
Making the Cut: A Bandit-based Approach to Tiered Interviewing

Candice Schumann* Zhi Lang* Jeffrey S. Foster† John P. Dickerson*

*University of Maryland †Tufts University

{schumann, zlang}@cs.umd.edu, jfoster@cs.tufts.edu, john@cs.umd.edu

Abstract

Given a huge set of applicants, how should a firm allocate sequential resume screenings, phone interviews, and in-person site visits? In a tiered interview process, later stages (e.g., in-person visits) are more informative, but also more expensive than earlier stages (e.g., resume screenings). Using accepted hiring models and the concept of structured interviews, a best practice in human resources, we cast tiered hiring as a combinatorial pure exploration (CPE) problem in the stochastic multi-armed bandit setting. The goal is to select a subset of arms (in our case, applicants) with some combinatorial structure. We present new algorithms in both the probably approximately correct (PAC) and fixed-budget settings that select a near-optimal cohort with provable guarantees. We show on real data from one of the largest US-based computer science graduate programs that our algorithms make better hiring decisions or use less budget than the status quo.

‘... nothing we do is more important than hiring and developing people. At the end of the day, you bet on people, not on strategies.’ – Lawrence Bossidy, The CEO as Coach (1995)

1 Introduction

Hiring workers is expensive and lengthy. The average cost-per-hire in the United States is \$4,129 [Society for Human Resource Management, 2016], and with over five million hires per month on average, total annual hiring cost in the United States tops hundreds of billions of dollars [United States Bureau of Labor Statistics, 2018]. In the past decade, the average length of the hiring process has doubled to nearly one month [Chamberlain, 2017]. At every stage, firms expend resources to learn more about each applicant’s true quality, and choose to either cut that applicant or continue interviewing with the intention of offering employment.

In this paper, we address the problem of a firm hiring a *cohort* of multiple workers, each with unknown true utility, over multiple stages of structured interviews. We operate under the assumption that a firm is willing to spend an increasing amount of resources—e.g., money or time—on applicants as they advance to later stages of interviews. Thus, the firm is motivated to aggressively “pare down” the applicant pool at every stage, culling low-quality workers so that resources are better spent in more costly later stages. This concept of tiered hiring can be extended to crowdsourcing or finding a cohort of trusted workers. At each successive stage, crowdsourced workers are given harder tasks.

Using techniques from the multi-armed bandit (MAB) and submodular optimization literature, we present two new algorithms—in the probably approximately correct (PAC) (§3) and fixed-budget settings (§4)—and prove upper bounds that select a near-optimal cohort in this restricted setting. We explore those bounds in simulation and show that the restricted setting is not necessary in practice (§5). Then, using real data from admissions to a large US-based computer science Ph.D. program, we show that our algorithms yield *better* hiring decisions at equivalent cost to the status quo—or *comparable* hiring decisions at lower cost (§5).

2 A Formal Model of Tiered Interviewing

In this section, we provide a brief overview of related work, give necessary background for our model, and then formally define our general multi-stage combinatorial MAB problem. Each of our n applicants is an arm a in the full set of arms A . Our goal is to select $K < n$ arms that maximize some objective w using a maximization oracle. We split up the review/interview process into m stages, such that each stage $i \in [m]$ has per-interview information gain s_i , cost j_i , and number of required arms K_i (representing the size of the “short list” of applicants who proceed to the next round). We want to solve this problem using either a confidence constraint (δ, ϵ) , or a budget constraint over each stage (T_i) . We rigorously define each of these inputs below.

Multi-armed bandits. The multi-armed bandit problem allows for modeling resource allocation during sequential decision making. Bubeck *et al.* [2012] provide a general overview of historic research in this field. In a MAB setting there is a set of n arms A . Each arm $a \in A$ has a true utility $u(a) \in [0, 1]$, which is unknown. When an arm $a \in A$ is pulled, a reward is pulled from a distribution with mean $u(a)$ and a σ -sub-Gaussian tail. These pulls give an empirical estimate $\hat{u}(a)$ of the underlying utility, and an uncertainty bound $rad(a)$ around the empirical estimate, i.e., $\hat{u}(a) - rad(a) < u(a) < \hat{u}(a) + rad(a)$ with some probability δ . Once arm a is pulled (e.g. an application is reviewed or an interview is performed), $\hat{u}(a)$ and $rad(a)$ are updated.

Top-K and subsets. Traditionally, MAB problems focus on selecting a single best (i.e., highest utility) arm. Recently, MAB formulations have been proposed that select an optimal subset of K arms. Bubeck *et al.* [2013] propose a budgeted algorithm (SAR) that successively accepts and rejects arms. We build on work by Chen *et al.* [2014], which generalizes SAR to a setting with a combinatorial objective. They also outline a fixed-confidence version of the combinatorial MAB problem. In the Chen *et al.* [2014] formulation, the overall goal is to choose an optimal cohort M^* from a decision class \mathcal{M} . In this work, we use decision class $\mathcal{M}_K(A) = \{M \subseteq A \mid |M| = K\}$. A cohort is optimal if it maximizes a linear objective function $w : \mathbb{R}^n \times \mathcal{M}_K(A) \rightarrow \mathbb{R}$. Chen *et al.* [2014] rely on a maximization oracle, as do we, defined as

$$Oracle_K(\hat{\mathbf{u}}, A) = \operatorname{argmax}_{M \in \mathcal{M}_K(A)} w(\hat{\mathbf{u}}, M). \quad (1)$$

Chen *et al.* [2014] define a *gap score* for each arm a in the optimal cohort M^* , which is the difference in combinatorial utility between M^* and the best cohort *without* arm a . For each arm a not in the optimal set M^* , the gap score is the difference in combinatorial utility between M^* and the best set *with* arm a . Formally, for any arm $a \in A$, the gap score Δ_a is defined as

$$\Delta_a = \begin{cases} w(M^*) - \max_{\{M \mid M \in \mathcal{M}_K \wedge a \in M\}} w(M), & \text{if } a \notin M^* \\ w(M^*) - \max_{\{M \mid M \in \mathcal{M}_K \wedge a \notin M\}} w(M), & \text{if } a \in M^*. \end{cases} \quad (2)$$

Using this gap score we estimate the hardness of a problem as the sum of inverse squared gaps:

$$\mathbf{H} = \sum_{a \in A} \Delta_a^{-2} \quad (3)$$

This helps determine how easy it is to differentiate between arms at the border of accept/reject.

Objectives. Cao *et al.* [2015] tighten the bounds of Chen *et al.* [2014] where the objective function is Top-K, defined as

$$w_{\text{TOP}}(\mathbf{u}, M) = \sum_{a \in M} u(a). \quad (4)$$

In this setting the objective is to pick the K arms with the highest utility. Jun *et al.* [2016] look at the Top-K MAB problem with batch arm pulls, and Singla *et al.* [2015a] look at the Top-K problem from a crowdsourcing point of view.

In this paper, we explore a different type of objective that balances both individual utility and the diversity of the set of arms returned. Research has shown that a more diverse workforce produces better products and increases productivity [Desrochers, 2001; Hunt *et al.*, 2015]. Thus, such an objective is of interest to our application of hiring workers. In the document summarization setting, Lin and Bilmes [2011] introduced a *submodular* diversity function where the arms are partitioned into q disjoint groups P_1, \dots, P_q :

$$w_{\text{DIV}}(\mathbf{u}, M) = \sum_{i=0}^q \sqrt{\sum_{a \in P_i \cap M} u(a)}. \quad (5)$$

Nemhauser *et al.* [1978] prove theoretical bounds for the simple greedy algorithm that selects a set that maximizes a submodular, monotone function. Krause and Golovin [2014] overview submodular optimization in general. Singla *et al.* [2015b] propose an algorithm for maximizing an unknown function, and Ashkan *et al.* [2015] introduce a greedy algorithm that optimally solves the problem of diversification if that diversity function is submodular and monotone. Radlinski *et al.* [2008] learn a diverse ranking from behavior patterns of different users by using multiple MAB instances. Yue and Guestrin [2011] introduce the linear submodular bandits problem to select diverse sets of content while optimizing for a class of feature-rich submodular utility models. Each of these papers uses submodularity to promote some notion of diversity. Using this as motivation, we empirically show that we can hire a diverse cohort of workers (Section 5).

Variable costs. In many real-world settings, there are different ways to gather information, each of which vary in cost and effectiveness. Ding *et al.* [2013] looked at a regret minimization MAB problem in which, when an arm is pulled, a random reward is received and a random cost is taken from the budget. Xia *et al.* [2016] extend this work to a batch arm pull setting. Jain *et al.* [2014] use MABs with variable rewards and costs to solve a crowdsourcing problem. While we also assume non-unit costs and rewards, our setting is different than each of these, in that we actively choose how much to spend on each arm pull.

Interviews allow firms to compare applicants. *Structured interviews* treat each applicant the same by following the same questions and scoring strategy, allowing for meaningful cross-applicant comparison. A substantial body of research shows that structured interviews serve as better predictors of job success and reduce bias across applicants when compared to traditional methods [Harris, 1989; Posthuma *et al.*, 2002]. As decision-making becomes more data-driven, firms look to demonstrate a link between hiring criteria and applicant success—and increasingly adopt structured interview processes [Kent and McCarthy, 2016; Levashina *et al.*, 2014]. Motivated by the structured interview paradigm, Schumann *et al.* [2019] introduced a concept of “weak” and “strong” pulls in the Strong Weak Arm Pull (SWAP) algorithm. SWAP probabilistically chooses to strongly or weakly pull an arm. Inspired by that work we associate pulls with a cost $j \geq 1$ and information gain $s \geq j$, where a pull receives a reward pulled from a distribution with a $\frac{\sigma}{\sqrt{s}}$ -sub-Gaussian tail, but incurs cost j . This type of arm pull is equivalent in information gained to doing a traditional unit cost pull s times while only costing j . Schumann *et al.* [2019] only allow for two types of arm pulls and they do not account for the structure of current tiered hiring frameworks; nevertheless, in Section 5, we extend (as best we can) their model to our setting and compare as part of our experimental testbed.

Generalizing to multiple stages. This paper, to our knowledge, gives the first computational formalization of tiered structured interviewing. We build on hiring models from the behavioral science literature [Vardarlier *et al.*, 2014; Breugh and Starke, 2000] in which the hiring process starts at recruitment and follows several stages, concluding with successful hiring. We model these m successive stages as having an increased cost—in-person interviews cost more than phone interviews, which in turn cost more than simple résumé screenings—but return additional information via the score given to an applicant. For each stage $i \in [m]$ the user defines a cost j_i and an information gain s_i for the type of pull (type of interview) being used in that stage. During each stage, K_i arms move on to the next stage (we cut off $K_{i-1} - K_i$ arms), where $n = K_0 > K_1 > \dots > K_{m-1} > K_m = K$. The user must therefore define K_i for each $i \in [m - 1]$. The arms chosen to move on to the next stage are denoted as $A_m \subset A_{m-1} \subset \dots \subset A_1 \subset A_0 = A$.

3 Probably Approximately Correct Hiring

In this section, we present Cutting Arms using a Combinatorial Oracle (CACO), the first of two multi-stage algorithms for selecting a cohort of arms with provable guarantees. CACO is a probably approximately correct (PAC) [Haussler and Warmuth, 1993] algorithm that performs interviews over m stages, for a user-supplied parameter m , before returning a final subset of K arms.

Algorithm 1 provides pseudocode for CACO. The algorithm requires several user-supplied parameters in addition to the standard PAC-style confidence parameters (δ - confidence probability, ϵ - error), including the total number of stages m ; pairs (s_i, j_i) for each stage $i \in [m]$ representing the information gain s_i and cost j_i associated with each arm pull; the number K_i of arms to remain at the end of each stage $i \in [m]$; and a maximization oracle. After each stage i is complete, CACO

removes all but K_i arms. The algorithm tracks these “active” arms, denoted by A_{i-1} for each stage i , the total cost $Cost$ that accumulates over time when pulling arms, and per-arm a information such as empirical utility $\hat{u}(a)$ and total information gain $T(a)$. For example, if arm a has been pulled once in stage 1 and twice in stage 2, then $T(a) = s_1 + 2s_2$.

Algorithm 1 Cutting Arms using a Combinatorial Oracle (CACO)

Require: Confidence $\delta \in (0, 1)$; Error $\epsilon \in (0, 1)$; *Oracle*; number of stages m ; (s_i, j_i, K_i) for each stage i

```

1:  $A_0 \leftarrow A$ 
2: for stage  $i = 1, \dots, m$  do
3:   Pull each  $a \in A_{i-1}$  once using the given  $s_i, j_i$  pair
4:   Update empirical means  $\hat{\mathbf{u}}$ 
5:    $Cost \leftarrow Cost + K_{i-1} \cdot j_i$ 
6:   for  $t = 1, 2, \dots$  do
7:      $A_i \leftarrow Oracle_{K_i}(\hat{\mathbf{u}})$ 
8:     for  $a \in A_{i-1}$  do
9:        $rad(a) \leftarrow \sigma \sqrt{\frac{2 \log(4|A| Cost^3) / \delta}{T(a)}}$ 
10:      if  $a \in A_i$  then  $\tilde{u}(a) \leftarrow \hat{u}(a) - rad_t(a)$ 
11:      else  $\tilde{u}(a) \leftarrow \hat{u}(a) + rad(a)$ 
12:       $\tilde{A}_i \leftarrow Oracle_{K_i}(\tilde{\mathbf{u}})$ 
13:      if  $|w(\tilde{A}_i) - w(A_i)| < \epsilon$  then break
14:       $p \leftarrow \arg \max_{a \in (\tilde{A}_i \setminus A_i) \cup (A_i \setminus \tilde{A}_i)} rad(a)$ 
15:      Pull arm  $p$  using the given  $s_i, j_i$  pair
16:      Update  $\hat{u}(p)$  with the observed reward
17:       $T(p) \leftarrow T(p) + s_i$ 
18:       $Cost \leftarrow Cost + j_i$ 
19:  $Out \leftarrow A_m$ ; return  $Out$ 

```

is pulled (line 14). At the end of the last stage m , CACO returns a final set of K active arms that approximately maximizes an objective function (line 19).

We prove a bound on CACO in Theorem 1. As a special case of this theorem, when only a single stage of interviewing is desired, and as $\epsilon \rightarrow 0$, then Algorithm 1 reduces to Chen *et al.* [2014]’s CLUCB, and our bound then reduces to their upper bound for CLUCB. This bound provides insights into the trade-offs of cost, information gain, problem hardness, and shortlist size. Given the cost and information gain parameters Theorem 1 provides a tighter bound than those for CLUCB.

Theorem 1. *Given any $\delta \in (0, 1)$, any $\epsilon \in (0, 1)$, any decision classes $\mathcal{M}_i \subseteq 2^{[n]}$ for each stage $i \in [m]$, any linear function w , and any expected rewards $u \in \mathbb{R}^n$, assume that the reward distribution φ_a for each arm $a \in [n]$ has mean $u(a)$ with a σ -sub-Gaussian tail. Let $M_i^* = \arg \max_{M \in \mathcal{M}_i}$ denote the optimal set in stage $i \in [m]$. Set $rad_t(a) = \sigma \sqrt{2 \log(\frac{4K_{i-1} Cost_{i,t}^3}{\delta}) / T_{i,t}(a)}$ for all $t > 0$ and $a \in [n]$. Then, with probability at least $1 - \delta$, the CACO algorithm (Algorithm 1) returns the set Out where $w(Out) - w(M_m^*) < \epsilon$ and*

$$T \leq O \left(\sigma^2 \sum_{i \in [m]} \left(\frac{j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{1}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log \left(\frac{\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{1}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \right) \right).$$

Theorem 1 gives a bound relative to problem-specific parameters such as the gap scores, inter-stage cohort sizes K_i , and so on. Figure 1 lends intuition as to how CACO changes with respect to these inputs, in terms of problem hardness (defined in Eq. 3). When a problem is easy, the min parts of the bound are dominated by gap scores Δ_a , and there is a smooth increase in total cost.

CACO begins with all arms active (line 1). Each stage i starts by pulling each active arm once using the given (s_i, j_i) pair to initialize or update empirical utilities (line 3). It then pulls arms until a confidence level is triggered, removes all but K_i arms, and continues to the next stage (line 13).

In a stage i , CACO proceeds in rounds indexed by t . In each round, the algorithm first finds a set A_i of size K_i using the maximization oracle and the current empirical means $\hat{\mathbf{u}}$ (line 7). Then, given a confidence radius (line 9), it computes pessimistic estimates $\tilde{u}(a)$ of the true utilities of each arm a and uses the oracle to find a set of arms \tilde{A}_i under these pessimistic assumptions (lines 10-12). If those two sets are “close enough” (ϵ away), CACO proceeds to the next stage (line 13). Otherwise, across all arms a in the symmetric difference between A_i and \tilde{A}_i , the arm p with the most uncertainty over its true utility—determined via $rad(a)$ —




Figure 1: Hardness (**H**) vs theoretical cost (T) as user-specified parameters to the CACO algorithm change.

When the problem gets harder, the mins are dominated by K_i^2/ϵ^2 and the cost is noisy but bounded below. When ϵ or δ increases, the lower bounds of the noisy section decrease—with the impact of ϵ dominating that of δ . A policymaker can use these high-level trade-offs to determine hiring mechanism parameters. For example, assume there are two interview stages. As the number K_1 of applicants who pass the first interview stage increases, so too does total cost T . However, if K_1 is too small (here, very close to the final cohort size K), then the cost also increases.

4 Hiring on a Fixed Budget with BRUTAS

Algorithm 2 Budgeted Rounds Updated Targets Successively (BRUTAS)

Require: Budgets $T_i \forall i \in [m]$; (s_i, j_i, \tilde{K}_i) for each stage i ; constrained oracle $COracle$

- 1: Define $\log(n) \triangleq \sum_{i=1}^n \frac{1}{i}$
- 2: $A_{0,1} \leftarrow \emptyset$; $B_{0,1} \leftarrow \emptyset$
- 3: **for** stage $i = 1, \dots, m$ **do**
- 4: $A_{i,1} \leftarrow A_{i-1, \tilde{K}_{i-1}+1}$; $B_{i,1} \leftarrow B_{i-1, \tilde{K}_{i-1}+1}$; $\tilde{T}_{i,0} \leftarrow 0$
- 5: **for** $t = 1, \dots, \tilde{K}_i$ **do**
- 6: $\tilde{T}_{i,t} \leftarrow \left\lceil \frac{T_i - \tilde{K}_i}{\log(\tilde{K}_i) j_i (\tilde{K}_i - t + 1)} \right\rceil$
- 7: **foreach** $a \in [n] \setminus (A_{i,t} \cup B_{i,t})$ **do**
- 8: Pull a $(\tilde{T}_{i,t} - \tilde{T}_{i,t-1})$ times; update $\hat{\mathbf{u}}_{i,t}(a)$
- 9: $M_{i,t} \leftarrow COracle(\hat{\mathbf{u}}_{i,t}, A_{i,t}, B_{i,t})$
- 10: **if** $M_{i,t} = \perp$ **then return** \perp
- 11: **foreach** $a \in [n] \setminus (A_{i,t} \cup B_{i,t})$ **do**
- 12: **if** $a \in M_{i,t}$ **then**
- 13: $\tilde{M}_{i,t,a} \leftarrow COracle(\hat{\mathbf{w}}_{i,t}, A_{i,t}, B_{i,t} \cup \{a\})$
- 14: **else**
- 15: $\tilde{M}_{i,t,a} \leftarrow COracle(\hat{\mathbf{w}}_{i,t}, A_{i,t} \cup \{a\}, B_{i,t})$
- 16: $p_{i,t} \leftarrow \operatorname{argmax}_{a \in [n] \setminus (A_{i,t} \cup B_{i,t})} w(M_{i,t}) - w(\tilde{M}_{i,t,a})$
- 17: **if** $p_{i,t} \in M_t$ **then**
- 18: $A_{i,t+1} \leftarrow A_{i,t} \cup \{p_{i,t}\}$; $B_{i,t+1} \leftarrow B_{i,t}$
- 19: **else**
- 20: $A_{i,t+1} \leftarrow A_{i,t}$; $B_{i,t+1} \leftarrow B_{i,t} \cup \{p_{i,t}\}$
- 21: **Out** $\leftarrow A_{m, \tilde{K}_{m+1}}$; **return** **Out**

represent how many decisions we need to make at stage $i \in [m]$; thus, $\sum_{i \in [m]} \tilde{K}_i = n$. The \tilde{K}_i s are independent of K , the final number of arms we want to accept, except that the total number of accept decisions across all \tilde{K} must sum to K .

The budgeted setting uses a constrained oracle $COracle : \mathbb{R}^n \times 2^{[n]} \times 2^{[n]} \rightarrow \mathcal{M} \cup \{\perp\}$ defined as

$$COracle(\hat{\mathbf{u}}, A, B) = \operatorname{argmax}_{\{M \in \mathcal{M}_{\mathcal{K}} \mid A \subseteq M \wedge B \cap M = \emptyset\}} w(\hat{\mathbf{u}}, M),$$

where A is the set of arms that have been accepted and B is the set of arms that have been rejected.

In each stage $i \in [m]$, BRUTAS starts by collecting the accept and reject sets from the previous stage. It then proceeds through \tilde{K}_i rounds, indexed by t , and selects a single arm to place in the accept set A or the reject set B . In a round t , it first pulls each *active* arm—arms not in A or B —a total of $\tilde{T}_{i,t} - \tilde{T}_{i,t-1}$ times using the appropriate s_i and j_i values. $\tilde{T}_{i,t}$ is set according to Line 6; note that $\tilde{T}_{i,0} = 0$. Once all the empirical means for each active arm have been updated, the constrained oracle is run to find the empirical best set $M_{i,t}$ (Line 9). For each active arm a , a new pessimistic set $\tilde{M}_{i,t,a}$ is found (Lines 11-15). a is placed in the accept set A if a is not in $M_{i,t}$, or in the reject set B if a is in $M_{i,t}$. This is done to calculate the gap that arm a creates (Equation 2). The arm $p_{i,t}$ with the largest gap is selected and placed in the accept set A if $p_{i,t}$ was included in $M_{i,t}$, or placed

In many hiring situations, a firm or committee has a fixed budget for hiring (number of phone interviews, total dollars to spend on hosting, and so on). With that in mind, in this section, we present Budgeted Rounds Updated Targets Successively (BRUTAS), a tiered-interviewing algorithm in the fixed-budget setting.

Algorithm 2 provides pseudocode for BRUTAS, which takes as input fixed budgets T_i for each stage $i \in [m]$, where $\sum_{i \in [m]} T_i = T$, the total budget. In this version of the tiered-interview problem, we also know how many decisions—whether to accept or reject an arm—we need to make in each stage. This is slightly different than in the CACO setting (§3), where we need to remove all but K_i arms at the conclusion of each stage i . We make this change to align with the CSAR setting of Chen *et al.* [2014], which BRUTAS generalizes. In this setting, let \tilde{K}_i

in the reject set B otherwise (Lines 16-20). Once all rounds are complete, the final accept set A is returned.

Theorem 2, provides an lower bound on the confidence that BRUTAS returns the optimal set. Note that if there is only a single stage, then Algorithm 2 reduces to Chen *et al.* [2014]’s CSAR algorithm, and our Theorem 2 reduces to their upper bound for CSAR. Again Theorem 2 provides tighter bounds than those for CSAR given the parameters for information gain s_b and arm pull cost j_b .

Theorem 2. *Given any T_i s such that $\sum_{i \in [m]} T_i = T > n$, any decision class $\mathcal{M}_K \subseteq 2^{[n]}$, any linear function w , and any true expected rewards $\mathbf{u} \in \mathbb{R}^n$, assume that reward distribution φ_a for each arm $a \in [n]$ has mean $u(a)$ with a σ -sub-Gaussian tail. Let $\Delta_{(1)}, \dots, \Delta_{(n)}$ be a permutation of $\Delta_1, \dots, \Delta_n$ (defined in Eq. 2) such that $\Delta_{(1)} \leq \dots \leq \Delta_{(n)}$. Define $\tilde{\mathbf{H}} \triangleq \max_{i \in [n]} i \Delta_{(i)}^{-2}$. Then, Algorithm 2 uses at most T_i samples per stage $i \in [m]$ and outputs a solution $\text{Out} \in \mathcal{M}_K \cup \{\perp\}$ such that*

$$\Pr[\text{Out} \neq M_*] \leq n^2 \exp\left(-\frac{\sum_{b=1}^m s_b(T_b - K_b)/(j_b \widetilde{\log}(K_b))}{72\sigma^2 \tilde{\mathbf{H}}}\right) \quad (6)$$

where $\widetilde{\log}(n) \triangleq \sum_{i=1}^n i^{-1}$, and $M_* = \operatorname{argmax}_{M \in \mathcal{M}_K} w(M)$.

When setting the budget for each stage, a policymaker should ensure there is sufficient budget for the number of arms in each stage i , and for the given exogenous cost values j_i associated with interviewing at that stage. There is also a balance between the number of decisions that must be made in a given stage i and the ratio $\frac{s_i}{j_i}$ of interview information gain and cost. Intuitively, giving higher budget to stages with a higher $\frac{s_i}{j_i}$ ratio makes sense—but one also would not want to make *all* accept/reject decisions in those stages, since more decisions corresponds to lower confidence. Generally, arms with high gap scores are accepted/rejected in the earlier stages, while arms with low gap scores are accepted/rejected in the later stages. If the problem is easy (most of the arms have high gap scores, i.e., many applicants are “clearly” low quality), then a policymaker should allocate more budget to the earlier stages. If the problem is hard (i.e., comparing applicants is less obvious), then a policymaker should allocate more budget in the later stages.




Figure 2: Comparison of $Cost$ vs information gain (s) as ϵ increases for CACO. Here, $\delta = 0.05$ and $\sigma = 0.2$. As ϵ increases, the cost of the algorithm also decreases. If the overall cost of the algorithm is low, then increasing s (while keeping j constant) provides diminishing returns.

5 Experiments

In this section, we experimentally evaluate BRUTAS and CACO in two different settings. The first setting uses data from a toy problem of Gaussian distributed arms. The second setting uses real admissions data from one of the largest US-based graduate computer science programs.

5.1 Gaussian Arm Experiments

We begin by using simulated data to test the tightness of our theoretical bounds. To do so, we instantiate a cohort of $n = 50$ arms whose true utilities, u_a , are sampled from a normal distribution. We aim to select a final cohort of size $K = 7$. When an arm is pulled during a stage with cost j and information gain s , the algorithm is charged a cost of j and a reward is pulled from a distribution with mean u_a and standard deviation of σ/\sqrt{s} . For simplicity, we present results in the setting of $m = 2$ stages.




Figure 3: Hardness (\mathbf{H}) vs $Cost$, comparing against Theorem 1.

CACO. To evaluate CACO, we vary δ , ϵ , σ , K_1 , and s_2 . We find that as δ increases, both cost and utility decrease, as expected. Similarly, Figure 2 shows that as ϵ increases, both cost and utility decrease. Higher values of σ increase the total cost, but do not affect utility. We also find diminishing returns from high information gain s values (x -axis of Figure 2). This makes sense—as s tends to infinity, the true utility is returned from a single arm pull. We also notice that if many “easy” arms (arms with very large gap scores) are allowed in higher stages, total cost rises substantially.

Although the bound defined in Theorem 1 assumes a linear function w , we empirically tested CACO using a submodular function w_{DIV} . We find that the cost of running CACO using this submodular function is significantly lower than the theoretical bound. This suggests that (i) the bound for CACO can be tightened and (ii) CACO could be run with submodular functions w .

BRUTAS. To evaluate BRUTAS, we varied σ and (\tilde{K}_i, T_i) pairs for two stages. Utility varies as expected from Theorem 2: when σ increases, utility decreases. There is also a trade-off between \tilde{K}_i and T_i values. If the problem is easy, a low budget and a high \tilde{K}_1 value is sufficient to get high utility. If the problem is hard (high \mathbf{H} value), a higher overall budget is needed, with more budget spent in the second stage. Figure 4 shows this escalating relationship between budget and utility based on problem hardness. Again we found that BRUTAS performed well when using a submodular function w_{DIV} .




Figure 4: Effect of an increasing budget on the overall utility of a cohort. As hardness (\mathbf{H}) increases, more budget is needed to produce a high quality cohort.

Finally, we compare CACO and BRUTAS to two baseline algorithms: UNIFORM and RANDOM, which uniformly and randomly respectively, pulls arms in each stage. In both algorithms, the maximization oracle is run after each stage to determine which arms should move on to the next stage. When given a budget of 2,750, BRUTAS achieves a utility of 244.0, which outperforms both the UNIFORM and RANDOM baseline utilities of 178.4 and 138.9, respectively. When CACO is run on the same problem, it finds a solution (utility of 231.0) that beats both UNIFORM and RANDOM at a roughly equivalent cost of 2,609. This qualitative behavior exists for other budgets.

5.2 Graduate Admissions Experiment

We evaluate how CACO and BRUTAS might perform in the real world by applying them to a graduate admissions dataset from one of the largest US-based graduate computer science programs. These experiments were approved by the university’s Institutional Review Board and did not affect any admissions decisions for the university. Our dataset consists of three years (2014–16) worth of graduate applications. For each application we also have graduate committee review scores (normalized to between 0 and 1) and admission decisions.

Experimental setup. Using information from 2014 and 2015, we used a random forest classifier [Pedregosa *et al.*, 2011], trained in the standard way on features extracted from the applications, to predict probability of acceptance. Features included numerical information such as GPA and GRE scores, topics from running Latent Dirichlet Allocation (LDA) on faculty recommendation letters [Schmader *et al.*, 2007], and categorical information such as region of origin and undergraduate school. In the testing phase, the classifier was run on the set of applicants A from 2016 to produce a probability of acceptance $P(a)$ for every applicant $a \in A$.

We mimic the university’s application process of two stages: a first review stage where admissions committee members review the application packet, and a second interview stage where committee members perform a Skype interview for a select subset of applicants. The committee members follow a structured interview approach. We determined that the time taken for a Skype interview is roughly 6 times as long as a packet review, and therefore we set the cost multiplier for the second stage $j_2 = 6$. We ran over a variety of s_2 values, and we determined σ by looking at the distribution of review scores from past years. When an arm $a \in A$ is pulled with information gain s and cost j , a reward is randomly pulled from the arm’s review scores (when $s_1 = 1$ and $j_1 = 1$, as in the first stage), or a reward is pulled from a Gaussian distribution with mean $P(a)$ and a standard deviation of $\frac{\sigma}{\sqrt{s}}$.




Figure 5: Utility vs Cost over four different algorithms (RANDOM, UNIFORM, SWAP, CACO, BRUTAS) and the actual admissions decisions made at the university. Both CACO and BRUTAS produce equivalent cohorts to the actual admissions process with lower cost, or produce high quality cohorts than the actual admissions process with equivalent cost.

We ran simulations for BRUTAS, CACO, UNIFORM, and RANDOM. In addition we compare to an adjusted version of Schumann *et al.* [2019]’s SWAP. SWAP uses a strong pull policy to probabilistically weak or strong pull arms. In this adjusted version we use a strong pull policy of always weak pulling arms until some threshold time t and strong pulling for the remainder of the algorithm. Note that this adjustment moves SWAP away from fixed confidence but not all the way to a budgeted algorithm like BRUTAS but fits into the tiered structure. For the budgeted algorithms BRUTAS, UNIFORM, and RANDOM, (as well as the pseudo-budgeted SWAP) if there are K_i arms in round i , the budget is $K_i \cdot x_i$ where $x_i \in \mathbb{N}$. We vary δ and ϵ to control CACO’s cost.

We compare the utility of the cohort selected by each of the algorithms to the utility from the cohort that was actually selected by the university. We maximize either objective w_{TOP} or w_{DIV} for each of the algorithms. We instantiate w_{DIV} , defined in Equation 5, in two ways: first, with self-reported gender, and second, with region of origin. Note that since the graduate admissions process is run entirely by humans, the committee does not explicitly maximize a particular function. Instead, the committee tries to find a good overall cohort while balancing areas of interest and general diversity.

Results. Figure 5 compares each algorithm to the actual admissions decision process performed by the real-world committee. In terms of utility, for both w_{TOP} and w_{DIV} , BRUTAS and CACO achieve similar gains to the actual admissions process (higher for w_{DIV} over region of origin) when using less cost/budget. When roughly the same amount of budget is used, BRUTAS and CACO are able to provide higher predicted utility than the true accepted cohort, for both w_{TOP} and w_{DIV} . As expected, BRUTAS and CACO outperform the baseline algorithms RANDOM, UNIFORM. The adjusted SWAP algorithm performs poorly in this restricted setting of tiered hiring. By limiting the strong pull policy of SWAP, only small incremental improvements can be made as Cost is increased.

6 Conclusions & Discussion of Future Research

We provided a formalization of tiered structured interviewing and presented two algorithms, CACO in the PAC setting and BRUTAS in the fixed-budget setting, which select a near-optimal cohort of applicants with provable bounds. We used simulations to quantitatively explore the impact of various parameters on CACO and BRUTAS and found that behavior aligns with theory. We showed empirically that both CACO and BRUTAS work well with a submodular function that promotes diversity. Finally, on a real-world dataset from a large US-based Ph.D. program, we showed that CACO and BRUTAS identify higher quality cohorts using equivalent budgets, or comparable cohorts using lower budgets, than the status quo admissions process. Moving forward, we plan to incorporate multi-dimensional feedback (e.g., with respect to an applicant’s technical, presentation, and analytical qualities) into our model; recent work due to Katz-Samuels and Scott [2018, 2019] introduces that feedback (in a single-tiered setting) as a marriage of MAB and constrained optimization, and we see this as a fruitful model to explore combining with our novel tiered system.

Discussion. The results support the use of BRUTAS and CACO in a practical hiring scenario. Once policymakers have determined an objective, BRUTAS and CACO could help reduce costs and produce better cohorts of employees. Yet, we note that although this experiment uses real data, it is still a simulation. The classifier is not a true predictor of utility of an applicant. Indeed, finding an estimate of utility for an applicant is a nontrivial task. Additionally, the data that we are using incorporates hu-

man bias in admission decisions, and reviewer scores [Schmader *et al.*, 2007; Angwin *et al.*, 2016]. Finally, defining an objective function on which to run CACO and BRUTAS is a difficult task. Recent advances in human value judgment aggregation [Freedman *et al.*, 2018; Noothigattu *et al.*, 2018] could find use in this decision-making framework.

7 Acknowledgements

Schumann and Dickerson were supported by NSF IIS RI CAREER Award #1846237. We thank Google for gift support, and the anonymous reviewers for helpful comments.

References

- Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing, 2016.
- Azin Ashkan, Branislav Kveton, Shlomo Berkovsky, and Zheng Wen. Optimal greedy diversity for recommendation. In *IJCAI*, 2015.
- James A Breugh and Mary Starke. Research on employee recruitment. *Journal of Management*, 2000.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 2012.
- Sébastien Bubeck, Tengyao Wang, and Nitin Viswanathan. Multiple identifications in multi-armed bandits. In *ICML*, 2013.
- Wei Cao, Jian Li, Yufei Tao, and Zhize Li. On top-k selection in multi-armed bandits and hidden bipartite graphs. In *NeurIPS*, 2015.
- Andrew Chamberlain. How long does it take to hire? <https://www.glassdoor.com/research/time-to-hire-in-25-countries/>, 2017.
- Shouyuan Chen, Tian Lin, Irwin King, Michael R Lyu, and Wei Chen. Combinatorial pure exploration of multi-armed bandits. In *NeurIPS*, 2014.
- Pierre Desrochers. Local diversity, human creativity, and technological innovation. *Growth and Change*, 2001.
- Wenkui Ding, Tao Qin, Xu-Dong Zhang, and Tie-Yan Liu. Multi-armed bandit with budget constraint and variable costs. In *AAAI*, 2013.
- Rachel Freedman, J Schaich Borg, Walter Sinnott-Armstrong, J Dickerson, and Vincent Conitzer. Adapting a kidney exchange algorithm to align with human values. In *AAAI*, 2018.
- Michael M Harris. Reconsidering the employment interview. *Personal Psychology*, 1989.
- David Haussler and Manfred Warmuth. The probably approximately correct (PAC) and other learning models. In *Foundations of Knowledge Acquisition*. 1993.
- Vivian Hunt, Dennis Layton, and Sara Prince. Diversity matters. *McKinsey & Company*, 2015.
- Shweta Jain, Sujit Gujar, Satyanath Bhat, Onno Zoeter, and Y Narahari. An incentive compatible mab crowdsourcing mechanism with quality assurance. *arXiv*, 2014.
- Kwang-Sung Jun, Kevin Jamieson, Robert Nowak, and Xiaojin Zhu. Top arm identification in multi-armed bandits with batch arm pulls. In *AISTATS*, 2016.
- Julian Katz-Samuels and Clayton Scott. Feasible arm identification. In *ICML*, 2018.
- Julian Katz-Samuels and Clayton Scott. Top feasible arm identification. In *AISTATS*, 2019.
- Julia D. Kent and Maureen Terese McCarthy. *Holistic Review in Graduate Admissions*. Council of Graduate Schools, 2016.

- Andreas Krause and Daniel Golovin. Submodular function maximization. In *Tractability*. 2014.
- Julia Levashina, Christopher J Hartwell, Frederick P Morgeson, and Michael A Campion. The structured employment interview. *Personnel Psychology*, 2014.
- Hui Lin and Jeff Bilmes. A class of submodular functions for document summarization. In *ACL*, 2011.
- George L Nemhauser, Laurence A Wolsey, and Marshall L Fisher. An analysis of approximations for maximizing submodular set functions. *Mathematical Programming*, 1978.
- Ritesh Noothigattu, Snehal Kumar Neil S. Gaikwad, Edmond Awad, Sohan D’Souza, Iyad Rahwan, Pradeep Ravikumar, and Ariel D. Procaccia. A voting-based system for ethical decision making. In *AAAI*, 2018.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn. *JMLR*, 2011.
- Richard A Posthuma, Frederick P Morgeson, and Michael A Campion. Beyond employment interview validity. *Personal Psychology*, 2002.
- Filip Radlinski, Robert Kleinberg, and Thorsten Joachims. Learning diverse rankings with multi-armed bandits. In *ICML*, 2008.
- Toni Schmader, Jessica Whitehead, and Vicki H. Wysocki. A linguistic comparison of letters of recommendation for male and female chemistry and biochemistry job applicants. *Sex Roles*, 2007.
- Candice Schumann, Samsara N. Counts, Jeffrey S. Foster, and John P. Dickerson. The diverse cohort selection problem. *AAMAS*, 2019.
- Adish Singla, Eric Horvitz, Pushmeet Kohli, and Andreas Krause. Learning to hire teams. In *HCOMP*, 2015.
- Adish Singla, Sebastian Tschischek, and Andreas Krause. Noisy submodular maximization via adaptive sampling with applications to crowdsourced image collection summarization. *AAAI*, 2015.
- Society for Human Resource Management. Human capital benchmarking report, 2016.
- United States Bureau of Labor Statistics. Job openings and labor turnover. <https://www.bls.gov/news.release/pdf/jolts.pdf>, 2018.
- Pelin Vardarlier, Yalcin Vural, and Semra Birgun. Modelling of the strategic recruitment process by axiomatic design principles. In *Social and Behavioral Sciences*, 2014.
- Yingce Xia, Tao Qin, Weidong Ma, Nenghai Yu, and Tie-Yan Liu. Budgeted multi-armed bandits with multiple plays. In *IJCAI*, 2016.
- Yisong Yue and Carlos Guestrin. Linear submodular bandits and their application to diversified retrieval. In *NeurIPS*, 2011.

A Table of Symbols

In this section, for expository ease and reference, we aggregate all symbols used in the main paper and give a brief description of their meaning and use. We note that each symbol is also defined explicitly in the body of the paper; Table 1 is provided as a reference.

B Proofs

In this section, we provide proofs for the theoretical results presented in the main paper. Appendix B.1 gives proofs for CACO, defined as Algorithm 1 in Section 3. Appendix B.2 gives proofs for BRUTAS, defined as Algorithm 2 in Section 4.

B.1 CACO

Theorem 1 requires lemmas from Chen *et al.* [2014]. We restate the theorem here for clarity and then proceed with the proof.

Theorem 1. *Given any $\delta \in (0, 1)$, any $\epsilon \in (0, 1)$, any decision classes $\mathcal{M}_i \subseteq 2^{[n]}$ for each stage $i \in [m]$ and any expected rewards $u \in \mathbb{R}^n$, assume that the reward distribution φ_a for each arm $a \in [n]$ has mean $u(e)$ with a σ -sub-Gaussian tail. Let $M_i^* = \operatorname{argmax}_{M \in \mathcal{M}_i}$ denote the optimal set in stage $i \in [m]$. Set $\operatorname{rad}_t(a) = \sigma \sqrt{2 \log(\frac{4K_{i-1} \operatorname{Cost}_{i,t}^3}{\delta})} / T_{i,t}(a)$ for all $t > 0$ and $a \in [n]$. Then, with probability at least $1 - \delta$, the CUT algorithm (Algorithm 1) returns the set Out where $\text{Out} - M_m^* < \epsilon$ and*

$$T \leq O \left(\sigma^2 \sum_{i \in [m]} \left(\frac{j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{1}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log \left(\frac{\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{1}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \right) \right) \quad (7)$$

Proof. Assume we are in some round i , and that we are at time t_a where some arm a is going to be pulled for the last time in round i . Set $\operatorname{rad}_{i,t}(a) = \sigma \sqrt{\frac{2 \log(4K_{i-1} \operatorname{Cost}_{i,t}^3 / \delta)}{T_{i,t}(a)}}$ Using Lemma 13 from Chen *et al.* [2014] we know that $\operatorname{rad}_{i,t} \geq \max \left\{ \frac{\Delta_a}{6}, \frac{\epsilon}{2K_i} \right\}$. Before arm a is pulled the following must be true:

$$\operatorname{rad}_{i,t_a} \geq \max \left\{ \frac{\Delta_a}{6}, \frac{\epsilon}{2K_i} \right\} \quad (8)$$

$$\begin{aligned} \operatorname{rad}_{i,t_a} &= \sigma \sqrt{\frac{2 \log(4K_{i-1} \operatorname{Cost}_{i,t_a}^3 / \delta)}{T_i(a) - s_i}} \\ &\leq \sigma \sqrt{\frac{2 \log(4K_{i-1} j_i^3 t_a^3 / \delta)}{T_i(a) - s_i}}. \end{aligned} \quad (9)$$

Equation 9 holds since $j_i > j_{i-1} > \dots > j_0$. Given equations 8 and equation 9 we have,

$$\begin{aligned} \max \left\{ \frac{\Delta_a}{6}, \frac{\epsilon}{2K_i} \right\} &\leq \sigma \sqrt{\frac{2 \log(2K_{i-1} j_i^3 t_a^3 / \delta)}{T_i(a) - s_i}} \\ &\leq \sigma \sqrt{\frac{2 \log(2K_{i-1} j_i^3 T_i^3 / \delta)}{T_i(a) - s_i}} \end{aligned}$$

Solving for $T_i(a)$ we have,

$$\begin{aligned} T_i(a) &\leq \sigma^2 \min \left\{ \frac{72}{\Delta_a^2}, \frac{16K_i^2}{\epsilon^2} \right\} \log(4K_{i-1} j_i^3 T_i^3 / \delta) \\ &\quad + s_i \end{aligned} \quad (10)$$

Note that

$$\begin{aligned} T_i(a) &\leq T(a) \\ \frac{j_i}{s_i} \sum_{a \in A_{i-1}} T_i(a) &= T_i. \end{aligned} \quad (11)$$

We will show later on in the proof

$$\begin{aligned}
T_i &\leq 499 \frac{\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
&\quad \log \left(\frac{4\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
&\quad + 2j_i K_{i-1}.
\end{aligned} \tag{12}$$

Summing up over equation 12 we have

$$\begin{aligned}
T &\leq 499 \sigma^2 \sum_{i \in [m]} \left(\frac{j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \right. \\
&\quad \left. \log \left(\frac{4\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \right)
\end{aligned} \tag{13}$$

which proves theorem 1.

Now we will go back to prove equation 12. If $K_{i-1} \geq \frac{1}{2}T_i$, then we see that $T_i \leq 2K_{i-1}$ and therefore equation 12 holds. Assume, then, that $K_{i-1} < \frac{1}{2}T_i$. Since $T_i > K_{i-1}$, we can write

$$\begin{aligned}
T &= C \frac{\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
&\quad \log \left(\frac{2K_{i-1} \sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right).
\end{aligned} \tag{14}$$

If $C < 499$ then equation 12 holds. Suppose then that $C > 499$. Using equation 11 and summing equation 13 for all active arms $a \in A_{i-1}$, we have

$$\begin{aligned}
T_i &\leq \frac{j_i}{s_i} (K_{i-1} s_i + \\
&\quad \sum_{a \in A_{i-1}} \sigma^2 \min \left\{ \frac{72}{\Delta_a^2}, \frac{16K_i^2}{\epsilon^2} \right\} \log \left(\frac{4K_{i-1} j_i^3 T_i^3}{\delta} \right)) \\
&\leq K_{i-1} j_i + \\
&\quad \frac{18\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log \left(\frac{4K_{i-1} j_i^3 T_i^3}{\delta} \right) \\
&= K_{i-1} j_i + \\
&\quad \frac{18\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log \left(\frac{4K_{i-1} j_i^3}{\delta} \right) \\
&\quad + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log(T) \\
&\leq K_{i-1} j_i +
\end{aligned}$$

$$\begin{aligned}
& \frac{18\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log \left(\frac{4K_{i-1}j_i^3}{\delta} \right) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \left(\frac{2C\sigma^2 j^4}{s_i \delta} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \right) \\
& \log \left(\frac{4K_{i-1}\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \tag{15} \\
= & K_{i-1}j_i + \\
& \frac{18\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log \left(\frac{4K_{i-1}j_i^3}{\delta} \right) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \log(2C) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \left(\frac{\sigma^2 j^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \log \left(\frac{4K_{i-1}\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
\leq & K_{i-1}j_i + \\
& \frac{18\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \left(\frac{4K_{i-1}\sigma^2 j^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log(2C) \log \left(\frac{4K_{i-1}\sigma^2 j^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right)
\end{aligned}$$

$$\begin{aligned}
& \log \left(\frac{4K_{i-1}\sigma^2 j^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& + \frac{54\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \left(\frac{4K_{i-1}\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
= & K_{i-1} j_i + \\
& (126 + 54 \log(2C)) \frac{\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \left(\frac{4K_{i-1}\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
\leq & K_{i-1} j_i \\
& + C \frac{\sigma^2 j_i}{s_i} \left(\sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \\
& \log \left(\frac{4K_{i-1}\sigma^2 j_i^4}{s_i \delta} \sum_{a \in A_{i-1}} \min \left\{ \frac{4}{\Delta_a^2}, \frac{K_i^2}{\epsilon^2} \right\} \right) \tag{16} \\
= & T_i \tag{17}
\end{aligned}$$

where equation 15 follows from equation 14 and the assumption that $K_{i-1} < \frac{1}{2}T_i$; equation 16 follows since $136 + 54 \log(2C) < C$ for all $C > 499$; and 17 is due to 14. Equation 17 is a contradiction. Therefore $C \leq 499$ and we have proved equation 12. \square

B.2 BRUTAS

In order to prove Theorem 2, we first need a few lemmas.

Lemma 1. *Let $\Delta_{(1)}, \dots, \Delta_{(n)}$ be a permutation of $\Delta_1, \dots, \Delta_n$ (defined in Eq. (2)) such that $\Delta_{(1)} \leq \dots \leq \Delta_{(n)}$. Given a stage $i \in [m]$, and a phase $t \in [K_i]$, we define random event $\tau_{i,t}$ as follows*

$$\tau_{i,t} = \left\{ \forall i \in [n] \setminus (A_t \cup B_t) \quad |\hat{u}_{i,t}(a) - u(a)| \right. \tag{18}$$

$$\left. < \frac{\Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)}}{6} \right\}. \tag{19}$$

Then, we have

$$\begin{aligned}
\tau &= \Pr \left[\bigcap_{i=1}^m \bigcap_{t=1}^{K_i} \tau_{i,t} \right] \\
&\geq 1 - n^2 \exp \left(- \frac{\sum_{b=1}^m s_b (T_b - K_b) / (j_i \widetilde{\log}(K_b))}{72\sigma^2 \mathbf{H}} \right). \tag{20}
\end{aligned}$$

Proof. In round i at phase t , arm a has been pulled $T(a)$ times. Therefore, by Hoeffding's inequality, we have

$$\Pr \left[|\hat{u}_{i,t}(a) - u(a)| \geq \frac{\Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)}}{6} \right]$$

$$\leq 2 \exp \left(- \frac{T_{i,t}(a) \Delta^2}{72\sigma^2} \right) \quad (21)$$

By using the definition of \tilde{T}_i, t , the quantity $\tilde{T}_{i,t} \Delta^2_{(n-\sum_{b=1}^{K_{i-1}} K_b - t + 1)}$ on the right-hand side of Eq. 21 can be further bounded by

$$\begin{aligned} & T_{i,t} \Delta_{(n-\sum_{b=1}^{K_{i-1}} K_b - t + 1)} \\ &= (s_i \tilde{T}_{i,t} + \sum_{b=1}^{i-1} s_b \tilde{T}_{b, K_b + 1}) \Delta^2_{(n-\sum_{b=1}^{K_{i-1}} K_b - t + 1)} \\ &\geq \left(\frac{s_i (T_{i,t} - K_i)}{j_i \widetilde{\log}(K_i) (K_i - t + 1)} \right. \\ &\quad \left. + \sum_{b=0}^{i-1} \frac{s_b (T_{b, K_b + 1} - K_b)}{j_b \widetilde{\log}(K_b) (K_b - K_b + 2)} \right) \Delta^2_{(n-\sum_{b=1}^{i-1} K_b - t + 1)} \\ &\geq \sum_{b=1}^i \frac{s_b (T_b - K_b)}{j_b \widetilde{\log}(K_b) \mathbf{H}}, \end{aligned}$$

where the last inequality follows from the definition of $\mathbf{H} = \max_{i \in [n]} i \Delta_{(i)}^{-2}$. By plugging the last inequality into Eq. 21, we have

$$\begin{aligned} & \Pr \left[|\hat{u}_{i,t}(a) - u(a)| \geq \frac{\Delta_{(n-\sum_{b=0}^{i-1} K_b - t + 1)}}{6} \right] \\ &\leq 2 \exp \left(- \frac{\sum_{b=1}^i (s_b (T_b - K_b) / (j_b \widetilde{\log}(K_b)))}{72\sigma^2 \mathbf{H}} \right) \quad (22) \end{aligned}$$

Now, using Eq. 22 and a union bound for all $i \in [m]$, all $t \in [K_i]$, and all $a \in [n] \setminus (A_{t,i} \cup B_{t,i})$, we have

$$\begin{aligned} & \Pr \left[\bigcap_{i=1}^m \bigcap_{t=1}^{K_i} \tau_{i,t} \right] \\ &\geq 1 - 2 \sum_{i=1}^m \sum_{t=1}^{K_i} (n - \sum_{b=0}^{i-1} K_b - t + 1) \\ &\quad \exp \left(- \frac{\sum_{b=1}^i (s_b (T_b - K_b) / (j_b \widetilde{\log}(K_b)))}{72\sigma^2 \mathbf{H}} \right) \\ &\geq 1 - n^2 \exp \left(- \frac{\sum_{b=1}^m s_b (T_b - K_b) / (j_b \widetilde{\log}(K_b))}{72\sigma^2 \mathbf{H}} \right). \end{aligned}$$

□

Lemma 2. Fix a stage $i \in [m]$, and a phase $t \in [K_i]$, suppose that random event $\tau_{i,t}$ occurs. For any vector $\mathbf{a} \in \mathbb{R}^n$, suppose that $\text{supp}(\mathbf{a}) \cap (A_{i,t} \cup B_{i,t}) = \emptyset$, where $\text{supp}(\mathbf{a}) \triangleq \{i | a(i) \neq 0\}$ is the support of vector \mathbf{a} . Then, we have

$$|\langle \tilde{\mathbf{u}}_{i,t}, \mathbf{a} \rangle - \langle \mathbf{u}_{i,t}, \mathbf{a} \rangle| < \frac{\Delta_{(n-\sum_{b=0}^{i-1} K_b - t + 1)}}{6} \|\mathbf{a}\|_1$$

Proof. Suppose that $\tau_{i,t}$ occurs. Then, we have

$$|\langle \tilde{\mathbf{u}}_{i,t}, \mathbf{a} \rangle - \langle \mathbf{u}_{i,t}, \mathbf{a} \rangle| \quad (23)$$

$$\begin{aligned}
&= |\langle \tilde{\mathbf{u}}_{i,t} - \mathbf{w}, \mathbf{a} \rangle| \\
&= \left| \sum_{b=1}^n (\tilde{u}_{t,i}(b) - u(b)) a(b) \right| \\
&\leq \left| \sum_{b \in [n] \setminus (A_{i,t} \cup B_{i,t})} (\tilde{u}_{i,t}(b) - u(b)) a(b) \right| \tag{24}
\end{aligned}$$

$$\begin{aligned}
&\leq \sum_{b \in [n] \setminus (A_{i,t} \cup B_{i,t})} |(\tilde{u}_{i,t}(b) - u(b)) a(b)| \\
&\leq \sum_{b \in [n] \setminus (A_{i,t} \cup B_{i,t})} |\tilde{u}_{i,t}(b) - u(i)| |a(i)| \\
&< \frac{\Delta_{(n - \sum_{b=0}^{i-1} K_i - t + 1)}}{6} \sum_{b \in [n] \setminus (A_{i,t} \cup B_{i,t})} |a(b)| \tag{25}
\end{aligned}$$

$$= \frac{\Delta_{(n - \sum_{b=0}^{i-1} K_i - t + 1)}}{6} \|\mathbf{a}\|_1 \tag{26}$$

where Eq. 24 follows from the assumption that \mathbf{a} is supported on $[n] \setminus (A_{i,t} \cup B_{i,t})$; Eq. 25 follows from the definition of $\tau_{i,t}$ (Eq. 18). \square

Lemma 3. Fix a stage $i \in [m]$, and a phase $t \in [K_i]$. Suppose that $A_{i,t} \subseteq M_*$ and $B_{i,t} \cap M_* = \emptyset$. Let M be a set such that $A_{i,t} \subseteq M$ and $B_{i,t} \cap M = \emptyset$. Let a and b be two sets satisfying $a \subseteq M \setminus M_*$, $b \subseteq M_* \setminus M$, and $a \cap b = \emptyset$. Then, we have

$$A_{i,t} \subseteq (M \setminus a \cup b)$$

and

$$B_{i,t} \cap (M \setminus a \cup b) = \emptyset$$

and

$$(a \cup b) \cap (A_{i,t} \cup B_{i,t}) = \emptyset.$$

Lemma 3 is due to Chen *et al.* [2014].

Lemma 4. Fix any stage $i \in [m]$, and any phase $t \in [K_i]$ such that $\sum_{b=0}^{i-1} K_i + t > 0$. Suppose that event $\tau_{i,t}$ occurs. Also assume that $A_{i,t} \subseteq M_*$ and $B_{i,t} \cap M_* = \emptyset$. Let $a \in [n] \setminus (A_{i,t} \cup B_{i,t})$ be an active arm such that $\Delta_{(n - \sum_{b=0}^{i-1} K_i - t + 1)} \leq \Delta_a$. Then, we have

$$\tilde{u}_{i,t}(M_{i,t}) - \tilde{u}_{i,t}(\tilde{M}_{i,t,a}) > \frac{2}{3} \Delta_{(n - \sum_{b=0}^{i-1} K_i - t + 1)}$$

Lemma 4 is due to Chen *et al.* [2014].

Lemma 5. Fix any stage $i \in [m]$, and any phase $t \in [K_i]$ such that $\sum_{b=0}^{i-1} K_i + t > 0$. Suppose that event $\tau_{i,t}$ occurs. Also assume that $A_{i,t} \subseteq M_*$ and $B_{i,t} \cap M_* = \emptyset$. Suppose an active arm $a \in [n] \setminus (A_{i,t} \cup B_{i,t})$ satisfies that $a \in (M_* \cap \neg M_{i,t}) \cup (\neg M_* \cap M_{i,t})$. Then, we have

$$\tilde{u}_{i,t}(M_{i,t}) - \tilde{u}_{i,t}(\tilde{M}_{i,t,a}) \leq \frac{1}{3} \Delta_{(n - \sum_{b=0}^{i-1} K_i - t + 1)}$$

Lemma 5 is due to Chen *et al.* [2014].

Now we can prove Theorem 2, restated below for clarity.

Theorem 2. Given any T_i 's where $\sum_{i \in [m]} T_i = T > n$, any decision class $\mathcal{M} \subseteq 2^{[n]}$, and any expected rewards $\mathbf{u} \in \mathbb{R}^n$. Assume that reward distribution φ_a for each arm $a \in [n]$ has mean $u(e)$ with a σ -sub-Gaussian tail. Let $\Delta_{(1)}, \dots, \Delta_{(n)}$ be a permutation of $\Delta_1, \dots, \Delta_n$ (defined in Eq. 2) such that $\Delta_{(1)} \leq \dots \leq \Delta_{(n)}$. Define $\mathbf{H} \triangleq \max_{i \in [n]} i \Delta_{(i)}^{-2}$. Then, algorithm 2 uses at most T_i samples per stage $i \in [m]$ and outputs a solution $\text{Out} \in \mathcal{M} \cup \{\perp\}$ such that

$$\Pr[\text{Out} \neq M_*] \leq n^2 \exp\left(-\frac{\sum_{b=1}^m s_b (T_b - K_b) / (j_b \widetilde{\log}(K_i))}{72\sigma^2 \mathbf{H}}\right) \tag{27}$$

where $\widetilde{\log}(n) \triangleq \sum_{i=1}^n i^{-1}$, $M_* = \operatorname{argmax}_{M \in \mathcal{M}} w(M)$.

Proof. First we show that the algorithm takes at most T_i samples in every stage $i \in [m]$. It is easy to see that exactly one arm is pulled for $\tilde{T}_i, 1$ times in stage i , one arm is pulled for $\tilde{T}_i, 2$ times in stage i, \dots , and one arm is pulled for $\tilde{T}_i, K_i - 1$ times in stage i . Therefore, the total number of samples used by the algorithm in stage $i \in [m]$ is bounded by

$$\begin{aligned} j_i \sum_{t=1}^{K_i} \tilde{T}_i, t &\leq j_i \sum_{t=1}^{K_i} \left(\frac{T_i - K_i}{\log(K_i) j_i (K_i - t + 1)} + 1 \right) \\ &= \frac{(T_i - K_i) j_i}{\log(K_i) j_i} \tilde{\log}(K_i) + K_i \\ &= T_i. \end{aligned}$$

By Lemma 1, we know that the event τ occurs with probability at least $1 - n^2 \exp\left(-\frac{\sum_{b=1}^m s_b (T_b - K_b) / (j_b \tilde{\log}(K_b))}{72\sigma^2 \mathbf{H}}\right)$. Therefore, we only need to prove that, under event τ , the algorithm outputs M_* . Assume that the event τ occurs in the rest of the proof.

We will use induction. Fix a stage $i \in [m]$ and phase $t \in [K_i]$. Suppose that the algorithm does not make any error before stage i and phase t , i.e. $A_{i,t} \subseteq M_*$ and $B_{i,t} \cap M_* = \emptyset$. We will show that the algorithm does not err at stage i , phase t .

At the beginning of phase t in stage i there are exactly $\sum_{b=0}^{i-1} K_b + t - 1$ inactive arms $|A_{i,t} \cup B_{i,t}| = \sum_{b=0}^{i-1} K_b + t - 1$. Therefore there must exist an active arm $e_{i,t} \in [n] \setminus (A_{i,t} \cup B_{i,t})$ such that $\Delta_{e_{i,t}} \geq \Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)}$. Hence, by Lemma 4, we have

$$\tilde{w}_{i,t}(M_{i,t}) - \tilde{w}_{i,t}(M_{i,t,e_{i,t}}) \geq \frac{2}{3} \Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)}. \quad (28)$$

Notice that the algorithm makes an error in phase t in stage i if and only if it accepts an arm $p_{i,t} \notin M_*$ or rejects an arm $p_{i,t} \in M_*$. On the other hand, arm $p_{i,t}$ is accepted when $p_{i,t} \in M_{i,t}$ and is rejected when $p_{i,t} \notin \tilde{M}_{i,t}$. Therefore, the algorithm makes an error in phase t in stage i if and only if $p_t \in (M_* \cap \neg M_{i,t}) \cup (\neg M_* \cap M_{i,t})$.

Suppose that $p_t \in (M_* \cap \neg M_{i,t}) \cup (\neg M_* \cap M_{i,t})$. Using Lemma 5, we see that

$$\tilde{w}_{i,t}(M_{i,t}) - \tilde{w}_{i,t}(\tilde{M}_{i,t,p_{i,t}}) \leq \frac{1}{3} \Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)}. \quad (29)$$

By combining Eq. 28 and Eq. 29, we see that

$$\tilde{w}_{i,t}(M_{i,t}) - \tilde{w}_{i,t}(\tilde{M}_{i,t,p_{i,t}}) \quad (30)$$

$$\leq \frac{1}{3} \Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)} \quad (31)$$

$$< \frac{2}{3} \Delta_{(n - \sum_{b=0}^{i-1} K_b - t + 1)} \quad (32)$$

$$\leq \tilde{w}_{i,t}(M_{i,t}) - \tilde{w}_{i,t}(\tilde{M}_{i,t,e_{i,t}}) \quad (33)$$

However, Eq. 30 is contradictory to the definition of $p_{i,t} \triangleq \operatorname{argmax}_{e \in [n] \setminus (A_{i,t} \cup B_{i,t})} \tilde{w}_{i,t}(M_{i,t}) - \tilde{w}_{i,t}(M_{i,t,e})$. This proves that $p_t \notin (M_* \cap \neg M_{i,t}) \cup (\neg M_* \cap M_{i,t})$. This means that the algorithm does not err at phase t in stage i , or equivalently $A_{i,t+1} \subseteq M_*$ and $B_{i,t+1} \cap M_* = \emptyset$.

Hence we have $A_{m,K_{m+1}} \subseteq M_*$ and $B_{m,K_{m+1}} \subseteq \neg M_*$ in the final phase of the final stage. Notice that $|A_{m,K_{m+1}}| + |B_{m,K_{m+1}}| = n$ and $A_{m,K_{m+1}} \cap B_{m,K_{m+1}} = \emptyset$. This means that $A_{m,K_{m+1}} = M_*$ and $B_{m,K_{m+1}} = \neg M_*$. Therefore the algorithm outputs $\text{Out} = A_{m,K_{m+1}} = M_*$ after phase K_m in stage m . \square

C Experimental Setup

The machines used for the experiments had 32GB RAM, 8 Intel SandyBridge CPU cores, and were initialized with Red Hat Enterprise Linux 7.3. A single run of SWAP over the graduate admissions data takes about 1 minute depending on the parameters.




Figure 6: Comparison of *Cost* over information gain (s) as σ increases for CACO. Here, $\delta = 0.05$ and $\epsilon = 0.05$.

D Additional Experimental Results

In this section, we present additional experimental results for CACO and BRUTAS. Table 3 supports the Gaussian simulation experiments of Section 5.1, specifically, the comparison of CACO and BRUTAS to two baseline pulling strategies.

Table 4 also supports the Gaussian simulation experiments from Section 5.1. Here, we vary δ instead of ϵ , as was done in Figure 2 of the main paper. As expected, when δ increases, the cost decreases. However, the magnitude of the effect is smaller than the effect from decreasing ϵ or varying K_1 . This is also expected, as discussed in the final paragraphs of Section 3, and shown in Figure 1.

Figure 6 shows that, as the standard deviation σ of the Gaussian distribution from which rewards are drawn increases, so too does the total cost of running CACO. The qualitative behavior shown in, e.g., Figure 2 of the main paper remains: as information gain s increases, overall cost decreases; as s increases substantially, we see a saturation effect; and, as final cohort size K increases, overall cost increases.

Figure 7 shows the behavior of CACO for different arm initializations, representing different utilities and groupings. We chose 4 representative initializations. For most initializations, when $K_1 = 10$, higher values of s_2 do not result in gains. This is because with $K_1 = 10$ and $K = 7$, there are only 3 decisions to make on which arms to cut and the information gain from the initial pull of all arms in stage 2 grants enough information, thus no additional pulls need to be made and cost is uniform across s_2 . However, if the problem of selecting from the short list is hard enough, additional resources must be spent to narrow the decisions down, as in the top left graph, where total costs decrease as s_2 increases for $K_1 = 10$ because additional pulls need to be made after the initial pulls of remaining arms in stage 2. This reflects real life well: usually, the short list can be cut down with one additional round of (more informative) interviews. However, in rare situations, some candidates are so close to each other that additional assessments need to be made about them. Another interesting result is that $K_1 = 10$ is not always the most cost effective option. If many of the initial candidates are close together in utility, it will be hard to narrow it down to a final 10 based on resume review alone: more candidates should be allowed to move onto the next round which has higher information gain. This can be seen in the bottom right graph.




Figure 7: Comparison of $Cost$ over information gain (s) for different sets of arms for CACO. Here, $\delta = 0.075$, $\epsilon = 0.05$, $\sigma = 0.2$.

E Limitations

This experiment uses real data but is still a simulation. The classifier is not a true predictor of utility of an applicant. Indeed, finding an estimate of utility for an applicant is a nontrivial task. Additionally, the data that we are using incorporates human bias in admission decisions, and reviewer scores. This means that the classifier—and therefore the algorithms—may produce a biased cohort. Training a human committee or using quantitative methods to (attempt to) mitigate the impact of human bias in review scoring is important future work. Similarly, CACO and BRUTAS require an objective function to run; recent advances in human value judgment aggregation [Freedman *et al.*, 2018; Noothigattu *et al.*, 2018] could find use in this decision-making framework. Additionally, although we were able to empirically show that both CACO and BRUTAS perform well using a submodular function w_{DIV} , there are no theoretical guarantees for submodular functions.

F Structured Interviews for Graduate Admissions

The goal of the interview is to help judge whether the applicant should be granted admission. The interviewer asks questions to provide insight into the applicant’s academic capabilities, research experience, perseverance, communication skills, and leadership abilities, among others.

Some example questions include:

- Describe a time when you have faced a difficult academic challenge or hurdle that you successfully navigated. What was the challenge and how did you handle it?
- What research experience have you had? What problem did you work on? What was most challenging? What did you learn most from the experience?
- Have you had any experiences where you were playing a leadership or mentoring role for others?
- What are your goals for graduate school? What do you want to do when you graduate?
- What concerns do you have about the program? What will your biggest challenge be? Is there anything else we should discuss?

The interviewer fills out an answer and score sheet during the interview. Each interviewer follows the same questions and is provided with the same answer and score sheet. This allows for consistency across interviews.

Symbol	Summary
n	Number of applicants/arms
A	Set of all arms (e.g., the set of all applicants)
a	An arm in A (e.g., an individual applicant)
K	Size of the required cohort
$\mathcal{M}_K(A)$	Decisions class or set of possible cohorts of size K
$u(a)$	True utility of arm a where $u(a) \in [0, 1]$
$\hat{u}(a)$	Empirical estimate of the utility of arm a
$rad(a)$	Uncertainty bound around the empirical estimate of the utility $\hat{u}(a)$ of arm a
w	Submodular and monotone objective function for a cohort where $w : \mathbb{R}^n \times \mathcal{M}_K(A) \rightarrow \mathbb{R}$
<i>Oracle</i>	Maximization oracle defined in Equation 1 and used by CACO
<i>COracle</i>	Constrained maximization oracle used by BRUTAS
M^*	Optimal cohort given the true utilities
Δ_a	The gap score of arm a defined in Equation 2
\mathbf{H}	The hardness of a problem defined in Equation 3
j_i	Cost of an arm pull at stage i
s_i	Information gain of an arm pull at stage i
m	Number of pulling stages (or interview stages)
K_i	Number of arms moving onto the next stage (stage $i + 1$)
A_i	The active arms that move onto the next stage (stage $i + 1$)
$T(a)$	Total information gain for arm a
$\bar{u}(a)$	Worst case estimate of utility of arm a
A_i	Best cohort chosen by using the worst case estimates of utility
ϵ	We want to return a cohort with total utility bounded by $w(M^*) - \epsilon$ for Algorithm 1
δ	The probability that we are within ϵ of the best cohort for Algorithm 1
T_i	Budget constraint for round i
T	Total budget
σ	Property of the σ -sub-Gaussian tailed normal distribution
p	The arm with the greatest uncertainty in CACO
K_i	Number of decisions to make in round i
$T_{i,t}$	Budget for BRUTAS in stage i , round t
$M_{i,t}$	Best cohort chosen in BRUTAS stage i , round t , using empirical utilities
$\tilde{M}_{i,t,a}$	Pessimistic estimate in BRUTAS stage i , round t , for arm a
$p_{i,t}$	Arm which results in largest gap in BRUTAS stage i , round t
\mathbf{H}	Hardness for BRUTAS
$P(a)$	Probability of acceptance for an arm (candidate), estimated by Random Forest Classifier
q	Number of groups for submodular diversity function
P_1, P_2, \dots, P_q	The groups for submodular diversity function

Table 1: List of symbols used in the main paper.

Parameter	Range
δ	0.3,0.2,0.1,0.075,0.05
ϵ	0.3,0.2,0.1,0.075,0.05
σ	0.1,0.2
j	6
s	7, ..., 20

Table 2: Parameters for graduate admissions experiments

Algorithm	Cost	Utility
Random	2750	138.9 (5.1)
Uniform	2750	178.4 (0.2)
CACO	2609	231.0 (0.1)
BRUTAS	2750	244.0 (0.1)

Table 3: Comparing CACO and BRUTAS to the baseline of Uniform and Random

δ	Cost			
	$K_1 = 10$	$K_1 = 13$	$K_1 = 18$	$K_1 = 29$
0.050	552.475	605.250	839.525	1062.725
0.075	542.425	582.675	827.025	1040.700
0.100	537.175	587.900	820.575	1078.975
0.200	503.650	568.300	801.525	1012.550

Table 4: Cost for CACO over various δ , for $\epsilon = 0.05$, $\sigma = 0.20$, $s_2 = 7$

Experiment		Cost	w_{TOP}	w_{DIV} over Gender	w_{DIV} over Region of Origin
Actual	–	~2,000	60.9	10.1	17.9
RANDOM	lower	1,359	40.2 (0.3)	9.7 (0.2)	16.9 (0.3)
	~equivalent	2,277	43.6 (0.5)	9.9 (0.1)	17.2 (0.2)
	higher	11,556	72.9 (4.9)	11.5 (0.1)	18.1 (3.5)
UNIFORM	lower	1,359	49.7 (0.3)	9.8 (0.1)	17.7 (0.1)
	~equivalent	2,277	54.7 (0.3)	9.9 (0.2)	18.3 (0.4)
	higher	11,556	79.5 (3.2)	11.9 (0.3)	19.6 (0.6)
SWAP	lower	1,400–1,500	58.7 (0.5)	10.1 (0.1)	19.0 (0.1)
	~equivalent	1,900–2,000	60.2 (0.4)	10.5 (0.1)	19.1 (0.1)
	higher	2,500–2,700	61.5 (0.5)	10.8 (0.2)	19.3 (0.1)
CACO	lower	1,400–1,460	61.1 (0.1)	10.1 (0.2)	18.9 (0.1)
	~equivalent	1,950–1,990	78.7 (0.2)	10.7 (0.1)	19.4 (0.2)
	higher	2,500–2,700	80.1 (0.4)	12.0 (0.3)	19.8 (0.3)
BRUTAS	lower	1,649	61.2 (0.2)	10.6 (0.1)	19.1 (0.2)
	~equivalent	2,038	79.3 (0.3)	10.7 (0.1)	19.8 (0.3)
	higher	2,510	80.2 (0.3)	12.0 (0.2)	19.9 (0.2)

Table 5: Utility vs Cost over five different algorithms (RANDOM, UNIFORM, SWAP, CACO, BRUTAS) and the actual admissions decisions made at our university. (Since CACO is a probabilistic method, the cost is given over a range of values.) For each of the algorithms, we give results assuming a cost/budget *lower*, roughly *equivalent*, and *higher* than that used by the real admissions committee. Both CACO and BRUTAS produce equivalent cohorts to the actual admissions process with lower cost, or produce high quality cohorts than the actual admissions process with equivalent cost. Our extension of SWAP to this multi-tiered setting also performs well relative to RANDOM and UNIFORM, but performs worse than both CACO and BRUTAS across the board.