Magicians Don't Move: An Easy Peasy OCRS

Mohammad Reza Aminian, Rad Niazadeh, Pranav Nuti

Abstract

Online contention resolution schemes (OCRSs) are a basic tool for converting ex-ante feasible fractional solutions to relaxations of problems into online policies that remain feasible in every realized sample path. They play a central role in prophet inequalities, Bayesian online allocation, and online mechanism design. In this note, we provide a particularly simple algorithm and analysis for the k-unit OCRS problem, obtaining the asymptotically optimal dependence on k.

1 Introduction

A recurring strategy in Bayesian online decision making problems is to first compute an *ex-ante feasible* plan that satisfies constraints only in expectation, and then to convert it into an online policy that preserves *ex-post* feasibility in every realization. OCRSs are a technical tool that provide a principled way to perform this conversion. The general OCRS framework for downward-closed feasibility systems was developed in Feldman et al. [5], and has since become a standard black-box rounding primitive for Bayesian selection problems.

In the k-unit online contention resolution scheme problem, also referred to as the magician's problem [1], we have a set of elements E, each active independently with known probability p_i, where $\sum_i p_i \leq k$. The elements arrive online in an arbitrary order (chosen potentially by an adversary with no knowledge of the activations and any random bits the algorithm uses), and an algorithm must (irrevocably) select at most k active elements. An algorithm is said to be α-selectable if every element $i \in E$ is selected with probability at least αp_i; equivalently with probability at least α conditional on being active. The goal is to make α as large as possible. An algorithm with a large selectability immediately provides several applications; for instance, an α-selectable OCRS yields an α competitive ratio for the k-unit prophet inequality problem (see [8] for a discussion of the relationship between OCRSs and prophet inequalities).

Multiple algorithms in the literature are known to obtain the optimal $1 - O(1/\sqrt{k})$ dependence on k for the maximal possible selectability. Alaei's magician algorithm [1] is characterized by a dynamic program, and later analysis showing the algorithm is optimal uses a factor-revealing LP, with the explicit selectability coming from the analysis of a differential equation in the vanishing-probability case [6, 9]. A more recent adaptive combinatorial scheme also obtains the $1 - O(1/\sqrt{k})$ guarantee with a simpler algorithm, but the analysis is still involved due to having to study an associated random walk, and the probability it reaches new integral heights [4].

In this short note, we present a different algorithm obtaining the asymptotically optimal $1 - O(1/\sqrt{k})$ selectability, together with the very short analysis developed in [2].

Theorem 1. *For every finite ground set E, every integer $k \geq 1$, and every activation vector $p \in [0,1]^E$ satisfying $\sum_{i \in E} p_i \leq k$, there is an OCRS that selects at most k active elements and is $1 - \sqrt{\frac{2}{k}} - \frac{1}{2k}$-selectable.*

Note that for the sake of simplifying the calculations, we make no attempt to optimize the constants for small values of k.

2 The algorithm and analysis

Fix a finite ground set E, capacity $k \geq 1$, and activation probabilities $p = (p_i)_{i \in E}$ with $p_i \in [0, 1]$ and $\sum_i p_i \leq k$. Let

$$r := \left\lceil \sqrt{k/2} \right\rceil, \quad \gamma := 1 - \frac{r}{k}.$$

Let Y_i be independent Bernoulli random variables with means $\mathbb{P}[Y_i = 1] = \gamma p_i$, and let $N := \sum_i Y_i$. Define μ to be the law of the random set $\{i : Y_i = 1\}$ conditioned on $N \leq k$. To start the algorithm, sample a random set R from μ.

Algorithm 1 Stationary policy with random discarding

Require: activation probabilities p_i, capacity k, arrival order

$r \leftarrow \left\lceil \sqrt{k/2} \right\rceil$ and $\gamma \leftarrow 1 - r/k$
Sample $R \sim \mu$, i.e. $R = \{i : Y_i = 1\}$ for independent $Y_i \sim \text{Bernoulli}(\gamma p_i)$ conditioned on $|R| \leq k$ ▷ R is the real selected set plus future predictions
for each arriving element i **do**
    $R \leftarrow R \setminus \{i\}$ ▷ forget the old prediction for i
    if i is active and $|R| < k$ **then**
        With probability γ, accept i and set $R \leftarrow R \cup \{i\}$
        Otherwise reject i
    end if
end for
return R ▷ at the end, all predictions have been replaced with selections

The set R should be interpreted as "selected elements from the past plus predictions for the future." When an element i arrives, the algorithm first deletes i from R, "forgetting its prediction for i". If element i is active and the remaining real-plus-predicted set R has room, the element is accepted and added into R with the fixed probability γ. Else the element is rejected. After all the elements are processed, the initial set of predictions is replaced entirely by a set of selections, which is then returned as output.

The main reason we refer to our algorithm as being stationary is the following proposition:

Proposition 1. *After every arrival in Algorithm 1, the set R continues to have distribution μ. Consequently, the final selected set has distribution μ.*

Proof. Initially $R \sim \mu$ by construction. Assume inductively that $R \sim \mu$ immediately before element i is processed, and set $T = R \setminus \{i\}$. The activation of i is independent of T and occurs with probability p_i. If $|T| < k$, Algorithm 1 inserts i into T with probability $p_i \gamma$. If $|T| = k$, it never inserts i. This is exactly the conditional distribution of the ith coordinate under μ given $R \setminus \{i\} = T$. The rest of the set (or coordinates) remains equal to T. Thus the post-update set again has law μ. Induction proves the claim about the final selected set. □

Note that the final selected set must have size at most k, since this is where μ is supported. Thus, to finish the proof, it suffices to show the following proposition:

Proposition 2. *If $R \sim \mu$, then*

$$\mathbb{P}[i \in R] \geq \left(1 - \sqrt{\frac{2}{k}} - \frac{1}{2k}\right) p_i.$$

Proof. Let Y_i and $N = \sum_i Y_i$ be the independent Bernoulli variables used to define μ, and write $N_{-i} := \sum_{j \neq i} Y_j$. Then

$$\begin{aligned}
\mathbb{P}[i \in R] &= \mathbb{P}[Y_i = 1 \mid N \leq k] \\
&= \gamma p_i \cdot \frac{\mathbb{P}[N_{-i} \leq k-1]}{\mathbb{P}[N \leq k]} \\
&\geq \gamma p_i \cdot \frac{\mathbb{P}[N < k]}{\mathbb{P}[N \leq k]}.
\end{aligned}$$

It remains to lower bound the last factor. Since $\sum_i p_i \leq k$, the mean of N satisfies

$$\mathbb{E}[N] = \gamma \sum_i p_i \leq \gamma k.$$

We use three classical facts about sums of independent Bernoulli random variables: such distributions are unimodal, the mode lies within a distance of 1 from the mean [3], and so does the median [7]. Since γk is an integer and $\mathbb{E}[N] \leq \gamma k$, γk is at least the median and the mode of N. Therefore

$$\begin{aligned}
\mathbb{P}[N \leq k] &\geq 2\mathbb{P}[\gamma k < N \leq k] && \text{(since } \gamma k \text{ is at least the median)} \\
&= 2 \sum_{\ell=\gamma k+1}^{k} \mathbb{P}[N = \ell] \\
&\geq 2(k - \gamma k)\mathbb{P}[N = k] && \text{(by unimodality, since } \gamma k \text{ is at least the mode)} \\
&= 2r\mathbb{P}[N = k].
\end{aligned}$$

Hence, we have

$$\frac{\mathbb{P}[N < k]}{\mathbb{P}[N \leq k]} = 1 - \frac{\mathbb{P}[N = k]}{\mathbb{P}[N \leq k]} \geq 1 - \frac{1}{2r}.$$

Thus

$$\gamma \frac{\mathbb{P}[N < k]}{\mathbb{P}[N \leq k]} \geq \left(1 - \frac{r}{k}\right)\left(1 - \frac{1}{2r}\right) = 1 - \frac{r}{k} - \frac{1}{2r} + \frac{1}{2k}.$$

Now note that $r = \lceil \sqrt{k/2} \rceil$. Since

$$\sqrt{k/2} \leq r \leq \sqrt{k/2} + 1,$$

we have

$$\frac{r}{k} + \frac{1}{2r} - \frac{1}{2k} \leq \frac{\sqrt{k/2}+1}{k} + \frac{1}{2\sqrt{k/2}} - \frac{1}{2k} = \sqrt{\frac{2}{k}} + \frac{1}{2k}.$$

This gives the claimed selectability. $\square$

3 Further discussion

The algorithm in this note is an instance of a more general idea from [2]. In the setting of this note, we first sample independent Bernoulli random variables with probabilities γp_i, condition on the event that the resulting set has size at most k, and then maintain this same conditional distribution throughout the online process by replacing predictions with real selections. Our proof shows that, once the right initial distribution is chosen, the online part of the algorithm is almost forced: when an element arrives, we must delete its prediction and resample whether it belongs to the maintained set, subject to the fact that the element must be active, and the distribution of the maintained set remains fixed.

The full "stationary OCRS" framework develops this principle for much more general feasibility constraints. Instead of conditioning only on the event $|R| \leq k$, one may condition on membership in an arbitrary downward-closed feasibility family. Similarly, instead of using the homogeneous probabilities γp_i, one may use more general probabilities for membership in the initial set. The same philosophy then gives an online algorithm: initialize from the chosen feasible distribution, and when an element arrives, update its coordinate according to the corresponding conditional probability. If the chosen distribution has sufficiently large marginals, this immediately gives a good OCRS.

A useful way to find such distributions is through the maximum-entropy method. Roughly speaking, the maximum-entropy distribution is the most product-like distribution over feasible sets with the desired marginals. This makes the conditional probabilities that arise in our algorithm simple to calculate, while moving most of the analysis to an "offline" question about the existence of a random feasible set with good properties.

In [2], we use the maximum-entropy viewpoint to obtain a k-unit OCRS with selectability $\frac{\Pr(\mathrm{Pois}(k)<k)}{\Pr(\mathrm{Pois}(k)\leq k)}$. This improves the asymptotic selectability from $1-\sqrt{\frac{2}{k}}$ to $1-\sqrt{\frac{2}{\pi k}}$, while also yielding the optimal factor of $1/2$ for $k=1$. We also obtain a simple, conjectured-to-be-optimal, $(3-\sqrt{5})/2$-selectable OCRS for bipartite matchings. The stationary distribution viewpoint can turn the design of online contention resolution schemes into the search for the right distribution over feasible sets, often yielding algorithms that are both explicit and easier to reason about than more indirect constructions.

References

[1] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[2] Mohammad Reza Aminian, Rad Niazadeh, and Pranav Nuti. Stationary online contention resolution schemes, 2026. URL https://arxiv.org/abs/2603.21532.

[3] John N Darroch. On the distribution of the number of successes in independent trials. *The Annals of Mathematical Statistics*, 35(3):1317–1321, 1964.

[4] Atanas Dinev and S Matthew Weinberg. Simple and optimal online contention resolution schemes for k-uniform matroids. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, page 39:1–39:23. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

[5] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM, 2016.

[6] Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tight guarantees for multiunit prophet inequalities and online stochastic knapsack. *Operations Research*, 73(3):1703–1721, 2025.

[7] Kumar Jogdeo and Stephen M Samuels. Monotone convergence of binomial probabilities and a generalization of Ramanujan's equation. *The Annals of Mathematical Statistics*, 39(4):1191–1195, 1968.

[8] Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112, page 57:1–57:14. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.

[9] Xinshang Wang, Van-Anh Truong, and David Bank. Online advance admission scheduling for services with customer preferences. *arXiv preprint arXiv:1805.10412*, 2018.