

# Multi-Armed Bandits

Given on November 21, 2024

Assignment 1

- The primary objective of this assignment is to evaluate your comprehension and application of multi-armed bandit algorithms. This task demonstrates your ability to understand the core concepts, implement the algorithms, and analyze their performance.
- The following assignment is just a modified version of the problem seen during the practical session.
- This is an individual assignment. You are not allowed to discuss the problems with other students.
- Write your code only in the .py files provided. Avoid creating new files. Also, do not change the input or output structure of the functions.
- Make sure you include Debabrota Basu in cc, and be sure to submit:
  1. A '.zip' file containing all your Python codes (in .py format).
  2. A '.pdf' file that contains your answers to the questions and generated plots
- You cannot use ChatGPT or any other code assistance tool for the programming part. If you use ChatGPT to edit the grammar in your report, you have to explicitly state it in the report.
- If you have questions regarding the assignment, you can send me an email.

## 1 Problem 1

### Aragorn's Trials: The Quest to Find Middle-earth's True Champion

In the twilight of the Third Age, as shadows deepen over Middle-earth, a dire threat looms over all free peoples. Sauron, the Dark Lord of Mordor, has regained much of his former strength, and his malice spreads across the lands like a rising tide. In these dark times, hope rests on a knife's edge, and the fate of the world depends on the courage and strength of its heroes. Among these heroes stands Aragorn, son of Arathorn, the rightful heir to the throne of Gondor. As a leader of the Free Peoples, Aragorn understands the gravity of the coming war. He knows that to stand against Sauron, Middle-earth must place its trust in the mightiest of champions, a hero capable of overcoming the most perilous quests and leading the forces of good to victory. But the paths of Middle-earth are fraught with danger, and the true strength of each hero is veiled in mystery. From the noble halls of Rivendell to the rugged wilds of Rohan, heroes of all races have answered the call—Men, Elves, Dwarves, and even Hobbits. Yet who among them is worthy to be named the Champion of Middle-earth? Aragorn must decide, for the time of reckoning draws near. Aragorn's goal is to identify the hero with the highest probability of successfully completing the most dangerous and crucial quests. However, the true potential of each hero is unknown at the outset. Through trials and quests, Aragorn must gather knowledge of their abilities, learning from their successes and failures, to ultimately choose the one hero who will lead the charge against Sauron's dark forces. As a trusted advisor to Aragorn, you have been entrusted with the sacred duty of devising a strategy to discover the greatest hero of Middle-earth. Through a series of trials, you will guide Aragorn in deciding which hero to send on quests, using both exploration to uncover hidden strengths and exploitation to maximize the power of known talents

#### Details:

- Heroes and Quests:
  1. A fellowship of  $N$  heroes, each a legend in their own right, has assembled to face the trials. Among them are the noblemen of Gondor, the wise Elves of Lothlórien, the stout Dwarves of the Lonely Mountain, and the brave Hobbits of the Shire.
  2. The true prowess of each hero is represented by a success probability  $\pi_i$ , reflecting their chances of triumph in any given quest. This probability is not known at the beginning and must be inferred from their performance in the trials.
  3. When a hero is chosen to undertake a quest—whether it be navigating the treacherous paths of Mirkwood, retrieving a lost relic from the depths of Moria, or defending the borders of Rohan from marauding Orcs—the outcome is determined by their true success probability  $\pi_i$ , drawn from a Bernoulli distribution. A success

means the hero has overcome the challenge, while a failure indicates they were not up to the task.

- **Quests and Outcomes:**
  1. In each round, Aragorn must choose one hero to send on a quest. After each quest, you will observe whether the hero succeeded or failed, using this knowledge to guide future decisions.
  2. The trials require a delicate balance between exploring the potential of lesser known heroes and exploiting the abilities of those who have already proven their worth. This balance is crucial, as it will allow Aragorn to gather enough information to make an informed decision about who can lead Middle-earth in the final battle.
  3. By the end of the trials, Aragorn must identify the hero with the highest estimated probability of success—the one who will become the Champion of Middle Earth, standing by his side as they march towards the ultimate confrontation with Sauron.

We will use the multi-armed bandit (MAB) algorithms to help Aragorn. Please use the code `_middle_earth.zip` project and implement or answer the questions below. Make sure to upload all the `.py` files when you are done.

## Heroes:

1. (5 points) Completing the Heroes class in `heroes.py`

The Heroes class has been partially implemented. Your task is to complete the `attempt_quest` method to fulfill the following requirements:

  - Update the total quests and successes for the given hero.
  - Return the reward of the quest: return 1 for success and 0 otherwise.

---

**Example Setup:** From now on, let's work with a specific set of heroes:

```
heroes = Heroes(total_quests=3000, true_probability_list=[0.35, 0.6, 0.1])
```

In this setup:

- (a) `total_quests=3000` specifies the total number of quests you want to simulate.
- (b) `true_probability_list=[0.35, 0.6, 0.1]` defines the true success probabilities for each hero. The first hero has a 35% chance of success, the second hero has a 60% chance, and the third hero has a 10% chance.

### $\epsilon$ -greedy Method

2. (20 points): Completing the `eps_greedy` method in `eps_greedy.py`
  - (a) (10 points) The `eps_greedy` method has been partially implemented. Your task is to complete the missing parts according to the following requirements:
    - Accurately define the values of optimal reward and optimal hero index based on the true success probabilities.
    - Implement the  $\epsilon$ -greedy action-selection strategy, ensuring that the exploration-exploitation balance is properly managed.
    - Return the following lists:
      - `rew_record`: A list of rewards received at each attempt.
      - `avg_ret_record`: The running average of rewards at each attempt.
      - `tot_reg_record`: The cumulative regret at each attempt, reflecting the difference between the optimal reward and the actual reward.
      - `opt_action_record`: The percentage of times the optimal hero was selected at each attempt.
  - (b) Execute the command `python eps_greedy.py`. Be aware that this process may take some time to complete.
  - (c) (5 points) The first experiment investigates the impact of different  $\epsilon$  values from the set `[0.2, 0.1, 0.01, 0.]`. Once the experiment is complete, locate and report the file `results/epsilon_greedy_various_epsilons.pdf`. Provide an analysis of the results, focusing on how varying  $\epsilon$  values influence exploration behavior and overall performance. (Provide your plots and explanation in the `.pdf` report file.)
  - (d) (5 points) The second experiment explores the effect of optimistic initial values with  $\epsilon = 0$ . After running the experiment, find and report the file `results/epsilon_greedy_various_init_values.pdf`. Discuss how different initial value estimates impact both exploration and the effectiveness of the policy. (Provide your plots and explanation in the `.pdf` report file.)

### Upper Confidence Bound Method (UCB)

3. (15 points): Completing the `ucb` method in `ucb.py`:
  - (a) (10 points) The `ucb` method has been partially implemented. Your task is to complete the missing parts according to the following requirements:
    - Accurately define the values of `optimal_reward` and `optimal_hero_index` based on the true success probabilities.
    - Implement the Upper-Confidence-Bound (UCB) action-selection strategy.

- Return the following lists:
  - `rew_record`: A list of rewards received at each attempt.
  - `avg_ret_record`: The running average of rewards at each attempt.
  - `tot_reg_record`: The cumulative regret at each attempt, reflecting the difference between the optimal reward and the actual reward.
  - `opt_action_record`: The percentage of times the optimal hero was selected at each attempt.
- (b) Execute the command `python ucb.py`. Be aware that this process may take some time to complete.
- (c) (5 points) Our experiment investigates the impact of different  $c$  (coefficient) values from the set  $[0.0, 0.5, 2]$ . Once the experiment is complete, locate and report the file `results/ucb_various_c_values.pdf`. Provide an analysis of the results, focusing on how varying  $c$  values influence exploration behavior and overall performance. (Provide your plots and explanation in the `.pdf` report file.)

## 2 Problem 2: Contextual bandits

A supervised classification problem  $(X, y)$  with  $X$  the dataset features in  $\mathbb{R}^{N \times d}$  and  $y$  the classes vector in  $\{1, \dots, K\}^N$  can be formulated as a *contextual bandits problem* formalized as follows. At time  $t$ , a learner observes a context vector  $x_n$  drawn at random from  $X$ . The learner then pulls an arm  $a \in \mathcal{A} = \{1, \dots, K\}$  based on  $x_n$  and receives a binary reward  $\mathbb{I}(a = y_n)$ . Learning a function  $f : X \rightarrow \mathcal{A}$  that maximizes the expected cumulative reward  $\mathbb{E}_{x_n \sim X} \left[ \sum_{t=0}^T \mathbb{I}(f(x_n) = y_n) \right]$  is learning a classifier that minimizes the misclassification rate.

For reference, see section 5.1 from Bandits Algorithms Book and following of this paper.  
Refer to the instructions provided in the notebook.