

Ranked Prioritization of Groups in Combinatorial Bandit Allocation

Lily Xu¹, Arpita Biswas¹, Fei Fang² and Milind Tambe¹

¹Center for Research on Computation and Society, Harvard University

²Institute for Software Research, Carnegie Mellon University

lily_xu@g.harvard.edu, arpitabiswas@seas.harvard.edu, feif@cs.cmu.edu, milind_tambe@harvard.edu

Abstract

Preventing poaching through ranger patrols is critical for protecting endangered wildlife. Combinatorial bandits have been used to allocate limited patrol resources, but existing approaches overlook the fact that each location is home to multiple species in varying proportions, so a patrol benefits each species to differing degrees. When some species are more vulnerable, we ought to offer more protection to these animals; unfortunately, existing combinatorial bandit approaches do not offer a way to prioritize important species. To bridge this gap, (1) We propose a novel combinatorial bandit objective that trades off between reward maximization and also accounts for prioritization over species, which we call *ranked prioritization*. We show this objective can be expressed as a weighted linear sum of Lipschitz-continuous reward functions. (2) We provide RankedCUCB, an algorithm to select combinatorial actions that optimize our prioritization-based objective, and prove that it achieves asymptotic no-regret. (3) We demonstrate empirically that RankedCUCB leads to up to 38% improvement in outcomes for endangered species using real-world wildlife conservation data. Along with adapting to other challenges such as preventing illegal logging and overfishing, our no-regret algorithm addresses the general combinatorial bandit problem with a weighted linear objective.

1 Introduction

More than a quarter of mammals assessed by the IUCN Red List are threatened with extinction [Gilbert, 2008]. As part of the UN Sustainable Development Goals, Target 15.7 focuses on ending poaching, and Target 15.5 directs us to halt the loss of biodiversity and prevent the extinction of threatened species. To efficiently allocate limited resources, multi-armed bandits (MABs), and in particular combinatorial bandits [Chen *et al.*, 2016; Cesa-Bianchi and Lugosi, 2012], have been widely used for a variety of tasks [Bastani and Bayati, 2020; Segal *et al.*, 2018; Misra *et al.*, 2019] including ranger patrols to prevent poaching [Xu *et al.*, 2021a]. For




Figure 1: Animal density distributions and poaching risk (expected number of snares) across a 10×10 km region in Srepok Wildlife Sanctuary, Cambodia. The distribution of the species of greatest conservation concern, the elephant, differs from that of other animals and overall poaching risk.

poaching prevention, the patrol planner is tasked with repeatedly and efficiently allocating a limited number of patrol resources across different locations within the park [Plumptre *et al.*, 2014; Fang *et al.*, 2016].

In past work, we worked with the World Wide Fund for Nature (WWF) to deploy machine learning approaches for patrol planning in Srepok Wildlife Sanctuary in Cambodia [Xu *et al.*, 2020; Xu *et al.*, 2021b]. Subsequent conversations with park managers and conservation biologists raised the importance of protecting *vulnerable* species during patrols. For example, they pointed to the Uganda Wildlife Authority which ranked threatened species to allocate patrol effort in national parks, using the IUCN Red List as a key standard [Rodrigues *et al.*, 2006]. In this paper, we address directly consider the conservation of vulnerable species to better align our algorithmic objectives with on-the-ground conservation priorities.

Looking at real-world animal distribution and snare locations in Srepok (Figure 1), we observe that the locations that maximize expected reward, defined as finding the most snares (darkest red in the risk map), are not necessarily the regions with high density of endangered animals (elephants). To effectively improve conservation outcomes, we must account for these disparate impacts, as the relatively abundant muntjac and wildpig would benefit most if we simply patrol the regions with most snares, neglecting the endangered elephant. Prioritization of species is well-grounded in existing conservation best practices [Regan *et al.*, 2008; Arponen, 2012; Dilkina *et al.*, 2017]. For example, Plumptre *et al.* [2014] propose planning ranger patrols by setting conservation targets for key species. We term the goal of preferentially allocating resources to maximize benefit to priority groups as *ranked prioritization*.

Some existing multi-armed bandit approaches have considered priority-based fairness [Joseph *et al.*, 2016; Kearns *et al.*, 2017; Schumann *et al.*, 2022]. However, these prior works only consider stochastic, not combinatorial, bandits. In our combinatorial bandit problem setup, we must determine some hours of ranger patrol to allocate across N locations at each timestep, subject to a budget constraint. The rewards obtained from these actions are unknown *a priori* so must be learned in an online manner. Existing combinatorial bandit approaches [Chen *et al.*, 2016; Xu *et al.*, 2021a] achieve no-regret guarantees with respect to the objective of maximizing the direct sum of rewards obtained across all arms. However, we wish to directly consider ranked prioritization of groups in our objective; learning combinatorial rewards in an online fashion while ensuring we justly prioritize vulnerable groups along with reward makes the problem more challenging. *How to trade off reward and prioritization in the combinatorial bandit setting* has so far remained an open question. We show experimentally that straightforward solutions fail to make the appropriate trade-off.

To improve wildlife conservation outcomes by addressing the need for species prioritization, we contribute to the literature on online learning for resource allocation across groups with the following:

1. We introduce a novel problem formulation that considers resource allocation across groups with ranked prioritization, which has significant implications for wildlife conservation, as well as illegal logging and overfishing. We show that an objective that considers both prioritization and reward can be expressed as a weighted linear sum, providing the useful result that the optimal action in hindsight can be found efficiently via a linear program.
2. We provide a no-regret learning algorithm, RankedCUCB, for the novel problem of selecting a combinatorial action (amount of patrol effort to allocate to each location) at each timestep to simultaneously attain high reward and group prioritization. We prove that RankedCUCB achieves sub-linear regret of $O(\frac{\ln T}{NT} + \frac{N^J}{T})$ for a setting with N locations, $O(J^N)$ combinatorial actions, and time horizon T .
3. Using real species density and poaching data from Srepok Wildlife Sanctuary in Cambodia along with synthetic data, we experimentally demonstrate that RankedCUCB achieves up to 38% improvement over existing approaches in ranked prioritization for endangered wildlife.

2 Problem Formulation

Consider a protected area with N locations and G groups of interest, with each group representing a species or a set of species in the context of poaching prevention. Let d_{gi} denote the (known) fraction of animals of a group $g \in [G]$ present in location $i \in [N]$. Note that $\sum_{i=1}^N d_{gi} = 1$ for all $g \in [G]$. We assume the groups are disjoint, i.e., each animal is a member of exactly one group, and the impact of an action on a location equally benefits all animals that are present there.

At each timestep t until the horizon T , the planner determines an action, which is an *effort vector* $\vec{\beta}^{(t)} =$

$(\beta_1, \dots, \beta_N)$ specifying the amount of effort (e.g., number of patrol hours) to allocate to each location. The total effort at each timestep is constrained by a budget B such that $\sum_{i=1}^N \beta_i \leq B$. To enable implementation in practice, we assume the effort is discretized into J levels, thus $\beta_i \in \Psi = \{\psi_1, \dots, \psi_J\}$ for all i .

After taking action $\vec{\beta}$, the decision maker receives some reward $\mu(\vec{\beta})$. We assume the reward is decomposable [Chen *et al.*, 2016], defined as the sum of the reward at each location:

$$\text{expected reward} = \mu(\vec{\beta}) = \sum_{i=1}^N \mu_i(\beta_i). \quad (1)$$

We assume $\mu_i \in [0, 1]$ for all i . In the poaching context, the reward μ_i corresponds to the true probability of detecting snares at a location i . The reward function is initially unknown, leading to a learning problem. Following the model of combinatorial bandit allocation from Xu *et al.* [2021a], we assume that the reward function $\mu_i(\cdot)$ is (1) Lipschitz continuous, i.e., $|\mu_i(\psi_j) - \mu_i(\psi_k)| \leq L \cdot |\psi_j - \psi_k|$ for some Lipschitz constant L and all pairs ψ_j, ψ_k , and (2) monotonically non-decreasing such that exerting greater effort never results in a decrease in reward, i.e., $\beta'_i > \beta_i$ implies $\mu_i(\beta'_i) \geq \mu_i(\beta_i)$.

Finally, we assume that an ordinal ranking is known to indicate which group is more vulnerable. The ranking informs the planner which groups are of relatively greater concern, as we often lack sufficient data to quantify the extent to which one group should be prioritized. Without loss of generality, assume the groups are numerically ordered with $g = 1$ being the most vulnerable group and $g = G$ being the least, so the true rank is $\text{rank} = \langle 1, \dots, G \rangle$.

2.1 Measuring Ranked Prioritization

To evaluate prioritization, we must measure how closely the outcome of a combinatorial action $\vec{\beta}$ aligns with the priority ranking over groups. Later in Section 6, we discuss how our approach can generalize to other prioritization definitions, such as the case where we have greater specificity over the relative prioritization of each group.

To formalize prioritization, we define $\text{benefit}(\vec{\beta}; g)$ which quantifies the benefit group g receives from action $\vec{\beta}$. As mentioned earlier, we assume that reward $\mu_i(\beta_i)$ obtained by taking an action β_i impacts all individuals at location i . Let η_{gi} be the number of individuals from group g at location i and η_g be the total number of individuals in group g . We define

$$\begin{aligned} \text{benefit}(\vec{\beta}; g) &:= \frac{\sum_{i=1}^N \eta_{gi} \mu_i(\beta_i)}{\eta_g} \\ &\triangleq \sum_{i=1}^N d_{gi} \mu_i(\beta_i). \end{aligned} \quad (2)$$

An action $\vec{\beta}$ perfectly follows ranked prioritization when

$$\text{benefit}(\vec{\beta}; 1) \geq \text{benefit}(\vec{\beta}; 2) \geq \dots \geq \text{benefit}(\vec{\beta}; G).$$

However, when $\vec{\beta}$ does not follow perfect rank prioritization, we need a metric to measure the extent to which groups are appropriately prioritized. This metric would quantify similarity between the ranking induced by the benefits of $\vec{\beta}$ and the true rank. One common approach to evaluate the similarity between rankings is the *Kendall tau* coefficient:

$$\frac{(\# \text{ concordant pairs}) - (\# \text{ discordant pairs})}{\binom{G}{2}}. \quad (3)$$

The number of concordant pairs is the number of pairwise rankings that agree, across all $\binom{G}{2}$ pairwise rankings; the discordant pairs are those that disagree. This metric yields a value $[-1, 1]$ which signals amount of agreement, where $+1$ is perfect agreement and -1 is complete disagreement (i.e., reverse ordering). However, a critical weakness of Kendall tau is that it is discontinuous, abruptly jumping in value when the effort β_i on an arm is perturbed, rendering optimization difficult.

Fortunately, we have at our disposal not just each pairwise ranking, but also the *magnitude* by which each pair is in concordance or discordance. Leveraging this information, we construct a convex function $\mathcal{P}(\vec{\beta})$ that quantifies group prioritization:

$$\mathcal{P}(\vec{\beta}) = \frac{\sum_{g,h \in [G]} \mathbb{1}(g < h) \cdot (\text{benefit}(\vec{\beta}; g) - \text{benefit}(\vec{\beta}; h))}{\binom{G}{2}} \quad (4)$$

where $\mathbb{1}(g < h) = 1$ if g is of higher priority than h according to the true rank, and 0 otherwise. When $g < h$, the difference in benefit to group g and h is positive if the pairs are concordant and negative if the pairs are discordant. Intuitively, for each concordant pair, we earn $(\text{benefit}(\vec{\beta}; g) - \text{benefit}(\vec{\beta}; h))$, and for each discordant pair, we are penalized this quantity. The value $\mathcal{P}(\vec{\beta})$ can be interpreted as the correlation between the priority outcomes of $\vec{\beta}$ and the true priority ranking, giving us a metric for evaluating the group prioritization of an action $\vec{\beta}$.

Our prioritization metric $\mathcal{P}(\vec{\beta})$ induces nice properties. First, this metric is robust to perturbations in $\mu_i(\beta_i)$ values whereas Kendall tau is not. As $\mathcal{P}(\cdot)$ is continuous and differentiable, we can use it directly in the objective function. Second, with slight modifications to \mathcal{P} we can seamlessly adapt to other useful settings, as we discuss in Section 6.

Leveraging the fixed ordering $\text{rank} = \langle 1, \dots, G \rangle$ over groups, we can simplify Equation (4) as

$$\begin{aligned} \mathcal{P}(\vec{\beta}) &= \frac{\sum_{g=1}^{G-1} \sum_{h=g+1}^G (\text{benefit}(\vec{\beta}; g) - \text{benefit}(\vec{\beta}; h))}{\binom{G}{2}} \\ &\triangleq \frac{\sum_{g=1}^{G-1} \sum_{h=g+1}^G \left(\sum_{i=1}^N \mu_i(\beta_i) d_{gi} - \sum_{i=1}^N \mu_i(\beta_i) d_{hi} \right)}{\binom{G}{2}} \\ &= \sum_{i=1}^N \mu_i(\beta_i) \frac{\sum_{g=1}^{G-1} \sum_{h=g+1}^G (d_{gi} - d_{hi})}{\binom{G}{2}}. \end{aligned} \quad (5)$$

2.2 Objective with Prioritization and Reward

We wish to take actions that maximize our expected reward (Eq. 1) while also distributing our effort across the various groups as effectively as possible (Eq. 5). Recognizing that targeting group prioritization requires us to sacrifice reward, we set up the objective to balance reward and prioritization with a parameter $\lambda \in (0, 1]$ that enables us to tune the degree to which we emphasize reward vs. prioritization:

$$\text{obj}(\vec{\beta}) = \lambda \mu(\vec{\beta}) + (1 - \lambda) \cdot \mathcal{P}(\vec{\beta}). \quad (6)$$

3 Approach

We show that the objective (Eq. 6) can be reformulated as a weighted linear combination of the reward of each individual arm, which we then solve using our RankedCUCB algorithm, producing a general-form solution for combinatorial bandits.

Using the prioritization metric from Equation (5), we can re-express the objective (6) as

$$\text{obj}(\vec{\beta}) = \sum_{i=1}^N \mu_i(\beta_i) \left(\lambda + (1 - \lambda) \frac{\sum_{g=1}^{G-1} \sum_{h=g+1}^G (d_{gi} - d_{hi})}{\binom{G}{2}} \right) \quad (7)$$

Observe that the large parenthetical is comprised only of known constants, so we can precompute the parenthetical value as Γ_i for each location i to yield

$$\text{obj}(\vec{\beta}) = \sum_{i=1}^N \mu_i(\beta_i) \Gamma_i, \quad (8)$$

providing the useful result that our objective is decomposable across the locations $i \in [N]$ — specifically that the objective is a linear combination of the reward for each component of the action.

3.1 Combinatorial UCB Approach

The expected reward functions $\mu_i(\cdot)$ are unknown in advance and stochastic, so our setting falls within the space of sequential learning problems that need to be learned over multiple timesteps. Therefore, an action (an effort vector $\vec{\beta}$) we take in each timestep can be viewed as an arm pull under a multi-armed bandit framework. Specifically, we use a combinatorial bandit approach [Chen *et al.*, 2016], where each *arm* is an effort level ψ_j at one location i , rendering NJ total arms. We track the reward of each arm separately then use an oracle to select a combined *action* $\vec{\beta}$, which is a combinatorial selection of arms. Our goal is to design an algorithm that determines an action at each time t to minimize regret, defined as the difference between the total expected reward of an optimal action and the reward achieved following actions recommended by the algorithm.

To balance exploration (learning the reward of actions with greater uncertainty) with exploitation (taking actions known to be high-reward), we employ an upper confidence bound (UCB) approach that considers for each action its estimated reward along with our uncertainty over that estimate. Let $\hat{\mu}_t(i, j) = \text{reward}_t(i, \psi_j) / n_t(i, \psi_j)$ be the average observed reward of location i at effort ψ_j given cumulative empirical

reward $\text{reward}_t(i, \psi_j)$ over $n_t(i, \psi_j)$ arm pulls by timestep t . The *confidence radius* r_t of an arm (i, j) is then a function of the number of times we have pulled that arm:

$$r_t(i, j) = \sqrt{\frac{3\Gamma_i^2 \log T}{2n_t(i, \psi_j)}}. \quad (9)$$

We distinguish between UCB and a term we call SELFUCB. The SELFUCB of an arm (i, j) representing location i with effort ψ_j at time t is the UCB of an arm based only on its own observations, given by

$$\text{SELFUCB}_t(i, j) = \Gamma_i \hat{\mu}_t(i, j) + r_t(i, j). \quad (10)$$

This definition of SELFUCB corresponds with the standard interpretation of confidence bounds from the standard UCB1 algorithm [Auer *et al.*, 2002]. The UCB of an arm is then computed by taking the minimum of the bounds of all SELFUCBs as applied to the arm. These bounds are determined by adding the distance between arm (i, j) and all other arms (u, v) to the SELFUCB:

$$\text{UCB}_t(i, j) = \min_{k \in [J]} \{ \text{SELFUCB}_t(i, k) + L \cdot \text{dist} \} \quad (11)$$

$$\text{dist} = \max\{0, \psi_j - \psi_k\} \quad (12)$$

which leverages the assumptions described in Section 2 that the expected rewards are L -Lipschitz continuous and monotonically nondecreasing.

Given these UCB estimates for each arm (i, j) , we must now select a combinatorial action $\vec{\beta}$ to take at each timestep. As the prioritization metric can be expressed as a linear combination of the reward (8), we can directly optimize the overall objective using an integer linear program \mathcal{LP} (in Appendix A), which selects an optimal action that respects the budget constraint.

3.2 Trading Off Prioritization and Learning

A key challenge with our prioritization metric is that it is defined with respect to the true expected reward functions μ_i which are initially unknown. Instead, we estimate per-round prioritization based on our current *belief* of the true reward, $\hat{\mu}_i$, related to *subjective fairness* from the algorithmic fairness literature [Dimitrakakis *et al.*, 2017]. However, this belief is nonsensical in early rounds when our reward estimates are extremely coarse, so we discount the weight of prioritization in our objective until our learning improves.

Inspired by decayed epsilon-greedy, we incorporate an ε coefficient to tune how much we want to prioritize rank order at each step (versus learning to maximize reward):

$$\text{obj}(\vec{\beta}) = \lambda \mu(\vec{\beta}) + (1 - \lambda)(1 - \varepsilon) \cdot \mathcal{P}(\vec{\beta}). \quad (13)$$

For example, epsilon-greedy methods often use exploration probability $\varepsilon_t \sim t^{-1/3}$, which gradually attenuates at a decreasing rate with each increased timestep. Our definition gives nice properties that $\varepsilon_t = 1$ at $t = 1$, so we do not care about ranked prioritization at all (we have no estimate of the reward values so we cannot reasonably estimate ranking) but at $t = \infty$ then $\varepsilon_t = 0$ and we fully care about ranking w.r.t. λ (since we have full knowledge of the reward thus can precisely estimate rank order).

All together, we call the approach described here RankedCUCB and provide pseudocode in Algorithm 1.

Algorithm 1 RankedCUCB

Input: time horizon T , budget B , discretization levels $\Psi = \{\psi_1, \dots, \psi_J\}$, arms $i \in [N]$ with unknown reward μ_i
Parameters: tuning parameter λ

- 1: Precompute Γ_i for each arm $i \in [N]$
 - 2: $n(i, \psi_j) = 0, \text{reward}(i, \psi_j) = 0 \quad \forall i \in [N], j \in [J]$
 - 3: **for** timestep $t = 1, 2, \dots, T$ **do**
 - 4: Let $\varepsilon_t = t^{-1/3}$
 - 5: Compute UCB_t for all arms using Eq. (11)
 - 6: Solve $\mathcal{LP}(\text{UCB}_t, \{\Gamma_i\}, B)$ to select super arm $\vec{\beta}$
 - 7: // *Execute action*
 - 8: Act on $\vec{\beta}'$ to observe rewards $X_1^{(t)}, X_2^{(t)}, \dots, X_N^{(t)}$
 - 9: **for** arm $i = 1, \dots, N$ **do**
 - 10: $\text{reward}(i, \beta'_i) = \text{reward}(i, \beta'_i) + X_i^{(t)}$
 - 11: $n(i, \beta_i) = n(i, \beta_i) + 1$
-

4 Regret Analysis

We now prove that our iterative algorithm RankedCUCB (Alg. 1) guarantees no regret with respect to the optimal solution for the objective (8) that jointly considers reward and prioritization. More formally, we show that $\text{Regret}_T \rightarrow 0$ as $T \rightarrow \infty$ where

$$\text{Regret}_T := \mu(\vec{\beta}^*) - \frac{1}{T} \sum_{t=1}^T \mu(\vec{\beta}^{(t)}). \quad (14)$$

Here, $\vec{\beta}^*$ is an optimal action and expected reward $\mu(\vec{\beta}) := \sum_{i=1}^N \mu_i(\beta_i) \Gamma_i$ for an effort vector $\vec{\beta} = \{\beta_1, \dots, \beta_N\}$. Note that if $\Gamma_i < 0$, any solution to the maximization problem (8) would allocate $\beta_i = 0$, and so would RankedCUCB. Hence, for the analysis we assume that we consider only those locations whose $\Gamma_i > 0$.

4.1 Convergence of Estimates

To prove the no-regret guarantee of RankedCUCB, we first establish Lemma 1 which states that, with high probability, the UCB estimate $\hat{\mu}_t(i, j) \Gamma_i + r_t(i, j)$ converges within a confidence bound of $r_t(i, j) = \sqrt{(3\Gamma_i^2 \log T)/(2n_t)}$.

Lemma 1. *Using RankedCUCB, after t timesteps, each $\hat{\mu}_t(i, j) \Gamma_i$ estimate converges to a value within a radius $r_t(i, j) = \sqrt{3\Gamma_i^2 \ln t / 2n_t(i, j)}$ of the corresponding true $\mu_t(i, j) \Gamma_i$ values with probability $1 - \frac{2NJ}{t^2}$ for all i, j .*

Proof sketch. Using the Chernoff–Hoeffding bound, we show that, at any timestep t , the probability that the difference between $\mu_t(i, j) \Gamma_i$ and $\hat{\mu}_t(i, j) \Gamma_i$ is greater than $r_t(i, j)$, is at most $2/t^3$. We then use union bound to show that $\hat{\mu}_t(i, j) \Gamma_i$ converges to a value within radius $r_t(i, j)$ of $\mu_t(i, j) \Gamma_i$ with probability $1 - \frac{2NJ}{t^2}$. The complete proof is given in Appendix B.2. \square

4.2 Achieving No Regret

Theorem 1. *The cumulative regret of RankedCUCB is $O(\frac{J \ln T}{N} + NJ)$ with N arms, J discrete effort values, and time horizon T .*

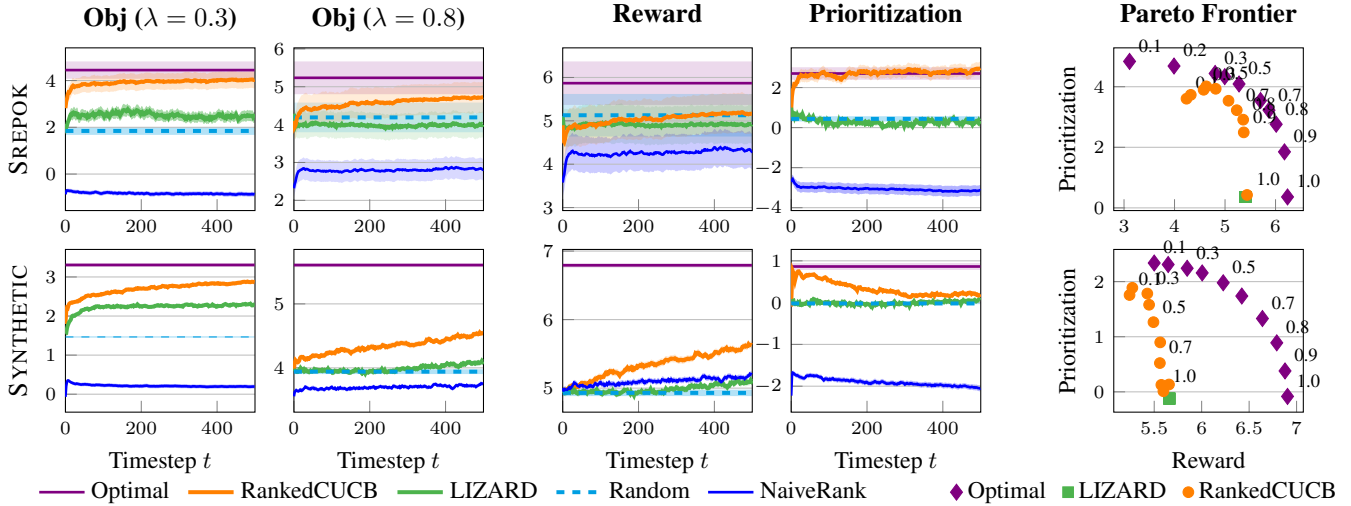


Figure 2: The performance of each approach. **LEFT** evaluates the objective with tuning parameter $\lambda = 0.3$ and $\lambda = 0.8$. Our approach, RankedCUCB, performs significantly better than baselines. **CENTER** evaluates reward and prioritization (at $\lambda = 0.8$), the two components of the combined objective. The reward-maximizing LIZARD algorithm rapidly maximizes reward but performs worse than random in terms of rank order. **RIGHT** visualizes the Pareto frontier trading off the two components of our objective. Labels represent different values of $\lambda \in \{0.1, 0.2, \dots, 1.0\}$. Each point plots the reward and ranked prioritization as the average of the final 10 timesteps. All results are averaged over 30 random seeds and smoothed with an exponential moving average.

Proof sketch. We first show the expected regret (Eq. 14) of RankedCUCB can be redefined as:

$$\frac{1}{T} \sum_{i=1}^N \sum_{j=1}^J \mathbb{E}[L_{i,j,T}] \zeta_{i,j}$$

where $L_{i,j,T}$ specifies the number of times the effort ψ_j selected for location i results in a suboptimal solution, and $\zeta_{i,j}$ denotes the minimum loss incurred due to the suboptimal selection. Then, by contradiction, we show that when all $\hat{\mu}_t(i, j)\Gamma_i$ estimates are within their confidence radius, RankedCUCB selects an optimal effort ψ_j (and not a suboptimal one) at time t . Using this fact and Lemma 1, we show that the $\mathbb{E}[L_{i,j,T}]$ can be upper bounded by a finite term, and hence the expected regret is $O(\frac{J \ln T}{N} + NJ)$. The complete proof is given in Appendix B.3. \square

Our analysis improves upon the regret analysis for standard (non-weighted) combinatorial bandits [Chen *et al.*, 2016] by not having to rely on the “bounded smoothness” assumption of reward functions. The technical tools we employ for this regret analysis can be of independent interest to the general multi-armed bandits community.

5 Experiments

We evaluate the performance of our RankedCUCB approach in real-world conservation data from Srepok Wildlife Sanctuary to demonstrate significant improvement over approaches that ignore prioritization.¹

Of the key species in Srepok on which we have species density data, we prioritize elephants, followed by banteng,

muntjac, and wild pig, corresponding to their IUCN Red List categories of critically endangered, endangered, least concern but decreasing, and least concern respectively. Thus our setting has $G = 4$ groups distributed across $N = 25$ locations within the park. We further evaluate our algorithm on a synthetic dataset with $G = 5$ groups randomly distributed across the park.

Baselines. We compare the performance of our *RankedCUCB* approach with naive priority-aware, priority-blind, random, and optimal benchmarks. *NaiveRank* takes the straightforward approach of directly solving for the objective that weighs each target by its individual ranked prioritarian metric, which accounts for prioritization induced by each target independently but ignores the coupled effect across targets. *LIZARD* [Xu *et al.*, 2021a] maximizes the combinatorial reward objective; this algorithm enhances standard UCB approaches by accounting for smoothness and monotonicity in the reward function, but ignores prioritization. *Optimal* solves for the action $\beta^* = \arg \max_{\beta} \text{obj}(\beta)$ that maximizes the objective, using the ground truth values of μ_i to directly compute the best action. *Random* selects an arbitrary subset of actions that satisfies the budget constraint.

See Figure 2 for a comparison of the approaches. RankedCUCB performs consistently the best on our overall objective, and the breakdown of the reward and prioritization components reveals that this gain is a result of the tradeoff between the two components: although LIZARD is able to learn high-reward actions, these actions lead to prioritization outcomes that are worse than random. LIZARD even achieves reward that exceeds that of the optimal action (which also considers rank priority), but that comes at the cost of poor

¹Code is available at <https://github.com/lily-x/rankedCUCB>

prioritization. Notably, NaiveRank performs worse than random, even measured on fairness: focusing on the individual targets with the best prioritization neglects group-wide patterns throughout the park.

The reward–prioritization tradeoff becomes more apparent when we analyze the Pareto frontier in Figure 2(right), which plots the reward and prioritization of each approach as we change the value of λ . The price of prioritization is clear: the more we enforce prioritization (smaller values of λ) the lower our possible reward. However, this tradeoff is not one-for-one; the steepness of the Pareto curve allows us to strategically select the value of λ . For example, we achieve a nearly two-fold increase in prioritization by going from $\lambda = 0.9$ to 0.7 while sacrificing less than 10% of the reward, indicating a worthwhile tradeoff.

6 Generalizability to Other Settings

Beyond wildlife conservation, our problem formulation applies to domains with the following characteristics: (1) repeated resource allocation, (2) multiple groups with some priority ordering, (3) an *a priori* unknown reward function that must be learned over time, and (4) actions that impact some subset of the groups to differing degrees. Our approach also adapts to the following related objectives.

Weighted rank. Gatmiry *et al.* [2021] suggest some specific metrics for how rangers should prioritize different species, noting ranger enforcement of poaching (a wild sheep) should be 2.5 stricter than red deer. Our approach could easily adapt to these settings where domain experts have more precise specification of desired outcomes. In this case, we introduce a parameter α_g for each group g that specifies the relative importance of group g , leading to the following prioritization metric:

$$\mathcal{P}_{\text{CF}}(\vec{\beta}) = \sum_{i=1}^N \mu_i(\beta_i) \frac{\sum_{g=1}^{G-1} \sum_{h=g+1}^G \alpha_g d_{gi} - \alpha_h d_{hi}}{\binom{G}{2}} \quad (15)$$

where we set $\alpha_g > \alpha_h$ when $g < h$ to raise our threshold of how strongly group g should be favored. From our example, we would set $\alpha_{\text{urial}} = 2.5\alpha_{\text{deer}}$. This metric aligns with the definition of *calibrated fairness* [Liu *et al.*, 2017].

Weighted reward. We can also accommodate the setting where the reward is a weighted combination, i.e., if reward is $\sum_{i=1}^N c_i \mu_i(\beta_i)$ for some coefficients $c_i \in \mathbb{R}$, as the coefficients c_i would be absorbed into the Γ_i term of the objective.

7 Related Work

Multi-armed bandits. MABs [Lattimore and Szepesvári, 2020] have been applied to resource allocation for healthcare [Bastani and Bayati, 2020], education [Segal *et al.*, 2018], and dynamic pricing [Misra *et al.*, 2019]. These papers solve various versions of the stochastic MAB problem [Auer *et al.*, 2002].

Several prior works consider resource allocation settings where each arm pull is costly, limiting the total number of pulls by a budget. Tran-Thanh *et al.* [2010] use an ϵ -based

approach to achieve regret linear in the budget B ; Tran-Thanh *et al.* [2012] provide a knapsack-based policy using UCB to improve to $O(\ln B)$ regret. Chen *et al.* [2016] extend UCB1 to a combinatorial setting, matching the regret bound of $O(\log T)$. Slivkins [2013] address ad allocation where each stochastic arm is limited by a budget constraint on the number of pulls. Our objective (8) is similar to theirs but they do not handle combinatorial actions, which makes our regret analysis even harder.

Prioritization notions in MABs. Addressing prioritization requires reward trade-offs, which is related to objectives from algorithmic fairness [Dwork *et al.*, 2012; Kleinberg *et al.*, 2018; Corbett-Davies *et al.*, 2017]. However, most MAB literature favor traditional reward-maximizing approaches [Lattimore *et al.*, 2015; Agrawal and Devanur, 2014; Verma *et al.*, 2019]. The closest metric to our prioritization objective is *meritocratic fairness* [Kearns *et al.*, 2017]. For MAB, meritocratic fairness was introduced by Joseph *et al.* [2016], to ensure that lower-reward arms are never favored above a high-reward one. Liu *et al.* [2017] apply calibrated fairness by enforcing a smoothness constraint on our selection of similar arms, using Thompson sampling to achieve an $\tilde{O}(T^{2/3})$ fairness regret bound. For contextual bandits, Chohlas-Wood *et al.* [2021] take a consequentialist (outcome-oriented) approach and evaluate possible outcomes on a Pareto frontier, which we also investigate. Wang *et al.* [2021] ensure that arms are proportionally represented according to their merit, a metric they call *fairness of exposure*. Others have considered fairness in multi-agent bandits [Hossain *et al.*, 2021], restless bandits [Herlihy *et al.*, 2021], sleeping bandits [Li *et al.*, 2019], contextual bandits with biased feedback [Schumann *et al.*, 2022], and infinite bandits [Joseph *et al.*, 2018]. However, none of these papers consider correlated groups, which is our focus.

8 Conclusion

We address the challenge of allocating limited resources to prioritize endangered species within protected areas where true snare distributions are unknown, closing a key gap identified by our conservation partners at Srepok Wildlife Sanctuary. Our novel problem formulation introduces the metric of *ranked prioritization* to measure impact across disparate groups. Our RankedCUCB algorithm offers a principled way of balancing the tradeoff between reward and ranked prioritization in an online learning setting. Notably, our theoretical guarantee bounding the regret of RankedCUCB applies to a broad class of combinatorial bandit problems with a weighted linear objective.

Acknowledgments

This work was supported in part by NSF grant IIS-2046640 (CAREER) and the Army Research Office (MURI W911NF-18-1-0208). Biswas supported by the Center for Research on Computation and Society (CRCS). Thank you to Andrew Plumptre for a helpful discussion on species prioritization in wildlife conservation; Jessie Finocchiaro for comments on an earlier draft; and all the rangers on the front lines of biodiversity efforts in Srepok and around the world.

References

- [Agrawal and Devanur, 2014] Shipra Agrawal and Nikhil R Devanur. Bandits with concave rewards and convex knapsacks. In *Economics and Computation*, 2014.
- [Arponen, 2012] Anni Arponen. Prioritizing species for conservation planning. *Biodiversity and Conservation*, 21(4), 2012.
- [Auer *et al.*, 2002] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 2002.
- [Bastani and Bayati, 2020] Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *OR*, 68(1), 2020.
- [Cesa-Bianchi and Lugosi, 2012] Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5), 2012.
- [Chen *et al.*, 2016] Wei Chen, Yajun Wang, Yang Yuan, and Qinshi Wang. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *JMLR*, 17(1), 2016.
- [Chohlas-Wood *et al.*, 2021] Alex Chohlas-Wood, Madison Coots, Emma Brunskill, and Sharad Goel. Learning to be fair: A consequentialist approach to equitable decision-making. *arXiv*, 2021.
- [Corbett-Davies *et al.*, 2017] Sam Corbett-Davies, Emma Pierson, Avi Feller, Sharad Goel, and Aziz Huq. Algorithmic decision making and the cost of fairness. In *KDD*, 2017.
- [Dilkina *et al.*, 2017] Bistra Dilkina, Rachel Houtman, et al. Tradeoffs and efficiencies in optimal budget-constrained multispecies corridor networks. *Conservation Biology*, 2017.
- [Dimitrakakis *et al.*, 2017] Christos Dimitrakakis, Yang Liu, David Parkes, and Goran Radanovic. Subjective fairness: Fairness is in the eye of the beholder. Technical report, 2017.
- [Dwork *et al.*, 2012] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *ITCS*, 2012.
- [Fang *et al.*, 2016] Fei Fang, Thanh Hong Nguyen, Rob Pickles, Wai Y Lam, Gopalasamy R Clements, Bo An, Amandeep Singh, Milind Tambe, et al. Deploying PAWS: Field optimization of the Protection Assistant for Wildlife Security. In *IAAI*, 2016.
- [Gatmiry *et al.*, 2021] Zohreh S Gatmiry, Ashkan Hafezalkotob, Roya Soltani, et al. A security game approach for strategic conservation against poaching considering food web complexities. *Ecological Complexity*, 2021.
- [Gilbert, 2008] Natasha Gilbert. A quarter of mammals face extinction. *Nature*, 455(7214):717–718, 2008.
- [Herlihy *et al.*, 2021] Christine Herlihy, Aviva Prins, Aravind Srinivasan, and John Dickerson. Planning to fairly allocate: Probabilistic fairness in the restless bandit setting. *arXiv*, 2021.
- [Hossain *et al.*, 2021] Safwan Hossain, Evi Micha, and Nisarg Shah. Fair algorithms for multi-agent multi-armed bandits. In *NeurIPS*, 2021.
- [Joseph *et al.*, 2016] Matthew Joseph, Michael Kearns, Jamie H Morgenstern, and Aaron Roth. Fairness in learning: Classic and contextual bandits. In *NeurIPS*, 2016.
- [Joseph *et al.*, 2018] Matthew Joseph, Michael Kearns, Jamie Morgenstern, Seth Neel, and Aaron Roth. Meritocratic fairness for infinite and contextual bandits. In *AIES*, 2018.
- [Kearns *et al.*, 2017] Michael Kearns, Aaron Roth, and Zhiwei Steven Wu. Meritocratic fairness for cross-population selection. In *ICML*, 2017.
- [Kleinberg *et al.*, 2018] Jon Kleinberg, Jens Ludwig, Sendhil Mullainathan, and Ashesh Rambachan. Algorithmic fairness. In *AEA Papers and Proceedings*, 2018.
- [Lattimore and Szepesvári, 2020] Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- [Lattimore *et al.*, 2015] Tor Lattimore, Koby Crammer, and Csaba Szepesvári. Linear multi-resource allocation with semi-bandit feedback. In *NeurIPS*, 2015.
- [Li *et al.*, 2019] Fengjiao Li, Jia Liu, and Bo Ji. Combinatorial sleeping bandits with fairness constraints. *IEEE Transactions on Network Science and Engineering*, 7(3), 2019.
- [Liu *et al.*, 2017] Yang Liu, Goran Radanovic, Christos Dimitrakakis, Debmalaya Mandal, and David C Parkes. Calibrated fairness in bandits. *FAT-ML Workshop*, 2017.
- [Misra *et al.*, 2019] Kanishka Misra, Eric M Schwartz, and Jacob Abernethy. Dynamic online pricing with incomplete information using multiarmed bandit experiments. *Marketing Science*, 2019.
- [Plumtre *et al.*, 2014] Andrew J Plumtre, Richard A Fuller, Aggrey Rwetsiba, Fredrick Wanyama, Deo Kujirakwinja, Margaret Driciru, et al. Efficiently targeting resources to deter illegal activities in protected areas. *J. of Appl. Ecology*, 51(3), 2014.
- [Regan *et al.*, 2008] Helen Regan, Lauren Hierl, et al. Species prioritization for monitoring and management in regional multiple species conservation plans. *Diversity and Distributions*, 2008.
- [Rodrigues *et al.*, 2006] Ana SL Rodrigues, John D Pilgrim, John F Lamoreux, Michael Hoffmann, et al. The value of the IUCN Red List for conservation. *Trends Ecol. Evol.*, 21(2), 2006.
- [Schumann *et al.*, 2022] Candice Schumann, Zhi Lang, Nicholas Mattei, and John P Dickerson. Group fairness in bandit arm selection. *AAMAS*, 2022.
- [Segal *et al.*, 2018] Avi Segal, Yossi Ben David, Joseph Jay Williams, Kobi Gal, and Yaar Shalom. Combining difficulty ranking with multi-armed bandits to sequence educational content. In *AIED*, 2018.
- [Slivkins, 2013] Aleksandrs Slivkins. Dynamic ad allocation: Bandits with budgets. *arXiv*, 2013.
- [Tran-Thanh *et al.*, 2010] Long Tran-Thanh, Archie Chapman, Enrique Munoz De Cote, Alex Rogers, et al. Epsilon-first policies for budget-limited multi-armed bandits. In *AAAI*, 2010.
- [Tran-Thanh *et al.*, 2012] Long Tran-Thanh, Archie Chapman, Alex Rogers, and Nicholas Jennings. Knapsack based optimal policies for budget-limited multi-armed bandits. In *AAAI*, 2012.
- [Verma *et al.*, 2019] Arun Verma, Manjesh K. Hanawal, Arun Rajkumar, et al. Censored semi-bandits: A framework for resource allocation with censored feedback. In *NeurIPS*, 2019.
- [Wang *et al.*, 2021] Lequn Wang, Yiwei Bai, Wen Sun, and Thorsten Joachims. Fairness of exposure in stochastic bandits. In *ICML*, 2021.
- [Xu *et al.*, 2020] Lily Xu, Shahrzad Gholami, Sara Mc Carthy, Bistra Dilkina, Andrew Plumtre, et al. Stay ahead of poachers: Illegal wildlife poaching prediction and patrol planning under uncertainty with field test evaluations. In *ICDE*, 2020.
- [Xu *et al.*, 2021a] Lily Xu, Elizabeth Bondi, Fei Fang, Andrew Perrault, Kai Wang, and Milind Tambe. Dual-mandate patrols: Multi-armed bandits for green security. In *AAAI*, 2021.
- [Xu *et al.*, 2021b] Lily Xu, Andrew Perrault, Fei Fang, Haipeng Chen, and Milind Tambe. Robust reinforcement learning under minimax regret for green security. In *UAI*, 2021.