Dey, Hayden S Helm, Will LeVine, Ronak D Mehta, Tyler M Tomita, Haoyin Xu, Ali Geisa, Qingyang Wang, Gido M van de Ven, et al. Representation ensembling for synergistic lifelong learning with quasilinear complexity. *arXiv preprint arXiv:2004.12908*, 2020.

R Jacob Vogelstein, Udayan Mallik, Joshua T Vogelstein, and Gert Cauwenberghs. Dynamically reconfigurable silicon array of spiking neurons with conductance-based synapses. *IEEE Trans. Neural Netw.*, 18(1):253–265, January 2007.

Jane X Wang. Meta-learning in natural and artificial intelligence, November 2020.

Yaqing Wang, Quanming Yao, James T Kwok, and Lionel M Ni. Generalizing from a few examples: A survey on few-shot learning. *ACM computing surveys (csur)*, 53(3):1–34, 2020.

Halbert White. Maximum Likelihood Estimation of Misspecified Models. *Econometrica*, 50(1):1–25, 1982.

Huaxiu Yao, Yingbo Zhou, Mehrdad Mahdavi, Zhenhui Jessie Li, Richard Socher, and Caiming Xiong. Online structured meta-learning. *Advances in Neural Information Processing Systems*, 33:6779–6790, 2020.

Huaxiu Yao, Caroline Choi, Bochuan Cao, Yoonho Lee, Pang Wei W Koh, and Chelsea Finn. Wild-time: A benchmark of in-the-wild distribution shift over time. *Advances in Neural Information Processing Systems*, 35:10309–10324, 2022.

Nanyang Ye, Kaican Li, Haoyue Bai, Runpeng Yu, Lanqing Hong, Fengwei Zhou, Zhenguo Li, and Jun Zhu. Ood-bench: Quantifying and understanding two dimensions of out-of-distribution generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7947–7958, 2022.

Yiming Ying and Massimiliano Pontil. Online gradient descent learning algorithms. *Foundations of Computational Mathematics*, 8:561–596, 2008.

Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International conference on machine learning*, pp. 3987–3995. PMLR, 2017.