

Problem Set 8

Due: Friday, June 3, by 11:59pm. (No extensions / late days allowed)

Instructions

Solutions format, collaboration policy, and late policy. See PSet 1 for further details. The same requirements and policies still apply. Also follow the typesetting instructions from the prior PSets.

Solutions submission. You must submit your solution via Gradescope. In particular:

- For the solutions to Tasks 1-5, and the written part of Task 6, submit under "PSet 8 [Written]" a *single* PDF file containing the solution to all tasks in the PSet. Each numbered task should be solved on its own page (or pages). Follow the prompt on Gradescope to link tasks to your pages. Do not write your name on the individual pages – Gradescope will handle that.
- For the programming part (sub-task **c**) of Task 6), submit your code under "PSet 8 [Coding]" as a file called `mcmc_knapsack.py`.

Task 1 – Elections

[12 pts]

Individuals in a certain country are voting in an election between 3 candidates: A , B and C . Suppose that each person makes their choice independent of others and votes for candidate A with probability θ_1 , for candidate B with probability $2\theta_2$ and for candidate C with probability $1 - \theta_1 - 2\theta_2$. (Thus, $0 \leq \theta_1 + 2\theta_2 \leq 1$.) The parameters θ_1, θ_2 are unknown.

Suppose that x_1, \dots, x_n are n independent, identically distributed samples from this distribution. (Let $n_A =$ number of x_i 's equal to A , let $n_B =$ number of x_i 's equal to B , and let $n_C =$ number of x_i 's equal to C .) What are the maximum likelihood estimates for θ_1 and θ_2 in terms of n_A, n_B , and n_C ?

(You don't need to check second order conditions.)

Task 2 – (Un)biased Estimation

[12 pts]

- Let x_1, \dots, x_n be independent samples from the Poisson distribution with parameter θ . In the concept check, we have seen that the MLE estimator for θ is $\hat{\theta} = \frac{1}{n} \sum_{i=1}^n x_i$. Is this estimator unbiased?
- Let x_1, \dots, x_n be independent samples from $\text{Unif}(0, \theta)$, the continuous uniform distribution on $[0, \theta]$. Then, consider the estimator $\hat{\theta}_{\text{first}} = 2x_1$, i.e., our estimator ignores the samples x_2, \dots, x_n and just outputs twice the value of the first sample. Is $\hat{\theta}_{\text{first}}$ unbiased?

Task 3 – Uniform Distribution

[24 pts]

Let x_1, x_2, \dots, x_n be independent samples from $\text{Unif}(0, \theta)$, the continuous uniform distribution on $[0, \theta]$. Here, θ is the unknown parameter.

- Write down the likelihood function $\mathcal{L}(x_1, \dots, x_n \mid \theta)$.
- Explain why $\hat{\theta} = \max\{x_1, \dots, x_n\}$ is the MLE estimator for θ .

To this end, explain why this value of $\hat{\theta}$ maximizes $\mathcal{L}(x_1, x_2, \dots, x_n \mid \theta)$ directly, by describing the behavior of the likelihood function (as a function of θ) and inferring the maximum from this description, rather than trying to use calculus.

- c) For the remaining parts of this question we suppose that this uniform distribution is on $[0, \theta]$ for a particular fixed $\theta > 0$, and let

$$\hat{\Theta}_n = \hat{\theta}(X_1, \dots, X_n) = \max\{X_1, \dots, X_n\}$$

be a *random variable*, where X_1, \dots, X_n are independent and follow the uniform distribution on $[0, \theta]$.

Compute the CDF $F_{\hat{\Theta}_n}(x)$.


Hint: Focus first on the interval $0 \leq x \leq \theta$, but when you're done with that, don't forget to also define $F_{\hat{\Theta}_n}(x)$ on the rest of the real numbers.

- d) From your answer to c) compute the probability density function $f_{\hat{\Theta}_n}(x)$ of $\hat{\Theta}$.
- e) From your answer to d), compute $\mathbb{E}[\hat{\Theta}_n]$. Why is $\hat{\theta}$ not an unbiased estimator of θ ?
- f) Starting from the value of $\mathbb{E}[\hat{\Theta}]$ you computed in e), show how to make a small modification to the MLE studied so far to produce an unbiased estimator of θ and explain why it is unbiased.
- Note:** There also are other unbiased estimators of θ , but we want you to produce one that directly uses what you have inferred in e).

Task 4 – Markov Chains

[16 pts]

In this problem we will explore the Markov chain shown in the following figure:



- a) Give the transition probability matrix (TPM) \mathbf{M} of the Markov chain in the above figure.
- b) Recall that $X^{(t)}$ denotes the state of the Markov chain in the t -th step. Use the law of total probability to compute $\mathbb{P}(X^{(2)} = 4 \mid X^{(0)} = 1)$. Show your work.
- c) Suppose that the initial state is uniformly distributed, i.e., we represent its distribution through the vector

$$\mathbf{q}^{(0)} = \left[\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4} \right].$$

Compute $\mathbf{q}^{(0)}\mathbf{M}$, the vector-matrix-product of $\mathbf{q}^{(0)}$ and the TPM \mathbf{M} as **4 simplified fractions**.


Describe what each component of $\mathbf{q}^{(0)}\mathbf{M}$ represents.

- d) Find the stationary distribution of the Markov chain, i.e., find a probability distribution π such that $\pi\mathbf{M} = \pi$.

Task 5 – Random Walk Analysis

[12 pts]

The following figure depicts undirected graph with 5 nodes.



A *random walk* on this graph is a Markov chain whose states are the nodes of this graph. If we are at a particular node i in the graph at some time step, then in the following time step, we will transition to each of its neighbors with equal probability. So if it has one neighbor, we transition there with probability 1. If it has 3 neighbors, we transition to each of those with probability $1/3$, and so on.

- a) Consider now a random walk on the above graph. Give the transition probability matrix of the corresponding Markov chain.
- b) Consider a random walk on an undirected graph with n vertices and m edges, where the degree of vertex i is d_i . (The degree of a vertex is the number of neighbors it has in the graph. So in the graph above, for example, $d_1 = 3$ and $d_4 = 2$.) Prove that its stationary distribution $\pi = [\pi_1, \dots, \pi_n]$ satisfies

$$\pi_i = \frac{d_i}{2m}$$

is the stationary distribution of the random walk. That is, show that if $\pi_i = \frac{d_i}{2m}$ for all i , it holds that $\sum_{i=1}^n \pi_i = 1$ and that $\pi = \pi \mathbf{M}$, where \mathbf{M} is the TPM of the Markov chain.

This shows, for example, that for the random walk on the graph shown above

$$\pi = \left[\frac{3}{14}, \frac{4}{14}, \frac{2}{14}, \frac{2}{14}, \frac{3}{14} \right].$$

However, make sure your proof that $\pi_i = \frac{d_i}{2m}$ works for an arbitrary undirected graph, and not just for our specific Markov Chain.

Hint: You can use the fact that $\sum_{i=1}^n d_i = 2m$ since the sum counts each edge twice, once from each end.

Task 6 – Knapsacks (Coding)

[24 pts]

Markov Chain Monte Carlo (MCMC) is a technique that can be used to heuristically and approximately solve otherwise hard optimization problems (among other things). The general strategy is as follows:

1. Define a Markov Chain with states being possible solutions, and (implicitly defined) transition probabilities that result in the stationary distribution π having higher probabilities on “good” solutions to our problem. We don’t actually compute π , but we just want to define the Markov Chain such that the stationary distribution would have higher probabilities on more desirable solutions.
2. Run MCMC, i.e., simulate the Markov Chain for many iterations until we reach a “good” state/solution.

In this question, there is a collection of n items, numbered 0 to $n-1$, available to us, and each has some *value* and some *weight*, both of which are positive real values. We want to find the optimal subset of items that maximizes the total value (the sum of the values of the items we take), subject to the total weight (the sum of the weights of the items we take) being less than some $W > 0$. (This is known as the **knapsack problem**). In `items.txt`,

you'll find a list of potential items with each row containing the name of the item (string), and its value and weight (positive floats).

You will implement an MCMC algorithm which also depends on a parameter T that is not part of the problem definition. Pseudocode is provided below, and a detailed explanation is provided immediately after.

Algorithm 1 MCMC for 0-1 Knapsack Problem

```
1: subset  $\leftarrow$  vector of  $n$  zeros (indexed by 0 to  $n - 1$ ), where subset is always a binary vector in  $\{0, 1\}^n$  that
   represents whether or not we have each item. (This means that we initially start with an empty knapsack).
2: best_subset  $\leftarrow$  subset
3: for  $t = 1, \dots, \text{NUM\_ITER}$  do
4:    $k \leftarrow$  a random uniform integer in  $\{0, 1, \dots, n - 1\}$ .
5:   new_subset  $\leftarrow$  subset but with subset[ $k$ ] flipped ( $0 \rightarrow 1$  or  $1 \rightarrow 0$ ).
6:    $\Delta \leftarrow \text{value}(\text{new\_subset}) - \text{value}(\text{subset})$ 
7:   if new_subset satisfies weight constraint (total weight  $\leq W$ ) then
8:     if  $\Delta > 0$  OR ( $T > 0$  AND  $\text{Unif}(0, 1) < e^{\Delta/T}$ ) then
9:       subset  $\leftarrow$  new_subset
10:  if  $\text{value}(\text{subset}) > \text{value}(\text{best\_subset})$  then
11:    best_subset  $\leftarrow$  subset
```

The extra parameter T in the MCMC algorithm represents a “temperature” that controls the trade-off between exploration and exploitation. The state space \mathcal{S} is the set of all subsets of n items. The algorithm starts with a state (current subset) corresponding to an empty knapsack. At each iteration, the algorithm proposes a new state (proposed subset) as follows: choose a random index k from $\{0, 1, \dots, n - 1\}$. If the item k is not already in the knapsack (current subset), then this proposed subset will just add item k , but if item k is already in the knapsack (current subset), the proposed subset will just remove item k from the knapsack (current subset).

- If the proposed subset is infeasible (doesn't fit in our knapsack because of the weight constraint), we return to the start of the loop and abandon the newly proposed subset.
- Suppose that the proposed subset is feasible. If the proposed subset has higher total value (is better) than the current subset, we will always transition to it (exploitation). Otherwise, if it is worse and $T > 0$, with probability $e^{\Delta/T}$, we update the current subset to the proposed subset, where $\Delta < 0$ is the decrease in total value. This allows us to transition to a “worse” subset occasionally (exploration), and get out of local optima! Repeat this for NUM.ITER transitions from the initial state (subset), and output the highest value subset found during the entire process (which may not be the final subset).

a) What is the size of the Markov Chain's state space \mathcal{S} (the number of possible subsets)?

b) Let's try to figure out what the temperature parameter T does.

(a) Suppose that $T = 0$. Will we ever get to a worse subset than before as we transition?

(b) Suppose $T > 0$. For a fixed T , does the probability of transitioning to a worse subset increase or decrease with larger absolute values of Δ (larger absolute values means “more negative” values, since $\Delta < 0$)? For a fixed Δ , does the probability of transitioning to a worse subset increase or decrease with larger values of T ? Explain briefly how the temperature parameter T controls the degree of exploration we do.

c) Implement the functions `value`, `weight`, and `mcmc` in `mcmc_knapsack.py`. To this end, you will use the [edstem lesson](#). Remember that only code submitted via Gradescope will be graded.

Hints: To get full score, you **must** use `np.random.rand()` to generate an uniform value in $[0, 1]$, and `np.random.randint(low (inclusive), high (exclusive))` to generate your random index(es). Make sure to read the documentation and hints provided!