

Homework 1, CS685 Spring 2025

This is due on **Sunday, Sept 14, 2025** (11:59pm), submitted via Gradescope as a PDF (File>Print>Save as PDF). 100 points total.

IMPORTANT: After copying this notebook to your Google Drive, please paste a link to it below. To get a publicly-accessible link, hit the *Share* button at the top right, then click "Get shareable link" and copy over the result. If you fail to do this, you will receive no credit for this homework!

LINK: <https://colab...> *replace this text with your link!*

How to do this problem set:

- Some questions require writing Python code and computing results, and the rest of them have written answers. For coding problems, you will have to fill out all code blocks that say `YOUR CODE HERE`.
 - For text-based answers, you should replace the text that says "Write your answer here..." with your actual answer.
 - This assignment is designed so that you can run all cells almost instantly. If it is taking longer than that, you have made a mistake in your code.
 - There is no penalty for using AI assistance on this homework as long as you fully disclose it in the final cell of this notebook (this includes storing any prompts that you feed to large language models). That said, anyone caught using AI assistance without proper disclosure will receive a zero on the assignment.
-

How to submit this problem set:

- Write all the answers in this Colab notebook. Once you are finished, generate a PDF via (File -> Print -> Save as PDF) and upload it to Gradescope.
 - **Important:** check your PDF before you submit to Gradescope to make sure it exported correctly. If Colab gets confused about your syntax, it will sometimes terminate the PDF creation routine early.
 - **Important:** on Gradescope, please make sure that you tag each page with the corresponding question(s). This makes it significantly easier for our graders to grade submissions. We may take off points for submissions that are not tagged.
 - When creating your final version of the PDF to hand in, please do a fresh restart and execute every cell in order. Then you'll be sure it's actually right. One handy way to do this is by clicking `Runtime -> Run ALL` in the notebook menu.
-

Academic honesty

- We will audit the Colab notebooks from a set number of students, chosen at random. The audits will check that the code you wrote actually generates the answers in your PDF. If you turn in correct answers on your PDF without code that actually generates those answers, we will consider this a serious case of cheating. See the course page for honesty policies.
 - We will also run automatic checks of Colab notebooks for plagiarism. Copying code from others is also considered a serious case of cheating.
-

✓ Question 1.1 (5 points)

Let's begin with a quick probability review. In the task of language modeling, we're interested in computing the **joint** probability of some text. Say we have a sentence s with n words $(w_1, w_2, w_3, \dots, w_n)$ and we want to compute the joint probability $P(w_1, w_2, w_3, \dots, w_n)$. Assume we are given a model that produces the conditional probability of the next word in a sentence given all preceding words: $P(w_i | w_1, w_2, \dots, w_{i-1})$. How can we use this model to compute the joint probability of sentence s ?

Write your answer here! Please include an equation(s) formatted in LaTeX in your answer. You can add LaTeX to your answer by wrapping it in $...$ signs; see the above cell for examples. If you've never used LaTeX before, please check out [this notebook](#) to get the hang of it!

✓ Question 1.2 (10 points)

Why would we ever want to compute the joint probability of a sentence? Provide **two** different reasons why this probability might be useful to solve an NLP task.

Write your answer here! Please keep it brief (i.e., 2-3 sentences).

✓ Question 1.3 (5 points)

Here is a simple way to build a language model: for any prefix w_1, w_2, \dots, w_{i-1} , retrieve all occurrences of that prefix in some huge text corpus (such as the [Common Crawl](#)) and keep count of the word w_i that follows each occurrence. I can then use this to estimate the conditional probability $P(w_i | w_1, w_2, \dots, w_{i-1})$ for any prefix. Explain why this method is completely impractical.

Write your answer here! Please keep it brief (i.e., 2-3 sentences).

Start coding or [generate](#) with AI.

✓ Question 1.4 (5 points)

In lecture, for a k -gram model, we discussed a theoretical maximum of V^k possible k -grams to keep counts over (where V is the vocabulary size). However, given a corpus of length N , there is another upper bound on the number of possible k -grams that is typically lower in most practical situations. What is it, and why is it often lower? How do the bounds compare for the examples in Figure 2.1 in the JM reading?

Write your answer here!

✓ Question 1.5 (10 points)

Consider a simple language in which each token is drawn from the vocabulary with probability $1/V$, independent of all other tokens. Given a corpus of size M , what is the expectation of the fraction of all possible bigrams that have zero count? Give the mathematical derivation, and implement a simulation to confirm.

Write your answer here!

simulation here

✓ Question 2.1 (5 points)

Let's switch over to coding! The below coding cell contains the opening paragraph of Daphne du Maurier's novel *Rebecca*. Write some code in this cell to compute the number of unique word **types** and total word **tokens** in this paragraph (watch the lecture videos if you're confused about what these terms mean!). To tokenize, please split on whitespace and punctuation (we recommend `re.split()`), and apply lowercase normalization. Be sure that the cell's output is visible in the PDF file you turn in on Gradescope.

```
paragraph = '''Last night I dreamed I went to Manderley again. It seemed to me
that I was passing through the iron gates that led to the driveway.
The drive was just a narrow track now, its stony surface covered
with grass and weeds. Sometimes, when I thought I had lost it, it
would appear again, beneath a fallen tree or beyond a muddy pool
formed by the winter rains. The trees had thrown out new
low branches which stretched across my way. I came to the house
suddenly, and stood there with my heart beating fast and tears
filling my eyes.''.lower() # lowercase normalization is often useful in NLP

types = 0
tokens = 0

# YOUR CODE HERE! POPULATE THE types AND tokens VARIABLES WITH THE CORRECT VALUES!

# DO NOT MODIFY THE BELOW LINE!
print('Number of word types: %d, number of word tokens:%d' % (types, tokens))
```

Question 2.2 (5 points)

Now let's look at the most frequently used word **types** in this paragraph. Write some code in the below cell to print out the 15 most frequently-occurring types. We have initialized a [Counter](#) object that you should use for this purpose. In general, Counters are very useful for text processing in Python.

```
from collections import Counter
c = Counter()

# YOUR CODE HERE!
```

Question 2.3 (5 points)

What do you notice about these words and their linguistic functions (i.e., parts-of-speech)? These words are sometimes called "stopwords" and are often removed from the text before any computational modeling is done. Give one reason why this may be a good idea, and give one reason why this may not be a good idea.

Write your answer here! Please keep it brief (i.e., 2-3 sentences).

Question 3.1 (10 points)

In *neural* language models, we represent words with low-dimensional vectors also called *embeddings*. We use these embeddings to compute a vector representation x of a given prefix, and then predict the probability of the next word conditioned on x . In the below cell, we use [PyTorch](#), a machine learning framework, to explore this setup. We provide embeddings for the prefix "Alice talked to"; your job is to combine them into a single vector representation x using [element-wise vector addition](#).

TIP: if you're finding the PyTorch coding problems difficult, you may want to run through [the 60 minutes blitz tutorial!](#)

```
import torch
torch.set_printoptions(sci_mode=False)
torch.manual_seed(0)

prefix = 'Alice talked to'

# spend some time understanding this code / reading relevant documentation!
# this is a toy problem with a 5 word vocabulary and 10-d embeddings
embeddings = torch.nn.Embedding(num_embeddings=5, embedding_dim=10)
vocab = {'Alice':0, 'talked':1, 'to':2, 'Bob':3, '.':4}

# we need to encode our prefix as integer indices (not words) that index
# into the embeddings matrix. the below line accomplishes this.
# note that PyTorch inputs are always Tensor objects, so we need
# to create a LongTensor out of our list of indices first.
indices = torch.LongTensor([vocab[w] for w in prefix.split()])
prefix_embs = embeddings(indices)
print('prefix embedding tensor size: ', prefix_embs.size())

# okay! we now have three embeddings corresponding to each of the three
# words in the prefix. write some code that adds them element-wise to obtain
# a representation of the prefix! store your answer in a variable named "x".

### YOUR CODE HERE!
x = torch.zeros(10)

### DO NOT MODIFY THE BELOW LINE
print('embedding sum: ', x)
```

Question 3.2 (5 points)

Most useful language models do not use element-wise addition to combine the different word embeddings in the prefix into a single representation (a process called *composition*). What is a major issue with element-wise functions that makes them unsuitable for use as

Write your answer here! Please keep it brief (i.e., 2-3 sentences).

Question 3.3 (5 points)

One very important function in neural language models (and for basically every task we'll look at this semester) is the [softmax](#), which is defined over an n -dimensional vector $\langle x_1, x_2, \dots, x_n \rangle$ as $\text{softmax}(x_i) = \frac{e^{x_i}}{\sum_{1 \leq j \leq n} e^{x_j}}$. Let's say we have our prefix representation x from before. We can use the softmax function, along with a linear projection using a matrix W , to go from x to a probability distribution p over the next word: $p = \text{softmax}(Wx)$. Let's explore this in the code cell below:

```
# remember, our goal is to produce a probability distribution over the
# next word, conditioned on the prefix representation x. This distribution
# is thus over the entire vocabulary (i.e., it is a 5-dimensional vector).
# take a look at the dimensionality of x, and you'll notice that it is a
# 10-dimensional vector. first, we need to project this representation
# down to 5-d. We'll do this using the below matrix:

W = torch.rand(10, 5)

# use this matrix to project x to a 5-d space, and then
# use the softmax function to convert it to a probability distribution.
# this will involve using PyTorch to compute a matrix/vector product.
# look through the documentation if you're confused (torch.nn.functional.softmax)
# please store your final probability distribution in the "probs" variable.

### YOUR CODE HERE
probs = torch.rand(5)

### DO NOT MODIFY THE BELOW LINE!
print('probability distribution', probs)
```

Question 3.4 (15 points)

So far, we have looked at just a single prefix ("Alice talked to"). In practice, it is common for us to compute many prefixes in one computation, as this enables us to take advantage of GPU parallelism and also obtain better gradient approximations (we'll talk more about the latter point later). This is called *batching*, where each prefix is an example in a larger batch. Here, you'll redo the computations from the previous cells, but instead of having one prefix, you'll have a batch of two prefixes. The final output of this cell should be a 2x5 matrix that contains two probability distributions, one for each prefix. **NOTE: YOU WILL LOSE POINTS IF YOU USE ANY LOOPS IN YOUR ANSWER!** Your code should be completely vectorized (a few large computations is faster than many smaller ones).

```
# for this problem, we'll just copy our old prefix over three times
# to form a batch. in practice, each example in the batch would be different.
batch_indices = torch.cat(2 * [indices]).reshape((2, 3))
batch_embs = embeddings(batch_indices)
print('batch embedding tensor size: ', batch_embs.size())

# now, follow the same procedure as before:
# step 1: compose each example's embeddings into a single representation
# using element-wise addition. HINT: check out the "dim" argument of the torch.sum function!

# step 2: project each composed representation into a 5-d space using matrix W
# step 3: use the softmax function to obtain a 2x5 matrix with the probability distributions

# please store this probability matrix in the "batch_probs" variable, which is
# currently initialized with random numbers.

batch_probs = torch.rand(2,5)

### DO NOT MODIFY THE BELOW LINE
print("batch probability distributions:", batch_probs)
```

Question 4 (15 points)

Choose one paper from ACL 2025 (for example, from [all papers](#), or perhaps [paper awards](#), or [ACL Anthology version with hyperlinks](#)) that is (1) related to language modeling and (2) of interest to you. A good way to do this is by searching for some keywords and then scanning the resulting titles and abstracts; there are many papers so take your time before selecting one! Then, write a summary in your own words of the paper you chose.

Your summary should answer the following questions: what is its motivation? Why should anyone care about it? Were there things in the paper that you didn't understand at all? What were they? Fill out the below cell, and make sure to write 2-4 paragraphs for the summary to receive full credit.

Title of paper:

Authors:

Conference name:

URL:

Your summary:

AI Disclosure

- Did you use any AI assistance to complete this homework? If so, please also specify what AI you used.
 - *your response here*
-

(only complete the below questions if you answered yes above)

- If you used a large language model to assist you, please paste *all* of the prompts that you used below. Add a separate bullet for each prompt, and specify which problem is associated with which prompt.
 - *your response here*
- **Free response:** For each problem for which you used assistance, describe your overall experience with the AI. How helpful was it? Did it just directly give you a good answer, or did you have to edit it? Was its output ever obviously wrong or irrelevant? Did you use it to get the answer or check your own answer?
 - *your response here*