
Midterm for 6.864
Name:

30	30	30	30

Good luck!

Question 1 (10 points)

We define a PCFG where non-terminal symbols are $\{S, A, B\}$, the terminal symbols are $\{a, b\}$, and the start non-terminal (the non-terminal always at the root of the tree) is S . The PCFG has the following rules:

Rule	Probability
$S \rightarrow S A$	0.1
$S \rightarrow A B$	0.4
$S \rightarrow B A$	0.5
$A \rightarrow a$	0.2
$A \rightarrow b$	0.8
$B \rightarrow a$	0.3
$B \rightarrow b$	0.7

For the input string $abab$, show two possible parse trees under this PCFG, and show how to calculate their probability.

Question 2 (10 points)

Consider an application of global linear models to parsing. In this scenario each input x is a sentence. We have a fixed context-free grammar; $\text{GEN}(x)$ returns the set of all parses allowed for x under the context-free grammar. The feature vector $\Phi(x, y)$ for any sentence x paired with a parse tree y is defined as

$$\Phi(x, y) = \sum_{\alpha \rightarrow \beta \in (x, y)} \bar{\phi}(\alpha \rightarrow \beta)$$

where $\bar{\phi}$ is a function that maps a context-free rule $\alpha \rightarrow \beta$ to a feature vector, and the notation $\sum_{\alpha \rightarrow \beta \in (x, y)}$ refers to a sum over all context-free rules in the parse tree defined by (x, y) .

We'd like $\Phi(x, y)$ to be a 3-dimensional feature vector, with the following values for its three components:

$$\begin{aligned}\Phi_1(x, y) &= \text{number of times } \mathbf{S} \rightarrow \mathbf{NP VP} \text{ is seen in } (x, y) \\ \Phi_2(x, y) &= \text{number of times } \mathbf{N} \rightarrow \mathbf{dog} \text{ is seen in } (x, y) \\ \Phi_3(x, y) &= \text{number of times } \mathbf{NP} \rightarrow \mathbf{NP NP} \text{ is seen in } (x, y)\end{aligned}$$

Give a definition for the function $\bar{\phi}$ which leads to this definition of $\Phi(x, y)$.

Question 3 (10 points)

In the lecture on text segmentation, we considered unsupervised methods for this task. An alternative approach is to model text segmentation as a supervised learning task. During training, we are provided with texts where sentence breaks that correspond to segment boundaries are annotated as positive examples. The rest of the sentence breaks are annotated as negative examples. Assume that a boundary location depends on the word on the left, the word on the right, and on the tag of the previous sentence break (i.e., whether it is a segment boundary or not).

We would like to apply a log-linear tagger to this problem. Define the features of a log-linear tagger which would implement the specification above.

Consider the following HMM model for tagging. We will assume that the vocabulary consists of three words, *the*, *dog*, *sleeps*. There are two possible part-of-speech tags, *D*, *N*.

One set of parameters in the HMM are probabilities of the form $P(\text{word}|\text{tag})$, for example $P(\text{the}|D)$, $P(\text{the}|N)$, etc. In our HMM these probabilities take the following values:

Parameter	Value
$P(\text{the} D)$	0.7
$P(\text{dog} D)$	0.2
$P(\text{sleeps} D)$	0.1
$P(\text{the} N)$	0.2
$P(\text{dog} N)$	0.5
$P(\text{sleeps} N)$	0.3

Another set of parameters in the HMM are of the form $P(\text{tag}|\text{previous-tag})$, for example $P(N|D)$, $P(D|N)$, etc. Note that we have a *bigram* HMM tagger, in that each tag depends only on the previous tag (in lecture we saw a *trigram* tagger, where each tag depends on the previous two tags). We take **START** to be a special tag which always appears at the start of a sentence, and **STOP** to be a tag that appears at the end of the sentence. In our model these parameters take the following values:

Parameter	Value
$P(D \text{START})$	0.1
$P(N \text{START})$	0.2
$P(\text{STOP} \text{START})$	0.7
$P(D D)$	0.6
$P(N D)$	0.2
$P(\text{STOP} D)$	0.2
$P(D N)$	0.4
$P(N N)$	0.5
$P(\text{STOP} N)$	0.1

Question 4 (30 points) We would like you to define a PCFG that is “equivalent” to the above HMM. By “equivalent” we mean the following:

- For any symbol sequence x and state sequence y which has probability $P_{HMM}(x, y) > 0$ under the HMM: 1) there should be a parse tree y' such

that the pair x, y' gets probability $P_{PCFG}(x, y')$ under the PCFG; 2) this probability should satisfy $P_{HMM}(x, y) = P_{PCFG}(x, y')$

- There should be a one-to-one function $f(y) = y'$ that maps a state sequence in the HMM to a parse tree generated by the PCFG, and that satisfies $P_{HMM}(x, y) = P_{PCFG}(x, f(y))$

In your solution you should: (a) write down your PCFG which is equivalent to the above HMM; (b) define the function $f(y)$ between state sequences and parse trees.

Note: you may find it useful to make use of productions in your PCFG where ϵ (the empty string) is on the right-hand-side of your rule. For example,

$$X \rightarrow \epsilon \quad \text{with } P(X \rightarrow \epsilon | X) = 0.1$$

states that the non-terminal X can rewrite to the empty string with probability 0.1. To elaborate further, the context-free grammar

$$\begin{aligned} S &\rightarrow a X \\ X &\rightarrow \epsilon \end{aligned}$$

Generates one string, i.e., the string a .

This question considers log-linear models. We'd like to build a model that estimates a distribution $P(\text{tag}|\text{word})$ using a log-linear model. The variable tag can take any one of three values, D , N , V . The variable word could potentially be any member of a set \mathcal{V} of possible words. The set \mathcal{V} contains the words *the*, *dog*, *sleeps*, as well as additional words (i.e., $|\mathcal{V}| > 3$). The distribution should give the following probabilities:

$$\begin{aligned}P(D|\textit{the}) &= 0.9 \\P(N|\textit{dog}) &= 0.9 \\P(V|\textit{sleeps}) &= 0.9 \\P(D|\textit{word}) &= 0.6 \text{ for any word other than } \textit{the}, \textit{dog} \text{ or } \textit{sleeps} \\P(N|\textit{word}) &= 0.3 \text{ for any word other than } \textit{the}, \textit{dog} \text{ or } \textit{sleeps} \\P(V|\textit{word}) &= 0.1 \text{ for any word other than } \textit{the}, \textit{dog} \text{ or } \textit{sleeps}\end{aligned}$$

Note that we have intentionally left the values for the following probabilities undefined: $P(N|\textit{the})$, $P(V|\textit{the})$, $P(D|\textit{dog})$, $P(V|\textit{dog})$, $P(D|\textit{sleeps})$, $P(N|\textit{sleeps})$. It is assumed that these probabilities could take any values such that $\sum_{\text{tag}} P(\text{tag}|\text{word}) = 1$ is satisfied for $\text{word} = \textit{the}, \textit{dog}, \textit{sleeps}$, or indeed any other word in \mathcal{V} .

Question 5 (10 points)

Define the features for a log-linear model that can model this distribution $P(\text{tag}|\text{word})$ perfectly. Each feature should be an indicator function: i.e., each feature $\phi(x, y)$ can take only the values 0 or 1 depending on the values of x and y . Your model should make use of as few features as possible. We will give you 10 points for using 6 features, and will penalise you for using more than 6 features. (It may be possible to use only 5 features, but a model with 6 features may be more straightforward to analyse.)

Question 6 (10 points)

Write an expression for each of the probabilities

$$P(D|\text{cat})$$

$$P(N|\text{laughs})$$

$$P(D|\text{dog})$$

$$P(V|\text{sleeps})$$

as a function of the parameters in your model. (Assume that the words *laughs* and *cat* are both members of the set \mathcal{V} .)

Question 7 (10 points)

What value do the parameters in your model take to give the distribution described above?

Consider the task of aligning two English translations of the same story (extracts from two translations of Andersen's tale are shown below.) In the sentence alignment task, the goal is to identify groups of sentences in one translation that correspond to groups of sentences in the other translation.

- | |
|---|
| 1) It was a Princess who was standing outside the door.
2) She was in a sad condition.
3) The water trickled down from her hair.
4) Her clothes clung to her body.
5) She said she was a real Princess. |
| A) It was a princess standing out there in front of the gate.
B) The water ran down from her hair and clothes, and her dress clung to her body.
C) And yet she said that she was a real princess. |

Question 8 (10 points)

We first need to introduce a function that compares similarity between two sentences. Recall the definition for the cosine distance between two vectors x and y , as seen in lecture:

$$\text{cosine}(x, y) = \frac{\sum_{i=1}^n x_i y_i}{\sqrt{\sum_{i=1}^n x_i^2} \sqrt{\sum_{i=1}^n y_i^2}}$$

How would you define a mapping from sentences s to vectors $f(s)$ such that the cosine distance $\text{cosine}(f(s_1), f(s_2))$ is a measure of the similarity between two sentences s_1 and s_2 ? We'll give full marks for any reasonable mapping.

Question 10 (10 points)

Next, we'll modify the alignment score to be the same as before, but to include a fixed penalty p each time a sentence is aligned to an empty sentence. p is a parameter, which is ≥ 0 , chosen by the user of the algorithm. Describe a modified dynamic programming method which takes this new penalty into account.