Universal Online Contention Resolution with Preselected Order

Junyao Zhao
IRIF, CNRS, Université Paris Cité, Paris, France

Abstract

Online contention resolution scheme (OCRS) is a powerful technique for online decision making, which – in the case of matroids – given a matroid and a prior distribution of active elements, selects a subset of active elements that satisfies the matroid constraint in an online fashion. OCRS has been studied mostly for product distributions in the literature. Recently, universal OCRS, that works even for correlated distributions, has gained interest, because it naturally generalizes the classic notion, and its existence in the random-order arrival model turns out to be equivalent to the matroid secretary conjecture. However, currently very little is known about how to design universal OCRSs for any arrival model. In this work, we consider a natural and relatively flexible arrival model, where the OCRS is allowed to preselect (i.e., non-adaptively select) the arrival order of the elements, and within this model, we design simple and optimal universal OCRSs that are computationally efficient. In the course of deriving our OCRSs, we also discover an efficient reduction from universal online contention resolution to the matroid secretary problem for any arrival model, answering a question posed in [8].

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Matroids, online contention resolution schemes, secretary problems

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.137

Category Track A: Algorithms, Complexity and Games

Related Version *Full paper*: `https://arxiv.org/abs/2504.16327`

Funding *Junyao Zhao*: Supported by a postdoctoral fellowship from the Fondation Sciences Mathématiques de Paris. Part of the work was done while the author was a Ph.D. student at Stanford University, where he was supported by NSF CCF-1954927.

1 Introduction

A contention resolution scheme (CRS) is an algorithm which given a set system $\mathcal{M} \subseteq 2^{[n]}$ and a set of *active* elements $A \subseteq [n]$ sampled from a known prior distribution $\mathcal{D}_A$, selects a subset $X \subseteq A$ that satisfies the feasibility constraint $X \in \mathcal{M}$. The design goal of CRS is to guarantee that every element in $[n]$ is selected with some constant probability α conditioned on it being active. Informally, when such CRS exists for set system $\mathcal{M}$ and prior distribution $\mathcal{D}_A$, we call $\mathcal{D}_A$ an α-*uncontentious* distribution for $\mathcal{M}$. CRS was introduced by [7] and has since been studied for various set systems. In this paper, we focus on a most studied set system in the literature – matroid (see Definition 4).

A particular class of CRSs, known as online contention resolution schemes (OCRSs), has found many applications in online decision making, such as prophet inequalities, stochastic probing, and sequential posted-price auctions (e.g., [15, 11, 1]). Specifically, an OCRS knows only $\mathcal{M}$ and $\mathcal{D}_A$ but not the set of active elements A at the beginning. Instead, elements in $[n]$ arrive one by one (their order depends on the specific arrival model), and upon the arrival of each element $i \in [n]$, it is revealed whether $i \in A$, and the OCRS must decide immediately and irrevocably whether to include i in its solution set X before the next element arrives.

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).
Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis
Article No. 137; pp. 137:1–137:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Many elegant OCRSs with strong guarantees (in various arrival models) have been discovered (e.g., [7, 11, 1, 19, 12]), but most of them were established exclusively for *product* distribution $\mathcal{D}_A$ (i.e., each element is active independently with some probability), except for [10] and [14], which studied OCRSs for pairwise independent distribution $\mathcal{D}_A$.

Recently, *universal* OCRSs [8], that work even for general *correlated* distribution $\mathcal{D}_A$, have started to gain interest. Informally, an OCRS is (α, β)-universal if it guarantees that every element is selected with some constant probability β conditioned on it being active, for any matroid $\mathcal{M}$ and (arbitrarily correlated) α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$. Universal OCRSs naturally generalize classic OCRSs for product distributions. Moreover, Dughmi [8, 9] proved that the existence of universal OCRSs in the random-order arrival model is equivalent to the matroid secretary conjecture posed by [3].

Currently, very little is known about how to design universal OCRSs, except that Dughmi [8] showed that for any arrival model[1], a universal OCRS exists if there is a constant-competitive matroid secretary algorithm for that model. For the free-order arrival model, where the algorithm is allowed to *adaptively* choose the arrival order of the remaining elements after observing any number of elements, [18] designed a constant-competitive matroid secretary algorithm. This algorithm, together with Dughmi's result, implies that there exists a universal OCRS in the free-order model. However, this approach is not known to be computationally efficient, making it difficult to understand the implied universal OCRS for specific problem instances. Indeed, Dughmi's result is information-theoretic, and whether it can be made computationally efficient was left as an open question [8, Section 6].

In this work, we strive to design universal OCRSs that are efficiently computable and simple to understand. We focus on a natural and relatively flexible arrival model in which the OCRS is allowed to preselect the arrival order of the elements (i.e., *non-adaptively* choose the arrival order of the elements given $\mathcal{M}$ and $\mathcal{D}_A$). We adopt the term *preselected order* to differentiate from the free-order model. The preselected-order model is a step toward the random-order model, and it has been studied for various online decision problems, including prophet inequalities [17, 2, 20, 21, 5], the multi-choice secretary problem [16], sequential posted-price auctions [6, 4], stochastic probing [15] and OCRSs for product distributions [7]. The main contribution of our work is the design and analysis of three different universal OCRSs in the preselected-order model.

1.1 Overview of our universal OCRSs

Now we give an overview of our universal OCRSs. All of our universal OCRSs are generalizations of *ordered* OCRSs for product distributions [7, 15] with necessary randomization. Briefly, an ordered OCRS [15, Definition 3.2] preselects the arrival order of the elements, and then upon the arrival of each element, it greedily adds the element to the solution set, provided that the element is *selectable* (i.e., if it is active and can be added to the solution set without violating the matroid constraint).

Our first two universal OCRSs (Algorithm 1 and Algorithm 3) generalize ordered OCRSs by *subsampling selectable elements*. Specifically, these two OCRSs first preselect the arrival order (based on their respective criteria) and sample a subset of elements $T \subseteq [n]$. Then, upon the arrival of each element, they add the element to the solution set if it is selectable and belongs to T. Algorithm 1 and Algorithm 3 use different subsampling methods of

[1] To be precise, Dughmi's result [8, Theorem 4.1] was stated for the random-order arrival model, but the proof of that result applies to any arrival model.

independent interest – Algorithm 1 includes each element in T independently with a certain probability, while Algorithm 3 employs a correlated subsampling method and achieves a slightly stronger universality guarantee.

▶ **Theorem 1** (Restatement of Theorem 16 and Theorem 21). *For all $\alpha \in [0,1]$, Algorithm 1 is an $(\alpha, \frac{\alpha^2}{4})$-universal OCRS, and Algorithm 3 is an $(\alpha, \frac{\alpha^2}{2})$-universal OCRS. Both algorithms preselect the arrival order of the elements before any element arrives.*

Instead of subsampling selectable elements, our third universal OCRS (presented in Section 5) *samples the arrival order* of the elements. Specifically, given matroid $\mathcal{M}$ and prior distribution $\mathcal{D}_A$, this OCRS first computes a distribution over permutations by solving a linear program (LP), which is similar to the LP used by [7] to compute offline CRSs for product distributions. Then, it samples the arrival order from this distribution, and upon the arrival of each element, it includes the element in the solution set if the element is selectable. The universality guarantee of this OCRS is nearly optimal[2].

▶ **Theorem 2** (Informal restatement of Theorem 25). *For any $\varepsilon > 0$, there exists a computationally efficient OCRS with preselected order, which is $(\alpha, (1-\varepsilon)\alpha)$-universal for all $\alpha \in [0,1]$.*

For comparison, our first two OCRSs have weaker universality guarantees, but they are easier to reason about for specific problem instances. In contrast, our third OCRS provides the strongest universality guarantee, though it is less intuitive because of the use of the LP.

1.2 Universal OCRSs from secretary algorithms

In the course of deriving our third universal OCRS, we discovered an LP-based *efficient reduction* from universal online contention resolution to the matroid secretary problem for any arrival model (see Section 6 for the setup of the matroid secretary problem), thereby answering the aforementioned question from [8].

▶ **Theorem 3** (Informal restatement of Theorem 28). *For any $c, \varepsilon > 0$, for any arrival model, if there is a computationally efficient c-competitive matroid secretary algorithm, then there exists a computationally efficient OCRS that is $(\alpha, (1-\varepsilon)c \cdot \alpha)$-universal for all $\alpha \in [0,1]$.*

Briefly, Dughmi's information-theoretic reduction [8, Theorem 4.1] was based on the separating hyperplane theorem. Our reduction replaces that with an LP duality argument, and then applies a technique from [19] to solve the corresponding LPs efficiently.

2 Preliminaries

2.1 Matroids

We start by introducing essential definitions and properties of *matroids*, and we refer interested readers to [22] for a comprehensive treatment of matroid theory. Essentially, a matroid is a set system with some independence structure, which is defined as follows.

▶ **Definition 4** (matroid). *A set system $\mathcal{M} \subseteq 2^{[n]}$ is a matroid if it satisfies the following properties:*

i. $\emptyset \in \mathcal{M}$.

[2] We note that the ε loss is only due to computation – we will show that (α, α)-universal OCRSs exist.

ii. *If $X \in \mathcal{M}$, then $Y \in \mathcal{M}$ for all $Y \subseteq X$.*
iii. *If $X, Y \in \mathcal{M}$ and $|Y| < |X|$, then there exists $i \in X \setminus Y$ such that $Y \cup \{i\} \in \mathcal{M}$.*
Given a matroid, we can associate a *(weighted) rank function* with it.

▶ **Definition 5** (rank). *Given a matroid $\mathcal{M} \subseteq 2^{[n]}$, the rank function $r_{\mathcal{M}} : 2^{[n]} \to \mathbb{Z}_{\geq 0}$ associated with $\mathcal{M}$ is $r_{\mathcal{M}}(X) := \max_{Y \subseteq X} |Y|$ s.t. $Y \in \mathcal{M}$. For convenience, for any sets $X, Y \subseteq [n]$, we denote $r_{\mathcal{M}}(X \mid Y) := r_{\mathcal{M}}(X \cup Y) - r_{\mathcal{M}}(Y)$.*

More generally, for any weight vector $w \in \mathbb{R}^n_{\geq 0}$, the weighted rank function $r_{\mathcal{M},w} : 2^{[n]} \to \mathbb{R}_{\geq 0}$ is $r_{\mathcal{M},w}(X) := \max_{Y \subseteq X} \sum_{i \in Y} w_i$ s.t. $Y \in \mathcal{M}$.

Moreover, we can define notions of *span* and *basis* for a matroid.

▶ **Definition 6** (span). *Given a matroid $\mathcal{M} \subseteq 2^{[n]}$ and its rank function $r_{\mathcal{M}} : 2^{[n]} \to \mathbb{Z}_{\geq 0}$, we say that an element $i \in [n]$ is spanned by a set $X \subseteq [n]$ if $r_{\mathcal{M}}(X \cup \{i\}) = r_{\mathcal{M}}(X)$, and we define the span of X as $\operatorname{span}_{\mathcal{M}}(X) := \{i \in [n] \mid i \text{ is spanned by } X\}$.*

▶ **Definition 7** (basis). *Given a matroid $\mathcal{M} \subseteq 2^{[n]}$ and a set $X \subseteq [n]$, we say that a subset $Y \subseteq X$ is a basis of X if $Y \in \mathcal{M}$ and for all $i \in X \setminus Y$, $Y \cup \{i\} \notin \mathcal{M}$.*

Furthermore, we define the *restriction* of a matroid to a set of elements.

▶ **Definition 8** (restriction). *Given a matroid $\mathcal{M} \subseteq 2^{[n]}$ and a set $X \subseteq [n]$, we define the restriction of $\mathcal{M}$ to X as $\mathcal{M}_X := \{Y \subseteq X \mid Y \in \mathcal{M}\}$.*

In the following lemma, we state several well-known properties of matroids.

▶ **Lemma 9** (see e.g., [22]). *Any matroid $\mathcal{M} \subseteq 2^{[n]}$ satisfies the following properties:*
i. *For any $X \subseteq [n]$, $r_{\mathcal{M}}(X) \leq |X|$.*
ii. *For any $X, Y \subseteq [n]$, $r_{\mathcal{M}}(X \mid Y) \leq \sum_{i \in X} r_{\mathcal{M}}(\{i\} \mid Y)$.*

2.2 Contention resolution schemes

Now we introduce *contention resolution schemes* (CRSs) for matroids. Given a matroid $\mathcal{M} \subseteq 2^{[n]}$ and a random set of *active* elements $A \subseteq [n]$ sampled from a prior distribution[3] $\mathcal{D}_A \in \Delta(2^{[n]})$, a CRS selects a subset of active elements $X \subseteq A$ such that $X \in \mathcal{M}$. Formally, we first let $\Phi_{\mathcal{M}}$ denote the family of maps that take an active set $A \subseteq [n]$ as input and output a subset $X \subseteq A$ such that $X \in \mathcal{M}$, i.e.,

$$\Phi_{\mathcal{M}} := \{\phi : 2^{[n]} \to \mathcal{M} \mid \text{for all } A \subseteq [n],\ \phi(A) \subseteq A\}.$$

Then, a CRS for $(\mathcal{M}, \mathcal{D}_A)$ is a random map sampled from a distribution $\mathcal{D}_\phi \in \Delta(\Phi_{\mathcal{M}})$ (since the choice of $\mathcal{D}_\phi$ fully determines the CRS, we also use the notation $\mathcal{D}_\phi$ to refer to the CRS itself). Moreover, for $\alpha \in [0, 1]$, we say that a CRS $\mathcal{D}_\phi$ for $(\mathcal{M}, \mathcal{D}_A)$ is *α-balanced* if it satisfies

$$\Pr_{A \sim \mathcal{D}_A,\, \phi \sim \mathcal{D}_\phi}[i \in \phi(A) \mid i \in A] \geq \alpha \text{ for all } i \in [n] \text{ s.t. } \Pr_{A \sim \mathcal{D}_A}[i \in A] > 0,$$

and we say that $\mathcal{D}_A$ is *α-uncontentious* for $\mathcal{M}$ if there exists an α-balanced CRS for $(\mathcal{M}, \mathcal{D}_A)$.

Furthermore, we use the notation $\mathcal{D}_A^S$ to denote marginal distributions of $\mathcal{D}_A$. That is, for any $S \subseteq [n]$, $\mathcal{D}_A^S$ is defined as follows: $\forall Z \subseteq S$, $\Pr_{A' \sim \mathcal{D}_A^S}[A' = Z] := \Pr_{A \sim \mathcal{D}_A}[A \cap S = Z]$. We note that if $\mathcal{D}_A$ is α-uncontentious for matroid $\mathcal{M}$, then $\mathcal{D}_A^S$ is α-uncontentious for the restriction $\mathcal{M}_S$.

[3] Throughout the paper, we assume w.l.o.g. that $\mathcal{D}_A$ satisfies that $\forall i \in [n]$, $\Pr_{A \sim \mathcal{D}_A}[i \in A] > 0$. This assumption will simplify the presentation, as it ensures that the probabilities conditioned on the event $i \in A$ are well-defined.

▶ **Lemma 10.** *Given any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, for any $S \subseteq [n]$, the marginal distribution $\mathcal{D}_A^S$ is α-uncontentious for the restriction $\mathcal{M}_S$.*

A self-contained proof of Lemma 10 is provided in the full paper.

Online contention resolution schemes with preselected order

In this paper, we mostly focus on *online contention resolution schemes* (OCRSs) that first *preselect* the arrival order of the elements and then select a subset of active elements in an *online* fashion. Specifically, an OCRS with preselected order is an algorithm ALG that works in three stages:

(1) Given oracle access[4] to a matroid $\mathcal{M} \subseteq 2^{[n]}$ and a prior distribution $\mathcal{D}_A \in \Delta(2^{[n]})$, ALG first selects a (random) permutation $\pi : [n] \to [n]$ and initializes an empty solution set X_{ALG}.

(2) A random set of active elements $A \subseteq [n]$ is sampled from $\mathcal{D}_A$ and is unknown to ALG.

(3) Then, ALG runs in n steps. At each step $i \in [n]$, it is revealed to ALG whether $\pi(i) \in A$. If $\pi(i)$ is *selectable* (i.e., $\pi(i) \in A$ and $X_{\mathsf{ALG}} \cup \{\pi(i)\} \in \mathcal{M}$), ALG must decide immediately and irrevocably whether to add $\pi(i)$ to X_{ALG}, and it can use randomness to make this decision.

Similar to general CRSs, for $\beta \in [0, 1]$, we say that ALG is *β-balanced* for $(\mathcal{M}, \mathcal{D}_A)$ if it satisfies

$$\Pr_{A\sim\mathcal{D}_A,\ \text{randomness of } \mathsf{ALG}}[i \in X_{\mathsf{ALG}} \mid i \in A] \geq \beta \text{ for all } i \in [n] \text{ s.t. } \Pr_{A\sim\mathcal{D}_A}[i \in A] > 0.$$

Moreover, for $0 \leq \beta \leq \alpha \leq 1$, we say that ALG is (α, β)-*universal* if for any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any α-uncontentious distribution $\mathcal{D}_A \in \Delta(2^{[n]})$ for $\mathcal{M}$, ALG is β-balanced for $(\mathcal{M}, \mathcal{D}_A)$. Intuitively, universality means that ALG is approximately as balanced as any CRS for any instance $(\mathcal{M}, \mathcal{D}_A)$. Furthermore, we say ALG is *computationally efficient* if it runs in $O\left(\mathsf{poly}\left(\frac{n}{p_{\min}}\right) \cdot (t_{\mathcal{D}_A} + t_{\mathcal{M}})\right)$ time, where $p_{\min} := \min_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i \in A]$, and $t_{\mathcal{D}_A}$, $t_{\mathcal{M}}$ are the time it takes to generate a sample from $\mathcal{D}_A$ and to check whether a set of elements belongs to $\mathcal{M}$ respectively. We note that the runtime must depend on $\frac{1}{p_{\min}}$, if ALG is (α, β)-universal for arbitrary constants $\alpha, \beta \in (0, 1)$.

▶ **Example 11.** We consider the 1-uniform matroid $\mathcal{M} = \{\emptyset, \{1\}, \{2\}, \ldots, \{n\}\}$. For any constant $\alpha \in (0, 1)$ and any $\delta \in \left(0, \frac{1}{n+1/\alpha-2}\right]$, we consider a family of prior distributions $\{\mathcal{D}_A^{(j)} \mid j \in [n]\}$, where each distribution $\mathcal{D}_A^{(j)}$ is defined as follows:

$$\Pr[A = \emptyset] = 1 - \delta \cdot \left(n + \frac{1}{\alpha} - 2\right), \quad \Pr[A = [n]] = \delta \cdot \left(\frac{1}{\alpha} - 1\right),$$
$$\Pr[A = \{i\}] = \delta \text{ for all } i \neq j.$$

Note that $p_{\min} = \Pr[j \in A] = \delta \cdot \left(\frac{1}{\alpha} - 1\right)$ for each prior distribution $\mathcal{D}_A^{(j)}$. Moreover, every prior distribution $\mathcal{D}_A^{(j)}$ is α-uncontentious for matroid $\mathcal{M}$, because the CRS, that selects element j if $A = [n]$ and selects element i if $A = \{i\}$ for any $i \neq j$, is α-balanced for $(\mathcal{M}, \mathcal{D}_A^{(j)})$.

[4] We assume that matroid $\mathcal{M}$ is given by a membership oracle that answers whether $S \in \mathcal{M}$ for any input set $S \subseteq [n]$, and prior distribution $\mathcal{D}_A$ is given by an oracle that outputs a fresh sample of $\mathcal{D}_A$ upon each query.

However, given instance $(\mathcal{M}, \mathcal{D}_A^{(j)})$ for an unknown j chosen uniformly at random from $[n]$, with only sample access to $\mathcal{D}_A^{(j)}$, if an algorithm ALG only draws $o(\frac{1}{p_{\min}}) = o(\frac{1}{\delta})$ samples, then with high probability, it cannot distinguish element j from most other elements. As a result, if the input set of active elements A is $[n]$, which is the only case where $j \in A$, ALG will not be able to select element j with constant probability. We prove this formally in the full paper.

2.3 Other useful notion

Finally, we define a subsampling operator as follows.

▶ **Definition 12** (subsampling operator $\mathcal{T}_\rho$)**.** *For any $\rho \in [0, 1]$, the random operator $\mathcal{T}_\rho$ takes any set $X \subseteq [n]$ as input and outputs a random subset $\mathcal{T}_\rho(X)$ of X such that each element in X appears in $\mathcal{T}_\rho(X)$ independently with probability ρ.*

3 A simple universal OCRS via independent subsampling

In this section, we design a simple universal OCRS with preselected order by subsampling selectable elements independently. Although this subsampling-based OCRS has a slightly weaker universality guarantee than the one in Section 4, its analysis is simpler and of independent interest.

3.1 Main structural lemma

Our universal OCRS is based on the following structural lemma, which says that given a matroid $\mathcal{M}$ and an uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, there exists an element i, such that if we sample a set of active elements A from $\mathcal{D}_A$ and then independently remove each element in A with some constant probability, there is a decent chance that the remaining elements do not span element i conditioned on i being active. This type of lemma was also central to the previous ordered/greedy OCRSs [7, 11] for product distributions.

▶ **Lemma 13.** *For any $\alpha \in [0, 1]$ and $\rho \in [0, \alpha]$, given any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, there exists $i \in [n]$ such that*

$$\Pr_{A\sim\mathcal{D}_A}[i \notin \operatorname{span}_{\mathcal{M}}(\mathcal{T}_\rho(A)) \mid i \in A] \geq \alpha - \rho.$$

We note that using subsampling is necessary for the above lemma, in the sense that even if A is sampled from an α-uncontentious distribution $\mathcal{D}_A$ for some strictly positive constant α, it is still possible that every element i is always spanned by $A \setminus \{i\}$ conditioned on $i \in A$.

▶ **Example 14.** Consider a simple matroid $\mathcal{M} = \{\emptyset, \{1\}, \{2\}\}$ with two elements and a distribution $\mathcal{D}_A$ that is specified by $\Pr[A = \{1, 2\}] = \frac{1}{2}$ and $\Pr[A = \emptyset] = \frac{1}{2}$. $\mathcal{D}_A$ is $\frac{1}{2}$-uncontentious because the CRS, that selects an element from 1 and 2 uniformly at random when $A = \{1, 2\}$, is $\frac{1}{2}$-balanced. However, we notice that $\Pr_{A\sim\mathcal{D}_A}[i \in \operatorname{span}_{\mathcal{M}}(A \setminus \{i\}) \mid i \in A] = 1$ for any $i \in \{1, 2\}$.

Before proving Lemma 13, we state Lemma 15, which is implied by the characterization of uncontentious distributions [8, Theorem 2.1]. We provide the proof in the full paper.

▶ **Lemma 15.** *For any $\alpha \in [0, 1]$, for any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, for any weight vector $w \in \mathbb{R}^n_{\geq 0}$, it holds that $\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M},w}(A)] \geq \alpha \cdot \mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in A} w_i]$, which in particular, implies that $\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A)] \geq \alpha \cdot \mathbb{E}_{A\sim\mathcal{D}_A}[|A|]$.*

Now we proceed to the proof of Lemma 13.

Proof of Lemma 13. First, we upper bound $\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A)]$ using basic matroid properties:

$$\begin{aligned}
&\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A)] && (1)\\
&= \mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A\cup\mathcal{T}_\rho(A))]\\
&= \mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A\mid\mathcal{T}_\rho(A))+r_{\mathcal{M}}(\mathcal{T}_\rho(A))]\\
&= \mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A\mid\mathcal{T}_\rho(A))]+\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(\mathcal{T}_\rho(A))]\\
&\le \mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in A} r_{\mathcal{M}}(\{i\}\mid\mathcal{T}_\rho(A))]+\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(\mathcal{T}_\rho(A))] && \text{(By Lemma 9-ii)}\\
&\le \mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in A} r_{\mathcal{M}}(\{i\}\mid\mathcal{T}_\rho(A))]+\mathbb{E}_{A\sim\mathcal{D}_A}[|\mathcal{T}_\rho(A)|] && \text{(By Lemma 9-i)}\\
&= \mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in A} r_{\mathcal{M}}(\{i\}\mid\mathcal{T}_\rho(A))]+\rho\cdot\mathbb{E}_{A\sim\mathcal{D}_A}[|A|] && \text{(By definition of } \mathcal{T}_\rho\text{).} \quad (2)
\end{aligned}$$

Since $r_{\mathcal{M}}(\{i\}\mid\mathcal{T}_\rho(A))$ equals 1 if $i\notin \mathrm{span}_{\mathcal{M}}(\mathcal{T}_\rho(A))$ and 0 otherwise, it follows that

$$\begin{aligned}
\mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in A} r_{\mathcal{M}}(\{i\}\mid\mathcal{T}_\rho(A))] &= \mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in[n]} \mathbb{1}(i\in A,\, i\notin \mathrm{span}_{\mathcal{M}}(\mathcal{T}_\rho(A)))]\\
&= \sum_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i\in A,\, i\notin \mathrm{span}_{\mathcal{M}}(\mathcal{T}_\rho(A))]. \qquad (3)
\end{aligned}$$

Combining Ineq. (1) and Ineq. (3), we have that

$$\mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M}}(A)] \le \sum_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i\in A,\, i\notin \mathrm{span}_{\mathcal{M}}(\mathcal{T}_\rho(A))]+\rho\cdot\mathbb{E}_{A\sim\mathcal{D}_A}[|A|].$$

By Lemma 15, this implies that

$$\sum_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i\in A,\, i\notin \mathrm{span}_{\mathcal{M}}(\mathcal{T}_\rho(A))] \ge (\alpha-\rho)\cdot\mathbb{E}_{A\sim\mathcal{D}_A}[|A|] = (\alpha-\rho)\cdot\sum_{i\in[n]}\Pr_{A\sim\mathcal{D}_A}[i\in A].$$

Therefore, there exists $i\in[n]$ such that

$$\Pr_{A\sim\mathcal{D}_A}[i\notin \mathrm{span}(\mathcal{T}_\rho(A))\mid i\in A] = \frac{\Pr_{A\sim\mathcal{D}_A}[i\notin \mathrm{span}_{\mathcal{M}}(\mathcal{T}_\rho(A)),\, i\in A]}{\Pr_{A\sim\mathcal{D}_A}[i\in A]} \ge \alpha-\rho. \qquad \blacktriangleleft$$

3.2 Constructing our universal OCRS with independent subsampling

Our universal OCRS with independent subsampling (Algorithm 1) first samples a set $T=\mathcal{T}_{\frac{\alpha}{2}}([n])$ and applies Lemma 13 iteratively to determine an order of the elements, and then following this order, it greedily selects each element that is selectable and belongs to T.

We state the universality guarantee of Algorithm 1 in Theorem 16 and defer the proof to the full paper, where we will also show that the universality guarantee in Theorem 16 is essentially tight for Algorithm 1.

▶ **Theorem 16.** *Algorithm 1 is an $(\alpha, \frac{\alpha^2}{4})$-universal OCRS for any $\alpha\in[0,1]$.*

Moreover, we note that Algorithm 1 does not specify how to find element $\pi(i)$ at Line 4. In the full paper, we implement this step by estimating probabilities $\Pr_{A'\sim\mathcal{D}_A^{S_i}}[j\in \mathrm{span}_{\mathcal{M}}(\mathcal{T}_{\frac{\alpha}{2}}(A'))\mid j\in A']$ for all $i\in[n]$ and $j\in S_i$ using Monte-Carlo sampling, which results in the following corollary.

Algorithm 1 UNIVERSAL-OCRS-WITH-INDEPENDENT-SUBSAMPLING.

Input : Matroid $\mathcal{M} \subseteq 2^{[n]}$, $\alpha \geq 0$, α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, and $A \sim \mathcal{D}_A$

Output: $X_n \subseteq A$ such that $X_n \in \mathcal{M}$

$S_n \leftarrow [n]$, $X_0 \leftarrow \emptyset$, and $T \leftarrow \mathcal{T}_{\frac{\alpha}{2}}([n])$;

for $i = n, \ldots, 1$ `// notice the reversed order`

do

  Find an element $\pi(i) \in S_i$ such that $\Pr_{A' \sim \mathcal{D}_A^{S_i}}[\pi(i) \notin \operatorname{span}_{\mathcal{M}}(\mathcal{T}_{\frac{\alpha}{2}}(A')) \mid \pi(i) \in A'] \geq \frac{\alpha}{2}$;

  $S_{i-1} \leftarrow S_i \setminus \{\pi(i)\}$;

end

for $i = 1, \ldots, n$ **do**

  $X_i \leftarrow X_{i-1}$;

  if $\pi(i) \in A \cap T$ **and** $X_{i-1} \cup \{\pi(i)\} \in \mathcal{M}$ **then**

    $X_i \leftarrow X_{i-1} \cup \{\pi(i)\}$;

  end

end

return X_n;

▶ **Corollary 17.** *For any $\alpha, \varepsilon \in (0, 1]$, there exists an $(\alpha, \frac{(1-\varepsilon)\alpha^2}{4})$-universal OCRS with preselected order, which given input matroid $\mathcal{M} \subseteq 2^{[n]}$ and α-uncontentious prior distribution $\mathcal{D}_A$ for $\mathcal{M}$, runs in $O\left(\frac{n \log(n/\varepsilon)}{\alpha^2 \varepsilon^2 p_{\min}} \cdot (t_{\mathcal{D}_A} + t_{\mathcal{M}} \cdot n)\right)$ time, where $p_{\min} := \min_{i \in [n]} \Pr_{A \sim \mathcal{D}_A}[i \in A]$, and $t_{\mathcal{D}_A}$, $t_{\mathcal{M}}$ are the time it takes to generate a sample from $\mathcal{D}_A$ and to check whether a set of elements belongs to $\mathcal{M}$ respectively.*

Furthermore, as we mentioned in the introduction, Algorithm 1 (in fact, all of our universal OCRSs) generalizes ordered OCRSs for product distributions [7, 15]. For product distributions, it is known that ordered OCRSs can be strengthened to obtain *greedy* OCRSs, which work even for the worst-case arrival model [11, Theorem 2.1]. One might wonder whether we can also apply our techniques to make those greedy OCRSs universal. Unfortunately, the answer is no, because this would yield a universal OCRS for the worst-case arrival model, which does not exist even for 1-uniform matroids [8, Theorem 5.7]. However, this does not preclude the possibility that there might be meaningful relaxations of greedy OCRSs which can be made universal.

4 An improved universal OCRS via correlated subsampling

In this section, we design a universal OCRS with preselected order using correlated subsampling, which achieves a slightly stronger universality guarantee than Algorithm 1. We first introduce some notations which we will use in this section. We will use symbols π, σ, τ to represent permutations. Given any permutation $\sigma : [n] \to [n]$ and any element $e \in [n]$, we let $\mathsf{prefix}(\sigma, e)$ denote the set of elements appearing before element e in permutation σ, namely, $\mathsf{prefix}(\sigma, e) := \{\sigma(j) \mid j \in [\sigma^{-1}(e) - 1]\}$. Moreover, for any $k \in [n]$, we let σ^k denote the set of the first k elements in permutation σ, namely, $\sigma^k := \{\sigma(j) \mid j \in [k]\}$, and we let $\sigma^0 := \emptyset$. Furthermore, for any set of elements S, we let $\mathcal{P}(S)$ denote the uniform distribution over all permutations of elements in S.

4.1 Main structural lemma

The main idea of our improved universal OCRS is replacing the independent subsampling operator $\mathcal{T}_\rho$ with a correlated subsampling method[5] that is inspired by Lemma 18. Essentially, Lemma 18 says that given a matroid $\mathcal{M}$ and an uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, there exists an element i, such that if we sample a set of active elements A from $\mathcal{D}_A$ and a uniformly random permutation σ, and then remove all elements in A except those appearing before element i in permutation σ, there is a decent chance that the remaining elements do not span i conditioned on i being active. We provide the proof of Lemma 18 in the full paper.

▶ **Lemma 18.** *For any $\alpha \in [0,1]$, given any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, there exists some $i \in [n]$ such that*

$$\Pr_{A\sim\mathcal{D}_A,\,\sigma\sim\mathcal{P}([n])}[i \notin \mathrm{span}(A \cap \mathsf{prefix}(\sigma, i)) \mid i \in A] \geq \alpha.$$

We can iteratively apply Lemma 18 to determine an order of elements, which will be the preselected order of our improved universal OCRS. We formally describe this procedure in Subroutine 2 and state its guarantee in Corollary 19. The proof of Corollary 19 is provided in the full paper.

Subroutine 2 ORDER-PRESELECTING.

Input : Matroid $\mathcal{M} \subseteq 2^{[n]}$, $\alpha \geq 0$, α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$
Output : Permutation $\pi : [n] \to [n]$

$S_n \leftarrow [n]$;
for $i = n, \ldots, 1$ `// notice the reversed order`
do
  Find an element $\pi(i) \in S_i$ such that
  $\Pr_{A'\sim\mathcal{D}_A^{S_i},\,\sigma\sim\mathcal{P}(S_i)}[\pi(i) \notin \mathrm{span}_{\mathcal{M}}(A' \cap \mathsf{prefix}(\sigma, \pi(i))) \mid \pi(i) \in A'] \geq \alpha$;
  $S_{i-1} \leftarrow S_i \setminus \{\pi(i)\}$
end
return π;

▶ **Corollary 19.** *For any $\alpha \in [0,1]$, given any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, Subroutine 2 outputs a permutation $\pi : [n] \to [n]$ such that for all $i \in [n]$,*

$$\Pr_{A\sim\mathcal{D}_A,\,\sigma\sim\mathcal{P}(\pi^i)}[\pi(i) \notin \mathrm{span}_{\mathcal{M}}(A \cap \mathsf{prefix}(\sigma, \pi(i))) \mid \pi(i) \in A] \geq \alpha. \tag{4}$$

4.2 Correlated subsampling

In order to make use of Corollary 19, given a permutation π generated by Subroutine 2, we need to sample a set $T \subseteq [n]$ such that for each $i \in [n]$, the subset $\pi^{i-1} \cap T$ follows the same distribution as the random set $\mathsf{prefix}(\sigma, \pi(i))$, where $\sigma \sim \mathcal{P}(\pi^i)$. In Lemma 20, we show that this can be achieved by first sampling a permutation $\sigma' \sim \mathcal{P}([n+1])$ and then setting T as $\mathsf{prefix}(\sigma', n+1)$.

[5] The author thanks an anonymous reviewer for suggesting this method and raising valuable questions.

▶ **Lemma 20.** *Suppose that we are given a permutation $\pi : [n] \to [n]$, and we sample a permutation $\sigma' \sim \mathcal{P}([n+1])$ and let $T = \mathsf{prefix}(\sigma', n+1)$. Then, for all $i \in [n]$, the subset $T_{i-1} := \pi^{i-1} \cap T$ follows the same distribution as the random set $\mathsf{prefix}(\sigma, \pi(i))$, where $\sigma \sim \mathcal{P}(\pi^i)$. Moreover, for all $i \in [n]$ and $S \subseteq \pi^{i-1}$, we have that*

$$\Pr[\pi(i) \in T \mid T_{i-1} = S] = \frac{|S|+1}{i+1}. \tag{5}$$

We defer the proof of Lemma 20 to the full paper.

4.3 Constructing our universal OCRS with correlated subsampling

We are ready to present our improved universal OCRS with correlated subsampling (Algorithm 3). Algorithm 3 first runs Subroutine 2 to preselect an order π, and then, it samples a set T according to Lemma 20. Finally, it iterates through all elements following order π, and for each $i \in [n]$, it selects element $\pi(i)$ if $\pi(i)$ is selectable and belongs to T. We state its universality guarantee in Theorem 21, which, as we will show in the full paper, is essentially tight for this algorithm.

Algorithm 3 UNIVERSAL-OCRS-WITH-CORRELATED-SUBSAMPLING.

Input : Matroid $\mathcal{M} \subseteq 2^{[n]}$, $\alpha \geq 0$, α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, and $A \sim \mathcal{D}_A$
Output: $X_n \subseteq A$ such that $X_n \in \mathcal{M}$

Run Subroutine 2 on input $(\mathcal{M}, \alpha, \mathcal{D}_A)$, which outputs a permutation $\pi : [n] \to [n]$;
Sample a permutation $\sigma' \sim \mathcal{P}([n+1])$;
$T \leftarrow \mathsf{prefix}(\sigma', n+1)$ and $X_0 \leftarrow \emptyset$;
for $i = 1, \dots, n$ **do**
  $X_i \leftarrow X_{i-1}$;
  if $\pi(i) \in A \cap T$ **and** $X_{i-1} \cup \{\pi(i)\} \in \mathcal{M}$ **then**
    $X_i \leftarrow X_{i-1} \cup \{\pi(i)\}$;
  end
end
return X_n;

▶ **Theorem 21.** *Algorithm 3 is an $(\alpha, \frac{\alpha^2}{2})$-universal OCRS for any $\alpha \in [0,1]$.*

We briefly note that unlike Algorithm 1, which is based on independent subsampling, Algorithm 3 employs a correlated subsampling method. Therefore, in Algorithm 3, the event that element $\pi(i)$ is subsampled (i.e., $\pi(i) \in T$) is correlated with the event that element $\pi(i)$ is spanned by the partial solution set X_{i-1} (which is a subset of T). Analyzing this correlation will be the main technicality of the proof of Theorem 21, which is provided in the full paper.

5 An optimal universal OCRS via linear programming

In this section, we present an LP that computes a universal OCRS with nearly optimal universality guarantee in the preselected-order model, and we solve this LP efficiently using the ellipsoid method. In fact, this approach was originally developed by [7] to compute optimal offline CRSs for product distributions. We show that it applies naturally to ordered OCRSs for (arbitrarily correlated) uncontentious distributions.

Algorithm 4 Deterministic ordered OCRS ϕ_π.

Input : Matroid $\mathcal{M} \subseteq 2^{[n]}$, prior distribution $\mathcal{D}_A$, and $A \sim \mathcal{D}_A$
Output : $X \subseteq A$ such that $X \in \mathcal{M}$

$X \leftarrow \emptyset$;
for $i = 1, \ldots, n$ **do**
  if $\pi(i) \in A$ **and** $X \cup \{\pi(i)\} \in \mathcal{M}$ **then**
    $X \leftarrow X \cup \{\pi(i)\}$;
  end
end
return X;

We start by introducing some notations which we will use in this section. We let $\mathcal{S}_n$ denote the set of all permutations of $[n]$. For each permutation $\pi \in \mathcal{S}_n$, we let ϕ_π denote the deterministic ordered OCRS with preselected order π (Algorithm 4). Moreover, we define a permutation π_w for each weight vector $w \in \mathbb{R}^n_{\geq 0}$ as follow:

▶ **Definition 22.** *Given any weight vector $w \in \mathbb{R}^n_{\geq 0}$, we assign weight w_i to each element $i \in [n]$, and we permute elements of $[n]$ in the decreasing order of their weights (with arbitrary tie-breaking for equal weights), and then we let π_w denote this permutation.*

In the following lemma, we state a key property of OCRS ϕ_π, which follows from the well-known greedy algorithm for selecting the maximum-weight basis of a matroid [22, Chapter 19]. We defer the proof to the full paper.

▶ **Lemma 23.** *For any matroid $\mathcal{M} \subseteq 2^{[n]}$ and any prior distribution $\mathcal{D}_A \in \Delta(2^{[n]})$, for any weight vector $w \in \mathbb{R}^n_{\geq 0}$ and any permutation $\pi \in \mathcal{S}_n$, we have that*

$$\mathbb{E}_{A\sim\mathcal{D}_A}\left[\sum\nolimits_{i\in\phi_\pi(A)} w_i\right] \leq \mathbb{E}_{A\sim\mathcal{D}_A}\left[\sum\nolimits_{i\in\phi_{\pi_w}(A)} w_i\right] = \mathbb{E}_{A\sim\mathcal{D}_A}[r_{\mathcal{M},w}(A)].$$

5.1 Formulating the LP

Now we formulate an LP that computes an α-balanced OCRS for any matroid $\mathcal{M}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$. The resulting OCRS will be a random mixture of OCRSs ϕ_π for various permutations $\pi \in \mathcal{S}_n$. Specifically, we let $x_i := \Pr_{A\sim\mathcal{D}_A}[i \in A]$ and $q_{i,\pi} := \Pr_{A\sim\mathcal{D}_A}[i \in \phi_\pi(A)]$ for all $i \in [n]$ and $\pi \in \mathcal{S}_n$, and we consider the following LP (LP) and its dual (DP).

$$\begin{aligned}
\text{(LP)} \qquad & \max_{\beta,\, \lambda_\pi} \beta \\
& \text{s.t.} \sum_{\pi\in\mathcal{S}_n} q_{i,\pi}\lambda_\pi \geq \beta x_i && \forall i \in [n] \\
& \qquad \sum_{\pi\in\mathcal{S}_n} \lambda_\pi = 1 \\
& \qquad \lambda_\pi \geq 0 && \forall \pi \in \mathcal{S}_n. \\
\text{(DP)} \qquad & \min_{\gamma,\, \mu_i} \gamma \\
& \text{s.t.} \sum_{i\in[n]} q_{i,\pi}\mu_i \leq \gamma && \forall \pi \in \mathcal{S}_n
\end{aligned}$$

$$\sum_{i\in[n]} x_i\mu_i = 1$$

$$\mu_i \geq 0 \qquad \forall i \in [n]. \tag{6}$$

We note that every feasible solution (β, λ_π) to (LP) corresponds to a β-balanced OCRS with preselected order for $(\mathcal{M}, \mathcal{D}_A)$. Indeed, because of the last two constraints in (LP), variables λ_π for all $\pi \in \mathcal{S}_n$ together specify a distribution $\mathcal{D}_\pi$ over $\mathcal{S}_n$. We observe that ϕ_π with $\pi \sim \mathcal{D}_\pi$ is a randomized OCRS with preselected order for $(\mathcal{M}, \mathcal{D}_A)$, and the first constraint in (LP) ensures that this OCRS is β-balanced.

The next lemma shows that if the prior distribution $\mathcal{D}_A$ is α-uncontentious for matroid $\mathcal{M}$, then the optimal value of (DP) is at least α. By LP duality, the optimal value of (LP) is also at least α, and hence, the optimal solution to (LP) corresponds to an α-balanced OCRS with preselected order for $(\mathcal{M}, \mathcal{D}_A)$.

▶ **Lemma 24.** *For any $\alpha \in [0, 1]$, if the prior distribution $\mathcal{D}_A$ is α-uncontentious for matroid $\mathcal{M}$, then any vector $\mu \in \mathbb{R}^n$ that satisfies the last two constraints in (DP) in Eq. (6) must also satisfy that $\sum_{i\in[n]} q_{i,\pi_\mu}\mu_i \geq \alpha$, where permutation π_μ is defined in Definition 22. This implies that the optimal value of (DP) is at least α.*

Lemma 24 follows from Lemma 15 and Lemma 23. We provide the proof in the full paper.

5.2 Solving the LP efficiently

We have shown that for any matroid $\mathcal{M}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, we can find an α-balanced OCRS for $(\mathcal{M}, \mathcal{D}_A)$ by solving (LP) in Eq. (6). Solving (LP) directly is not computationally efficient because there are super-exponentially many variables λ_π. However, we can use the ellipsoid method [13] to solve (DP) in Eq. (6).

To apply the ellipsoid method to solve (DP), we need to construct a separation oracle, which given any $\gamma \in \mathbb{R}$ and $\mu \in \mathbb{R}^n_{\geq 0}$ such that $\sum_{i\in[n]} x_i\mu_i = 1$, checks whether there is a permutation $\pi \in \mathcal{S}_n$ such that $\sum_{i\in[n]} q_{i,\pi}\mu_i > \gamma$, and if so, outputs the constraint $\sum_{i\in[n]} q_{i,\pi}\mu_i \leq \gamma$. We notice that by Lemma 23, $\mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in\phi_\pi(A)} \mu_i] \leq \mathbb{E}_{A\sim\mathcal{D}_A}[\sum_{i\in\phi_{\pi_\mu}(A)} \mu_i]$ for all permutations $\pi \in \mathcal{S}_n$, which is equivalent to $\sum_{i\in[n]} q_{i,\pi}\mu_i \leq \sum_{i\in[n]} q_{i,\pi_\mu}\mu_i$ for all $\pi \in \mathcal{S}_n$. Therefore, we can implement the separation oracle efficiently as follows: Given $\gamma \in \mathbb{R}$ and $\mu \in \mathbb{R}^n_{\geq 0}$ such that $\sum_{i\in[n]} x_i\mu_i = 1$, the oracle checks whether $\sum_{i\in[n]} q_{i,\pi_\mu}\mu_i > \gamma$, and if so, outputs the constraint $\sum_{i\in[n]} q_{i,\pi_\mu}\mu_i \leq \gamma$.

Given this separation oracle, we can solve (DP) in Eq. (6) efficiently using the ellipsoid method, which identifies a polynomial number of dual constraints that certify the optimal value of (DP). Then, we can solve (LP) efficiently by restricting it to the variables that correspond to the dual constraints identified by the ellipsoid method. The only issue is that the coefficients x_i and $q_{i,\pi}$ in the LP constraints are not necessarily known. However, this can be addressed by estimating the coefficients *on demand* using Monte-Carlo sampling (i.e., we estimate coefficients x_i for all $i \in [n]$ at the beginning, and we estimate coefficients q_{i,π_μ} for all $i \in [n]$ only when the ellipsoid method queries the separation oracle with input $\mu \in \mathbb{R}^n_{\geq 0}$ and certain $\gamma \in \mathbb{R}$) and solving the approximate versions of the LPs. We defer the details to the full paper and state the guarantee of the resulting OCRS in Theorem 25.

▶ **Theorem 25.** *For any $\alpha, \varepsilon \in (0, 1]$, there exists an $(\alpha, (1-\varepsilon)\alpha)$-universal OCRS with preselected order, which given matroid $\mathcal{M} \subseteq 2^{[n]}$ and α-uncontentious prior distribution $\mathcal{D}_A$ for $\mathcal{M}$, runs in $O\left(\mathsf{poly}\left(\frac{n}{\alpha\cdot\varepsilon\cdot p_{\min}}\right)\cdot(t_{\mathcal{D}_A} + t_{\mathcal{M}})\right)$ time, where $p_{\min} := \min_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i \in A]$, and $t_{\mathcal{D}_A}$, $t_{\mathcal{M}}$ are the time it takes to generate a sample from $\mathcal{D}_A$ and to check whether a set of elements belongs to $\mathcal{M}$ respectively.*

6 From secretary algorithms to universal OCRSs (efficiently)

In this section, we show how to use linear programming to efficiently compute a universal OCRS for any arrival model, given a constant-competitive matroid secretary algorithm for that model. Briefly, in the *matroid secretary problem*, we are given a matroid $\mathcal{M} \subseteq 2^{[n]}$ and a weight vector $w \in \mathbb{R}^n_{\geq 0}$. At the beginning, a matroid secretary algorithm knows[6] only $\mathcal{M}$ but not w. Then, elements in $[n]$ arrive in a certain order according to the arrival model. Upon the arrival of each element $i \in [n]$, its weight w_i is revealed, and the algorithm must decide immediately and irrevocably whether to select the element. The goal of the algorithm is to select a set of elements $X \in \mathcal{M}$ with maximum total weight. We say that the algorithm is c-competitive if for any input matroid $\mathcal{M} \subseteq 2^{[n]}$ and weight vector $w \in \mathbb{R}^n_{\geq 0}$, it guarantees that $\mathbb{E}[\sum_{i\in X} w_i] \geq c \cdot r_{\mathcal{M},w}([n])$, where the expectation is taken over the randomness of the algorithm (and possibly the arrival model).

Given a matroid secretary algorithm ALG in any arrival model (we assume w.l.o.g. that ALG only selects elements with strictly positive weights), for each weight vector $w \in \mathbb{R}^n_{\geq 0}$, we construct an OCRS $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ in the same arrival model as ALG: Given input matroid $\mathcal{M} \subseteq 2^{[n]}$ and a set of active elements $A \subseteq [n]$, OCRS $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ provides $\mathcal{M}$ as the input matroid to algorithm ALG. Suppose that elements in $[n]$ arrive in an order $\pi : [n] \to [n]$ according to the arrival model of ALG (it is possible that π is chosen randomly and adaptively by ALG). For each $i \in [n]$, when element $\pi(i)$ arrives, OCRS $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ checks whether $\pi(i) \in A$. If $\pi(i) \in A$, it presents element $\pi(i)$ with weight $w_{\pi(i)}$ to algorithm ALG; otherwise it presents $\pi(i)$ with weight 0 to ALG. OCRS $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ selects element $\pi(i)$ if and only if algorithm ALG selects $\pi(i)$.

This OCRS was originally constructed in the proof of [8, Theorem 4.1]. We state its key property in the following lemma.

▶ **Lemma 26** ([8, Lemma 4.3]). *For any $\alpha, c \in [0,1]$, if the matroid secretary algorithm ALG is c-competitive, then given any input matroid $\mathcal{M} \subseteq 2^{[n]}$ and α-uncontentious prior distribution $\mathcal{D}_A$ for $\mathcal{M}$, for any weight vector $w \in \mathbb{R}^n_{\geq 0}$, OCRS $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ guarantees that*

$$\mathbb{E}_{A\sim\mathcal{D}_A,\phi\sim\mathcal{D}_\phi^{(\mathsf{ALG},w)}}\left[\sum\nolimits_{i\in\phi(A)} w_i\right] \geq c \cdot \alpha \cdot \mathbb{E}_{A\sim\mathcal{D}_A}\left[\sum\nolimits_{i\in A} w_i\right].$$

6.1 Formulating the LP

Now we formulate an LP that given a c-competitive matroid secretary algorithm ALG, computes a nearly $(c\cdot\alpha)$-balanced OCRS for any matroid $\mathcal{M}$ and any α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$. The resulting OCRS will be a random mixture of OCRSs $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ for various weight vectors $w \in \mathbb{R}^n_{\geq 0}$. Specifically, we let $x_i := \Pr_{A\sim\mathcal{D}_A}[i \in A]$ and $q_{i,w} := \Pr_{A\sim\mathcal{D}_A,\phi\sim\mathcal{D}_\phi^{(\mathsf{ALG},w)}}[i \in \phi(A)]$ for all $i \in [n]$ and $w \in \mathbb{R}^n_{\geq 0}$, and we define $W_\varepsilon := \left\{\frac{\varepsilon\cdot i}{n} \mid i \in \left\{0,\ldots,\left\lceil\frac{n}{\varepsilon\cdot p_{\min}}\right\rceil\right\}\right\}$, where $\varepsilon > 0$ is a parameter which we can choose arbitrarily, and $p_{\min} := \min_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i \in A]$ (we assume that $p_{\min} > 0$ as in the preliminary). We consider the following LP (LP1) and its dual (DP1).

$$\text{(LP1)} \qquad \max_{\beta,\,\lambda_w} \beta$$

[6] We note that many matroid secretary algorithms in the literature do not require full knowledge of $\mathcal{M}$ from the outset. Our result in this section also applies to these algorithms.

$$
\begin{aligned}
&\text{s.t.} \quad \sum_{w\in W_\varepsilon^n} q_{i,w}\lambda_w \geq \beta x_i && \forall i\in[n]\\
&\qquad \sum_{w\in W_\varepsilon^n} \lambda_w = 1 &&\\
&\qquad \lambda_w \geq 0 && \forall w\in W_\varepsilon^n.
\end{aligned}
$$

$$
\begin{aligned}
\text{(DP1)} \qquad &\min_{\gamma,\,\mu_i} \gamma &&\\
&\text{s.t.} \quad \sum_{i\in[n]} q_{i,w}\mu_i \leq \gamma && \forall w\in W_\varepsilon^n\\
&\qquad \sum_{i\in[n]} x_i\mu_i = 1 &&\\
&\qquad \mu_i \geq 0 && \forall i\in[n].
\end{aligned} \tag{7}
$$

We note that every feasible solution (β, λ_w) to (LP1) corresponds to a β-balanced OCRS for $(\mathcal{M}, \mathcal{D}_A)$ in the same arrival model as algorithm ALG. Indeed, because of the last two constraints in (LP1), variables λ_w for all $w \in W_\varepsilon^n$ together specify a distribution $\mathcal{D}_w$ over W_ε^n. We observe that $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ with $w \sim \mathcal{D}_w$ is an OCRS for $(\mathcal{M}, \mathcal{D}_A)$ in the same arrival model as algorithm ALG, and the first constraint in (LP1) ensures that this OCRS is β-balanced.

The next lemma shows that if the matroid secretary algorithm ALG is c-competitive, and the prior distribution $\mathcal{D}_A$ is α-uncontentious for matroid $\mathcal{M}$, then the optimal value of (DP1) is at least $(1-\varepsilon)c\cdot\alpha$. By LP duality, the optimal value of (LP1) is also at least $(1-\varepsilon)c\cdot\alpha$, and hence, the optimal solution to (LP1) corresponds to an $((1-\varepsilon)c\cdot\alpha)$-balanced OCRS for $(\mathcal{M}, \mathcal{D}_A)$ in the same arrival model as algorithm ALG.

▶ **Lemma 27.** *For any $\alpha, c, \varepsilon \in (0,1]$, if the matroid secretary algorithm ALG is c-competitive, then for any matroid $\mathcal{M} \subseteq 2^{[n]}$ and α-uncontentious prior distribution $\mathcal{D}_A$ for $\mathcal{M}$, any vector $\mu \in \mathbb{R}^n$ that satisfies the last two constraints in (DP1) must also satisfy that $\sum_{i\in[n]} q_{i,\mu'}\mu_i \geq (1-\varepsilon)c\cdot\alpha$, where $\mu' \in W_\varepsilon^n$ is defined as follows:*

$$\mu_i' = \max\{y \in W_\varepsilon \mid y \leq \mu_i\} \text{ for all } i \in [n]. \tag{8}$$

This implies that the optimal value of (DP1) is at least $(1-\varepsilon)c\cdot\alpha$.

Lemma 27 follows from Lemma 26. We provide the proof in the full paper.

6.2 Reducing the number of dual constraints

We will not directly solve (DP1) in Eq. (7) using the ellipsoid method (because here we do not have a straightforward implementation of the separation oracle). Instead, we will use the ellipsoid method to reduce the number of constraints in (DP1) such that its optimal value remains at least $(1-2\varepsilon)c\cdot\alpha$ (this technique was also used by [19] to compute optimal OCRSs for product distributions). Specifically, we consider the following polytope Q_ε:

$$Q_\varepsilon := \{\mu \in \mathbb{R}^n_{\geq 0} \mid \sum_{i\in[n]} x_i\mu_i = 1, \sum_{i\in[n]} q_{i,w}\mu_i \leq (1-2\varepsilon)c\cdot\alpha \text{ for all } w \in W_\varepsilon^n\}.$$

If $\alpha, c, \varepsilon \in (0,1]$, then by Lemma 27, any vector $\mu \in \mathbb{R}^n_{\geq 0}$ such that $\sum_{i\in[n]} x_i\mu_i = 1$ must satisfy that $\sum_{i\in[n]} q_{i,\mu'}\mu_i \geq (1-\varepsilon)c\cdot\alpha$, where $\mu' \in W_\varepsilon^n$ is defined in Eq. (8). Therefore, polytope Q_ε is empty, and moreover, we can construct an efficient separation oracle for Q_ε as follows: Given any $\mu \in \mathbb{R}^n_{\geq 0}$ such that $\sum_{i\in[n]} x_i\mu_i = 1$, the oracle outputs the violated constraint $\sum_{i\in[n]} q_{i,\mu'}\mu_i \le (1-2\varepsilon)c\cdot\alpha$ for $\mu' \in W_\varepsilon^n$ given by Eq. (8). Using this separation oracle, we can apply the ellipsoid method to identify a polynomial-size subset $W' \subseteq W_\varepsilon^n$ such that the following polytope Q'_ε is empty:

$$Q'_\varepsilon := \{\mu \in \mathbb{R}^n_{\ge 0} \mid \sum_{i\in[n]} x_i\mu_i = 1,\ \sum_{i\in[n]} q_{i,w}\mu_i \le (1-2\varepsilon)c\cdot\alpha \text{ for all } w \in W'\}.$$

Now we consider the following LP (LP2) and its dual (DP2), which are the reduced versions of (LP1) and (DP1) respectively:

$$\begin{aligned}
\text{(LP2)} \qquad & \max_{\beta,\,\lambda_w} \beta \\
& \text{s.t.} \sum_{w\in W'} q_{i,w}\lambda_w \ge \beta x_i && \forall i \in [n] \\
& \qquad \sum_{w\in W'} \lambda_w = 1 \\
& \qquad \lambda_w \ge 0 && \forall w \in W'. \\
\text{(DP2)} \qquad & \min_{\gamma,\,\mu_i} \gamma \\
& \text{s.t.} \sum_{i\in[n]} q_{i,w}\mu_i \le \gamma && \forall w \in W' \\
& \qquad \sum_{i\in[n]} x_i\mu_i = 1 \\
& \qquad \mu_i \ge 0 && \forall i \in [n].
\end{aligned} \tag{9}$$

Because polytope Q'_ε is empty, the optimal value of (DP2) is greater than $(1-2\varepsilon)c\cdot\alpha$, and by LP duality, the optimal value of (LP2) is also greater than $(1-2\varepsilon)c\cdot\alpha$.

Furthermore, note that (LP2) has polynomially many variables and constraints, and hence, its optimal solution, which we denote by (β^*, λ_w^*), can be computed in polynomial time. By the last two constraints in (LP2), variables λ_w^* for all $w \in W'$ together specify a distribution $\mathcal{D}_w^*$ over W'. We observe that $\mathcal{D}_\phi^{(\mathsf{ALG},w)}$ with $w \sim \mathcal{D}_w^*$ is an OCRS for $(\mathcal{M}, \mathcal{D}_A)$ in the same arrival model as the matroid secretary algorithm ALG. This OCRS is β^*-balanced because of the first constraint in (LP2). Thus, it is $((1-2\varepsilon)c\cdot\alpha)$-balanced since $\beta^* > (1-2\varepsilon)c\cdot\alpha$.

Finally, similar to Subsection 5.2, here we also need to estimate the coefficients x_i and $q_{i,w}$ using Monte-Carlo sampling because they are not necessarily known. We defer the details to the full paper and summarize the result in Theorem 28.

▶ **Theorem 28.** *For any $\alpha, c, \varepsilon \in (0,1]$, for any arrival model, if there is a c-competitive matroid secretary algorithm ALG, then there exists an $(\alpha, (1-\varepsilon)c\cdot\alpha)$-universal OCRS, which given input matroid $\mathcal{M} \subseteq 2^{[n]}$ and α-uncontentious distribution $\mathcal{D}_A$ for $\mathcal{M}$, runs in $O\left(\mathsf{poly}\left(\frac{n}{\alpha\cdot c\cdot\varepsilon\cdot p_{\min}}\right)\cdot(t_{\mathsf{ALG}} + t_{\mathcal{D}_A})\right)$ time, where $p_{\min} := \min_{i\in[n]} \Pr_{A\sim\mathcal{D}_A}[i \in A]$, and t_{ALG} is the worst-case runtime of algorithm ALG on matroid secretary problem instances specified by matroid $\mathcal{M}$ and weight vectors $w \in W_\varepsilon^n$, and $t_{\mathcal{D}_A}$ is the time it takes to generate a sample from $\mathcal{D}_A$.*

References

1 Marek Adamczyk and Michał Włodarczyk. Random order contention resolution schemes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 790–801. IEEE, 2018.

2 Shipra Agrawal, Jay Sethuraman, and Xingyu Zhang. On optimal ordering in the optimal stopping problem. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 187–188, 2020. doi:10.1145/3391403.3399484.

3 Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Symposium on Discrete Algorithms (SODA'07)*, pages 434–443, 2007. URL: http://dl.acm.org/citation.cfm?id=1283383.1283429.

4 Hedyeh Beyhaghi, Negin Golrezaei, Renato Paes Leme, Martin Pal, and Balasubramanian Sivan. Improved approximations for free-order prophets and second-price auctions. *arXiv preprint arXiv:1807.03435*, 2018. arXiv:1807.03435.

5 Archit Bubna and Ashish Chiplunkar. Prophet inequality: Order selection beats random order. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 302–336, 2023. doi:10.1145/3580507.3597687.

6 Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010. doi:10.1145/1806689.1806733.

7 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014. doi:10.1137/110839655.

8 Shaddin Dughmi. The outer limits of contention resolution on matroids and connections to the secretary problem. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2020.

9 Shaddin Dughmi. Matroid secretary is equivalent to contention resolution. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

10 Shaddin Dughmi, Yusuf Hakan Kalayci, and Neel Patel. Limitations of stochastic selection problems with pairwise independent priors. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 479–490, 2024. doi:10.1145/3618260.3649718.

11 Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM Journal on Computing*, 50(2):255–300, 2021. doi:10.1137/18M1226130.

12 Hu Fu, Pinyan Lu, Zhihao Gavin Tang, Hongxun Wu, Jinzhao Wu, and Qianfan Zhang. Sample-based matroid prophet inequalities. In *Proceedings of the 25th ACM Conference on Economics and Computation*, 2024.

13 Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric algorithms and combinatorial optimization*, volume 2. Springer Science & Business Media, 2012.

14 Anupam Gupta, Jinqiao Hu, Gregory Kehne, and Roie Levin. Pairwise-independent contention resolution. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 196–209. Springer, 2024. doi:10.1007/978-3-031-59835-7_15.

15 Anupam Gupta and Viswanath Nagarajan. A stochastic probing problem with applications. In *Integer Programming and Combinatorial Optimization: 16th International Conference, IPCO 2013, Valparaíso, Chile, March 18-20, 2013. Proceedings 16*, pages 205–216. Springer, 2013. doi:10.1007/978-3-642-36694-9_18.

16 Mohammad Taghi Hajiaghayi, Dariusz R Kowalski, Piotr Krysta, and Jan Olkowski. Optimal algorithms for free order multiple-choice secretary. *arXiv preprint arXiv:2207.10703*, 2022. doi:10.48550/arXiv.2207.10703.

17 TP Hill. Prophet inequalities and order selection in optimal stopping problems. *Proceedings of the American Mathematical Society*, 88(1):131–137, 1983.

18 Patrick Jaillet, José A Soto, and Rico Zenklusen. Advances on matroid secretary problems: Free order model and laminar case. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 254–265. Springer, 2013. doi:10.1007/978-3-642-36694-9_22.

19 Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018.

20 Allen Liu, Renato Paes Leme, Martin Pál, Jon Schneider, and Balasubramanian Sivan. Variable decomposition for prophet inequalities and optimal ordering. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 692–692, 2021.

21 Bo Peng and Zhihao Gavin Tang. Order selection prophet inequality: From threshold optimization to arrival time design. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 171–178. IEEE, 2022. `doi:10.1109/FOCS54457.2022.00023`.

22 Dominic JA Welsh. *Matroid theory*. Courier Corporation, 2010.