

## Online Selection of Diverse Committees

Virginie Do<sup>1,2\*</sup>, Jamal Atif<sup>1</sup>, Jérôme Lang<sup>1</sup> and Nicolas Usunier<sup>2</sup>

<sup>1</sup>LAMSADE, Université PSL, Université Paris-Dauphine, CNRS, France

<sup>2</sup>Facebook AI Research

virginiedo@fb.com, jamal.atif@lamsade.dauphine.fr, lang@lamsade.dauphine.fr, usunier@fb.com

### Abstract

Citizens' assemblies need to represent subpopulations according to their proportions in the general population. These large committees are often constructed in an online fashion by contacting people, asking for the demographic features of the volunteers, and deciding to include them or not. This raises a trade-off between the number of people contacted (and the incurring cost) and the representativeness of the committee. We study three methods, theoretically and experimentally: a greedy algorithm that includes volunteers as long as proportionality is not violated; a non-adaptive method that includes a volunteer with a probability depending only on their features, assuming that the joint feature distribution in the volunteer pool is known; and a reinforcement learning based approach when this distribution is not known a priori but learnt online.

### 1 Introduction

Forming a representative committee consists in selecting a set of individuals, who agree to serve, in such a way that every part of the population, defined by specific features, is represented proportionally to its size. As a paradigmatic example, the Climate Assembly in the UK and the Citizens' Convention for Climate in France brought together 108 and 150 participants respectively, representing sociodemographic categories such as gender, age, education level, professional activity, residency, and location, in proportion to their importance in the wider society. Beyond citizens' deliberative assemblies, proportional representation often has to be respected when forming an evaluation committee, selecting a diverse pool of students or employees, and so on.

Two key criteria for evaluating the committee formation process are the representativeness of the final selection and the number of persons contacted (each of these incurring a cost). The trade-off is that the higher the number of people contacted, the more proportional the resulting committee.

A first possibility is to use an offline strategy (as for the UK assembly): invitations are sent to a large number of peo-

ple (30,000), and the final group is selected among the pool of volunteers. An alternative setting which is common in hiring is to consider an online process: the decision-maker is given a stream of candidates and has to decide at each timestep whether or not to admit the candidate to the final committee. This work focuses on the latter setting.

A further difficulty is that the distribution of *volunteers* is not necessarily known in advance. For example, although the target is to represent distinct age groups proportionally to their distribution in the wider population, it may be the case that older people are predominant among volunteers.

Multi-attribute proportional representation in committee selection in an off-line setting usually assumes full access to a finite (typically large) database of candidates. This assumption is impractical in a variety of real-world settings: first, the database does not exist beforehand and constructing it would require contacting many more people than necessary; second, in some domains, the decision to hire someone should be made immediately so that people don't change their mind in the meantime (which is typical in professional contexts).

An online strategy must achieve a good trade-off between sample complexity, i.e. the number of timesteps needed to construct a full committee, and the quality of the final committee, as measured by its distance to the target distribution.

We focus on the online setting. We introduce a new model and offer three different strategies, which rely on different assumptions on the input (and the process). The *greedy* strategy selects volunteers as long as their inclusion does not jeopardize the size and representation constraints; it does not assume any prior distribution on the volunteer pool. The *nonadaptive* strategy, based on constrained Markov decision processes, repeatedly chooses a random person, and decides whether to include or not a volunteer with a probability that depends only on their features; it assumes the joint distribution in the volunteer pool is known; it can be parallelised. Finally, the *reinforcement learning* strategy assumes this distribution is not known a priori but can be learnt online.

Which of these strategies are interesting depends on domain specificities. For each, we study bounds for expected quality and sample complexity, and perform experiments using real data from the UK Citizens' Assembly on Brexit.

The outline of the paper is as follows. We discuss related work in Section 2, define the problem in Section 3, define and study our three strategies in Sections 3.2, 4 and 5, analyse our

\*Contact Author.

†Full version available at <https://arxiv.org/abs/2105.09295>.

experiments in Section 6 and conclude in Section 7.

## 2 Related Work

### Diversity and representation in committee (s)election.

The problem of selecting a diverse set of candidates from a candidate database, where each candidate is described by a vector of attribute values, has been considered in several places. In [Lang and Skowron, 2018], the goal is to find a committee of a fixed size whose distribution of attribute values is as close as possible to a given target distribution. In [Celis *et al.*, 2018; Bredereck *et al.*, 2018], each candidate has a score, obtained from a set of votes, and some constraints on the proportion of selected candidates with a given attribute value are specified; the goal is to find a fixed-size committee of maximal score satisfying the constraints. In the same vein, [Aziz, 2019] considers soft constraints, and [Bei *et al.*, 2020] do not require the size of the committee to be fixed.<sup>1</sup>

Our online setting shifts the difficulty of the multi-attribute representation problem from computational complexity analyses, to the need for probabilistic guarantees on the tradeoffs between sample complexity and achieved proportionality.

**Representative and fair sortition.** Finding a representative committee (typically, a panel of citizens) with respect to a set of attributes, using *sortition*, is the topic of at least two recent papers. [Benadè *et al.*, 2019] show that stratification (random selection from small subgroups defined by attribute values, rather than from the larger group) only helps marginally. [Flanigan *et al.*, 2020] go further and consider this three-stage selection process: (1) letters are sent to a large number of random individuals (the *recipients*); (2) these recipients answer whether they agree to participate, and if so, give their features; those individuals constitute the *pool*; (3) a sampling algorithm is used to select the final *panel* from the pool. As the probability of willingness to participate is different across demographic groups, each person is selected with a probability that depends on their features, so as to correct this self-selection bias. This guarantees that the whole process be fair to all individuals of the population, with respect of going from the initial population to the panel.<sup>2</sup>

The main differences between this work and ours are: (1) (once again) our process is online; (2) we do not consider individual fairness, only group representativeness; (3) we care about minimizing the number of people contacted. Moreover, unlike off-line processes, our process can be applied in contexts where hiring a person just interviewed cannot be delayed; this may not be crucial for citizens’ assemblies (although someone who volunteers at first contact may change their mind if the delay until the final selection is long), but this is definitely so when hiring a diverse team of employees.

<sup>1</sup>Note that *diversity* and *proportional representation* are often used with a different meaning in multiwinner elections, namely, in the sense that each voter should feel represented in an elected committee, regardless of attributes. A good entry to this literature is the survey [Faliszewski *et al.*, 2017].

<sup>2</sup>Fairness guarantees are pushed further in following (yet unpublished) work by the authors: see <https://youtu.be/x.1Ce1kT7vc>.

**Online selection problems.** Generalized secretary problems [Babaioff *et al.*, 2008] are optimal stopping problems where the goal is to hire the best possible subset of persons, assuming that persons arrive one at a time, their value is observed at that time, and the decision to hire or not them must be taken immediately. The problem has been generalized to finding a set of items maximizing a submodular value function [Bateni *et al.*, 2013; Badanidiyuru *et al.*, 2014] While the latter models do not deal with diversity constraints, [Stoyanovich *et al.*, 2018] aims at selecting a group of people arriving in a streaming fashion from a finite pool, with the goal of optimizing their overall quality subject to diversity constraints. The common point with our approach is the online nature of the selection process. The main differences are that they consider only one attribute, the size of the pool is known, and yet more importantly, what is optimized is the intrinsic quality values of the candidates and not the number of persons interviewed. Closer to our setting is [Panigrahi *et al.*, 2012] who consider diversity along multiple features in online selection of search results, regardless of item quality. They only seek to maximise diversity, and do not consider trade-offs with the number of items observed.

The diverse hiring setting of [Schumann *et al.*, 2019] is very different. At each time step, the decision-maker chooses which candidate to interview and only decides on which subset to hire after multiple rounds, whereas in our setting, candidates arrive one by one and decisions are made immediately.

## 3 Formal Setting

### 3.1 Problem Definition

Let  $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_d$  be the product space of  $d$  finite domains, each of size  $D_i = |\mathcal{X}_i|$ , and where we identify  $\mathcal{X}_i$  with  $[D_i] = \{1, \dots, D_i\}$ . Each candidate is represented by a *characteristic vector*  $x \in \mathcal{X}$  with  $d$  features. Let  $x^i \in \mathcal{X}_i$  denote the value of the  $i$ -th feature. For each  $i \in [d]$ , we consider a *target vector*  $\rho^i \in (0, 1)^{D_i}$  with  $\sum_{j=1}^{D_i} \rho_j^i = 1$ .

The candidate database is infinite and the horizon as well. At each timestep  $t \geq 1$ , the agent observes a candidate  $x_t$  drawn i.i.d. from a stationary distribution  $p$  over  $\mathcal{X}$ , i.e.  $x_t \sim p$ . The decision-maker must immediately decide between two actions: *accept* or *reject* the candidate, which we respectively denote as  $a_t = 1$  and  $a_t = 0$ .

The goal is to select a *committee*  $C$  of  $K$  candidates that matches the target vectors as closely as possible, while minimizing the number of candidates screened.

For some set  $C$ , let  $\lambda(C) \in \prod_{i=1}^d [0, 1]^{D_i}$  be the *representation profile* of  $C$ , where  $\lambda_j^i(C) = \frac{|\{x \in C : x^i = j\}|}{|C|}$ . We define the *representation loss* as  $\|\lambda(C) - \rho\|_\infty = \max_{i \in [d], j \in [D_i]} |\lambda_j^i(C) - \rho_j^i|$ . We evaluate how much  $C$  matches the target  $\rho$  by the  $\ell_\infty$  metric, because it is harsher than  $\ell_1, \ell_2$  on committees that are unacceptable in our applications (e.g. committees with no women that achieve perfect representation on all other categories than gender).

Let  $C_t = \{x_{t'} : t' \leq t, a_{t'} = 1\}$  denote the set of all accepted candidates at the end of step  $t$ . The agent stops at  $\tau$ , where  $\tau$  is the first time when  $K$  candidates have been accepted, i.e. the total number of candidates screened.

gender \ age	S	J
M	$1/2 - \epsilon'$	$1/4$
F	$1/4$	$\epsilon'$

 Table 1: Example candidate distribution  $p$  with 2 binary features.

The agent following a (possibly randomized) algorithm  $ALG$  must minimize the *sample complexity*  $\mathbb{E}^{p, ALG}[\tau]$ .

Importantly, we consider two settings: whether the candidate distribution  $p$  is *known* or *unknown*.

*Remark.* In this model, we simply ignore non-volunteers, since the agent only needs to make decisions for volunteers, which from now on we call *candidates*. The joint distribution of characteristic vectors in the population of candidates is  $p$ .

### 3.2 Greedy Strategy

We describe a first simple strategy. In **Greedy**, the agent greedily accepts any candidate as long as the number of people in the committee with  $x^i = j$  does not exceed the quota  $\lceil \rho_j^i K \rceil + \frac{\epsilon K}{(D_i - 1)}$  for any  $i, j$ , where  $\epsilon > 0$  is some tolerance parameter for the representation quality.

**Proposition 1.** *The representation loss incurred by Greedy is bounded as follows:*

$$\|\lambda(C_\tau) - \rho\|_\infty \stackrel{a.s.}{\leq} \left( \frac{\max_{i \in [d]} D_i - 1}{K} + \epsilon \right).$$

This method is simple to interpret and implement, and can even be used when the candidate distribution  $p$  is unknown. However, in the following example, we see that **Greedy** may be inefficient because it requires interacting with an arbitrarily large number of candidates to recruit a full committee.

**Example 1.** *Let  $\epsilon' > 0, \ll 1$ . There are 2 binary features, gender and age, with domains  $\mathcal{X}_{\text{gender}} = \{M, F\}$  and  $\mathcal{X}_{\text{age}} = \{S, J\}$ . The candidates are distributed as  $p$  given in Table 1. We want a committee of size  $K = 4$  (e.g., a thesis committee) and the target is  $\rho^{\text{gender}} = (1/2, 1/2)$  and  $\rho^{\text{age}} = (3/4, 1/4)$ .*

*Let  $A$  be the event that in the first 3 timesteps, the agent observes candidates with characteristic vectors  $\{FS, MS, MS\}$  in any order. Then **Greedy** accepts all of them, i.e.  $A = \{C_3 = \{FS, MS, MS\}\}$ . We have:  $\mathbb{P}[A] = 1/4(1/2 - \epsilon')^2 \times 3! = 3/2(1/2 - \epsilon')^2 \geq 3/2(1/3)^2 = 1/6$ .*

*Under event  $A$ , **Greedy** can only stop upon finding  $FJ$  in order to satisfy the representation constraints. Therefore,  $\tau|A$  follows a geometric distribution with success probability  $\epsilon'$ , hence its expectation is  $1/\epsilon'$ , and  $\mathbb{E}^{p, \text{Greedy}}[\tau] \geq \mathbb{E}[\tau|A] \times \mathbb{P}[A] = 1/6\epsilon'$ . Therefore, the sample complexity of **Greedy** in this example is arbitrarily large.*

This example shows the limits of directly applying a naive strategy to our online selection problem, where the difficulty arises from considering multiple features simultaneously, even when there are only 2 binary features. We further discuss the strengths and weaknesses of **Greedy**, and its sensitivity to the tolerance  $\epsilon$  in our experiments in Section 6.

The greedy strategy is adaptive, in the sense that decisions are made based on the current candidate and candidates accepted in the past. In the following section, we present, with theoretical guarantees, an efficient yet non-adaptive algorithm

based on constrained MDPs for the setting in which the candidate distribution is known. We then adapt this approach to the case when this distribution is unknown, using techniques for efficient exploration / exploitation in constrained MDPs relying on the principle of optimism in the face of uncertainty.

## 4 $p$ Is Known: Constrained MDP Strategy

In this section, we assume the distribution  $p$  is known, and we place ourselves in the limit where we would select a committee of infinite size, and aim to maximize the rate at which candidates are selected, under the constraint that the proportion of accepted candidates per feature value is controlled by  $\rho$ . One advantage of this approximation is that the optimal policy is stationary, thus simple to represent. Moreover, as stationary policies can be very well parallelized, in the case where multiple candidates can be interviewed simultaneously. To apply this approach to the finite-size committee selection problem, one needs to interrupt the agent when  $K$  candidates have been selected. We showcase a high probability bound of  $O(\sqrt{1/K})$  on the representation loss, which guarantees that for large enough values of  $K$ , the resulting committee is representative.

From now on, we assume that any feature vector can be observed, i.e.,  $p(x) > 0$  for all  $x$ , so that proportional representation constraints can be satisfied.

### 4.1 Our Model

Fundamentally, our problem could be seen as a contextual bandit with stochastic contexts  $x_t \sim p$  and two actions  $a_t = 0$  or 1. However, the type of constraints incurred by proportional representation are well studied in constrained MDPs (CMDPs) [Altman, 1999], whereas the contextual bandits literature focused on other constraints (e.g., knapsack constraints [Agrawal and Devanur, 2016]). We show how we can efficiently leverage the CMDP framework for our online committee selection problem.

Formally, we introduce an MDP  $M = (\mathcal{X}, \mathcal{A}, P, r)$ , where the set of states is the  $d$ -dimensional candidate space  $\mathcal{X}$ , the set of actions is  $\mathcal{A} = \{0, 1\}$ , and the (deterministic) reward is  $r(x, a) = \mathbb{1}_{\{a=1\}}$ . The transition kernel  $P$ , which defines the probability to be in state  $x'$  given that the previous state was  $x$  and the agent took action  $a$ , is very simple in our case: we simply have  $P(x'|x, a) = p(x')$  since candidates are drawn i.i.d regardless of the previous actions and candidates.

We consider the *average reward* setting in which the performance of a policy  $\pi : \mathcal{X} \times \mathcal{A} \rightarrow [0, 1]$  is measured by its *gain*  $g^{p, \pi}$ , defined as:

$$g^{p, \pi}(x) = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^{p, \pi} \left[ \sum_{t=1}^T r(x_t, a_t) \mid x_1 = x \right].$$

We simply write  $g^{p, \pi} := g^\pi$  when the underlying transition is  $p$  without ambiguity.

We include proportional representation constraints following the framework of CMDPs, where the set of allowed policies is restricted by a set of additional constraints specified by reward functions. In our case, for  $i \in [d], j \in [D_i]$ , we introduce  $r_j^i(x, a) = \mathbb{1}_{\{x^i=j, a=1\}}$ , and let  $\xi_j^i = r_j^i - \rho_j^i r$  be the

reward function for the constraint indexed by  $i, j$ . Similarly to the gain, we define  $h_j^{i\pi} = \lim_{T \rightarrow \infty} \frac{1}{T} \mathbb{E}^\pi \left[ \sum_{t=1}^T \xi_j^i(x_t, a_t) \right]$ . The CMDP is defined by:

$$\max_{\pi} \{g^\pi \mid \forall i \in [d], \forall j \in [D_i], h_j^{i\pi} = 0\}. \quad (1)$$

Given the simplicity of the transition kernel, and since the MDP is ergodic by the assumption  $p > 0$ , the gain is constant, i.e.  $\forall x \in \mathcal{X}, g^\pi(x) = g^\pi$ , and problem (1) is well defined. From now on, we only write  $g^\pi$  and  $\xi_j^{i\pi}$ . Moreover, the optimal policy for the CMDP (1) is denoted  $\pi^*$  and is *stationary* [Altman, 1999].

**Lemma 1.**  $g^\pi$  is the selection rate under policy  $\pi$ :

$$g^\pi = \sum_x p(x) \pi(x, 1) = \mathbb{P}^{p, \pi}[a = 1]$$

Moreover, if  $\pi$  is feasible for CMDP (1), then:

$$\forall i \in [d], \forall j \in [D_i], \mathbb{P}^{p, \pi}[x^i = j \mid a = 1] = \rho_j^i.$$

Lemma 1 implies that (a)  $\pi^*$  maximises the selection rate of candidates, and (b) the constraints of (1) force candidates  $x$  with  $x^i = j$  to be accepted in proportions given by  $\rho_j^i$ .

The CMDP can be expressed as the linear program:

$$\begin{aligned} & \max_{\pi \in \mathbb{R}_+^{\mathcal{X} \times \mathcal{A}}} \sum_{x, a} \pi(x, a) p(x) r(x, a) \\ \text{u.c.} \quad & \forall x \in \mathcal{X}, \sum_a \pi(x, a) = 1 \\ & \forall i, j, \sum_{x, a} \pi(x, a) p(x) \xi_j^i(x, a) = 0. \end{aligned} \quad (2)$$

Notice that problem (2) is feasible by the assumption that  $\forall x \in \mathcal{X}, p(x) > 0$ . Next we study how well the proportional selection along features is respected when we shift from infinite to finite-sized committee selection.

## 4.2 Theoretical Guarantees

We analyze the CMDP-based strategy where at each timestep, the agent observes candidates  $x_t \sim p$ , decides to accept  $x_t$  by playing  $a_t \sim \pi^*(\cdot \mid x_t)$  and stops when  $K$  candidates have been accepted. We later refer to it as CMDP for brevity.

First, we formally relate the gain  $g^\pi$  that we optimize for in (1) to the quantity of interest  $\mathbb{E}^{p, \pi}[\tau]$ .

**Lemma 2.** For any stationary policy  $\pi$ ,  $\mathbb{E}^{p, \pi}[\tau] = \frac{K}{g^\pi}$ .

Lemma 2 is a direct consequence of the fact that  $\tau + K$  follows a negative binomial distribution with parameters  $K$  and  $1 - g^\pi$ , which are respectively the number of successes and the probability of failure, i.e. of rejecting a candidate under  $\pi$ . Note that this is only true because in our case the transition structure of the MDP ensures constant gain. A quick sanity check shows that if the agent systematically accepts all candidates, i.e.  $g^\pi = 1$ , then  $\mathbb{E}^{p, \pi}[\tau] = K$ , and that maximizing  $g^\pi$  is equivalent to minimizing  $\mathbb{E}^{p, \pi}[\tau]$ .

We exhibit a bound on the representation loss of CMDP which follows the optimal stationary policy  $\pi^*$  of CMDP (1). Let  $\tilde{d} = \sum_{i=1}^d (D_i - 1)$ . ( $\tilde{d} = d$  when all features are binary.)

---

### Algorithm 1: RL-CMDP algorithm.

---

**input :** confidence  $\delta$ , committee size  $K$ , targets  $\rho$   
**output:** committee  $C_\tau$

```

1  $t \leftarrow 0, C_0 \leftarrow \emptyset;$ 
2 while  $|C_t| < K$  do
3   for episode  $l = 1, 2, \dots$  do
4      $\tau_l = t + 1;$ 
5      $\pi_l \leftarrow \text{sol. of (4) via the extended LP (5);}$ 
6     while  $n_t(x_t) < 2n_{\tau_l-1}(x_t)$  do
7        $t \leftarrow t + 1, \text{Execute } \pi_l;$ 
8     end
9   end
10 end
11 return  $C_t$ 
```

---

**Proposition 2.** Let  $\pi^*$  be an optimal stationary policy for CMDP (1). Let  $\delta > 0$ . Then,

$$\mathbb{P}^{p, \pi^*} \left[ \|\lambda(C_\tau) - \rho\|_\infty \leq \sqrt{\frac{\log(\frac{2\tilde{d}}{\delta})}{2K}} \right] \geq 1 - \delta.$$

All proofs of this section are available in Appendix ??.

The upper bound on the representation loss of CMDP decreases with the committee size in  $\sqrt{1/K}$ . This shows that the stationary policy  $\pi^*$  works well for larger committees, although it acts independently from previously accepted candidates. The intuition is that for larger committees, adding a candidate has less impact on the current representation vector.

**Example 2.** We take the same attributes and same distribution as in Table 1, with  $\epsilon' = 1/6$ . Here, the target vectors are  $\rho^{\text{gender}} = (1/2, 1/2)$  and  $\rho^{\text{age}} = (1/2, 1/2)$ : an ideal committee contains as many women as men, as many senior as junior.

With the optimal policy for LP (2), each time the current volunteer is a senior male, we select him with probability  $1/2$ ; all other volunteers are selected with probability 1. The expected final composition of the pool is 30% of junior male, 30% of senior female, 20% of junior female and 20% of senior male. As the policy selects in average  $3/5$  of the volunteers, the expected time until we select  $K$  candidates is  $\mathbb{E}^{p, \pi^*}[\tau] = (6/5)K$ .

## 5 $p$ Is Unknown: Optimistic CMDP Strategy

We now tackle the committee selection problem when the candidate distribution  $p$  is unknown and must be learned online. Let  $g^* = g^{\pi^*}$  be the value of (1), which is the optimal gain of the CMDP when the distribution  $p$  is known. We evaluate a learning algorithm by:

1. the performance regret:  $R(T) = \sum_{t=1}^T (g^* - r(x_t, a_t))$ ,
2. the cost of constraint violations:  
 $R^c(T) = \max_{i, j} \left| \sum_{t=1}^T \xi_j^i(x_t, a_t) \right|$ .

We propose an algorithm that we call RL-CMDP (Reinforcement Learning in CMDP, Alg. 1). It is an adaptation of the optimistic algorithm UCRL2 [Jaksch *et al.*, 2010], and it also builds on the algorithm OptCMDP proposed by [Efroni *et al.*,

2020] for finite-horizon CMDPs. Learning in average-reward CMDPs involves different challenges, because there is no guarantee that the policy at each episode has constant gain. It does not matter in our case, since as we noted in Sec. 4, the simple structure of the transition kernel ensures constant gain, and does not require to use the Bellman equation. The few works on learning in average-reward CMDPs make unsuitable assumptions for our setting [Zheng and Ratliff, 2020; Singh *et al.*, 2020].

RL-CMDP proceeds in episodes, which end each time the number of observations for some candidate  $x$  doubles. During each episode  $l$ , observed candidates  $x_t$  are accepted on the basis of a single stationary policy  $\pi_l$ .

Let  $\tau_l$  denote the start time of episode  $l$  and  $E_l = [\tau_l, \tau_{l+1}]$ . Let  $n_t(x) = \sum_{t'=1}^t \mathbb{1}_{\{x_{t'}=x\}}$  and  $N(t) = |C_{t-1}| = \sum_{t'=1}^{t-1} \mathbb{1}_{\{a_{t'}=1\}}$ . Let  $N_j^i(t) = \sum_{t'=1}^{t-1} \mathbb{1}_{\{x_{t'}^i=j, a_{t'}=1\}}$  be the number of accepted candidates  $x$  such that  $x^i = j$  before  $t$ .

At each episode  $l$ , the algorithm estimates the true candidate distribution by the empirical distribution  $\hat{p}_l(x) = \frac{n_{\tau_l-1}(x)}{\tau_l-1}$  and maintains confidence sets  $B_l$  on  $p$ . As in UCRL2, these are built using the inequality on the  $\ell_1$ -deviation of  $p$  and  $\hat{p}_l$  from [Weissman *et al.*, 2003]:

**Lemma 3.** *With probability  $\geq 1 - \frac{\delta}{3}$ ,*

$$\|\hat{p}_l - p\|_1 \leq \sqrt{\frac{2|\mathcal{X}| \log(6|\mathcal{X}|\tau_l(\tau_l-1)/\delta)}{\tau_l-1}} := \beta_l \quad (3)$$

Let  $B_l = \{\tilde{p} \in \Delta(\mathcal{X}) : \|\hat{p}_l - \tilde{p}\|_1 \leq \beta_l\}$  be the confidence set for  $p$  at episode  $l$ . The associated set of compatible CMDPs is then  $\{M = (\mathcal{X}, \mathcal{A}, \tilde{p}, r, \xi) : \tilde{p} \in B_l\}$ . At the beginning of each episode, RL-CMDP finds the optimum of:

$$\max_{\pi \in \Pi, \tilde{p} \in B_l} \{g^{\tilde{p}, \pi} \mid \forall i, j, h_j^{i, \tilde{p}, \pi} = 0\}. \quad (4)$$

**Extended LP.** In order to optimize this problem, we rewrite (4) as an extended LP. Following [Rosenberg and Mansour, 2019] and the CMDP literature, we introduce the state-action occupation measure  $\mu(x, a) = \pi(x, a)p(x)$  and variables  $\beta(x)$  to linearize the  $\ell_1$  constraint induced by the confidence set:

$$\begin{aligned} & \max_{\substack{\mu \in \mathbb{R}^{\mathcal{X} \times \mathcal{A}} \\ \beta \in \mathbb{R}^{\mathcal{X}}}} \sum_{x, a} \mu(x, a) r(x, a) \\ \text{u.c.} \quad & \mu \geq 0, \sum_{x, a} \mu(x, a) = 1 \\ & \forall x, \sum_a \mu(x, a) \leq \hat{p}_l(x) + \beta(x) \\ & \forall x, \sum_a \mu(x, a) \geq \hat{p}_l(x) - \beta(x) \\ & \forall x, a, \sum_y \beta(y) \leq \mu(x, a) \beta_l \\ & \forall i, j, \sum_{x, a} \mu(x, a) \xi_j^i(x, a) = 0. \end{aligned} \quad (5)$$

The last constraint is the proportional representation constraint. The second to fourth constraints enforce the compatibility of  $\mu$  with the  $\ell_1$  confidence set. We retrieve the distribution as  $\hat{p}_l(x) = \sum_a \mu(x, a)$ , and the policy as:

$$\pi_l(x, a) = \begin{cases} \frac{\mu(x, a)}{\hat{p}_l(x)} & \text{if } \hat{p}_l(x) \neq 0 \\ \frac{1}{2} & \text{otherwise.} \end{cases}$$

Precisely, if some  $\hat{p}_l(x) = 0$ , we may set the policy  $\pi_l(a|x)$  arbitrarily. Since the MDP induced by  $\tilde{p}$  is still weakly communicating, and in particular any policy is unichain, the optimal gain in this CMDP is not affected.

We now provide regret and representativeness guarantees.

**Theorem 1.** *With probability  $\geq 1 - \delta$ , the regret of RL-CMDP satisfies:*

$$\begin{aligned} R(T) &= O(\sqrt{|\mathcal{X}|T \log(|\mathcal{X}|T/\delta)}) \\ R^c(T) &= O(\sqrt{|\mathcal{X}|T \log(|\mathcal{X}|T/\delta)}). \end{aligned}$$

*Moreover, with probability  $1 - \delta$ , the representation loss of RL-CMDP at horizon  $T$  satisfies:*

$$\|\lambda(C_T) - \rho\|_\infty = O\left(\frac{1}{g^*} \sqrt{\frac{|\mathcal{X}| \log(|\mathcal{X}|T/\delta)}{T}}\right).$$

It relies on decomposing regret over episodes, bounding the error on  $p$  which decreases over episodes as the confidence sets are refined, and leveraging martingale inequalities on the cumulative rewards.

Since  $\frac{R(T)}{T} = g^* - \frac{N(T)}{T}$ , it means that with high probability, the difference between the optimal selection rate and the selection rate of RL-CMDP decreases in  $\sqrt{\log(T)/T}$  w.r.t. the horizon  $T$ . The representation loss decreases at the same speed, meaning that the agent should see enough candidates to accurately estimate  $p$ , and accept candidates at little cost for representativeness.

Compared to the bound from Proposition 2, the cost of not knowing  $p$  on representativeness is a  $\sqrt{|\mathcal{X}| \log(|\mathcal{X}|)}$  factor. This is due to the estimation of  $p$  in the worst case, which is controlled by Lemma 3. As we show in our experiments (Sec. 6), the impact of  $|\mathcal{X}|$  on performance regret (and in turn on sample complexity) is not problematic in our typical citizens' assembly scenario: since there are only a handful of features, our algorithm selects candidates quickly in practice (though representativeness is weakened by not knowing  $p$ ). For specific structures of  $p$ , we obtain bounds with better scaling in  $|\mathcal{X}|$ , by controlling each entry of  $p$  with Bernstein bounds [Maurer and Pontil, 2009], instead the  $\ell^1$ -norm.

Interestingly, the representation loss is also inversely proportional to  $g^*$ , the optimal selection rate in the true CMDP. The reason is that the CMDP constraints do not control the ratios  $\lambda_j^i(C_T) = \frac{N_j^i(T)}{N(T)}$ , but  $N_j^i(T)$  instead (by definition of  $R^c(T)$  and  $\xi_j^i$ ). If  $N(T)$  is small, i.e. due to a small selection rate  $g$ , then  $R_j^i(T) = |N_j^i(T) - \rho_j^i N(T)|$  is small, but not necessarily  $|\frac{N_j^i(T)}{N(T)} - \rho_j^i|$ : the committee is too small to be representative.

## 6 Experiments

The goal of these experiments is to answer the following: **(Q1)** In practice, for which range of committee sizes do our strategies achieve satisfying sample complexity and representation loss? **(Q2)** What is the cost of not knowing the distribution  $p$  for the sample complexity and representation loss?

**Experimental setting.** To answer these questions, we use summary data from the 2017 Citizens’ Assembly on Brexit. The participants were recruited in an offline manner: volunteers could express interest in a survey, and then 53 citizens were drawn from the pool of volunteers using stratified sampling, in order to construct an assembly that reflects the diversity of the UK electorate. We use summary statistics published in the report [Renwick *et al.*, 2017] to simulate an on-line recruitment process.

There are  $d = 6$  features: the organisers expressed target quotas for 2 ethnicity groups, 2 social classes, 3 age groups, 8 regions, 2 gender groups and 2 Brexit vote groups (remain, leave). The report also includes the number of people contacted per feature group (e.g., women, or people who voted to remain) and the volunteering rate for each feature group, which we use as probability of volunteering given a feature group. We use Bayes’ rule to compute the probabilities of feature groups among volunteers, and use them as the marginal distributions  $\Pr[x^i = j | \text{volunteers}]$  (since we only consider the population of volunteers). Since we only have access to the marginals, we compute the joint distribution as if the features were independent, although our model is agnostic to the dependence structure of the joint distribution.

We study Greedy with tolerance  $\epsilon = 0.02, 0.05$ . We run experiments for  $K = 50, 100, 150, 250, 500, 1000$ , averaged over 50 simulations.

**(A1).** We compare Greedy and CMDP, when the distribution  $p$  is *known*. Figure 1 shows that the greedy strategy with  $\epsilon = 0.05$  requires 10 times more samples than CMDP, and its representation loss is higher as soon as  $K \geq 250$ . Greedy with lower tolerance  $\epsilon = 0.02$  achieves better representation than CMDP for smaller committees ( $K \leq 100$ ), but the margin quickly decreases with  $K$ . However, even for small committees, it requires about 100 times more samples, which is prohibitively expensive. Figure 1 shows that for CMDP, the sample complexity grows linearly in the committee size, with a reasonable slope (we need to find  $\tau \approx 500$  volunteers for a committee of size  $K \approx 200$ ).

**(A2).** To corroborate the previously discussed effect of  $|\mathcal{X}|$  when  $p$  is *unknown*, we evaluate RL-CMDP on different configurations: (1) using only the features ethnicity, social class, and gender ( $d = 3, |\mathcal{X}| = 8$ ), (2) using all features except regions ( $d = 5, |\mathcal{X}| = 48$ ). Fig. 2 shows that unlike CMDP which has full knowledge of  $p$ , it is for large committee sizes that RL-CMDP reaches low representation loss (below 0.05 for  $K \geq 1500$  in the configuration(1)). This is because RL-CMDP needs to collect more samples to estimate  $p$ , as discussed in Th. 1. For known  $p$ , the CMDP approach achieves the same representativeness for middle-sized committees (repr. loss  $\leq 0.05$  for  $K \approx 250$ ). Hence, comparing the cases of known (Fig. 1) and unknown distribution  $p$  (Fig. 2), the ignorance




Figure 1: Effect of committee size  $K$  on sample complexity and representation loss for different strategies, in the UK Brexit Assembly experiment, using all features.  $p$  is **known**.




Figure 2: Effect of committee size  $K$  on sample complexity and representation loss for RL-CMDP, on data simulated from the UK Brexit Assembly, using 3 and 5 features.  $p$  is **unknown**.

of  $p$  is not costly for sample complexity, but rather for the representation loss which decreases more slowly.

Consistently with Th. 1, we observe that the representation loss is higher when  $\mathcal{X}$  is larger ( $d = 5$ ). For small and middle-sized committees, the loss of RL-CMDP is much worse than Greedy’s which also works for unknown  $p$ . For large committees though, the margin is only 0.05 when  $K \gtrsim 2000$  and  $\tau \approx 3500$  for RL-CMDP (which is  $\times 3$  more sample efficient than Greedy). In absolute terms, the theoretical regret bounds have a large constant  $\sqrt{|\mathcal{X}|}$ . This constant is likely unavoidable asymptotically because it comes from Lem. 3, but our experiments suggest that in the non-asymptotic regime, RL-CMDP performs better than the bound suggests.

## 7 Conclusion

We formalised the problem of selecting a diverse committee with multi-attribute proportional representation in an online setting. We addressed the case of known candidate distributions with constrained MDPs, and leveraged exploration-exploitation techniques to address unknown distributions.

## Acknowledgements

This work was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “Investissements d’avenir” program ANR-19-P3IA-0001 (PRAIRIE 3IA Institute). We thank Matteo Pirota for his helpful suggestions and feedback.

## References

- [Agrawal and Devanur, 2016] Shipra Agrawal and Nikhil Devanur. Linear contextual bandits with knapsacks. In *Advances in Neural Information Processing Systems*, pages 3450–3458, 2016.
- [Altman, 1999] Eitan Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.
- [Aziz, 2019] Haris Aziz. A rule for committee selection with soft diversity constraints. *Group Decision and Negotiation*, 28:1193–1200, 2019.
- [Babaioff *et al.*, 2008] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *ACM SIGecom Exchanges*, 7(2):1–11, 2008.
- [Badanidiyuru *et al.*, 2014] Ashwinkumar Badanidiyuru, Baharan Mirzasoleiman, Amin Karbasi, and Andreas Krause. Streaming submodular maximization: Massive data summarization on the fly. In *Proceedings of the 20th ACM SIGKDD*, pages 671–680, 2014.
- [Bateni *et al.*, 2013] MohammadHossein Bateni, Mohammadtaghi Hajiaghayi, and Morteza Zadimoghaddam. Submodular secretary problem and extensions. *ACM Transactions on Algorithms (TALG)*, 9(4):1–23, 2013.
- [Bei *et al.*, 2020] Xiaohui Bei, Shengxin Liu, Chung Keung Poon, and Hongao Wang. Candidate selections with proportional fairness constraints. In *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20, Auckland, New Zealand, May 9-13, 2020*, pages 150–158, 2020.
- [Benadè *et al.*, 2019] Gerdus Benadè, Paul Gözl, and Ariel D Procaccia. No stratification without representation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 281–314, 2019.
- [Bredereck *et al.*, 2018] Robert Bredereck, Piotr Faliszewski, Ayumi Igarashi, Martin Lackner, and Piotr Skowron. Multiwinner elections with diversity constraints. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [Celis *et al.*, 2018] L. Elisa Celis, Lingxiao Huang, and Nisheeth K. Vishnoi. Multiwinner voting with fairness constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence*, pages 144–151, 2018.
- [Efroni *et al.*, 2020] Yonathan Efroni, Shie Mannor, and Matteo Pirota. Exploration-exploitation in constrained mdps. *arXiv preprint arXiv:2003.02189*, 2020.
- [Faliszewski *et al.*, 2017] Piotr Faliszewski, Piotr Skowron, Arkadii Slinko, and Nimrod Talmon. Multiwinner voting: A new challenge for social choice theory. *Trends in computational social choice*, 74:27–47, 2017.
- [Flanigan *et al.*, 2020] Bailey Flanigan, Paul Gözl, Anupam Gupta, and Ariel D Procaccia. Neutralizing self-selection bias in sampling for sortition. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 6528–6539. Curran Associates, Inc., 2020.
- [Jaksch *et al.*, 2010] Thomas Jaksch, Ronald Ortner, and Peter Auer. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research*, 11(4), 2010.
- [Lang and Skowron, 2018] Jérôme Lang and Piotr Skowron. Multi-attribute proportional representation. *Artificial Intelligence*, 263:74–106, 2018.
- [Maurer and Pontil, 2009] Andreas Maurer and Massimiliano Pontil. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- [Panigrahi *et al.*, 2012] Debmalya Panigrahi, Atish Das Sarma, Gagan Aggarwal, and Andrew Tomkins. Online selection of diverse results. In *Proceedings of the fifth ACM international conference on Web search and data mining*, pages 263–272, 2012.
- [Renwick *et al.*, 2017] A Renwick, S Allan, W Jennings, R McKee, M Russell, and G Smith. A considered public voice on brexit: The report of the citizens’ assembly on brexit. 2017.
- [Rosenberg and Mansour, 2019] Aviv Rosenberg and Yishay Mansour. Online convex optimization in adversarial markov decision processes. *arXiv preprint arXiv:1905.07773*, 2019.
- [Schumann *et al.*, 2019] Candice Schumann, Samsara N Counts, Jeffrey S Foster, and John P Dickerson. The diverse cohort selection problem. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 601–609. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- [Singh *et al.*, 2020] Rahul Singh, Abhishek Gupta, and Ness B Shroff. Learning in markov decision processes under constraints. *arXiv preprint arXiv:2002.12435*, 2020.
- [Stoyanovich *et al.*, 2018] Julia Stoyanovich, Ke Yang, and HV Jagadish. Online set selection with fairness and diversity constraints. In *Proceedings of the EDBT Conference*, 2018.
- [Weissman *et al.*, 2003] Tsachy Weissman, Erik Ordentlich, Gadiel Seroussi, Sergio Verdu, and Marcelo J Weinberger. Inequalities for the  $l_1$  deviation of the empirical distribution. *Hewlett-Packard Labs, Tech. Rep.*, 2003.
- [Zheng and Ratliff, 2020] Liyuan Zheng and Lillian J Ratliff. Constrained upper confidence reinforcement learning. *arXiv preprint arXiv:2001.09377*, 2020.