

6.036 Lecture Notes

Contents

1	Introduction	4
1	Problem class	5
1.1	Supervised learning	5
1.1.1	Classification	5
1.1.2	Regression	6
1.2	Unsupervised learning	6
1.2.1	Density estimation	6
1.2.2	Clustering	6
1.2.3	Dimensionality reduction	6
1.3	Reinforcement learning	6
1.4	Sequence learning	7
1.5	Other settings	7
2	Assumptions	7
3	Evaluation criteria	8
4	Model type	9
4.1	No model	9
4.2	Prediction rule	9
5	Model class and parameter fitting	9
6	Algorithm	10
2	Linear classifiers	11
1	Classification	11
2	Learning algorithm	12
3	Linear classifiers	12
4	Learning linear classifiers	13
5	Evaluating a learning algorithm	14
3	The Perceptron	15
1	Algorithm	15
2	Offset	16
3	Theory of the perceptron	17
3.1	Linear separability	18
3.2	Convergence theorem	18

4	Feature representation	21
1	Polynomial basis	22
2	Hand-constructing features for real domains	24
2.1	Discrete features	25
2.2	Text	26
2.3	Numeric values	26
5	Margin Maximization	27
1	Machine learning as optimization	27
2	Regularization	27
3	Maximizing the margin	28
6	Gradient Descent	33
1	One dimension	33
2	Multiple dimensions	35
3	Application to SVM objective	35
4	Stochastic Gradient Descent	36
7	Regression	38
1	Analytical solution: ordinary least squares	39
2	Regularization	40
3	Optimization via gradient descent	41
8	Neural Networks	43
1	Basic element	44
2	Networks	44
2.1	Single layer	45
2.2	Many layers	46
3	Choices of activation function	46
4	Error back-propagation	49
5	Training	51
6	Loss functions and activation functions	52
6.1	Two-class classification and log likelihood	52
6.2	Multi-class classification and log likelihood	53
7	Optimizing neural network parameters	54
7.1	Batches	54
7.2	Adaptive step-size	55
7.2.1	Running averages	55
7.2.2	Momentum	56
7.2.3	Adadelta	56
7.2.4	Adam	57
8	Regularization	58
8.1	Methods related to ridge regression	58
8.2	Dropout	58
8.3	Batch Normalization	59
9	Convolutional Neural Networks	60
1	Filters	61
2	Max Pooling	63
3	Typical architecture	63

10 Sequential models	66
1 State machines	66
2 Markov decision processes	68
2.1 Finite-horizon solutions	68
2.1.1 Evaluating a given policy	69
2.1.2 Finding an optimal policy	69
2.2 Infinite-horizon solutions	70
2.2.1 Evaluating a policy	71
2.2.2 Finding an optimal policy	71
2.2.3 Theory	72
11 Reinforcement learning	73
1 Bandit problems	74
2 Sequential problems	75
2.1 Model-based RL	75
2.2 Policy search	76
2.3 Value function learning	76
2.3.1 Q-learning	76
2.3.2 Function approximation	78
2.3.3 Fitted Q-learning	79
12 Recurrent Neural Networks	80
1 RNN model	80
2 Sequence-to-sequence RNN	81
3 Back-propagation through time	82
4 Training a language model	85
5 Vanishing gradients and gating mechanisms	85
5.1 Simple gated recurrent networks	86
5.2 Long short-term memory	86
13 Recommender systems	88
1 Content-based recommendations	88
2 Collaborative filtering	89
2.1 Optimization	92
2.1.1 Alternating least squares	93
2.1.2 Stochastic gradient descent	94
14 Non-parametric methods	95
1 Intro	95
2 Trees	95
2.1 Regression	96
2.1.1 Building a tree	96
2.1.2 Pruning	97
2.2 Classification	97
3 Bagging	99
3.1 Random Forests	100
4 Nearest Neighbor	100

CHAPTER 1

Introduction

The main focus of machine learning is *making decisions or predictions based on data*. There are a number of other fields with significant overlap in technique, but difference in focus: in economics and psychology, the goal is to discover underlying causal processes and in statistics it is to find a model that fits a data set well. In those fields, the end product is a model. In machine learning, we often fit models, but as a means to the end of making good predictions or decisions.

This story paraphrased from a post on 9/4/12 at andrewgelman.com

As machine-learning (ML) methods have improved in their capability and scope, ML has become the best way, measured in terms of speed, human engineering time, and robustness, to make many applications. Great examples are face detection and speech recognition and many kinds of language-processing tasks. Almost any application that involves understanding data or signals that come from the real world can be best addressed using machine learning.

One crucial aspect of machine learning approaches to solving problems is that human engineering plays an important role. A human still has to *frame* the problem: acquire and organize data, design a space of possible solutions, select a learning algorithm and its parameters, apply the algorithm to the data, validate the resulting solution to decide whether it's good enough to use, etc. These steps are of great importance.

and often undervalued

The conceptual basis of learning from data is the *problem of induction*: Why do we think that previously seen data will help us predict the future? This is a serious philosophical problem of long standing. We will operationalize it by making assumptions, such as that all training data are IID (independent and identically distributed) and that queries will be drawn from the same distribution as the training data, or that the answer comes from a set of possible answers known in advance.

Bertrand Russell is my hero. -lpk

In general, we need to solve these two problems:

- **estimation:** When we have data that are noisy reflections of some underlying quantity of interest, we have to aggregate the data and make estimates or predictions about the quantity. How do we deal with the fact that, for example, the same treatment may end up with different results on different trials? How can we predict how well an estimate may compare to future results?
- **generalization:** How can we predict results of a situation or experiment that we have never encountered before in our data set?

We can describe problems and their solutions using six characteristics, three of which characterize the problem and three of which characterize the solution:

1. **Problem class:** What is the nature of the training data and what kinds of queries will be made at testing time?
2. **Assumptions:** What do we know about the source of the data or the form of the solution?
3. **Evaluation criteria:** What is the goal of the prediction or estimation system? How will the answers to individual queries be evaluated? How will the overall performance of the system be measured?
4. **Model type:** Will an intermediate model be made? What aspects of the data will be modeled? How will the model be used to make predictions?
5. **Model class:** What particular parametric class of models will be used? What criterion will we use to pick a particular model from the model class?
6. **Algorithm:** What computational process will be used to fit the model to the data and/or to make predictions?

Without making some assumptions about the nature of the process generating the data, we cannot perform generalization. In the following sections, we elaborate on these ideas.

1 Problem class

There are many different *problem classes* in machine learning. They vary according to what kind of data is provided and what kind of conclusions are to be drawn from it. Five standard problem classes are described below, to establish some notation and terminology.

In this course, we will focus on classification and regression (two examples of supervised learning), and will touch on reinforcement learning and sequence learning.

1.1 Supervised learning

The idea of *supervised* learning is that the learning system is given inputs and told which specific outputs should be associated with them. We divide up supervised learning based on whether the outputs are drawn from a small finite set (classification) or a large finite or continuous set (regression).

1.1.1 Classification

Training data \mathcal{D}_n is in the form of a set of pairs $\{(x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)})\}$ where $x^{(i)}$ represents an object to be classified, most typically a d -dimensional vector of real and/or discrete values, and $y^{(i)}$ is an element of a discrete set of values. The y values are sometimes called *target values*.

A classification problem is *binary* or *two-class* if $y^{(i)}$ is drawn from a set of two possible values; otherwise, it is called *multi-class*.

The goal in a classification problem is ultimately, given a new input value $x^{(n+1)}$, to predict the value of $y^{(n+1)}$.

Classification problems are a kind of *supervised learning*, because the desired output (or class) $y^{(i)}$ is specified for each of the training examples $x^{(i)}$.

Don't feel you have to memorize all these kinds of learning, etc. We just want you to have a very high-level view of (part of) the breadth of the field.

Many textbooks use x_i and t_i instead of $x^{(i)}$ and $y^{(i)}$. We find that notation somewhat difficult to manage when $x^{(i)}$ is itself a vector and we need to talk about its elements. The notation we are using is standard in some other parts of the machine-learning literature.

1.1.2 Regression

Regression is like classification, except that $y^{(i)} \in \mathbb{R}^k$.

1.2 Unsupervised learning

Unsupervised learning doesn't involve learning a function from inputs to outputs based on a set of input-output pairs. Instead, one is given a data set and generally expected to find some patterns or structure inherent in it.

1.2.1 Density estimation

Given samples $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^D$ drawn IID from some distribution $\Pr(X)$, the goal is to predict the probability $\Pr(x^{(n+1)})$ of an element drawn from the same distribution. Density estimation sometimes plays a role as a "subroutine" in the overall learning method for supervised learning, as well.

IID stands for *independent and identically distributed*, which means that the elements in the set are related in the sense that they all come from the same underlying probability distribution, but not in any other ways.

1.2.2 Clustering

Given samples $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^D$, the goal is to find a partitioning (or "clustering") of the samples that groups together samples that are similar. There are many different objectives, depending on the definition of the similarity between samples and exactly what criterion is to be used (e.g., minimize the average distance between elements inside a cluster and maximize the average distance between elements across clusters). Other methods perform a "soft" clustering, in which samples may be assigned 0.9 membership in one cluster and 0.1 in another. Clustering is sometimes used as a step in density estimation, and sometimes to find useful structure in data.

1.2.3 Dimensionality reduction

Given samples $x^{(1)}, \dots, x^{(n)} \in \mathbb{R}^D$, the problem is to re-represent them as points in a d -dimensional space, where $d < D$. The goal is typically to retain information in the data set that will, e.g., allow elements of one class to be discriminated from another.

Dimensionality reduction is a standard technique which is particularly useful for visualizing or understanding high-dimensional data. If the goal is ultimately to perform regression or classification on the data after the dimensionality is reduced, it is usually best to articulate an objective for the overall prediction problem rather than to first do dimensionality reduction without knowing which dimensions will be important for the prediction task.

1.3 Reinforcement learning

In reinforcement learning, the goal is to learn a mapping from input values x to output values y , but without a direct supervision signal to specify which output values y are best for a particular input. There is no training set specified *a priori*. Instead, the learning problem is framed as an agent interacting with an environment, in the following setting:

- The agent observes the current state, $x^{(0)}$.
- It selects an action, $y^{(0)}$.
- It receives a reward, $r^{(0)}$, which depends on $x^{(0)}$ and possibly $y^{(0)}$.
- The environment transitions probabilistically to a new state, $x^{(1)}$, with a distribution that depends only on $x^{(0)}$ and $y^{(0)}$.

- The agent observes the current state, $x^{(1)}$.
- ...

The goal is to find a policy π , mapping x to y , (that is, states to actions) such that some long-term sum or average of rewards r is maximized.

This setting is very different from either supervised learning or unsupervised learning, because the agent's action choices affect both its reward and its ability to observe the environment. It requires careful consideration of the long-term effects of actions, as well as all of the other issues that pertain to supervised learning.

1.4 Sequence learning

In sequence learning, the goal is to learn a mapping from *input sequences* x_0, \dots, x_n to *output sequences* y_1, \dots, y_m . The mapping is typically represented as a *state machine*, with one function f used to compute the next hidden internal state given the input, and another function g used to compute the output given the current hidden state.

It is supervised in the sense that we are told what output sequence to generate for which input sequence, but the internal functions have to be learned by some method other than direct supervision, because we don't know what the hidden state sequence is.

1.5 Other settings

There are many other problem settings. Here are a few.

In *semi-supervised learning*, we have a supervised-learning training set, but there may be an additional set of $x^{(i)}$ values with no known $y^{(i)}$. These values can still be used to improve learning performance if they are drawn from $\Pr(X)$ that is the marginal of $\Pr(X, Y)$ that governs the rest of the data set.

In *active learning*, it is assumed to be expensive to acquire a label $y^{(i)}$ (imagine asking a human to read an x-ray image), so the learning algorithm can sequentially ask for particular inputs $x^{(i)}$ to be labeled, and must carefully select queries in order to learn as effectively as possible while minimizing the cost of labeling.

In *transfer learning* (also called *meta-learning*), there are multiple tasks, with data drawn from different, but related, distributions. The goal is for experience with previous tasks to apply to learning a current task in a way that requires decreased experience with the new task.

2 Assumptions

The kinds of assumptions that we can make about the data source or the solution include:

- The data are independent and identically distributed.
- The data are generated by a Markov chain.
- The process generating the data might be adversarial.
- The "true" model that is generating the data can be perfectly described by one of some particular set of hypotheses.

The effect of an assumption is often to reduce the "size" or "expressiveness" of the space of possible hypotheses and therefore reduce the amount of data required to reliably identify an appropriate hypothesis.

3 Evaluation criteria

Once we have specified a problem class, we need to say what makes an output or the answer to a query good, given the training data. We specify evaluation criteria at two levels: how an individual prediction is scored, and how the overall behavior of the prediction or estimation system is scored.

The quality of predictions from a learned model is often expressed in terms of a *loss function*. A loss function $L(g, a)$ tells you how much you will be penalized for making a guess g when the answer is actually a . There are many possible loss functions. Here are some frequently used examples:

- **0-1 Loss** applies to predictions drawn from finite domains.

$$L(g, a) = \begin{cases} 0 & \text{if } g = a \\ 1 & \text{otherwise} \end{cases}$$

If the actual values are drawn from a continuous distribution, the probability they would ever be equal to some predicted g is 0 (except for some weird cases).

- **Squared loss**

$$L(g, a) = (g - a)^2$$

- **Linear loss**

$$L(g, a) = |g - a|$$

- **Asymmetric loss** Consider a situation in which you are trying to predict whether someone is having a heart attack. It might be much worse to predict “no” when the answer is really “yes”, than the other way around.

$$L(g, a) = \begin{cases} 1 & \text{if } g = 1 \text{ and } a = 0 \\ 10 & \text{if } g = 0 \text{ and } a = 1 \\ 0 & \text{otherwise} \end{cases}$$

Any given prediction rule will usually be evaluated based on multiple predictions and the loss of each one. At this level, we might be interested in:

- Minimizing expected loss over all the predictions (also known as risk)
- Minimizing maximum loss: the loss of the worst prediction
- Minimizing or bounding regret: how much worse this predictor performs than the best one drawn from some class
- Characterizing asymptotic behavior: how well the predictor will perform in the limit of infinite training data
- Finding algorithms that are probably approximately correct: they probably generate a hypothesis that is right most of the time.

There is a theory of rational agency that argues that you should always select the action that *minimizes the expected loss*. This strategy will, for example, make you the most money in the long run, in a gambling setting. Expected loss is also sometimes called *risk* in the machine-learning literature, but that term means other things in economics or other parts of decision theory, so be careful...it's risky to use it. We will, most of the time, concentrate on this criterion.

Of course, there are other models for action selection and it's clear that people do not always (or maybe even often) select actions that follow this rule.

4 Model type

Recall that the goal of a machine-learning system is typically to estimate or generalize, based on data provided. Below, we examine the role of model-making in machine learning.

4.1 No model

In some simple cases, in response to queries, we can generate predictions directly from the training data, without the construction of any intermediate model. For example, in regression or classification, we might generate an answer to a new query by averaging answers to recent queries, as in the *nearest neighbor* method.

4.2 Prediction rule

This two-step process is more typical:

1. "Fit" a model to the training data
2. Use the model directly to make predictions

In the *prediction rule* setting of regression or classification, the model will be some hypothesis or prediction rule $y = h(x; \theta)$ for some functional form h . The idea is that θ is a vector of one or more parameter values that will be determined by fitting the model to the training data and then be held fixed. Given a new $x^{(n+1)}$, we would then make the prediction $h(x^{(n+1)}; \theta)$.

The fitting process is often articulated as an optimization problem: Find a value of θ that minimizes some criterion involving θ and the data. An optimal strategy, if we knew the actual underlying distribution on our data, $\Pr(X, Y)$ would be to predict the value of y that minimizes the *expected loss*, which is also known as the *test error*. If we don't have that actual underlying distribution, or even an estimate of it, we can take the approach of minimizing the *training error*: that is, finding the prediction rule h that minimizes the average loss on our training data set. So, we would seek θ that minimizes

$$\mathcal{E}_n(\theta) = \frac{1}{n} \sum_{i=1}^n L(h(x^{(i)}; \theta), y^{(i)}) ,$$

where the loss function $L(g, a)$ measures how bad it would be to make a guess of g when the actual value is a .

We will find that minimizing training error alone is often not a good choice: it is possible to emphasize fitting the current data too strongly and end up with a hypothesis that does not generalize well when presented with new x values.

5 Model class and parameter fitting

A model *class* \mathcal{M} is a set of possible models, typically parameterized by a vector of parameters Θ . What assumptions will we make about the form of the model? When solving a regression problem using a prediction-rule approach, we might try to find a linear function $h(x; \theta, \theta_0) = \theta^T x + \theta_0$ that fits our data well. In this example, the parameter vector $\Theta = (\theta, \theta_0)$.

For problem types such as discrimination and classification, there are huge numbers of model classes that have been considered...we'll spend much of this course exploring these model classes, especially neural networks models. We will almost completely restrict our

We write $f(a; b)$ to describe a function that is usually applied to a single argument a , but is a member of a parametric family of functions, with the particular function determined by parameter value b . So, for example, we might write $h(x; p) = x^p$ to describe a function of a single argument that is parameterized by p .

attention to model classes with a fixed, finite number of parameters. Models that relax this assumption are called “non-parametric” models.

How do we select a model class? In some cases, the machine-learning practitioner will have a good idea of what an appropriate model class is, and will specify it directly. In other cases, we may consider several model classes. In such situations, we are solving a *model selection* problem: model-selection is to pick a model class \mathcal{M} from a (usually finite) set of possible model classes; *model fitting* is to pick a particular model in that class, specified by parameters θ .

6 Algorithm

Once we have described a class of models and a way of scoring a model given data, we have an algorithmic problem: what sequence of computational instructions should we run in order to find a good model from our class? For example, determining the parameter vector θ which minimizes $\mathcal{E}_n(\theta)$ might be done using a familiar least-squares minimization algorithm, when the model h is a function being fit to some data x .

Sometimes we can use software that was designed, generically, to perform optimization. In many other cases, we use algorithms that are specialized for machine-learning problems, or for particular hypotheses classes.

Some algorithms are not easily seen as trying to optimize a particular criterion. In fact, the first algorithm we study for finding linear classifiers, the perceptron algorithm, has this character.

1 Classification

A binary *classifier* is a mapping from $\mathbb{R}^d \rightarrow \{-1, +1\}$. We'll often use the letter h (for hypothesis) to stand for a classifier, so the classification process looks like:

$$x \rightarrow \boxed{h} \rightarrow y .$$

Actually, general classifiers can have a range which is any discrete set, but we'll work with this specific case for a while.

Real life rarely gives us vectors of real numbers; the x we really want to classify is usually something like a song, image, or person. In that case, we'll have to define a function $\varphi(x)$, whose domain is \mathbb{R}^d , where φ represents *features* of x , like a person's height or the amount of bass in a song, and then let the $h : \varphi(x) \rightarrow \{-1, +1\}$. In much of the following, we'll omit explicit mention of φ and assume that the $x^{(i)}$ are in \mathbb{R}^d , but you should always have in mind that some additional process was almost surely required to go from the actual input examples to their feature representation.

In *supervised learning* we are given a training data set of the form

$$\mathcal{D}_n = \left\{ \left(x^{(1)}, y^{(1)} \right), \dots, \left(x^{(n)}, y^{(n)} \right) \right\} .$$

We will assume that each $x^{(i)}$ is a $d \times 1$ *column vector*. The intended meaning of this data is that, when given an input $x^{(i)}$, the learned hypothesis should generate output $y^{(i)}$.

What makes a classifier useful? That it works well on *new* data; that is, that it makes good predictions on examples it hasn't seen. But we don't know exactly what data this classifier might be tested on when we use it in the real world. So, we have to *assume* a connection between the training data and testing data; typically, they are drawn independently from the same probability distribution.

My favorite analogy is to problem sets. We evaluate a student's ability to *generalize* by putting questions on the exam that were not on the homework (training set).

Given a training set \mathcal{D}_n and a classifier h , we can define the *training error* of h to be

$$\mathcal{E}_n(h) = \frac{1}{n} \sum_{i=1}^n \begin{cases} 1 & h(x^{(i)}) \neq y^{(i)} \\ 0 & \text{otherwise} \end{cases} .$$

For now, we will try to find a classifier with small training error (later, with some added criteria) and hope it *generalizes well* to new data, and has a small *test error*

$$\mathcal{E}(h) = \frac{1}{n'} \sum_{i=n+1}^{n+n'} \begin{cases} 1 & h(x^{(i)}) \neq y^{(i)} \\ 0 & \text{otherwise} \end{cases}$$

on n' new examples that were not used in the process of finding the classifier.

2 Learning algorithm

A *hypothesis class* \mathcal{H} is a set (finite or infinite) of possible classifiers, each of which represents a mapping from $\mathbb{R}^d \rightarrow \{-1, +1\}$.

A *learning algorithm* is a procedure that takes a data set \mathcal{D}_n as input and returns an element h of \mathcal{H} ; it looks like

$$\mathcal{D}_n \longrightarrow \boxed{\text{learning alg } (\mathcal{H})} \longrightarrow h$$

We will find that the choice of \mathcal{H} can have a big impact on the test error of the h that results from this process. One way to get h that generalizes well is to restrict the size, or “expressiveness” of \mathcal{H} .

3 Linear classifiers

We’ll start with the hypothesis class of *linear classifiers*. They are (relatively) easy to understand, simple in a mathematical sense, powerful on their own, and the basis for many other more sophisticated methods.

A linear classifier in d dimensions is defined by a vector of parameters $\theta \in \mathbb{R}^d$ and scalar $\theta_0 \in \mathbb{R}$. So, the hypothesis class \mathcal{H} of linear classifiers in d dimensions is the *set* of all vectors in \mathbb{R}^{d+1} . We’ll assume that θ is an $n \times 1$ column vector.

Given particular values for θ and θ_0 , the classifier is defined by

$$h(x; \theta, \theta_0) = \text{sign}(\theta^T x + \theta_0) = \begin{cases} +1 & \text{if } \theta^T x + \theta_0 > 0 \\ -1 & \text{otherwise} \end{cases} .$$

Remember that we can think of θ, θ_0 as specifying a hyperplane. It divides \mathbb{R}^d , the space our $x^{(i)}$ points live in, into two half-spaces. The one that is on the same side as the normal vector is the *positive* half-space, and we classify all points in that space as positive. The half-space on the other side is *negative* and all points in it are classified as negative.

Let’s be careful about dimensions. We have assumed that x and θ are both $n \times 1$ column vectors. So $\theta^T x$ is 1×1 , which in math (but not necessarily numpy) is the same as a scalar.


Example: Let h be the linear classifier defined by $\theta = \begin{bmatrix} -1 \\ 1.5 \end{bmatrix}$, $\theta_0 = 3$.

The diagram below shows several points classified by h . In particular, let $x^{(1)} = \begin{bmatrix} 3 \\ 2 \end{bmatrix}$ and $x^{(2)} = \begin{bmatrix} 4 \\ -1 \end{bmatrix}$.

$$h(x^{(1)}; \theta, \theta_0) = \text{sign} \left(\begin{bmatrix} -1 & 1.5 \end{bmatrix} \begin{bmatrix} 3 \\ 2 \end{bmatrix} + 3 \right) = \text{sign}(3) = +1$$

$$h(x^{(2)}; \theta, \theta_0) = \text{sign} \left(\begin{bmatrix} -1 & 1.5 \end{bmatrix} \begin{bmatrix} 4 \\ -1 \end{bmatrix} + 3 \right) = \text{sign}(-2.5) = -1$$

Thus, $x^{(1)}$ and $x^{(2)}$ are given positive and negative classifications, respectively.



Study Question: What is green vector normal to the hyperplane? Specify it as a column vector.

Study Question: What change would you have to make to θ, θ_0 if you wanted to have the separating hyperplane in the same place, but to classify all the points labeled '+' in the diagram as negative and all the points labeled '-' in the diagram as positive?

4 Learning linear classifiers

Now, given a data set and the hypothesis class of linear classifiers, our objective will be to find the linear classifier with the smallest possible training error.

This is a well-formed optimization problem. But it's not computationally easy!

We'll start by considering a very simple learning algorithm. The idea is to generate k possible hypotheses by generating their parameter vectors at random. Then, we can evaluate the training-set error on each of the hypotheses and return the hypothesis that has the lowest training error (breaking ties arbitrarily).

It's a good idea to think of the "stupidest possible" solution to a problem, before trying to get clever. Here's a fairly (but not completely) stupid algorithm.

RANDOM-LINEAR-CLASSIFIER(\mathcal{D}_n, k, d)

- 1 **for** $j = 1$ **to** k
- 2 randomly sample $(\theta^{(j)}, \theta_0^{(j)})$ from $(\mathbb{R}^d, \mathbb{R})$
- 3 $j^* = \arg \min_{j \in \{1, \dots, k\}} \mathcal{E}_n(\theta^{(j)}, \theta_0^{(j)})$
- 4 **return** $(\theta^{(j^*)}, \theta_0^{(j^*)})$

A note about notation.

Study Question: What do you think happens to $\mathcal{E}_n(h)$, where h is the hypothesis returned by RANDOM-LINEAR-CLASSIFIER, as k is increased?

Study Question: What properties of \mathcal{D}_n do you think will have an effect on $\mathcal{E}_n(h)$?

This might be new notation: $\arg \min_x f(x)$ means the value of x for which $f(x)$ is the smallest. Sometimes we write $\arg \min_{x \in \mathcal{X}} f(x)$ when we want to explicitly specify the set \mathcal{X} of values of x over which we want to minimize.

5 Evaluating a learning algorithm

How should we evaluate the performance of a *classifier* h ? The best method is to measure *test error* on data that was not used to train it.

How should we evaluate the performance of a *learning algorithm*? This is trickier. There are many potential sources of variability in the possible result of computing test error on a learned hypothesis h :

- Which particular *training examples* occurred in \mathcal{D}_n
- Which particular *testing examples* occurred in \mathcal{D}_n
- Randomization inside the learning *algorithm* itself

Generally, we would like to execute the following process multiple times:

- Train on a new training set
- Evaluate resulting h on a testing set *that does not overlap the training set*

Doing this multiple times controls for possible poor choices of training set or unfortunate randomization inside the algorithm itself.

One concern is that we might need a lot of data to do this, and in many applications data is expensive or difficult to acquire. We can re-use data with *cross validation* (but it's harder to do theoretical analysis).

CROSS-VALIDATE(\mathcal{D}, k)

- 1 divide \mathcal{D} into k chunks $\mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_k$ (of roughly equal size)
- 2 **for** $i = 1$ **to** k
- 3 train h_i on $\mathcal{D} \setminus \mathcal{D}_i$ (withholding chunk \mathcal{D}_i)
- 4 compute "test" error $\mathcal{E}_i(h_i)$ on withheld data \mathcal{D}_i
- 5 **return** $\frac{1}{k} \sum_{i=1}^k \mathcal{E}_i(h_i)$

It's very important to understand that cross-validation neither delivers nor evaluates a single particular hypothesis h . It evaluates the *algorithm* that produces hypotheses.

CHAPTER 3

The Perceptron

First of all, the coolest algorithm name! It is based on the 1943 model of neurons made by McCulloch and Pitts and by Hebb. It was developed by Rosenblatt in 1962. At the time, it was not interpreted as attempting to optimize any particular criteria; it was presented directly as an algorithm. There has, since, been a huge amount of study and analysis of its convergence properties and other aspects of its behavior.

Well, maybe “neocognitron,” also the name of a real ML algorithm, is cooler.

1 Algorithm

Recall that we have a training dataset \mathcal{D}_n with $x \in \mathbb{R}^d$, and $y \in \{-1, +1\}$. The Perceptron algorithm trains a binary classifier $h(x; \theta, \theta_0)$ using the following algorithm to find θ and θ_0 using τ iterative steps:

We use Greek letter τ here instead of T so we don't confuse it with transpose!

PERCEPTRON(τ, \mathcal{D}_n)

```
1  $\theta = [0 \ 0 \ \dots \ 0]^T$ 
2  $\theta_0 = 0$ 
3 for  $t = 1$  to  $\tau$ 
4   for  $i = 1$  to  $n$ 
5     if  $y^{(i)} (\theta^T x^{(i)} + \theta_0) \leq 0$ 
6        $\theta = \theta + y^{(i)} x^{(i)}$ 
7        $\theta_0 = \theta_0 + y^{(i)}$ 
8 return  $\theta, \theta_0$ 
```

Intuitively, on each step, if the current hypothesis θ, θ_0 classifies example $x^{(i)}$ correctly, then no change is made. If it classifies $x^{(i)}$ incorrectly, then it moves θ, θ_0 so that it is “closer” to classifying $x^{(i)}, y^{(i)}$ correctly.

Let's check dimensions. Remember that θ is $n \times 1$, $x^{(i)}$ is $n \times 1$, and $y^{(i)}$ is a scalar. Does everything match?

Note that if the algorithm ever goes through one iteration of the loop on line 4 without making an update, it will never make any further updates (verify that you believe this!) and so it should just terminate at that point.

Study Question: What is true about \mathcal{E}_n if that happens?

Example: Let h be the linear classifier defined by $\theta^{(0)} = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$, $\theta_0^{(0)} = 1$. The diagram below shows several points classified by h . However, in this case, h (represented by the bold line) misclassifies the point $x^{(1)} = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$ which has label $y^{(1)} = 1$. Indeed,


$$y^{(1)} (\theta^T x^{(1)} + \theta_0) = [1 \quad -1] \begin{bmatrix} 1 \\ 3 \end{bmatrix} + 1 = -1 < 0$$

By running an iteration of the Perceptron algorithm, we update

$$\theta^{(1)} = \theta^{(0)} + y^{(1)} x^{(1)} = \begin{bmatrix} 2 \\ 2 \end{bmatrix}$$

$$\theta_0^{(1)} = \theta_0^{(0)} + y^{(1)} = 2$$

The new classifier (represented by the dashed line) now correctly classifies each point.



A really important fact about the perceptron algorithm is that, if there is a linear classifier with 0 training error, then this algorithm will (eventually) find it! We'll look at a proof of this in detail, next.

2 Offset

Sometimes, it can be easier to implement or analyze classifiers of the form

$$h(x; \theta) = \begin{cases} +1 & \text{if } \theta^T x > 0 \\ -1 & \text{otherwise.} \end{cases}$$

Without an explicit offset term (θ_0), this separator must pass through the origin, which may appear to be limiting. However, we can convert any problem involving a linear separator *with* offset into one with *no* offset (but of higher dimension)!

Consider the d -dimensional linear separator defined by $\theta = [\theta_1 \ \theta_2 \ \dots \ \theta_d]$ and offset θ_0 .

- to each data point $x \in \mathcal{D}$, append a coordinate with value +1, yielding

$$x_{\text{new}} = [x_1 \ \dots \ x_d \ +1]^T$$

- define

$$\theta_{\text{new}} = [\theta_1 \ \dots \ \theta_d \ \theta_0]$$

Then,

$$\begin{aligned} \theta_{\text{new}}^T \cdot x_{\text{new}} &= \theta_1 x_1 + \dots + \theta_d x_d + \theta_0 \cdot 1 \\ &= \theta^T x + \theta_0 \end{aligned}$$

Thus, θ_{new} is an equivalent $((d + 1)$ -dimensional) separator to our original, but with no offset.

Consider the data set:

$$\begin{aligned} X &= [[1], [2], [3], [4]] \\ Y &= [[+1], [+1], [-1], [-1]] \end{aligned}$$

It is linearly separable in $d = 1$ with $\theta = [-1]$ and $\theta_0 = 2.5$. But it is not linearly separable through the origin! Now, let

$$X_{\text{new}} = \begin{bmatrix} [1] & [2] & [3] & [4] \\ [1] & [1] & [1] & [1] \end{bmatrix}$$

This new dataset is separable through the origin, with $\theta_{\text{new}} = [-1, 2.5]^T$.

We can make a simplified version of the perceptron algorithm if we restrict ourselves to separators through the origin:

PERCEPTRON-THROUGH-ORIGIN(τ, \mathcal{D}_n)

```

1   $\theta = [0 \ 0 \ \dots \ 0]^T$ 
2  for  $t = 1$  to  $\tau$ 
3      for  $i = 1$  to  $n$ 
4          if  $y^{(i)} (\theta^T x^{(i)}) \leq 0$ 
5               $\theta = \theta + y^{(i)} x^{(i)}$ 
6  return  $\theta$ 
```

We list it here because this is the version of the algorithm we'll study in more detail.

3 Theory of the perceptron

Now, we'll say something formal about how well the perceptron algorithm really works. We start by characterizing the set of problems that can be solved perfectly by the perceptron algorithm, and then prove that, in fact, it can solve these problems. In addition, we provide a notion of what makes a problem difficult for perceptron and link that notion of difficulty to the number of iterations the algorithm will take.

3.1 Linear separability

A training set \mathcal{D}_n is *linearly separable* if there exist θ, θ_0 such that, for all $i = 1, 2, \dots, n$:

$$y^{(i)} \left(\theta^T x^{(i)} + \theta_0 \right) > 0 .$$

Another way to say this is that all predictions on the training set are correct:

$$h(x^{(i)}; \theta, \theta_0) = y^{(i)} .$$

And, another way to say this is that the training error is zero:

$$\mathcal{E}_n(h) = 0 .$$

3.2 Convergence theorem

The basic result about the perceptron is that, if the training data \mathcal{D}_n is linearly separable, then the perceptron algorithm is guaranteed to find a linear separator.

We will more specifically characterize the linear separability of the dataset by the *margin* of the separator. We'll start by defining the margin of a point with respect to a hyperplane.

First, recall that the distance of a point x to the hyperplane θ, θ_0 is

$$\frac{\theta^T x + \theta_0}{\|\theta\|} .$$

Then, we'll define the *margin* of a *labeled point* (x, y) with respect to hyperplane θ, θ_0 to be

$$y \cdot \frac{\theta^T x + \theta_0}{\|\theta\|} .$$

This quantity will be positive if and only if the point x is classified as y by the linear classifier represented by this hyperplane.

Study Question: What sign does the margin have if the point is incorrectly classified? Be sure you can explain why.

Now, the *margin* of a *dataset* \mathcal{D}_n with respect to the hyperplane θ, θ_0 is the *minimum* margin of any point with respect to θ, θ_0 :


$$\min_i \left(y^{(i)} \cdot \frac{\theta^T x^{(i)} + \theta_0}{\|\theta\|} \right) .$$

The margin is positive if and only if all of the points in the data-set are classified correctly. In that case (only!) it represents the distance from the hyperplane to the closest point.

If the training data is *not* linearly separable, the algorithm will not be able to tell you for sure, in finite time, that it is not linearly separable. There are other algorithms that can test for linear separability with run-times $O(n^{d/2})$ or $O(d^{2n})$ or $O((n^{d-1} \log n))$.

Example: Let h be the linear classifier defined by $\theta = \begin{bmatrix} 1 \\ -1 \end{bmatrix}$, $\theta_0 = 1$.

The diagram below shows several points classified by h , one of which is misclassified. We compute the margin for each point:



$$y^{(1)} \cdot \frac{\theta^T x^{(1)} + \theta_0}{\|\theta\|} = 1 \cdot \frac{-2 + 1}{\sqrt{2}} = -\frac{\sqrt{2}}{2}$$

$$y^{(2)} \cdot \frac{\theta^T x^{(2)} + \theta_0}{\|\theta\|} = 1 \cdot \frac{1 + 1}{\sqrt{2}} = \sqrt{2}$$

$$y^{(3)} \cdot \frac{\theta^T x^{(3)} + \theta_0}{\|\theta\|} = -1 \cdot \frac{-3 + 1}{\sqrt{2}} = \sqrt{2}$$

Note that since point $x^{(1)}$ is misclassified, its margin is negative. Thus the margin for the whole data set is given by $-\frac{\sqrt{2}}{2}$.

Theorem 3.1 (Perceptron Convergence). For simplicity, we consider the case where the linear separator must pass through the origin. If the following conditions hold,

- (a) there exists θ^* such that $y^{(i)} \frac{\theta^{*T} x^{(i)}}{\|\theta^*\|} \geq \gamma$ for all $i = 1, \dots, n$ and for some $\gamma > 0$
- (b) all the examples have bounded magnitude: $\|x^{(i)}\| \leq R$ for all $i = 1, \dots, n$

then the perceptron algorithm will make at most $\left(\frac{R}{\gamma}\right)^2$ mistakes.

Proof. We initialize $\theta^{(0)} = 0$, and let $\theta^{(k)}$ define our hyperplane after the perceptron algorithm has made k mistakes. We are going to think about the angle between the hypothesis we have now, $\theta^{(k)}$ and the assumed good separator θ^* . Since they both go through the origin, if we can show that the angle between them is decreasing usefully on every iteration, then we will get close to that separator.

So, let's think about the cos of the angle between them, and recall, by the definition of dot product:

$$\cos(\theta^{(k)}, \theta^*) = \frac{\theta^{(k)} \cdot \theta^*}{\|\theta^*\| \|\theta^{(k)}\|}$$

We'll divide this up into two factors,

$$\cos(\theta^{(k)}, \theta^*) = \left(\frac{\theta^{(k)} \cdot \theta^*}{\|\theta^*\|} \right) \cdot \left(\frac{1}{\|\theta^{(k)}\|} \right), \quad (3.1)$$

and start by focusing on the first factor.

Without loss of generality, assume that the k^{th} mistake occurs on the i^{th} example $(x^{(i)}, y^{(i)})$.

$$\begin{aligned} \frac{\theta^{(k)} \cdot \theta^*}{\|\theta^*\|} &= \frac{(\theta^{(k-1)} + y^{(i)}x^{(i)}) \cdot \theta^*}{\|\theta^*\|} \\ &= \frac{\theta^{(k-1)} \cdot \theta^*}{\|\theta^*\|} + \frac{y^{(i)}x^{(i)} \cdot \theta^*}{\|\theta^*\|} \\ &\geq \frac{\theta^{(k-1)} \cdot \theta^*}{\|\theta^*\|} + \gamma \\ &\geq k\gamma \end{aligned}$$

where we have first applied the margin condition from (a) and then applied simple induction.

Now, we'll look at the second factor in equation 3.1. We note that since $(x^{(i)}, y^{(i)})$ is classified incorrectly, $y^{(i)}(\theta^{(k-1)T}x^{(i)}) \leq 0$. Thus,

$$\begin{aligned} \|\theta^{(k)}\|^2 &= \|\theta^{(k-1)} + y^{(i)}x^{(i)}\|^2 \\ &= \|\theta^{(k-1)}\|^2 + 2y^{(i)}\theta^{(k-1)T}x^{(i)} + \|x^{(i)}\|^2 \\ &\leq \|\theta^{(k-1)}\|^2 + R^2 \\ &\leq kR^2 \end{aligned}$$

where we have additionally applied the assumption from (b) and then again used simple induction.

Returning to the definition of the dot product, we have

$$\cos(\theta^{(k)}, \theta^*) = \frac{\theta^{(k)} \cdot \theta^*}{\|\theta^{(k)}\| \|\theta^*\|} = \left(\frac{\theta^{(k)} \cdot \theta^*}{\|\theta^*\|} \right) \frac{1}{\|\theta^{(k)}\|} \geq (k\gamma) \cdot \frac{1}{\sqrt{k}R} = \sqrt{k} \cdot \frac{\gamma}{R}$$

Since the value of the cosine is at most 1, we have

$$\begin{aligned} 1 &\geq \sqrt{k} \cdot \frac{\gamma}{R} \\ k &\leq \left(\frac{R}{\gamma} \right)^2. \end{aligned}$$

□


This result endows the margin γ of \mathcal{D}_n with an operational meaning: when using the Perceptron algorithm for classification, at most $(R/\gamma)^2$ classification errors will be made, where R is an upper bound on the magnitude of the training vectors.

CHAPTER 4

Feature representation


Linear classifiers are easy to work with and analyze, but they are a very restricted class of hypotheses. If we have to make a complex distinction in low dimensions, then they are unhelpful.

Our favorite illustrative example is the “exclusive or” (XOR) data set, the drosophila of machine-learning data sets:



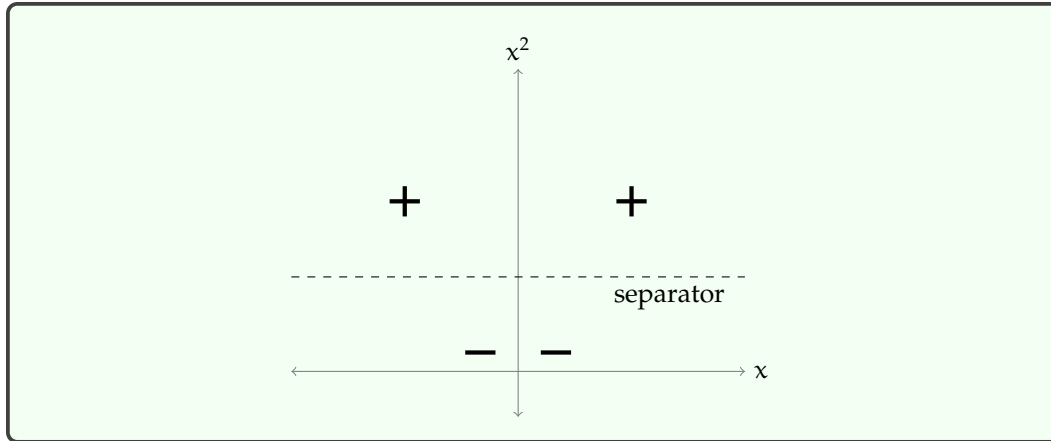
D. Melanogaster is a species of fruit fly, used as a simple system in which to study genetics, since 1910.

There is no linear separator for this two-dimensional dataset! But, we have a trick available: take a low-dimensional data set and move it, using a non-linear transformation into a higher-dimensional space, and look for a linear separator there. Let’s look at an example data set that starts in 1-D:




These points are not linearly separable, but consider the transformation $\phi(x) = [x, x^2]$. Putting the data in ϕ space, we see that it is now separable. There are lots of possible separators; we have just shown one of them here.

What’s a linear separator for data in 1D? A point!



A linear separator in ϕ space is a nonlinear separator in the original space! Let's see how this plays out in our simple example. Consider the separator $x^2 - 1 = 0$, which labels the half-plane $x^2 - 1 > 0$ as positive. What separator does it correspond to in the original 1-D space? We have to ask the question: which x values have the property that $x^2 - 1 = 0$. The answer is $+1$ and -1 , so those two points constitute our separator, back in the original space. And we can use the same reasoning to find the region of 1D space that is labeled positive by this separator.



This is a very general and widely useful strategy. It's the basis for *kernel methods*, a powerful technique that we won't study in this class, and can be seen as a motivation for multi-layer neural networks.

There are many different ways to construct ϕ . Some are relatively systematic and domain independent. We'll look at the *polynomial basis* in section 1 as an example of that. Others are directly related to the semantics (meaning) of the original features, and we construct them deliberately with our domain in mind. We'll explore that strategy in section 2.

1 Polynomial basis

If the features in your problem are already naturally numerical, one systematic strategy for constructing a new feature space is to use a *polynomial basis*. The idea is that, if you are using the k th-order basis (where k is a positive integer), you include a feature for every possible product of k different dimensions in your original input.

Here is a table illustrating the k th order polynomial basis for different values of k .

Order	$d = 1$	in general
0	[1]	[1]
1	[1, x]	[1, x_1, \dots, x_d]
2	[1, x, x^2]	[1, $x_1, \dots, x_d, x_1^2, x_1x_2, \dots$]
3	[1, x, x^2, x^3]	[1, $x_1, \dots, x_1^2, x_1x_2, \dots, x_1x_2x_3, \dots$]
\vdots	\vdots	\vdots

So, what if we try to solve the XOR problem using a polynomial basis as the feature transformation? We can just take our two-dimensional data and transform it into a higher-

dimensional data set, by applying ϕ . Now, we have a classification problem as usual, and we can use the perceptron algorithm to solve it.

Let's try it for $k = 2$ on our XOR problem. The feature transformation is


$$\phi((x_1, x_2)) = (1, x_1, x_2, x_1^2, x_1x_2, x_2^2) .$$

Study Question: If we use perceptron to train a classifier after performing this feature transformation, would we lose any expressive power if we let $\theta_0 = 0$ (i.e. trained without offset instead of with offset)?

After 4 iterations, perceptron finds a separator with coefficients $\theta = (0, 0, 0, 0, 4, 0)$ and $\theta_0 = 0$. This corresponds to

$$0 + 0x_1 + 0x_2 + 0x_1^2 + 4x_1x_2 + 0x_2^2 + 0 = 0$$


and is plotted below, with the shaded and un-shaded regions showing the classification results:



Study Question: Be sure you understand why this high-dimensional hyperplane is a separator, and how it corresponds to the figure.

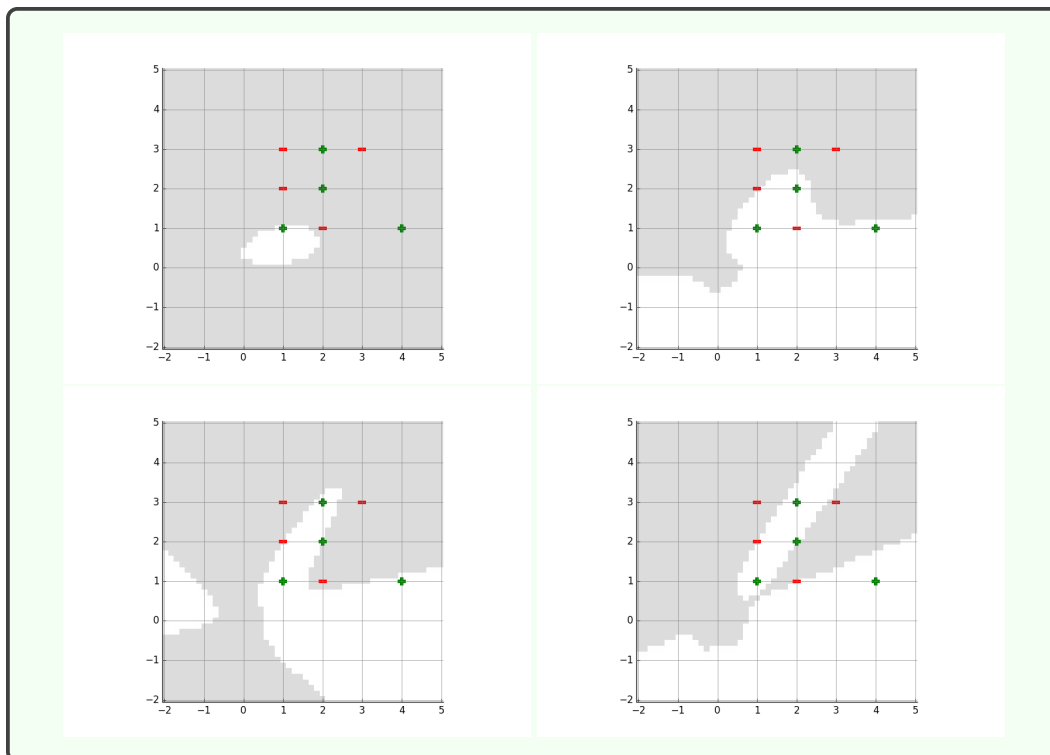
For fun, we show some more plots below. Here is the result of running perceptron on XOR, but where the data are put in a different place on the plane. After 65 mistakes (!) it arrives at these coefficients: $\theta = (1, -1, -1, -5, 11, -5)$, $\theta_0 = 1$, which generate this separator:

The jaggedness in the plotting of the separator is an artifact of a lazy lpk strategy for making these plots—the true curves are smooth.



Study Question: It takes many more iterations to solve this version. Apply knowledge of the convergence properties of the perceptron to understand why.

Here is a harder data set. After 200 iterations, we could not separate it with a second or third-order basis representation. Shown below are the results after 200 iterations for bases of order 2, 3, 4, and 5.



2 Hand-constructing features for real domains

In many machine-learning applications, we are given descriptions of the inputs with many different types of attributes, including numbers, words, and discrete features. An impor-

tant factor in the success of an ML application is the way that the features are chosen to be encoded by the human who is framing the learning problem.

2.1 Discrete features

Getting a good encoding of discrete features is particularly important. You want to create “opportunities” for the ML system to find the underlying regularities. Although there are machine-learning methods that have special mechanisms for handling discrete inputs, all the methods we consider in this class will assume the input vectors x are in \mathbb{R}^d . So, we have to figure out some reasonable strategies for turning discrete values into (vectors of) real numbers.

We’ll start by listing some encoding strategies, and then work through some examples. Let’s assume we have some feature in our raw data that can take on one of k discrete values.

- **Numeric** Assign each of these values a number, say $1.0/k, 2.0/k, \dots, 1.0$. We might want to then do some further processing, as described in section 8.3. This is a sensible strategy *only* when the discrete values really do signify some sort of numeric quantity, so that these numerical values are meaningful.
- **Thermometer code** If your discrete values have a natural ordering, from $1, \dots, k$, but not a natural mapping into real numbers, a good strategy is to use a vector of length k binary variables, where we convert discrete input value $0 < j \leq k$ into a vector in which the first j values are 1.0 and the rest are 0.0. This does not necessarily imply anything about the spacing or numerical quantities of the inputs, but does convey something about ordering.
- **Factored code** If your discrete values can sensibly be decomposed into two parts (say the “make” and “model” of a car), then it’s best to treat those as two separate features, and choose an appropriate encoding of each one from this list.
- **One-hot code** If there is no obvious numeric, ordering, or factorial structure, then the best strategy is to use a vector of length k , where we convert discrete input value $0 < j \leq k$ into a vector in which all values are 0.0, except for the j th, which is 1.0.
- **Binary code** It might be tempting for the computer scientists among us to use some binary code, which would let us represent k values using a vector of length $\log k$. *This is a bad idea!* Decoding a binary code takes a lot of work, and by encoding your inputs this way, you’d be forcing your system to *learn* the decoding algorithm.

As an example, imagine that we want to encode blood types, which are drawn from the set $\{A+, A-, B+, B-, AB+, AB-, O+, O-\}$. There is no obvious linear numeric scaling or even ordering to this set. But there is a reasonable *factoring*, into two features: $\{A, B, AB, O\}$ and $\{+, -1\}$. And, in fact, we can reasonably factor the first group into $\{A, \text{not}A\}, \{B, \text{not}B\}$. So, here are two plausible encodings of the whole set:

- Use a 6-D vector, with two dimensions to encode each of the factors using a one-hot encoding.
- Use a 3-D vector, with one dimension for each factor, encoding its presence as 1.0 and absence as -1.0 (this is sometimes better than 0.0). In this case, $AB+$ would be $(1.0, 1.0, 1.0)$ and $O-$ would be $(-1.0, -1.0, -1.0)$.

It is sensible (according to Wikipedia!) to treat O as having neither feature A nor feature B .

Study Question: How would you encode $A+$ in both of these approaches?

2.2 Text

The problem of taking a text (such as a tweet or a product review, or even this document!) and encoding it as an input for a machine-learning algorithm is interesting and complicated. Much later in the class, we'll study sequential input models, where, rather than having to encode a text as a fixed-length feature vector, we feed it into a hypothesis word by word (or even character by character!).

There are some simpler encodings that work well for basic applications. One of them is the *bag of words* (BOW) model. The idea is to let d be the number of words in our vocabulary (either computed from the training set or some other body of text or dictionary). We will then make a binary vector (with values 1.0 and 0.0) of length d , where element j has value 1.0 if word j occurs in the document, and 0.0 otherwise.

2.3 Numeric values

If some feature is already encoded as a numeric value (heart rate, stock price, distance, etc.) then you should generally keep it as a numeric value. An exception might be a situation in which you know there are natural "breakpoints" in the semantics: for example, encoding someone's age in the US, you might make an explicit distinction between under and over 18 (or 21), depending on what kind of thing you are trying to predict. It might make sense to divide into discrete bins (possibly spacing them closer together for the very young) and to use a one-hot encoding for some sorts of medical situations in which we don't expect a linear (or even monotonic) relationship between age and some physiological features.

If you choose to leave a feature as numeric, it is typically useful to *scale* it, so that it tends to be in the range $[-1, +1]$. Without performing this transformation, if you have one feature with much larger values than another, it will take the learning algorithm a lot of work to find parameters that can put them on an equal basis. So, we might perform transformation $\phi(x) = \frac{x - \bar{x}}{\sigma}$, where \bar{x} is the average of the $x^{(i)}$, and σ is the standard deviation of the $x^{(i)}$. The resulting feature values will have mean 0 and standard deviation 1. This transformation is sometimes called *standardizing* a variable.

Then, of course, you might apply a higher-order polynomial-basis transformation to one or more groups of numeric features.

Such standard variables are often known as "z-scores," for example, in the social sciences.

Study Question: Percy Eptron has a domain with 4 numeric input features, (x_1, \dots, x_4) . He decides to use a representation of the form

$$\phi(x) = \text{PolyBasis}((x_1, x_2), 3) \frown \text{PolyBasis}((x_3, x_4), 3)$$

where $a \frown b$ means the vector a concatenated with the vector b . What is the dimension of Percy's representation? Under what assumptions about the original features is this a reasonable choice?

1 Machine learning as optimization

The perceptron algorithm was originally written down directly via cleverness and intuition, and later analyzed theoretically. Another approach to designing machine learning algorithms is to frame them as optimization problems, and then use standard optimization algorithms and implementations to actually find the hypothesis.

We begin by writing down an *objective function* $J(\theta)$, where θ stands for *all* the parameters in our model. We also often write $J(\theta; \mathcal{D})$ to make clear the dependence on the data \mathcal{D} . The objective function describes how we feel about possible hypotheses θ : we will generally look for values for parameters θ that minimize the objective function:

$$\theta^* = \arg \min_{\theta} J(\theta) .$$

You can think about θ^* here as “the theta that minimizes J ”.

A very common form for an ML objective is

$$J(\theta) = \left(\frac{1}{n} \sum_{i=1}^n \underbrace{L(x^{(i)}, y^{(i)}, \theta)}_{\text{loss}} \right) + \underbrace{\lambda}_{\text{constant}} \underbrace{R(\theta)}_{\text{regularizer}} .$$

The *loss* tells us how unhappy we are about the prediction $h(x^{(i)}; \theta)$ that θ makes for $(x^{(i)}, y^{(i)})$. A common example is the 0-1 loss, introduced in chapter 1:

$$L_{01}(x, y, \theta) = \begin{cases} 0 & \text{if } y = h(x; \theta) \\ 1 & \text{otherwise} \end{cases}$$

which gives a value of 0 for a correct prediction, and a 1 for an incorrect prediction. In the case of linear separators, this becomes:

$$L_{01}(x, y, \theta, \theta_0) = \begin{cases} 0 & \text{if } y(\theta^T x + \theta_0) > 0 \\ 1 & \text{otherwise} \end{cases}$$

We will sometimes write $J(\theta, \theta_0)$ because when studying linear classifiers, we have used these two names for our whole collection of parameters.


2 Regularization

If all we cared about was finding a hypothesis with small loss on the training data, we would have no need for regularization, and could simply omit the second term in the ob-

jective. But remember that our objective is to *perform well on input values that we haven't trained on!* It may seem that this is an impossible task, but humans and machine-learning methods do this successfully all the time. What allows *generalization* to new input values is a belief that there is an underlying regularity that governs both the training and testing data. We have already discussed one way to describe an assumption about such a regularity, which is by choosing a limited class of possible hypotheses. Another way to do this is to provide smoother guidance, saying that, within a hypothesis class, we prefer some hypotheses to others. The regularizer articulates this preference and the constant λ says how much we are willing to trade off loss on the training data versus preference over hypotheses.

This trade-off is illustrated in the figure below. Hypothesis h_1 has 0 loss, but is very complicated. Hypothesis h_2 mis-classifies two points, but is very simple. In absence of other beliefs about the solution, it is often better to prefer that the solution be “simpler”, and so we might prefer h_2 over h_1 , expecting it to perform better on future examples drawn from this same distribution. Another nice way of thinking about regularization is that we would like to prevent our hypothesis from being too dependent on the particular training data that we were given: we would like for it to be the case that if the training data were changed slightly, the hypothesis would not change.

To establish some vocabulary, we might say that h_1 is *overfit* to the training data.



A common strategy for specifying a *regularizer* is to use the form

$$R(\theta) = \|\theta - \theta_{\text{prior}}\|^2$$

when you have some idea in advance that θ ought to be near some value θ_{prior} . In the absence of such knowledge a default is to *regularize toward zero*:

$$R(\theta) = \|\theta\|^2 .$$

Learn about Bayesian methods in machine learning to see the theory behind this and cool results!

3 Maximizing the margin

One criterion to use for judging separators that classify all examples correctly (that is, that have a 0-1 loss value of 0) is to prefer ones that have a *large margin*. Recall that the margin


of a labeled point (x, y) with respect to the hyperplane θ, θ_0 is

$$\gamma(x, y, \theta, \theta_0) = \frac{y (\theta^T x + \theta_0)}{\|\theta\|} .$$

The margin is positive if a point is classified correctly, and the absolute value of the margin is the perpendicular distance of x to the hyperplane. The margin of a dataset \mathcal{D} with respect to θ, θ_0 is

$$\min_{(x^{(i)}, y^{(i)}) \in \mathcal{D}} \gamma(x^{(i)}, y^{(i)}, \theta, \theta_0) .$$

The figure below illustrates the margin of two linear separators, each with 0 loss, with respect to the same data set. The separator, h_1 , with the larger margin γ_1 feels intuitively like it will have better generalization performance. Separator h_2 has a very small margin and would make errors if the data were perturbed by a small amount.



There is a lot of interesting theoretical work justifying a preference for separators with large margin. There is an important class of algorithms called *support vector machines* that focus on finding large-margin separators with some other interesting properties. Before the current boom in neural networks, they were the “go-to” algorithms for classification. We don’t have time to study them in more detail in this class, but we recommend that you read about them some time.

If our data is linearly separable, we can frame the problem of finding the maximum-margin separator as

$$\theta^*, \theta_0^* = \arg \max_{\theta, \theta_0} \min_i \gamma(x^{(i)}, y^{(i)}, \theta, \theta_0) ,$$

which we can find by minimizing the objective

$$J(\theta, \theta_0) = - \min_i \gamma(x^{(i)}, y^{(i)}, \theta, \theta_0) .$$

However, this form of the objective can be tricky to optimize, and is not useful when the data are non-separable.

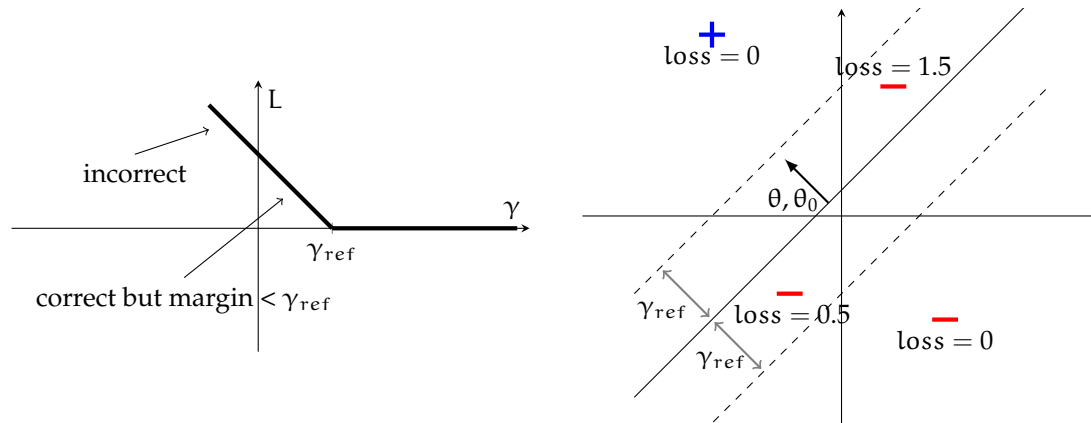
We will develop another possible formulation. Let’s start by assuming that we know a good value for the target margin, $\gamma_{ref} > 0$. We would like

- All points to have margin $> \gamma_{ref}$ and
- the target margin γ_{ref} to be big.

To make this more precise, we start by defining a new loss function called the *hinge loss*:

$$L_h(\gamma/\gamma_{ref}) = \begin{cases} 1 - \gamma/\gamma_{ref} & \text{if } \gamma < \gamma_{ref} \\ 0 & \text{otherwise} \end{cases} .$$

The figure below and to the left plots the hinge loss as a function of the margin, γ , of a point. It is zero if the margin is greater than γ_{ref} and increases linearly as the margin decreases.



Study Question: Plot 0-1 loss on the axes of the figure on the left.

In the figure on the right, we illustrate a hyperplane θ, θ_0 . Parallel to that hyperplane, but offset in either direction by γ_{ref} are two more hyperplanes (dotted lines in the figure) representing the margins. Any correctly classified point outside the margin has loss 0. Any correctly classified point inside the margin has loss in $(0, 1)$. Any incorrectly classified point has loss ≥ 1 .

Study Question: Be sure you can compute the loss values shown on this figure.

Also, try to visualize a plot in 3 dimensions that shows the loss function for positive points and the loss function for negative points going in the Z dimension (up out of the page) on top of this plot.

Now, let's formulate an objective function, in terms of the parameters $(\theta, \theta_0, \gamma_{ref})$. We will express our desire for a large margin by employing a regularization term in the objective function, namely $R(\theta, \theta_0, \gamma_{ref}) = \frac{1}{\gamma_{ref}^2}$. Then, we have

$$J(\theta, \theta_0, \gamma_{ref}) = \frac{1}{n} \sum_{i=1}^n L_h \left(\frac{\gamma(x^{(i)}, y^{(i)}, \theta, \theta_0)}{\gamma_{ref}} \right) + \lambda \left(\frac{1}{\gamma_{ref}^2} \right) .$$

We see that the two terms in the objective function have opposing effects, favoring small and large γ_{ref} , respectively, with λ governing the trade-off.

Study Question: You should, by now, be asking yourself "How do we pick λ ?" You can think of the different objective functions that result from different choices of hyperparameter λ as different *learning algorithms*. What do you know about choosing among algorithms? How would you use that to pick λ ?

Now, in order to slightly simplify our problem and to connect up with more standard existing methods, we are going to make a somewhat unintuitive step. We actually have one more parameter in our set of parameters $(\theta, \theta_0, \gamma_{ref})$ than is really necessary for describing this objective, so we are going to rewrite it.

Remember that any linear scaling of θ, θ_0 represents the same separator. So, without losing any ability to represent separators or describe their margins, we can scale θ, θ_0 so that

$$\|\theta\| = \frac{1}{\gamma_{ref}} .$$

Note that, because our target margin scales inversely with $\|\theta\|$, wanting the margin to be large is the same as wanting $\|\theta\|$ to be small.

If you are thinking to yourself that we could express the desire for large margin by setting $R = -\gamma_{ref}$ or $R = 1/\gamma_{ref}$ or any of a variety of other things, you would be right. We'll insist on this choice for now, because it will turn out to have useful consequences later, but you have no way of seeing that at this point.

Because it's not a parameter of the hypothesis, but rather a parameter of the method we are using to choose the hypothesis, we call λ a *hyperparameter*.

This is kind of tricky. Stew about it until it makes sense to you.

It's like we're secretly encoding our target margin as the inverse of the norm of θ .

With this trick, we don't need γ_{ref} as a parameter any more, and we get the following objective, which is the objective for support vector machines, and we'll often refer to as the SVM objective:

$$J(\theta, \theta_0) = \left(\frac{1}{n} \sum_{i=1}^n L_h \left(\mathbf{y}^{(i)} (\theta^T \mathbf{x} + \theta_0) \right) \right) + \lambda \|\theta\|^2 .$$

Here are some observations:

1. If $\lambda = 0$, for any separator that correctly classifies all the points, θ, θ_0 can always be chosen so that the objective function evaluates to 0.
2. If $\lambda > 0$ but is very small, we will pick θ with smallest $\|\theta\|$ while still maintaining separation of the data.
3. If λ is large, we tolerate errors in favor of having a "simpler" (smaller norm) separator.

Study Question: Be sure you can make a detailed explanation of each of these points. In point 1 above, would we need increase or decrease the magnitude of θ to make the objective go to zero? How does point 3 relate to the idea of regularizing toward zero?

At the optimum, for separable data with very small λ :

- $\mathbf{y}^{(i)} (\theta^T \mathbf{x}^{(i)} + \theta_0) \geq 1$ for all i , since the hinge loss evaluates to 0.
- $\mathbf{y}^{(i)} (\theta^T \mathbf{x}^{(i)} + \theta_0) = 1$ for at least one i , because the regularizer will drive $\|\theta\|$ to be as small as possible with the loss still remaining at 0.

Points for which $\mathbf{y}^{(i)} (\theta^T \mathbf{x}^{(i)} + \theta_0) = 1$ have margin exactly equal to γ_{ref} (otherwise we could decrease $\|\theta\|$ to obtain a larger margin). For these points that lie along the margin, we use their distance to θ, θ_0 to compute the margin of this separator with respect to the data set:

$$\frac{\mathbf{y}^{(i)} (\theta^T \mathbf{x}^{(i)} + \theta_0)}{\|\theta\|} = \text{margin} ,$$

and so

$$\text{margin} = \frac{1}{\|\theta\|} .$$

Study Question: Be sure you can derive this last step!

Note that this last assertion (that the margin is the inverse of the norm of θ) is only true under the assumptions listed at the top of this paragraph: when the data is separable, when λ is very small, and when θ is the optimum of the SVM objective.

Gradient Descent

Now we have shown how to describe an interesting objective function for machine learning, but we need a way to find the optimal $\theta^* = \arg \min_{\theta} J(\theta)$? There is an enormous, fascinating, literature on the mathematical and algorithmic foundations of optimization, but for this class, we will consider one of the simplest methods, called *gradient descent*.

Intuitively, in one or two dimensions, we can easily think of $J(\theta)$ as defining a surface over θ ; that same idea extends to higher dimensions. Now, our objective is to find the θ value at the lowest point on that surface. One way to think about gradient descent is that you start at some arbitrary point on the surface, look to see in which direction the “hill” goes down most steeply, take a small step in that direction, determine the direction of steepest descent from where you are, take another small step, etc.

Which you should consider studying some day!

1 One dimension

We will start by considering gradient descent in one dimension. Assume $\theta \in \mathbb{R}$, and that we know both $J(\theta)$ and its first derivative with respect to θ , $J'(\theta)$. Here is pseudocode for gradient descent on an arbitrary function f . Along with f and f' , we have to specify the initial value for parameter θ , a *step-size* parameter η , and an *accuracy* parameter ϵ :

1D-GRADIENT-DESCENT($\theta_{init}, \eta, f, f', \epsilon$)

```

1  $\theta^{(0)} = \theta_{init}$ 
2  $t = 0$ 
3 repeat
4    $t = t + 1$ 
5    $\theta^{(t)} = \theta^{(t-1)} - \eta f'(\theta^{(t-1)})$ 
6 until  $|f'(\theta^{(t)})| < \epsilon$ 
7 return  $\theta^{(t)}$ 
```


Note that there are many other reasonable ways to decide to terminate, including when $|\theta^{(t)} - \theta^{(t-1)}| < \epsilon$ or when $|f(\theta^{(t)}) - f(\theta^{(t-1)})| < \epsilon$.

If J is convex, for any desired accuracy ϵ , there is some step size η such that gradient descent will converge to within ϵ of the optimal θ . However, we must be careful when choosing the step size to prevent slow convergence, oscillation around the minimum, or divergence.

Here's a very old-school humorous description of gradient descent and other optimization algorithms using analogies involving kangaroos:
<ftp://ftp.sas.com/pub/neural/kangaroos.txt>


Woo hoo! We have a convergence guarantee, of sorts

The following plot illustrates a convex function $f(x) = (x-2)^2$, starting gradient descent at $\theta_{\text{init}} = 4.0$ with a step-size of $1/2$. It is very well-behaved!



Study Question: What happens in this example with very small η ? With very big η ?

If J is non-convex, where gradient descent converges to depends on θ_{init} . When it reaches a value of θ where $f'(\theta) = 0$ and $f''(\theta) > 0$, but it is not a minimum of the function, it is called a *local minimum* or *local optimum*.



2 Multiple dimensions

The extension to the case of multi-dimensional θ is straightforward. Let's assume $\theta \in \mathbb{R}^m$, so $J : \mathbb{R}^m \rightarrow \mathbb{R}$. The gradient of J with respect to θ is

$$\nabla_{\theta} J = \begin{bmatrix} \partial J / \partial \theta_1 \\ \vdots \\ \partial J / \partial \theta_m \end{bmatrix}$$

The algorithm remains the same, except that the update step in line 5 becomes

$$\theta^{(t)} = \theta^{(t-1)} - \eta \nabla_{\theta} J$$

and we have to change the termination criterion. The easiest thing is to replace the test in line 6 with $|f(\theta^{(t)}) - f(\theta^{(t-1)})| < \epsilon$, which is sensible no matter the dimensionality of θ .

3 Application to SVM objective

There are two slight “wrinkles” involved in applying gradient descent to the SVM objective.

We begin by stating the objective and the gradient necessary for doing gradient descent. In our problem, the entire parameter vector is described by parameter vector θ and scalar θ_0 and so we will have to adjust them both and compute gradients of J with respect to each of them. The objective and gradient (note we have replaced the constant λ with $\frac{\lambda}{2}$ for convenience), are

$$J(\theta, \theta_0) = \frac{1}{n} \sum_{i=1}^n L_h \left(\mathbf{y}^{(i)} (\theta^T \mathbf{x}^{(i)} + \theta_0) \right) + \frac{\lambda}{2} \|\theta\|^2$$

$$\nabla_{\theta} J = \frac{1}{n} \sum_{i=1}^n L'_h \left(\mathbf{y}^{(i)} (\theta^T \mathbf{x}^{(i)} + \theta_0) \right) \mathbf{y}^{(i)} \mathbf{x}^{(i)} + \lambda \theta$$

The following step requires passing familiarity with matrix derivatives. A foolproof way of computing them is to compute partial derivative of J with respect to each component θ_i of θ .

Study Question: Convince yourself that the dimensions of all these quantities are correct, under the assumption that θ is $d \times 1$.

Recall the hinge-loss

$$L_h(v) = \begin{cases} 1 - v & \text{if } v < 1 \\ 0 & \text{otherwise} \end{cases} .$$

This loss is not differentiable, since the derivative at $v = 1$ doesn't exist! So we consider the subgradient

$$L'_h(v) = \begin{cases} -1 & \text{if } v < 1 \\ 0 & \text{otherwise} \end{cases} .$$

You don't have to really understand the idea of a subgradient, just that it has value 0 at $v = 1$ here.

This gives us a complete definition of $\nabla_{\theta} J$. But we also have to go down the gradient with respect to θ_0 , so we find

$$\frac{\partial J}{\partial \theta_0} = \frac{1}{n} \sum_{i=1}^n L'_h \left(\mathbf{y}^{(i)} (\theta^T \mathbf{x} + \theta_0) \right) \mathbf{y}^{(i)} .$$

Finally, our gradient descent algorithm becomes

SVM-GRADIENT-DESCENT($\theta_{init}, \theta_{0init}, \eta, J$)

```

1   $\theta^{(0)} = \theta_{init}$ 
2   $\theta_0^{(0)} = \theta_{0init}$ 
3   $t = 0$ 
4  repeat
5       $t = t + 1$ 
6       $\theta^{(t)} = \theta^{(t-1)} - \eta \left( \frac{1}{n} \sum_{i=1}^n \begin{cases} -1 & \text{if } \mathbf{y}^{(i)} (\theta^{(t-1)T} \mathbf{x}^{(i)} + \theta_0^{(t-1)}) < 1 \\ 0 & \text{otherwise} \end{cases} \mathbf{y}^{(i)} \mathbf{x}^{(i)} + \lambda \theta^{(t-1)} \right)$ 
7       $\theta_0^{(t)} = \theta_0^{(t-1)} - \eta \left( \frac{1}{n} \sum_{i=1}^n \begin{cases} -1 & \text{if } \mathbf{y}^{(i)} (\theta^{(t-1)T} \mathbf{x}^{(i)} + \theta_0^{(t-1)}) < 1 \\ 0 & \text{otherwise} \end{cases} \mathbf{y}^{(i)} \right)$ 
8  until  $|J(\theta^{(t)}, \theta_0^{(t)}) - J(\theta^{(t-1)}, \theta_0^{(t-1)})| < \epsilon$ 
9  return  $\theta^{(t)}, \theta_0^{(t)}$ 

```

Study Question: Is it okay that λ doesn't appear in line 7?

4 Stochastic Gradient Descent

When the form of the gradient is a sum, rather than take one big(ish) step in the direction of the gradient, we can, instead, randomly select one term of the sum, and take a very small step in that direction. This seems sort of crazy, but remember that all the little steps would average out to the same direction as the big step if you were to stay in one place. Of course, you're not staying in that place, so you move, in expectation, in the direction of the gradient.

The word "stochastic" means probabilistic, or random; so does "aleatoric," which is a very cool word. Look up aleatoric music, sometime.

Most objective functions in machine learning can end up being written as a sum over data points, in which case, stochastic gradient descent (SGD) is implemented by picking a data point randomly out of the data set, computing the gradient as if there were only that one point in the data set, and taking a small step in the negative direction.

Here is pseudocode for applying SGD to ridge regression:

SGD FOR REGRESSION

```

1   $\theta^{(0)} = 0$ 
2  for  $t = 1$  to  $T$ 
3      randomly select  $i \in \{1, 2, \dots, n\}$ 
4       $\theta^{(t+1)} = \theta^{(t)} - \eta_t \left( \left( \theta^{(t)T} x^{(i)} - y^{(i)} \right) x^{(i)} + \frac{\lambda}{n} \theta^{(t)} \right)$ 
5  return  $\theta$ 

```

Note that now instead of a fixed value of η , it is indexed by the iteration of the algorithm, t . For SGD to converge to a local optimum as t increases, the step size has to decrease as a function of time. To guarantee convergence, the following relationships must hold:

$$\sum_{t=1}^{\infty} \eta_t = \infty \quad \text{and} \quad \sum_{t=1}^{\infty} \eta_t^2 < \infty .$$

One "legal" way of setting the step size is to make $\eta_t = 1/t$ but people often use rules that decrease more slowly, and so don't strictly satisfy the criteria for convergence.

Study Question: If you start a long way from the optimum, would making η_t decrease more slowly tend to make you move more quickly or more slowly to the optimum?

It can be interesting to re-organize the terms in the gradient update (line 4 of the pseudocode) in this way:

$$\theta^{(t+1)} = \underbrace{\left(1 - \frac{\lambda \eta_t}{n} \right)}_{\lambda \text{ pressures } \theta \text{ to be small}} \theta^{(t)} + \eta_t \underbrace{\left(y^{(i)} - \theta^{(t)T} x^{(i)} \right)}_{\text{error}} x^{(i)}$$

The first term will tend to move θ toward 0, and the second will tend to reduce the error on the current data point.

CHAPTER 7

Regression

Now we will turn to a slightly different form of machine-learning problem, called *regression*. It is still supervised learning, so our data will still have the form

$$S_n = \left\{ \left(x^{(1)}, y^{(1)} \right), \dots, \left(x^{(n)}, y^{(n)} \right) \right\} .$$

But now, instead of the y values being discrete, they will be real-valued, and so our hypotheses will have the form

$$h : \mathbb{R}^d \rightarrow \mathbb{R} .$$

This is a good framework when we want to predict a numerical quantity, like height, stock value, etc.

The first step is to pick a loss function, to describe how to evaluate the quality of the predictions our hypothesis is making, when compared to the “target” y values in the data set. The choice of loss function is part of modeling your domain. In the absence of additional information about a regression problem, we typically use *squared error* (SE):

$$\text{Loss}(\text{guess}, \text{actual}) = (\text{guess} - \text{actual})^2 .$$

It penalizes guesses that are too high the same amount as guesses that are too low, and has a good mathematical justification in the case that your data are generated from an underlying linear hypothesis, but with Gaussian-distributed noise added to the y values.

We will consider the case of a linear hypothesis class,

$$h(x; \theta, \theta_0) = \theta^\top x + \theta_0 ,$$

remembering that we can get a rich class of hypotheses by performing a non-linear feature transformation before doing the regression. So, $\theta^\top x + \theta_0$ is a linear function of x , but $\theta^\top \varphi(x) + \theta_0$ is a non-linear function of x if φ is a non-linear function of x .

We will treat regression as an optimization problem, in which, given a data set \mathcal{D} , we wish to find a linear hypothesis that minimizes mean square error. Our objective, often called *mean squared error* is to find values for θ, θ_0 that minimize

$$J(\theta, \theta_0) = \frac{1}{n} \sum_{i=1}^n \left(\theta^\top x^{(i)} + \theta_0 - y_i \right)^2 ,$$

making the solution be

$$\theta^*, \theta_0^* = \arg \min_{\theta, \theta_0} J(\theta, \theta_0) . \tag{7.1}$$

“Regression,” in common parlance, means moving backwards. But this is forward progress!

1 Analytical solution: ordinary least squares

One very interesting aspect of the problem finding a linear hypothesis that minimizes mean squared error (this general problem is often called *ordinary least squares* (OLS) is that we can find a closed-form formula for the answer!

Everything is easier to deal with if we assume that the $x^{(i)}$ have been augmented with an extra input dimension (feature) that always has value 1, so we may ignore θ_0 . (See chapter 3, section 2 for a reminder about this strategy).

We will approach this just like a minimization problem from calculus homework: take the derivative of J with respect to θ , set it to zero, and solve for θ . There is an additional step required, to check that the resulting θ is a minimum (rather than a maximum or an inflection point) but we won't work through that here. It is possible to approach this problem by:

- Finding $\partial J / \partial \theta_k$ for k in $1, \dots, d$,
- Construct a set of k equations of the form $\partial J / \partial \theta_k = 0$, and
- Solving the system for values of θ_k .

That works just fine. To get practice for applying techniques like this to more complex problems, we will work through a more compact (and cool!) matrix view.

Study Question: Work through this and check your answer against ours below.

We can think of our training data in terms of matrices X and Y , where each column of X is an example, and each "column" of Y is the corresponding label:

$$X = \begin{bmatrix} x_1^{(1)} & \dots & x_1^{(n)} \\ \vdots & \ddots & \vdots \\ x_d^{(1)} & \dots & x_d^{(n)} \end{bmatrix} \quad Y = \begin{bmatrix} y^{(1)} & \dots & y^{(n)} \end{bmatrix} .$$

In most textbooks, they think of an individual example $x^{(i)}$ as a row, rather than a column. So that we get an answer that will be recognizable to you, we are going to define a new matrix and vector, W and T , which are just transposes of our X and Y , and then work with them:

$$W = X^T = \begin{bmatrix} x_1^{(1)} & \dots & x_d^{(1)} \\ \vdots & \ddots & \vdots \\ x_1^{(n)} & \dots & x_d^{(n)} \end{bmatrix} \quad T = Y^T = \begin{bmatrix} y^{(1)} \\ \vdots \\ y^{(n)} \end{bmatrix} .$$

Now we can write

$$J(\theta) = \frac{1}{n} \underbrace{(W\theta - T)^T}_{1 \times n} \underbrace{(W\theta - T)}_{n \times 1}$$

and using facts about matrix/vector calculus, we get

$$\nabla_{\theta} J = \frac{2}{n} \underbrace{W^T}_{d \times n} \underbrace{(W\theta - T)}_{n \times 1} .$$

What does "closed form" mean? Generally, that it involves direct evaluation of a mathematical expression using a fixed number of "typical" operations (like arithmetic operations, trig functions, powers, etc.). So equation 7.1 is not in closed form, because it's not at all clear what operations one needs to perform to find the solution.

We will use d here for the total number of features in each $x^{(i)}$, including the added 1.

Setting to 0 and solving, we get:

$$\begin{aligned}\frac{2}{n}W^T(W\theta - T) &= 0 \\ W^TW\theta - W^T T &= 0 \\ W^TW\theta &= W^T T \\ \theta &= (W^TW)^{-1}W^T T\end{aligned}$$

And the dimensions work out!

$$\theta = \underbrace{(W^TW)^{-1}}_{d \times d} \underbrace{W^T}_{d \times n} \underbrace{T}_{n \times 1}$$

So, given our data, we can directly compute the linear regression that minimizes mean squared error. That's pretty awesome!

2 Regularization

Well, actually, there are some kinds of trouble we can get into. What if (W^TW) is not invertible?

Study Question: Consider, for example, a situation where the data-set is just the same point repeated twice: $x^{(1)} = x^{(2)} = (1, 2)^T$. What is W in this case? What is W^TW ? What is $(W^TW)^{-1}$?

Another kind of problem is *overfitting*: we have formulated an objective that is just about fitting the data as well as possible, but as we discussed in the context of margin maximization, we might also want to *regularize* to keep the hypothesis from getting *too* attached to the data.

We address both the problem of not being able to invert $(W^TW)^{-1}$ and the problem of overfitting using a mechanism called *ridge regression*. We add a regularization term $\|\theta\|^2$ to the OLS objective, with trade-off parameter λ .

Study Question: When we add a regularizer of the form $\|\theta\|^2$, what is our most "preferred" value of θ , in the absence of any data?

$$J_{\text{ridge}}(\theta, \theta_0) = \frac{1}{n} \sum_{i=1}^n \left(\theta^T x^{(i)} + \theta_0 - y^{(i)} \right)^2 + \lambda \|\theta\|^2$$

Larger λ values pressure θ values to be near zero. Note that we don't penalize θ_0 ; intuitively, θ_0 is what "floats" the regression surface to the right level for the data you have, and so you shouldn't make it harder to fit a data set where the y values tend to be around one million than one where they tend to be around one. The other parameters control the orientation of the regression surface, and we prefer it to have a not-too-crazy orientation.

There is an analytical expression for the θ, θ_0 values that minimize J_{ridge} , but it's a little bit more complicated to derive than the solution for OLS because θ_0 needs special treatment. If we decide not to treat θ_0 specially (so we add a 1 feature to our input vectors), then we get:

$$\nabla_{\theta} J_{\text{ridge}} = \frac{2}{n} W^T(W\theta - T) + 2\lambda\theta .$$

Setting to 0 and solving, we get:

$$\begin{aligned}\frac{2}{n}W^T(W\theta - T) + 2\lambda\theta &= 0 \\ \frac{1}{n}W^TW\theta - \frac{1}{n}W^TT + \lambda\theta &= 0 \\ \frac{1}{n}W^TW\theta + \lambda\theta &= \frac{1}{n}W^TT \\ W^TW\theta + n\lambda\theta &= W^TT \\ (W^TW + n\lambda I)\theta &= W^TT \\ \theta &= (W^TW + n\lambda I)^{-1}W^TT\end{aligned}$$

Whew! So,

$$\theta_{\text{ridge}} = (W^TW + n\lambda I)^{-1}W^TT$$

which becomes invertible when $\lambda > 0$.

Study Question: Derive the ridge regression formula.

This is called “ridge” regression because we are adding a “ridge” of λ values along the diagonal of the matrix before inverting it.

Talking about regularization In machine learning in general, not just regression, it is useful to distinguish two ways in which a hypothesis $h \in \mathcal{H}$ might engender errors on test data. We have

Structural error: This is error that arises because there is no hypothesis $h \in \mathcal{H}$ that will perform well on the data, for example because the data was really generated by a sin wave but we are trying to fit it with a line.

Estimation error: This is error that arises because we do not have enough data (or the data are in some way unhelpful) to allow us to choose a good $h \in \mathcal{H}$.

When we increase λ , we tend to increase structural error but decrease estimation error, and vice versa.

Study Question: Consider using a polynomial basis of order k as a feature transformation ϕ on your data. Would increasing k tend to increase or decrease structural? What about estimation error?

There are technical definitions of these concepts that are studied in more advanced treatments of machine learning. Structural error is referred to as *bias* and estimation error is referred to as *variance*.

3 Optimization via gradient descent

Inverting the $d \times d$ matrix (W^TW) takes $O(d^3)$ time, which makes the analytic solution impractical for large d . If we have high-dimensional data, we can fall back on gradient descent.

Study Question: Why is having large n not as much of a computational problem as having large d ?

Recall the ridge objective

$$J_{\text{ridge}}(\theta, \theta_0) = \frac{1}{n} \sum_{i=1}^n (\theta^T x^{(i)} + \theta_0 - y^{(i)})^2 + \lambda \|\theta\|^2$$

and its gradient with respect to θ

$$\nabla_{\theta} J = \frac{2}{n} \sum_{i=1}^n (\theta^T x^{(i)} + \theta_0 - y^{(i)}) x^{(i)} + 2\lambda\theta$$

Well, actually, Gauss-Jordan elimination, a popular algorithm, takes $O(d^3)$ arithmetic operations, but the bit complexity of the intermediate results can grow exponentially! There are other algorithms with polynomial bit complexity. (If this just made no sense to you, don't worry.)

and partial derivative with respect to θ_0

$$\frac{\partial J}{\partial \theta_0} = \frac{2}{n} \sum_{i=1}^n \left(\theta^T x^{(i)} + \theta_0 - y^{(i)} \right) .$$

Armed with these derivatives, we can do gradient descent, using the regular or stochastic gradient methods from chapter 6.

Even better, the objective functions for OLS and ridge regression are *convex*, which means they have only one minimum, which means, with a small enough step size, gradient descent is *guaranteed* to find the optimum.

CHAPTER 8

Neural Networks

Unless you live under a rock with no internet access, you've been hearing a lot about "neural networks." Now that we have several useful machine-learning concepts (hypothesis classes, classification, regression, gradient descent, regularization, etc.) we are completely well equipped to understand neural networks in detail.

This, in some sense, the "third wave" of neural nets. The basic idea is founded on the 1943 model of neurons of McCulloch and Pitts and learning ideas of Hebb. There was a great deal of excitement, but not a lot of practical success: there were good training methods (e.g., perceptron) for linear functions, and interesting examples of non-linear functions, but no good way to train non-linear functions from data. Interest died out for a while, but was re-kindled in the 1980s when several people came up with a way to train neural networks with "back-propagation," which is a particular style of implementing gradient descent, which we will study here. By the mid-90s, the enthusiasm waned again, because although we could train non-linear networks, the training tended to be slow and was plagued by a problem of getting stuck in local optima. Support vector machines (SVMs) (regularization of high-dimensional hypotheses by seeking to maximize the margin) and kernel methods (an efficient and beautiful way of using feature transformations to non-linearly transform data into a higher-dimensional space) provided reliable learning methods with guaranteed convergence and no local optima.

However, during the SVM enthusiasm, several groups kept working on neural networks, and their work, in combination with an increase in available data and computation, has made them rise again. They have become much more reliable and capable, and are now the method of choice in many applications. There are many, many variations of neural networks, which we can't even begin to survey. We will study the core "feed-forward" networks with "back-propagation" training, and then, in later chapters, address some of the major advances beyond this core.

We can view neural networks from several different perspectives:

View 1: An application of stochastic gradient descent for classification and regression with a potentially very rich hypothesis class.

View 2: A brain-inspired network of neuron-like computing elements that learn distributed representations.

View 3: A method for building applications that make predictions based on huge amounts

As with many good ideas in science, the basic idea for how to train non-linear neural networks with gradient descent, was independently developed by more than one researcher.

The number increases daily, as may be seen on arxiv.org.


of data in very complex domains.

We will mostly take view 1, with the understanding that the techniques we develop will enable the applications in view 3. View 2 was a major motivation for the early development of neural networks, but the techniques we will study do not seem to actually account for the biological learning processes in brains.

Some prominent researchers are, in fact, working hard to find analogues of these methods in the brain

1 Basic element

The basic element of a neural network is a “neuron,” pictured schematically below. We will also sometimes refer to a neuron as a “unit” or “node.”



It is a non-linear function of an input vector $x \in \mathbb{R}^m$ to a single output value $a \in \mathbb{R}$. It is parameterized by a vector of *weights* $(w_1, \dots, w_m) \in \mathbb{R}^m$ and an *offset or threshold* $w_0 \in \mathbb{R}$. In order for the neuron to be non-linear, we also specify an *activation function* $f : \mathbb{R} \rightarrow \mathbb{R}$, which can be the identity ($f(x) = x$), but can also be any other function, though we will only be able to work with it if it is differentiable.

The function represented by the neuron is expressed as:

$$a = f(z) = f\left(\sum_{j=1}^m x_j w_j + w_0\right) = f(w^T x + w_0) .$$

Before thinking about a whole network, we can consider how to train a single unit. Given a loss function $L(\text{guess}, \text{actual})$ and a dataset $\{(x^{(1)}, y^{(1)}), \dots, (x^{(n)}, y^{(n)})\}$, we can do (stochastic) gradient descent, adjusting the weights w, w_0 to minimize

$$\sum_i L\left(\text{NN}(x^{(i)}; w, w_0), y^{(i)}\right) .$$

where NN is the output of our neural net for a given input.

We have already studied two special cases of the neuron: linear classifiers with hinge loss and regressors with quadratic loss! Both of these have activation functions $f(x) = x$.

Study Question: Just for a single neuron, imagine for some reason, that we decide to use activation function $f(z) = e^z$ and loss function $L(g, a) = (g - a)^2$. Derive a gradient descent update for w and w_0 .

Sorry for changing our notation here. We were using d as the dimension of the input, but we are trying to be consistent here with many other accounts of neural networks. It is impossible to be consistent with all of them though—there are many different ways of telling this story.

This should remind you of our θ and θ_0 for linear models.

2 Networks

Now, we'll put multiple neurons together into a *network*. A neural network in general takes in an input $x \in \mathbb{R}^m$ and generates an output $a \in \mathbb{R}^n$. It is constructed out of multiple


neurons; the inputs of each neuron might be elements of x and/or outputs of other neurons. The outputs are generated by n *output units*.

In this chapter, we will only consider *feed-forward* networks. In a feed-forward network, you can think of the network as defining a function-call graph that is *acyclic*: that is, the input to a neuron can never depend on that neuron's output. Data flows, one way, from the inputs to the outputs, and the function computed by the network is just a composition of the functions computed by the individual neurons.

Although the graph structure of a neural network can really be anything (as long as it satisfies the feed-forward constraint), for simplicity in software and analysis, we usually organize them into *layers*. A layer is a group of neurons that are essentially "in parallel": their inputs are outputs of neurons in the previous layer, and their outputs are the input to the neurons in the next layer. We'll start by describing a single layer, and then go on to the case of multiple layers.

2.1 Single layer

A *layer* is a set of units that, as we have just described, are not connected to each other. The layer is called *fully connected* if, as in the diagram below, the inputs to each unit in the layer are the same (i.e. x_1, x_2, \dots, x_m in this case). A layer has input $x \in \mathbb{R}^m$ and output (also known as *activation*) $a \in \mathbb{R}^n$.



Since each unit has a vector of weights and a single offset, we can think of the weights of the whole layer as a matrix, W , and the collection of all the offsets as a vector W_0 . If we have m inputs, n units, and n outputs, then

- W is an $m \times n$ matrix,
- W_0 is an $n \times 1$ column vector,
- X , the input, is an $m \times 1$ column vector,
- Z , the pre-activation, is an $n \times 1$ column vector,
- A , the activation, is an $n \times 1$ column vector,

and we can represent the output vector as follows

$$A = f(Z) = f(W^T X + W_0) .$$

The activation function f is applied element-wise to the pre-activation values Z .

What can we do with a single layer? We have already seen single-layer networks, in the form of linear separators and linear regressors. All we can do with a single layer is make a linear hypothesis (with some possible linear transformation on the output). The whole reason for moving to neural networks is to move in the direction of *non-linear* hypotheses. To do this, we will have to consider multiple layers.

2.2 Many layers

A single neural network generally combines multiple layers, most typically by feeding the outputs of one layer into the inputs of another


We have to start by establishing some nomenclature. We will use l to name a layer, and let m^l be the number of inputs to the layer and n^l be the number of outputs from the layer. Then, W^l and W_0^l are of shape $m^l \times n^{l-1}$ and $n^l \times 1$, respectively. Let f^l be the activation function of layer l . Then, the pre-activation outputs are the $n^l \times 1$ vector

$$Z^l = W^l A^{l-1} + W_0^l$$

and the activated outputs are simply the $n^l \times 1$ vector

$$A^l = f^l(Z^l) .$$

Here's a diagram of a many-layered network, with two blocks for each layer, one representing the linear part of the operation and one representing the non-linear activation function. We will use this structural decomposition to organize our algorithmic thinking and implementation.



It is technically possible to have different activation functions within the same layer, but, again, for convenience in specification and implementation, we generally have the same activation function within a layer.

3 Choices of activation function

There are many possible choices for the activation function. We will start by thinking about whether it's really necessary to have an f at all.

What happens if we let f be the identity? Then, in a network with L layers (we'll leave out W_0 for simplicity, but it won't change the form of this argument),

$$A^L = W^L A^{L-1} = W^L W^{L-1} \dots W^1 X$$

So, multiplying out the weight matrices, we find that

$$A^L = W^{\text{total}} X ,$$

which is a *linear* function of X ! Having all those layers did not change the representational capacity of the network: the non-linearity of the activation function is crucial.

Study Question: Convince yourself that any function representable by any number of linear layers (where f is the identity function) can be represented by a single layer.

Now that we are convinced we need a non-linear activation, let's examine a few common choices.

Step function

$$\text{step}(z) = \begin{cases} 0 & \text{if } z < 0 \\ 1 & \text{otherwise} \end{cases}$$

Rectified linear unit

$$\text{ReLU}(z) = \begin{cases} 0 & \text{if } z < 0 \\ z & \text{otherwise} \end{cases} = \max(0, z)$$

Sigmoid function Also known as a *logistic* function, can be interpreted as probability, because for any value of z the output is in $[0, 1]$


$$\sigma(z) = \frac{1}{1 + e^{-z}}$$

Hyperbolic tangent Always in the range $[-1, 1]$

$$\tanh(z) = \frac{e^z - e^{-z}}{e^z + e^{-z}}$$

Softmax function Takes a whole a vector $Z \in \mathbb{R}^n$ and generates as output a vector $A \in [0, 1]^n$ with the property that $\sum_{i=1}^n A_i = 1$, which means we can interpret it as a probability distribution over n items:

$$\text{softmax}(z) = \begin{bmatrix} \exp(z_1) / \sum_i \exp(z_i) \\ \vdots \\ \exp(z_n) / \sum_i \exp(z_i) \end{bmatrix}$$



The original idea for neural networks involved using the **step** function as an activation, but because the derivative is discontinuous, we won't be able to use gradient-descent methods to tune the weights in a network with step functions, so we won't consider them further. They have been replaced, in a sense, by the sigmoid, relu, and tanh activation functions.

Study Question: Consider sigmoid, relu, and tanh activations. Which one is most like a step function? Is there an additional parameter you could add to a sigmoid that would make it be more like a step function?

Study Question: What is the derivative of the relu function? Are there some values of the input for which the derivative vanishes?

ReLU's are especially common in internal ("hidden") layers, and sigmoid activations are common for the output for binary classification and softmax for multi-class classification (see section 6.1 for an explanation).

4 Error back-propagation

We will train neural networks using gradient descent methods. It's possible to use *batch* gradient descent, in which we sum up the gradient over all the points (as in section 2 of chapter 6) or stochastic gradient descent (SGD), in which we take a small step with respect to the gradient considering a single point at a time (as in section 4 of chapter 6).

Our notation is going to get pretty hairy pretty quickly. To keep it as simple as we can, we'll focus on taking the gradient with respect to a single point, for SGD; you can simply sum up these gradients over all the data points if you wish to do batch descent.

So, to do SGD for a training example (x, y) , we need to compute $\nabla_W \text{Loss}(\text{NN}(x; W), y)$, where W represents all weights W^l, W_0^l in all the layers $l = (1, \dots, L)$. This seems terrifying, but is actually quite easy to do using the chain rule.

Remember that we are always computing the gradient of the loss function *with respect to the weights* for a particular value of (x, y) . That tells us how much we want to change the weights, in order to reduce the loss incurred on this particular training example.

First, let's see how the loss depends on the weights in the final layer, W^L . Remembering that our output is A^L , and using the shorthand loss to stand for $\text{Loss}(\text{NN}(x; W), y)$ which is equal to $\text{Loss}(A^L, y)$, and finally that $A^L = f^L(Z^L)$ and $Z^L = W^{L^T} A^{L-1}$, we can use the chain rule:

$$\frac{\partial \text{loss}}{\partial W^L} = \underbrace{\frac{\partial \text{loss}}{\partial A^L}}_{\text{depends on loss function}} \cdot \underbrace{\frac{\partial A^L}{\partial Z^L}}_{f^{L'}} \cdot \underbrace{\frac{\partial Z^L}{\partial W^L}}_{A^{L-1}}.$$

To actually get the dimensions to match, we need to write this a bit more carefully, and note that it is true for any l , including $l = L$:

$$\underbrace{\frac{\partial \text{loss}}{\partial W^l}}_{m^l \times n^l} = \underbrace{A^{l-1}}_{m^l \times 1} \underbrace{\left(\frac{\partial \text{loss}}{\partial Z^l} \right)^T}_{1 \times n^l} \quad (8.1)$$

Yay! So, in order to find the gradient of the loss with respect to the weights in the other layers of the network, we just need to be able to find $\partial \text{loss} / \partial Z^l$.

If we repeatedly apply the chain rule, we get this expression for the gradient of the loss

Remember the chain rule! If $a = f(b)$ and $b = g(c)$ (so that $a = f(g(c))$), then $\frac{da}{dc} = \frac{da}{db} \cdot \frac{db}{dc} = f'(b)g'(c) = f'(g(c))g'(c)$.

It might reasonably bother you that $\partial Z^L / \partial W^L = A^{L-1}$. We're somehow thinking about the derivative of a vector with respect to a matrix, which seems like it might need to be a three-dimensional thing. But note that $\partial Z^L / \partial W^L$ is really $\partial W^{L^T} A^{L-1} / \partial W^L$ and it seems okay in at least an informal sense that it's A^{L-1} .

with respect to the pre-activation in the first layer:

$$\frac{\partial \text{loss}}{\partial Z^1} = \underbrace{\frac{\partial \text{loss}}{\partial A^L} \cdot \frac{\partial A^L}{\partial Z^L} \cdot \frac{\partial Z^L}{\partial A^{L-1}} \cdot \frac{\partial A^{L-1}}{\partial Z^{L-1}} \cdots \frac{\partial A^2}{\partial Z^2} \cdot \frac{\partial Z^2}{\partial A^1} \cdot \frac{\partial A^1}{\partial Z^1}}_{\partial \text{loss} / \partial A^1} \quad (8.2)$$

This derivation was informal, to show you the general structure of the computation. In fact, to get the dimensions to all work out, we just have to write it backwards! Let's first understand more about these quantities:

- $\partial \text{loss} / \partial A^L$ is $n^L \times 1$ and depends on the particular loss function you are using.
- $\partial Z^L / \partial A^{L-1}$ is $m^L \times n^L$ and is just W^L (you can verify this by computing a single entry $\partial Z_i^L / \partial A_j^{L-1}$).
- $\partial A^L / \partial Z^L$ is $n^L \times n^L$. It's a little tricky to think about. Each element $a_i^L = f^L(z_i^L)$. This means that $\partial a_i^L / \partial z_j^L = 0$ whenever $i \neq j$. So, the off-diagonal elements of $\partial A^L / \partial Z^L$ are all 0, and the diagonal elements are $\partial a_i^L / \partial z_i^L = f^{L'}(z_i^L)$.


Now, we can rewrite equation 8.2 so that the quantities match up as

$$\frac{\partial \text{loss}}{\partial Z^1} = \frac{\partial A^L}{\partial Z^L} \cdot W^{L+1} \cdot \frac{\partial A^{L+1}}{\partial Z^{L+1}} \cdots W^{L-1} \cdot \frac{\partial A^{L-1}}{\partial Z^{L-1}} \cdot W^L \cdot \frac{\partial A^L}{\partial Z^L} \cdot \frac{\partial \text{loss}}{\partial A^L} \quad (8.3)$$

Using equation 8.3 to compute $\partial \text{loss} / \partial Z^1$ combined with equation 8.1, lets us find the gradient of the loss with respect to any of the weight matrices.

Study Question: Apply the same reasoning to find the gradients of loss with respect to W_0^L .

This general process is called *error back-propagation*. The idea is that we first do a *forward pass* to compute all the a and z values at all the layers, and finally the actual loss on this example. Then, we can work backward and compute the gradient of the loss with respect to the weights in each layer, starting at layer L and going back to layer 1.



If we view our neural network as a sequential composition of modules (in our work so far, it has been an alternation between a linear transformation with a weight matrix, and a component-wise application of a non-linear activation function), then we can define a simple API for a module that will let us compute the forward and backward passes, as well as do the necessary weight updates for gradient descent. Each module has to provide the following "methods." We are already using letters a, x, y, z with particular meanings, so here we will use u as the vector input to the module and v as the vector output:

- forward: $u \rightarrow v$
- backward: $u, v, \partial L / \partial v \rightarrow \partial L / \partial u$
- weight grad: $u, \partial L / \partial v \rightarrow \partial L / \partial W$ only needed for modules that have weights W

In homework we will ask you to implement these modules for neural network components, and then use them to construct a network and train it as described in the next section.

I like to think of this as "blame propagation". You can think of loss as how mad we are about the prediction that the network just made. Then $\partial \text{loss} / \partial A^L$ is how much we blame A^L for the loss. The last module has to take in $\partial \text{loss} / \partial A^L$ and compute $\partial \text{loss} / \partial Z^L$, which is how much we blame Z^L for the loss. The next module (working backwards) takes in $\partial \text{loss} / \partial Z^L$ and computes $\partial \text{loss} / \partial A^{L-1}$. So every module is accepting its blame for the loss, computing how much of it to allocate to each of its inputs, and passing the blame back to them.

5 Training

Here we go! Here's how to do stochastic gradient descent training on a feed-forward neural network. After this pseudo-code, we motivate the choice of initialization in lines 2 and 3. The actual computation of the gradient values (e.g. $\partial \text{loss} / \partial A^L$) is not directly defined in this code, because we want to make the structure of the computation clear.

Study Question: What is $\partial Z^L / \partial W^L$?

Study Question: Which terms in the code below depend on f^L ?

```

TRAIN-NEURAL-NET( $\mathcal{D}_n, T, L, (m^1, \dots, m^L), (f^1, \dots, f^L)$ )
1  for l = 1 to L
2     $W_{ij}^l \sim \text{Gaussian}(0, 1/m^l)$ 
3     $W_{0j}^l \sim \text{Gaussian}(0, 1)$ 
4  for t = 1 to T
5    i = random sample from  $\{1, \dots, n\}$ 
6     $A^0 = x^{(i)}$ 
7    // forward pass to compute the output  $A^L$ 
8    for l = 1 to L
9       $Z^l = W^{lT} A^{l-1} + W_0^l$ 
10      $A^l = f^l(Z^l)$ 
11    loss = Loss( $A^L, y^{(i)}$ )
12    for l = L to 1:
13      // error back-propagation
14       $\partial \text{loss} / \partial A^l = \text{if } l < L \text{ then } \partial \text{loss} / \partial Z^{l+1} \cdot \partial Z^{l+1} / \partial A^l \text{ else } \partial \text{loss} / \partial A^L$ 
15       $\partial \text{loss} / \partial Z^l = \partial \text{loss} / \partial A^l \cdot \partial A^l / \partial Z^l$ 
16      // compute gradient with respect to weights
17       $\partial \text{loss} / \partial W^l = \partial \text{loss} / \partial Z^l \cdot \partial Z^l / \partial W^l$ 
18       $\partial \text{loss} / \partial W_0^l = \partial \text{loss} / \partial Z^l \cdot \partial Z^l / \partial W_0^l$ 
19      // stochastic gradient descent update
20       $W^l = W^l - \eta(t) \cdot \partial \text{loss} / \partial W^l$ 
21       $W_0^l = W_0^l - \eta(t) \cdot \partial \text{loss} / \partial W_0^l$ 

```

Initializing W is important; if you do it badly there is a good chance the neural network training won't work well. First, it is important to initialize the weights to random values. We want different parts of the network to tend to "address" different aspects of the problem; if they all start at the same weights, the symmetry will often keep the values from moving in useful directions. Second, many of our activation functions have (near) zero slope when the pre-activation z values have large magnitude, so we generally want to keep the initial weights small so we will be in a situation where the gradients are non-zero, so that gradient descent will have some useful signal about which way to go.

One good general-purpose strategy is to choose each weight at random from a Gaussian (normal) distribution with mean 0 and standard deviation $(1/m)$ where m is the number of inputs to the unit.

Study Question: If the input x to this unit is a vector of 1's, what would the expected pre-activation z value be with these initial weights?

We write this choice (where \sim means "is drawn randomly from the distribution")

$$W_{ij}^l \sim \text{Gaussian}\left(0, \frac{1}{m^l}\right).$$

It will often turn out (especially for fancier activations and loss functions) that computing

$$\frac{\partial \text{loss}}{\partial Z^L}$$

is easier than computing

$$\frac{\partial \text{loss}}{\partial A^L} \text{ and } \frac{\partial A^L}{\partial Z^L} .$$

So, we may instead ask for an implementation of a loss function to provide a backward method that computes $\partial \text{loss} / \partial Z^L$ directly.

6 Loss functions and activation functions

Different loss functions make different assumptions about the range of inputs they will get as input and, as we have seen, different activation functions will produce output values in different ranges. When you are designing a neural network, it's important to make these things fit together well. In particular, we will think about matching loss functions with the activation function in the last layer, f^L . Here is a table of loss functions and activations that make sense for them:

Loss	f^L
squared	linear
hinge	linear
NLL	sigmoid
NLLM	softmax

But what is NLL?

6.1 Two-class classification and log likelihood

For classification, the natural loss function is 0-1 loss, but we have already discussed the fact that it's very inconvenient for gradient-based learning because its derivative is discontinuous. Hinge loss gives us a way, for binary classification problems, to make a smoother objective. An alternative loss function that has a nice probabilistic interpretation, is in popular use, and extends nicely to multi-class classification is called *negative log likelihood* (NLL). We will discuss it first in the two-class case, and then generalize to multiple classes.

Let's assume that the activation function on the output layer is a sigmoid and that there is a single unit in the output layer, so the output of the whole neural network is a scalar, a^L . Because f^L is a sigmoid, we know $a^L \in [0, 1]$, and we can interpret it as the probability that the input x is a positive example. Let us further assume that the labels in the training data are $y \in \{0, 1\}$, so they can also be interpreted as probabilities.

We might want to pick the parameters of our network to maximize the probability that the network assigns the correct labels to all the points. That would be

$$\prod_{i=1}^n \begin{cases} a^{(i)} & \text{if } y^{(i)} = 1 \\ 1 - a^{(i)} & \text{otherwise} \end{cases} ,$$

under the assumption that our predictions are independent. This can be cleverly rewritten as

$$\prod_{i=1}^n a^{(i)y^{(i)}} (1 - a^{(i)})^{1-y^{(i)}} .$$

Study Question: Be sure you can see why these two expressions are the same.

Now, because products are kind of hard to deal with, and because the log function is monotonic, the W that maximizes the log of this quantity will be the same as the W that maximizes the original, so we can try to maximize

$$\sum_{i=1}^n y^{(i)} \log a^{(i)} + (1 - y^{(i)}) \log(1 - a^{(i)}) ,$$

which we can write in terms of a loss function

$$\sum_{i=1}^n \mathcal{L}_{\text{nll}}(a^{(i)}, y^{(i)})$$

where \mathcal{L}_{nll} is the *negative log likelihood* loss function:

$$\mathcal{L}_{\text{nll}}(\text{guess}, \text{actual}) = -(\text{actual} \cdot \log(\text{guess}) + (1 - \text{actual}) \cdot \log(1 - \text{guess})) .$$

This loss function is also sometimes referred to as the *log loss* or *cross entropy*.

You can use any base for the logarithm and it won't make any real difference. If we ask you for numbers, use log base e .

6.2 Multi-class classification and log likelihood

We can extend this idea directly to multi-class classification with K classes, where the training label is represented with the one-hot vector $\mathbf{y} = [y_1, \dots, y_K]^T$, where $y_k = 1$ if the example is of class k . Assume that our network uses *softmax* as the activation function in the last layer, so that the output is $\mathbf{a} = [a_1, \dots, a_K]^T$, which represents a probability distribution over the K possible classes. Then, the probability that our network predicts the correct class for this example is $\prod_{k=1}^K a_k^{y_k}$ and the log of the probability that it is correct is $\sum_{k=1}^K y_k \log a_k$, so

$$\mathcal{L}_{\text{nllm}}(\text{guess}, \text{actual}) = - \sum_{k=1}^K \text{actual}_k \cdot \log(\text{guess}_k) .$$

We'll call this NLLM for *negative log likelihood multiclass*.

Study Question: Show that $\mathcal{L}_{\text{nllm}}$ for $K = 2$ is the same as \mathcal{L}_{nll} .

7 Optimizing neural network parameters

Because neural networks are just parametric functions, we can optimize loss with respect to the parameters using standard gradient-descent software, but we can take advantage of the structure of the loss function and the hypothesis class to improve optimization. As we have seen, the modular function-composition structure of a neural network hypothesis makes it easy to organize the computation of the gradient. As we have also seen earlier, the structure of the loss function as a sum over terms, one per training data point, allows us to consider stochastic gradient methods. In this section we'll consider some alternative strategies for organizing training, and also for making it easier to handle the step-size parameter.

7.1 Batches

Assume that we have an objective of the form

$$J(W) = \sum_{i=1}^n \mathcal{L}(h(x^{(i)}; W), y^{(i)}) ,$$

where h is the function computed by a neural network, and W stands for all the weight matrices and vectors in the network.

When we perform *batch* gradient descent, we use the update rule

$$W := W - \eta \nabla_W J(W) ,$$

which is equivalent to

$$W := W - \eta \sum_{i=1}^n \nabla_W \mathcal{L}(h(x^{(i)}; W), y^{(i)}) .$$

So, we sum up the gradient of loss at each training point, with respect to W , and then take a step in the negative direction of the gradient.

In *stochastic* gradient descent, we repeatedly pick a point $(x^{(i)}, y^{(i)})$ at random from the data set, and execute a weight update on that point alone:

$$W := W - \eta \nabla_W \mathcal{L}(h(x^{(i)}; W), y^{(i)}) .$$

As long as we pick points uniformly at random from the data set, and decrease η at an appropriate rate, we are guaranteed to converge to at least a local optimum.

These two methods have offsetting virtues. The batch method takes steps in the exact gradient direction but requires a lot of computation before even a single step can be taken, especially if the data set is large. The stochastic method begins moving right away, and can sometimes make very good progress before looking at even a substantial fraction of the whole data set, but if there is a lot of variability in the data, it might require a very small η to effectively average over the individual steps moving in "competing" directions.

An effective strategy is to "average" between batch and stochastic gradient descent by using *mini-batches*. For a mini-batch of size k , we select k distinct data points uniformly at random from the data set and do the update based just on their contributions to the gradient

$$W := W - \eta \sum_{i=1}^k \nabla_W \mathcal{L}(h(x^{(i)}; W), y^{(i)}) .$$

Most neural network software packages are set up to do mini-batches.

Study Question: For what value of k is mini-batch gradient descent equivalent to stochastic gradient descent? To batch gradient descent?

Picking k unique data points at random from a large data-set is potentially computationally difficult. An alternative strategy, if you have an efficient procedure for randomly shuffling the data set (or randomly shuffling a list of indices into the data set) is to operate in a loop, roughly as follows:

MINI-BATCH-SGD(NN, data, k)

```

1  n = length(data)
2  while not done:
3      RANDOM-SHUFFLE(data)
4      for i = 1 to n/k
5          BATCH-GRADIENT-UPDATE(NN, data[(i - 1)k : ik])

```

7.2 Adaptive step-size

Picking a value for η is difficult and time-consuming. If it's too small, then convergence is slow and if it's too large, then we risk divergence or slow convergence due to oscillation. This problem is even more pronounced in stochastic or mini-batch mode, because we know we need to decrease the step size for the formal guarantees to hold.

It's also true that, within a single neural network, we may well want to have different step sizes. As our networks become *deep* (with increasing numbers of layers) we can find that magnitude of the gradient of the loss with respect to the weights in the last layer, $\partial \text{loss} / \partial W_L$, may be substantially different from the gradient of the loss with respect to the weights in the first layer $\partial \text{loss} / \partial W_1$. If you look carefully at equation 8.3, you can see that the output gradient is multiplied by all the weight matrices of the network and is "fed back" through all the derivatives of all the activation functions. This can lead to a problem of *exploding* or *vanishing* gradients, in which the back-propagated gradient is much too big or small to be used in an update rule with the same step size.

So, we'll consider having an independent step-size parameter *for each weight*, and updating it based on a local view of how the gradient updates have been going.

7.2.1 Running averages

We'll start by looking at the notion of a *running average*. It's a computational strategy for estimating a possibly weighted average of a sequence of data. Let our data sequence be a_1, a_2, \dots ; then we define a sequence of running average values, A_0, A_1, A_2, \dots using the equations

$$\begin{aligned}
 A_0 &= 0 \\
 A_t &= \gamma_t A_{t-1} + (1 - \gamma_t) a_t
 \end{aligned}$$

where $\gamma_t \in (0, 1)$. If γ_t is a constant, then this is a *moving average*, in which

$$\begin{aligned}
 A_T &= \gamma A_{T-1} + (1 - \gamma) a_T \\
 &= \gamma (A_{T-2} + (1 - \gamma) a_{T-1}) + (1 - \gamma) a_T \\
 &= \sum_{t=0}^T \gamma^{T-t} (1 - \gamma) a_t
 \end{aligned}$$

So, you can see that inputs a_t closer to the end of the sequence have more effect on A_t than early inputs.

This section is very strongly influenced by Sebastian Ruder's excellent blog posts on the topic: ruder.io/optimizing-gradient-descent

If, instead, we set $\gamma_t = (t - 1)/t$, then we get the actual average.

Study Question: Prove to yourself that the previous assertion holds.

7.2.2 Momentum

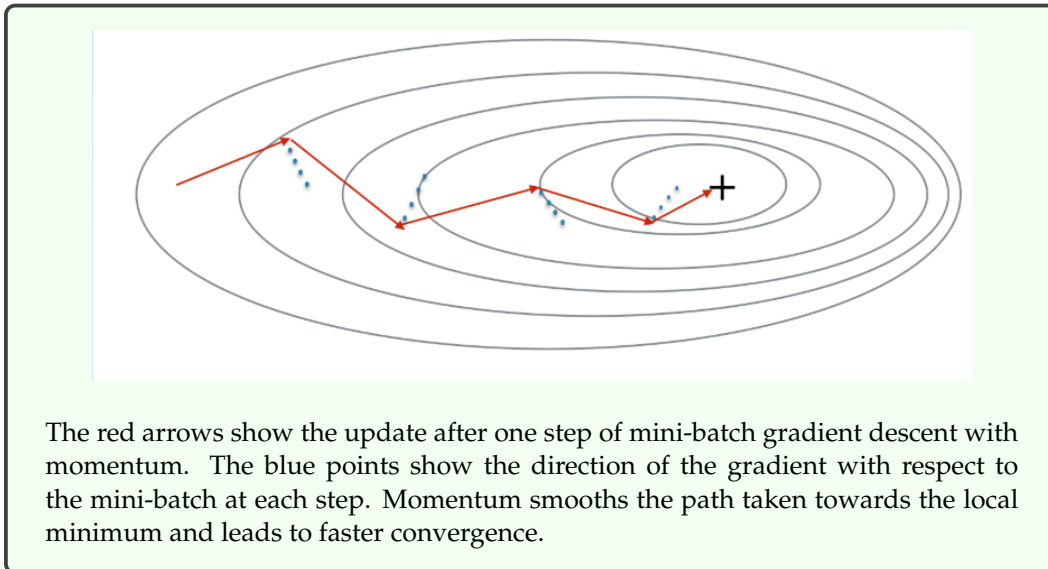
Now, we can use methods that are a bit like running averages to describe strategies for computing η . The simplest method is *momentum*, in which we try to “average” recent gradient updates, so that if they have been bouncing back and forth in some direction, we take out that component of the motion. For momentum, we have

$$\begin{aligned}V_0 &= 0 \\V_t &= \gamma V_{t-1} + \eta \nabla_W J(W_{t-1}) \\W_t &= W_{t-1} - V_t\end{aligned}$$

This doesn't quite look like an adaptive step size. But what we can see is that, if we let $\eta = \eta'(1 - \gamma)$, then the rule looks exactly like doing an update with step size η' on a moving average of the gradients with parameter γ :

$$\begin{aligned}V_0 &= 0 \\M_t &= \gamma M_{t-1} + (1 - \gamma) \nabla_W J(W_{t-1}) \\W_t &= W_{t-1} - \eta' M_t\end{aligned}$$

We will find that V_t will be bigger in dimensions that consistently have the same sign for ∇_θ and smaller for those that don't. Of course we now have *two* parameters to set, but the hope is that the algorithm will perform better overall, so it will be worth trying to find good values for them. Often γ is set to be something like 0.9.



Study Question: If you set $\gamma = 0.1$, would momentum have more of an effect or less of an effect than if you set it to 0.9?

7.2.3 Adadelta

Another useful idea is this: we would like to take larger steps in parts of the space where $J(W)$ is nearly flat (because there's no risk of taking too big a step due to the gradient

being large) and smaller steps when it is steep. We'll apply this idea to each weight independently, and end up with a method called *adadelta*, which is a variant on *adagrad* (for adaptive gradient). Even though our weights are indexed by layer, input unit and output unit, for simplicity here, just let W_j be any weight in the network (we will do the same thing for all of them).

$$\begin{aligned} g_{t,j} &= \nabla_W J(W_{t-1})_j \\ G_{t,j} &= \gamma G_{t-1,j} + (1 - \gamma) g_{t,j}^2 \\ W_{t,j} &= W_{t-1,j} - \frac{\eta}{\sqrt{G_{t,j} + \epsilon}} g_{t,j} \end{aligned}$$

The sequence $G_{t,j}$ is a moving average of the square of the j th component of the gradient. We square it in order to be insensitive to the sign—we want to know whether the magnitude is big or small. Then, we perform a gradient update to weight j , but divide the step size by $\sqrt{G_{t,j} + \epsilon}$, which is larger when the surface is steeper in direction j at point W_{t-1} in weight space; this means that the step size will be smaller when it's steep and larger when it's flat.

7.2.4 Adam

Adam has become the default method of managing step sizes neural networks. It combines the ideas of momentum and adadelta. We start by writing moving averages of the gradient and squared gradient, which reflect estimates of the mean and variance of the gradient for weight j :

$$\begin{aligned} g_{t,j} &= \nabla_W J(W_{t-1})_j \\ m_{t,j} &= B_1 m_{t-1,j} + (1 - B_1) g_{t,j} \\ v_{t,j} &= B_2 v_{t-1,j} + (1 - B_2) g_{t,j}^2 \end{aligned}$$

A problem with these estimates is that, if we initialize $m_0 = v_0 = 0$, they will always be biased (slightly too small). So we will correct for that bias by defining

$$\begin{aligned} \hat{m}_{t,j} &= \frac{m_{t,j}}{1 - B_1^t} \\ \hat{v}_{t,j} &= \frac{v_{t,j}}{1 - B_2^t} \\ W_{t,j} &= W_{t-1,j} - \frac{\eta}{\sqrt{\hat{v}_{t,j} + \epsilon}} \hat{m}_{t,j} \end{aligned}$$

Note that B_1^t is B_1 raised to the power t , and likewise for B_2^t . To justify these corrections, note that if we were to expand $m_{t,j}$ in terms of $m_{0,j}$ and $g_{0,j}, g_{1,j}, \dots, g_{t,j}$ the coefficients would sum to 1. However, the coefficient behind $m_{0,j}$ is B_1^t and since $m_{0,j} = 0$, the sum of coefficients of non-zero terms is $1 - B_1^t$, hence the correction. The same justification holds for $v_{t,j}$.

Now, our update for weight j has a step size that takes the steepness into account, as in adadelta, but also tends to move in the same direction, as in momentum. The authors of this method propose setting $B_1 = 0.9, B_2 = 0.999, \epsilon = 10^{-8}$. Although we now have even more parameters, Adam is not highly sensitive to their values (small changes do not have a huge effect on the result).

Study Question: Define \hat{m}_j directly as a moving average of $g_{t,j}$. What is the decay?

Even though we now have a step-size for each weight, and we have to update various quantities on each iteration of gradient descent, it's relatively easy to implement by maintaining a matrix for each quantity ($m_t^\ell, v_t^\ell, g_t^\ell, g_t^{2\ell}$) in each layer of the network.

Although, interestingly, it may actually violate the convergence conditions of SGD:
arxiv.org/abs/1705.08292

8 Regularization

So far, we have only considered optimizing loss on the training data as our objective for neural network training. But, as we have discussed before, there is a risk of overfitting if we do this. The pragmatic fact is that, in current deep neural networks, which tend to be very large and to be trained with a large amount of data, overfitting is not a huge problem. This runs counter to our current theoretical understanding and the study of this question is a hot area of research. Nonetheless, there are several strategies for regularizing a neural network, and they can sometimes be important.

8.1 Methods related to ridge regression

One group of strategies can, interestingly, be shown to have similar effects: early stopping, weight decay, and adding noise to the training data.

Early stopping is the easiest to implement and is in fairly common use. The idea is to train on your training set, but at every *epoch* (pass through the whole training set, or possibly more frequently), evaluate the loss of the current W on a *validation set*. It will generally be the case that the loss on the training set goes down fairly consistently with each iteration, the loss on the validation set will initially decrease, but then begin to increase again. Once you see that the validation loss is systematically increasing, you can stop training and return the weights that had the lowest validation error.

Another common strategy is to simply penalize the norm of all the weights, as we did in ridge regression. This method is known as *weight decay*, because when we take the gradient of the objective

$$J(W) = \sum_{i=1}^n \text{Loss}(\text{NN}(x^{(i)}, y^{(i)}; W) + \lambda \|W\|^2$$

we end up with an update of the form

$$\begin{aligned} W_t &= W_{t-1} - \eta \left(\left(\nabla_W \text{Loss}(\text{NN}(x^{(i)}, y^{(i)}; W_{t-1}) \right) + \lambda W_{t-1} \right) \\ &= W_{t-1}(1 - \lambda\eta) - \eta \left(\nabla_W \text{Loss}(\text{NN}(x^{(i)}, y^{(i)}; W_{t-1}) \right) . \end{aligned}$$

This rule has the form of first “decaying” W_{t-1} by a factor of $(1 - \lambda\eta)$ and then taking a gradient step.

Finally, the same effect can be achieved by perturbing the $x^{(i)}$ values of the training data by adding a small amount of zero-mean normally distributed noise before each gradient computation. It makes intuitive sense that it would be more difficult for the network to overfit to particular training data if they are changed slightly on each training step.

8.2 Dropout

Dropout is a regularization method that was designed to work with deep neural networks. The idea behind it is, rather than perturbing the data every time we train, we’ll perturb the network! We’ll do this by randomly, on each training step, selecting a set of units in each layer and prohibiting them from participating. Thus, all of the units will have to take a kind of “collective” responsibility for getting the answer right, and will not be able to rely on any small subset of the weights to do all the necessary computation. This tends also to make the network more robust to data perturbations.

During the training phase, for each training example, for each unit, randomly with probability p temporarily set $a_j^\ell := 0$. There will be no contribution to the output and no gradient update for the associated unit.

Result is due to Bishop, described in his textbook and here doi.org/10.1162/neco.1995.7.1.108.

Study Question: Be sure you understand why, when using SGD, setting an activation value to 0 will cause that unit's weights not to be updated on that iteration.

When we are done training and want to use the network to make predictions, we multiply all weights by p to achieve the same average activation levels.

Implementing dropout is easy! In the forward pass during training, we let

$$a^\ell = f(z^\ell) * d^\ell$$

where $*$ denotes component-wise product and d^ℓ is a vector of 0's and 1's drawn randomly with probability p . The backwards pass depends on a^ℓ , so we do not need to make any further changes to the algorithm.

It is common to set p to 0.5, but this is something one might experiment with to get good results on your problem and data.

8.3 Batch Normalization

A slightly more modern alternative to dropout, which achieves somewhat better performance, is *batch normalization*. It was originally developed to address a problem of *covariate shift*: that is, if you consider the second layer of a two-layer neural network, the distribution of its input values is changing over time as the first layer's weights change. Learning when the input distribution is changing is extra difficult: you have to change your weights to improve your predictions, but also just to compensate for a change in your inputs (imagine, for instance, that the magnitude of the inputs to your layer is increasing over time—then your weights will have to decrease, just to keep your predictions the same).

So, when training with mini-batches, the idea is to *standardize* the input values for each mini-batch, just in the way that we did it in section of chapter 4, subtracting off the mean and dividing by the standard deviation of each input dimension. This means that the scale of the inputs to each layer remains the same, no matter how the weights in previous layers change. However, this somewhat complicates matters, because the computation of the weight updates will need to take into account that we are performing this transformation. In the modular view, batch normalization can be seen as a module that is applied to a^ℓ , interposed after the output of f^ℓ and before the product with $W^{\ell+1}$.

Batch normalization ends up having a regularizing effect for similar reasons that adding noise and dropout do: each mini-batch of data ends up being mildly perturbed, which prevents the network from exploiting very particular values of the data points.

For more details see
arxiv.org/abs/1502.03167.

Convolutional Neural Networks

So far, we have studied what are called *fully connected* neural networks, in which all of the units at one layer are connected to all of the units in the next layer. This is a good arrangement when we don't know anything about what kind of mapping from inputs to outputs we will be asking the network to learn to approximate. But if we *do* know something about our problem, it is better to build it into the structure of our neural network. Doing so can save computation time and significantly diminish the amount of training data required to arrive at a solution that generalizes robustly.

One very important application domain of neural networks, where the methods have achieved an enormous amount of success in recent years, is signal processing. Signals might be spatial (in two-dimensional camera images or three-dimensional depth or CAT scans) or temporal (speech or music). If we know that we are addressing a signal-processing problem, we can take advantage of *invariant* properties of that problem. In this chapter, we will focus on two-dimensional spatial problems (images) but use one-dimensional ones as a simple example. Later, we will address temporal problems.

Imagine that you are given the problem of designing and training a neural network that takes an image as input, and outputs a classification, which is positive if the image contains a cat and negative if it does not. An image is described as a two-dimensional array of *pixels*, each of which may be represented by three integer values, encoding intensity levels in red, green, and blue color channels.

A *pixel* is a "picture element."

There are two important pieces of prior structural knowledge we can bring to bear on this problem:

- **Spatial locality:** The set of pixels we will have to take into consideration to find a cat will be near one another in the image.
- **Translation invariance:** The pattern of pixels that characterizes a cat is the same no matter where in the image the cat occurs.

We will design neural network structures that take advantage of these properties.

So, for example, we won't have to consider some combination of pixels in the four corners of the image, in case they encode cat-ness.

Cats don't look different if they're on the left or the right side of the image.

1 Filters

We begin by discussing *image filters*. An image filter is a function that takes in a local spatial neighborhood of pixel values and detects the presence of some pattern in that data.

Let's consider a very simple case to start, in which we have a 1-dimensional binary "image" and a filter F of size two. The filter is a vector of two numbers, which we will move along the image, taking the dot product between the filter values and the image values at each step, and aggregating the outputs to produce a new image.

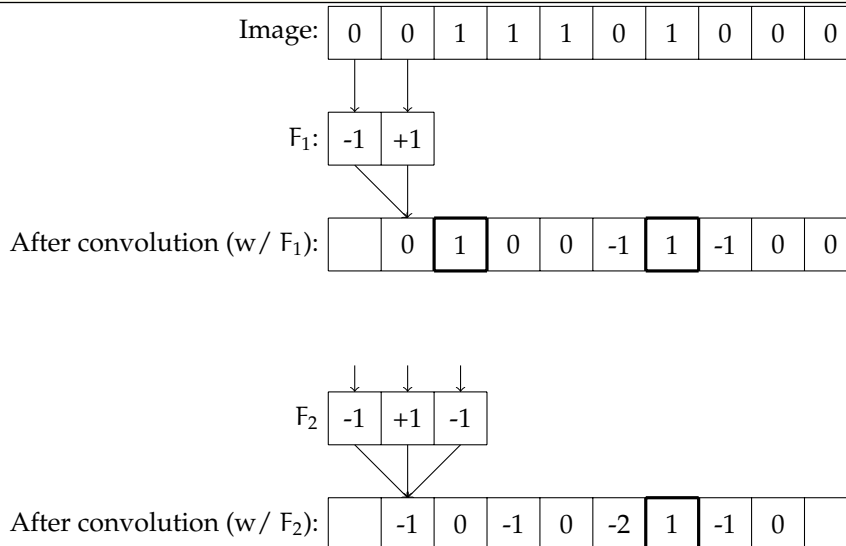
Let X be the original image, of size d ; then the output image is specified by

$$Y_i = F^T(X_{i-1}, X_i)^T$$

To ensure that the output image is also of dimension d , we will generally "pad" the input image with 0 values if we need to access pixels that are beyond the bounds of the input image. This process of applying the filter to the image to create a new image is called "convolution."

Here is a concrete example. Let the filter $F_1 = (-1, +1)$. Then given the first image below, we can convolve it with filter F_1 to obtain the second image. You can think of this filter as a detector for "left edges" in the original image—to see this, look at the places where there is a 1 in the output image, and see what pattern exists at that position in the input image. Another interesting filter is $F_2 = (-1, +1, -1)$. The third image below shows the result of convolving the first image with F_2 .


Study Question: Convince yourself that this filter can be understood as a detector for isolated positive pixels in the binary image.



Two-dimensional versions of filters like these are thought to be found in the visual cortex of all mammalian brains. Similar patterns arise from statistical analysis of natural images. Computer vision people used to spend a lot of time hand-designing *filter banks*. A filter bank is a set of sets of filters, arranged as shown in the diagram below.

Unfortunately in AI/ML/CS/Math, the word "filter" gets used in many ways: in addition to the one we describe here, it can describe a temporal process (in fact, our moving averages are a kind of filter) and even a somewhat esoteric algebraic structure.

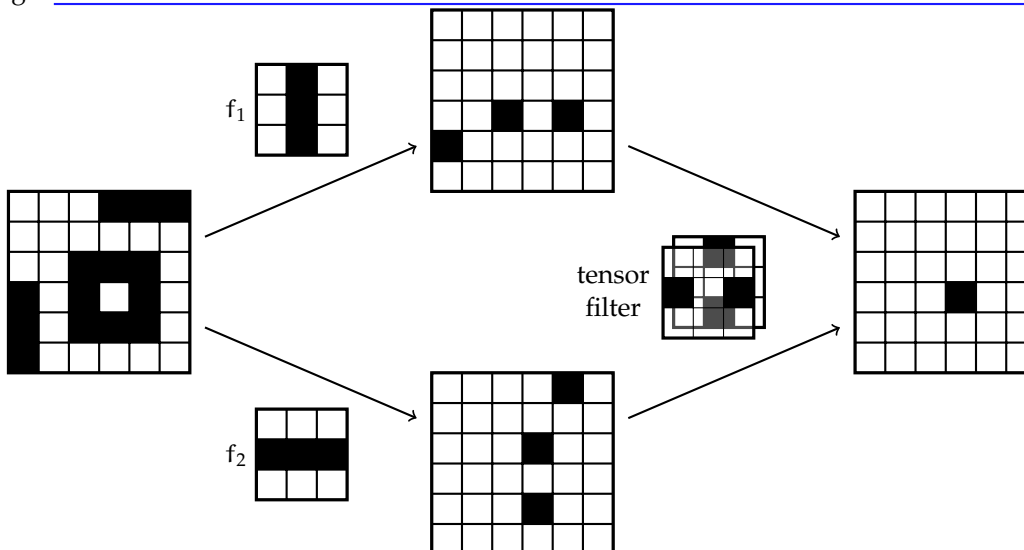
And filters are also sometimes called *convolutional kernels*.



All of the filters in the first group are applied to the original image; if there are k such filters, then the result is k new images, which are called *channels*. Now imagine stacking all these new images up so that we have a cube of data, indexed by the original row and column indices of the image, as well as by the channel. The next set of filters in the filter bank will generally be *three-dimensional*: each one will be applied to a sub-range of the row and column indices of the image and to all of the channels.

These 3D chunks of data are called *tensors*. [The algebra of tensors is fun, and a lot like matrix algebra, but we won't go into it in any detail.](#)

Here is a more complex example of two-dimensional filtering. We have two 3×3 filters in the first layer, f_1 and f_2 . You can think of each one as “looking” for three pixels in a row, f_1 vertically and f_2 horizontally. Assuming our input image is $n \times n$, then the result of filtering with these two filters is an $n \times n \times 2$ tensor. Now we apply a tensor filter (hard to draw!) that “looks for” a combination of two horizontal and two vertical bars (now represented by individual pixels in the two channels), resulting in a single final $n \times n$ image.



We will use a popular piece of neural-network software called *TensorFlow* because it makes operations on tensors easy.

When we have a color image as input, we treat it as having 3 channels, and hence as an $n \times n \times 3$ tensor.

We are going to design neural networks that have this structure. Each “bank” of the filter bank will correspond to a neural-network layer. The numbers in the individual filters will be the “weights” of the network, which we will train using gradient descent. What makes this interesting and powerful (and somewhat confusing at first) is that the same weights are used many many times in the computation of each layer. This *weight sharing* means that we can express a transformation on a large image with relatively few parameters; it also means we'll have to take care in figuring out exactly how to train it!

We will define a filter layer l formally with:

- number of filters m^l ;
- size of filters $k^l \times k^l \times m^{l-1}$;

For simplicity, we are assuming that all images and filters are square (having the same number of rows and columns). That is in no way necessary, but is usually fine and definitely simplifies our notation.

- *stride* s^l is the spacing at which we apply the filter to the image; in all of our examples so far, we have used a stride of 1, but if we were to “skip” and apply the filter only at odd-numbered indices of the image, then it would have a stride of two (and produce a resulting image of half the size);
- *input tensor size* $n^{l-1} \times n^{l-1} \times m^{l-1}$

This layer will produce output tensor of size $n^l \times n^l \times m^l$, where $n^l = \lfloor n^{l-1}/s^l \rfloor$. The weights are the values defining the filter: there will be m^l different $k^l \times k^l \times m^{l-1}$ tensors of weight values.

This may seem complicated, but we get a rich class of mappings that exploit image structure and have many fewer weights than a fully connected layer would.

Study Question: How many weights are in a convolutional layer specified as above?

Study Question: If we used a fully-connected layer with the same size inputs and outputs, how many weights would it have?

2 Max Pooling

It is typical to structure filter banks into a *pyramid*, in which the image sizes get smaller in successive layers of processing. The idea is that we find local patterns, like bits of edges in the early layers, and then look for patterns in those patterns, etc. This means that, effectively, we are looking for patterns in larger pieces of the image as we apply successive filters. Having a stride greater than one makes the images smaller, but does not necessarily aggregate information over that spatial range.

Both in engineering and in nature

Another common layer type, which accomplishes this aggregation, is *max pooling*. A max pooling layer operates like a filter, but has no weights. *You can think of it as a pure functional layer, like a ReLU layer in a fully connected network.* It has a filter size, as in a filter layer, but simply returns the maximum value in its field. Usually, we apply max pooling with the following traits:

We sometimes use the term *receptive field* or just *field* to mean the area of an input image that a filter is being applied to.

- $\text{stride} > 1$, so that the resulting image is smaller than the input image; and
- $k \geq \text{stride}$, so that the whole image is covered.

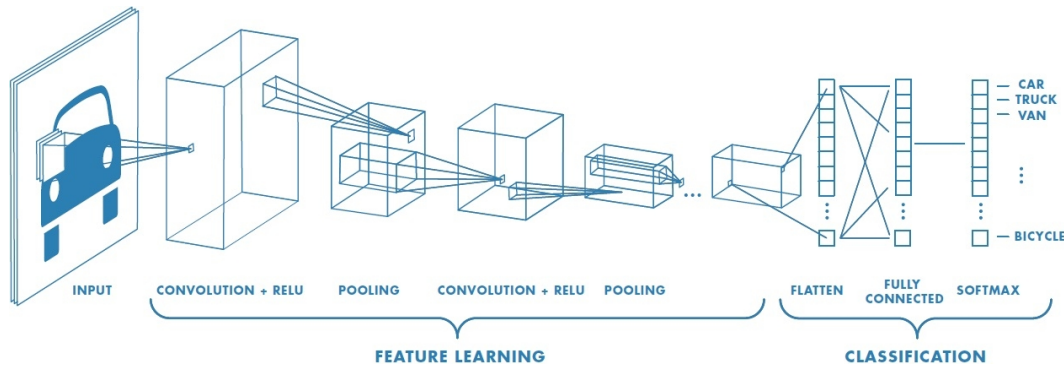
As a result of applying a max pooling layer, we don't keep track of the precise location of a pattern. This helps our filters to learn to recognize patterns independent of their location.

Consider a max pooling layer of $\text{stride} = k = 2$. This would map a $64 \times 64 \times 3$ image to a $32 \times 32 \times 3$ image.

Study Question: Maximilian Poole thinks it would be a good idea to add two max pooling layers of size k to their network. What single layer would be equivalent?

3 Typical architecture

Here is the form of a typical convolutional network:




Source: <https://www.mathworks.com/solutions/deep-learning/convolutional-neural-network.html>

After each filter layer there is generally a ReLU layer; there may be multiple filter/ReLU layers, then a max pooling layer, then some more filter/ReLU layers, then max pooling. Once the output is down to a relatively small size, there is typically a last fully-connected layer, leading into an activation function such as softmax that produces the final output. The exact design of these structures is an art—there is not currently any clear theoretical (or even systematic empirical) understanding of how these various design choices affect overall performance of the network.

The critical point for us is that this is all just a big neural network, which takes an input and computes an output. The mapping is a differentiable function of the weights, which means we can adjust the weights to decrease the loss by performing gradient descent, and we can compute the relevant gradients using back-propagation!

Well, the derivative is not continuous, both because of the ReLU and the max pooling operations, but we ignore that fact.

Let's work through a *very* simple example of how back-propagation can work on a convolutional network. The architecture is shown below. Assume we have a one-dimensional single-channel image, of size $n \times 1 \times 1$ and a single $k \times 1 \times 1$ filter in the first convolutional layer. Then we pass it through a ReLU layer and a fully-connected layer with no additional activation function on the output.



For simplicity assume k is odd, let the input image $X = A^0$, and assume we are using

squared loss. Then we can describe the forward pass as follows:

$$\begin{aligned} Z_i^1 &= W^{1T} \cdot A_{[i-\lfloor k/2 \rfloor : i+\lfloor k/2 \rfloor]}^0 \\ A^1 &= \text{ReLU}(Z^1) \\ A^2 &= W^{2T} A^1 \\ L(A^2, y) &= (A^2 - y)^2 \end{aligned}$$

Study Question: For a filter of size k , how much padding do we need to add to the top and bottom of the image?

How do we update the weights in filter W^1 ?

$$\frac{\partial \text{loss}}{\partial W^1} = \frac{\partial Z^1}{\partial W^1} \cdot \frac{\partial A^1}{\partial Z^1} \cdot \frac{\partial \text{loss}}{\partial A^1}$$

- $\partial Z^1 / \partial W^1$ is the $k \times n$ matrix such that $\partial Z_i^1 / \partial W_j^1 = X_{i-\lfloor k/2 \rfloor + j - 1}$. So, for example, if $i = 10$, which corresponds to column 10 in this matrix, which illustrates the dependence of pixel 10 of the output image on the weights, and if $k = 5$, then the elements in column 10 will be $X_8, X_9, X_{10}, X_{11}, X_{12}$.
- $\partial A^1 / \partial Z^1$ is the $n \times n$ diagonal matrix such that

$$\partial A_i^1 / \partial Z_i^1 = \begin{cases} 1 & \text{if } Z_i^1 > 0 \\ 0 & \text{otherwise} \end{cases}$$

- $\partial \text{loss} / \partial A^1 = \partial \text{loss} / \partial A^2 \cdot \partial A^2 / \partial A^1 = 2(A^2 - y)W^2$, an $n \times 1$ vector

Multiplying these components yields the desired gradient, of shape $k \times 1$.

CHAPTER 10

Sequential models

So far, we have limited our attention to domains in which each output y is assumed to have been generated as a function of an associated input x , and our hypotheses have been “pure” functions, in which the output depends only on the input (and the parameters we have learned that govern the function’s behavior). In the next few weeks, we are going to consider cases in which our models need to go beyond functions.

- In *recurrent neural networks*, the hypothesis that we learn is not a function of a single input, but of the whole sequence of inputs that the predictor has received.
- In *reinforcement learning*, the hypothesis is either a *model* of a domain (such as a game) as a recurrent system or a *policy* which is a pure function, but whose loss is determined by the ways in which the policy interacts with the domain over time.

Before we engage with those forms of learning, we will study models of sequential or recurrent systems that underlie the learning methods.

1 State machines

A *state machine* is a description of a process (computational, physical, economic) in terms of its potential sequences of *states*.

The *state* of a system is defined to be all you would need to now about the system to predict its future trajectories as well as possible. It could be the position and velocity of an object or the locations of your pieces on a board game, or the current traffic densities on a highway network.

Formally, we define a *state machine* as $(\mathcal{S}, \mathcal{X}, \mathcal{Y}, s_0, f, g)$ where

- \mathcal{S} is a finite or infinite set of possible states;
- \mathcal{X} is a finite or infinite set of possible inputs;
- \mathcal{Y} is a finite or infinite set of possible outputs;
- $s_0 \in \mathcal{S}$ is the initial state of the machine;
- $f : \mathcal{S} \times \mathcal{X} \rightarrow \mathcal{S}$ is a *transition function*, which takes an input and a previous state and produces a next state;

This is such a pervasive idea that it has been given many names in many subareas of computer science, control theory, physics, etc., including: *automaton*, *transducer*, *dynamical system*, *system*, etc.

There are a huge number of major and minor variations on the idea of a state machine. We’ll just work with one specific one in this section and another one in the next, but don’t worry if you see other variations out in the world!

- $g : \mathcal{S} \rightarrow \mathcal{Y}$ is an *output function*, which takes a state and produces an output.


The basic operation of the state machine is to start with state s_0 , then iteratively compute:

In some cases, we will pick a starting state.

$$s_t = f(s_{t-1}, x_t)$$

$$y_t = g(s_t)$$

The diagram below illustrates this process. Note that the “feedback” connection of s_t back into f has to be buffered or delayed by one time step—otherwise what it is computing would not generally be well defined.



So, given a sequence of inputs x_1, x_2, \dots the machine generates a sequence of outputs


$$\underbrace{g(f(x_1, s_0))}_{y_1}, \underbrace{g(f(x_2, f(x_1, s_0)))}_{y_2}, \dots$$

We sometimes say that the machine *transduces* sequence x into sequence y . The output at time t can have dependence on inputs from steps 1 to t .

One common form is *finite state machines*, in which $\mathcal{S}, \mathcal{X}, \mathcal{Y}$ are all finite sets. They are often described using *state transition diagrams* such as the one below, in which nodes stand for states and arcs indicate transitions. Nodes are labeled by which output they generate and arcs are labeled by which input causes the transition.

All computers can be described, at the digital level, as finite state machines. Big, but finite!

One can verify that the state machine below reads binary strings and determines the parity of the number of zeros in the given string. Check for yourself that all inputted binary strings end in state S_1 if and only if they contain an even number of zeros.



Another common structure that is simple but powerful and used in signal processing and control is *linear time-invariant (LTI) systems*. In this case, $\mathcal{S} = \mathbb{R}^m$, $\mathcal{X} = \mathbb{R}^1$ and $\mathcal{Y} = \mathbb{R}^n$, and f and g are linear functions of their inputs. In discrete time, they can be defined by a linear difference equation, like

$$y[t] = 3y[t-1] + 6y[t-2] + 5x[t] + 3x[t-2],$$

and can be implemented using state to store relevant previous input and output information.

We will study *recurrent neural networks* which are a lot like a non-linear version of an LTI system, with transition and output functions

$$f(s, x) = f_1(W^{sx}x + W^{ss}s + W_0^{ss})$$

$$g(s) = f_2(W^0s + W_0^0)$$

defined by weight matrices

$$W^{sx} : m \times \ell$$

$$W^{ss} : m \times m$$

$$W_0^{ss} : m \times 1$$

$$W^0 : n \times m$$

$$W_0^0 : n \times 1$$

and activation functions f_1 and f_2 . We will see that it's actually possible to learn weight values using gradient descent.

2 Markov decision processes

A *Markov decision process* (MDP) is a variation on a state machine in which:

- The transition function is *stochastic*, meaning that it defines a probability distribution over the next state given the previous state and input, but each time it is evaluated it draws a new state from that distribution.
- The output is equal to the state (that is g is the identity function).
- Some states (or state-action pairs) are more desirable than others.

An MDP can be used to model interaction with an outside "world," such as a single-player game.

We will focus on the case in which \mathcal{S} and \mathcal{X} are finite, and will call the input set \mathcal{A} for actions (rather than \mathcal{X}). The idea is that an agent (a robot or a game-player) can model its environment as an MDP and try to choose actions that will drive the process into states that have high scores.

Formally, an MDP is $\langle \mathcal{S}, \mathcal{A}, T, R, \gamma \rangle$ where:

- $T : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow \mathbb{R}$ is a *transition model*, where $T(s, a, s') = P(S_t = s' | S_{t-1} = s, A_{t-1} = a)$, specifying a conditional probability distribution;
- $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a *reward function*, where $R(s, a)$ specifies how desirable it is to be in state s and take action a ; and
- $\gamma \in [0, 1]$ is an optimal *discount factor*.

A *policy* is a function $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that specifies what action to take in each state.

2.1 Finite-horizon solutions

Given an MDP our goal is typically to find a policy that is optimal in the sense that it gets as much total reward as possible, in expectation over the stochastic transitions that the domain makes. In this section, we will consider the case where there is a finite *horizon* H , indicating the total number of steps of interaction that the agent will have with the MDP.

Recall that stochastic is another word for *probabilistic*; we don't say "random" because that can be interpreted in two ways, both of which are incorrect. We don't pick the transition function itself at random from a distribution. The transition function doesn't pick its output *uniformly* at random.

There is an interesting variation on MDPs, called a *partially observable* MDP, in which the output is also drawn from a distribution depending on the state.

And there is an interesting, direct extension to two-player zero-sum games, such as Chess and Go.

The notation here uses capital letters, like S to stand for random variables and small letters to stand for concrete values. So S_t here is a random variable that can take on elements of \mathcal{S} as values.

2.1.1 Evaluating a given policy

Before we can talk about how to find a good policy, we have to specify a measure of the goodness of a policy. We will do so by defining the horizon h value of a state given a policy, $V_{\pi}^h(s)$ for an MDP. We do this by induction on the horizon, which is the *number of steps left to go*.

The base case is when there are no steps remaining, in which case, no matter what state we're in, the value is 0, so

$$V_{\pi}^0(s) = 0 .$$

Then, the value of a policy in state s at horizon $h + 1$ is equal to the reward it will get in state s plus the expected horizon h value of the next state. So, starting with horizons 1 and 2, and then moving to the general case, we have:

$$\begin{aligned} V_{\pi}^1(s) &= R(s, \pi(s)) + 0 \\ V_{\pi}^2(s) &= R(s, \pi(s)) + \sum_{s'} T(s, \pi(s), s') \cdot R(s', \pi(s')) \\ &\vdots \\ V_{\pi}^h(s) &= R(s, \pi(s)) + \sum_{s'} T(s, \pi(s), s') \cdot V_{\pi}^{h-1}(s') \end{aligned}$$

The sum over s' is an *expected value*: it considers all possible next states s' , and computes an average of their $(h - 1)$ -horizon values, weighted by the probability that the transition function from state s with the action chosen by the policy, $\pi(s)$, assigns to arriving in state s' .

Study Question: What is $\sum_{s'} T(s, a, s')$ for any particular s and a ?

Then we can say that one policy is better than another one for horizon h , $\pi_1 >_h \pi_2$, if and only if for all $s \in \mathcal{S}$, $V_{\pi_1}^h(s) \geq V_{\pi_2}^h(s)$ and there exists at least one $s \in \mathcal{S}$ such that $V_{\pi_1}^h(s) > V_{\pi_2}^h(s)$.

2.1.2 Finding an optimal policy

How can we go about finding an optimal policy for an MDP? We could imagine enumerating all possible policies and calculating their value functions as in the previous section and picking the best one...but that's too much work!

The first observation to make is that, in a finite-horizon problem, the best action to take depends on the current state, but also on the horizon: imagine that you are in a situation where you could reach a state with reward 5 in one step or a state with reward 10 in two steps. If you have at least two steps to go, then you'd move toward the reward-10 state, but if you only have step left to go, you should go in the direction that will allow you to gain 5!

One way to find an optimal policy is to compute an *optimal action-value function*, Q . We define $Q^h(s, a)$ to be the expected value of

- starting in state s
- executing action a
- continuing for $h - 1$ more steps executing an optimal policy for the appropriate horizon on each step.

Similar to our definition of V for evaluating a policy, we define the Q function recursively according to the horizon. The only difference is that, on each step with horizon h , rather than selecting an action specified by a given policy, we select the value of a that will maximize the expected Q^h value of the next state.

$$\begin{aligned} Q^0(s, a) &= 0 \\ Q^1(s, a) &= R(s, a) + 0 \\ Q^2(s, a) &= R(s, a) + \sum_{s'} T(s, a, s') \max_{a'} R(s', a') \\ &\vdots \\ Q^h(s, a) &= R(s, a) + \sum_{s'} T(s, a, s') \max_{a'} Q^{h-1}(s', a') \end{aligned}$$

We can solve for the values of Q with a simple recursive algorithm called *value iteration* which just computes Q^h starting from horizon 0 and working backward to the desired horizon H . Given Q , an optimal policy is easy to find:

$$\pi_h^*(s) = \arg \max_a Q^h(s, a) .$$

There may be multiple possible optimal policies.

2.2 Infinite-horizon solutions

It is actually more typical to work in a regime where the actual finite horizon is not known. This is called the *infinite horizon* version of the problem, when you don't know when the game will be over! However, if we tried to simply take our definition of Q^h above and set $h = \infty$, we would be in trouble, because it could well be that the Q^∞ values for all actions would be infinite, and there would be no way to select one over the other.

There are two standard ways to deal with this problem. One is to take a kind of *average* over all time steps, but this can be a little bit tricky to think about. We'll take a different approach, which is to consider the *discounted* infinite horizon. We select a discount factor $0 < \gamma < 1$. Instead of trying to find a policy that maximizes expected finite-horizon undiscounted value,

$$\mathbb{E} \left[\sum_{t=0}^h R_t \mid \pi, s_0 \right]$$

we will try to find one that maximizes expected *infinite horizon discounted value*, which is

$$\mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t R_t \mid \pi, s_0 \right] = \mathbb{E} [R_0 + \gamma R_1 + \gamma^2 R_2 + \dots \mid \pi, s_0] .$$

Note that the t indices here are not the number of steps to go, but actually the number of steps forward from the starting state (there is no sensible notion of "steps to go" in the infinite horizon case).

There are two good intuitive motivations for discounting. One is related to economic theory and the present value of money: you'd generally rather have some money today than that same amount of money next week (because you could use it now or invest it). The other is to think of the whole process terminating, with probability $1 - \gamma$ on each step of the interaction. This value is the expected amount of reward the agent would gain under this terminating model.

2.2.1 Evaluating a policy

We will start, again, by evaluating a policy, but now in terms of the expected discounted infinite-horizon value that the agent will get in the MDP if it executes that policy. We define the value of a state s under policy π as

$$V_{\pi}(s) = \mathbb{E}[R_0 + \gamma R_1 + \gamma^2 R_2 + \dots \mid \pi, S_0 = s] = \mathbb{E}[R_0 + \gamma(R_1 + \gamma(R_2 + \gamma \dots))] \mid \pi, S_0 = s] .$$

Because the expectation of a linear combination of random variables is the linear combination of the expectations, we have

$$\begin{aligned} V_{\pi}(s) &= \mathbb{E}[R_0 \mid \pi, S_0 = s] + \gamma \mathbb{E}[R_1 + \gamma(R_2 + \gamma \dots)] \mid \pi, S_0 = s] \\ &= R(s, \pi(s)) + \gamma \sum_{s'} T(s, \pi(s), s') V_{\pi}(s') \end{aligned}$$

You could write down one of these equations for each of the $n = |\mathcal{S}|$ states. There are n unknowns $V_{\pi}(s)$. These are linear equations, and so it's easy to solve them using Gaussian elimination to find the value of each state under this policy.

2.2.2 Finding an optimal policy

The best way of behaving in an infinite-horizon discounted MDP is not time-dependent: at every step, your expected future lifetime, given that you have survived until now, is $1/(1 - \gamma)$.

An important theorem about MDPs is: there exists a stationary optimal policy π^* (there may be more than one) such that for all $s \in \mathcal{S}$ and all other policies π we have

$$V_{\pi^*}(s) \geq V_{\pi}(s) .$$

There are many methods for finding an optimal policy for an MDP. We will study a very popular and useful method called *value iteration*. It is also important to us, because it is the basis of many *reinforcement-learning* methods.

Define $Q^*(s, a)$ to be the expected infinite horizon discounted value of being in state s , executing action a , and executing an optimal policy π^* thereafter. Using similar reasoning to the recursive definition of V_{π} , we can define

$$Q^*(s, a) = R(s, a) + \gamma \sum_{s'} T(s, a, s') \max_{a'} Q^*(s', a') .$$

This is also a set of equations, one for each (s, a) pair. This time, though, they are not linear, and so they are not easy to solve. But there is a theorem that says they have a unique solution!

If we knew the optimal value function, then we could derive the optimal policy π^* as

$$\pi^*(s) = \arg \max_a Q^*(s, a) .$$

We can iteratively solve for the Q^* values with the value iteration algorithm, shown below:

VALUE-ITERATION($\mathcal{S}, \mathcal{A}, T, R, \gamma, \epsilon$)

```

1  for  $s \in \mathcal{S}, a \in \mathcal{A}$  :
2       $Q_{old}(s, a) = 0$ 
3  while True:
4      for  $s \in \mathcal{S}, a \in \mathcal{A}$  :
5           $Q_{new}(s, a) = R(s, a) + \gamma \sum_{s'} T(s, a, s') \max_{a'} Q_{old}(s', a')$ 
6      if  $\max_{s, a} |Q_{old}(s, a) - Q_{new}(s, a)| < \epsilon$  :
7          return  $Q_{new}$ 
8       $Q_{old} := Q_{new}$ 

```

This is *so* cool! In a discounted model, if you find that you survived this round and landed in some state s' , then you have the same expected future lifetime as you did before. So the value function that is relevant in that state is exactly the same one as in state s .

Stationary means that it doesn't change over time; the optimal policy in a finite-horizon MDP is *non-stationary*.

2.2.3 Theory

There are a lot of nice theoretical results about value iteration. For some given (not necessarily optimal) Q function, define $\pi_Q(s) = \arg \max_a Q(s, a)$.

- After executing value iteration with parameter ϵ , $\|V_{\pi_{Q_{\text{new}}}} - V_{\pi^*}\|_{\max} < \epsilon$.

- There is a value of ϵ such that

$$\|Q_{\text{old}} - Q_{\text{new}}\|_{\max} < \epsilon \implies \pi_{Q_{\text{new}}} = \pi^*$$

- As the algorithm executes, $\|V_{\pi_{Q_{\text{new}}}} - V_{\pi^*}\|_{\max}$ decreases monotonically on each iteration.
- The algorithm can be executed asynchronously, in parallel: as long as all (s, a) pairs are updated infinitely often in an infinite run, it still converges to optimal value.

This is new notation! Given two functions f and f' , we write $\|f - f'\|_{\max}$ to mean $\max_x |f(x) - f'(x)|$. It measures the maximum absolute point of disagreement between the two functions.

This is very important for reinforcement learning.


CHAPTER 11

Reinforcement learning

So far, all the learning problems we have looked at have been *supervised*: that is, for each training input $x^{(i)}$, we are told which value $y^{(i)}$ should be the output. A very different problem setting is *reinforcement learning*, in which the learning system is not directly told which outputs go with which inputs. Instead, there is an interaction of the form:

- Learner observes *input* $s^{(i)}$
- Learner generates *output* $a^{(i)}$
- Learner observes *reward* $r^{(i)}$
- Learner observes *input* $s^{(i+1)}$
- Learner generates *output* $a^{(i+1)}$
- Learner observes *reward* $r^{(i+1)}$
- ...

The learner is supposed to find a *policy*, mapping s to a , that maximizes expected reward over time.



This problem setting is equivalent to an *online* supervised learning under the following assumptions:

1. The space of possible outputs is binary (e.g. $\{+1, -1\}$) and the space of possible rewards is binary (e.g. $\{+1, -1\}$);
2. $s^{(i)}$ is independent of all previous $s^{(j)}$ and $a^{(j)}$; and
3. $r^{(i)}$ depends only on $s^{(i)}$ and $a^{(i)}$.

In this case, for any experience tuple $(s^{(i)}, a^{(i)}, r^{(i)})$, we can generate a supervised training example, which is equal to $(s^{(i)}, a^{(i)})$ if $r^{(i)} = +1$ and $(s^{(i)}, -a^{(i)})$ otherwise.

Study Question: What supervised-learning loss function would this objective correspond to?

Reinforcement learning is more interesting when these properties do not hold. When we relax assumption 1 above, we have the class of *bandit problems*, which we will discuss in section 1. If we relax assumption 2, but assume that the environment that the agent is interacting with is an MDP, so that $s^{(i)}$ depends only on $s^{(i-1)}$ and $a^{(i-1)}$ then we are in the classical *reinforcement-learning* setting, which we discuss in section 2. Weakening the assumptions further, for instance, not allowing the learner to observe the current state completely and correctly, makes the problem substantially more difficult, and beyond the scope of this class.

1 Bandit problems

A basic bandit problem is given by

- A set of actions \mathcal{A} ;
- A set of reward values \mathcal{R} ; and
- A probabilistic reward function $R : \mathcal{A} \rightarrow \mathbb{R}$ where $R(a) = P(R | A = a)$ is a probability distribution over possible reward values in \mathcal{R} conditioned on which action is selected.

The most typical bandit problem has $\mathcal{R} = \{0, 1\}$ and $|\mathcal{A}| = k$. This is called a *k-armed bandit problem*. There is a lot of mathematical literature on optimal strategies for k-armed bandit problems under various assumptions. The important question is usually one of *exploration versus exploitation*. Imagine that you have tried each action 10 times, and now you have an estimate \hat{p}_j for the expected value of $R(a_j)$. Which arm should you pick next? You could

exploit your knowledge, and choose the arm with the highest value of \hat{p}_j on all future trials; or

explore further, by trying some or all actions more times, hoping to get better estimates of the p_j values.

The theory ultimately tells us that, the longer our horizon H (or, similarly, closer to 1 our discount factor), the more time we should spend exploring, so that we don't converge prematurely on a bad choice of action.

Study Question: Why is it that "bad" luck during exploration is more dangerous than "good" luck? Imagine that there is an action that generates reward value 1 with probability 0.9, but the first three times you try it, it generates value 0. How might that cause difficulty? Why is this more dangerous than the situation when an action that generates reward value 1 with probability 0.1 actually generates reward 1 on the first three tries?

Note that what makes this a very different kind of problem from the batch supervised learning setting is that:

- The agent gets to influence what data it gets (selecting a_j gives it another sample from r_j), and
- The agent is penalized for mistakes it makes while it is learning (if it is trying to maximize the expected sum of r_t it gets while behaving).

In a *contextual* bandit problem, you have multiple possible states, drawn from some set \mathcal{S} , and a separate bandit problem associated with each one.

Bandit problems will be an essential sub-component of reinforcement learning.

Why? Because in English slang, "one-armed bandit" is a name for a slot machine (an old-style gambling machine where you put a coin into a slot and then pull its arm to see if you get a payoff.) because it has one arm and takes your money! What we have here is a similar sort of machine, but with k arms.

There is a setting of supervised learning, called *active learning*, where instead of being given a training set, the learner gets to select values of x and the environment gives back a label y ; the problem of picking good x values to query is interesting, but the problem of deriving a hypothesis from (x, y) pairs is the same as the supervised problem we have been studying.

2 Sequential problems

In the more typical (and difficult!) case, we can think of our learning agent interacting with an MDP, where it knows \mathcal{S} and \mathcal{A} , but not $T(s, a, s')$ or $R(s, a)$. The learner can interact with the environment by selecting actions. So, this is somewhat like a contextual bandit problem, but more complicated, because selecting an action influences not only what the immediate reward will be, but also what state the system ends up in at the next time step and, therefore, what additional rewards might be available in the future.

A *reinforcement-learning RL algorithm* is a kind of a policy that depends on the whole history of states, actions, and rewards and selects the next action to take. There are several different ways to measure the quality of an RL algorithm, including:

- Ignoring the r_t values that it gets *while* learning, but consider how many interactions with the environment are required for it to learn a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that is nearly optimal.
- Maximizing the expected discounted sum of total rewards while it is learning.

Most of the focus is on the first criterion, because the second one is very difficult. The first criterion is reasonable when the learning can take place somewhere safe (imagine a robot learning, inside the robot factory, where it can't hurt itself too badly) or in a simulated environment.

Approaches to reinforcement-learning differ significantly according to what kind of hypothesis or model they learn. In the following sections, we will consider several different approaches.

2.1 Model-based RL

The conceptually simplest approach to RL is to estimate R and T from the data we have gotten so far, and then use those estimates, together with an algorithm for solving MDPs (such as value iteration) to find a policy that is near-optimal given the current model estimates.

Assume that we have had some set of interactions with the environment, which can be characterized as a set of tuples of the form $(s^{(t)}, a^{(t)}, r^{(t)}, s^{(t+1)})$.

We can estimate $T(s, a, s')$ using a simple counting strategy,

$$\hat{T}(s, a, s') = \frac{\#(s, a, s') + 1}{\#(s, a) + |\mathcal{S}|}.$$

Here, $\#(s, a, s')$ represents the number of times in our data set we have the situation where $s_t = s, a_t = a, s_{t+1} = s'$ and $\#(s, a)$ represents the number of times in our data set we have the situation where $s_t = s, a_t = a$.

Study Question: Prove to yourself that $\#(s, a) = \sum_{s'} \#(s, a, s')$.

Adding 1 and $|\mathcal{S}|$ to the numerator and denominator, respectively, are a form of smoothing called the *Laplace correction*. It ensures that we never estimate that a probability is 0, and keeps us from dividing by 0. As the amount of data we gather increases, the influence of this correction fades away.

We also estimate the reward function $R(s, a)$:

$$\hat{R}(s, a) = \frac{\sum r \mid s, a}{\#(s, a)}$$

where

$$\sum r \mid s, a = \sum_{\{t \mid s_t = s, a_t = a\}} r^{(t)}.$$

This is just the average of the observed rewards for each s, a pair.

We can now solve the MDP $(\mathcal{S}, \mathcal{A}, \hat{T}, \hat{R})$ to find an optimal policy using value iteration, or use a finite-depth expecti-max search to find an action to take for a particular state.

This technique is effective for problems with small state and action spaces, where it is not too hard to get enough experience to estimate T and R well but it is difficult to generalize to continuous (or very large discrete) state spaces, and is a topic of current research.

2.2 Policy search

A very different strategy is to search directly for a good policy, without estimating the model. The strategy here is to define a functional form $f(s; \theta) = a$ for the policy, where θ represents the parameters we learn from experience. We choose f to be differentiable, and often let $f(s; \theta) = P(a)$, a probability distribution over our possible actions.

Now, we can train the policy parameters using gradient descent:

- When θ has relatively low dimension, we can compute a numeric estimate of the gradient by running the policy multiple times for $\theta \pm \epsilon$, and computing the resulting rewards.
- When θ has higher dimensions (e.g., it is a complicated neural network), there are more clever algorithms, e.g., one called REINFORCE, but they can often be difficult to get to work reliably.

Policy search is a good choice when the policy has a simple known form, but the model would be much more complicated to estimate.

2.3 Value function learning

The most popular class of algorithms learns neither explicit transition and reward models nor a direct policy, but instead concentrates on learning a value function. It is a topic of current research to describe exactly under what circumstances value-function-based approaches are best, and there are a growing number of methods that combine value functions, transition and reward models and policies into a complex learning algorithm in an attempt to combine the strengths of each approach.

We will study two variations on value-function learning, both of which estimate the Q function.

2.3.1 Q-learning

This is the most typical way of performing reinforcement learning. Recall the value-iteration update:

$$Q(s, a) = R(s, a) + \gamma \sum_{s'} T(s, a, s') \max_{a'} Q(s', a')$$

We will adapt this update to the RL scenario, where we do not know the transition function T or reward function R .

The thing that most students seem to get confused about is when we do value iteration and when we do Q learning. Value iteration assumes you know T and R and just need to *compute* Q . In Q learning, we don't know or even directly estimate T and R : we estimate Q directly from experience!

Q-LEARNING($\mathcal{S}, \mathcal{A}, s_0, \gamma, \alpha$)

```

1  for  $s \in \mathcal{S}, a \in \mathcal{A}$  :
2       $Q[s, a] = 0$ 
3   $s = s_0$  // Or draw an  $s$  randomly from  $\mathcal{S}$ 
4  while True:
5       $a = \text{select\_action}(s, Q)$ 
6       $r, s' = \text{execute}(a)$ 
7       $Q[s, a] = (1 - \alpha)Q[s, a] + \alpha(r + \gamma \max_{a'} Q[s', a'])$ 
8       $s = s'$ 

```

Here, α represents the “learning rate,” which needs to decay for convergence purposes, but in practice is often set to a constant.

Note that the update can be rewritten as

$$Q[s, a] = Q[s, a] - \alpha \left(Q[s, a] - (r + \gamma \max_{a'} Q[s', a']) \right),$$

which looks something like a gradient update! This is often called *temporal difference learning* method, because we make an update based on the difference between the current estimated value of taking action a in state s , which is $Q[s, a]$, and the “one-step” sampled value of taking a in s , which is $r + \gamma \max_{a'} Q[s', a']$.

It is actually not a gradient update, but later, when we consider function approximation, we will treat it as if it were.

You can see this method as a combination of two different iterative processes that we have already seen: the combination of an old estimate with a new sample using a running average with a learning rate α , and the dynamic-programming update of a Q value from value iteration.


Our algorithm above includes a procedure called *select_action*, which, given the current state s , has to decide which action to take. If the Q value is estimated very accurately and the agent is behaving in the world, then generally we would want to choose the apparently optimal action $\arg \max_{a \in \mathcal{A}} Q(s, a)$. But, during learning, the Q value estimates won't be very good and exploration is important. However, exploring completely at random is also usually not the best strategy while learning, because it is good to focus your attention on the parts of the state space that are likely to be visited when executing a good policy (not a stupid one).

A typical action-selection strategy is the ϵ -greedy strategy:

- with probability $1 - \epsilon$, choose $\arg \max_{a \in \mathcal{A}} Q(s, a)$
- with probability ϵ , choose the action $a \in \mathcal{A}$ uniformly at random

Q-learning has the surprising property that it is *guaranteed* to converge to the actual optimal Q function under fairly weak conditions! Any exploration strategy is okay as long as it tries every action infinitely often on an infinite run (so that it doesn't converge prematurely to a bad action choice).

Q-learning can be very sample-inefficient: imagine a robot that has a choice between moving to the left and getting a reward of 1, then returning to its initial state, or moving to the right and walking down a 10-step hallway in order to get a reward of 1000, then returning to its initial state.



The first time the robot moves to the right and goes down the hallway, it will update the Q value for the last state on the hallway to have a high value, but it won't yet understand that moving to the right was a good choice. The next time it moves down the hallway it updates the value of the state before the last one, and so on. After 10 trips down the hallway, it now can see that it is better to move to the right than to the left.

More concretely, consider the vector of Q values $Q(0 : 10, \text{right})$, representing the Q values for moving right at each of the positions $0, \dots, 9$. Then, for $\alpha = 1$ and $\gamma = 0.9$,

$$Q(i, \text{right}) = R(i, \text{right}) + 0.9 \cdot \max_a Q(i + 1, a)$$

Starting with Q values of 0,

$$Q^{(0)}(0 : 10, \text{right}) = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0]$$

Since the only nonzero reward from moving right is $R(9, \text{right}) = 1000$, after our robot makes it down the hallway once, our new Q vector is

$$Q^{(1)}(0 : 10, \text{right}) = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1000 \ 0]$$

After making its way down the hallway again, $Q(8, \text{right}) = 0 + 0.9 \cdot Q(9, \text{right}) = 900$ updates:

$$Q^{(2)}(0 : 10, \text{right}) = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 900 \ 1000 \ 0]$$

Similarly,

$$Q^{(3)}(0 : 10, \text{right}) = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 810 \ 900 \ 1000 \ 0]$$

$$Q^{(4)}(0 : 10, \text{right}) = [0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 729 \ 810 \ 900 \ 1000 \ 0]$$

\vdots

$$Q^{(10)}(0 : 10, \text{right}) = [387.4 \ 420.5 \ 478.3 \ 531.4 \ 590.5 \ 656.1 \ 729 \ 810 \ 900 \ 1000 \ 0],$$

and the robot finally sees the value of moving right from position 0.

Study Question: Determine the Q value functions that will result from updates due to the robot always executing the "move left" policy.

2.3.2 Function approximation

In our Q-learning algorithm above, we essentially keep track of each Q value in a table, indexed by s and a . What do we do if S and/or A are large (or continuous)?

We can use a function approximator like a neural network to store Q values. For example, we could design a neural network that takes in inputs s and a , and outputs $Q(s, a)$. We can treat this as a regression problem, optimizing the squared Bellman error, with loss:

$$\left(Q(s, a) - (r + \gamma \max_{a'} Q(s', a')) \right)^2,$$

where $Q(s, a)$ is now the output of the neural network.

There are actually several different architectural choices for using a neural network to approximate Q values:

- One network for each action a_j , that takes s as input and produces $Q(s, a_j)$ as output;
- One single network that takes s as input and produces a vector $Q(s, \cdot)$, consisting of the Q values for each action; or

We are violating our usual notational conventions here, and writing $Q^{(i)}$ to mean the Q value function that results after the robot runs all the way to the end of the hallway, when executing the policy that always moves to the right.

We can see how this interacts with the exploration/exploitation dilemma: from the perspective of s_0 , it will seem, for a long time, that getting the immediate reward of 1 is a better idea, and it would be easy to converge on that as a strategy without exploring the long hallway sufficiently.

- One single network that takes s, a concatenated into a vector (if a is discrete, we would probably use a one-hot encoding, unless it had some useful internal structure) and produces $Q(s, a)$ as output.

The first two choices are only suitable for discrete (and not too big) action sets. The last choice can be applied for continuous actions, but then it is difficult to find $\arg \max_{a'} Q(s, a')$.

There are not many theoretical guarantees about Q-learning with function approximation and, indeed, it can sometimes be fairly unstable (learning to perform well for a while, and then getting suddenly worse, for example). But it has also had some significant successes.

One form of instability that we do know how to guard against is *catastrophic forgetting*. In standard supervised learning, we expect that the training x values were drawn independently from some distribution. But when a learning agent, such as a robot, is moving through an environment, the sequence of states it encounters will be temporally correlated. This can mean that while it is in the dark, the neural-network weight-updates will make the Q function “forget” the value function for when it’s light. We can handle this with *experience replay*, where we save our (s, a, r, s') experiences in a buffer and mix into our learning process some updates from the historic buffer.

For continuous action spaces, it is increasingly popular to use a class of methods called *actor-critic* methods, which combine policy and value-function learning. We won't get into them in detail here, though.

And, in fact, we routinely shuffle their order in the data file, anyway.

For example, it might spend 12 hours in a dark environment and then 12 in a light one.

2.3.3 Fitted Q-learning

An alternative strategy for learning the Q function that is somewhat more robust than the standard Q-learning algorithm is a method called *fitted Q*.

FITTED-Q-LEARNING($\mathcal{A}, s_0, \gamma, \alpha, \epsilon, m$)

```

1   $s = s_0$  // Or draw an  $s$  randomly from  $\mathcal{S}$ 
2   $\mathcal{D} = \{ \}$ 
3  initialize neural-network representation of Q
4  while True:
5       $\mathcal{D}_{\text{new}}$  = experience from executing  $\epsilon$ -greedy policy based on Q for  $m$  steps
6       $\mathcal{D} = \mathcal{D} \cup \mathcal{D}_{\text{new}}$  represented as  $(s, a, r, s')$  tuples
7       $\mathcal{D}_{\text{sup}} = \{ (x^{(i)}, y^{(i)}) \}$  where  $x^{(i)} = (s, a)$  and  $y^{(i)} = r + \gamma \max_{a' \in \mathcal{A}} Q(s', a')$ 
8          for each tuple  $(s, a, r, s')^{(i)} \in \mathcal{D}$ 
9      re-initialize neural-network representation of Q
10      $Q = \text{supervised\_NN\_regression}(\mathcal{D}_{\text{sup}})$ 
```

Here, we alternate between using the policy induced by the current Q function to gather a batch of data \mathcal{D}_{new} , adding it to our overall data set \mathcal{D} , and then using supervised neural-network training to learn a representation of the Q value function on the whole data set. This method does not mix the dynamic-programming phase (computing new Q values based on old ones) with the function approximation phase (training the neural network) and avoids catastrophic forgetting. The regression training in line 9 typically uses squared error as a loss function and would be trained until the fit is good (possibly measured on held-out data).

Recurrent Neural Networks

In chapter 8 we studied neural networks and how we can train the weights of a network, based on data, so that it will adapt into a function that approximates the relationship between the (x, y) pairs in a supervised-learning training set. In section 1 of chapter 10, we studied state-machine models and defined *recurrent neural networks* (RNNs) as a particular type of state machine, with a multidimensional vector of real values as the state. In this chapter, we'll see how to use gradient-descent methods to train the weights of an RNN so that it performs a *transduction* that matches as closely as possible a training set of input-output sequences.


1 RNN model

Recall that the basic operation of the state machine is to start with some state s_0 , then iteratively compute:

$$s_t = f(s_{t-1}, x_t)$$

$$y_t = g(s_t)$$

as illustrated in the diagram below (remembering that there needs to be a delay on the feedback loop):



So, given a sequence of inputs x_1, x_2, \dots the machine generates a sequence of outputs

$$\underbrace{g(f(x_1, s_0))}_{y_1}, \underbrace{g(f(x_2, f(x_1, s_0)))}_{y_2}, \dots$$

A recurrent neural network is a state machine with neural networks constituting functions f and g :

$$\begin{aligned} f(s, x) &= f_1(W^{sx}x + W^{ss}s + W_0^{ss}) \\ g(s) &= f_2(W^O s + W_0^O) \end{aligned}$$

. The inputs, outputs, and states are all vector-valued:

$$\begin{aligned} x_t &: \ell \times 1 \\ s_t &: m \times 1 \\ y_t &: v \times 1 \end{aligned}$$

The weights in the network, then, are

$$\begin{aligned} W^{sx} &: m \times \ell \\ W^{ss} &: m \times m \\ W_0^{ss} &: m \times 1 \\ W^O &: v \times m \\ W_0^O &: v \times 1 \end{aligned}$$

with activation functions f_1 and f_2 . Finally, the operation of the RNN is described by

$$\begin{aligned} s_t &= f_1(W^{sx}x_t + W^{ss}s_{t-1} + W_0) \\ y_t &= f_2(W^O s_t + W_0^O) \end{aligned}$$

Study Question: Check dimensions here to be sure it all works out. Remember that we apply f_1 and f_2 elementwise.

We are very sorry! This course material has evolved from different sources, which used $W^T x$ in the forward pass for regular feed-forward NNs and Wx for the forward pass in RNNs. I only just noticed this inconsistency and it's too late to change for this term. Of course it doesn't make any technical difference, but is a potential source of confusion.

2 Sequence-to-sequence RNN

Now, how can we train an RNN to model a transduction on sequences? This problem is sometimes called *sequence-to-sequence* mapping. You can think of it as a kind of regression problem: given an input sequence, learn to generate the corresponding output sequence.

A training set has the form $[(x^{(1)}, y^{(1)}), \dots, (x^{(q)}, y^{(q)})]$, where

- $x^{(1)}$ and $y^{(1)}$ are length $n^{(q)}$ sequences;
- sequences in the *same pair* are the same length; and sequences in different pairs may have different lengths.

Next, we need a loss function. We start by defining a loss function on sequences. There are many possible choices, but usually it makes sense just to sum up a per-element loss function on each of the output values:

$$\text{Loss}_{\text{seq}}(p^{(i)}, y^{(i)}) = \sum_{j=1}^{n^{(q)}} \text{Loss}_{\text{elt}}(p_j^{(i)}, y_j^{(i)})$$

The per-element loss function Loss_{elt} will depend on the type of y_t and what information it is encoding, in the same way as for a supervised network.. Then, letting $\theta =$

One way to think of training a sequence classifier is to reduce it to a transduction problem, where $y_t = 1$ if the sequence x_1, \dots, x_t is a positive example of the class of sequences and -1 otherwise.

So it could be NLL, hinge loss, squared loss, etc.

($W^{sx}, W^{ss}, W^O, W_0, W_0^O$), our overall objective is to minimize

$$J(\theta) = \sum_{i=1}^q \text{Loss}_{\text{seq}} \left(\text{RNN}(x^{(i)}; \theta), y^{(i)} \right) ,$$

where $\text{RNN}(x; \theta)$ is the output sequence generated, given input sequence x .

It is typical to choose f_1 to be *tanh* but any non-linear activation function is usable. We choose f_2 to align with the types of our outputs and the loss function, just as we would do in regular supervised learning.

Remember that it looks like a sigmoid but ranges from -1 to +1.

3 Back-propagation through time

Now the fun begins! We can find θ to minimize J using gradient descent. We will work through the simplest method, *back-propagation through time* (BPTT), in detail. This is generally not the best method to use, but it's relatively easy to understand. In section 5 we will sketch alternative methods that are in much more common use.

Calculus reminder Most of us are not very careful about the difference between the *partial derivative* and the *total derivative*. We are going to use a nice example from the Wikipedia article on partial derivatives to illustrate the difference.

The volume of a cone depends on its height and radius:

$$V(r, h) = \frac{\pi r^2 h}{3} .$$

The partial derivatives of volume with respect to height and radius are

$$\frac{\partial V}{\partial r} = \frac{2\pi r h}{3} \quad \text{and} \quad \frac{\partial V}{\partial h} = \frac{\pi r^2}{3} .$$

They measure the change in V assuming everything is held constant except the single variable we are changing. But! in a cone, the radius and height are not independent, and so we can't really change one without changing the other. In this case, we really have to think about the *total derivative*, which sums the "paths" along which r might influence V :

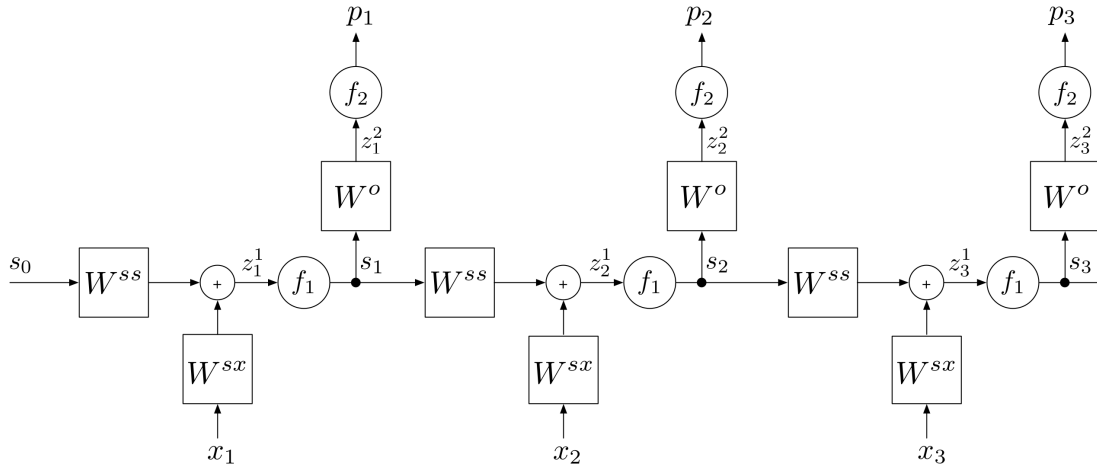
$$\begin{aligned} \frac{dV}{dr} &= \frac{\partial V}{\partial r} + \frac{\partial V}{\partial h} \frac{dh}{dr} \\ &= \frac{2\pi r h}{3} + \frac{\pi r^2}{3} \frac{dh}{dr} \\ \frac{dV}{dh} &= \frac{\partial V}{\partial h} + \frac{\partial V}{\partial r} \frac{dr}{dh} \\ &= \frac{\pi r^2}{3} + \frac{2\pi r h}{3} \frac{dr}{dh} \end{aligned}$$

Just to be completely concrete, let's think of a right circular cone with a fixed angle $\alpha = \tan r/h$, so that if we change r or h then α remains constant. So we have $r = h \tan^{-1} \alpha$; let constant $c = \tan^{-1} \alpha$, so now $r = ch$. Now, we know that

$$\begin{aligned} \frac{dV}{dr} &= \frac{2\pi r h}{3} + \frac{\pi r^2}{3} \frac{1}{c} \\ \frac{dV}{dh} &= \frac{\pi r^2}{3} + \frac{2\pi r h}{3} c \end{aligned}$$

The BPTT process goes like this:

- (1) Sample a training pair of sequences (x, y) ; let their length be n .
- (2) "Unroll" the RNN to be length n (picture for $n = 3$ below), and initialize s_0 :



Now, we can see our problem as one of performing what is almost an ordinary back-propagation training procedure in a feed-forward neural network, but with the difference that the weight matrices are shared among the layers. In many ways, this is similar to what ends up happening in a convolutional network, except in the conv-net, the weights are re-used spatially, and here, they are re-used temporally.

- (3) Do the *forward pass*, to compute the predicted output sequence p :

$$\begin{aligned} z_t^1 &= W^{sx}x_t + W^{ss}s_{t-1} + W_0 \\ s_t &= f_1(z_t^1) \\ z_t^2 &= W^o s_t + W_0^o \\ p^t &= f_2(z_t^2) \end{aligned}$$

- (4) Do *backward pass* to compute the gradients. For both W^{ss} and W^{sx} we need to find

$$\frac{dL_{\text{seq}}}{dW} = \sum_{u=1}^n \frac{dL_u}{dW}$$

Letting $L_u = L_{\text{elt}}(p_u, y_u)$ and using the *total derivative*, which is a sum over all the ways in which W affects L_u , we have

$$= \sum_{u=1}^n \sum_{t=1}^n \frac{\partial L_u}{\partial s_t} \cdot \frac{\partial s_t}{\partial W}$$

Re-organizing, we have

$$= \sum_{t=1}^n \frac{\partial s_t}{\partial W} \cdot \sum_{u=1}^n \frac{\partial L_u}{\partial s_t}$$

Because s_t only affects L_t, L_{t+1}, \dots, L_n ,

$$\begin{aligned} &= \sum_{t=1}^n \frac{\partial s_t}{\partial W} \cdot \sum_{u=t}^n \frac{\partial L_u}{\partial s_t} \\ &= \sum_{t=1}^n \frac{\partial s_t}{\partial W} \cdot \left(\frac{\partial L_t}{\partial s_t} + \underbrace{\sum_{u=t+1}^n \frac{\partial L_u}{\partial s_t}}_{\delta^{s_t}} \right) \end{aligned} \quad (12.1)$$

δ^{s_t} is the dependence of the loss on steps after t on the state at time t .

We can compute this backwards, with t going from n down to 1. The trickiest part is figuring out how early states contribute to later losses. We define *future loss*

That is, δ^{s_t} is how much we can blame state s_t for all the future element losses.

$$F_t = \sum_{u=t+1}^n \text{Loss}_{\text{elt}}(p_u, y_u) ,$$

so

$$\delta^{s_t} = \frac{\partial F_t}{\partial s_t} .$$

At the last stage, $F_n = 0$ so $\delta^{s_n} = 0$.

Now, working backwards,

$$\begin{aligned} \delta^{s_{t-1}} &= \frac{\partial}{\partial s_{t-1}} \sum_{u=t}^n \text{Loss}_{\text{elt}}(p_u, y_u) \\ &= \frac{\partial s_t}{\partial s_{t-1}} \cdot \frac{\partial}{\partial s_t} \sum_{u=t}^n \text{Loss}_{\text{elt}}(p_u, y_u) \\ &= \frac{\partial s_t}{\partial s_{t-1}} \cdot \frac{\partial}{\partial s_t} \left[\text{Loss}_{\text{elt}}(p_t, y_t) + \sum_{u=t+1}^n \text{Loss}_{\text{elt}}(p_u, y_u) \right] \\ &= \frac{\partial s_t}{\partial s_{t-1}} \cdot \left[\frac{\partial \text{Loss}_{\text{elt}}(p_t, y_t)}{\partial s_t} + \delta^{s_t} \right] \end{aligned}$$

Now, we can use the chain rule again to find the dependence of the element loss at time t on the state at that same time,

$$\underbrace{\frac{\partial \text{Loss}_{\text{elt}}(p_t, y_t)}{\partial s_t}}_{(m \times 1)} = \underbrace{\frac{\partial z_t^2}{\partial s_t}}_{(m \times v)} \cdot \underbrace{\frac{\partial \text{Loss}_{\text{elt}}(p_t, y_t)}{\partial z_t^2}}_{(v \times 1)} ,$$

and the dependence of the state at time t on the state at the previous time, noting that we are performing an *elementwise* multiplication between W_{ss}^T and the vector of $f^{1'}$ values, $\partial s_t / \partial z_t^1$:

$$\underbrace{\frac{\partial s_t}{\partial s_{t-1}}}_{(m \times m)} = \underbrace{\frac{\partial z_t^1}{\partial s_{t-1}}}_{(m \times m)} \cdot \underbrace{\frac{\partial s_t}{\partial z_t^1}}_{(m \times 1)} = \underbrace{W_{ss}^T * f^{1'}(z_t^1)}_{\text{not dot!}} .$$

Putting this all together, we end up with

$$\delta^{s_{t-1}} = \underbrace{\frac{\partial s_t}{\partial s_{t-1}}}_{W_{ss}^T * f^{1'}(z_t^1)} \cdot \left(\underbrace{W_0^T \frac{\partial L_t}{\partial z_t^2}}_{\frac{\partial F_{t-1}}{\partial s_t}} + \delta^{s_t} \right)$$

Last Updated: 02/01/19 01:38:02

There are two ways to think about $\partial s_t / \partial z_t^1$: here, we take the view that it is an $m \times 1$ vector and we multiply each column of W^T by it. Another, equally good, view, is that it is an $m \times m$ diagonal matrix, with the values along the diagonal, and then this operation is a matrix multiply. Our software implementation will take the first view.

We're almost there! Now, we can describe the actual weight updates. Using equation 12.1 and recalling the definition of $\delta^{s_t} = \partial F_t / \partial s_t$, as we iterate backwards, we can accumulate the terms in equation 12.1 to get the gradient for the whole loss:

$$\frac{dL_{\text{seq}}}{dW^{ss}} + = \frac{\partial F_{t-1}}{\partial W^{ss}} = \frac{\partial z_t^1}{\partial W^{ss}} \frac{\partial s_t}{\partial z_t^1} \frac{\partial F_{t-1}}{\partial s_t}$$

$$\frac{dL_{\text{seq}}}{dW^{sx}} + = \frac{\partial F_{t-1}}{\partial W^{sx}} = \frac{\partial z_t^1}{\partial W^{sx}} \frac{\partial s_t}{\partial z_t^1} \frac{\partial F_{t-1}}{\partial s_t}$$

We can handle W^O separately; it's easier because it does not effect future losses in the way that the other weight matrices do:

$$\frac{\partial L_{\text{seq}}}{\partial W^O} = \sum_{t=1}^n \frac{\partial L_t}{\partial W^O} = \sum_{t=1}^n \frac{\partial L_t}{\partial z_t^2} \cdot \frac{\partial z_t^2}{\partial W^O}$$

Assuming we have $\frac{\partial L_t}{\partial z_t^2} = (p_t - y_t)$, (which ends up being true for squared loss, softmax-NLL, etc.), then on each iteration

$$\underbrace{\frac{\partial L_{\text{seq}}}{\partial W^O}}_{v \times m} + = \underbrace{(p_t - y_t)}_{v \times 1} \cdot \underbrace{s_t^T}_{1 \times m}$$

Whew!

Study Question: Derive the updates for the offsets W_0 and W_0^O .

4 Training a language model

A *language model* is just trained on a set of input sequences, $(c_1^{(i)}, c_2^{(i)}, \dots, c_{n_i}^{(i)})$, and is used to predict the next character, given a sequence of previous tokens: _____

A "token" is generally a character or a word.

$$c_t = \text{RNN}(c_1, c_2, \dots, c_{t-1})$$

We can convert this to a sequence-to-sequence training problem by constructing a data set of (x, y) sequence pairs, where we make up new special tokens, start and end, to signal the beginning and end of the sequence:

$$x = (\langle \text{start} \rangle, c_1, c_2, \dots, c_n)$$

$$y = (c_1, c_2, \dots, \langle \text{end} \rangle)$$

5 Vanishing gradients and gating mechanisms

Let's take a careful look at the backward propagation of the gradient along the sequence:

$$\delta^{s_{t-1}} = \frac{\partial s_t}{\partial s_{t-1}} \cdot \left[\frac{\partial \text{Loss}_{\text{elt}}(p_t, y_t)}{\partial s_t} + \delta^{s_t} \right]$$

Consider a case where only the output at the end of the sequence is incorrect, but it depends critically, via the weights, on the input at time 1. In this case, we will multiply the loss at step n by

$$\frac{\partial s_2}{\partial s_1} \cdot \frac{\partial s_3}{\partial s_2} \dots \frac{\partial s_n}{\partial s_{n-1}} \cdot$$

In general, this quantity will either grow or shrink exponentially with the length of the sequence, and make it very difficult to train.

Study Question: The last time we talked about exploding and vanishing gradients, it was to justify per-weight adaptive step sizes. Why is that not a solution to the problem this time?

An important insight that really made recurrent networks work well on long sequences was the idea of *gating*.

5.1 Simple gated recurrent networks

A computer only ever updates some parts of its memory on each computation cycle. We can take this idea and use it to make our networks more able to retain state values over time and to make the gradients better-behaved. We will add a new component to our network, called a *gating network*. Let g_t be a $m \times 1$ vector of values and let W^{g^x} and W^{g^s} be $m \times l$ and $m \times m$ weight matrices, respectively. We will compute g_t as

$$g_t = \text{sigmoid}(W^{g^x}x_t + W^{g^s}s_{t-1})$$

It can have an offset, too, but we are omitting it for simplicity.

and then change the computation of s_t to be

$$s_t = (1 - g_t) * s_{t-1} + g_t * f_1(W^{s^x}x_t + W^{s^s}s_{t-1} + W_0) ,$$


where $*$ is component-wise multiplication. We can see, here, that the output of the gating network is deciding, for each dimension of the state, how much it should be updated now. This mechanism makes it much easier for the network to learn to, for example, “store” some information in some dimension of the state, and then not change it during future state updates.

Study Question: Why is it important that the activation function for g be a sigmoid?

5.2 Long short-term memory

The idea of gating networks can be applied to make a state-machine that is even more like a computer memory, resulting in a type of network called an LSTM for “long short-term memory.” We won’t go into the details here, but the basic idea is that there is a memory cell (really, our state vector) and three (!) gating networks. The *input* gate selects (using a “soft” selection as in the gated network above) which dimensions of the state will be updated with new values; the *forget* gate decides which dimensions of the state will have its old values moved toward 0, and the *output* gate decides which dimensions of the state will be used to compute the output value. These networks have been used in applications like language translation with really amazing results. A diagram of the architecture is shown below:

Yet another awesome name for a neural network!



Recommender systems

The problem of choosing items from a large set to recommend to a user comes up in many contexts, including music services, shopping, and online advertisements. As well as being an important application, it is interesting because it has several formulations, some of which take advantage of a particular interesting structure in the problem.

Concretely, we can think about a company like Netflix, which recommends movies to its users. Netflix knows the ratings given by many different people to many different movies, and your ratings on a small subset of all possible movies. How should it use this data to recommend a movie for you to watch tonight?

There are two prevailing approaches to this problem. The first, content-based recommendation, is formulated as a supervised learning problem. The second, collaborative filtering, introduces a new learning problem formulation.

1 Content-based recommendations

In content-based recommendation, we try to learn a predictor, f , that uses the movies that you have rated so far as training data, finds a hypothesis that maps a movie into a prediction of what rating you would give it, and then returns some movies with high predicted ratings.

The first step is designing representations for the input and output.

It's actually pretty difficult to design a good feature representation for movies. Reasonable approaches might construct features based on the movie's genre, length, main actors, director, location, or even ratings given by some standard critics or aggregation sources. This design process would yield

$$\phi : \text{movie} \rightarrow \text{vector} .$$

Movie ratings are generally given in terms of some number of stars, so the output domain might be $\{1, 2, 3, 4, 5\}$. It's not appropriate for one-hot encoding on the output, and pretending that these are real values is also not entirely sensible. Nevertheless, we will treat the output as if it's in \mathbb{R} .

Study Question: Why is one-hot not a good choice? Why is \mathbb{R} not a good choice?

Thermometer coding might be reasonable, but it's hard to say without trying it. Some more advanced techniques try to predict rankings (would I prefer movie A over movie B) rather than raw ratings.

Now that we have an encoding, we can make a training set based on *your* previous ratings of movies

$$\left\{ \left(\phi(m^{(1)}), \text{rating}(m^{(1)}) \right), \left(\phi(m^{(2)}), \text{rating}(m^{(2)}) \right), \dots \right\}$$

The next step is to pick a loss function. This is closely related to the choice of output encoding. Since we decided to treat the output as a real, we can formulate the problem as a regression from $\phi \rightarrow \mathbb{R}$, with $\text{Loss}(p, y) = \frac{1}{2}(y - p)^2$. We will generally need to regularize because we typically have a very small amount of data (unless you really watch a lot of movies!).

Finally, we need to pick a hypothesis space. The simplest thing would be to make it linear, but you could definitely use something fancier, like a neural network.

If we put all this together, with a linear hypothesis space, we end up with the objective

$$J(\theta) = \frac{1}{2} \sum_{i \in D_a} (y^{(i)} - \theta^T x^{(i)} - \theta_0)^2 + \frac{\lambda}{2} \|\theta\|^2 .$$

This is our old friend, ridge regression, and can be solved analytically or with gradient descent.

2 Collaborative filtering

There are two difficulties with content-based recommendation systems:

- It's hard to design a good feature set to represent movies.
- They only use your previous movie ratings, but don't have a way to use the rest of their data.

In collaborative filtering, we'll try to use *all* the ratings that other people have made of movies to help make better predictions for you.

Intuitively, we can see this process as finding the kinds of people who like the kinds of movies I like, and then predicting that I will like other movies that they like.

Formally, we will start by constructing a *data matrix* Y , where Y_{ai} represents the score given by user a to movie i . So, if we have n users and m movies, Y has shape $n \times m$.

In fact, there's a third strategy that is really directly based on this idea, in which we concretely try to find other users who are our "nearest neighbors" in movie preferences, and then predict movies they like. The approach we discuss here has similar motivations but is more robust.

We will in fact not *actually* represent the whole data matrix explicitly—it would be too big. But it's useful to think about.

	m movies					
	5		3			
		1				
			2			...
n users	4					
			⋮			

Y is very sparse (most entries are empty). So, we will think of our training data-set as $D = \{(a, i, r)\}$, a set of tuples, where a is the index assigned to a particular user, i is the index assigned to a particular movie, and r is user a 's rating of movie i .

In the Netflix challenge data set, there are 400,000 users and 17,000 movies. Only 1% of the data matrix is filled.

We are going to try to find a way to use D to predict values for missing entries. Let X be our predicted matrix of ratings. Now, we need to find a loss function that relates X and Y , so that we can try to optimize it to find a good predictive model.

Idea #1 Following along with our previous approaches to designing loss functions, we might want to say that our predictions X_{ai} should agree with our data Y_{ai} , and then add some regularization, yielding loss function

$$\text{Loss}(X, Y) = \frac{1}{2} \sum_{(a,i) \in D} (Y_{ai} - X_{ai})^2 + \sum_{\text{all } (a,i)} X_{ai}^2 .$$

This is a **bad idea!** It will set $X_{ai} = 0$ for all $(a, i) \notin D$.

Study Question: Convince yourself of that!

We need to find a different kind of regularization that will force some generalization to unseen entries.

Linear algebra idea: The *rank* of a matrix is the maximum number of linearly independent rows in the matrix (which is equal to the maximum number of linearly independent columns in the matrix).

If an $n \times m$ matrix X is rank 1, then there exist U and V of shapes $n \times 1$ and $m \times 1$, respectively, such that

$$X = UV^T .$$

If X is rank k , then there exist U and V of shape $n \times k$ and $m \times k$, respectively, such that

$$X = UV^T .$$

Idea #2 Find the rank 1 matrix X that fits the entries in Y as well as possible. This is a much lower-dimensional representation (it has $m + n$ parameters rather than $m \cdot n$ parameters) and the same parameter is shared among many predictions, so it seems like it might have better generalization properties than our previous idea.

So, we would need to find vectors U and V such that

$$UV^T = \begin{bmatrix} \mathbf{u}^{(1)} \\ \vdots \\ \mathbf{u}^{(n)} \end{bmatrix} \begin{bmatrix} \mathbf{v}^{(1)} & \dots & \mathbf{v}^{(m)} \end{bmatrix} = \begin{bmatrix} \mathbf{u}^{(1)}\mathbf{v}^{(1)} & \dots & \mathbf{u}^{(1)}\mathbf{v}^{(m)} \\ \vdots & \ddots & \vdots \\ \mathbf{u}^{(n)}\mathbf{v}^{(1)} & \dots & \mathbf{u}^{(n)}\mathbf{v}^{(m)} \end{bmatrix} = X .$$

And, since we're using squared loss, our objective function would be

$$J(U, V) = \frac{1}{2} \sum_{(a,i) \in D} (\mathbf{u}^{(a)}\mathbf{v}^{(i)} - Y_{ai})^2 .$$

Now, how can we find the optimal values of U and V ? We could take inspiration from our work on linear regression and see that the gradients of J are with respect to the parameters in U and V . For example,

$$\frac{\partial J}{\partial \mathbf{u}^{(a)}} = \sum_{i, \text{ for } (a,i) \in D} (\mathbf{u}^{(a)}\mathbf{v}^{(i)} - Y_{ai})\mathbf{v}^{(i)} .$$


We could get an equation like this for each parameter $U^{(a)}$ or $V^{(i)}$. We don't know how to get an immediate analytic solution to this set of equations because the parameters U and V are multiplied by one another in the predictions, so the model does not have a linear dependence on the parameters. We could approach this problem using gradient descent, though, and we'll do that with a related model in the next section.

But, before we talk about optimization, let's think about the expressiveness of this model. It has one parameter per user (the elements of U) and one parameter per movie (the elements of V), and the predicted rating is the product of these two. It can really represent only each user's general enthusiasm and each movie's general popularity, and predict the user's rating of the movie to be the product of these values.

Study Question: What if we had two users, 1 and 2, and two movies, A and B. Can you find U, V that represents well the data set $(1, A, 1), (1, B, 5), (2, A, 5), (2, B, 1)$?

Idea #3 If using a rank 1 decomposition of the matrix is not expressive enough, maybe we can try a *rank* k decomposition! In this case, we would try to find an $n \times k$ matrix U and an $m \times k$ matrix V that minimize

$$J(U, V) = \frac{1}{2} \sum_{(a,i) \in D} (\mathbf{u}^{(a)} \cdot \mathbf{v}^{(i)} - Y_{ai})^2 .$$



Here, the length k vector $u^{(a)}$ is the a^{th} row of U , and represents the k “features” of person a . Likewise, the length k vector $v^{(i)}$ is the i^{th} row of V , and represents the k “features” of movie i . Performing the matrix multiplication $X = UV^T$, we see what the prediction for person a and movie i is $x_{ai} = u^{(a)} \cdot v^{(i)}$.

The total number of parameters that we have is $nk + mk$. But, it is a redundant representation. We have 1 extra scaling parameter when $k = 1$, and k^2 in general. So, we really effectively have $nk + mk - k^2$ “degrees of freedom.”

Study Question: Imagine $k = 3$. If we were to take the matrix U and multiply the first column by 2, the second column by 3 and the third column by 4, to make a new matrix U' , what would we have to do to V to get a V' so that $U'V'^T = UV^T$? How does this question relate to the comments above about redundancy?

It is still useful to add offsets to our predictions, so we will include an $n \times 1$ vector b_U and an $m \times 1$ vector b_V of offset parameters, and perform regularization on the parameters in U and V . So our final objective becomes

$$J(U, V) = \frac{1}{2} \sum_{(a,i) \in D} (u^{(a)} \cdot v^{(i)} + b_U^{(a)} + b_V^{(i)} - Y_{ai})^2 + \frac{\lambda}{2} \sum_{a=1}^n \|u^{(a)}\|^2 + \frac{\lambda}{2} \sum_{i=1}^m \|v^{(i)}\|^2.$$

Study Question: What would be an informal interpretation of $b_U^{(a)}$? Of $b_V^{(i)}$?

2.1 Optimization

Now that we have an objective, it’s time to optimize! There are two reasonable approaches to finding U , V , b_U , and b_V that optimize this objective: alternating least squares (ALS), which builds on our analytical solution approach for linear regression, and stochastic gradient descent (SGD), which we have used in the context of neural networks and other models.

2.1.1 Alternating least squares

One interesting thing to notice is that, if we were to fix U and b_U , then finding the minimizing V and b_V is a linear regression problem that we already know how to solve. The same is true if we were to fix V and b_V , and seek U and b_U . So, we will consider an algorithm that takes alternating steps of this form: we fix U, b_U , initially randomly, find the best V, b_V ; then fix those and find the best U, b_U , etc.

This is a kind of optimization sometimes called “coordinate descent,” because we only improve the model in one (or, in this case, a set of) coordinates of the parameter space at a time. Generally, coordinate ascent has similar kinds of convergence properties as gradient descent, and it cannot guarantee that we find a global optimum. It is an appealing choice in this problem because we know how to directly move to the optimal values of one set of coordinates given that the other is fixed.

More concretely, we:

1. Initialize V and b_V at random
2. For each a in $1, 2, \dots, n$
 - Construct a linear regression problem to find $U^{(a)}$ to minimize

$$\frac{1}{2} \sum_{\{i|(a,i) \in D\}} \left(U^{(a)} \cdot V^{(i)} + b_U^{(a)} + b_V^{(i)} - Y_{ai} \right)^2 + \frac{\lambda}{2} \|U^{(a)}\|^2$$

- Recall minimizing the least squares objective (we are ignoring the offset and regularizer in the following so you can see the basic idea):

$$(W\theta - T)^T(W\theta - T) .$$

In this scenario,

- $\theta = U^{(a)}$ is the $k \times 1$ parameter vector that we are trying to find,
- T is a $m_a \times 1$ vector of target values (for the m_a movies a has rated), and
- W is the $m_a \times k$ matrix whose rows are the $V^{(i)}$ where a has rated movie i .

The solution to the least squares problem using ridge regression is our new $U^{(a)}$ and $b_U^{(a)}$.

3. For each i in $1, 2, \dots, m$
 - Construct a linear regression problem to find $V^{(i)}$ and $b_V^{(i)}$ to minimize

$$\frac{1}{2} \sum_{\{i|(a,i) \in D\}} \left(U^{(a)} \cdot V^{(i)} + b_U^{(a)} + b_V^{(i)} - Y_{ai} \right)^2 + \frac{\lambda}{2} \|V^{(i)}\|^2$$

- Now, $\theta = V^{(i)}$ is a $k \times 1$ parameter vector, T is a $n_i \times 1$ target vector (for the n_i users that have rated movie i), and W is the $n_i \times k$ matrix whose rows are the $U^{(a)}$ where i has been rated by user a .

Again, we solve using ridge regression for a new value of $V^{(i)}$ and $b_V^{(i)}$.

4. Alternate between steps 2 and 3, optimizing U and V , and stop after a fixed number of iterations or when the difference between successive parameter estimates is small.

2.1.2 Stochastic gradient descent

Finally, we can approach this problem using stochastic gradient descent. It's easier to think about if we reorganize the objective function to be

$$J(\mathbf{U}, \mathbf{V}) = \frac{1}{2} \sum_{(a,i) \in \mathcal{D}} \left(\left(\mathbf{U}^{(a)} \cdot \mathbf{V}^{(i)} + b_{\mathbf{U}}^{(a)} + b_{\mathbf{V}}^{(i)} - Y_{ai} \right)^2 + \lambda_{\mathbf{U}}^{(a)} \|\mathbf{U}^{(a)}\|^2 + \lambda_{\mathbf{V}}^{(i)} \|\mathbf{V}^{(i)}\|^2 \right)$$

where

$$\lambda_{\mathbf{U}}^{(a)} = \frac{\lambda}{\# \text{ times } (a, _) \in \mathcal{D}} = \frac{\lambda}{\sum_{\{i | (a,i) \in \mathcal{D}\}} 1}$$

$$\lambda_{\mathbf{V}}^{(i)} = \frac{\lambda}{\# \text{ times } (_, i) \in \mathcal{D}} = \frac{\lambda}{\sum_{\{a | (a,i) \in \mathcal{D}\}} 1}$$

Then,

$$\frac{\partial J(\mathbf{U}, \mathbf{V})}{\partial \mathbf{U}^{(a)}} = \sum_{\{i | (a,i) \in \mathcal{D}\}} \left[\left(\mathbf{U}^{(a)} \cdot \mathbf{V}^{(i)} + b_{\mathbf{U}}^{(a)} + b_{\mathbf{V}}^{(i)} - Y_{ai} \right) \mathbf{V}^{(i)} + \lambda_{\mathbf{U}}^{(a)} \mathbf{U}^{(a)} \right]$$

$$\frac{\partial J(\mathbf{U}, \mathbf{V})}{\partial b_{\mathbf{U}}^{(a)}} = \sum_{\{i | (a,i) \in \mathcal{D}\}} \left(\mathbf{U}^{(a)} \cdot \mathbf{V}^{(i)} + b_{\mathbf{U}}^{(a)} + b_{\mathbf{V}}^{(i)} - Y_{ai} \right)$$

We can similarly obtain gradients with respect to $\mathbf{V}^{(i)}$ and $b_{\mathbf{V}}^{(i)}$.

Then, to do gradient descent, we draw an example (a, i, Y_{ai}) from \mathcal{D} at random, and do gradient updates on $\mathbf{U}^{(a)}$, $b_{\mathbf{U}}^{(a)}$, $\mathbf{V}^{(i)}$, and $b_{\mathbf{V}}^{(i)}$.

Study Question: Why don't we update the other parameters, such as $\mathbf{U}^{(a')}$ for some other user a' or $\mathbf{V}^{(i')}$ for some other movie i' ?

1 Intro

We will continue to broaden the class of models that we can fit to our data. Neural networks have adaptable complexity, in the sense that we can try different structural models and use cross validation to find one that works well on our data.

We now turn to models that automatically adapt their complexity to the training data. The name *non-parametric methods* is misleading: it is really a class of methods that does not have a fixed parameterization in advance. Some non-parametric models, such as decision trees, which we might call *semi-parametric methods*, can be seen as dynamically constructing something that ends up looking like a more traditional parametric model, but where the actual training data affects exactly what the form of the model will be. Other non-parametric methods, such as nearest-neighbor, rely directly on the data to make predictions and do not compute a model that summarizes the data.

The semi-parametric methods tend to have the form of a composition of simple models. We'll look at:

- *Tree models*: partition the input space and use different simple predictions on different regions of the space; this increases the hypothesis space.
- *Additive models*: train several different classifiers on the whole space and average the answers; this decreases the estimation error.

Boosting is a way to construct an additive model that decreases both estimation and structural error, but we won't address it in this class.

2 Trees

The idea here is that we would like to find a partition of the input space and then fit very simple models to predict the output in each piece. The partition is described using a (typically binary) "decision tree," which recursively splits the space.

These methods differ by:

- The class of possible ways to split the space at each node; these are generally linear splits, either aligned with the axes of the space, or more general.

- The class of predictors within the partitions; these are often simply constants, but may be more general classification or regression models.
- The way in which we control the complexity of the hypothesis: it would be within the capacity of these methods to have a separate partition for each individual training example.
- The algorithm for making the partitions and fitting the models.

The primary advantage of tree models is that they are easily interpretable by humans. This is important in application domains, such as medicine, where there are human experts who often ultimately make critical decisions and who need to feel confident in their understanding of recommendations made by an algorithm.

We'll concentrate on the CART/ID3 family of algorithms, which were invented independently in the statistics and the artificial intelligence communities. They work by greedily constructing a partition, where the splits are *axis aligned* and by fitting a *constant* model in the leaves. The interesting questions are how to select the splits and how to control capacity. The regression and classification versions are very similar.

2.1 Regression

The predictor is made up of

- A partition function, π , mapping elements of the input space into exactly one of M regions, R_1, \dots, R_M .
- A collection of M output values, O_m , one for each region.

If we already knew a division of the space into regions, we would set \hat{y}_m , the constant output for region R_m , to be the average of the training output values in that region; that is:

$$O_m = \text{average}_{\{i | x^{(i)} \in R_m\}} y^{(i)} .$$

Define the error in a region as

$$E_m = \sum_{\{i | x^{(i)} \in R_m\}} (y^{(i)} - O_m)^2 .$$

Ideally, we would select the partition to minimize

$$\lambda M + \sum_{m=1}^M E_m ,$$

for some regularization constant λ . It is enough to search over all partitions of the training data (not all partitions of the input space) to optimize this, but the problem is NP-complete.

2.1.1 Building a tree

So, we'll be greedy. We establish a criterion, given a set of data, for finding the best single split of that data, and then apply it recursively to partition the space. We will select the partition of the data that *minimizes the sum of the mean squared errors of each partition*.

Given a data set D , let

- $R_{j,s}^+(D) = \{x \in D \mid x_j \geq s\}$
- $R_{j,s}^-(D) = \{x \in D \mid x_j < s\}$

- $\hat{y}_{j,s}^+ = \text{average}_{\{i | x^{(i)} \in R_{j,s}^+(D)\}} y^{(i)}$
- $\hat{y}_{j,s}^- = \text{average}_{\{i | x^{(i)} \in R_{j,s}^-(D)\}} y^{(i)}$

BuildTree(D):

- If $|D| \leq k$: return Leaf(D)
- Find the variable j and split point s that minimizes:

$$E_{R_{j,s}^+(D)} + E_{R_{j,s}^-(D)} .$$

- Return **Node**($j, s, \text{BuildTree}(R_{j,s}^+(D)), \text{BuildTree}(R_{j,s}^-(D))$)

Each call to **BuildTree** considers $O(dn)$ splits (only need to split between each data point in each dimension); each requires $O(n)$ work.

2.1.2 Pruning

It might be tempting to regularize by stopping for a somewhat large k , or by stopping when splitting a node does not significantly decrease the error. One problem with short-sighted stopping criteria is that they might not see the value of a split that is, essentially, two-dimensional. So, we will tend to build a tree that is much too large, and then prune it back.

Define *cost complexity* of a tree T , where m ranges over its leaves as

$$C_\alpha(T) = \sum_{m=1}^{|T|} E_m(T) + \alpha|T| .$$

For a fixed α , find a T that (approximately) minimizes $C_\alpha(T)$ by “weakest-link” pruning. Create a sequence of trees by successively removing the bottom-level split that minimizes the increase in overall error, until the root is reached. Return the T in the sequence that minimizes the criterion.

Pick α using cross validation.

2.2 Classification

The strategy for building and pruning classification trees is very similar to the one for regression trees.

The output is now the majority of the output values in the leaf:

$$O_m = \text{majority}_{\{i | x^{(i)} \in R_m\}} y^{(i)} .$$

Define the error in a region as the number of data points that do not have the value O_m :

$$E_m = \left| \{i | x^{(i)} \in R_m \text{ and } y^{(i)} \neq O_m\} \right| .$$

Define the *empirical probability* of an item from class k occurring in region m as:

$$\hat{P}_{mk} = \hat{P}(R_m)(k) = \frac{|\{i | x^{(i)} \in R_m \text{ and } y^{(i)} = k\}|}{N_m} ,$$

where N_m is the number of training points in region m . We’ll define the empirical probabilities of split values, as well, for later use.

$$\hat{P}_{mjv} = \hat{P}(R_{mj})(v) = \frac{|\{i | x^{(i)} \in R_m \text{ and } x_j^{(i)} \geq v\}|}{N_m}$$

Splitting criteria Minimize “impurity” in child nodes. Some measures include:

- *Misclassification error*:

$$Q_m(T) = \frac{E_m}{N_m} = 1 - \hat{p}_{m0}$$

- *Gini index*:

$$Q_m(T) = \sum_k \hat{p}_{mk}(1 - \hat{p}_{mk})$$

- *Entropy*:

$$Q_m(T) = H(R_m) = - \sum_k \hat{p}_{mk} \log \hat{p}_{mk}$$


Choosing the split that minimizes the entropy of the children is equivalent to maximize the *information gain* of the test $X_j = v$, defined by

$$\text{infoGain}(X_j = v, R_m) = H(R_m) - (\hat{p}_{mjv} H(R_{j,v}^+) + (1 - \hat{p}_{mjv}) H(R_{j,v}^-))$$

In the two-class case, all the criteria have the values

$$\begin{cases} 0.0 & \text{when } \hat{p}_{m0} = 0.0 \\ 0.0 & \text{when } \hat{p}_{m0} = 1.0 \end{cases}$$


The respective impurity curves are shown below, where $p = \hat{p}_{m0}$:



There used to be endless haggling about which one to use. It seems to be traditional to use:

- Entropy to select which node to split while growing the tree
- Misclassification error in the pruning criterion

As a concrete example, consider the following images:



The left image depicts a set of labeled data points and the right shows a partition into regions by a decision tree.

Points about trees There are many variations on this theme:

- Linear regression or other regression or classification method in each leaf
- Non-axis-parallel splits: e.g., run a perceptron for a while to get a split.

What's good about trees:

- Easily interpretable
- Easy to handle multi-class classification
- Easy to handle different loss functions (just change predictor in the leaves)

What's bad about trees:

- High estimation error: small changes in the data can result in very big changes in the hypothesis.
- Often not the best predictions

Hierarchical mixture of experts Make a "soft" version of trees, in which the splits are probabilistic (so every point has some degree of membership in every leaf). Can be trained with a form of gradient descent.

3 Bagging

Bootstrap aggregation is a technique for reducing the estimation error of a non-linear predictor, or one that is adaptive to the data.

- Construct B new data sets of size n by sampling them with replacement from \mathcal{D}
- Train a predictor on each one: \hat{f}^b
- *Regression case*: bagged predictor is

$$\hat{f}_{\text{bag}}(x) = \frac{1}{B} \sum_{b=1}^B \hat{f}^b(x)$$

- *Classification case*: majority bagged predictor: let $\hat{f}^b(x)$ be a "one-hot" vector with a single 1 and $K - 1$ zeros, so that $\hat{y}^b(x) = \arg \max_k \hat{f}^b(x)_k$. Then

$$\hat{f}_{\text{bag}}(x) = \frac{1}{B} \sum_{b=1}^B \hat{f}^b(x),$$

which is a vector containing the proportion of classifiers that predicted each class k for input x ; and the predicted output is

$$\hat{y}_{\text{bag}}(x) = \arg \max_k \hat{f}_{\text{bag}}(x)_k .$$

There are theoretical arguments showing that bagging does, in fact, reduce estimation error.

However, when we bag a model, any simple predictability is lost.

3.1 Random Forests

Random forests are collections of trees that are constructed to be de-correlated, so that using them to vote gives maximal advantage.

For $b = 1..B$

- Draw a bootstrap sample \mathcal{D}_b of size n from \mathcal{D}
- Grow a tree on data \mathcal{D}_b by recursively repeating these steps:
 - Select m variables at random from the d variables
 - Pick the best variable and split point among them
 - Split the node
- return tree T_b

Given the ensemble of trees, vote to make a prediction on a new x .

4 Nearest Neighbor

In nearest-neighbor models, we don't do any processing of the data at training time—we just remember it! All the work is done at prediction time.

Input values x can be from any domain \mathcal{X} (\mathbb{R}^d , documents, tree-structured objects, etc.). We just need a distance metric, $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$, which satisfies the following, for all $x, x', x'' \in \mathcal{X}$:

$$\begin{aligned} d(x, x) &= 0 \\ d(x, x') &= d(x', x) \\ d(x, x'') &\leq d(x, x') + d(x', x'') \end{aligned}$$

Given a data-set $\mathcal{D} = \{(x^{(i)}, y^{(i)})\}_{i=1}^n$, our predictor for a new $x \in \mathcal{X}$ is


$$h(x) = y^{(i)} \quad \text{where } i = \arg \min_i d(x, x^{(i)}) ,$$

that is, the predicted output associated with the training point that is closest to the query point x .

This same algorithm works for regression *and* classification!

The nearest neighbor prediction function can be described by a Voronoi partition (dividing the space up into regions whose closest point is each individual training point) as shown below:

It's a floor wax *and* a dessert topping!



In each region, we predict the associated y value.

There are several useful variations on this method. In *k-nearest-neighbors*, we find the k training points nearest to the query point x and output the majority y value for classification or the average for regression. We can also do *locally weighted regression* in which we fit locally linear regression models to the k nearest points, possibly giving less weight to those that are farther away. In large data-sets, it is important to use good data structures (e.g., ball trees) to perform the nearest-neighbor look-ups efficiently (without looking at all the data points each time).