

Chapter 2

Language Models

2.1 Motivation: Machine Translation

To start the course, we're going to look at the problem of *machine translation*, that is, translating sentences from one language to another (traditionally, a French sentence f to an English sentence e). A translation system has to try to do two things at once: first, it has to generate an e that means the same thing as f (*adequacy*), and second, it has to generate an e that is good English (*fluency*). We could produce a perfectly adequate but not fluent translation by outputting f itself; we could produce a perfectly fluent but not adequate translation by always outputting "My hovercraft is full of eels." Doing both at once is what makes the problem nontrivial.

Neural networks are rather good at combining two kinds of information like this. But older machine learning methods were not as good at doing this, so they divided the work between two submodels, in the following way. Warren Weaver first proposed, in 1947, to treat translation as a decoding problem:

One naturally wonders if the problem of translation could conceivably be treated as a problem in cryptography. When I look at an article in Russian, I say: 'This is really written in English, but it has been coded in some strange symbols. I will now proceed to decode.'

In other words, when a Russian speaker speaks Russian, he first thinks in English, but then, as he expresses his thoughts, they somehow become encoded in Russian. Then (leaving aside how absurd this is), the task of translating a Russian sentence f is to recover the original English sentence e that the speaker was thinking before he said it. Mathematically,

$$P(f, e) = P(e)P(f | e).$$

Then, if we are given a sentence \mathbf{f} , we can reconstruct \mathbf{e} by finding:

$$\begin{aligned}\hat{\mathbf{e}} &= \arg \max_{\mathbf{e}} P(\mathbf{e} | \mathbf{f}) \\ &= \arg \max_{\mathbf{e}} \frac{P(\mathbf{e}, \mathbf{f})}{P(\mathbf{f})} \\ &= \arg \max_{\mathbf{e}} P(\mathbf{e}, \mathbf{f}) \\ &= \arg \max_{\mathbf{e}} P(\mathbf{e}) P(\mathbf{f} | \mathbf{e}).\end{aligned}$$

So the model is divided into two parts. The term $P(\mathbf{f} | \mathbf{e})$ is called the *translation model*, which says how similar the two sentences are in meaning (adequacy). The term $P(\mathbf{e})$ is the *language model*. It says what kinds of sentences are more likely than others (fluency).

In this chapter, we will focus on the language model. With apologies, we're going to switch notation; since we're leaving translation aside for the moment, we don't need to distinguish between languages, and we simply call the sentence $\mathbf{w} = w_1 \cdots w_N$, where each w_i is a character, or a word, or something in between. (How we cut up a sentence into segments depends on the application, but the techniques are the same in any case.)

2.2 n-gram Models

2.2.1 Definition

The simplest kind of language model is the n -gram language model, in which each word depends on the $(n - 1)$ previous words. In a 1-gram or *unigram* language model, each word is generated independently:

$$P(w_1 \cdots w_N) = p(w_1) \cdots p(w_N) p(\text{EOS}). \quad (2.1)$$

The symbol EOS stands for “end of sentence,” and is needed in order to make the probabilities of all sentences of all lengths sum to one. Imagine rolling a die with one word written on each face to generate random sentences; in order to know when to stop rolling, you need (at least) one face of the die to say EOS which means “stop rolling.”

In a *bigram* (2-gram) language model, each word's probability depends on the previous word:

$$P(w_1 \cdots w_N) = p(w_1 | \text{BOS}) \left(\prod_{i=2}^N p(w_i | w_{i-1}) \right) p(\text{EOS} | w_N). \quad (2.2)$$

where we now also have a special symbol BOS for the beginning of the sentence, because the first word doesn't have a real previous word to condition on.

A general n -gram language model is:

$$P(w_1 \cdots w_N) = \prod_{i=1}^{N+1} p(w_i | w_{i-n+1} \cdots w_{i-1}), \quad (2.3)$$

where we pretend that $w_i = \text{BOS}$ for $i \leq 0$, and $w_{N+1} = \text{EOS}$.

2.2.2 Training

Training an n -gram model is easy. To estimate the probabilities of a unigram language model, just count the number of times each word occurs and divide it by the total number of words:

$$p(w) = \frac{c(w)}{\sum_{w'} c(w')}.$$

And for general n -grams,

$$p(w | \mathbf{u}) = \frac{c(\mathbf{u}w)}{\sum_{w'} c(\mathbf{u}w')}$$

where \mathbf{u} ranges over $(n - 1)$ -grams.

2.2.3 Smoothing

A never-ending challenge in all machine learning settings is the *bias-variance* tradeoff, or the tradeoff between *underfitting* and *overfitting*. In language modeling, underfitting usually means that the model probability of a word doesn't sufficiently take into account the context of the word. For example, a unigram language model would think that "the the the" is a very good sentence. In the world of n -gram language models, the antidote to underfitting is to increase n .

Overfitting usually means that the model overestimates the probability of words or word sequences seen in data and underestimates the probability of words or word sequences not seen in data. With n -gram language models, the classic solution is *smoothing*, which tries to take some probability mass away from seen n -grams and give it to unseen n -grams.

Recall that the unsmoothed probability estimate of an n -gram is

$$p(w | \mathbf{u}) = \frac{c(\mathbf{u}w)}{\sum_{w'} c(\mathbf{u}w')}. \quad (2.4)$$

There are basically two ways to take probability mass away: multiply the probability by $\lambda < 1$, or subtract $d > 0$ from the numerator. Then the probability distribution doesn't sum to one, so we add probability mass back proportional to some other distribution \bar{p} .

$$\text{Multiplying:} \quad p(w | \mathbf{u}) = \lambda \frac{c(\mathbf{u}w)}{\sum_{w'} c(\mathbf{u}w')} + (1 - \lambda)\bar{p}(w | \mathbf{u}) \quad (2.5)$$

$$\text{Subtracting:} \quad p(w | \mathbf{u}) = \frac{\max(0, c(\mathbf{u}w) - d)}{\sum_{w'} c(\mathbf{u}w')} + \alpha\bar{p}(w | \mathbf{u}). \quad (2.6)$$

where:

- λ and d must be chosen carefully; for much more information on how to do that, see the technical report by Chen and Goodman (1998).

- In (2.6), α is chosen to make the distribution sum to one.
- The distribution \bar{p} is typically simpler than p . If p is an n -gram model where $n > 1$, then \bar{p} is typically an $(n - 1)$ -gram model. If p is a unigram model, then \bar{p} is typically the uniform distribution over words.

A special case of (2.5) is simply to add 1 to the count of every n -gram, including unseen n -grams.

$$p(w | \mathbf{u}) = \frac{c(\mathbf{u}w) + 1}{\sum_{w'} (c(\mathbf{u}w') + 1)}, \quad (2.7)$$

which is called *add-one smoothing*. (How is this a special case of (2.5)?) In many situations (and particularly for language modeling), this is a terrible smoothing method. But it's good to know because it's so easy.

2.3 Unknown Words

Natural languages probably don't have a finite vocabulary, and even if they do, the distribution of word frequencies has such a long tail that, in any data outside the training data, *unknown* or *out-of-vocabulary* (OOV) words are rather common. Unknown words are problematic for all language models, and we have a few techniques for handling them.

2.3.1 Smoothing

Above, we saw that smoothing decreases probability estimates of seen n -grams and increases probability estimates of unseen n -grams, including zero-probability n -grams. If we add a pseudo-word UNK to our model's vocabulary, it'll have a count of zero, but smoothing will give it a nonzero probability.

2.3.2 Limiting the vocabulary

A simpler method is to pretend that some word types seen in the training data are unknown. For example, we might limit the vocabulary to 10,000 word types, and all other word types are changed to UNK. Or, we might limit the vocabulary just to those seen two or more times, and all other word types are changed to UNK.

2.3.3 Subword segmentation

Another idea is to break words into smaller pieces, usually using a method like *byte pair encoding* (Sennrich, Haddow, and Birch, 2016). This breaks unknown words down into smaller, known, pieces. However, this only alleviates the unknown word problem; it does not solve it. So one of the above techniques is still needed.

2.4 Evaluation

Whenever we build any kind of model, we always have to think about how to evaluate it.

2.4.1 Generating random sentences

One popular way of demonstrating a language model is using it to generate random sentences. While this is entertaining and can give a qualitative sense of what kinds of information a language model does and doesn't capture, but it is *not* a rigorous way to evaluate language models. Why? Imagine a language model that just memorizes the sentences in the training data. This model would randomly generate perfectly-formed sentences. But if you gave it a sentence w not seen in the training data, it would give w a probability of zero.

2.4.2 Extrinsic evaluation

The best way to evaluate language models is extrinsically. Language models are usually used as part of some larger system, so to compare two language models, compare how much they help the larger system.

2.4.3 Perplexity

For intrinsic evaluation, the standard way to evaluate a language model is how well it fits some *held-out* data (that is, data that is different from the training data). There are various ways to measure this:

$$\text{likelihood} = P(w_1 \cdots w_N \text{ EOS}) \quad (2.8)$$

$$\text{cross-entropy} = -\frac{1}{N} \log_2 \text{likelihood} \quad (2.9)$$

$$\text{perplexity} = 2^{\text{cross-entropy}} \quad (2.10)$$

Perplexity is the standard metric. The best (lowest) possible perplexity is 1, meaning that the model always knows what the next word is. The worst (highest) possible perplexity is the vocabulary size, meaning that if the model had to guess the next word, it would be choosing randomly and uniformly from the vocabulary.

Held-out data is always going to have unknown words, which require some special care. Above, we handled unknown words by mapping them to a special token UNK, but if we compare two language models, they must map exactly the same subset of word types to UNK. (If not, can you think of a way to cheat and get a perplexity of 1?)

2.5 Weighted Automata

If you've taken *Theory of Computing*, you should be quite familiar with finite automata; if not, you may be familiar with regular expressions, which are equivalent to finite automata. Many models in NLP can be thought of as finite automata, or variants of finite automata, including n -gram language models. Although this

may feel like overkill at first, we'll soon see that formalizing models as finite automata makes it much easier to combine models in various ways.


2.5.1 Definitions

A *finite automaton (FA)* is an imaginary machine that reads in a string and outputs an answer, either “accept” or “reject.” (For example, a FA could accept only words in an English dictionary and reject all other strings.) At any given time, the machine is in one *state* out of a finite set of possible states. It has rules, called *transitions*, that tell it how to move from one state to another.

A FA is typically represented by a directed graph. We draw nodes to represent the various states that the machine can be in. The node can be drawn with or without the state's name inside. The machine starts in the *initial state* (or *start state*), which we draw as:



The edges of the graph represent transitions, for example:




which means that if the machine is in state q and the next input symbol is a , then it can read in a and move to state r . The machine also has zero or more *final states* (or *accept states*), which we draw as:



If the machine can reach the end of the string while in a final state, then it accepts the string. Otherwise, it rejects.

We say that a FA is *deterministic* (or a DFA) if every state has the property that, for each label, there is exactly one outgoing transition with that label. Otherwise, it is *nondeterministic* (or an NFA).

A *weighted finite automaton* adds a *weight* to each transition. A transition on symbol a with weight p is written




The weight of a path through a weighted FA is the product of the weights of the transitions along the path. A weighted FA defines a weighted language, or a distribution over strings, in which the weight of a string is the sum of the weights of all accepting paths of the string.

In a *probabilistic FA*, the weight of a string is a probability (that is, the weights of all strings sum to one). To make this happen, we impose the following conditions:


- Each state has the property that the weights of all of the outgoing transitions sum to one.
- There is a special *stop symbol*, which we write as EOS, that we assume appears at the end of every string, and only at the end of every string.

2.5.2 Language models as automata


An n -gram language model is a probabilistic DFA with a very simple structure. A bigram model with an alphabet $\Sigma = \{a, b\}$ looks like this:



In general, we need a state for every observed context, that is, one for BOS, which we call q_{BOS} , and one for each word type a , which we call q_a . And we need a final state q_{EOS} . For all a, b , there is a transition




and for every state q_a , there is a transition




Generalizing to n -grams, we need a state for every $(n - 1)$ -gram. It would be messy to actually draw the diagram, but we can describe how to construct it:


- For all $\mathbf{u} \in \Sigma^{n-1}$, there is a state $q_{\mathbf{u}}$.
- The start state is $q_{\text{BOS}^{n-1}}$.
- The accept state is q_{EOS} .
- For all $a \in \Sigma, \mathbf{u} \in \Sigma^{n-2}, b \in \Sigma$, there's a transition



- For all $\mathbf{u} \in \Sigma^{n-1}$, there's a transition




One can imagine designing other kinds of language models as well. For example, a *trie* is often used for storing lists like dictionaries:



The portion shown accepts the strings {a, aa, ab, ba}.

Here’s an example of a probabilistic NFA, known as a *hidden Markov model (HMM)*.



where each transition probability is defined in terms of two smaller steps:

$$p(r, a | q) = t(r | q) o(a | r). \tag{2.12}$$

Notice how a single string can have multiple accepting paths. For example, if the input symbols are English, then we could set the transition probabilities so that the NFA goes to q_1 when reading a noun and q_2 when reading a verb (and we could create more states for more parts of speech). In the sentence “I saw her duck,” the word “duck” could be either a noun or a verb, so it would be appropriate for the NFA to have more than one path that accepts this sentence.

2.5.3 Training

If we are given a collection of strings \mathcal{D} and a DFA M , we can learn weights very easily. For each string $w \in \mathcal{D}$ (which we assume to end in EOS), run M on w and

count, for each state q and each word $a \in \Sigma \cup \{\text{EOS}\}$, the number of times $c(q, a)$ that M is in state q and reads an a . Then the weight of transition $\textcircled{q} \xrightarrow{a} \textcircled{r}$ is $\frac{c(q,a)}{\sum_{a'} c(q,a')}$. This is the weighting of M that maximizes the likelihood of \mathcal{D} .

If the automaton is not deterministic, the above won't work. This is because, for a given string, there might be more than one path that accepts it, and we don't know which path's transitions to count.

In other words, we want to maximize the log-likelihood,

$$\begin{aligned} L &= \log \prod_{w \in \mathcal{D}} P(w) \\ &= \log \prod_{w \in \mathcal{D}} \sum_{\text{paths } \pi \text{ for } w} P(\pi) \\ &= \log \prod_{w \in \mathcal{D}} \sum_{\text{paths for } w} \prod_{\text{transitions in path}} (\text{weight of path}). \end{aligned}$$

That summation looks intractable, but in fact we can compute it efficiently using dynamic programming. Let $\alpha[u, q]$ be the total weight of all paths from the start state q_0 to state q that the NFA can take while reading u . We can compute this recursively like so:

$$\alpha[\epsilon, q] = \begin{cases} 1 & q = q_0 \\ 0 & \text{otherwise} \end{cases} \quad (2.13)$$

$$\alpha[ua, r] = \sum_{q \in Q} \alpha[u, q] \cdot \text{weight} \left(\textcircled{q} \xrightarrow{a} \textcircled{r} \right). \quad (2.14)$$

Then the total weight of all accepting paths for w is $\sum_{q \in F} \alpha[w, q]$. So we now have

$$L = \log \prod_{w \in \mathcal{D}} \sum_{q \in F} \alpha[w, q].$$

Unfortunately, we can't maximize this by (say) setting its derivative to zero and solving for the transition weights. Instead, we must use some kind of iterative approximation.


One way to do this is called *expectation-maximization*, the traditional way to train an HMM. The other way is to use gradient-based optimization, which we'll cover later when we talk about neural networks (Section 2.6.4).

2.6 Recurrent Neural Networks


For a long time, researchers tried to find language models that were better than n -gram models and failed, but in recent years, neural networks have become powerful enough to retire n -grams at last. One way of defining a language model as a neural network is as a *recurrent neural network* (RNN).

2.6.1 From finite automata . . .

Let's start with a bigram language model, M_{bigram} . Recall that there are states q_{BOS} and q_{EOS} , and a state q_a for every $a \in \Sigma$. To keep the diagram simple, we show what the transitions look like for two symbols $a, b \in \Sigma$, and we disallow the empty string.




In contrast to the bigram model, if you, a human, had to predict the next word in a sentence, you might use information from further back in the string. For example, if a sentence starts with a lowercase letter, then it's probably all lowercase (even the proper names). To model this, we want every symbol to depend on the first symbol:



How would we combine the bigram and look-at-first-symbol model? The easy

but wrong answer is to take their union:



The trouble with this is that the probability of a string under the combined model is a weighted sum of the probability under the two original models. If one model loves a string and the other model hates it, then the combined model loves it even more – there’s no way for one to veto the other.

Another idea would be to intersect them. The resulting model would simulate both the original models at the same time, which would be very powerful, but also very large: the number of states in the intersection is the *product* of the number of states in the two original automata.

To get around this explosion in the number of states, we look back to the original definition of finite automata given by Kleene (1951). Under the definition you are familiar with (Rabin and Scott, 1959), an NFA can enter state r if there’s *at least one* incoming transition to r that the NFA can take. More formally, let $\text{exists} \left(\begin{array}{c} \textcircled{q} \xrightarrow{a} \textcircled{r} \end{array} \right) = 1$ if there is a transition $\textcircled{q} \xrightarrow{a} \textcircled{r}$, and 0 otherwise. Then the *activation* of state r after reading string u is

$$\alpha[\epsilon, q] = \mathbb{I}[q = q_0] \quad (2.18)$$

$$\alpha[ua, r] = \mathbb{I} \left[\sum_{q \in Q} \alpha[u, q] \cdot \text{exists} \left(\begin{array}{c} \textcircled{q} \xrightarrow{a} \textcircled{r} \end{array} \right) \geq 1 \right] \quad (2.19)$$

where $\mathbb{I}[\cdot]$ is 1 if the thing inside the square brackets is true; 0 otherwise. Compare with equations (2.13–2.14).

Let’s generalize this rule by letting $\alpha[ua, r]$ be an arbitrary Boolean-valued function of the $\alpha[u, q]$ and a . For example, in the NFA above (2.17), we can change




Figure 2.1: The sigmoid function is 0 for very negative values, 1 for very positive values, and smooth in between.

the behavior of q_{EOS} to

$$\alpha[ua, q_{\text{EOS}}] = \mathbb{I} \left[\sum_{q \in \{q_a, q_b, r_a, r_b\}} \alpha[u, q] \geq 2 \right]. \quad (2.20)$$

Note that the threshold is 2 instead of 1. This means that a string is now accepted by the automaton iff it is accepted by *both* of the original automata. (The automaton can never be in q_a and q_b at once, nor r_a and r_b .) So we get the power of intersection without the huge number of states that intersection usually creates.

(Kleene's more general definition does not increase the power of finite automata at all in terms of the languages recognized. Any automaton under this definition can be converted into an equivalent DFA, by a procedure very similar to converting an NFA to a DFA.)

2.6.2 ...to recurrent neural networks

Kleene's definition was a generalization of the original neural networks, defined by McCulloch and Pitts (1943). In a McCulloch-Pitts neural network, the behavior of every state has a definition of the form

$$\alpha[ua, r] = \mathbb{I} \left[\sum_q A(q, r) \alpha[u, q] + B(a, r) + c \geq 0 \right] \quad (2.21)$$

where A and B can be any functions returning natural numbers, and c is any natural number or $-\infty$. In some ways this is not as powerful as NFAs (since there is no direct dependence between q and a) and in some ways this is more powerful (because we can do things similar to intersection).




Figure 2.2: A simple RNN, shown for a string of length 3. Each rectangle is a vector that is either input to or computed by the network.

A so-called *simple* (or *Elman*) RNN makes a McCulloch-Pitts neural network continuous-valued, which makes them much easier to learn. The α 's can be not only 0 or 1, but any real number in between. And the step function, $z \mapsto \mathbb{I}[z \geq 0]$, is replaced with the *sigmoid* function,

$$\text{sigmoid}(z) = \frac{1}{1 + \exp(-z)}$$

which is a smooth version of the step function (see Figure 2.1).

RNNs nowadays are defined as a system of matrix equations. So let's now switch to this more standard notation. Number the symbols of the alphabet, starting from 1. The ordering is completely arbitrary. For example, if the alphabet is $\{a, b, \text{EOS}\}$, we could number them: $a = 1, b = 2, \text{EOS} = 3$. From now on, we will use a symbol and its number interchangeably.

See Figure 2.3 for a picture of an RNN. If the input string is $w = w_1 \cdots w_n$ and $w_{n+1} = \text{EOS}$, define a sequence of vectors $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$. (The superscripts are written with parentheses to make it clear that this isn't exponentiation.) Each vector $\mathbf{x}^{(i)}$ encodes w_i as a *one-hot* vector, which means that $\mathbf{x}^{(i)}$ is a vector with all 0's except for a 1 at position w_i . For example, if $w = \text{aba}$, then the input vectors would be

$$\mathbf{x}^{(1)} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad \mathbf{x}^{(2)} = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \quad \mathbf{x}^{(3)} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} \quad \mathbf{x}^{(4)} = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

From these vectors, the RNN computes a sequence of vectors, which we previously called α :

$$\mathbf{h}^{(i)} = \text{sigmoid}(\mathbf{A}\mathbf{h}^{(i-1)} + \mathbf{B}\mathbf{x}^{(i)} + \mathbf{c}) \quad i = 1, \dots, n \quad (2.22)$$

where

$$\mathbf{h}^{(0)} \in \mathbb{R}^d \quad (2.23)$$

$$\mathbf{A} \in \mathbb{R}^{d \times d} \quad (2.24)$$

$$\mathbf{B} \in \mathbb{R}^{d \times |\Sigma|} \quad (2.25)$$

$$\mathbf{c} \in \mathbb{R}^d \quad (2.26)$$

are parameters of the model. They don't have a definition or an intuitive interpretation (like probabilities do); they will be learned during the training process, as described in Section 2.6.4.

At any point in time $i = 1, \dots, n$, the RNN can make a prediction about the next symbol:

$$\mathbf{y}^{(i)} = \text{softmax}(\mathbf{D}\mathbf{h}^{(i)} + \mathbf{e}) \quad (2.27)$$

$$P(w_{i+1} \mid w_1 \cdots w_i) \approx y_{w_{i+1}}^{(i)} \quad (2.28)$$

where

$$\mathbf{D} \in \mathbb{R}^{|\Sigma| \times d} \quad (2.29)$$

$$\mathbf{e} \in \mathbb{R}^{|\Sigma|} \quad (2.30)$$

are more parameters of the model. See Section 1.4.2 for a definition of the softmax function. For example, if

$$\mathbf{y}^{(1)} = \begin{bmatrix} 0.6 \\ 0.2 \\ 0.4 \end{bmatrix},$$

that means

$$P(w_2 = a \mid \text{BOS}) = 0.6$$

$$P(w_2 = b \mid \text{BOS}) = 0.2$$

$$P(w_2 = \text{EOS} \mid \text{BOS}) = 0.4.$$

2.6.3 Example

Figure 2.3 shows a run of a simple RNN with 30 hidden units trained on the Wall Street Journal portion of the Penn Treebank, a common toy dataset for neural language models. When we run this model on a new sentence, we can visualize what each of its hidden units is doing at each time step. The units have been sorted by how rapidly they change.

The first unit seems to be unchanging; maybe it's useful for other units to compute their values. The second unit is blue on the start symbol, then becomes deeper and deeper red as the end of the sentence approaches. This unit seems to be measuring the position in the sentence and/or trying to predict the end of the sentence. The third unit is red for the first part of the sentence, usually the subject, and turns blue for the second part, usually the predicate. The rest of the units are unfortunately difficult to interpret. But we can see that the model is learning something about the large-scale structure of a sentence, without being explicitly told anything about sentence structure.

Other kinds of RNNs that perform better than this simple RNN have been shown to have units that perform various functions (Karpathy, Johnson, and Fei-Fei, 2016).




Figure 2.3: Visualization of a simple RNN language model on English text.

2.6.4 Training

We are given a set of training examples, each of which can be converted into a sequence of vectors, $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(n)}$, and $\mathbf{x}^{(n)} = \text{EOS}$. We write \mathbf{x} for the whole sequence of vectors.

We write θ for the collection of all the parameters of the model, flattened into a single vector: $\theta = (\mathbf{h}^{(0)}, \mathbf{A}, \mathbf{B}, \mathbf{c}, \mathbf{D}, \mathbf{e})$. For each training example and each time step i , the RNN predicts the probability of word $i + 1$ as a vector $\mathbf{y}^{(i)}$, which we now write as $\mathbf{y}^{(i)}(\mathbf{x}^{(1, \dots, i)}, \theta)$ to make its dependence on \mathbf{x} and θ explicit.

During training, our goal is to find the parameter values that maximize the log-likelihood,¹

$$L(\theta) = \log \prod_{\mathbf{x} \in \text{data}} P(\mathbf{x}^{(1)} \dots \mathbf{x}^{(n+1)}; \theta) \quad (2.31)$$

$$= \sum_{\mathbf{x} \in \text{data}} \log P(\mathbf{x}^{(1)} \dots \mathbf{x}^{(n+1)}; \theta) \quad (2.32)$$

$$= \sum_{\mathbf{x} \in \text{data}} \sum_{i=0}^n \mathbf{x}^{(i+1)} \cdot \log \mathbf{y}^{(i)}(\mathbf{x}^{(1, \dots, i)}, \theta) \quad (2.33)$$

where the log is elementwise and \cdot is a vector dot product (inner product). Since each $\mathbf{x}^{(i)}$ is a one-hot vector, dotting it with another vector selects a single component from that other vector, which in this case is the log-probability of the i th word.

To maximize this function, there are lots of different methods. We're going to look at the easiest (but still very practical) method, *stochastic gradient ascent*.² It is also known as the method of steepest ascent. Imagine that the log-likelihood is an infinite, many-dimensional surface. Each point on the surface corresponds to a setting of the θ 's, and the altitude of the point is the log-likelihood for that setting of the θ 's. We want to find the highest point on the surface. We start at an arbitrary location and then repeatedly move a little bit in the steepest uphill direction.

In pseudocode, gradient ascent looks like this:

initialize parameters θ randomly

repeat

$$\theta \leftarrow \theta + \eta \frac{\partial}{\partial \theta} L(\theta)$$

until done

The randomness of the initialization is important, because there are many situations where if two parameters are initialized to the same value, they'll always have the same value and therefore be redundant.

The function $\frac{\partial}{\partial \theta} L$ is the gradient of L and gives the direction, at θ , that goes uphill the steepest. These days, it's uncommon to need to figure out what the gradient is by hand, because there are numerous automatic differentiation packages that do this for you.

¹Since "likelihood," "log-likelihood," and "loss function" all start with L , it's common to write L for all three. Here, it stands for "log-likelihood."

²If we're minimizing a function, then we use stochastic gradient *descent*, and this is the name that the method is more commonly known by.

The *learning rate* $\eta > 0$ controls how far we move at each step. (What happens if η is too small? too big?) To guarantee convergence, η should decrease over time (for example, $\eta = 1/t$), but it's also common in practice to leave it fixed. See below for another common trick.

In *stochastic* gradient ascent, we work on just one sentence at a time. Let $L(\mathbf{x}, \theta)$ be the log-likelihood of just one sentence, that is,

$$L(\mathbf{x}, \theta) = \sum_{i=0}^n \mathbf{x}^{(i+1)} \cdot \log y^{(i)}(\mathbf{x}^{(1,\dots,i)}, \theta).$$

It could be thought of as an approximation to the full log-likelihood $L(\theta)$. Then stochastic gradient ascent goes like this:

initialize parameters θ to random numbers

repeat

for each sentence \mathbf{x} **do**

$$\theta \leftarrow \theta + \eta \frac{\partial}{\partial \theta} L(\mathbf{x}, \theta)$$

end for

until done

Each pass through the training data is called an *epoch*. This method has two advantages and one disadvantage compared to full gradient ascent:

- + Computing the gradient for one sentence uses much less memory.
- + Updating the model after every sentence instead of waiting until the end of the data means that the model can get better faster.
- Because the per-sentence log-likelihoods are only an approximation to the full log-likelihood, the updates can temporarily take us in the wrong direction.

The next section talks about one way to mitigate this disadvantage.

2.6.5 Tricks

There are a number of tricks that are important for training well. This is not a complete list, but these are the most essential and/or easiest tricks.

Validation. The above pseudocode doesn't specify how to choose the learning rate η or when to stop. There are many ways to do this, but one tried-and-true method is to look at the score (likelihood or some other metric) on held-out data (also known as development or validation data). At the end of each epoch, run on the validation data and compute the score. If it got worse, multiply the learning rate by $\frac{1}{2}$ and continue. Usually, the validation score will start to go up again. If the learning rate goes below some threshold (say, after a certain number of halvings), stop training.

Shuffling. Because stochastic gradient ascent updates the model based on one sentence at a time, it will have a natural tendency to remember the recent sentences most. To mitigate this effect, before each epoch, randomly shuffle the order of the training sentences or minibatches.

Minibatching. To speed up training and/or to reduce random variations between sentences, it's standard to train on a small number (10–1000) of sentences at a time instead of a single sentence at a time. If we can process the sentences in one minibatch in parallel, we get a huge speedup. For example, if the model contains the matrix-vector product $\mathbf{A}\mathbf{h}$ where \mathbf{A} is a parameter matrix and \mathbf{h} is a vector that depends on the input sentence, then with minibatching, \mathbf{h} becomes a matrix (one row for each sentence), and $\mathbf{A}\mathbf{h}$ can become a matrix-matrix product, which is much faster than a bunch of matrix-vector products. You just have to make sure that the indices match up correctly: $\mathbf{h}\mathbf{A}^\top$ or in PyTorch, `A @ h[:, :, None]`.

However, a major nuisance is that the sentences are all different lengths. The typical solution goes like this:

- Sort all the sentences by length.
- Divide up the sentences into minibatches. Because of the sorting, each minibatch contains sentences with similar lengths.
- In each minibatch, equalize the lengths of sentences by appending a special symbol PAD.
- When computing L , mask out the PAD symbols to avoid biasing the model towards predicting PAD (not to mention wasting training time).

Gradient clipping. When using SGA on RNNs, a common problem is known as the *vanishing gradient* problem and its evil twin, the *exploding gradient* problem. What happens is that L is a very long chain of functions (n times a constant). When we differentiate L , then by the chain rule, the partial derivatives are products of the partial derivatives of the functions in the chain. Suppose these partial derivatives are small numbers (less than 1). Then the product of many of them will be a vanishingly small number, and the gradient update will not have very much effect. Or, suppose these partial derivatives are large numbers (greater than 1). Then the product of many of them will explode into a very large number, and the gradient update will be very damaging. This is definitely the more serious problem, and preventing it is important. There are fancier learning methods than SGA that alleviate this problem (currently, the most popular is probably Adam), but for SGA, the simplest fix is *gradient clipping*: just check if the norm of the gradient is bigger than 5, and if so, scale it so that its norm is just 5.

Bibliography

- Chen, Stanley F. and Joshua Goodman (1998). *An Empirical Study of Smoothing Techniques for Language Modeling*. Tech. rep. TR-10-98. Harvard University Center for Research in Computing Technology. URL: <http://nrs.harvard.edu/urn-3:HUL.InstRepos:25104739>.
- Karpathy, Andrej, Justin Johnson, and Li Fei-Fei (2016). “Visualizing and Understanding Recurrent Neural Networks”. In: *Proc. ICLR*. URL: <https://arxiv.org/abs/1506.02078>.
- Kleene, S. C. (1951). *Representation of Events in Nerve Nets and Finite Automata*. Tech. rep. RM-704. RAND. URL: https://www.rand.org/content/dam/rand/pubs/research_memoranda/2008/RM704.pdf.
- McCulloch, Warren S. and Walter Pitts (1943). “A Logical Calculus of the Ideas Immanent in Nervous Activity”. In: *Bulletin of Mathematical Biophysics* 5, pp. 115–133. URL: <https://doi.org/10.1007/BF02478259>.
- Rabin, M. O. and D. Scott (1959). “Finite Automata and Their Decision Problems”. In: *IBM Journal of Research and Development* 3.2, pp. 114–125. URL: <https://doi.org/10.1147/rd.32.0114>.
- Sennrich, Rico, Barry Haddow, and Alexandra Birch (2016). “Neural Machine Translation of Rare Words with Subword Units”. In: *Proc. ACL*, pp. 1715–1725. doi: 10.18653/v1/P16-1162.