



Online Covering with Multiple Experts

Enikő Kevi, Nguyễn Kim Thắng

► To cite this version:

| Enikő Kevi, Nguyễn Kim Thắng. Online Covering with Multiple Experts. 2023. hal-04297794

HAL Id: hal-04297794

<https://hal.science/hal-04297794>

Preprint submitted on 22 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NonCommercial 4.0 International License

Online Covering with Multiple Experts

Kevi Enikő and Nguyễn Kim Thắng

LIG, University Grenoble-Alpes, France

November 22, 2023

Abstract

Designing online algorithms with machine learning predictions is a recent technique beyond the worst-case paradigm for various practically relevant online problems (scheduling, caching, clustering, ski rental, etc.). While most previous learning-augmented algorithm approaches focus on integrating the predictions of a single oracle, we study the design of online algorithms with *multiple* experts. To go beyond the popular benchmark of a static best expert in hindsight, we propose a new *dynamic* benchmark (linear combinations of predictions that change over time). We present a competitive algorithm in the new dynamic benchmark with a performance guarantee of $O(\log K)$, where K is the number of experts, for 0–1 online optimization problems. Furthermore, our multiple-expert approach provides a new perspective on how to combine in an online manner several online algorithms - a long-standing central subject in the online algorithm research community.

1 Introduction

The domain of algorithms with predictions [24] - or learning augmented algorithms - emerged recently and grown immensely at the intersection of (discrete) algorithm design and machine learning (ML). Combining ML techniques with traditional algorithm design methods enables online algorithms to benefit from predictions that can infer future information from patterns in past data. Online algorithms with predictions can obtain performance guarantees beyond the worst-case analysis and provide fine-tuned solutions to various problems. In the literature, many significant problems have new learning-augmented results, for example, scheduling [20, 23], caching (paging) [21, 25, 4], ski rental [14, 19, 3], counting sketches [16], bloom filters [18, 22], and metric task systems [5].

Even though predictions provide a glimpse of the future, there is no mathematical guarantee of their accuracy. Adjusting the algorithm’s trust in the predictions is a significant challenge since online algorithms must make irrevocable decisions at each time step. Ideally, if the predictions are accurate, the algorithm should perform well compared to the offline setting. In contrast, if the predictions are misleading, the algorithm should maintain a competitive solution, similar to the online setting where no predictive information is available. In other words, online algorithms with predictions are expected to bring the best of both worlds: mathematical performance guarantees of classical algorithms and good future prediction capabilities of machine learning methods.

Predictions can come from multiple sources (heuristics, oracles, randomized methods, etc.), but we ignore their nature and call all of them *experts*. An algorithm’s consistency with the experts’ suggestions is typically measured by comparing the algorithm’s result with the solution of the *best* expert. A representative example is the popular notion of regret in online learning, which fueled the development of many powerful algorithms and techniques.

A natural research question is whether it is possible to design competitive algorithms with mathematical performance guarantees with a stronger benchmark than the best expert. Comparing an algorithm with a stronger benchmark could provide deeper insights into the learning process and give better ways of exploiting the experts' predictions.

Taking a broader view, we can study whether combining predictions of several experts is similar to combining multiple online algorithms and whether we can expect to achieve better solutions with the combination. Assuming that we do not know in advance which of the given algorithms would perform best on the upcoming requests, can we combine the algorithms in some generic way to obtain a competitive online strategy? This has been a long-standing question in the community of online algorithms [6, 8]. To find an answer, it is crucial to understand to what extent such an online strategy can benefit from the input of multiple algorithms and what is a suitable benchmark to evaluate its performance.

While in a completely general setting such an online strategy and a corresponding benchmark may not exist, in our paper we propose an algorithm for online linear problems with covering constraints that is competitive with a new benchmark (informally the *best linear combination* of the experts). Therefore, our paper partially addresses the question we raised in the previous paragraph.

1.1 Model and Problem

Covering problem with experts. We have n resources and each resource i has a cost per unit c_i that we know in advance ($1 \leq i \leq n$). Let x_i be a non-negative variable representing the amount chosen from resource i . The total cost of a solution $(x_i)_{i=1}^n$ is $\sum_{i=1}^n c_i x_i$. The problem includes K experts and the problem's (covering type) constraints are revealed online (one by one). At each time $t \geq 1$, we receive a covering constraint $\sum_{i=1}^n a_i^t x_i \geq 1$ (where $a_i^t \geq 0$) and each expert k (where $1 \leq k \leq K$) provides a solution $(s_{i,k}^t)_{i=1}^n$. An algorithm can observe the experts' solutions and afterwards it must update its own solution (denoted as $(x_i^t)_{i=1}^n$) to satisfy the new constraint, while maintaining the satisfaction of the previous ones. This algorithm must update its solution in the sense of online algorithms, so it cannot modify the previously made decisions. Formally, $x_i^t \geq x_i^{t-1} \forall i, t$. Our goal is to design such an algorithm and minimize $\sum_{i=1}^n c_i x_i^T$ subject to all online covering constraints t , where $1 \leq t \leq T$. The value T is the last time a constraint is released, and it is not known by the algorithm.

Experts. In our model, the experts' predictions are also online solutions. In other words, the experts' solutions fulfill the following properties:

1. for every expert k and for every time t the solution $(s_{i,k}^t)_{i=1}^n$ is feasible, therefore, every constraint t' where $1 \leq t' \leq t$ is satisfied;
2. for every expert k and for every time t and for every resource i , the previous expert solutions are irrevocable, therefore $s_{i,k}^t \geq s_{i,k}^{t'}$ for all $t' \leq t$.

These properties can be verified online. If some experts do not satisfy them, we simply ignore those experts both in the decision-making and in the benchmark. A crucial remark: we do *not* assume that the experts' solutions must be tight at each constraint t , meaning that $\sum_{i=1}^n a_i^t s_{i,k}^t = 1 \forall t, k$. This assumption is unrealistic and cannot be maintained in an online manner (see the discussion in Appendix A). Besides, assuming tight constraint satisfaction would simplify the problem, while intuitively, the difficulty of designing competitive algorithms comes from the lack of obvious ways to distinguish good expert solutions from (probably many) non-efficient/misleading ones.

$$\begin{aligned}
& \min \sum_{i=1}^n c_i x_i^T = \min \sum_{i=1}^n c_i \sum_{t=1}^T (x_i^t - x_i^{t-1}) \\
& \text{subject to} \quad \sum_{k=1}^K w_k^t = 1 \quad \forall t \\
& \quad \quad \quad x_i^t \geq \sum_{k=1}^K w_k^t s_{i,k}^t \quad \forall i, t \\
& \quad \quad \quad x_i^t \geq x_i^{t-1} \quad \forall i, t \\
& \quad \quad \quad w_k^t \geq 0 \quad \forall t, k
\end{aligned}$$

where $1 \leq t \leq T$ and $1 \leq i \leq n$.

Figure 1: Formulation of the LIN-COMB benchmark

Benchmark. We consider a dynamic benchmark that intuitively captures the *best linear combination* of all experts' solutions *over time*. Informally, at any online time step, the benchmark can take a linear combination of the experts' solutions. The linear combination can be changed over time, and it can be different from previous combinations. However, the benchmark's decisions are also online, so it cannot decrease the value of the decision variables (x_i). We refer to our benchmark with the name LIN-COMB from now on.

The LIN-COMB benchmark's formal description is a linear program, visible on Figure 1. Let $w_k^t \geq 0$ be the weight assigned by the LIN-COMB benchmark to expert k (where $1 \leq k \leq K$) at time t . Since we consider a linear combination, the constraint $\sum_{k=1}^K w_k^t = 1$ must hold. The solution of LIN-COMB at time t is ideally $x_i^t = \sum_{k=1}^K w_k^t s_{i,k}^t$, however, x_i^t must be larger than x_i^{t-1} . Therefore, we set $x_i^t = \max\{\sum_{k=1}^K w_k^t s_{i,k}^t, x_i^{t-1}\}$. In other words, given the chosen weights, if $\sum_{k=1}^K w_k^t s_{i,k}^t < x_i^{t-1}$ then $x_i^t \leftarrow x_i^{t-1}$, otherwise $x_i^t \leftarrow \sum_{k=1}^K w_k^t s_{i,k}^t$.

Since every expert's solution is feasible by our assumptions, at each time t and for all resource i (where $1 \leq i \leq n$), the constructed solution $x_i^t \geq \sum_{k=1}^K w_k^t s_{i,k}^t$ constitutes a feasible solution to the covering constraints of the original covering problem. Formally, for every constraint t' with $t' \leq t$,

$$\sum_{i=1}^n a_i^{t'} x_i^t \geq \sum_{i=1}^n a_i^{t'} \left(\sum_{k=1}^K w_k^t s_{i,k}^t \right) = \sum_{k=1}^K w_k^t \left(\sum_{i=1}^n a_i^{t'} s_{i,k}^t \right) \geq \sum_{k=1}^K w_k^t \geq 1$$

where the second inequality holds due to the feasibility of the experts' solutions. We highlight that the best-expert benchmark is included in LIN-COMB. By setting $w_{k^*}^t = 1$ for all t , where $1 \leq t \leq T$, and $w_k^t = 0$ for all $k \neq k^*$ where k^* we get the best expert in hindsight (so $x_i^t = s_{i,k^*}^t$ for all i, t).

1.2 Our approach and contribution

To design competitive algorithms with the new benchmark, we consider a primal-dual approach. First, we relax the linear program formulation of LIN-COMB (visible on Figure 1), which serves as a lower bound. Then, we take the dual of the relaxation, which is a lower bound on the relaxation. Therefore, following the chain of lower bounds, the dual problem is a lower bound on the LIN-COMB benchmark.

At every time step during the execution, our algorithm constructs decisions based on the solution of a convex program. Our approach is inspired by the convex regularization method of [11]. The objective of the convex program is a shifted entropy function. These functions have been widely used, in particular in the recent breakthrough related to k -server [10, 12] and metrical tasks problems [9], in which the entropy functions are shifted by constant parameters. A *novel* point in our approach is that the entropy function is shifted by the average of the experts' solutions. Moreover, regarding the constraints of the convex program, instead of using the experts' solutions directly, we define auxiliary solutions that guarantee tight constraint satisfaction and use them in the constraints. Intuitively, this eliminates/modifies the malicious experts' solutions.

Let ρ be the maximum ratio between the experts' solutions on the resources. Formally,

$$\rho := \max_i \max_{t', t''} \left\{ \frac{\sum_{k=1}^K s_{i,k}^{t'}}{\sum_{k=1}^K s_{i,k}^{t''}} \right\} \text{ s.t. } \sum_{k=1}^K s_{i,k}^{t''} > 0.$$

Informally, ρ represents the discrepancy across the experts' predictions. Our main result is an algorithm that has an objective cost at most $O(\ln(K\rho))$ times the cost of the LIN-COMB benchmark. In particular, for 0-1 optimization problems, where the experts provide integer (deterministic or randomized) solutions, our algorithm is $O(\ln K)$ -competitive with LIN-COMB. An interesting feature of our algorithm is its resilience to the fluctuation of the quality of predictions (as discussed in the section below and illustrated in the experiments).

1.3 Related work and discussions

Much of the research focusing on surpassing worst-case performance guarantees is motivated by the spectacular advances of machine learning (ML). Specifically, ML methods can detect patterns among the arriving input requests and provide valuable insights for the online algorithms regarding future requests. [21] introduced a general framework to integrate ML predictions into classical algorithm designs to surpass the worst-case performance limit. As a result, many practically relevant online problems were revisited to enhance existing classical algorithms with ML predictions (see the aforementioned [20, 23, 21, 25, 4, 14, 19, 3, 16, 18, 22, 5]).

On a high-level view, we aim to design algorithms that are robust (competitive) to the offline optimal solution and also consistent with the expert's predictions. Ideally, the performance of the designed algorithm should surpass previous bounds whenever the predictions are reliable (low errors). However, most learning-augmented algorithms suffer when the error rates are neither very low nor very high, resulting in prediction confidence that is neither very low nor very high. Figure 2 provides a general picture of the performance of an algorithm with predictions, which is representative for many problems (for example, [7, 17]). In the figure, η indicates the confidence in the predictions (or equivalently the error rate of predictions). The learning-augmented algorithm's performance bound is the maximum value of the green and orange curves (gray shaded area on the figure). We can observe that when $0.4 \leq \eta \leq 0.9$, the algorithm's performance guarantee is worse than the classical worst-case guarantee (that can be achieved by simply ignoring all predictions). Intuitively, in the case of neither very low nor very high confidence in the predictions, the algorithm




Figure 2: Robustness-Consistency

has a hard time deciding if it should follow the predictions or the best-known standard algorithm in the worst-case paradigm. It naturally raises the question of whether one can surely guarantee to achieve at least a constant factor of the worst-case guarantee (where the constant is as close to 1 as possible), assuring the resilience of the output solutions despite the quality of the predictions. Our algorithm, together with the new benchmark, provides an answer to this question.

The paper of [2] is closely related to ours and studies the design of algorithms with multiple experts. They consider a DYNAMIC benchmark that is intuitively the minimum cost solution that is supported by at least one expert solution at each step. Formally:

$$\text{DYNAMIC} = \min_{\hat{\mathbf{x}} \in \hat{X}} \sum_{i=1}^n c_i \hat{x}_i, \text{ where}$$

$$\hat{X} = \{\hat{\mathbf{x}} : \forall i \in [n], \forall t \in [T], \exists k \in [K] \text{ such that } s_{i,k}^t \leq \hat{x}_i\}$$

Our benchmark, LIN-COMB, is included in DYNAMIC, since every solution x_i^t in LIN-COMB satisfies:

$$x_i^t \geq \sum_k s_{i,k}^t w_k^t \geq \min_k \{s_{i,k}^t\}$$

therefore, for any i and t , there exists k such that $x_i^t \geq s_{i,k}^t$. However, the inverse is not true: a solution $\hat{\mathbf{x}}^t \in \hat{X}$ in DYNAMIC is not necessarily a linear combination of the experts' solutions. The DYNAMIC benchmark in [2] relied on the assumption that at every time step the experts' solutions are tight. This assumption is unrealistic and impossible to maintain in online solutions (see Appendix A). [2] claimed an $O(\log K)$ -competitive algorithm in the DYNAMIC benchmark. Unfortunately, this is *incorrect*; we show an example in Appendix B in which their algorithm's performance guarantee is unbounded in the DYNAMIC benchmark.

Integrating multiple predictions into the online algorithm design was a topic of other papers as well. As an example, [15] studied the ski rental problem with multiple predictions. The authors defined a consistency metric, which compares the performance of their algorithm to the optimal solution, given that at least one prediction (among the k predictions) is optimal. [1] also considered multiple predictions in the online facility location problem. They compared the performance of their algorithm to the best possible solution obtained on the union of the suggestions. Recently, [13] studied the use of multiple predictors for several problems such as matching, load balancing, and non-clairvoyant scheduling. They provided algorithms competitive to the best predictor for such problems. An important remark: all the above benchmarks are captured within LIN-COMB.

Furthermore, [5] proposed an algorithm with multiple experts for the metrical task system problem. Their benchmark allows switching from one expert to another at each time step, but it does not allow combinations of experts or any solution not suggested by one of the experts. In our LIN-COMB benchmark, the linear combinations that evolve over time could result in a solution that is not suggested by one of the experts and potentially they can be much more efficient. In [5] there is a cost for state transitions, which is appropriate for their problems, but in many other problems, the smooth transition with additional costs from previous decisions to new ones is not allowed (past decisions are immutable). Therefore, the results of [5] are not applicable to our setting.

Combining online algorithms into a new algorithm to achieve better results than the individual input algorithms has been a long-standing online algorithm design question [6, 8]. Its intrinsic difficulty is similar to the issue we mentioned earlier: when the performance of the given input algorithms (or heuristics) is unclear (especially in the online setting), it is challenging to create a combination that can surpass the performance of the included algorithms. Following the current development of online algorithm design techniques with multiple predictions, this subject has been renewed with different machine learning approaches. Our paper contributes to this line of research.

2 Online covering with multiple experts

Our proposed algorithm solves online covering problems by creating linear combinations of the solutions proposed by K experts in an online manner. Recall that we evaluate the performance of our algorithm with the **LIN-COMB** benchmark (formalized on Figure 1), which consists of the best linear combination of the experts' solution at each step.

Since our **LIN-COMB** benchmark is a linear combination of the experts' solutions, the equality $\sum_{k=1}^K w_k^t = 1$ must hold, where $w_k^t \geq 0$ is the weight assigned to expert k (where $1 \leq k \leq K$) at time t . In the following, we formulate a relaxed version of the **LIN-COMB** formulation, where $\sum_{k=1}^K w_k^t \geq 1$. Additionally, the relaxed formulation enables us to avoid the (online) hard constraint requiring $w_k^t s_{i,k}^t \geq w_k^{t-1} s_{i,k}^{t-1}$ to hold, and instead, we introduce a new variable, y_i^t , to represent the increase of x_i^t compared to x_i^{t-1} . When $w_k^t s_{i,k}^t < w_k^{t-1} s_{i,k}^{t-1}$ during the execution, we set the contribution of i at time t to be 0, and therefore, $y_i^t = 0$. The relaxed formulation is visible on Figure 3.

Due to the relaxed constraint, the optimal solution of the relaxed linear program is a lower bound of our **LIN-COMB** benchmark. The dual of the relaxation is displayed on Figure 4.

$$\begin{array}{ll}
 (\alpha^t) & \min \sum_{t=1}^T \sum_{i=1}^n c_i y_i^t \\
 & \sum_{k=1}^K w_k^t \geq 1 \quad \forall t \\
 (\beta_i^t) & \sum_{k=1}^K \left(w_k^t s_{i,k}^t - w_k^{t-1} s_{i,k}^{t-1} \right) \leq y_i^t \quad \forall i, t \\
 & w_k^t, y_i^t \geq 0 \quad \forall i, t, k
 \end{array}$$

Figure 3: Formulation of the relaxation of the **LIN-COMB** benchmark

$$\begin{array}{ll}
 & \max \sum_{t=1}^T \alpha^t \\
 (x_k^t) & \alpha^t + \sum_{i=1}^n s_{i,k}^t (\beta_i^{t+1} - \beta_i^t) \leq 0 \quad \forall k, t \\
 (y_i^t) & \beta_i^t \leq c_i \quad \forall i, t \\
 & \alpha_i^t, \beta_i^t \geq 0 \quad \forall i, t
 \end{array}$$

Figure 4: Dual formulation of the relaxation of the **LIN-COMB** benchmark

According to the theorem of weak duality, any feasible solution of the dual program lower bounds any feasible solution of the primal program, and therefore, any feasible dual solution also lower bounds our **LIN-COMB** benchmark. Following the chain of lower bounds, our approach to design a competitive algorithm is as follows. At every time step t , we build solutions for all x_i^t

together with the solutions for the dual problem (α^t, β_i^t) . Then, we bound the cost of the algorithm to that of the dual. It is important to emphasize that the designed solution for every x_i^t must be feasible to the covering constraints, but it may *not necessarily* be a linear combination of the experts' solutions.

2.1 Competitive Algorithm

Preprocessing. Recall that by our assumptions, the experts' solutions are always feasible and non-decreasing. At the arrival of the t^{th} constraint, expert k (where $1 \leq k \leq K$) provides a feasible solution $s_k^t = (s_{i,k}^t)_{i=1}^n$, such that $s_{i,k}^t \geq s_{i,k}^{t'}$ for all $t' \leq t$ and all i where $1 \leq i \leq n$. These assumptions do not exclude the possibility for the experts to provide malicious solutions that instruct the algorithm to use an unnecessarily large amount of resources. Note that contrary to the assumption in [2], we can *not* expect the experts' solutions to be always tight. (In Appendix A we show an example that tight solutions cannot be maintained in an online manner.)

To circumvent this issue, we preprocess the experts' solutions at each iteration. During the preprocessing, every solution s_k^t is scaled down to make it as tight as possible on the t^{th} constraint, while always maintaining $s_{i,k}^t \geq s_{i,k}^{t-1}$ for all i . Additionally, after the down-scaling, we create an auxiliary solution \hat{s}_k^t that is tight for the t^{th} constraint. This solution is useful for our algorithm, and we create it with the following procedure.

After the down-scaling, do the following for each expert k .

1. If $(s_{i,k}^t)_{i=1}^n$ is tight on the t^{th} constraint, then set $\hat{s}_{i,k}^t \leftarrow s_{i,k}^t$ for every i .
2. Let $\hat{s}_{i,k}^{t-1}$ be the auxiliary solution of expert k at time $t-1$, meaning that, $\sum_{i=1}^n a_i^{t-1} \hat{s}_{i,k}^{t-1} = 1$.

Given $I := \{i : s_{i,k}^t > \hat{s}_{i,k}^{t-1} \cdot \frac{a_i^{t-1}}{a_i^t}\}$, we set $\hat{s}_{i,k}^t \leftarrow s_{i,k}^t$ if $i \notin I$ and set $\hat{s}_{i,k}^t$ to be some value in $[\hat{s}_{i,k}^{t-1} \cdot \frac{a_i^{t-1}}{a_i^t}, s_{i,k}^t]$ if $i \in I$, s.t. the solution $\hat{s}_{i,k}^t$ becomes tight on the t^{th} constraint.

Lemma 1 *Following the preprocessing procedure, we can always obtain the solutions $\hat{s}_{i,k}^t$ such that $\hat{s}_{i,k}^t \leq s_{i,k}^t$ and $\sum_{i=1}^n a_i^t \hat{s}_{i,k}^t = 1$.*

Proof Let us fix an expert k . We prove the lemma by induction. At time step $t = 1$, one can always scale down the solution $s_{i,k}^1 \geq 0$ such that the first constraint becomes tight. Assume that the lemma holds until $t-1$, $\sum_{i=1}^n a_i^{t-1} \hat{s}_{i,k}^{t-1} = 1$ and $\hat{s}_{i,k}^{t-1} \leq s_{i,k}^{t-1}$. Consider time t . If after scaling down (at the first step in the procedure) the t^{th} constraint becomes tight, then we are done. Otherwise, we have

$$\begin{aligned} 1 &< \sum_{i=1}^n a_i^t s_{i,k}^t = \sum_{i \in I} a_i^t s_{i,k}^t + \sum_{i \notin I} a_i^t s_{i,k}^t, \\ 1 &= \sum_{i=1}^n a_i^{t-1} \hat{s}_{i,k}^{t-1} = \sum_{i=1}^n a_i^t \left(\hat{s}_{i,k}^{t-1} \cdot \frac{a_i^{t-1}}{a_i^t} \right) \\ &\geq \sum_{i \in I} a_i^t \left(\hat{s}_{i,k}^{t-1} \cdot \frac{a_i^{t-1}}{a_i^t} \right) + \sum_{i \notin I} a_i^t s_{i,k}^t \end{aligned}$$

Hence, there exists $\hat{s}_{i,k}^t \in [\hat{s}_{i,k}^{t-1} \cdot \frac{a_i^{t-1}}{a_i^t}, s_{i,k}^t]$ for every i , where $1 \leq i \leq n$, such that $\sum_{i=1}^n a_i^t \hat{s}_{i,k}^t = 1$. \square

Algorithm. At the arrival of the t^{th} constraint,

1. solve the following convex program and set w^t to be the obtained optimal solution

$$\begin{aligned}
& \min_w \left\{ \sum_{i=1}^n c_i \left[\left(\sum_{k=1}^K s_{i,k}^t w_{i,k} + \delta_i^t \right) \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k} + \delta_i^t}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right) - \sum_{k=1}^K s_{i,k}^t w_{i,k} \right] \right\} \\
(\gamma^t) \quad & \sum_{i=1}^n a_i^t \left(\sum_{k=1}^K \hat{s}_{i,k}^t w_{i,k} \right) \geq 1 \quad \forall t \\
(\lambda_i^t) \quad & \sum_{k=1}^K w_{i,k} \geq 1 \quad \forall i \\
(\mu_i^t) \quad & \sum_{k=1}^K s_{i,k}^t w_{i,k} \geq 0 \quad \forall i, t
\end{aligned}$$

where $\delta_i^t = \frac{1}{K} \sum_k s_{i,k}^t$. Note that in this program, we use the auxiliary solution $\hat{s}_{i,k}^t$ in the first constraint. For every i where $s_{i,k}^t = 0$ for all k , the term related to i is not included in the objective function of the convex program. (We can set $w_{i,k} = 0$ for all k beforehand.)

2. For all i if $\sum_{k=1}^K w_{i,k}^t s_{i,k}^t > x_i^{t-1}$ then set $x_i^t \leftarrow \sum_{k=1}^K w_{i,k}^t s_{i,k}^t$; otherwise set $x_i^t \leftarrow x_i^{t-1}$.

Note: To avoid the possible division by 0 in the denominator of the objective function's logarithm, we can use a dummy expert, who sets each variable to some small value and then follows the greedy heuristic to solve the problem at each arriving constraint. The presence of this expert only changes the competitive ratio to $O(\log K + 1)$. Additionally, upon the arrival of the first constraint, we treat the denominator as 1.

2.2 Analysis

As w^t is the optimal solution of the convex program and $(\gamma^t, \lambda_i^t, \mu_i^t)$ is the optimal solution of its dual, the following Karush-Kuhn-Tucker (KKT) and complementary slackness conditions hold.

$$\begin{aligned}
& \left[\sum_{i=1}^n a_i^t \left(\sum_k \hat{s}_{i,k}^t w_{i,k}^t \right) - 1 \right] \gamma^t = 0 \quad \forall t \\
& \left[\sum_{k=1}^K w_{i,k}^t - 1 \right] \lambda_i^t = 0 \quad \forall i, t \\
& \left[\sum_{k=1}^K s_{i,k}^t w_{i,k}^t \right] \mu_i^t = 0 \quad \forall i, t \\
& c_i s_{i,k}^t \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \delta_i^t}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right) - a_i^t \hat{s}_{i,k}^t \gamma^t - \lambda_i^t - s_{i,k}^t \mu_i^t = 0 \quad \forall i, k, t \\
& \gamma^t, \lambda_i^t, \mu_i^t \geq 0 \quad \forall i, t
\end{aligned}$$

Moreover, if $\sum_{k=1}^K w_{i,k}^t s_{i,k}^t > 0$, meaning that $\mu_i^t = 0$, then

$$c_i s_{i,k}^t \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \delta_i^t}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right) - a_i^t \hat{s}_{i,k}^t \gamma^t - \lambda_i^t = 0 \quad (1)$$

Dual variables and feasibility. We set the dual variables of the linear program relaxation of our LIN-COMB benchmark based on the dual variables of the convex program used inside the algorithm.

$$\alpha^t = \frac{1}{\ln(K\rho)} \left(\gamma^t + \sum_i \lambda_i^t \right),$$

$$\beta_i^t = \frac{1}{\ln(K\rho)} c_i \ln \left(\frac{(1 + 1/K) \cdot \max_{t'} \sum_{k=1}^K s_{i,k}^{t'}}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right)$$

where recall that $\rho = \max_{i,t',t''} \left\{ \frac{\sum_{k=1}^K s_{i,k}^{t'}}{\sum_{k=1}^K s_{i,k}^{t''}} : \sum_{k=1}^K s_{i,k}^{t''} > 0 \right\}$.

Lemma 2 *The x_i^t solutions set by the algorithm for the original covering problem and the dual variables (α^t, β_i^t) of the LIN-COMB benchmark's linear program relaxation are feasible.*

Proof We first prove that the x_i^t variables satisfy the covering constraints by induction. At time 0, no constraint has been released yet, and every variable is set to 0. This all-zero solution is feasible. Let us assume that the algorithm provides feasible solutions up to time $t - 1$. At time t , the algorithm maintains the inequality $x_i^t \geq x_i^{t-1}$, so all constraints t' where $t' < t$ are satisfied. Besides, x_i^t is always at least $\sum_k w_{i,k}^t s_{i,k}^t$, which is larger than $\sum_k w_{i,k}^t \hat{s}_{i,k}^t$ since $s_{i,k}^t \geq \hat{s}_{i,k}^t$ for all i, k by the preprocessing step. Hence, the constraint t is also satisfied, formally,

$$\sum_{i=1}^n a_i^t x_i^t \geq \sum_{i=1}^n a_i^t \left(\sum_k \hat{s}_{i,k}^t w_{i,k}^t \right) \geq 1.$$

In the remaining part of the proof, we show the feasibility of α^t and every β_i^t . Since $\gamma^t \geq 0$ and $\lambda_i^t \geq 0$ for all i and t , we get that $\alpha^t \geq 0$. In the definition of β_i^t , the nominator of the logarithm term is always larger than the denominator, and it is smaller than $K\rho$ times the denominator. Consequently, $0 \leq \beta_i^t \leq c_i$. Furthermore,

$$\beta_i^{t+1} - \beta_i^t = -\frac{1}{\ln(K\rho)} c_i \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \delta_i^t}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right).$$

Since $\sum_i a_i^t \hat{s}_{i,k}^t = 1$, using the KKT conditions, we get:

$$\begin{aligned} \alpha^t + \sum_{i=1}^n s_{ik}^t (\beta_i^{t+1} - \beta_i^t) &= \frac{1}{\ln(K\rho)} \left(\gamma^t + \sum_i \lambda_i^t \right) - \frac{1}{\ln(K\rho)} \sum_{i=1}^n s_{i,k}^t c_i \ln \left(\frac{\sum_{k=1}^K s_{ik}^t w_{i,k}^t + \delta_i^t}{\sum_{k=1}^K s_{ik}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right) \\ &= \frac{1}{\ln(K\rho)} \left[\gamma^t + \sum_{i=1}^n \lambda_i^t - \sum_{i=1}^n (a_i^t \hat{s}_{i,k}^t \gamma^t + \lambda_i^t + s_{i,k}^t \mu_i^t) \right] \\ &\leq 0 \end{aligned}$$

□

Theorem 1 *The algorithm's cost is at most $O(\ln(K\rho))$ -competitive in the LIN-COMB benchmark.*

Proof Lemma 2 proved that our algorithm creates feasible solutions for the dual problem of the LIN-COMB benchmark relaxation and for the original covering problem. We show that the algorithm's solution increases the primal objective value of the original covering problem by at most $O(\ln(K\rho))$ times the value of the dual solution, which serves as the lower bound on the LIN-COMB benchmark - the best linear combination of the experts' solutions.

$$\begin{aligned} \sum_{i=1}^n c_i(x_i^t - x_i^{t-1}) &= \sum_{i: x_i^t > x_i^{t-1}} c_i(x_i^t - x_i^{t-1}) \\ &\leq \sum_{i: x_i^t > x_i^{t-1}} c_i(x_i^t + \delta_i^t) \ln \frac{x_i^{t-1} + \delta_i^t}{x_i^{t-1} + \delta_i^t} \end{aligned} \quad (2)$$

$$\leq \sum_{i: x_i^t > x_i^{t-1}} c_i(x_i^t + \delta_i^t) \ln \frac{x_i^{t-1} + \delta_i^t}{x_i^{t-1} + \delta_i^{t-1}} \quad (3)$$

$$= \sum_{i: x_i^t > x_i^{t-1}} c_i \left[\left(\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \frac{1}{K} \sum_{k=1}^K s_{i,k}^t \right) \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \delta_i^t}{x_i^{t-1} + \delta_i^{t-1}} \right) \right] \quad (4)$$

$$\leq \sum_{i: x_i^t > x_i^{t-1}} c_i \left[\left(\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \frac{1}{K} \sum_{k=1}^K s_{i,k}^t \right) \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \delta_i^t}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right) \right] \quad (5)$$

$$= \sum_{i: x_i^t > x_i^{t-1}} \sum_{k=1}^K (w_{i,k}^t + 1/K) c_i s_{i,k}^t \ln \left(\frac{\sum_{k=1}^K s_{i,k}^t w_{i,k}^t + \delta_i^t}{\sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1} + \delta_i^{t-1}} \right)$$

$$= \sum_{i: x_i^t > x_i^{t-1}} \sum_{k=1}^K (w_{i,k}^t + 1/K) \left(a_i^t \hat{s}_{i,k}^t \gamma^t + \lambda_i^t \right) \quad (6)$$

$$\begin{aligned} &\leq \sum_{i=1}^n \sum_{k=1}^K (w_{i,k}^t + 1/K) \left(a_i^t \hat{s}_{i,k}^t \gamma^t + \lambda_i^t \right) \\ &= \sum_{i=1}^n a_i^t \left(\sum_{k=1}^K w_{i,k}^t \hat{s}_{i,k}^t \right) \gamma^t + \sum_{i=1}^n \left(\sum_{k=1}^K w_{i,k}^t \right) \lambda_i^t + \frac{1}{K} \sum_{k=1}^K \left(\sum_{i=1}^n a_i^t \hat{s}_{i,k}^t \right) \gamma^t + \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^n \lambda_i^t \\ &= 2\gamma^t + 2 \sum_{i=1}^n \lambda_i^t = \ln(K\rho) \alpha^t \end{aligned} \quad (7)$$

The above corresponding transformations hold since:

- (2) follows from the inequality $a - b \leq a \ln(a/b)$ for all $0 < b \leq a$;
- (3) holds since $\delta_i^t \geq \delta_i^{t-1}$ (because $s_{i,k}^t \geq s_{i,k}^{t-1}$ for all i, k, t);
- (4) is valid because $x_i^t > x_i^{t-1}$, so $x_i^t = \sum_{k=1}^K s_{i,k}^t w_{i,k}^t$;
- (5) is by the design of the algorithm: $x_i^{t-1} \geq \sum_{k=1}^K s_{i,k}^{t-1} w_{i,k}^{t-1}$;
- (6) since given that $x_i^t > x_i^{t-1} \geq 0$ (so $\sum_{k=1}^K s_{i,k}^t w_{i,k}^t = x_i^t > 0$), the KKT condition (1) applies;
- (7) is true due to the complementary slackness conditions and that $\sum_{i=1}^n a_i^t \hat{s}_{i,k}^t = 1$.

□

Corollary 1 *For 0-1 optimization problems in which experts provide integer (deterministic or randomized) solutions, the algorithm is $O(\ln K)$ -competitive in the LIN-COMB benchmark. Subsequently, there exists an algorithm such that its performance is $O(\ln K)$ -competitive in the LIN-COMB benchmark and is up to a constant factor to the best guarantee in the worst-case benchmark*

Proof

If the value of $s_{i,k}^t$ is in $\{0, 1\}$ for every i, k, t , then

$$\rho = \max_i \max_{t', t''} \left\{ \frac{\sum_{k=1}^K s_{i,k}^{t'}}{\sum_{k=1}^K s_{i,k}^{t''}} : \sum_{k=1}^K s_{i,k}^{t''} > 0 \right\} \leq \frac{K}{1}$$

Therefore, the competitive ratio of the main algorithm in the LIN-COMB benchmark is $O(\log K \rho) = O(\log K^2)$.

To obtain an algorithm that is competitive in both the LIN-COMB and the worst-case benchmarks, we proceed as follows (an illustration in Figure 5). We first apply the main algorithm on the K experts' predictions to obtain an online algorithm, named A . Algorithm A is $O(\ln K)$ -competitive in the LIN-COMB benchmark. Let B be the algorithm with the best worst-case guarantee. One applies the main algorithm one more on two algorithms, A and B . The final algorithm is $O(\ln 2)$ -competitive to both A and B . In other words, its performance is $O(\ln K)$ -competitive in the LIN-COMB benchmark and is up to a constant factor worse than the best guarantee in the worst-case benchmark.

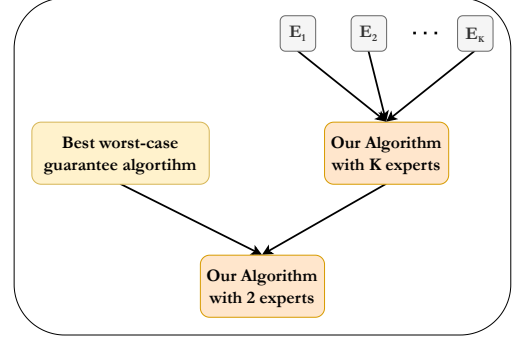


Figure 5: Structural overview of the algorithm's components. E_1, E_2, \dots, E_K correspond to the experts of the online problem. On the second layer, we integrate the best standard online algorithm with our algorithm.

□

By Corollary 1, given a 0-1 optimization problem, if there are K deterministic online algorithms, then we can design an algorithm that has a cost at most $O(\log K)$ times that of the best linear combination of those algorithms at any time. Similarly, if K given online algorithms are randomized (they output 0-1 solutions with probabilities), then our algorithm has an expected cost (randomization over the product of the distributions of those solutions) at most $O(\log K)$ times that of the best linear combination of those algorithms at any time. Many practical problems admit 0-1 solutions, for which our algorithm is of interest. Consider problems like network design, ski rental, TCP acknowledgement, facility location, etc. Given the fractional solutions constructed by our algorithm, we can apply existing online rounding schemes to obtain integral solutions for such problems.

3 Experiments

Implementation. The first step of our proposed algorithm is to solve a convex program. In the experiments, we approximate the optimal solution of this program using a vanilla Frank-Wolfe implementation. The linear minimization step within Frank-Wolfe is solved with the Gurobi optimizer.

Comparison. The best standard online algorithm for general covering problems without experts is the online multiplicative weight update (MWA) algorithm. In the experiments we compare our algorithm with the MWA algorithm. When a new constraint arrives in the online problem, the MWA algorithm increases each variable x_i in the constraint with a rate of $\frac{a_i^t}{c_i}(x_i + 1/n)$, where n is

the total number of variables. We also compare our results with the optimum offline solution (that knows the whole instance in advance) and the average solution of the experts.

Input. First, we evaluated the result of our algorithm on the pathological input of the MWA algorithm. This instance includes n variables and n constraints with uniform costs and coefficients. Each arriving constraint in this pathological example includes one less variable. While the optimal solution is 1, the worst-case guarantee of MWA is $O(\log n)$. For our algorithm we provided n experts, where $(n - 1)$ experts suggest an adversarial trivial solution to set all variables to 1, while 1 expert suggests the optimal offline solution. The result of this experiment is visible on Figure 6. An important highlight: our algorithm managed to identify the good expert among the majority of adversaries, obtaining a better objective value, than MWA. Second, we experimented with the counter example (in Appendix B) to show that the algorithm proposed by [2] has an incorrect performance proof. The result is visible on Figure 6. Finally, we generated some instances to observe the performance of our algorithm on non-specific inputs. The specification for the instance generation includes several parameters, which we detail on Figure 7.

Result. Multiplicative weight update is a simple and well-performing algorithm in practice. On its pathological worst-case example, our algorithm performs better, however on most instances the expert suggestions were significantly worse than MWA, which impacted the performance of our algorithm. To some extent, our algorithm can detect the good experts and is robust against even many adversaries. We think that given a real-life problem, it is possible to construct well-performing experts (for example using past data) and with a fine-tuned convex program solver our algorithm can be of interest for various use-cases.

Algo name	Worst-case for MWA	Counter-ex. for [2]	Inst. 1	Inst. 2	Inst. 3	Inst. 4
OPT Offline	1.0	1.0	1.3	1.5	10.6	31.3
MWA Online	2.9	2.3	2.0	1.7	28.1	63.7
Our Algo	2.2	4.4	2.0	1.7	26.7	61.7
Avg of experts	9.1	3.5	16.4	23.3	1897.47	96.9

Figure 6: Objective value of the experiment instances

Input Generation Parameters	Instance 1	Instance 2	Instance 3	Instance 4
Number of variables	10	10	44	30
Number of constraints	10	25	2	15
Min objective coefficient	1	10	1	1
Max objective coefficient	10	25	100	100
Min constraint coefficient	1	10	1	1
Max constraint coefficient	10	25	1	1
Min number of zero coefficient	0	1	11	5
Max number of zero coefficient	5	5	22	20
Number of perfect experts	1	0	0	2
Number of online experts	2	1	1	2
Number of random experts	1	1	11	0
Number of adversaries	1	1	0	0

Figure 7: Parameters of the generated experiment instances

4 Conclusion

We introduce a dynamic **LIN-COMB** benchmark in the setting of multiple expert predictions beyond the traditional static benchmark of the best expert in hindsight and give a competitive algorithm for the online covering problem in this benchmark. Our approach can provide valuable insights into the learning processes related to predictions, in particular, in aggregating information from predictions to improve the performance of existing algorithms, and how to combine online algorithms, an important subject in the online algorithm design community [6, 8]. The experiments support the fact that our algorithm can differentiate between good and adversarial experts to some extent.

An interesting open question is to design competitive algorithms in the **LIN-COMB** benchmark for different classes of problems, such as packing problems and problems with non-linear objectives.

References

- [1] Matteo Almanza, Flavio Chierichetti, Silvio Lattanzi, Alessandro Panconesi, and Giuseppe Re. Online facility location with multiple advice. In *Advances in Neural Information Processing Systems*, volume 34, pages 4661–4673, 2021.
- [2] Keerti Anand, Rong Ge, Amit Kumar, and Debmalya Panigrahi. Online algorithms with multiple predictions. In *Proc. 39th International Conference on Machine Learning*, 2022.
- [3] Spyros Angelopoulos, Christoph Dürr, Shendan Jin, Shahin Kamali, and Marc Renault. Online Computation with Untrusted Advice. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151, pages 52:1–52:15, 2020.
- [4] Antonios Antoniadis, Christian Coester, Marek Elias, Adam Polak, and Bertrand Simon. On-line metric algorithms with untrusted predictions. In *International Conference on Machine Learning*, pages 345–355, 2020.
- [5] Antonios Antoniadis, Christian Coester, Marek Eliáš, Adam Polak, and Bertrand Simon. Mixing predictions for online metric algorithms, 2023. [arXiv:2304.01781](https://arxiv.org/abs/2304.01781).
- [6] Yossi Azar, Andrei Z Broder, and Mark S Manasse. On-line choice of on-line algorithms. In *Symposium of Discrete Algorithms (SODA)*, pages 432–440, 1993.
- [7] Etienne Bamas, Andreas Maggiori, and Ola Svensson. The primal-dual method for learning augmented algorithms. In *Advances in Neural Information Processing Systems*, volume 33, pages 20083–20094, 2020.
- [8] Avrim Blum and Carl Burch. On-line learning and the metrical task system problem. *Machine Learning*, 39(1):35–58, 2000.
- [9] Sébastien Bubeck, Michael B Cohen, James R Lee, and Yin Tat Lee. Metrical task systems on trees via mirror descent and unfair gluing. *SIAM Journal on Computing*, 50(3):909–923, 2021.
- [10] Sébastien Bubeck, Michael B Cohen, Yin Tat Lee, James R Lee, and Aleksander Madry. K-server via multiscale entropic regularization. In *Proc. 50th Symposium on Theory of Computing*, pages 3–16, 2018.
- [11] Niv Buchbinder, Shahar Chen, and Joseph (Seffi) Naor. Competitive analysis via regularization. In *Proc. 25th Symposium on Discrete Algorithms*, pages 436–444, 2014.
- [12] Niv Buchbinder, Anupam Gupta, Marco Molinaro, and Joseph Naor. k-servers with a smile: Online algorithms via projections. In *Proc. 30th Symposium on Discrete Algorithms*, pages 98–116, 2019.
- [13] Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Algorithms with prediction portfolios. In *Advances in Neural Information Processing Systems*, 2022.
- [14] Sreenivas Gollapudi and Debmalya Panigrahi. Online algorithms for rent-or-buy with expert advice. In *International Conference on Machine Learning*, pages 2319–2327, 2019.

- [15] Sreenivas Gollapudi and Debmalya Panigrahi. Online algorithms for rent-or-buy with expert advice. In *Proceedings of the 36th International Conference on Machine Learning*, 2019.
- [16] Chen-Yu Hsu, Piotr Indyk, Dina Katabi, and Ali Vakilian. Learning-based frequency estimation algorithms. In *Proc. Conference on Learning Representations*, 2019.
- [17] Enikő Kevi and Nguyen Kim Thang. Primal-dual algorithms with predictions for online bounded allocation and ad-auctions problems. In *International Conference on Algorithmic Learning Theory*, pages 891–908, 2023.
- [18] Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proc. Conference on Management of Data*, pages 489–504, 2018.
- [19] Ravi Kumar, Manish Purohit, and Zoya Svitkina. Improving online algorithms via ML predictions. In *Proc. 32nd Conference on Neural Information Processing Systems*, pages 9684–9693, 2018.
- [20] Silvio Lattanzi, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Online scheduling via learned weights. In *Proc. Symposium on Discrete Algorithms*, pages 1859–1877, 2020.
- [21] Thodoris Lykouris and Sergei Vassilvitskii. Competitive caching with machine learned advice. In *International Conference on Machine Learning*, pages 3296–3305, 2018.
- [22] Michael Mitzenmacher. A model for learned bloom filters, and optimizing by sandwiching. In *Proc. Conference on Neural Information Processing Systems*, pages 464–473, 2018.
- [23] Michael Mitzenmacher. Scheduling with predictions and the price of misprediction. In *Proc. 11th Innovations in Theoretical Computer Science Conference*, 2020.
- [24] Michael Mitzenmacher, Sergei Vassilvitskii, and Tim Roughgarden. *Beyond the Worst-Case Analysis of Algorithms*, chapter Algorithms with Predictions. Cambridge University Press, 2020.
- [25] Dhruv Rohatgi. Near-optimal bounds for online caching with machine learned advice. In *Proc. Symposium on Discrete Algorithms*, pages 1834–1845, 2020.

Appendix

A Counter example for tight online expert solutions

The following example shows that we cannot expect online expert solutions (in the sense of online algorithms) to be tight on the arriving constraints. In the example below, we display the experts' solutions after each constraint.

$$\begin{array}{rcl}
 \min & x_1 + & x_2 \\
 & x_1 + \frac{1}{2} x_2 \geq 1 \\
 \text{Expert}_1 : & 1 & 0 \\
 \text{Expert}_2 : & 0 & 2 \\
 & x_2 \geq 1 \\
 \text{Expert}_1 : & 1 & 1 \\
 \text{Expert}_2 : & 0 & 2
 \end{array}$$

To have tight a suggestion from Expert₂ on the second constraint, Expert₂ not only has to decrease its value of x_2 (which is not allowed), but even increase the value of x_1 for the first constraint. In other words, Expert₂ has to completely modify its past decisions.

B Counter example for the performance of the algorithm of [2]

Anand, Ge, Kumar and Panigrahi [2] recently proposed online algorithms for online covering problems with multiple expert solutions. We show here a counter example that contradicts Theorem 2.1 presented in Section 3 of their paper. In the proof of Theorem 2.1 the authors state that *the total cost of the algorithm is at most 3 times the potential ϕ at the beginning, i.e., at most $O(\log K)$ times the DYNAMIC benchmark*. However, in our counter example the total cost of their algorithm is $O(L \log(K))$ times the DYNAMIC benchmark, where L is an arbitrary large number.

B.1 Setting

Algorithm 1 (from [2]) receives solutions from K experts. The authors denote with $x_i(j, s)$ the solution from expert s for variable i on constraint j . They assume that the expert solutions are tight, formally:

$$\sum_{i=1}^n a_{ij} x_i(j, s) = 1 \quad \forall s \in [K]$$

The algorithm's performance is compared to the DYNAMIC benchmark, which is the minimum cost solution that is supported by at least one expert at each step, formally:

$$\text{DYNAMIC} = \min_{\hat{\mathbf{x}} \in \hat{X}} \sum_{i=1}^n c_i \hat{x}_i, \text{ where}$$

$$\hat{X} = \{\hat{\mathbf{x}} : \forall i \in [n], \forall j \in [m], \exists s \in [K] \text{ where the solution } x_i(j, s) \leq \hat{x}_i\}$$

While a constraint is not satisfied, their algorithm updates each variable with an increasing rate of

$$\frac{dx_i}{dt} = \frac{a_{ij}}{c_i}(x_i + \delta_{ij})$$

where $\delta_{ij} = \frac{1}{K} \sum_{s=1}^K x_i(j, s)$ is the average of the experts' solutions for x_i at the arrival of constraint j . Algorithm 1 of [2] scales down the problem with 0.5, so it does not increase any variable above 0.5 and satisfies each constraint with value 0.5. The exact solution is obtained by doubling the variables at the end of the execution. (This descaling is an important aspect in the authors' proof.)

B.2 Counter example

In the following example we reveal in an online manner a linear program parametrized by L with K experts and observe the behavior of Algorithm 1 (from [2]).

Objective. The example has $(L \cdot K + 1)$ variables with uniform cost:

$$\min x_1 + x_2 + \cdots + x_K + \cdots + x_{2K} + \cdots + x_{LK} + x_{LK+1}$$

Constraints. There are L batches of $(K - 1)$ constraints. The first constraint of each batch has $(K + 1)$ variables. The last variable (x_{LK+1}) is present in every constraint in every batch, but none of the experts suggests to use this variable. Within a batch, each consecutive constraint has one less variable. The experts set each variable that appears in later batches to 0. The first batch:

$$\begin{array}{rcl} & x_1 + x_2 + \cdots + x_{(K-1)} + x_K + x_{LK+1} \geq 1 \\ \text{Expert}_1 : & 1 & 0 \quad \dots \quad 0 \quad 0 \quad 0 \\ \text{Expert}_2 : & 0 & 1 \quad \dots \quad 0 \quad 0 \quad 0 \\ & \vdots & \\ \text{Expert}_{K-1} : & 0 & 0 \quad \dots \quad 1 \quad 0 \quad 0 \\ \text{Expert}_K : & 0 & 0 \quad \dots \quad 0 \quad 1 \quad 0 \\ & x_2 + \cdots + x_{(K-1)} + x_K + x_{LK+1} \geq 1 \\ \text{Expert}_1 : & 1 & 1 \quad \dots \quad 0 \quad 0 \quad 0 \\ \text{Expert}_2 : & 0 & 1 \quad \dots \quad 0 \quad 0 \quad 0 \\ & \vdots & \\ \text{Expert}_{K-1} : & 0 & 0 \quad \dots \quad 1 \quad 0 \quad 0 \\ \text{Expert}_K : & 0 & 0 \quad \dots \quad 0 \quad 1 \quad 0 \\ & \vdots & \\ & x_{(K-1)} + x_K + x_{LK+1} \geq 1 \\ \text{Expert}_1 : & 1 & 1 \quad \dots \quad 1 \quad 0 \quad 0 \\ \text{Expert}_2 : & 0 & 1 \quad \dots \quad 1 \quad 0 \quad 0 \\ & \vdots & \\ \text{Expert}_{K-1} : & 0 & 0 \quad \dots \quad 1 \quad 0 \quad 0 \\ \text{Expert}_K : & 0 & 0 \quad \dots \quad 0 \quad 1 \quad 0 \end{array}$$

During the first constraint of every batch, the experts' solutions form an identity matrix. With each disappearing variable in the consecutive constraints, experts who suggested to use variables which are no longer available, choose to set the variable with the smallest index. Consequently, $(K - 1)$ experts suggest to use variable $x_{(K-1)}$ and one expert suggests to use x_K during the last constraint in the first batch. The pattern of the experts' solutions are identical for each batch. The constraints of the l^{th} batch ($1 \leq l \leq L$) are:

$$\begin{aligned}
x_{(l-1)K+1} + x_{(l-1)K+2} + \cdots + x_{(l-1)K+(K-1)} + x_{lK} + x_{LK+1} &\geq 1 \\
x_{(l-1)K+2} + \cdots + x_{(l-1)K+(K-1)} + x_{lK} + x_{LK+1} &\geq 1 \\
&\vdots \\
x_{(l-1)K+(K-1)} + x_{lK} + x_{LK+1} &\geq 1
\end{aligned}$$

Claim 1 *The objective value of Algorithm 1 (from [2]) on our example is $O(L \log(K))$ times the DYNAMIC benchmark.*

Proof The optimal solution \mathbf{x}^* of the DYNAMIC benchmark is the solution in which $x_{LK+1}^* = 1$ and $x_i^* = 0$ for $i \neq LK + 1$. We verify that $\mathbf{x}^* \in \hat{X}$. For each $i \neq lK$ where $1 \leq l \leq L$, for each constraint m , $x_i^* \geq 0 = x_i(m, K)$. For $i = lK$ where $1 \leq l \leq L$, for each constraint m , $x_i^* \geq 0 = x_i(m, 1) = x_i(m, 2) = \dots = x_i(m, K - 1)$. Moreover, \mathbf{x}^* satisfies all constraints (since variable x_{LK+1} appears in all constraints). Hence, $\mathbf{x}^* \in \hat{X}$. Subsequently, the objective value of the DYNAMIC benchmark is 1.

By the design of Algorithm 1, the increasing rate of x_{LK+1} is zero throughout the execution, and the variables which are not part of the current constraint are not increased. During the first constraint of each batch, the increasing rate of the first K variables in the batch is $1/K$, since the increasing rate of variable x_i is $(x_i + \frac{1}{K} \sum_{s=1}^K x_i(j, s))$ and initially every variable is set to zero. At the second constraint, the increasing rate of the second variable in the batch is higher than the other variables' increasing rate, because the first expert also uses this variable in its solution. Therefore, the increasing rate of the second variable is $(x_{(l-1)K+2} + 2/K)$, while the other remaining expert variables in the constraint have an increasing rate of $(x_i + 1/K)$. Following the same reasoning (apart from the first constraint in the batch), the variable with the smallest index in the constraint has a higher increasing rate, than the other variables. During the last constraint of each batch, the increasing rate of the last two remaining expert variables are $(x_{(l-1)K+(K-1)} + (K - 1)/K)$ and $(x_{lK} + 1/K)$. Keeping the increasing rates and the constraint satisfaction in mind, we can lower

bound the value of each variable:

$$\begin{aligned}
\frac{1}{K} &\leq x_{(l-1)K+1} \\
\frac{1}{K-1} &\leq x_{(l-1)K+2} \\
\frac{1}{K-2} &\leq x_{(l-1)K+3} \\
&\vdots \\
\frac{1}{3} &\leq x_{(l-1)K+(K-2)} \\
\frac{1}{2} &\leq x_{(l-1)K+(K-1)} \\
\frac{1}{K} &\leq x_{lK}
\end{aligned}$$

Summing the terms together, we get that the objective value increases at least with $O(\log K)$ during each batch. There are L batches, so the total cost of Algorithm 1 is at least $O(L \log(K))$, while the total cost of the DYNAMIC benchmark is 1, which concludes the proof. \square

B.3 Comparison

In this specific counter-example, the LIN-COMB benchmark is equivalent to the static best-expert benchmark, i.e., the solution of Expert $_K$. The objective value of LIN-COMB is L (since the optimal solution sets x_{lK} variables for $1 \leq l \leq L$ to one and other variables to 0). In this counter-example, the objective value of our algorithm is $O(L \log K)$. Consequently, our proposed algorithm is $O(\log K)$ competitive in the LIN-COMB benchmark.