Fast rates for prediction with limited expert advice

Anonymous Author(s)
Affiliation
Address
email

Abstract

We investigate the problem of minimizing the excess generalization error with respect to the best expert prediction in a finite family in the stochastic setting, under limited access to information. We assume that the learner only has access to a limited number of expert advices per training round, as well as for prediction. Assuming that the loss function is Lipschitz and strongly convex, we show that if we are allowed to see the advice of only one expert per round for T rounds in the training phase, or to use the advice of only one expert for prediction in the test phase, the worst-case excess risk is $\Omega(1/\sqrt{T})$ with probability lower bounded by a constant. However, if we are allowed to see at least two actively chosen expert advices per training round and use at least two experts for prediction, the fast rate $\mathcal{O}(1/T)$ can be achieved. We design novel algorithms achieving this rate in this setting, and in the setting where the learner has a budget constraint on the total number of observed experts advices, and give precise instance-dependent bounds on the number of training rounds and queries needed to achieve a given generalization error precision.

1 Introduction

In many real-world applications, a forecaster has access to many expert advices in order to make a prediction. In some situations, it may be overly expensive to query the advice of all experts at each time. For example, it is unrealistic to ask for the advice of all available doctors or to run a large battery of tests on each patient. One then wants to select a small number of opinions and make decisions based on their outcome. In a more typical machine learning scenario, each "expert" might be a fixed prediction method which has been already trained, albeit based on different sets of parameters or methodology; the goal then amounts to predictor selection or aggregation, in a situation where the computation of each single prediction constitutes the bottleneck cost, rather than data acquisition. This is an instance of *efficient model selection* problem.

In this paper we investigate the setting where the decision maker can only observe m out of K experts' advice and use p out of K advices for prediction. The motivation for this setting is not only minimizing monetary cost. In certain situations, such as in web search engine where a ranking of pages must be generated for each user incoming query, there is a constraint on the response time. It is therefore possible to only compute the prediction of a small number of models. Furthermore, in a recommender system, where experts may be seen as the user's attributes, it is important to develop an attribute efficient algorithm in order to prevent the acquisition of full information about each user for privacy-preserving reasons. Our emphasis in this work is on obtaining *fast convergence rates* guarantees on the excess risk (which are possible under a strong convexity assumption of the loss),

Submitted to 35th Conference on Neural Information Processing Systems (NeurIPS 2021). Do not distribute.

with high probability, and with precise efficiency guarantees on the number of training expert queries needed to achieve a given precision.

2 Related Work and Main Contributions

Games with limited feedback (slow rates): Our work investigates what happens between the full information and single-point feedback games. Learning with a restricted access to information was considered under various settings in [6], [19], [12], [20], [5]. A setting close to ours was considered in [21], where the agent chooses in each round a subset of experts to observe their advice, then follows the prediction of one expert. To minimize the cumulative regret in the adversarial setting, they used an extension of the Exp3 algorithm, which allows to have an excess risk of $\mathcal{O}(\sqrt{1/T})$ in the limited feedback setting and $\mathcal{O}(\sqrt{\log(C)/C})$ in the budgeted case with a budget C.

The differences in the setting considered here is that (a) we are interested in the generalization error in the stochastic setting rather than the cumulative regret in an adversarial setting and (b) our assumptions of the convexity of the loss allow for the possibility of fast excess risk convergence. Moreover, we consider the more general case where the player is allowed to combine p out of K experts for prediction. The possibility of playing a subset of arms was considered in the literature of Multiple Play Multi-armed bandits. It was treated with a budget constraint in [26] for example (see also [24]), where at each round, exactly p out of K possible arms have to be played. In addition to observing the individual rewards for each arm played, the player also learns a vector of costs which has to be covered with an a-priori defined budget C. In the stochastic setting, a UCB-type procedure gives a bound for the cumulative regret of $\mathcal{O}(\Delta_{\min}^{-1}\log(C)/C)$ that holds only in expectation, where $\Delta_{\min}^{-1}$ denotes the gap between the best choice of arms and the second best choice. This bound leads to an instance dependent bound of $\mathcal{O}(\sqrt{\log(C)/C})$ in the worst case. In the adversarial setting, an extension of Exp3 procedure gives a bound of $\mathcal{O}(\sqrt{\log(C)/C})$ for the cumulative regret that holds with high probability. In another online problem, where the objective is to minimize the cumulative regret in an adversarial setting with a small effective range of losses, [11] have shown the impossibility of regret scaling with the effective range of losses in the bandit setting, while [23] showed that it is possible to circumvent this impossibility result if the player is allowed one additional observation per round. However, in settings considered, it is impossible to achieve a regret dependence on T better than the rate of $\mathcal{O}(1/\sqrt{T})$.

Fast rates in the full information setting: The learning task of doing as well as the best expert of a finite family in the sense of generalization error has been studied quite extensively in the full information case. In an adversarial setting, it is well known that under suitable assumptions on the loss function (typically related to strong convexity), an appropriately tuned exponential weighted average (EWA) strategy has cumulative regret bounded by the "fast rate" $\mathcal{O}(\log(K)/T)$ [13, 9, 4], which, combined with the online-to-batch conversion principle [8, 4] (also known as progressive mixture rule, [7, 25]), yields a bound of the same order for the *expected* excess prediction risk in the stochastic case. However, it was shown that progressive mixture type rules are *deviation suboptimal* for prediction [2], that is, their excess risk takes a value larger than $c/\sqrt{T}$ with constant positive probability over the training phase. To lift the apparent contradiction between the two last statements, consider that the excess risk of the EWA can take *negative* values, since it is an *improper* learning rule. Thus negative and positive "large" deviations can compensate each other so that the expectation is small. The inefficiency of EWA in deviation is a significant drawback, and alternatives to the EWA progressive mixture rule that achieve $\mathcal{O}(\log(K)/T)$ excess prediction risk with high probability were proposed by [17] and [3]. In [17], the strategy consists in whittling down the set of experts by elimination of obviously suboptimal experts, and performing empirical risk minimization (ERM) over the convex combinations of the remaining experts. In [3], the *empirical star* algorithm consists in performing an ERM over all segments consisting of a two-point convex combination of the ERM expert and any other expert. Note that the empirical star algorithm has the advantage that the final prediction rule is a convex combination of (at most) *two* experts.

Linear regression with partially observed attributes: Other related work is that of [10], and [14] on learning linear regression models with partially observed attributes. The most related setting to ours is the local budget setting, where the learner is allowed to output a linear combination of features for prediction. The key idea is to use the observed attributes in order to build an unbiased estimate of the full information sample, then to use an optimization procedure to minimize the penalized empirical loss. In our setting, the minimization of penalized empirical loss was shown to be suboptimal (see [16]). Moreover, while we want to predict as well as the best expert, in [10], the objective is to be as good as the best linear combination of features with a small additive term (the optimal rate, in this case, is $\mathcal{O}\big(1/\sqrt{T}\big)$). Finally, we consider that the restriction on observed attributes (experts advice) does not apply only to the training samples but also to the testing data.

Online convex optimization with limited feedback: The idea of using multiple point feedback to achieve faster rates appeared in the online convex optimization literature (see [1], and [22]). It was shown that in the setting where the adversary chooses a loss function in each round if the player is allowed to query this function in two points, it is possible to achieve minimax rates that are close to those achievable in the full information setting. The key idea is to build a randomized estimate of the gradients, which are then fed into standard first-order algorithms. These ideas are not convertible into our setting because we consider a non-convex set of experts.

Our contributions: The present work focuses on excess generalization error with respect to the best expert in a finite set. Motivated by high probability bounds on the excess risk, we focus on the stochastic setting, where the target variable and experts' advice follow a fixed unknown distribution. We develop and analyse novel algorithms that have fast rates with high probability. We suppose that the learner has limited access to data. Two different formalisms are considered: the budgeted setting, where we assume that the learner has a global budget constraint on the total number of experts' predictions, and the limited experts' prediction setting, where there is a constraint on the number of observed experts advice per round. We develop high probability distribution dependent and distribution independent bounds for each setting. Furthermore, we show that having access to more than one expert per round and combining at least two experts for prediction is necessary to have fast rates ($\mathcal{O}(1/T)$).

The remainder of this paper is organized as follows; in the next section, we formalize our setting and present the game protocol. Section 4 gives a new algorithmic idea to achieve fast rates in the full information case. We extend this idea to the restricted access to information setting (under global budget constraint in Section 5 and under local budget constraint in Section 6). Finally, we provide lower bounds for some particular settings, where fast rates cannot be achieved. The proofs of the results are found in the supplemental material; references to the supplemental are indicated with the prefix S-.

3 Setting

We denote $[\![n]\!] = \{1, \dots, n\}$ for a nonnegative integer n. Given a finite set $\mathcal{G} := \{F_1, \dots, F_K\}$ of experts and a target variable Y, we consider the task of predicting (at least) as well as the best expert in $\mathcal{G}$. In each training round t, we suppose that $Z_t := (Y_t, F_{1,t}, \dots, F_{K,t})$ is drawn independently from the past from the same probability distribution denoted $\mathcal{D}$ on $\mathcal{Y} \times \mathcal{X}^K$, where $\mathcal{Y} \subseteq \mathbb{R}^d$ and $\mathcal{X} \subseteq \mathbb{R}^d$ are closed and convex. In each round, the agent has access to Y_t and only actively chosen $m \geq 1$ experts' predictions. When the game is halted after T rounds, the forecaster chooses a subset $\widehat{I} \subset [\![K]\!]$ experts, $|\widehat{I}| = p \geq 1$, an aggregation function $\hat{g} : \mathcal{X}^p \to \mathcal{X}$, and outputs the prediction, on a new example indexed $(T+1)$, given by $\widehat{F} = \hat{g}((F_{i,T+1})_{i \in \widehat{I}})$. The quality of the agent's output is measured by its generalization risk

$$R(\widehat{F}) := \mathbb{E}\big[l(\widehat{F}, Y_{T+1})\big],$$

where $l : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}_+$ is a loss function.

The agent's objective is to achieve a risk as close as possible to the risk of the best expert in $\mathcal{G}$, $R^* = \min_{i\in[\![K]\!]} R(F_i)$, where $R(F_i) = \mathbb{E}[l(F_i, Y)]$, and we measure the performance of a prediction rule via its excess risk with respect to that best expert.

We consider that a function $f : \mathcal{X} \to \mathbb{R}$ is L-Lipschitz if: $\forall x, y \in \mathcal{X}$:$|f(x) - f(y)| \leq L\|x - y\|$. Furthermore, we consider that a function f is ρ-strongly convex if the function: $x \to f(x) - \frac{\rho^2}{2}\|x\|^2$ is convex.

Assumption 1. *$\forall y \in \mathcal{Y}$: $l(., y)$ is L-Lipschitz and ρ-strongly convex.*

Remarks. This assumption implies that the diameter of $\mathcal{X}$ is bounded by $8L/\rho^2$ and the quantity $\sup_{x,x'\in\mathcal{X},y\in\mathcal{Y}}|l(x, y) - l(x', y)|$ is bounded by $B := 8L^2/\rho^2$. Consequently, without loss of generality we can assume that the loss is bounded by B (see Lemma S-1 and subsequent discussion for details). It is satisfied, for example, in the following setting: least square loss $l(x, y) = (y - x)^2$ where $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ with $\mathcal{X}$ and $\mathcal{Y}$ are bounded subsets of $\mathbb{R}^d$. Prior knowledge on ρ is not necessary if L and an upper bound on the the l_∞ norm of target variable Y and the experts are known.

Algorithm 1 summarizes the game protocol.

Algorithm 1 The Game Protocol

Parameters: a budget on the number of observed experts m, a budget on the number of experts for prediction p, ρ and L.
for each round $t = 1, 2, \ldots$ **do**
  The environment secretly draws y_t.
  Each expert $i \in [\![K]\!]$ gives a prediction $F_{i,t}$ (the learner has not access to these predictions).
  The forecaster chooses a subset of experts of size m and observes y_t and their predictions.
end for
On interrupt: The forecaster chooses a subset of experts of size p, and p positive weights summing to 1, and outputs as a prediction for future rounds the corresponding convex combination of experts.

4 The full information case ($m = p = K$)

In this section, we revisit the "classical" case where there is no constraint on the number of expert queries per observation round; assume the output of all experts are observed for T rounds (in other words, T i.i.d. training examples), which is the full information or "batch" setting. We want to output a final prediction rule with prediction risk controlled with high probability over the training phase.

We start with putting forward an apparently new rule , simpler than existing ones [17, 3], for the full information setting which, like the empirical star [3], outputs a convex combination of two experts. In contrast to the latter, our rule does not need any optimization over a union of segments. The underlying principle will guide us to construct a budget efficient expert selection rule in the sequel.

Define $\hat{R}(F_i) := T^{-1}\sum_{t=1}^{T} l(F_{i,t}, Y_t)$ the empirical loss of expert i, and $\hat{d}_{ij} := (T^{-1}\sum_{t=1}^{T}(F_{i,t} - F_{j,t})^2)^{\frac{1}{2}}$ the empirical L_2 distance between experts i and j over T rounds. Finally let $\alpha = \alpha(\delta) := (\log(4K\delta^{-1})/T)^{\frac{1}{2}}$, where $\delta \in (0, 1)$ is a fixed confidence parameter. Define

$$\Delta_{ij} := \hat{R}(F_j) - \hat{R}(F_i) - 6\alpha \max\left\{L\hat{d}_{ij}, B\alpha\right\}. \tag{1}$$

The quantity Δ_{ij} can be interpreted as a test statistic: if $\Delta_{ij} > 0$, then we have a guarantee that $R(F_j) > R(F_i)$, so that expert j is sub-optimal; this guarantee holds for all (i, j) uniformly with probability $(1 - \delta)$. It therefore makes sense to reduce the set of candidates to

$$S := \left\{j \in [\![K]\!] : \sup_{j\in[\![K]\!]} \Delta_{ij} \leq 0\right\}. \tag{2}$$

Our new full information setting rule is the following:

$$\text{choose } \bar{k} \in S \text{ arbitrarily ;} \quad \text{pick } \bar{j} \in \operatorname*{Arg\,Max}_{j\in S} \hat{d}_{\bar{k}j}; \quad \text{predict } \widehat{F} := \frac{1}{2}(F_{\bar{k}} + F_{\bar{j}}). \tag{3}$$

In words, the above rule consists in eliminating all experts that are manifestly outperformed by another one, and, among the remaining experts, pick two that are as different as possible (in terms of empirical L^2 distance) and output their simple average for prediction. The next theorem establishes fast convergence rate for the excess risk of this rule:

Theorem 1. If Assumption 1 holds, the prediction rule $\widehat{F}$ defined by (3) then with probability $1-3\delta$ over the training phase (c is an absolute constant):

$$R(\widehat{F}) \leq R^* + cB\frac{\log(4K\delta^{-1})}{T}.$$

Proof. Let $d_{ij}^2 = \mathbb{E}\big[(F_i - F_j)^2\big]$. The result hinges on the following high confidence control of risk differences, established in Corollary S-4 as a direct consequence of the empirical Bernstein's inequality: with probability at least $1-3\delta$, it holds:

$$\text{For all } i,j \in [\![K]\!]: \qquad \Delta_{ij} \leq (R_j - R_i) \leq \Delta_{ij} + 32\alpha \max(Ld_{ij}, B\alpha). \tag{4}$$

Let $i^* \in \operatorname{Arg\,Min}_{i\in[\![K]\!]} R_i$ be an optimal expert. Since $R_{i^*} - R_j \leq 0$ for all $j \in [\![K]\!]$, it follows that if (4) holds, then $i^* \in S$, from the definition of S. So if (4) holds, we have

$$\begin{aligned}
R\left(\frac{F_{\bar{k}} + F_{\bar{j}}}{2}\right) &\leq \frac{1}{2}\big(R_{\bar{k}} + R_{\bar{j}}\big) - \frac{\rho^2}{8} d_{\bar{k}\bar{j}}^2 \\
&= R^* + \frac{1}{2}\big((R_{\bar{k}} - R_{i^*}) + (R_{\bar{j}} - R_{i^*})\big) - \frac{\rho^2}{8} d_{\bar{k}\bar{j}}^2 \\
&\leq R^* + \frac{1}{2}\big(\Delta_{\bar{k}i^*} + \Delta_{\bar{j}i^*}\big) + 16\alpha\big(\max\big(Ld_{\bar{j}i^*}, B\alpha\big) + \max(Ld_{\bar{k}i^*}, B\alpha)\big) - \frac{\rho^2}{8} d_{\bar{k}\bar{j}}^2 \\
&\leq R^* + 32B\alpha^2 + 48L\alpha d_{\bar{k}\bar{j}} - \frac{\rho^2}{8} d_{\bar{k}\bar{j}}^2;
\end{aligned}$$

where we have used strong convexity of the loss (and therefore of $R(.)$ with respect to the L^2 distance) in the first line; the right-hand side of (4) in the third line; and, in the last line, the fact that $\bar{j}, \bar{k}, i^*$ are all in S along with $d_{\bar{j}i^*} \leq d_{\bar{j}\bar{k}} + d_{\bar{k}i^*} \leq 2d_{\bar{j}\bar{k}}$ by construction of $\bar{j}$. Finally upper bounding the value of the last bound by its maximum possible value as a function of $d_{\bar{k}\bar{j}}$ and recalling $B = 8L^2/\rho^2$, we obtain the statement. □

5 Budgeted Setting

In this section, we consider the budgeted setting. More precisely, given an a-priori defined budget C, at each round the decision-maker selects an arbitrary subset of experts to see their predictions. The player then pays a unit for each observed expert's advice. The game finishes when the budget is exhausted, at which point the player output a convex combination of experts for prediction.

We convert the batch rule defined in the full information setting to an "online" rule by performing the test $\Delta_{ji} > 0$ for each pair (i,j) after each allocation. If at any round an expert $i \in [\![K]\!]$ fails any of these tests (i.e $\exists j : \Delta_{ji} > 0$), it is no longer queried. This extension allows us to derive instance dependent bounds, which cover the rates obtained in the batch setting in the worst case.

Since the tests $\Delta_{ij} > 0$ are performed after each allocation, we introduce the following modification on the definition of Δ_{ij}, for concentration inequalities to hold uniformly over the runtime of the procedure. We define $\Delta_{ij}(t,\delta)$ as follows:

$$\Delta_{ij}(t,\delta) := \hat{R}(j,t) - \hat{R}(i,t) - 6\alpha(t, \delta/(t(t+1)) \max\big\{L\hat{d}_{ij}(t), B\alpha(t, \delta/(t(t+1))\big\}.$$

Let $i^* \in \operatorname{Arg\,Min}_{i\in[\![K]\!]} R(F_i)$. For $i,j \in [\![K]\!]$, we denote by $d_{ij} := (\mathbb{E}[(F_i - F_j)^2])^{1/2}$ the L_2 distance between the experts F_i and F_j. For $i \in [\![K]\!]$, we introduce the following quantity:

$$\Lambda_i := \max\left\{\frac{L^2 d_{ii^*}^2}{|R(F_i) - R(F_{i^*})|^2}; \frac{B}{R(F_i) - R(F_{i^*})}\right\}.$$

Algorithm 2 Budgeted aggregation

Input δ, L and ρ.
Initialization: $S \leftarrow [\![K]\!]$.
for $T = 1, 2, \ldots$ **do**
  Jointly query all the experts in S and update $\Delta_{ij} > 0$ for all i, j.
  For all $i, j \in [\![K]\!]$, if $\Delta_{ij} > 0$, eliminate j: $S \leftarrow S \setminus \{j\}$.
  if the budget is consumed **then**
    let $\bar{k} \in S$, and $\bar{l} \leftarrow \operatorname*{argmax}_{j \in S} \hat{d}_{\bar{k}j}$.
    Return $\frac{1}{2}(F_{\bar{k}} + F_{\bar{l}})$.
  end if
end for

Define the following set of experts:

$$\mathcal{S}_\epsilon = \left\{ i \in [\![K]\!] : \Lambda_i > \frac{1}{\epsilon} \right\},$$

and let $\mathcal{S}_\epsilon^c$ be its complementary.

Theorem 2. (Instance dependent bound) Suppose Assumption 1 holds. Let $C \geq K$ denote the global budget on queries and denote $\hat{g}$ the output of Algorithm 2 with inputs (δ, L, ρ) when the budget C runs out. For any $\epsilon \geq 0$, if

$$C > 578 \left(\sum_{i \in \mathcal{S}_\epsilon^c} \Lambda_i + \frac{|\mathcal{S}_\epsilon|}{\epsilon} \right) \log \left(K \delta^{-1} \left(\sum_{i \in \mathcal{S}_\epsilon^c} \Lambda_i + \frac{|\mathcal{S}_\epsilon|}{\epsilon} \right) \right),$$

then, with probability at least $1 - \delta$:

$$R(\hat{g}) \leq R^* + cB\,\epsilon,$$

where c is an absolute constant.

Remarks. Observe that the above result gives in particular a query budget bound for the problem of best expert identification in our setting, by taking $\epsilon = 0$, in which case the required expert query budget is of order $\sum_{i:\Lambda_i < +\infty} \Lambda_i$ up to logarithmic terms. We can compare this to the problem of best arm identification in a bandit setting (one arm pull/query per round); our setting can be cast into that framework by considering each expert as an arm and only recording the information of the loss of the asked expert. The known optimal query bound for best arm identification in the classical multi-armed bandits setting with loss/reward bounded by B is of order $\sum_{i:\Lambda_i < +\infty} \widetilde{\Lambda}_i$ [15], where $\widetilde{\Lambda}_i = B^2(R(F_i) - R(F_{i^*}))^{-2}$. Since the diameter of $\mathcal{X}$ is bounded by B/L (see Lemma S-1), it holds $\Lambda_i \leq \widetilde{\Lambda}_i$. Hence, for best expert identification, the bound of Theorem 2 improves upon the best arm identification bound, potentially by a significant margin (in particular concerning the contribution of suboptimal but close to optimal experts for which $d_{ii^*} \ll B/L$ and $R_i - R_{i^*} \ll B$). Again, the improvement is due to the Assumption 1 on the loss and the possibility to query several experts per round, which are not used when casting the problem as a classical bandit setting.

6 Two queries per round ($m = p = 2$)

In this section, we suppose that the decision-maker is constrained to see only two experts' advice per round ($m = 2$). We suppose that the horizon is unknown; when the game is halted, the player outputs a convex combination of at most two experts ($p = 2$). We will show that the rates obtained are as good as in the full information case in its dependence on the number of rounds T.

Algorithm 3 works as follows. To circumvent the limitation of observing only two experts per round, we select at the beginning an arbitrary expert. We call such an expert *the pivot* (denoted $\hat{k}$), and we let S denote the set containing the remainder of the experts. In each round the advice of the

pivot is jointly queried with an expert i from S, chosen using a deterministic uniform rule. Then the tests $\Delta'_{\hat{k}i} \leq 0$ and $\Delta'_{i\hat{k}} \leq 0$ are performed, where Δ'_{ij} is defined by (5). If i fails the test, which means that i is a suboptimal expert, it is eliminated from S. If $\hat{k}$ fails the test, which suggests that $R(F_i) < R(F_{\hat{k}})$, it is eliminated and i occupies the position of the pivot.

Finally, when the algorithm is halted, depending on the number of allocated samples, we choose either an empirical risk minimizer over the non-eliminated experts or the mean of two experts that were jointly queried the most. This rule allows our algorithm's output to enjoy the best of converge rates of the two methods.

We introduce the following notations: In round t, denote $T_{ij}(t)$ the number of samples where predictions of experts i and j were jointly queried and $T_i(t)$ the number of rounds where the prediction of expert i was queried. Denote $\hat{R}_{ij}(j,t)$ the empirical loss of expert i calculated using only the $T_{ij}(t)$ samples queried for (i,j) jointly. We define $\alpha_{ij}(t,\delta) := \sqrt{\frac{\log(4K\delta^{-1})}{T_{ij}(t)}}$ if $T_{ij}(t) > 0$ and $\alpha_{ij}(t) = \infty$ otherwise. Let $\hat{d}_{ij}(t)$ be the empirical L_2 distance between experts i and j based on the $T_{ij}(t)$ queried samples. Denote $\delta_t := \delta/(t(t+1))$. For $i,j \in [\![K]\!]$ we define:

$$\Delta'_{ij}(t,\delta) := \hat{R}_{ij}(j,t) - \hat{R}_{ij}(i,t) - 6\max\Big\{L\alpha_{ij}(t,\delta_t)\hat{d}_{ij}(t), B\alpha^2_{ij}(t,\delta_t)\Big\}. \tag{5}$$

Let $S \subseteq [\![K]\!]$ and $i \in [\![K]\!]$, we denote $T_{iS} = \max_{j\in S} T_{ij}$. Observe that by construction of Algorithm 3, we have: $\forall j, j' \in S_T : |T_{ij}(T) - T_{ij'}(T)| \leq 1$.

Algorithm 3 Two-point feedback

Input δ, L and ρ.
Initialization: $\hat{k} \leftarrow 1$, $S \leftarrow [\![K]\!] \setminus \{1\}$.
for $T = 1, 2, \ldots$ **do**
  Let $j \in \operatorname{Arg\,Min}_i T_{\hat{k}i}$.
  Query the experts $(\hat{k}, j)$ and update: $\Delta'_{\hat{k}j}, \Delta'_{j\hat{k}}, \hat{d}_{\hat{k}j}, T_{\hat{k}}$ and T_j.
  If $\Delta'_{\hat{k}j} > 0$ eliminate expert j: $S \leftarrow S \setminus \{j\}$.
  If $\Delta'_{j\hat{k}} > 0$, set j as the new pivot: $\hat{k} \leftarrow j$, $S \leftarrow S \setminus \{j\}$.
end for
On interrupt: Let $\bar{k}$ be the pivot satisfying: $\bar{k} \leftarrow \operatorname{argmax}_i T_{iS}$ and let $\bar{l} \leftarrow \operatorname*{argmax}_{j\in S} \hat{d}_{\bar{k}j}$.
Let q denote the empirical risk minimizer on $S \cup \{\hat{k}\}$.
if $T_{\bar{k}\bar{l}} > \sqrt{T_q \log(KT\delta^{-1})}$ **then**
  Return $\frac{1}{2}(F_{\bar{k}} + F_{\bar{l}})$.
else
  Return F_q.
end if

Observe that the instruction $j \in \operatorname{Arg\,Min}_i T_{\hat{k}i}(t)$ is equivalent to a sequential query of all the non-eliminated experts. Algorithm 3 is computationally efficient, its runtime is $\mathcal{O}(T)$ and its space complexity is $\mathcal{O}(K)$.

Our first result in this setting is an empirical bound. At any interruption time, it gives a bound on the excess risk, only depending on quantities available to the user, using the number of queries resulting from the querying strategy in Algorithm 3. We then use a worst-case bound on these quantities to develop an instance independent bound in Corollary 4.

Theorem 3. (Empirical bound) Suppose assumption 1 holds. Let $T \geq 1$, and denote $\hat{g}$ the output of Algorithm 3 with inputs (δ, L, ρ) in round T. Then with probability at least $1 - 3\delta$:

$$R(\hat{g}) \leq R^* + c\, B \min\left\{\frac{\log\big(TK\delta^{-1}\big)}{T_{\bar{k}\bar{l}}(T)}, \sqrt{\frac{\log(TK\delta^{-1})}{T_q(T)}}\right\}, \tag{6}$$

where $\bar{k}, \bar{l}$ and $\bar{q}$ are the experts in Algorithm 3 and c is an absolute constant.

Proof Sketch of Theorem 3 We start by noting that when running algorithm 3, the optimal experts ($i^* \in \operatorname{Arg\,Min}_{i\in[\![K]\!]} R(F_i)$) are never eliminated with high probability (Lemma S-5). This shows in particular, that when the procedure is terminated, we have either $\hat{k} = i^*$ or $i^* \in S_T$.

Then we show the following key result: in each round $t \leq T$, denote by $\hat{k}_t$ the pivot and by S_t the set of non-eliminated experts. For any expert $i \in [\![K]\!]$, let $j \in \operatorname{Arg\,Max}_{l\in S_t} \hat{d}_{il}(t)$, we have with probability at least $1-\delta$:

$$R\left(\frac{F_i + F_j}{2}\right) \leq R^* + cB\,\frac{\log(K\delta_t^{-1})}{T_{ij}(t)}.$$

For this bound to be meaningful, either i or j have to be a pivot at some point (otherwise $T_{ij}(t) = 0$). In fact, this inequality is satisfied for all pairs (i, j), where i is a pivot and $j \in \operatorname{Arg\,Max}_{l\in S_t} \hat{d}_{il}(t)$. The sharpest bound is therefore obtained for the pivot that was jointly queried the most with elements of S_T. For the second bound, recall that i^* belongs to $S_T \cup \{\hat{k}_T\}$ with high probability. Therefore, performing an empirical risk minimization over the set of non-eliminated experts leads to the bound $\sqrt{\frac{\log(KT\delta^{-1})}{T_q(T)}}$, through a simple concentration argument using Hoeffding's inequality.

Corollary 4. (Instance independent bound) Suppose assumption 1 holds. Let $T \geq 1$, and denote $\hat{g}$ the output of algorithm 3 with inputs (δ, L, ρ) in round T. Then with probability at least $1-3\delta$:

$$R(\hat{g}) \leq R^* + c\,B \min\left\{\frac{K^2\log\big(TK\delta^{-1}\big)}{T}, \sqrt{\frac{K\log(TK\delta^{-1})}{T}}\right\},$$

where c is an absolute constant.

Remarks. Observe that in all the considered settings (full information, budgeted and limited advice), the number of jointly sampled pairs (F_i, F_j) to attain an excess risk of $\mathcal{O}(\epsilon)$ is of the order of $\mathcal{O}(K^2/\epsilon)$.

Theorem 5. (Instance dependent bound) Suppose assumption 1 holds. Let $\hat{g}$ denote the output of Algorithm 3 with input (δ, L, ρ) and T denote the total number of rounds. If :

$$T > 578\,K\log\left(\frac{K}{\delta}\min\left\{\frac{1}{\epsilon}, \Lambda^*\right\}\right)\left(\sum_{i\in\mathcal{S}_\epsilon^c}\Lambda_i + \min\left\{\frac{1}{\epsilon}, \Lambda^*\right\}\right),$$

where $\Lambda^* = \max\{\Lambda_i; \Lambda_i < +\infty\}$. We have, with probability at least $1-\delta$:

$$R(\hat{g}) \leq R^* + cB\,\epsilon,$$

where c is an absolute constant.

Remarks. This instance dependent bound, which recovers the bound in Corollary 4 in the worst case, is a consequence of the allocation strategy in Algorithm 3. Had we used a uniform allocation strategy over all pairs (F_i, F_j) and apply the full-information bound, we would lose the ability of our algorithm to adapt to the case where many experts are clearly suboptimal.
If the algorithm is allowed to query $m > 2$ experts advices per round, then it can be modified to attain an improved excess risk. We present this extension in Section S-8 in the appendix, and prove that it leads to a rate of $\mathcal{O}\Big(\frac{(K/m)^2}{T}\log(KT/\delta)\Big)$, which interpolates for intermediate values of m.

Proof Sketch of Theorem 5 First, we develop instance-dependent upper and lower bound for $T_{ij}(t)$, for any $i, j \in [\![K]\!]$ such that: $R(F_i) \neq R(F_j)$. To do this we introduce the following lemma (see lemma S-7 in the appendix):

Lemma 6. Let $i, j \in [\![K]\!]$ such that $R(F_i) \neq R(F_j)$. With probability at least $1-4\delta$, for all $t \geq 1$:

If:

$$T_{ij}(t) \geq 289\log\big(K\delta_t^{-1}\big)\max\left\{\frac{L^2 d_{ij}^2}{|R(F_i)-R(F_j)|^2}; \frac{B}{|R(F_i)-R(F_j)|}\right\}.$$

Then we have either $\Delta'_{ij} > 0$ or $\Delta'_{ji} > 0$.

Furthermore, if:

$$T_{ij}(t) \leq 3\log\big(K\delta_t^{-1}\big)\max\left\{\frac{L^2 d_{ij}^2}{|R(F_i)-R(F_j)|^2}; \frac{B}{|R(F_i)-R(F_j)|}\right\}.$$

Then we have: $\Delta'_{ij} \leq 0$ and $\Delta'_{ji} \leq 0$.

This lemma shows in particular that if an expert $i \in [\![K]\!]$ satisfying $\Lambda_i > \frac{1}{\epsilon}$ occupies the position of the pivot at a certain point, then it will be queried jointly with i^* (and therefore with the experts in S_T) at least $\Omega\Big(\frac{\log(KT/\delta)}{\epsilon}\Big)$ rounds. This implies in particular that the output satisfies $T_{\bar{k}\bar{l}} = \Omega(1/\epsilon)$. To conclude, we only need to have enough rounds to make sure all the experts j satisfying $\Lambda_j < \frac{1}{\epsilon}$ where eliminated.

7 Lower Bounds for $m = 1$ or $p = 1$

This section considers the case where the agent is restricted to selecting one expert at the end of the procedure ($p = 1$), and the case where the learner is restricted to see only one feedback per round ($m = 1$). We show that in both cases it is impossible to do better than $\mathcal{O}\big(1/\sqrt{T}\big)$ in deviation.

Lemma 7 is a direct consequence of a more general lower bound in [18], which proved that if the closure of the experts class is non-convex, and a single expert must be picked at the end ("proper" learning rule), then even under full information access during training the best achievable rate with high probability is $\mathcal{O}\big(1/\sqrt{T}\big)$.

Lemma 7. ($p = 1$) Consider the squared loss function. For $K = m = 2$ and $p = 1$, for any $T > 0$, and for any output $\hat{g}$ following the game protocol 1 for T training rounds, there exists a probability distribution for experts $\{F_1, F_2\}$ and target variable Y (all bounded by 1) such that, with probability at least 0.1,

$$\hat{R}_T(\hat{g}) - R^* \geq \frac{c_1}{\sqrt{T}},$$

where $c_1 > 0$ is an absolute constant.

The second result shows that the same lower bound holds for the bandit feedback ($m = 1$) setting, even if the learner is allowed to predict using a convex combination of all the experts at the end. To the best of our knowledge, this is the first lower bound for deviations in this setting.

Lemma 8. ($m = 1$) Consider the squared loss function. For $K = p = 2$, and $m = 1$, for any $T > 0$, for any training observation strategy and output $\hat{g}$ following the game protocol 1, there exists a probability distribution for experts $\{F_1, F_2\}$ and target variable Y (all bounded by 1) such that with probability at least 0.1,

$$\hat{R}_T(\hat{g}) - R^* \geq \frac{1}{2\sqrt{T}}.$$

8 Conclusion

We discussed the impact of restricted access to information in generalization error minimization with respect to the best expert. As many classical methods, such as progressive mixture rules are deviation suboptimal, we came with a new procedure achieving fast rates with high probability. We focused on the global budget setting, where a constraint on the total number of experts' prediction is made, and the local budget, where a limited number of experts' advice are shown per round. Moreover, we proved fast rates are impossible to achieve if the agent is allowed to see just one expert advice per round or choose just one expert for prediction.

An interesting future direction is allowing experts to learn from data during the process. In this case, the i.i.d. assumption on the loss sequence is dropped, which necessitates deriving a new concentration for the key quantities.

Broader Impact

We believe that a relevant discussion on the broader impact is not applicable to this paper.

References

[1] A. Agarwal, O. Dekel, and L. Xiao. Optimal algorithms for online convex optimization with multi-point bandit feedback. In *COLT*, pages 28–40, 2010.

[2] J.-Y. Audibert. Progressive mixture rules are deviation suboptimal. In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20, 2008.

[3] J.-Y. Audibert. Progressive mixture rules are deviation suboptimal / Supplemental "Proof of the optimality of the empirical star algorithm". In J. Platt, D. Koller, Y. Singer, and S. Roweis, editors, *Advances in Neural Information Processing Systems*, volume 20, 2008.

[4] J.-Y. Audibert. Fast learning rates in statistical inference through aggregation. *The Annals of Statistics*, 37(4):1591–1646, 2009.

[5] J.-Y. Audibert and S. Bubeck. Regret bounds and minimax policies under partial monitoring. *The Journal of Machine Learning Research*, 11:2785–2836, 2010.

[6] S. Ben-David and E. Dichterman. Learning with restricted focus of attention. *Journal of Computer and System Sciences*, 56(3):277–298, 1998.

[7] O. Catoni. A mixture approach to universal model selection. Technical Report LMENS-97-30, Ecole Normale Supérieure, 1997.

[8] N. Cesa-Bianchi, A. Conconi, and C. Gentile. On the generalization ability of on-line learning algorithms. *IEEE Transactions on Information Theory*, 50(9):2050–2057, 2004.

[9] N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

[10] N. Cesa-Bianchi, S. Shalev-Shwartz, and O. Shamir. Efficient learning with partially observed attributes. *Journal of Machine Learning Research*, 12(10), 2011.

[11] S. Gerchinovitz and T. Lattimore. Refined lower bounds for adversarial bandits. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 1198–1206, 2016.

[12] S. Guha and K. Munagala. Approximation algorithms for budgeted learning problems. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 104–113, 2007.

[13] D. Haussler, J. Kivinen, and M. K. Warmuth. Sequential prediction of individual sequences under general loss functions. *IEEE Transactions on Information Theory*, 44(5):1906–1925, 1998.

[14] E. Hazan and T. Koren. Optimal algorithms for ridge and lasso regression with partially observed attributes. *arXiv preprint arXiv:1108.4559*, 2011.

[15] E. Kaufmann, O. Cappé, and A. Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.

[16] G. Lecué. Suboptimality of penalized empirical risk minimization in classification. In *International Conference on Computational Learning Theory*, pages 142–156. Springer, 2007.

[17] G. Lecué and S. Mendelson. Aggregation via empirical risk minimization. *Probability theory and related fields*, 145(3-4):591–613, 2009.

[18] W. S. Lee, P. Bartlett, and R. Williamson. The importance of convexity in learning with squared loss. *IEEE Transactions on Information Theory*, 44(5):1974–1980, 1998.

[19] O. Madani, D. J. Lizotte, and R. Greiner. Active model selection. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 357–365, 2004.

[20] S. Mannor and O. Shamir. From bandits to experts: on the value of side-observations. In *Proceedings of the 24th International Conference on Neural Information Processing Systems*, pages 684–692, 2011.

[21] Y. Seldin, P. Bartlett, K. Crammer, and Y. Abbasi-Yadkori. Prediction with limited advice and multiarmed bandits with paid observations. In *International Conference on Machine Learning*, pages 280–287. PMLR, 2014.

[22] O. Shamir. An optimal algorithm for bandit and zero-order convex optimization with two-point feedback. *The Journal of Machine Learning Research*, 18(1):1703–1713, 2017.

[23] T. S. Thune and Y. Seldin. Adaptation to easy data in prediction with limited advice. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 2914–2923, 2018.

[24] Y. Xia, T. Qin, W. Ma, N. Yu, and T.-Y. Liu. Budgeted multi-armed bandits with multiple plays. In *IJCAI*, pages 2210–2216, 2016.

[25] Y. Yang and A. Barron. Information-theoretic determination of minimax rates ofconvergence. *The Annals of Statistics*, 27(5):1564 – 1599, 1999.

[26] D. Zhou and C. Tomlin. Budget-constrained multi-armed bandits with multiple plays. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] The assumptions and settings considered are stated in section 3
   (b) Did you include complete proofs of all theoretical results? [Yes] We included some proof sketches in the main paper, all the remaining and complete proofs are presented in the supplementary material
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [N/A]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]

5. If you used crowdsourcing or conducted research with human subjects...

   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]

   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]

   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]