

CSC 446 Lecture Notes

March 12, 2015

Contents

| | | |
|-----------|---|-----------|
| 1 | What Is Machine Learning? | 1 |
| 2 | Probability Theory | 1 |
| 3 | Concentration Bounds | 3 |
| 4 | Maximum Likelihood Estimation | 3 |
| 5 | Entropy | 5 |
| 5.1 | Bounds on Entropy for a Discrete Random Variable | 5 |
| 5.2 | Further Entropy Equations | 5 |
| 6 | Mutual Information | 6 |
| 6.1 | Covariance | 6 |
| 6.2 | KL divergence | 6 |
| 6.3 | Lower Bound for KL divergence | 6 |
| 6.4 | L_1 norm | 7 |
| 7 | The Gaussian Distribution | 7 |
| 7.1 | Maximum Likelihood Estimation | 8 |
| 7.2 | Maximum Entropy for fixed Variance | 8 |
| 7.3 | Marginalization | 8 |
| 7.4 | Conditioning | 8 |
| 8 | Linear Regression | 9 |
| 9 | Smoothing | 9 |
| 9.1 | Prior Distributions | 10 |
| 9.2 | Dirichlet Prior | 10 |
| 9.3 | Gamma Function | 10 |
| 9.4 | Justifying the Dirichlet Prior | 11 |
| 10 | Comparison - Bayesian vs. MLE vs. MAP | 12 |
| 10.1 | Bayesian | 12 |
| 11 | Perceptrons | 12 |
| 11.1 | Proof of Convergence | 14 |
| 11.2 | Perceptron in Stochastic Gradient Descent perspective | 15 |
| 12 | Multi Layer Perceptron | 16 |
| 12.1 | Training A Network: Error Backpropagation | 16 |

| | |
|---|-----------|
| 13 Support Vector Machines | 17 |
| 13.1 Training Linear SVMs | 17 |
| 13.2 Convex Optimization Review | 19 |
| 13.3 Karush-Kuhn-Tucker (KKT) Conditions | 20 |
| 14 Kernel Functions | 20 |
| 14.1 Review Support Vector Machines | 20 |
| 14.2 Kernel Function | 21 |
| 14.3 Proof that ϕ exists | 22 |
| 14.4 Regression | 23 |
| 15 Graphical Models | 24 |
| 15.1 Example | 24 |
| 15.2 Factor Graph | 25 |
| 15.3 Message Passing (Belief Propagation) | 26 |
| 15.4 Running Time | 27 |
| 16 Junction Tree | 27 |
| 16.1 Max-Sum | 28 |
| 16.2 Tree Decomposition | 29 |
| 17 Expectation Maximization | 31 |
| 17.1 Parameter Setting: An Example | 31 |
| 17.2 Expectation-Maximization Algorithm | 32 |
| 17.3 EM Algorithm in General | 33 |
| 17.4 Gradient Ascent ($\frac{\partial L}{\partial \theta}$) | 35 |
| 17.5 Newton's Method | 36 |
| 17.6 Variational Method | 36 |
| 17.7 Mixture of Gaussians | 36 |
| 18 Sampling | 38 |
| 18.1 How to Sample a Continuous Variable: Basics | 38 |
| 18.2 The Metropolis-Hastings Algorithm | 38 |
| 18.3 Proof of the method | 39 |
| 18.4 Gibbs Sampling | 41 |
| 18.5 Gibbs sampling with Continuous Variables. | 42 |
| 18.6 EM with Gibbs Sampling. | 42 |
| 18.6.1 Some problems. | 42 |
| 18.6.2 Some advantages. | 43 |
| 19 Error Bounds | 44 |
| 19.1 Leave One Out | 44 |
| 19.2 One Way to Compute an Error Bound | 44 |
| 19.3 Another Way to Compute an Error Bound | 45 |
| 20 Logistic Regression a.k.a. Maximum Entropy | 47 |
| 21 Hidden Markov Models | 50 |
| 22 LBFGS | 51 |
| 22.1 Preliminary | 52 |
| 22.2 The BFGS Algorithm | 52 |
| 22.3 Proof of the method | 53 |
| 22.4 L-BFGS Algorithm | 53 |

| | |
|---|-----------|
| 23 Reinforcement Learning | 54 |
| 23.1 Markov Decision Processes | 54 |
| 23.2 Value Iteration | 54 |
| 23.3 Q-Learning | 55 |
| 23.4 Temporal Difference Learning | 56 |
| 24 Game Theory | 56 |
| 24.1 Definition of the Problem | 57 |
| 24.2 Some Simplification | 57 |
| 24.3 Nash Equilibrium | 58 |
| 24.4 Proof of the LP Duality | 59 |

1 What Is Machine Learning?

Machine learning is about automatically analyzing data; it mostly focuses on the problems of classification and regression. In this class, we will learn multiple methods for solving each of these problems.

- **Classification** is the problem of assigning datapoints to discrete categories; the goal is to pick the best category for each datapoint.
- **Regression** is the problem of learning a function from datapoints to numbers; fitting a line or a curve to the data is an example of a regression problem.

One example of a classification problem would be: given heights and weights, classify people by sex. We are given a number of training datapoints, whose heights, weights, and sexes we know; we can plot these datapoints in a two-dimensional space. Our goal is to learn a rule from these datapoints that will allow us to classify other people whose heights and weights are known but whose sex is unknown. This rule will take the form of a curve in our two-dimensional space: then, when confronted with future datapoints, we will classify those datapoints which fall below the curve as female those which fall above the curve as male. So how should we draw this curve?

One idea would be to draw a winding curve which carefully separates the datapoints, assuring that all males are on one side and all females are on the other. But this is a very complicated rule, and it's likely to match our training data too closely and not generalize well to new data. Another choice would be to draw a straight line; this is a much simpler rule which is likely to do better on new data, but it does not classify all of the training datapoints correctly. This is an example of a fundamental tradeoff in machine learning, that of **overfitting** vs. **generalization**. We will return to this tradeoff many times during this class, as we learn methods of preventing overfitting.

An example of a regression problem would be: given weights of people, predict their heights. We can apply the nearest neighbor model to solve this problem. The nearest neighbor model remembers the weights and corresponding heights of the people in the training data. Then for a new test weight, it looks up the person with the closest weight in the training data, and returns the corresponding height. This results in a piece-wise constant function that may be affected by outliers and may result in overfitting. Another choice would be to fit a straight line.

2 Probability Theory

This section contains a quick review of basic concepts from probability theory.

Let X be a **random variable**, i.e., a variable that can take on various values, each with a certain probability. Let x be one of those values. Then we denote the probability that $X = x$ as $P(X = x)$. (We will often write this less formally, as just $P(x)$, leaving it implicit which random variable we are discussing. We will also use $P(X)$ to refer to the entire probability distribution over possible values of X .)

In order for $P(X)$ to be a valid probability distribution, it must satisfy the following properties:

- For all x , $P(X = x) \geq 0$.

- $\sum_x P(X = x) = 1$ or $\int P(x)dx = 1$, depending on whether the probability distribution is discrete or continuous.

If we have two random variables, X and Y , we can define the **joint distribution** over X and Y , denoted $P(X = x, Y = y)$. The comma is like a logical “and”; this is the probability that both $X = x$ and $Y = y$. Analogously to the probability distributions for a single random variable, the joint distribution must obey the properties that for all x and for all y , $P(X = x, Y = y) \geq 0$ and either $\sum_{x,y} P(X = x, Y = y) = 1$ or $\int \int P(x, y)dydx = 1$, depending on whether the distribution is discrete or continuous.

From the joint distribution $P(X, Y)$, we can **marginalize** to get the distribution $P(X)$: namely, $P(X = x) = \sum_y P(X = x, Y = y)$. We can also define the **conditional probability** $P(X = x|Y = y)$, the probability that $X = x$ given that we already know $Y = y$. This is $P(X = x|Y = y) = \frac{P(X=x, Y=y)}{P(Y=y)}$, which is known as the **product rule**. Through two applications of the product rule, we can derive **Bayes rule**:

$$P(X = x|Y = y) = \frac{P(Y = y|X = x)P(X = x)}{P(Y = y)}$$

Two random variables X and Y are **independent** if knowing the value of one of the variables does not give us any further clues as to the value of the other variable. Thus, for X and Y independent, $P(X = x|Y = y) = P(X = x)$, or, written another way, $P(X = x, Y = y) = P(X = x)P(Y = y)$.

The **expectation** of a random variable X with respect to the probability distribution $P(X)$ is defined as $E_P[X] = \sum_x P(X = x)x$ or $E_P[X] = \int P(x)x dx$, depending on whether the random variable is discrete or continuous. The expectation is a weighted average of the values that a random variable can take on. Of course, this only makes sense for random variables which take on numerical values; this would not work in the example from earlier where the two possible values of the “sex” random variable were “male” and “female”.

We can also define the **conditional expectation**, the expectation of a random variable with respect to a conditional distribution: $E_{P(X|Y)}[X] = \sum_x P(X = x|Y = y)x$. This is also sometimes written as $E_P[X|Y]$. Lastly, we are not restricted to taking the expectations of random variables only; we can also take the expectation of functions of random variables: $E_P[f(X)] = \sum_x P(X = x)f(x)$. Note that we will often leave the probability distribution implicit and write the expectation simply as $E[X]$.

Expectation is **linear**, which means that the expectation of the sum is the sum of the expectations, i.e., $E[X + Y] = E[X] + E[Y]$, or, more generally, $E[\sum_{i=1}^N X_i] = \sum_{i=1}^N E[X_i]$. For the two-variable case, this can be proven as follows:

$$\begin{aligned} E[X + Y] &= \sum_{x,y} P(X = x, Y = y)(x + y) \\ &= \sum_{x,y} P(X = x, Y = y)x + \sum_{x,y} P(X = x, Y = y)y \\ &= \sum_x P(X = x)x + \sum_y P(Y = y)y \\ &= E[X] + E[Y] \end{aligned}$$

The N -variable case follows from this by induction.

For readability, let $\bar{x} = E[X]$. Then we can define the **variance**, $Var[X] = E[(x - \bar{x})^2]$. In words, the variance is the weighted average of the distance from the mean squared. Why is the distance squared? Well, if we take out the square, we get that $Var[X] = E[x - \bar{x}]$, which by linearity of expectation equals $E[x] - E[\bar{x}] = E[X] - \bar{x} = 0$, so we put the square in to keep that from happening. The reason we do not use absolute value instead is that the absolute value function is nondifferentiable. As a result of squaring, the variance penalizes further outliers more.

Unlike expectation, variance is not linear; that means in general $Var[X + Y] \neq Var[X] + Var[Y]$. The covariance of X and Y is defined as: $Cov[X, Y] = E[(X - E[X])(Y - E[Y])]$. We can show that $Var[aX] = a^2 Var[X]$ and $Var[X + Y] = Var[X] + Var[Y] + 2Cov[X, Y]$. If X and Y are independent, then $Cov[X, Y] = 0$.

3 Concentration Bounds

Markov's Inequality For a non-negative random variable, $X > 0$, and for any $\delta > 0$,

$$P(X \geq \delta E[X]) \leq \frac{1}{\delta}$$

or equivalently,

$$P(X \geq a) \leq \frac{E[X]}{a}$$

Proof: Your homework.

Chebyshev's Inequality For any random variable with finite variance σ^2 , and for any $k > 0$

$$P(|X - E[X]| \geq k\sigma) \leq \frac{1}{k^2}$$

Proof: Your homework.

4 Maximum Likelihood Estimation

Consider a set of N **independent and identically distributed** random variables X_1, \dots, X_N . “Independent and identically distributed”, which is usually abbreviated as i.i.d., means that the random variables are pairwise independent (i.e., for each i, j such that $i \neq j$, X_i and X_j are independent) and that they are all distributed according to the same probability distribution, which we will call P . Much of the data that we will look at in this class is i.i.d. Since our goal is to automatically infer the probability distribution P that is the best description of our data, it is essential to assume that all of our datapoints were actually generated by the same distribution. One of the implications of i.i.d. is that the joint probability distribution over all of the random variables decomposes as follows: $P(X_1, \dots, X_N) = \prod_{n=1}^N P(X_n)$.

Again, our task is to automatically infer the probability distribution P which best describes our data. But how can we quantify which probability distribution gives the best description? Suppose that our random variables are discrete and have K possible outcomes such that each datapoint X takes on a value in $\{1, \dots, K\}$. We can describe P with a K -dimensional vector that we will call θ , letting $\theta_k = P(X = k)$ for each k ; θ is called the **parameters** of the distribution. It's useful to describe the probabilities in our distribution using a vector, because then we can employ the powerful tools of vector calculus to help us solve our problems.

Now the question of finding the best probability distribution becomes a question of finding the optimal setting of θ . A good idea would be to pick the value of θ which assigns the highest probability to the data:

$$\theta^* = \underset{\theta}{\operatorname{argmax}} P(X_1, \dots, X_N; \theta)$$

This method of estimating θ is called **maximum likelihood estimation**, and we will call the optimal setting of the parameters θ_{MLE} . It is a constrained optimization problem that we can solve using the tools of vector calculus, though first we will introduce some more convenient notation. For each k , let $c(k) = \sum_{n=1}^N I(X_n = k)$ be the number of datapoints with value k . Here, I is an indicator variable which is 1 when the statement in the parentheses is true, and 0 when it is false.

Using these counts, we can rewrite the probability of our data as follows:

$$\begin{aligned} P(X_1, \dots, X_N | \theta) &= \prod_{n=1}^N \theta_{x_n} \\ &= \prod_{k=1}^K \theta_k^{c(k)} \end{aligned}$$

This switch in notation is very important, and we will do it quite frequently. Here we have grouped our data according to outcome rather than ordering our datapoints sequentially.

Now we can proceed with the optimization. Our goal is to find $\operatorname{argmax}_{\theta} \prod_{k=1}^K P(X = k)^{c(k)}$ such that $\sum_{k=1}^K \theta_k = 1$. (We need to add this constraint to assure that whichever θ we get describes a valid probability distribution.) If this were an unconstrained optimization problem, we would solve it by setting the derivative to 0 and then solving for θ . But since this is a constrained optimization problem, we must use a **Lagrange multiplier**.

In general, we might want to solve a constrained optimization problem of the form $\max_{\vec{x}} f(\vec{x})$ such that $g(\vec{x}) = c$. Here, $f(\vec{x})$ is called the **objective function** and $g(\vec{x})$ is called the **constraint**. We form the **Lagrangian**

$$\nabla f(\vec{x}) + \lambda \nabla g(\vec{x}) = 0$$

and then solve for both \vec{x} and λ .

Now we have all the tools required to solve this problem. First, however, we will transform the objective function a bit to make it easier to work with, using the convenient fact that the logarithm is monotonic increasing, and thus does not affect the solution.

$$\begin{aligned} \max_{\theta} \prod_{k=1}^K P(X = k)^{c(k)} &= \max_{\theta} \log \left(\prod_{k=1}^K \theta_k^{c(k)} \right) \\ &= \max_{\theta} \sum_{k=1}^K \log(\theta_k^{c(k)}) \\ &= \max_{\theta} \sum_{k=1}^K c(k) \log(\theta_k) \end{aligned}$$

We get the gradient of this objective function by, for each θ_k , taking the partial derivative with respect to θ_k :

$$\frac{\partial}{\partial \theta_k} \left[\sum_{j=1}^K c(j) \log(\theta_j) \right] = \frac{c(k)}{\theta_k}$$

(To get this derivative, observe that all of the terms in the sum are constant with respect to θ_k except for the one term containing θ_k ; taking the derivative of that term gives the result, and the other terms' derivatives are 0.)

Thus, we get that

$$\nabla f = \frac{\partial}{\partial \theta} \left[\sum_{j=1}^K c(j) \log(\theta_j) \right] = \left(\frac{c(1)}{\theta_1}, \dots, \frac{c(K)}{\theta_K} \right)$$

In a similar vein,

$$\nabla g = \frac{\partial}{\partial \theta} \left[\sum_{j=1}^K \theta_j \right] = (1, \dots, 1)$$

Now we substitute these results into the Lagrangian $\nabla f + \lambda \nabla g = 0$. Solving this equation, we discover that for each k , $\frac{c(k)}{\theta_k} = -\lambda$, or $\theta_k = -\frac{c(k)}{\lambda}$. To solve for λ , we substitute this back into our constraint, and discover that $\sum_{k=1}^K \theta_k = -\frac{1}{\lambda} \sum_{k=1}^K c(k)$, and thus $-\lambda = \sum_{k=1}^K c(k)$. This is thus our normalization constant.

In retrospect, this formula seems completely obvious. The probability of outcome k is the fraction of times outcome k occurred in our data. The math accords perfectly with our intuitions; why would we ever want to do anything else? The problem is that this formula overfits our data, like the curve separating the male datapoints from the female datapoints at the beginning of class. For instance, suppose we never see outcome k in our data. This formula would have us set $\theta_k = 0$. But we probably don't want to assume that outcome k will never, ever happen. In the next lecture, we will look at how to avoid this issue.

5 Entropy

Entropy is:

$$\begin{aligned} H(X) &= \sum_x P(x) \log \frac{1}{P(x)} \\ &= \int P(x) \log \frac{1}{P(x)} dx \end{aligned}$$

We can think of this as a measure of information content. An example of this idea of information content is seen in Huffman coding. High frequency letters have short encodings while rarer letters have longer encodings. This forms a binary tree where the letters are at the leaves and edges to the left are 0 bits and edges to the right are 1 bits. If the probabilities for the letters are all equal then this tree is balanced.

In the case of entropy we notice that $\log \frac{1}{P(x)}$ is a non-integer, so it is like an expanded Huffman coding.

5.1 Bounds on Entropy for a Discrete Random Variable

If the variable is discrete $H(X)$ is maximized when the distribution is uniform since $P(x) = \frac{1}{K}$, we see:

$$H(X) = \sum_{i=1}^K \frac{1}{K} \log K = \log K$$

If K is 2^n then $H(X) = \log 2^n = n$. Part of Homework 2 will be to prove that entropy on a discrete random variable is maximized by a uniform distribution ($\max_{\theta} H(X)$ where $\sum_n \theta_n = 1$ using the Lagrange equation).

To minimize $H(X)$ we want $P(x_i) = 1$ for some i (with all other $P(x_j)$ being zero¹) giving $H(X) = \sum_{1 \leq j \leq K, j \neq i} 0 \log \frac{1}{0} + 1 \log 1 = 0$. We see then that:

$$0 \leq H(X) \leq \log K$$

If we consider some distribution we can see that if we cut up the “shape” of the distribution and add in gaps that the gaps that are added do not contribute to $P(x) \log \frac{1}{P(x)}$.

5.2 Further Entropy Equations

$$\begin{aligned} H(X, Y) &= \sum P(x, y) \log \frac{1}{P(x, y)} \\ H(X|Y) &= \sum_{x,y} P(x|y)P(y) \log \frac{1}{P(x|y)} \\ &= E_{XY} \left[\log \frac{1}{P(x|y)} \right] \\ &= \sum_{x,y} P(x, y) \log \frac{1}{P(x|y)} \end{aligned}$$

¹What about $0 \cdot \log \frac{1}{0}$? It is standard to define this as equal to zero (justified by the limit being zero).

6 Mutual Information

Mutual information attempts to measure how correlated two variables are with each other:

$$I(X; Y) = \sum_{x,y} P(x, y) \log \frac{P(x, y)}{P(x)P(y)}$$

Consider communicating the values of two variables. The mutual information of these two variables is the difference between the entropy of communicating these variables individually and the entropy if we can send them together. For example if X and Y are the same then $H(X) + H(Y) = 2H(X)$ while $H(X, Y) = H(X)$ (since we know Y if we are given X). So $I(X; Y) = 2H(X) - H(X) = H(X)$.

6.1 Covariance

A number version of mutual information is covariance:

$$\begin{aligned} \text{Covar}[X, Y] &= \sum_{x,y} P(x, y)(x - \bar{X})(y - \bar{Y}) \\ \text{Covar}[X, X] &= \text{Var}[X] \end{aligned}$$

Covariance indicates the high level trend, so if both X and Y are generally increasing, or both generally decreasing, then the covariance will be positive. If one is generally increasing, but the other is generally decreasing, then the covariance will be negative. Two variables can have a high amount of mutual information but no general related trend and the covariance will not indicate much (probably be around zero).

6.2 KL divergence

Kullback–Leibler (KL) divergence compares two distributions over some variable:

$$\begin{aligned} D(P \parallel Q) &= \sum_x P(x) \log \frac{P(x)}{Q(x)} \\ &= E_P \left[\log \frac{1}{Q(x)} - \log \frac{1}{P(x)} \right] \\ &= \underbrace{H_P(Q)}_{\text{Cross Entropy}} - \underbrace{H(P)}_{\text{Entropy}} \end{aligned}$$

If we have the same distribution then there is no divergence $D(P \parallel P) = 0$. In general the KL divergence is non-symmetric $D(P \parallel Q) \neq D(Q \parallel P)$. If neither distribution is “special” the average $\frac{1}{2}[D(P \parallel Q) + D(Q \parallel P)]$ is sometimes used and is symmetric. The units of KL divergence are log probability.

The cross entropy has an information interpretation quantifying how many bits are wasted by using the wrong code:

$$H_P(Q) = \sum_x \underbrace{P(x)}_{\text{Sending } P} \log \overbrace{\frac{1}{Q(x)}}^{\text{code for } Q}$$

6.3 Lower Bound for KL divergence

We will show that KL divergence is always greater or equal to zero using Jensen’s inequality. First we need a definition of convex. A function f is convex if for all x_1, x_2 and θ where $0 \leq \theta \leq 1$, $f(\theta x_1 + (1 - \theta)x_2) \leq$

$\theta f(x_1) + (1 - \theta)f(x_2)$. This is saying that any chord on the function is above the function itself on the same interval.

Some examples of convex include a straight line and $f(x) = x^2$. If the Hessian exists for a function then $\nabla^2 f \succeq 0$ (the Hessian is positive semidefinite) indicates that f is convex. This works for a line, but not something like $f(x) = |x|$.

Jensen's inequality states that if f is convex then $E[f(X)] \geq f(E[X])$.

Proof.

$$\begin{aligned} D(P \parallel Q) &= E_P \left[\log \frac{P(x)}{Q(x)} \right] \\ &= E_P \left[-\log \frac{Q(x)}{P(x)} \right] \end{aligned}$$

To apply Jensen's inequality we will let $-\log$ be our function and $\frac{Q(x)}{P(x)}$ be our x (note that this ratio is a number so we can push the E_P inside).

$$\begin{aligned} E_P \left[-\log \frac{Q(x)}{P(x)} \right] &\geq -\log E_P \left[\frac{Q(x)}{P(x)} \right] \\ &= -\log \sum_x P(x) \frac{Q(x)}{P(x)} \\ &= -\log 1 = 0 \end{aligned}$$

□

Thinking of our information interpretation, we see that we always pay some cost for using the wrong code. Also note that $\log \frac{P(x)}{Q(x)}$ is sometimes positive and sometimes negative (P and Q both sum to one), yet $D(P \parallel Q) \geq 0$.

6.4 L_1 norm

The L_1 norm is defined as:

$$\|P - Q\|_1 = \sum_x |P(x) - Q(x)|$$

It can be thought of as "how much earth has to be moved" to match the distributions.

Because P and Q sum to one we quickly see that $0 \leq \|P - Q\|_1 \leq 2$. This property can be advantageous when bounds are needed.

7 The Gaussian Distribution

$$p(x; \mu, \sigma^2) = \frac{1}{\sqrt{2\pi}\sigma} \exp \left(-\frac{(x - \mu)^2}{2\sigma^2} \right)$$

$$p(x; \mu, \Sigma) = \frac{1}{2\pi^{d/2}|\Sigma|^{1/2}} \exp \left(-\frac{1}{2}(x - \mu)^T \Sigma^{-1}(x - \mu) \right)$$

7.1 Maximum Likelihood Estimation

$$\mu = \frac{1}{N} \sum_{n=1}^N x^{(n)}$$
$$\Sigma_{ij} = \frac{1}{N} \sum_{n=1}^N (x_i^{(n)} - \mu_i)(x_j^{(n)} - \mu_j)$$

7.2 Maximum Entropy for fixed Variance

$$\begin{aligned} \max_{p(x)} \quad & - \int p(x) \log p(x) dx \\ \text{s.t.} \quad & \int p(x)(x - \mu)^2 dx = 1 \\ & \int p(x) dx = 1 \end{aligned}$$

Solving with langrange multipliers:

$$p(x) = \frac{1}{Z} \exp(\lambda(x - \mu)^2) \tag{1}$$

7.3 Marginalization

If

$$\begin{aligned} x &= \begin{bmatrix} x_a & x_b \end{bmatrix} \\ x &\sim N(X; \mu \Sigma) \\ \mu &= \begin{bmatrix} \mu_a \\ \mu_b \end{bmatrix} \\ \Sigma &= \begin{bmatrix} \Sigma_{aa} & \Sigma_{ab} \\ \Sigma_{ba} & \Sigma_{bb} \end{bmatrix} \end{aligned}$$

then the marginal for x_a is itself gaussian:

$$x_a \sim N(x_a; \mu_a, \Sigma_{aa})$$

7.4 Conditioning

With x_a, x_b as above, the conditional distribution $P(X_a | x_b)$ is gaussian:

$$\begin{aligned} \mu_{a|b} &= \mu_a + \Sigma_{ab} \Sigma_{bb}^{-1} (x_b - \mu_b) \\ \Sigma_{a|b} &= \Sigma_{aa} - \Sigma_{ab} \Sigma_{bb}^{-1} \Sigma_{ba} \end{aligned}$$

8 Linear Regression

Let our prediction $\hat{y} = w^T x$.

$$\begin{aligned} \min_w \sum_n (\hat{y} - y_n)^2 \\ \min_w \sum_n (w^T x^{(n)} - y_n)^2 \\ \min_w \|w^T X - y\|^2 \end{aligned}$$

$$\begin{aligned} 0 &= \frac{\partial}{\partial w} \|w^T X - y\|^2 \\ &= 2X^T X w - 2X^T y \\ w &= (X^T X)^{-1} X^T y \end{aligned}$$

with regularization

$$w = (X^T X + \lambda I)^{-1} X^T y$$

9 Smoothing

If we use MLE to train a classifier, all of our probabilities are based on counts, any unseen combination of a single feature x and the class label y results in

$$\begin{aligned} P(x|y) &= \frac{c(x, y)}{c(y)} \\ &= 0. \end{aligned}$$

These zeros can ruin the entire classifier. For example, say there's one bill where all the Republicans we know about voted "no". Now, say we are trying to classify an unknown politician who followed the Republican line on every other bill, but voted "yes" on this bill. The classifier will say that there is zero probability of this person being a Republican, since it has never seen the combination (Republican, voted yes) for that bill. It gives that single feature way too much power. To get rid of that, we can use a technique called smoothing, and modify the probabilities a little :

$$\begin{aligned} P(x = k|y) &= \frac{c(x = k, y) + \alpha}{c(y) + K\alpha} \\ k &\in \{1, \dots, K\} \end{aligned}$$

Basically we are taking a little bit of the probability mass from things with high probability and giving it to things with otherwise zero probability. (Republicans might veto this technique, since it's like redistribution of wealth!) Note that these probabilities must still sum to 1. This seems great - we've gotten rid of things with zero probability. But doesn't this contradict what we proved earlier? That is, last week we said that we can best infer the probability distribution by solving

$$\begin{aligned} \operatorname{argmax}_{\theta} \prod_{n=1}^N P_{\theta}(x_n) \\ \text{s.t. } \sum_{k=1}^K \theta_k = 1 \end{aligned}$$

which results in the count-based distribution

$$\theta_k^* = \frac{c(k)}{N}.$$

How then can we mathematically justify our smoothed probabilities?

9.1 Prior Distributions

We can treat θ as a random variable itself with some probability distribution $P(\theta)$. Recall that θ is a vector of probabilities for each type of event k , so

$$\theta = [\theta_1, \theta_2, \dots, \theta_K]^T$$

and $\sum_{k=1}^K \theta_k = 1$

Suppose that we have a coin with two outcomes, heads or tails ($K=2$). We can picture the θ_1 and θ_2 which we could pick for the probability distribution of these two outcomes. A fair coin has $\theta_1 = 1/2$ and $\theta_2 = 1/2$. An weighted coin might have $\theta_1 = 2/3$ and $\theta_2 = 1/3$. Since we are treating θ as a random variable, its probability $P(\theta)$ is describing the probability that it takes on these values. $P(\theta)$ is called a prior, since it's what we believe about θ before we even have any observations. For example, we might tend to believe that the coin will be pretty fair, so we could have $P(\theta)$ be a normal curve with the peak where $\theta_1 = 1/2$ and $\theta_2 = 1/2$.

9.2 Dirichlet Prior

One useful prior distribution is the Dirichlet Prior :

$$\begin{aligned} P(\theta) &= \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)} \prod_{k=1}^K \theta_k^{\alpha_k - 1} \\ &= \frac{1}{Z} \prod_{k=1}^K \theta_k^{\alpha_k - 1} \end{aligned}$$

This is also written as $P(\theta; \alpha)$, $P_\alpha(\theta)$, or $P(\theta|\alpha)$. α is a vector with the same size as θ , and it is known as a “hyperparameter”. The choice of α determines the shape of θ 's distribution, which you can see by varying it. If α is simply a vector of ones, we just get a uniform distribution; all θ s are equally probable. In the case of two variables, we can have $\alpha_1=100$, and $\alpha_2=50$ and we see a sharp peak around $2/3$. The larger α_1 , the more sharply peaked it gets around $\frac{\alpha_1}{\alpha_1 + \alpha_2}$.

At this point, we are tactfully ignoring that Γ in the Dirichlet distribution. What is that function, and what does it do?

9.3 Gamma Function

$$\Gamma(x) = \int_0^\infty e^{-t} t^{x-1} dt$$

This function occurs often in difficult, nasty integrals. However, it has the nice property of being equivalent to the factorial function:

$$\Gamma(n) = (n-1)!$$

We can prove this using integration by parts:

$$\begin{aligned}
\Gamma(x) &= \int_0^\infty e^{-t} t^{x-1} dt \\
&= [-t^{x-1} e^{-t}]_0^\infty + \int_0^\infty e^{-t} (x-1) t^{x-2} dt \\
&= 0 + (x-1) \int_0^\infty e^{-t} t^{x-2} dt \\
&= (x-1) \Gamma(x-1)
\end{aligned}$$

Further noting that $\Gamma(1) = 1$, we can conclude that $\Gamma(n) = (n-1)!$. This function is used in the normalization constant of our Dirichlet prior in order to guarantee that:

$$\int_{\sum_k \theta_k = 1} P(\theta) d\theta = 1.$$

9.4 Justifying the Dirichlet Prior

How can we use this prior to compute probabilities?

$$\begin{aligned}
P(x = k | \theta) &= \theta_k \\
P(x = k) &= \int_{\sum_k \theta_k = 1} P(x | \theta) P(\theta) d\theta \\
&= \int \theta_k \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)} \prod_{k=1}^K \theta_k^{\alpha_k - 1} d\theta \\
&= \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)} \int \prod_{k'=1}^K \theta_k^{\alpha_{k'} - 1 + I(k'=k)} d\theta \\
&= \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)} \frac{\prod_{k'} \Gamma(\alpha_{k'} + I(k'=k))}{\Gamma(\sum_{k'} \alpha_{k'} + I(k'=k))} \\
&= \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\Gamma(\sum \alpha_k + 1)} \frac{\Gamma(\alpha_k + 1)}{\Gamma(\alpha_{k'})}
\end{aligned}$$

Now we use $\Gamma(x) = (x-1)\Gamma(x-1)$:

$$= \frac{\alpha_k}{\sum_{k'} \alpha_{k'}}$$

Most of the time, all of the α_k 's are set to the same number. So, we just showed that

$$P(x) = \frac{\alpha_k}{\sum_{k'} \alpha_{k'}}$$

But what about

$$\begin{aligned}
P(X_{N+1}|X_1^N) &= \int P(X_{N+1}, \theta | X_1^N) d\theta \\
&= \int P(X_{N+1} | \theta, X_1^N) P(\theta | X_1^N) d\theta \\
&= \int \theta_k \frac{P(X_1^N | \theta) P(\theta)}{P(X_1^N)} d\theta \\
&= \frac{1}{Z} \int \theta_k \prod_n \theta_{X_n} \frac{1}{Z'} \prod_k \theta_k^{\alpha_k - 1} d\theta \\
&\dots \\
&= \frac{c(k) + \alpha_k}{N + \sum_k \alpha_k}
\end{aligned}$$

10 Comparison - Bayesian vs. MLE vs. MAP

10.1 Bayesian

The quantity we just computed is known as the Bayesian:

$$P(X_{N+1}|X_1^N) = \frac{c(k) + \alpha_k}{N + \sum_k \alpha_k}$$

We can compare it to the MLE that we did before:

$$\begin{aligned}
P(x_{N+1}) \text{ follows } \theta^* \\
\theta^* = \underset{\theta}{\operatorname{argmax}} P_{\theta}(X_1^N)
\end{aligned}$$

And a third alternative is the MAP, or Maximum A Posteriori:

$$\begin{aligned}
P(x_{N+1}) \text{ follows } \theta^* \\
\theta^* = \underset{\theta}{\operatorname{argmax}} P(\theta) P(X_1^N | \theta)
\end{aligned}$$

This is simpler since it does not require an integral. Using the same Lagrange Multipliers technique as we did before:

$$\begin{aligned}
&\operatorname{argmax} \frac{1}{Z} \prod_k \theta_k^{\alpha_k - 1} \prod_k \theta_k^c(k) \\
&\text{s.t. } \sum_k \theta_k = 1
\end{aligned}$$

Then we get the result:

$$\theta_k^* = \frac{c(k) + \alpha_k - 1}{N + (\sum_{k'} \alpha_{k'}) - K}$$

11 Perceptrons

A Perceptron is a linear classifier that determines a decision boundary through successive changes to the slope of a line according to a binary feature. The classifier finds values for the weight vector \mathbf{w}^T to solve the equation

$$\mathbf{0} = \mathbf{w}^T \mathbf{x} + \mathbf{b}$$

We define our classification function sign as

$$\text{sign}(x) = \begin{cases} -1 & : x < 0 \\ 0 & : x = 0 \\ 1 & : x > 0 \end{cases}$$

We can remove \mathbf{b} from the equation by adding it as an element of \mathbf{w} and adding a 1 to \mathbf{x} in the same spot.

$$\mathbf{w}' = \begin{bmatrix} w_1 \\ \vdots \\ w_N \\ b \end{bmatrix} \quad \mathbf{x}' = \begin{bmatrix} x_1 \\ \vdots \\ x_N \\ 1 \end{bmatrix}$$

$$\mathbf{w}'^T \mathbf{x}' = \mathbf{w}^T \mathbf{x} + \mathbf{b}$$

The next question is, how do we pick \mathbf{w} ? We have \mathbf{x} as the vector of data points, with \mathbf{x}^n as the n th data point. We have t^n as the classifier output (1 or -1) for the n th data point.

$$t^n = \text{sign}(\mathbf{w}^T \mathbf{x}^n)$$

y^n is the true label for the n th data point.

Our goal is to solve the following equation:

$$\underset{\mathbf{w}}{\text{argmax}} \sum_n I(y^n \neq t^n)$$

Or, reformulated as minimizing error,

$$\underset{\mathbf{w}}{\text{argmin}} \sum_n I(y^n \neq t^n)$$

To solve this equation, we want to add up all of the positive points to get a vector in that direction. This gives us the Perceptron algorithm:

```

repeat
  for  $n = 1 \dots N$  do
    if  $t^n \neq y^n$  then
       $\mathbf{w} \leftarrow \mathbf{w} + y^n \mathbf{x}^n$ 
    end if
  end for
until  $\forall n \ t^n = y^n$  or maxiters

```

While this algorithm will completely separate linearly separable data, it may not be the best separation (it may not accurately represent the separating axis of the data).

We can solve this problem by replacing our original error function,

$$E = \sum_n I(y^n \neq t^n)$$

with the following error function:

$$E = \sum_n \frac{1}{2} (t^n - y^n)^2 = \sum_n E^n$$

Now, the function E w.r.t. \mathbf{w} is not convex, so we have to use an iterative method. We use Gradient Descent to solve for \mathbf{w} , which is the successive application of

$$\mathbf{w} = \mathbf{w} - \frac{\partial E}{\partial \mathbf{w}}$$




Figure 1: Linear Separability

$$\frac{\partial E}{\partial \mathbf{w}} = \sum_n \frac{\partial E^n}{\partial \mathbf{w}} = \sum_n \frac{\partial E^n}{\partial t} \frac{\partial t}{\partial \mathbf{w}}$$

However, our sign function is not differentiable at 0, so we replace it with an activation function g . A good candidate for g is tanh:

$$t^n = g(a) = \tanh(a) = \frac{e^a - e^{-a}}{e^a + e^{-a}}$$

where

$$a = \mathbf{w}^T \mathbf{x}$$

Plugging in, we get

$$\frac{\partial E}{\partial \mathbf{w}} = \sum_n (t^n - y^n) \frac{\partial t}{\partial a} \frac{\partial a}{\partial \mathbf{w}} \quad (2)$$

$$= \sum_n (t^n - y^n) g'(\mathbf{w}^T \mathbf{x}) \mathbf{x} \quad (3)$$

An alternative solution, which begins making small adjustments immediately, is Stochastic Gradient Descent. Having defined a learning rate η , the algorithm is only slightly different:

```

repeat
  for  $n = 1 \dots N$  do
     $\mathbf{w} \leftarrow \mathbf{w} - \eta \frac{\partial E^n}{\partial \mathbf{w}}$ 
  end for
until maxiters

```

11.1 Proof of Convergence

As mentioned before, perceptron algorithm will converge eventually if the data is linearly separable, but why? Let's first formally write down the problem,

Definition: $(x^n, y_n), n \in 1, 2, \dots, N$ is linearly separable, iff $\exists (u, \delta), \|u\| = 1, \delta > 0, s.t. \forall n, y_n u^T x^{(n)} \geq \delta$. The vector u is called an *oracle* vector that separates everything correctly.

Theorem: If $x^{(n)}$ is bounded by R , i.e., $\forall n, \|x^{(n)}\| \leq R$, then the perceptron algorithm makes at most $\frac{R^2}{\delta^2}$ updates. (for a vector $\|v\|$ denotes the Euclidean norm of v , i.e., $\|v\| = \sqrt{\sum_i v_i^2}$)

Proof: As we keep updating the weights w in the algorithm, a sequence of $w^{(k)}$ are generated.

Let $w^{(1)} = \mathbf{0}$.

Each time we encountered a misclassification, we use it to update the weights w . Suppose we use data point (x, y) to update $w^{(k)}$, which means Equation (4) holds,

$$w^{(k+1)} = w^{(k)} + yx \quad (4)$$

Bear in mind that since we are using (x, y) to update $w^{(k)}$, $w^{(k)}$ misclassified (x, y) , which means Equation (5) holds,

$$y(w^{(k)})^T x < 0 \quad (5)$$

Now we can use Equation (4), (5) and $w^{(1)} = \mathbf{0}$ to prove that $k \leq \frac{R^2}{\delta^2}$. Hence, the convergence holds. From Equation (6), we can get a upper bound of $\|w^{(k+1)}\|$, and from Equation (7), we can get its lower bound.

$$\begin{aligned} w^{(k+1)} &= w^{(k)} + yx \\ \Rightarrow u^T w^{(k+1)} &= u^T w^{(k)} + yu^T x && \text{Multiply oracle } u \text{ to both sides} \\ \Rightarrow u^T w^{(k+1)} &\geq u^T w^{(k)} + \delta && \text{Definition of } (u, \delta) \\ \Rightarrow u^T w^{(k+1)} &\geq k\delta && \text{Induction and } w^{(1)} = \mathbf{0} \\ \Rightarrow \|u^T w^{(k+1)}\|^2 &\geq k^2\delta^2 && \text{Both sides positive} \\ \Rightarrow \|u^T\|^2 \|w^{(k+1)}\|^2 &\geq k^2\delta^2 && \text{Cauchy Schwarz: } \|a\| \times \|b\| \geq |a^T b| \\ \Rightarrow \|w^{(k+1)}\|^2 &\geq k^2\delta^2 && \|u\| = 1 \end{aligned} \quad (6)$$

$$\begin{aligned} w^{(k+1)} &= w^{(k)} + yx \\ \Rightarrow \|w^{(k+1)}\|^2 &= \|w^{(k)} + yx\|^2 && \text{Apply Euclidean norm to both sides} \\ \Rightarrow \|w^{(k+1)}\|^2 &= \|w^{(k)}\|^2 + \|y\|^2 \|x\|^2 + 2yxw^{(k)} && \text{Expansion} \\ \Rightarrow \|w^{(k+1)}\|^2 &\leq \|w^{(k)}\|^2 + \|y\|^2 \|x\|^2 && \text{Equation (5)} \\ \Rightarrow \|w^{(k+1)}\|^2 &\leq \|w^{(k)}\|^2 + R^2 && \|y\| = 1, \|x\| \leq R \\ \Rightarrow \|w^{(k+1)}\|^2 &\leq kR^2 && \text{Induction and } w^{(1)} = \mathbf{0} \end{aligned} \quad (7)$$

Combining results of Equation (6) and (7), we get $k^2\delta^2 \leq \|w^{(k+1)}\| \leq kR^2$. Thus, $k^2\delta^2 \leq kR^2$ and $k \leq \frac{R^2}{\delta^2}$.

Since the number of updates is bounded by $\frac{R^2}{\delta^2}$, the perceptron algorithm will eventually converge to somewhere no updates are needed. ■

11.2 Perceptron in Stochastic Gradient Descent perspective

The perceptron algorithm can be analyzed in a more general framework, i.e., stochastic gradient descent for a convex optimization problem.

The ultimate goal for perceptron algorithm is to find w , such that $\forall (X^k, y_k), y_k w^T X^k \geq 0$. Therefore, a natural penalty for misclassification is $[-y_k w^T X^k]_+$. For scalar number S , $[S]_+$ is defined in Equation (8), which is called *hinge* function in class.

$$[S]_+ = \begin{cases} S & \text{if } S \geq 0 \\ 0, & \text{if } S < 0 \end{cases} \quad (8)$$

Thus, the optimization problem can be written down as Equation (9).

$$\underset{w}{\operatorname{argmin}} f(w) \triangleq \frac{1}{N} \sum_k f_k(w) \triangleq \frac{1}{N} \sum_k [-y_k w^T X^k]_+ \quad (9)$$

One way to optimize the convex function in Equation (9) is called *Gradient Descent*. Essentially, Gradient Descent keeps updating the weights w , $w \leftarrow w - \alpha \nabla_w f(w)$, in which α is called *learning rate*. The gradient

$\nabla_w f(w)$ can be carried out by $\nabla_w f(w) = \frac{1}{N} \sum_k \nabla_w f_k(w)$, and $\nabla_w f_k(w)$ is computed by Equation (10).

$$\nabla_w f_k(w) = \begin{cases} -y_k X^k & \text{if } y_k w^T X^k < 0 \\ 0, & \text{if } y_k w^T X^k \geq 0 \end{cases} \quad (10)$$

Because the gradient $\nabla_w f(w)$ is a summation of local gradients $\nabla_w f_k(w)$, we can also do Stochastic Gradient Descent by using one data instance a time.

1. Randomly pick a data instance, (X^k, y_k)
2. Compute local gradient on it, $\nabla_w f_k(w)$ as Equation (10).
3. Update weights using the local gradient, $w \leftarrow w - \alpha \nabla_w f_k(w)$. This is exactly the update in perceptron algorithm.

12 Multi Layer Perceptron

In a two layer neural network, the input, hidden, and output variables are represented by nodes, and the weight parameters are represented by links between the nodes. The arrow of the links indicate the direction of information flow through the network during forward propagation. The overall two stage network function takes the form

$$t_k(\mathbf{x}, \mathbf{w}) = \sigma \left(\sum_{j=0} w_{jk} g \left(\sum_{i=0} w_{ij} x_i \right) \right), \quad (11)$$

where

$$z_j = g \left(\sum_{i=0} w_{ij} x_i \right)$$

are known as the hidden units, w_{ij} are the weights, x_i are the input variables, and the σ and g are activation functions. Non-linear functions are usually chosen for activation functions such as tanh and sigmoid functions.

12.1 Training A Network: Error Backpropagation

Given a training set of input vector \mathbf{x}^n and its target vector \mathbf{y}^n for $n = 1 \dots N$, we want to minimize the error function

$$E^n(\mathbf{w}) = \frac{1}{2} \sum_{i=1}^K \| t_i^n - y_i^n \|^2 = \sum_{i=1}^K E_i^n. \quad (12)$$

The response function t_i^n take take form of (11) but to keep notation uncluttered we will omit the input parameters. To minimize the error function of a no hidden layer network, we apply the chain rule for the partial derivatives with respect to (w.r.t.) the weights w_{ij} to get

$$\begin{aligned} \frac{\partial E^n}{\partial w_{ij}} &= \sum_{j'} \frac{\partial E^n}{\partial t_{j'}^n} \frac{\partial t_{j'}^n}{\partial w_{ij}} \\ &= (t_j^n - y_j^n) \frac{\partial t_j^n}{\partial a_j} \frac{\partial a_j}{\partial w_{ij}} \\ &= (t_j^n - y_j^n) g'(a_j) x_i \\ &= \underbrace{(t_j^n - y_j^n) g'(a_j)}_{\delta_j} x_i \\ &= \delta_j x_i \end{aligned} \quad (13)$$

where

$$a_j = \sum_i w_{ij} x_i$$

is the sum of the weighted input feeds. The derivative respect to the matrix \mathbf{W} is

$$\frac{\partial E^n}{\partial \mathbf{W}} = \left[\frac{\partial E^n}{\partial w_{ij}} \right]_{ij} = [\delta_j x_i]_{ij} = \mathbf{x} \delta^T.$$

We now focus on training on a neural network with a hidden layer (multi layer) that takes the same form as (12) and minimize the error w.r.t. the weights w_{ij} in this form

$$\begin{aligned} \frac{\partial E^n}{\partial w_{ij}} &= \sum_k \frac{\partial E^n}{\partial t_k^n} \frac{\partial t_k^n}{\partial w_{ij}} = \sum_k \frac{\partial E^n}{\partial t_k^n} \frac{\partial t_k^n}{\partial a_k} \frac{\partial a_k}{\partial w_{ij}} \\ &= \sum_k (t_k - y_k) \left(g'(a_k) \frac{\partial}{\partial w_{ij}} \sum_{j'} w_{j'k} g(a_{j'}) \right) \\ &= \sum_k (t_k - y_k) \left(g'(a_k) w_{jk} \frac{\partial}{\partial w_{ij}} g(a_j) \right) \\ &= \sum_k (t_k - y_k) (g'(a_k) w_{jk} g'(a_j) x_i) \\ &= \sum_k \underbrace{(t_k - y_k) g'(a_k) w_{jk} g'(a_j)}_{\delta_j} x_i \\ &= \delta_j x_i \end{aligned}$$

The derivatives with respect to the first and second layer weights are given by

$$\frac{\partial E^n}{\partial w_{ij}} = \delta_j x_i \quad \text{and} \quad \frac{\partial E_n}{\partial w_{jk}} = \delta_k z_j.$$

Thus the training algorithm consists of calculating weight for n

- Calculate weight for n
 $\mathbf{W} \leftarrow \mathbf{W} + \eta \frac{\partial E^n}{\partial \mathbf{W}}.$

13 Support Vector Machines

The Support Vector Machine (SVM) is one of the most widely used classification methods. The SVM is different from other classifiers that we have covered so far. The SVM cares only about the data points near the class boundary and finds a hyperplane that maximizes the margin between the classes.

13.1 Training Linear SVMs

Let the input be a set of N training vectors $\{\mathbf{x}_n\}_{n=1}^N$ and corresponding class labels $\{y_n\}_{n=1}^N$, where $\mathbf{x}_n \in \mathbb{R}^D$ and $y_n \in \{-1, 1\}$. Initially we assume that the two classes are linearly separable. The hyperplane separating the two classes can be represented as:

$$\mathbf{w}^T \mathbf{x}_n + b = 0,$$

such that:

$$\begin{aligned} \mathbf{w}^T \mathbf{x}_n + b &\geq 1 \quad \text{for } y_n = +1, \\ \mathbf{w}^T \mathbf{x}_n + b &\leq -1 \quad \text{for } y_n = -1. \end{aligned}$$




Figure 2: The figure shows a linear SVM classifier for two linearly separable classes. The hyperplane $\mathbf{w}^T \mathbf{x} + b$ is the solid line between H_1 and H_2 , and the margin is M .

Let H_1 and H_2 be the two hyperplanes (Figure 2) separating the classes such that there is no other data point between them. Our goal is to maximize the margin M between the two classes. The objective function:

$$\begin{aligned} \max_{\mathbf{w}} \quad & M \\ \text{s.t.} \quad & y^n(\mathbf{w}^T \mathbf{x}_n + b) \geq M, \\ & \mathbf{w}^T \mathbf{w} = 1. \end{aligned}$$

The margin M is equal to $\frac{2}{\|\mathbf{w}\|}$. We can rewrite the objective function as:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \mathbf{w}^T \mathbf{w} \\ \text{s.t.} \quad & y^n(\mathbf{w}^T \mathbf{x}_n + b) \geq 1 \end{aligned}$$

Now, let's consider the case when the two classes are not linearly separable. We introduce slack variables $\{\xi_n\}_{n=1}^N$ and allow few points to be on the wrong side of the hyperplane at some cost. The modified objective function:

$$\begin{aligned} \min_{\mathbf{w}} \quad & \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{n=1}^N \xi_n \\ \text{s.t.} \quad & y^n(\mathbf{w}^T \mathbf{x}_n + b) + \xi_n \geq 1, \\ & \xi_n \geq 0, \quad \forall n. \end{aligned}$$

The parameter C can be tuned using development set. This is the primal optimization problem for SVM.

The Lagrangian for the primal problem:

$$L(\mathbf{w}, b, \xi, \alpha, \mu) = \frac{1}{2} \mathbf{w}^T \mathbf{w} + C \sum_{n=1}^N \xi_n - \sum_n \alpha_n [y_n(\mathbf{w}^T \mathbf{x}_n + b)] - \sum_n \alpha_n \xi_n + \sum_n \alpha_n - \sum_n \mu_n \xi_n,$$

where α_n and μ_n , $1 \leq n \leq N$ are Lagrange multipliers.

Differentiating the Lagrangian with respect to the variables:

$$\frac{\partial}{\partial \mathbf{w}} L(\mathbf{w}, b, \xi, \alpha, \mu) = \mathbf{w} - \sum_n \alpha_n y_n \mathbf{x}_n = 0$$

$$\frac{\partial}{\partial b} L(\mathbf{w}, b, \xi, \alpha, \mu) = - \sum_n \alpha_n y_n = 0$$

$$\frac{\partial}{\partial \xi_n} L(\mathbf{w}, b, \xi, \alpha, \mu) = C - \alpha_n - \mu_n = 0$$

Solving these equations, we get:

$$\mathbf{w} = \sum_n \alpha_n y_n \mathbf{x}_n \quad (14)$$

$$\sum_n \alpha_n y_n = 0$$

$$\alpha_n = C - \mu_n \quad (15)$$

We now plug-in these values to get the dual function and cancelling out some terms:

$$\begin{aligned} g(\alpha, \mu) &= \frac{1}{2} \sum_n \sum_m \alpha_n \alpha_m y_n y_m \mathbf{x}_n^T \mathbf{x}_m - \sum_n \sum_m \alpha_n \alpha_m y_n y_m \mathbf{x}_n^T \mathbf{x}_m + \sum_n \alpha_n \\ &= \sum_n \alpha_n - \frac{1}{2} \sum_n \sum_m \alpha_n \alpha_m y_n y_m \mathbf{x}_n^T \mathbf{x}_m \end{aligned} \quad (16)$$

Using the equation (15) and (16) and the KKT conditions, we obtain the dual optimization problem:

$$\begin{aligned} \max_{\alpha} \quad & \sum_n \alpha_n - \frac{1}{2} \sum_n \sum_m \alpha_n \alpha_m y_n y_m \mathbf{x}_n^T \mathbf{x}_m \\ \text{s.t.} \quad & 0 \leq \alpha_n \leq C. \end{aligned}$$

The dual optimization problem is concave and easy to solve. The dual variables (α_n) lie within a box with side C . We usually vary two values α_i and α_j at a time and numerically optimize the dual function. Finally, we plug in the values of the α_n^* 's to the equations (14) to obtain the primal solution \mathbf{w}^* .

13.2 Convex Optimization Review

Let we are given an optimization problem:

$$\begin{aligned} \min_x \quad & f_0(x) \\ \text{s.t.} \quad & f_i(x) \leq 0, \text{ for } i \in 1, 2, \dots, K, \end{aligned}$$

where f_0 and f_i ($i \in \{1, 2, \dots, K\}$) are convex functions. We call this optimization problem the 'primal' problem.

The Lagrangian is:

$$L(x, \lambda) = f_0(x) + \sum_{i=1}^K \lambda_i f_i(x)$$

The Lagrange dual function:

$$g(\lambda) = \min_x L(x, \lambda)$$

The dual function $g(\lambda)$ is concave and hence easy to solve. We can obtain the minima of a convex primal optimization problem by maximizing the dual function $g(\lambda)$. The dual optimization problem:

$$\begin{aligned} \max_{\lambda} \quad & g(\lambda) \\ \text{s.t.} \quad & \lambda_i \geq 0, \text{ for } i \in 1, 2, \dots, K. \end{aligned}$$

13.3 Karush-Kuhn-Tucker (KKT) Conditions

The Karush-Kuhn-Tucker (KKT) conditions are the conditions for optimality in primal and dual functions. If f_0 and f_i 's are convex, differentiable, and the feasible set has some interior points (satisfies Slater condition), the x^* and λ_i^* 's are the optimal solutions of the primal and dual problems if and only if they satisfy the following conditions:

$$\begin{aligned} f_i(x^*) &\leq 0 \\ \lambda_i^* &\geq 0, \forall i \in 1, \dots, K \\ \frac{\partial}{\partial x} L(x^*, \lambda_1^*, \dots, \lambda_K^*) &= 0 \\ \lambda_i^* f_i(x^*) &= 0 \end{aligned}$$

14 Kernel Functions

14.1 Review Support Vector Machines

Goal: To solve equation:

$$\begin{aligned} \min_w \quad & \left(\frac{1}{2} \|w\|^2 + C \sum_n \xi_n \right) \\ \text{s.t.} \quad & y^n (w^T x^n + b) + \xi_n \geq 1 \\ & \xi_n \geq 0 \end{aligned}$$

where

$$x^n = [x_1, x_2, \dots, x_K]^T, n \in 1, \dots, N$$

This is a K-dimensional problem, which means the more features the data has, the more complicated to solve this problem.

In the meantime, this equation is equal to

$$\begin{aligned} \max_{\alpha} \quad & \left(-\frac{1}{2} \sum_n \sum_m \alpha_n \alpha_m y^n y^m x^{nT} x^m + \sum_n \alpha_n \right) \\ \text{s.t.} \quad & \alpha_n \geq 0 \\ & \alpha_n \leq C \end{aligned}$$


This is a N-dimensional problem, which means the more data points we include in the training set, the more time it takes to find the optimized classifier.

To train the classifier is actually to solve this problem inside the box of alpha.

According to KKT,

$$\begin{aligned}\lambda_i f_i(x) &= 0 \\ \lambda_i &\geq 0 \\ f_i(x) &\leq 0\end{aligned}$$

As shown in the figure below,



$$w = \sum_n \alpha_n y^n x_n$$


Points on the right side but not on the margin contribute nothing because alpha equals to 0. (The green point)

For points on the wrong side (the red point), alpha equals to C, and

$$\xi_n > 0$$

so they along with points on the margin contribute to the vector, but no point is allowed to contribute more than C.

SVM can train classifier better than naive bayes in the most of time, but since its still binary classification it is not able to deal with situation like this one below:



14.2 Kernel Function

Now when we look back, the classification formula is

$$\text{Sign}(w^T x) = \text{Sign} \left(\left(\sum_n \alpha_n y^n x_n \right)^T x \right) = \text{Sign} \left(\sum_n \alpha_n y^n (x^{nT} x) \right)$$

We can introduce Kernel Function K now, the simplest one is:

$$x^{nT} x = K(x^n, x)$$

Now the problem is transformed into:

$$\max_{\alpha} \left(-\frac{1}{2} \sum_n \sum_m \alpha_n \alpha_m y^n y^m K(x^n, x^m) + \sum_n \alpha_n \right)$$

where

$$K(x, y) = \phi(x)^T \phi(y)$$

for some ϕ .

The most commonly seen Kernel Functions are:

$$K(x, y) = x^T y$$

$$K(x, y) = (x^T y)^m$$

$$K(x, y) = e^{-c\|x-y\|^2}$$

Generally, Kernel function is a measure of how x and y are similar, then they are the same, it has the peak output.

14.3 Proof that ϕ exists

For a two dimensional

$$x = [x_1, x_2]^T y = [y_1, y_2]^T$$

$$K(x, y) = (x_1 y_1 + x_2 y_2)^m$$


Let $m = 2$, then

$$K(x, y) = (x_1 y_1)^2 + (x_2 y_2)^2 + 2(x_1 y_1 x_2 y_2) = \phi(x)^T \phi(y)$$

Thus, we can conclude that

$$\phi(x) = [\sqrt{2}x_1 x_2, x_1^2, x_2^2]^T$$

Basically, ϕ transforms x from a linear space to a multi nominal space like shown below:



so that the points can be classified.

For


$$K(x, y) = e^{-\|x-y\|^2}$$

because we have

$$e^x = 1 + x + \frac{x^2}{2} + \dots$$

it transforms feature into a infinite dimensional space. Generally Kernel Functions lead to more dimension of w which is K -dimensional so solve dual is more practical.


For perceptron, its error term and w have a relationship drawn as below: (not convex) so that we cant



do the same thing to perceptron.

14.4 Regression

When we are predicting output we actually have a space like this: The line is the prediction line, the points



around it are the data set we have. We predict y with formula:

$$\hat{y} = w^T x$$

$$w = (X^T X)^{-1} X^T \vec{y}$$

its known as linear regression. The goal is to

$$\min_w \sum_n \frac{1}{2} \|\hat{y}^n - y^n\|^2$$

which leads us to Support Vector Regression:

$$\min_w \frac{1}{2} \|w\|^2 + C \sum_n (\xi_n + \hat{\xi}_n)$$

$$\text{s.t. } y^n - w^T x - \xi_n \leq \epsilon$$

$$-(y^n - w^T x) - \hat{\xi}_n \leq \epsilon$$

$$\xi_n \geq 0$$

$$\hat{\xi}_n \geq 0$$

15 Graphical Models

A directed graphical model, also known as a Bayes net or a belief net, is a joint distribution over several variables specified in terms of a conditional distribution for each variable:

$$P(X_1, X_2, \dots, X_N) = \prod_i P(X_i | \text{Parents}(X_i))$$

We draw a Bayes net as a graph with a node for each variable, and edges to each node from its parents. This graph expresses the independence relations implicit in the choice of parents for each node. The parent-child edges must form an acyclic graph.

An undirected graphical model is also a distribution specified in terms of a set of functions of specific variables:

$$P(X_1, X_2, \dots, X_N) = \frac{1}{Z} \prod_m f_m(\mathbf{X}_m)$$

where each \mathbf{X}_m is a subset of $\{X_1, X_2, \dots, X_N\}$. There are no normalization constraints on the individual factors f_m , and the normalization constant Z ensures that the entire joint distribution sums to one:


$$Z = \sum_{X_1, \dots, X_N} \prod_m f_m(\mathbf{X}_m)$$

Suppose that we wish to find the marginal probability of a variable X_i in a directed graphical model:

$$P(X_i) = \sum_{X_1, \dots, X_{i-1}, X_{i+1}, \dots, X_N} \prod_i P(X_i | \text{Parents}(X_i))$$

For binary variables, there are 2^{N-1} terms in this sum. Our goal in this section is to compute this probability more efficiently by using the structure of the network, thus taking advantage of the independence assumptions of the network. The techniques apply to both directed and undirected graphical models. They also apply to the problem of computing conditional probabilities where some variables are known, and we must marginalize over the others.

15.1 Example



To compute $P(X_7|X_2)$, we have

$$P(x_7|x_2) = \frac{1}{Z} \sum_{X_3} \sum_{X_4} \sum_{X_5} \sum_{X_6} P(X_3|x_2)P(X_4|X_3)P(X_5|X_4)P(X_6|X_5)P(x_7|X_6)$$

Suppose every variable X_i is binary, then the summation has $2^4 = 16$ terms. On the other hand, we can use the same trick in dynamic programming by recording every probabilities we have computed for reuse. For example, in above example, if we define

$$f_5(x_5) = \sum_{X_6} P(X_6|X_5 = x_5)P(x_7|X_6) \quad (17)$$

$$f_4(x_4) = \sum_{X_5} P(X_5|X_4 = x_4)f_5(X_5) \quad (18)$$

$$f_3(x_3) = \sum_{X_4} P(X_4|X_3 = x_3)f_4(X_4) \quad (19)$$

$$f_2(x_2) = \sum_{X_3} P(X_3|X_2 = x_2)f_3(X_3) \quad (20)$$

Then the probability above can be computed as

$$P(X_7 = x_7|X_2 = x_2) = \frac{1}{Z} \sum_{X_3} \sum_{X_4} \sum_{X_5} \sum_{X_6} P(X_3|X_2 = x_2)P(X_4|X_3)P(X_5|X_4)P(X_6|X_5)P(X_7 = x_7|X_6) \quad (21)$$

$$= \frac{1}{Z} \sum_{X_3} \sum_{X_4} \sum_{X_5} P(X_3|x_2)P(X_4|X_3)P(X_5|X_4)f_5(X_5) \quad (22)$$

$$= \frac{1}{Z} \sum_{X_3} \sum_{X_4} P(X_3|x_2)P(X_4|X_3)f_4(X_4) \quad (23)$$


$$= \frac{1}{Z} \sum_{X_3} P(X_3|x_2)f_3(X_3) \quad (24)$$

$$= \frac{1}{Z} f_2(x_2) \quad (25)$$


There are 4 sums and each sum needs to compute 2x2 probabilities, so a total of 16 steps.

15.2 Factor Graph

Factor graph is an undirected bipartite graph. There are two types of vertex in a factor graph, factor vertices and variable vertices. Factor vertices correspond to the function f_m in the above example, and each distinct variable vertex corresponds to a distinct variable. If factor function f_m is a function of X_i , its factor vertex is connected to X_i . So the factor graph for above example is,



For more examples,



Note that in the figures above, factor graphs illustrate that the shadowed variable nodes block the information flow from one variable node to another except the last one. In the last example, the two parent nodes are independent, although this cannot be seen from the graph structure. However, the blockage can be read from the table of the factor node in the center. Also note that the last two graphs have same undirected shape, but their factor graphs are different.

15.3 Message Passing (Belief Propagation)

We assume that the factor graph is a tree here. For each variable vertex n and its neighboring factor vertex f_m , the information propagated from n to f_m is,

$$q_{n \rightarrow m}(X_n) = \prod_{m' \in M(n) \setminus \{m\}} r_{m' \rightarrow n}(X_n)$$

where $M(n)$ is the set of factors touching X_n . This table contains the information propagated from variable n to its neighboring factor vertex f_m . For each factor vertex f_m and its neighboring variable n , the

information propagated from f_m to n is,


$$r_{m \rightarrow n}(X_n) = \sum_{\vec{X_m} \setminus X_n} f_m(\vec{X_m}) \prod_{n' \in N(m) \setminus \{n\}} q_{n' \rightarrow m}(X_{n'})$$

where $N(m)$ is the set of variables touching f_m . $\sum_{\vec{X_m} \setminus X_n}$ is the sum is over all variables connected to f_m except X_n . This table contains the information propagated from factor f_m to its neighbor variable n . Note that if variable vertex n is a leaf, $q_{n \rightarrow m} = 1$, and if factor vertex m is a leaf, $r_{m \rightarrow n} = f_m(X_n)$.

The procedure of message passing or belief propagation is first to propagate the information from leaf vertices to the center (i.e., from leaves to internal nodes) by filling in the tables for each message. Once all the messages variable x_n have been computed, the marginal probability of x_n is computed by combining the incoming messages:

$$P(X_n) = \frac{1}{Z} \prod_{m \in M(n)} r_{m \rightarrow n}(X_n)$$

To compute marginal probabilities for all variables, the information is propagated from center back to leaves.



15.4 Running Time

Suppose in a factor graph, there are N variable vertices and M factor vertices. For every variable vertex n , $|M(n)| < k$ and for every factor vertex f_m , $|N(m)| < l$, the running time is,

$$\mathcal{O}((N + M)(k + l)2^{l-1})$$

16 Junction Tree

From last class, we know that

$$q_{n \rightarrow m}(x_n) = \prod_{m' \in M(n) \setminus m} r_{m' \rightarrow n}(x_n)$$

$$r_{m \rightarrow n}(x_n) = \sum_{\vec{x_m} \setminus x_n} f_m(\vec{x_m}) \prod_{n' \in N(m) \setminus \{n\}} q_{n' \rightarrow m}(x_{n'})$$




Figure 3: An example of *max-product*

$q_{n \rightarrow m}(x_n)$ means the information propagated from variable node n to factor node f_m ; $r_{m \rightarrow n}(x_n)$ is the information propagated from factor node f_m to variable node n . And our goal is to compute the marginal probability for each variable x_n :

$$P(x_n) = \frac{1}{Z} \prod_{m \in M(n)} r_{m \rightarrow n}(x_n).$$

The joint distribution of two variables can be found by, for each joint assignment to both variables, performing message passing to marginalize out all other variables, and then renormalizing the result:

$$P(x_i, x_j) = \frac{1}{Z_{\{i,j\}}} \left(\prod_{m \in M(i)} r_{m \rightarrow i}(x_i) \right) \left(\prod_{m \in M(j)} r_{m \rightarrow j}(x_j) \right)$$

In the original problem, the marginal probability of variable x_n is obtained by summing the joint distribution over all the variables except x_n :

$$P(x_n) = \frac{1}{Z} \sum_{x_1} \cdots \sum_{x_{n-1}} \sum_{x_{n+1}} \cdots \sum_{x_N} \prod_m f_m(\vec{x}_m).$$

And by pushing summations inside the products, we obtain the efficient algorithm above.

16.1 Max-Sum

In practice, sometimes we wish to find the set of variables that maximizes the joint distribution $P(x^N) = \frac{1}{Z} \prod_m f_m(\vec{x}_m)$. Removing the constant factor, it can be expressed as

$$\begin{aligned} & \max_{x_1, \dots, x_N} \prod_m f_m(\vec{x}_m) \\ &= \max_{x_1} \dots \max_{x_N} \prod_m f_m(\vec{x}_m) \end{aligned}$$

Figure 3 shows an example, in which the shadowed variables x_j , x_k , and x_l block the outside information flow. So to compute $P(x_i | x_j, x_k, x_l)$, we can forget everything outside them, and just find assignments for inside variables:

$$\max_{\text{inside var}} \prod_m f_m(\vec{x}_m).$$




Figure 4: An example of *tree decomposition*

Like the *sum-product* algorithm, we can also make use of the distributive law for multiplication and push maxs inside the products to obtain an efficient algorithm. We can put max whenever we see \sum in the *sum-product* algorithm to get the *max-sum* algorithm, which now actually is *max-product* (Viterbi) algorithm. For example,

$$r_{m \rightarrow n}(x_n) = \max_{\vec{x}_m \setminus x_n} f_m(\vec{x}_m) \prod_{n' \in N(m) \setminus \{n\}} q_{n' \rightarrow m}(x_{n'})$$

Since products of many small probabilities may lead to numerical underflow, we take the logarithm of the joint distribution, replacing the products in the *max-product* algorithm with sums, so we obtain the *max-sum* algorithm.

$$\max \prod f_m \longrightarrow \max \log \prod f_m \longrightarrow \max \sum \log f_m$$

16.2 Tree Decomposition

If we consider a decision problem instead of a numerical version, the original *max-product* algorithm will be:

$$\text{find } x_1, \dots, x_N \text{ s.t. } \bigwedge_m f_m(\vec{x}_m).$$

We need to find some assignments to make it 1, which can be seen as a reduction from the 3-SAT problem (*constraint satisfaction*). So the problem is \mathcal{NP} -complete in general.

To solve the problem, we force the graph to look like a tree, which is *tree decomposition*. Figure 4 shows an example.

Given a *Factor Graph*, we first need to make a new graph (*Dependency Graph*) by replacing each factor with a clique, shown in Figure 5. Then we apply the *tree decomposition*.

Tree decomposition can be explained as: given graph $G = (V, E)$, we want to find $(\{X_i\}, T)$, $X_i \subseteq V$, $T = \text{tree over } \{X_i\}$. It should satisfy 3 conditions:




Figure 5: Dependency graph for *tree decomposition* (vertex for each variables)




Figure 6: The procedure of *tree decomposition* on a directed graphical model (we can directly get *Dependency Graph* by *moralization*)

1. $\bigcup_i X_i = V$, which means the new graph should cover all the vertex;
2. For $(u, v) \in E$, $\exists X_i$ such that $u, v \in X_i$;
3. If j is on the path from i to k in T , then $(X_i \cap X_k) \subseteq X_j$ (*running intersection property*).

Using this method, we can get the new graph in Figure 5 with $X_1 = \{A, B\}$, $X_2 = \{B, C, D\}$, and $X_3 = \{D, E\}$. The complexity of original problem is $O((N + M)(k + l)2^{l-1})$, with $l = \max_m |N(m)|$. By *tree decomposition*, we can obtain $l = \max_i |X_i|$. Figure 6 shows the procedure to do *tree decomposition* on a directed graphical model.

A new concept is the *treewidth* of a graph:

$$treewidth(G) = \min_{\{X_i\}, T} \max_i |X_i| - 1$$

For example, $treewidth(tree) = 1$, $treewidth(cycle) = 2$, and the worst case, $treewidth(K_n) = n - 1$ (K_n is a complete graph with n vertices). If the treewidth of the original graph is high, the *tree decomposition* becomes impractical.

Actually, finding the best *tree decomposition* is \mathcal{NP} -complete. One practical way is *Vertex Elimination*:

1. choose vertex v (heuristically, choose v with fewest neighbors);
2. create X_i for v and its neighbors;
3. remove v ;
4. connect v 's neighbors;
5. repeat the first four steps until no new vertex.

Vertex Elimination cannot ensure to find the optimum solution. Figure 7 shows an example of this method on a single cycle.





Figure 7: An example of *Vertex Elimination* on a single cycle

Another way to do *tree decomposition* is *Triangulation*:

1. find cycle without chord (shortcut);
 2. add chord;
 3. repeat the first two steps until triangulated (no cycles without chords).
- The cliques in the new graph are X_i in the *tree decomposition*.


17 Expectation Maximization

In last lecture, we introduced Tree Decomposition. Till now, we have covered a lot as regards how to do inference in a graphical model. In this lecture, we will move back to the learning part. We will consider how to set parameters for the variables.



17.1 Parameter Setting: An Example


In the discrete case, we set the parameters just by counting how often each variable occurs. However, we may not know the value of some variables. Thus, in the following, we will discuss learning with hidden (latent) variables. The simplest model is shown below. This model has been used in the Aspect Model



for probabilistic Latent Semantic Analysis (pLSA). The variable's value is called an aspect. pLSA has been widely used in information retrieval. In this example, let x_1 and x_2 respectively denote the document ID and word ID. Then, there is a sequence of pair (x_1, x_2) , e.g., (1, "the"), (1, "green"), ..., (1000, "the"). In this context, we may have various tasks, e.g., to find the words which co-occur, to find the words on the same topic, or to find the documents containing the same words.

Now, we will introduce the term *cluster*. The reasons why we need the cluster representation are as followed: (1) There are a large amount (e.g., 10,000) of documents, each of which is formed of a large amount (e.g., 10,000) of words. Without a cluster representation, we have to handle a huge query table with too many (say, $10,000^2$) entries. That makes any query difficult. (2) If we still set the parameters just by counting how often each variable occurs, then there is a risk of over-fitting the individuals (i.e., the pair of (document, word)). Because of them, we need to do something smart: clustering. Recall the graphical model displayed above, the hidden (latent) variable z is just the cluster ID.

Note that $x_1 \perp x_2 | z$. Therefore, the joint distribution $p(x_1, x_2) = \sum_z p(z)p(x_1|z)p(x_2|z)$. As shown in the figure below, now we will not directly compute each entry to obtain the $10,000 \times 10,000$ query table $P(x_1, x_2)$. Instead, we maintain low-rank matrices $P(z)$, $P(x_1|z)$ and $P(x_2|z)$.



Now, we have a set of observed variables $\mathbf{X} = \{(x_1^1, x_2^1), (x_1^2, x_2^2), \dots\}$, a set of hidden variables $\mathbf{Z} = \{z_1, z_2, \dots\}$ and a set of parameters $\theta = \{\theta_z, \theta_{x_1|z}, \theta_{x_2|z}\}$. Note that x_1^i are i.i.d. variables, and the same for x_2^i . To choose θ , we maximize the likelihoods (MLE): $\max_{\theta} P(\mathbf{X}; \theta)$.

$$\theta = \operatorname{argmax}_{\theta} \prod_n P_{\theta}(x_1^n, x_2^n) = \operatorname{argmax}_{\theta} \prod_n \sum_z p(z)p(x_1^n|z)p(x_2^n|z) = \operatorname{argmax}_{\theta} \sum_n \log \sum_z p(z)p(x_1^n|z)p(x_2^n|z) \quad (26)$$

If there is no hidden variable z , we will just count, instead of summing over z . However now, we need to sum over z and find the maximum of the above objective function, which is not a closed-form expression. Thus, it is not feasible to directly set the derivative to zero. To solve this tough optimization problem, we will introduce the Expectation-Maximization (EM) algorithm.

17.2 Expectation-Maximization Algorithm

The EM algorithm is an elegant and powerful method for finding maximum likelihood solutions for models with hidden (latent) variables. It breaks down the potentially tough problem of MLE into two stages. The basic idea is shown below.

E-step:

Guess \mathbf{z}

M-step:

MLE to fit θ to \mathbf{X}, \mathbf{Z}

Following the above example, we present the EM algorithm in detail.

REPEAT

E-step:

for $n = 1 \dots N$

for $z = 1 \dots K$

$p(z, n) = \theta_z(z) \cdot \theta_{x_1|z}(x_1^n|z) \cdot \theta_{x_2|z}(x_2^n|z)$

$sum += p(z, n)$

for z

$$\begin{aligned}
& p(z, n) \leftarrow \frac{p(z, n)}{sum} \\
& \text{(An alternative:} \\
& \quad ec(z) += p(z, n) \\
& \quad ec(z, x_1^n) += p(z, n) \\
& \quad ec(z, x_2^n) += p(z, n) \quad \text{)}
\end{aligned}$$

M-step:

$$\begin{aligned}
& \text{for } z \\
& \quad ec(z) = \sum_{n=1}^N p(z, n) \\
& \quad \theta_z \leftarrow \frac{ec(z)}{N} \\
& \text{for } z, x_1 \\
& \quad ec(z, x_1) = \sum_{n=1}^N I(x_1^n = x_1) p(z, n) \\
& \quad \theta_{x_1|z} = \frac{ec(z, x_1)}{ec(z)} \\
& \text{for } z, x_2 \\
& \quad ec(z, x_2) = \sum_{n=1}^N I(x_2^n = x_2) p(z, n) \\
& \quad \theta_{x_2|z} = \frac{ec(z, x_2)}{ec(z)}
\end{aligned}$$

UNTIL convergence

where *sum* is for normalization and *ec*(\cdot) denotes the expected count, which is not a real count but an average on what we think z is. Namely, this count is probabilistic. The intuition is to assign some credit to each possible value. Also note that $I(\cdot)$ is an indicator function (return 1/0 if the condition is true/false). In the following, we will derive how to approximate the maximum of the likelihood by maximizing the joint probability's log likelihood iteratively through E-M steps. For the example present in Sec.2, now let us go further using the same formulation with Eqn. (26).

$$\begin{aligned}
& \theta = \operatorname{argmax}_{\theta} Q(\theta; \theta^{old}) \\
& = \operatorname{argmax}_{\theta} E_{p(z|x, \theta^{old})} \log p(\mathbf{X}, \mathbf{Z}) \\
& = \operatorname{argmax}_{\theta} E_{p(z|x, \theta^{old})} [\log \prod_n p(x_1^n, x_2^n, z^n)] \\
& = \operatorname{argmax}_{\theta} E [\log \prod_n p(z^n) \cdot p(x_1^n | z^n) \cdot p(x_2^n | z^n)] \\
& = \operatorname{argmax}_{\theta} E [\sum_n \log p(z^n) + \sum_n \log p(x_1^n | z^n) + \sum_n \log p(x_2^n | z^n)] \\
& \left(= \operatorname{argmax}_{\theta} E [\sum_n \log \theta_{z^n} + \sum_n \log \theta_{x_1^n | z^n} + \sum_n \log \theta_{x_2^n | z^n}] \right) \\
& = \operatorname{argmax}_{\theta} E [\sum_k \sum_n I(z^n = k) \log p(z = k) + \sum_k \sum_n I(x_1^n = x_1 | z^n = k) \log p(x_1 | z = k) \\
& \quad + \sum_k \sum_n I(x_2^n = x_2 | z^n = k) \log p(x_2 | z = k)] \\
& \text{(sum over values } \mathbf{z} \text{ can take and individually group together)} \\
& = \operatorname{argmax}_{\theta} E [\sum_k c(z = k) \log p(z = k) + \sum_k c(x_1, z = k) \log p(x_1 | z = k) + \sum_k c(x_2, z = k) \log p(x_2 | z = k)] \\
& \left(= \operatorname{argmax}_{\theta} \sum_k E[c(z = k)] \log \theta_{z=k} + \sum_k E[c(x_1, z = k)] \log \theta_{x_1|z=k} + \sum_k E[c(x_2, z = k)] \log \theta_{x_2|z=k} \right. \\
& \quad \left. = \operatorname{argmax}_{\theta} \sum_k ec(z) \log \theta_{z=k} + \sum_k ec(x_1, z) \log \theta_{x_1|z=k} + \sum_k ec(x_2, z) \log \theta_{x_2|z=k} \right)
\end{aligned} \tag{27}$$

Therefore, $\theta = \frac{1}{sum_0} ec(z)$, $\theta_{x_1|z} = \frac{1}{sum_1} ec(x_1, z)$, and $\theta_{x_2|z} = \frac{1}{sum_2} ec(x_2, z)$, but make sure that normalization is done (sum to 1). Notably, $c(\cdot)$ denotes the count and $ec(\cdot)$ denotes the expected count. Also note that $E[c(z = k)] = ec(z)$ which we have mentioned in the EM algorithm flow, and similar for $ec(x_1, z)$ and $ec(x_2, z)$.

17.3 EM Algorithm in General

Now, we will give a general derivation for the EM algorithm. The denotation will be the same with the above. Similarly, we have $\theta = \operatorname{argmax}_{\theta} Q(\theta; \theta^{old}) = \operatorname{argmax}_{\theta} E_{p(z|x, \theta^{old})} \log p(\mathbf{X}, \mathbf{Z})$. Let us focus on the


objective function Q .

$$\begin{aligned}
Q(\theta; \theta^{old}) &= E_{p(z|x, \theta^{old})} \log p(\mathbf{X}, \mathbf{Z}) \\
&= E_{p(z|x, \theta^{old})} \log [p(\mathbf{Z}|\mathbf{X}) \cdot p(\mathbf{X})] \\
&= E_{p(z|x, \theta^{old})} [\log p(\mathbf{Z}|\mathbf{X}) + \log p(\mathbf{X})] \\
&= E \left[\log \frac{p(\mathbf{Z}|\mathbf{X})}{p_{\theta^{old}}(\mathbf{Z}|\mathbf{X})} \cdot p_{\theta^{old}}(\mathbf{Z}|\mathbf{X}) \right] + \log p(\mathbf{X}) \\
&\quad (\text{make it look like } K-L \text{ divergence}) \\
&= E \left[-\log \frac{p_{\theta^{old}}(\mathbf{Z}|\mathbf{X})}{p(\mathbf{Z}|\mathbf{X})} \right] - E [-\log p_{\theta^{old}}(\mathbf{Z}|\mathbf{X})] + \log p(\mathbf{X}) \\
&= -D(\mathbf{Z}|\mathbf{X}, \theta^{old} \parallel \mathbf{Z}|\mathbf{X}, \theta) - H(\mathbf{Z}|\mathbf{X}, \theta^{old}) + L(\theta)
\end{aligned} \tag{28}$$

where D, H, L are respectively the K-L divergence, the entropy and the likelihood. Note that our objective is to maximize the likelihood $\log p(\mathbf{X})$. It does not have Z inside, so it can be put out of $E(\cdot)$. Now, we write down $L(\theta)$ with simplified notations:


$$L(\theta) = Q(\theta; \theta^{old}) + H(\theta^{old}) + D(\theta^{old} \parallel \theta) \tag{29}$$

where Q, H and D are all dynamic functions. The approximation can be illustrated in the space of parameter θ , as shown schematically in the figure below. Here the red curve depicts $L(\theta)$ (incomplete data) which we



wish to maximize. We start with some initial parameter value θ^{old} , and in the first E step we evaluate the distribution of $Q(\theta; \theta^{old}) + H(\theta^{old})$, as shown by the blue curve. Since the K-L divergence $D(\theta^{old} \parallel \theta)$ is always positive, the blue curve gives a lower bound to the red curve $L(\theta)$. And $D(\theta^{old} \parallel \theta)$ just gives the gap between the two curves. Note that the bound makes a tangent contact with $L(\theta)$ at θ^{old} , so that both curves have the same gradient and $D(\theta^{old} \parallel \theta^{old}) = 0$. Thus, $L(\theta^{old}) = Q(\theta^{old}; \theta^{old}) + H(\theta^{old})$. Besides, the bound is a convex function having a unique maximum at $\theta^{new} = \arg\max_{\theta} [Q(\theta; \theta^{old})]$. In the M step, the bound is maximized giving the value θ^{new} , which also gives a larger value of $L(\theta)$ than θ^{old} : $L(\theta^{new}) = Q(\theta^{new}; \theta^{old}) + H(\theta^{old}) + D(\theta^{old} \parallel \theta^{new})$. In practice, during the beginning iterations, this point is usually still far away from the maximum of $L(\theta)$. However, if we run one more iteration, the result will get better. The subsequent E step then constructs a bound that is tangential at θ^{new} as shown by the green curve. Iteratively, the maximum will be accessed in the end, in a manner kind of similar with gradient ascent. In short, there are two properties for the EM algorithm: (1) the performance gets better step by step. (2) it will converge. At last, it should also be emphasized that EM is not guaranteed to find the global maximum, for there will generally be multiple local maxima. Being sensitive to the initialization, EM may not find the largest of these maxima.

In this lecture, we quickly go through the details of the EM algorithm, which maximizes L through maximizing Q at each iteration step. Then, the remaining problems are how to compute Q , and exactly how to compute $p(\mathbf{Z}|\mathbf{X}, \theta^{old})$.



17.4 Gradient Ascent ($\frac{\partial L}{\partial \theta}$)


Gradient Ascent $\frac{\partial L}{\partial \theta}$ means the gradient of likelihood function w.r.t. parameters.

Using the gradient ascent, our new parameter θ^{new} is given by :

$$\theta^{new} = \theta^{old} + \eta \frac{\partial L}{\partial \theta}$$

EM (expectation maximization) generally finds a local optimum of θ more quickly, because in this method we optimize the q-step before making a jump. It is similar to saying that we make a big jump first and then take smaller steps accordingly. But for Gradient Ascent method, the step size varies with the likelihood gradient, and so it requires more steps when the gradient is not that steep.


We are required to find $E_{P(Z|X, \theta)}[Z|X]$



$$\text{Our } \theta = \begin{cases} P(Z) \\ P(X_1|Z) \\ P(X_2|Z) \end{cases}$$

which is a long list of all the possible probabilities, after unfolding each of them.

We should make sure that the parameters $(\theta = [\theta_1, \theta_2], \text{ for only 2 parameters})$ always stays within the straight line shown below. It should never go out of the line.



$$P(Z = k) = \theta_k = \frac{e_k^\lambda}{\sum_{k'} e_{k'}^\lambda}$$

In this case, the gradient ascent is used to find the new value of λ

$$\lambda_{new} = \lambda_{old} + \eta \frac{\partial L}{\partial \lambda}$$


17.5 Newton's Method

$$\theta^{new} = \theta^{old} + (\nabla_{\theta}^2 L)^{-1} (\nabla_{\theta} L)$$

In Newton's Method, we approximate the curve with a quadratic function, and we jump to the maximum of the approximation in each step.

Differences between Newton's Method and EM :

1. Newton's Method takes a lot of time, because we need to calculate the Hessian Matrix, which is the 2nd derivative.
2. Since there is no KL divergence in Newton's Method, there is always a chance that we take a big jump, and arrive at a point far away from the global optimum and in fact worse than where we started.



17.6 Variational Method

In this method, at first we trace the Likelihood function by a new parameter, then fix and optimize that parameter, and then once again start the iteration, until we get the optimum value of the Likelihood function.

17.7 Mixture of Gaussians

$$\begin{aligned} P(X) &= \sum_k P(Z = k) N(X | \mu_k, \Sigma_k) \\ &= \sum_k \lambda_k N(X | \mu_k, \Sigma_k), \end{aligned}$$

where μ and Σ are the mean vector and co-variance matrix respectively.

$$N(X | \mu, \Sigma) = \frac{e^{-0.5(X-\mu)^T \Sigma^{-1} (X-\mu)}}{(2\pi)^{D/2} |\Sigma|^{0.5}} \quad (30)$$



Here X and μ are vectors, Σ is a 2-D matrix, and D is the dimensionality of data X . The left side side of equation (1) refers to :

$$N \left(\begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_N \end{bmatrix} \middle| \mu, \Sigma \right)$$

For 1-D data,

$$N(X | \mu, \sigma) = \frac{e^{-0.5(X-\mu)^2 / \sigma^2}}{(2\pi)^{0.5} \sigma}$$

Equation (1) is similar to writing as $f(X) = X^T A X$; wherein we are stretching the vector space.



For continuous variables the equation :

$$r_{m \rightarrow n} = \sum_{\vec{X}_m \setminus X_n} f(\vec{X}_m) \prod_{n \in N(m) \setminus n} q_{n \rightarrow m}$$

changes to

$$r_{m \rightarrow n} = \int f(\vec{X}_m) \prod_{n \in N(m) \setminus n} q_{n \rightarrow m} d(\vec{X}_m \setminus X_n)$$

The MLE estimates of the parameters of a Gaussian are as follows:

$$\begin{aligned} \mu^* &= \frac{\sum_n X_n}{N} \\ \sigma^* &= \sqrt{\frac{\sum_n (X_n - \mu^*)^2}{N}} \\ \Sigma &= \frac{\sum_n (X_n - \mu^*)(X_n - \mu^*)^T}{N}, \text{ where } \Sigma \text{ is the covariance matrix} \\ \Sigma_{ij} &= \frac{\sum_n (X_n^i - \mu_i)(X_n^j - \mu_j)}{N} = \text{covariance}(X_i, X_j) \end{aligned}$$

Θ (parameters of the model) = λ, μ_k, Σ , where λ and μ_k are vectors

E-step

for $n = 1, 2, \dots, N$

$$\begin{aligned} P(Z|X^n) &= \frac{P(Z)P(X^n|Z)}{P(X^n)} \\ &= \frac{(\lambda_k/z_k) \exp[-\frac{1}{2}(X^n - \mu_k)^T \Sigma^{-1}(X^n - \mu_k)]}{\sum_{k'} (1/z_{k'}) \exp[-\frac{1}{2}(X^n - \mu_{k'})^T \Sigma^{-1}(X^n - \mu_{k'})]} \end{aligned}$$

M-step

for $k = 1, 2, \dots, K$ (total number of hidden variables)

$$\begin{aligned}\lambda_k &= \frac{\sum_{n=1}^N P(Z = k|X^n)}{N} \\ \mu_k &= \frac{\sum_n P(Z = k|X^n) X^n}{\sum_n P(Z = k|X^n)} \\ \Sigma_k &= \frac{\sum_n P(Z = k|X^n) (X^n - \mu_k)(X^n - \mu_k)^T}{\sum_n P(Z = k|X^n)}\end{aligned}$$

18 Sampling

We have already studied how to calculate the probability of a variable or variables using the message passing method. However, there are some times when the structure of the graph is too complicated to be calculated. The relation between the diseases and symptoms is a good example, where the variables are all mixed together and brings the graph a high tree width. Another case is that of continuous variable, where during the message passing,

$$r_{m \rightarrow n} = \int f(\vec{x}_m) \prod_{n'} q_{n' \rightarrow m}(x_{n'}) d(\vec{x}_m \setminus x_n).$$

If this integration can not be calculated, what can we do to evaluate the probability of variables? This is what sampling is used for.

18.1 How to Sample a Continuous Variable: Basics

Now let us forget the above for a moment, say if we want to sample for a continuous variable, how can we ensure that the points we pick up satisfy the distribution of that variable? This question is easy for variables with uniform distribution, since we can generate random numbers directly using a computer. For some complicated distributions, we could use the inverse of cumulative distribution function (CDF) to map the uniform distribution onto the required distribution to generate samples, where the CDF for a distribution with probability distribution function (PDF) of P is

$$\text{CDF}(x) = \int_{-\infty}^x P(t) dt.$$

For example, if we want to sample from a variable with standard normal distribution, the points we pick up are calculated from

$$X = \text{erf}^{-1}(x),$$

where x is drawn from a uniform distribution, and

$$\text{erf}(x) = \int_0^x \mathcal{N}(t, 0, 1) dt,$$

We could play the same trick for many other distributions. However, there are some distributions which do not have a closed-form integral to calculate their CDF, which makes the above method fail. Under such conditions, we could turn to a framework called *Markov chain Monte Carlo* (MCMC).

18.2 The Metropolis-Hastings Algorithm

Before discussing this method in more detail, let us review some basic properties of Markov chains. A first-order Markov chain is a series of random variables such that each variable depends only on its previous state, that is,

$$x^t \sim P(x^t | x^{t-1}).$$

Our goal is to find a Markov chain which has a distribution similar to a given distribution which we want to sample from, so that by running the Markov chain, we get results as if we were sampling from the original distribution. In other words, we want to have the Markov chain that eventually be able to 1) explore over the entire space of the original distribution, 2) reflect the original PDF.

The general algorithm for generating the samples is called *the Metropolis-Hastings algorithm*. Such an algorithm draws a candidate

$$x' \sim Q(x'; x^t),$$

and then accepts it with probability

$$\min \left\{ 1, \frac{P(x')Q(x^t; x')}{P(x^t)Q(x'; x^t)} \right\}.$$

The key here is the function Q , called *proposed distribution* which is used to reduce the complexity of the original distribution. Therefore, we have to select a Q that is easy to sample from, for instance, a Gaussian function. Note that there is a trade-off on choosing the variance of the Gaussian, which determines the step size of the Markov chain. If it is too small, it will take a long time, or even make it impossible for the states of the variable to go over the entire space. However, if the variance is too large, the probability of accepting the new candidate will become small, and thus it is possible that the variable will stay on the same state for ever. All these extremes will make the chain fail to simulate the original PDF.

If we sample from P directly, that is $Q(x'; x^t) = P(x')$, we have

$$\frac{P(x')Q(x^t; x')}{P(x^t)Q(x'; x^t)} = 1,$$

which means that the candidate we draw will always be accepted. This tells us that Q should approximate P .

By the way, how do we calculate $P(x)$? There are two cases.

- Although we cannot get the integration of $P(x)$, $P(x)$ itself is easy to compute.
- $P(x) = f(x)/Z$, where $Z = \int f(x)dx$ is what we do not know. But since we know $f(x) = ZP(x)$, we could just substitute $f(x)$ instead of $P(x)$ in calculating the probability of acceptance of a candidate.

18.3 Proof of the method

In this section, our goal is to prove that the Markov chain generated by the Metropolis-Hastings algorithm has a unique stationary distribution. We will first introduce some basics about the definition of the stationary distribution, and the method to prove this “stationary”. Then we will apply those knowledge to accomplish our goal.

1. Stationary distribution

A distribution with respect to a Markov chain is said to be *stationary* if the distribution remains the same before and after taking one step in the chain, which could be denoted as

$$\Pi^t = T \times \Pi^{t-1} = \Pi^{t-1},$$

or

$$\Pi_i = \sum_j T_{ij} \Pi_j, \quad \forall i,$$

where Π is a vector which contains the stationary distribution of the state of the variable in each step with its element $\Pi_i = P(x = i)$, and T is the transition probability matrix where its element $T_{ij} = P(x^t = i | x^{t-1} = j)$ denotes the probability that the variable transits from state j to i . For example, the two Markov chains in [8a](#) and [8b](#) all have a stationary distribution $\Pi = \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix}$.




Figure 8: Example Markov Chains

The stationary distribution of a Markov chain could be calculated by solving the equation

$$\begin{cases} T\Pi = \Pi \\ \sum_i \Pi_i = 1 \end{cases}.$$

Note that there might be more than one stationary distribution for a Markov chain. A rather simple example would be a chain with a identity transition matrix shown in 8c.

If a Markov chain has a stationary distribution and the stationary distribution is unique, it is ensured that it will converge eventually to that distribution no matter what the original state the chain is.

2. Detailed Balance: Property to ensure the stationary distribution

Once we know a Markov chain is uniquely stationary, then we can use it to sample from a given distribution. Now, we will see a sufficient (but not necessary) condition for ensuring a Π is stationary, which is a property of the transition matrix called *Detailed Balance*. The definition of such a property is

$$\forall i, j, \quad T_{ij}\Pi_j = T_{ji}\Pi_i,$$

which means $P_{i \rightarrow j} = P_{j \rightarrow i}$, and is also called reversibility due to the symmetry of the structure.

Starting from such definition, we have

$$\forall i, \quad \sum_j T_{ij}\Pi_j = \sum_j T_{ji}\Pi_i = \Pi_i \sum_j T_{ji}.$$

Note that $\sum_j T_{ji} = 1$, we come up with

$$\forall i, \quad \sum_j T_{ij}\Pi_j = \Pi_i \cdot 1 = \Pi_i,$$

which is exactly the second definition of stationary distribution we have just discussed. Therefore, if a distribution makes the transition matrix of a Markov chain satisfy detailed balance, that distribution is the stationary distribution of that chain. Note that although a periodic Markov chain like that shown in 8d satisfies detailed balance, we do not call it stationary. This because it will not truly converge and thus is not guaranteed to approximate the original PDF. What is more, it is often the case that we add a probability like shown in 8e to avoid such a periodic circumstance.

Note that the Detailed Balance does not ensure the uniqueness of the stationary distribution of a Markov chain. However, such uniqueness is necessary, or the Markov chain would not go to the PDF we want. What we could do is that, when we construct the chain at the very beginning, we make the chain such that 1) any state is reachable for any other and 2) the chain is aperiodic. Under that condition, we could ensure the uniqueness of the stationary distribution.

3. Final proof

Now, let us be back to the Metropolis-Hastings algorithm and prove that the transition matrix of its Markov chain has the detailed balance property. If we can prove that, it is obvious that such a Markov chain has a unique stationary distribution.

According to the Metropolis-Hastings algorithm, the transition probability of the Markov chain of the algorithm is

$$T(x'; x) = Q(x'; x) \cdot \min \left\{ 1, \frac{P(x')Q(x; x')}{P(x)Q(x'; x)} \right\}$$

If $x' = x$, then it is automatically detailed balancing due to the symmetry of the definition of detailed balance. To be specific, the condition of detailed balance, which is

$$\forall ij, T_{ij}\Pi_j = T_{ji}\Pi_i,$$

will always be valid if $i = j$, which is just the case of $x' = x$.

For the circumstances that $x' \neq x$, by using the distributive property of multiplication, the transition probability is derived as,

$$T(x'; x) = \min \left\{ Q(x'; x), \frac{P(x')Q(x; x')}{P(x)} \right\}.$$

Multiply both sides by $P(x)$, it turns out that

$$\begin{aligned} T(x'; x)P(x) &= \underbrace{\min \{Q(x'; x)P(x), P(x')Q(x; x')\}}_{\text{symmetric for } x \text{ \& } x'} \\ &= T(x; x')P(x') \end{aligned}$$

Therefore, we proved the detailed balance of the transition matrix, and thus the Markov chain of the Metropolis-Hastings algorithm does have a stationary distribution, which means that we could use such a Markov chain to simulate the original PDF.

18.4 Gibbs Sampling

Now, back to the very first problem of this class, we want to get the result of

$$P(x_k) = \sum_{\vec{x} \setminus x_k} \frac{1}{Z} \prod_m f(\vec{x}_m)$$

without knowing Z . We could use the Gibbs Sampling, shown in Algorithm 1, where x_{-k} means all the variables x except x_k . Note that the Gibbs Sampling is actually a particular instance of the Metropolis-Hastings algorithm where the new candidate is always accepted. This is proved as follows.

Algorithm 1: Gibbs Sampling

for $k = 1 \dots K$ **do**
 $x_k \sim P(x_k | x_{\neg k}) = \frac{1}{Z'} \prod_{m \in M(k)} f(\vec{x}_m);$

Substitute

$$\begin{cases} P(x) = P(x_{\neg k})P(x_k | x_{\neg k}) \\ P(x') = P(x'_{\neg k})P(x'_k | x'_{\neg k}) \\ Q(x'; x) = P(x'_k | x_{\neg k}) \\ Q(x; x') = P(x_k | x'_{\neg k}) \\ x'_{\neg k} = x_{\neg k} \end{cases}$$

into the probability of the accepting the new candidate, we have

$$\begin{aligned} \frac{P(x')Q(x; x')}{P(x)Q(x'; x)} &= \frac{P(x'_{\neg k})P(x'_k | x'_{\neg k})P(x_k | x'_{\neg k})}{P(x_{\neg k})P(x_k | x_{\neg k})P(x'_k | x_{\neg k})} \\ &= 1. \end{aligned}$$

Therefore, the Gibbs Sampling algorithm will always accept the candidate.

18.5 Gibbs sampling with Continuous Variables.

What about the continuous case particularly where sampling is hard? Here we can have a second sampling step:

repeat

 for $k \leftarrow 1 \dots K$

$$\vec{x}_k \sim P(\vec{x}_k | \vec{x}_{\neg k}) = \begin{cases} \frac{1}{Z} \prod_{m \in M(k)} f_m(\vec{x}_m) \\ \text{Metropolis-Hastings} \end{cases}$$

18.6 EM with Gibbs Sampling.

E-step Sample each variable.

M-step Use hard (fixed value of the sample) assignments from sampling in:

$$\lambda_k = \frac{\sum_n I(z^n = k)}{N} = \frac{N_k}{N} \quad (31)$$

In this I is counting how many points are assigned k and we denote this sum as N_k .

$$\mu_k = \frac{\sum_n I(z^n = k) \vec{x}^n}{\sum_n I(z^n = k)} = \frac{\sum_n I(z^n = k) \vec{x}^n}{N_k} \quad (32)$$

The computation of Σ is similar.

18.6.1 Some problems.

What we have above is an approximation to what EM is doing. If we put the computations in expectation it is the same as EM. One problem with EM in general is that if a probability of a cluster hits zero it never comes back. In sampling we can get unlucky and get all zeros and never get that cluster back.

18.6.2 Some advantages.

With sampling we can apply EM to complicated models. For example, a factor graph with cycles or highly connected components. Sampling can be improved by sampling L times in the E-step.

In practice a short cut is taken and we combine the E and the M steps and take advantage of samples immediately:

for $n \leftarrow 1 \dots N$
 sample z^n

$$\lambda_k \leftarrow \lambda_k + \frac{1}{N} I(z^n = k) - \frac{1}{N} I(z_{\text{old}}^n = k)$$

or

$$\lambda_k \leftarrow \frac{N_k}{N}$$

$$\hat{\mu} \leftarrow \hat{\mu} + I(z^n = k) \vec{x}^n - I(z_{\text{old}}^n = k) \vec{x}^n$$

$$\mu_k \leftarrow \frac{\hat{\mu}}{\lambda_k N}$$

We are keeping a running count and updating it as we sample. New observations move data points from their current cluster to a new cluster so we subtract the old observation and add the new one. This technique is widely used and easy to implement.

If we run this long enough it should correspond to the real distribution of hidden variables $P(z^1 \dots z^N | \vec{x}^1 \dots \vec{x}^N)$ — but what does that mean here? Although we have defined $P(z^n | \vec{x}^n; \lambda, \mu, \Sigma)$, the parameters λ , μ , and Σ are changing as sampling progresses. If we run long enough we know that, because the problem is symmetric, $P(z^n = k) = \frac{1}{K}$.

To make sense of this we will make λ , Σ , and μ variables in our model with some distribution $P(\lambda)$. These variables have no parents so we can pick this distribution. We have seen this before and choose to use the Dirichlet distribution so we let $P(\lambda) = \text{Dir}(\alpha)$. Any point is just as likely to be μ so we let $P(\mu) = 1$ and similarly $P(\Sigma) = 1$.

We can take $P(\lambda)$ into account when we sample:

$$z^n \sim \frac{\lambda_k \mathcal{N}(\vec{x}; \mu_k \Sigma_k)}{Z}$$

When we have seen data the probability of the next is:

$$P(z^{N+1} | z^1 \dots z^N) = \frac{c(k) + \alpha}{N + K\alpha}$$

Applying this we have:

$$z^n \sim \frac{\frac{N_k + \alpha}{N + K\alpha} \mathcal{N}(\vec{x}; \mu_k \Sigma_k)}{Z}$$

$$\lambda_k = \frac{N_k + \alpha}{N + K\alpha}$$

Now λ_k can never go all the way to zero. Now if we run long enough we will converge to the real distribution. We have a legitimate Gibbs sampler (states with just relabeling have equal probability). We are sampling with

$$\lambda_k = \frac{1}{Z} \int P(z^1 \dots z^N | \lambda) P(\lambda) d\lambda$$

The λ is integrated out, giving what is called a collapsed Gibbs sampler.

A remaining question is: when should we stop? If we plot $P(x)$ as we iterate we should see a general trend upward with some small dips and then the curve levels off. But, it could be that in just a few steps a better state could continue the upward trend.

The real reason for doing this Gibbs sampling is to handle a complicated model. One example of that could be a factor graph of diseases and symptoms due to its high tree width. If we try to use EM directly

we have exponential computation for the expected values of the hidden variables. With Gibbs we avoid that problem.

19 Error Bounds

19.1 Leave One Out

Assume the data points are linear separable and iid. And the training and test data are D dimensions, while D is unknown, meaning our problem setting is distribution free.

$$Q(x, a) = \begin{cases} 1 & y \neq \hat{y}(x) \\ 0 & otherwise \end{cases}$$

Where Q is called loss. Risk is defined as expected loss:

$$R(a) = \int Q(x, a)P(x)dx$$

Now we need to know the maximum value of the expectation of the risk over possible training sets. This means we need to calculate $E(R(a)) \leq ?$; In the past lecture notes, we know $W = \phi = \sum_i \alpha_i X_i Y_i$ where X_i is the i th data point.

$$Y_i = \begin{cases} -1 \\ 1 \end{cases}$$

In order to estimate the risk for the function $Q(z, a)$, we need to use the following statistics: exclude the first vector z_1 from the sequence and obtain the function that minimizes the empirical risk for the remaining $l - 1$ elements of the sequence for the given sequence z_1, z_2, \dots, z_l ;

Let the function be $Q(z_1, a_{l-1}|z_1)$. In this notation, we indicate that the vector z_1 was excluded from the sequence. We use this excluded vector for computing the value $Q(z_1, a_{l-1}|z_1)$.

Now we excluded the second z_2 from the sequence (while the first vector is retained) and compute the value $Q(z_2, a_{l-1}|z_2)$.

In this manner, we compute the values for all the vectors and calculate the number of errors in the leave one out procedure;

$$L(z_1, z_2, \dots, z_l) = \sum_{i=1}^l Q(z_i, a_{l-1}|z_i)$$

19.2 One Way to Compute an Error Bound

We will use L as an estimate for the expectation of the function $Q(z, a_l)$ that minimizes the empirical risk.

First we need to prove $E(L(z_1, \dots, z_{l+1})/(l+1)) = ER(a_l)$.

The proof consists of the following chain of transformations:

$$\begin{aligned} EL(z_1, z_2, \dots, z_{l+1})/(l+1) &= \int 1/(l+1) \sum_{i=1}^{l+1} Q(z_i, a_l|z_i) dP(z_1), \dots, dP(z_{l+1}) \\ &= \int 1/(l+1) \sum_{i=1}^{l+1} (\int Q(z_i, a_l|z_i) dP(z_i)) \\ &= E(1/(l+1) \cdot \sum_{i=1}^{l+1} R(a_l|z_i)) \\ &= ER(a_l) \end{aligned}$$

We will introduce the essential support vectors: Essential support vectors are the vectors that, if removed, would result in learning a different SVM. Indeed, if the vector x_i is not an essential support vector, then there exists an expansion of the vector ϕ that defines the optimal hyperplane that doesn't contain the vector x_i .

Since the optimal hyperplane is unique, removing this vector from the training set doesn't change it. Therefore in the leave one out method it will be recognized correctly.

Thus the leave one out method recognizes correctly all the vectors. Therefore the number of $L(z_1, \dots, z_{l+1})$ of errors in the leave one out method doesn't exceed K_{l+1} , the number of the essential support vectors; that is,

$$L(z_1, \dots, z_{l+1}) \leq K_{l+1}$$

Then the equation we will get:

$$ER(a_l) = EL(z_1, z_2, \dots, z_{l+1})/(l+1) \leq EK_{l+1}/(l+1)$$

19.3 Another Way to Compute an Error Bound

In order to use another way to prove the bound of the number of errors in the leave on out estimator for the optimal hyperplanes, we use the following method.

Suppose we are given the training set:

$$(x_1, y_1), \dots, (x_{l+1}, y_{l+1})$$

And the maximum of $W(a)$ in the area $a \geq 0$ is achieved at the vector $a^0 = (a_1^0, \dots, a_l^0)$. Let the vector

$$\phi_0 = \sum_{i=1}^a a_i^0 x_i y_i$$

Define the optimal hyperplane passing through the original, where we enumerate the support vectors with $i = 1 : a$

Let us denote by a^p the vector providing the maximum for the functional $W(a)$ under the constraint $a_p = 0$

$$a_i \geq 0 \quad \text{where } (i \neq p)$$

Let the vector $\phi_p = \sum_{i=1}^a a_i^p x_i y_i$

The above defines the coefficients of the corresponding separating hyperplane passing through the origin.

Now denote by W_0^p the value of the $W(a)$ for:

$$a_i = a_i^0 \quad \text{where } (i \neq p)$$

$$a_p = 0;$$

Consider the vector a^p that maximize the function $W(a)$ under the above constraint. The following obvious inequality is valid:

$$W_0^p \leq W(a^p)$$

On the other hand the following inequality is true:

$$W(a^p) \leq W(a^0)$$

Therefore the inequality:

$$W(a^0) - W(a^p) \leq W(a^0) - W_0^p$$

is valid.

Now let us rewrite the right-hand side of the inequality in the explicit form;

$$\begin{aligned} W(a^0) - W(a^p) &= \sum_{i=1}^a a_i^0 - 1/2 * (\phi_0 * \phi_0) - \left(\sum_{i=1}^a a_i^0 - a_p^0 - 1/2 * ((\phi_0 - a_p^0 * y_p * x_p) * (\phi_0 - a_p^0 * y_p * x_p)) \right) \\ &= a_p^0 - a_p^0 y_p (x_p \phi_0) + 1/2 * (a_p^0)^2 * |x_p|^2 \end{aligned}$$

Taking into account that x_p is a support vector, we have

$$W(a^0) - W(a^p) = 1/2 * (a_p^0)^2 * |x_p|^2$$

Suppose the optimal hyperplane passing through the origin recognizes the vector x_p incorrectly. This means that the inequality:

$$y_p(x_p * \phi_0) \leq 0$$

is valid. This is possible only if the vector x_p is an essential support vector. Now let us make one step in maximization the function $W(a)$ by fix $a_i, i \neq p$ and change only one parameter $a_p > 0$. We obtain:

$$W(a) = W(a^p) + a_p(1 - y_p(x_p \phi_p)) - 1/2 * (a_p^0)^2 * |x_p|^2$$

From the equality we obtain the best value of a_p :

$$a_p = (1 - y_p(x_p \phi_p)) \div |x_p|^2$$

The increment of the $W(a)$ at this moment equals

$$\Delta W_p = 1/2 * (1 - y_p(x_p \phi_p))^2 \div |x_p|^2$$

Since ΔW_p does not exceed the increment of the function $W(a)$ for the complete maximization, we obtain:

$$W(a^0) - W(a^p) \geq \Delta W_p = 1/2 * (1 - y_p(x_p \phi_p))^2 \div |x_p|^2$$

Combining the equations above, finally we get:

$$1/2 * (a_p^0)^2 * |x|^2 \geq 1/2 * (1 - y_p(x_p \phi_p))^2 \div |x_p|^2$$

From this equation and use the equation above, we could obtain:

$$a_p^0 \geq (1 - y_p(x_p \phi_p)) / |x_p|^2 \geq 1 / |x_p|^2$$

Taking into account that $|x_p| \leq D_{l+1}$, we obtain

$$a_p^0 \geq 1 / D_{l+1}^2$$

Thus if the optimal hyperplane makes the error classifying vector x_p in the leave one out procedure, then the inequality holds. Therefore

$$\sum_{i=1}^a a_i^0 \geq L_{l+1} / D_{l+1}^2$$

Where $L((x_1, y_1), \dots, (x_{l+1}, y_{l+1}))$ is the number of errors in the leave one out procedure on the sample $(x_1, y_1), \dots, (x_{l+1}, y_{l+1})$.

Now let us recall the properties of the optimal hyperplane

$$(\phi_0 * \phi_0) = \sum_{i=1}^a a_i^0$$

and

$$(\phi_0 * \phi_0) = 1 / p_{l+1}^2$$

Combing the above equations, we conclude that the inequality:

$$L_{l+1} \leq D_{l+1}^2 / p_{l+1}^2$$

is true with probability 1. Finally we get:

$$ER(a_l) = EL_{l+1} / (l + 1) \leq E \frac{D_{l+1}^2 / p_{l+1}^2}{l + 1}$$

20 Logistic Regression a.k.a. Maximum Entropy

$$P(y|x) = \frac{1}{Z_x} e^{\sum_i \lambda_i f_i(x,y)}$$

$$Z_x = \sum_y e^{\sum_i \lambda_i f_i(x,y)}$$

e.g., In NLP, we may use

$$f_{100}(x, y) = \begin{cases} 1 & \text{if } x=\text{word that ends with 'tion'} \text{ and } y=\text{Noun} , \\ 0 & \text{if otherwise .} \end{cases}$$

The above is more general binary classification where we are only deciding whether an example belongs to a class or NOT. Here, we can have features contributing to multiple classes according to their weights. For binary classification, the decision boundary is linear, as with perceptron or SVM. A major difference from SVMs is that, during training, every example contributes to the objective function, whereas in SVMs only the examples close to the decision boundary matter.

If we plot this function, we get a sigmoid-like graph. We can draw analogy between maximum entropy and neural network, and consider features as the input nodes in the neural network.

If we take the log of equation (33), we get a linear equation

$$\log P = \sum_i \lambda_i f_i + c$$

What should λ be?

$$\begin{aligned} \max_{\lambda} \log \left(\prod_{n=1}^N P(y_n|x_n) \right)^{\frac{1}{N}} &= \max_{\lambda} \sum_n \frac{1}{N} \log \left(\frac{1}{Z_x} e^{\sum_i \lambda_i f_i} \right) \\ &= \max_{\lambda} \frac{1}{N} \sum_n \left(\sum_i \lambda_i f_i - \log Z_x \right) \end{aligned}$$

In the above equation, maximizing λ_i is easy, but maximizing Z_x is not. To maximize for λ_i , we turn it into a concave form and find the point where the derivative w.r.t. λ is zero (hill climbing)

$$L = \frac{1}{N} \sum_n \sum_i \lambda_i f_i - \log \sum_y e^{\sum_i \lambda_i f_i(x_n,y)} \quad (33)$$

$$\frac{\partial L}{\partial \lambda_j} = \frac{1}{N} \sum_n f_j - \frac{1}{Z_x} \frac{\partial}{\partial \lambda_j} \left(\sum_y e^{\sum_i \lambda_i f_i} \right) \quad (34)$$

$$= \frac{1}{N} \sum_n \left[f_j - \frac{1}{Z_x} \sum_y f_j e^{\sum_i \lambda_i f_i(y,x_n)} \right] \quad (35)$$

$$= \frac{1}{N} \sum_n f_j - \sum_y f_j P(y|x_n) \quad (36)$$

$$= \frac{1}{N} \sum_n f_j(x_n, y_n) - \sum_y f_j P(y|x_n) \quad (37)$$

We can morph eq. 37 into expectation form by defining joint probability as follows:

$$P(y, x) = P(y|x) \tilde{P}(x) \quad (38)$$

$$\tilde{P}(x) = \frac{1}{N} \sum_n I(x_n = x) \quad (39)$$

$$\tilde{P}(y|x) = \frac{c(x_n = x, y_n = y)}{c(x_n = x)} \quad (40)$$

Rewriting eq. 37 in expectation form, we get:

$$\frac{\partial L}{\partial \lambda_j} = E_{\tilde{P}}[f_j] - E_P[f_j] \quad (41)$$

where the first term (before the minus) is a constant, and the complexity of calculating the second term depends on the number of classes in the problem. Now we have:

$$\lambda \leftarrow \lambda + \eta \frac{\partial L}{\partial \lambda} \quad (42)$$

We will justify why we chose log linear form instead of something else. Assume we want to find the maximum entropy subject to constraints on the feature expectations:

$$\max_{P(y|x)} H(y|x) \quad (43)$$

$$\min_{P(y|x)} -H(y|x) \quad (44)$$

$$\text{s.t. } E_{\tilde{P}}[f_i] = E_P[f_i] \quad \forall i \quad (45)$$

$$\sum_y P(y|x) = 1 \quad \forall x \quad (46)$$

To put this into words, we want to build a model such that for each feature, our model should match the training data. We have

$$H(y|x) = \sum_{x,y} \tilde{P}(x) P(y|x) \log \frac{1}{P(y|x)} \quad (47)$$

Find the maximum entropy of the above equation as follows

$$L(P, \lambda, \mu) = f_0 + \sum_j \lambda_j f_j \quad (48)$$

$$= \sum_{x,y} \tilde{P}(x) P(y|x) \log P(y|x) \quad (49)$$

$$+ \sum_i \lambda_i \left(\sum_{x,y} \tilde{P}(x) \tilde{P}(y|x) f_i - \tilde{P}(x) P(y|x) f_i \right) \quad (50)$$

$$+ \sum_x \tilde{P}(x) \mu_x \left(\sum_y P(y|x) - 1 \right) \quad (51)$$

$$\frac{\partial L}{\partial P(y|x)} = \tilde{P}(x) (\log P(y|x) + 1) - \sum_i \lambda_i \tilde{P}(x) f_i + \tilde{P}(x) \mu_x = 0 \quad (52)$$

$$\log P(y|x) = -1 + \sum_i \lambda_i f_i - \mu_x \quad (53)$$

$$P(y|x) = e^{-1-\mu_x} e^{\sum_i \lambda_i f_i} \quad (54)$$

$$= \frac{1}{Z_x} e^{\sum_i \lambda_i f_i} \quad (55)$$

The above result shows that maximum entropy has log-linear form. If we solve the dual of the problem

$$g(\lambda, \mu) = E_P \left[\sum_i \lambda_i f_i - 1 - \mu_x \right] + E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] - E_P \left[\sum_i \lambda_i f_i \right] + E_P [\mu_x] - \sum_x \tilde{P}(x) \mu_x \quad (56)$$

$$= E_P [-1 - \mu_x] + E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] + E_P [\mu_x] - \sum_x \tilde{P}(x) \mu_x \quad (57)$$

$$= E_P [-1] + E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] - \sum_x \tilde{P}(x) \mu_x \quad (58)$$

$$= - \sum_{x,y} \tilde{P}(x) e^{-1-\mu_x} e^{\sum_i \lambda_i f_i} + E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] - \sum_x \tilde{P}(x) \mu_x \quad (59)$$

$$(60)$$

Solving analytically for μ_x that maximizes g :

$$0 = \frac{\partial g}{\partial \mu_x} = \tilde{P}(x) \left(\sum_y -e^{-1-\mu_x + \sum_i \lambda_i f_i} - 1 \right) \quad (61)$$

$$\mu_x = \log \sum_y e^{\sum_i \lambda_i f_i} - 1 \quad (62)$$

Substituting μ_x into g :

$$g(\lambda, \mu) = - \sum_{x,y} \tilde{P}(x) \frac{1}{\sum_y e^{\sum_i \lambda_i f_i}} e^{\sum_i \lambda_i f_i} + E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] - \sum_x \tilde{P}(x) \left(\log \sum_y e^{\sum_i \lambda_i f_i} - 1 \right) \quad (63)$$

$$= -1 + E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] - \sum_x \tilde{P}(x) \log \left(\sum_y e^{\sum_i \lambda_i f_i} \right) + 1 \quad (64)$$

$$= E_{\tilde{P}} \left[\sum_i \lambda_i f_i \right] - \sum_x \tilde{P}(x) \log \left(\sum_y e^{\sum_i \lambda_i f_i} \right) \quad (65)$$

$$= E_{\tilde{P}} \left[\sum_i \lambda_i f_i - \log \left(\sum_y e^{\sum_i \lambda_i f_i} \right) \right] \quad (66)$$

$$= E_{\tilde{P}} \left[\sum_i \lambda_i f_i - \log Z_x \right] \quad (67)$$

$$= E_{\tilde{P}} [\log P(y|x)] \quad (68)$$

$$= L \quad (69)$$

Thus solving the dual of the entropy maximization problem consists of maximizing the likelihood of the training data with a log-linear functional form for $P(y|x)$.

21 Hidden Markov Models

A Hidden Markov Model (HMM) is a Markov Chain (a series of states with probabilities of transitioning from one state to another) where the states are hidden (latent) and each state has an emission as a random variable. The model is described as follows:

- Ω : the set of states, with $y_i \in \Omega$ denoting a particular state
- Σ : the set of possible emissions with $x_i \in \Sigma$ denoting a particular emission
- $P \in \mathbb{R}_{[0,1]}^{\Omega \times \Omega}$: the matrix with each element giving the probability of a transition
- $Q \in \mathbb{R}_{[0,1]}^{\Omega \times \Sigma}$: the matrix with each element giving the probability of an emission
- Π : the matrix with each element giving the probability of starting in each state

The probability distribution of an HMM can be decomposed as follows:

$$P(x_1, \dots, x_n, y_1, \dots, y_n) = \Pi(y_1) \prod_{i=1}^{n-1} P(y_i, y_{i+1}) \prod_{i=1}^n Q(y_i, x_i)$$

An example HMM is given:

$$\begin{aligned}\Omega &= \{1, 2\} \\ \Sigma &= \{a, b, c\} \\ P &= \begin{pmatrix} \frac{1}{3} & \frac{2}{3} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \\ Q &= \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ \frac{1}{3} & \frac{1}{2} & \frac{1}{6} \end{pmatrix}\end{aligned}$$

One possible sequence of observations would be:

$$\begin{array}{cccccccc}1 & 2 & 2 & 1 & 1 & 2 & 1 & 1 & 2 \\ \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow \\ a & b & c & a & a & a & a & a & b\end{array}$$

We can consider multiple problems relating to HMMs.

1. Decoding I: Given $x_1, \dots, x_n, P, Q, \Pi$, determine the sequence $y_1 \dots y_n$ that maximizes $P(Y_1, \dots, Y_n | X_1, \dots, X_n)$.
2. Decoding II: Given x_1, \dots, x_n and t , determine the distribution of y_K , that is, for all values a of y_t , $P(y_t = a | X_1, \dots, X_n)$.
3. Evaluation: Given x_1, \dots, x_n , determine $P(X_1, \dots, X_n)$.
4. Learning: Given a sequence of observations, $x_1^{(1)}, \dots, x_n^{(1)}, \dots, x_1^{(k)}, \dots, x_n^{(k)}$, learn P, Q, Π that maximize the likelihood of the observed data.

We define two functions, α and β .

$$\alpha^t(a) := P(X_1 = \hat{x}_1, \dots, X_t = \hat{x}_t, Y_t = a)$$

$$\beta^t(a) := P(X_{t+1} = \hat{x}_{t+1}, \dots, X_n = \hat{x}_n | Y_t = a)$$

which are also recursively defined as follows:

$$\alpha^{t+1}(a) = \sum_{c \in \Omega} \alpha^t(c) P(c, a) Q(a, \hat{x}_{t+1})$$

$$\beta^{t-1}(a) = \sum_{c \in \Omega} Q(c, \hat{x}_t) \beta^t(c) P(a, c)$$

We return to the Decoding II problem. Given x_1, \dots, x_n and t , determine the distribution of Y_K , that is, for all values a of Y_t , $P(Y_t = a | X_1, \dots, X_n)$. To do this, we rewrite the equation as follows:

$$P(y_t = a | X_1, \dots, X_n) = \frac{P(X_1, \dots, X_n, Y_t = a)}{P(X_1, \dots, X_n)}.$$

However, we need to calculate $P(X_1, \dots, X_n)$. We can do this using either α or β .

$$\begin{aligned} P(X_1, \dots, X_n) &= \sum_{a \in \Omega} \alpha^n(a) \\ &= \sum_{a \in \Omega} \beta^1(a) \Pi(a) Q(a, \hat{x}_1) \end{aligned}$$

The Decoding I problem can be solved with Dynamic Programming. (Given $x_1, \dots, x_n, P, Q, \Pi$, determine the sequence $y_1 \dots y_n$ that maximizes $P(Y_1, \dots, Y_n | X_1, \dots, X_n)$.) We can fill in a table with the following values:

$$T[t, a] = \max_{y_1 \dots y_t, y_t = a} P(y_1, \dots, y_t | X_1, \dots, X_t)$$

which means that each value is the probability of the most likely sequence at time t with the last emission being a . This can be computed using earlier values with the following formula:

$$T[t+1, a] = \max_{c \in \Omega} T[t, c] P(c, a) Q(a, \hat{x}_{t+1})$$

To compute the most likely sequence, we simply solve

$$\max_{a \in \Omega} T[n, a]$$

The learning problem can be solved using EM. Given the number of internal states, and x_1, \dots, x_n , we want to figure out P, Q , and Π . In the E step, we want to compute an expectation over hidden variables:

$$L(\theta, q) = \sum_y q(y|x) \log \frac{P(X, Y|\theta)}{q(Y|X)}$$

For HMM's, the number of possible hidden state sequences is exponential, so we use dynamic programming to compute expected counts of individual transitions and emissions:

$$P(a, b) \propto \sum_{i=1}^{n-1} q(Y_i = a, Y_{i+1} = b | X_1 \dots X_n) \quad (70)$$

$$Q(a, b) \propto \sum_{i=1}^n q(Y_i = a | X_1 \dots X_n) I(X = w) \quad (71)$$

The new P is defined as:

$$P^{\text{new}}(a, b) \propto \sum_{i=1}^{n-1} \alpha^i(a) P^{\text{old}}(a, b) \beta^{i+1} Q^{\text{old}}(b, \hat{x}_{i+1})$$

22 LBFGS

To train maximum entropy (logistic regression) models, we maximized the probability of the training data over possible feature weights λ :

$$\max_{\lambda} \prod_n P(Y_n | X_n)$$

It is to maximize

$$L = \log \prod_n \frac{1}{Z_{X_n}} e^{\sum_i \lambda_i f_i}$$

Of course we can solve it by using gradient ascend, but we today will talk about using an approximation of Newton's method.

22.1 Preliminary

For quadratic objective function

$$f(x) = \frac{1}{2} x^\top A x + b^\top x + c$$

Newton's iteration is given by

$$x_{k+1} = x_k + (\nabla^2 f(x_k))^{-1} \nabla f(x_k)$$

However, the exact version of Newton's method involves a few problems as follows:

- We need to compute Hessian $\nabla^2 f$, which is expensive.
- We also need to invert Hessian, which is also a costly operation.
- Furthermore, we need to store Hessian, which is expensive in terms of space.

So in order to compute Newton's iteration in a fast but relatively accurate way, an approximation should be developed. L-BFGS is one of them.

22.2 The BFGS Algorithm

Let B_k denote our approximation of the Hessian $\nabla^2 f(x_k)$ and then we can write Newton's iteration as

$$x_{k+1} = x_k - \alpha B_k^{-1} \nabla f(x_k)$$

Because the Hessian can be seen as the second order derivative of f , we wish to choose B_k such that:

$$B_k(x_{k+1} - x_k) = \nabla f(x_{k+1}) - \nabla f(x_k)$$

Let

$$\begin{aligned} s_k &= x_{k+1} - x_k \\ y_k &= \nabla f(x_{k+1}) - \nabla f(x_k) \end{aligned}$$

then we have

$$B_k s_k = y_k$$

This is to say that, our approximation is a solution of above equation. Consider

$$B_k = \frac{y_k y_k^\top}{s_k^\top y_k}$$

Because

$$B_k s_k = \frac{y_k y_k^\top s_k}{s_k^\top y_k} = \frac{y_k (y_k^\top s_k)}{s_k^\top y_k} = y_k$$

Further, let H_k be our approximation of $(\nabla^2 f(x_k))^{-1}$, the inverse of Hessian. We will have:

$$s_k = H_k y_k$$

H_k is a solution of above equation. One such H_k is given by

$$H_k = \frac{s_k s_k^\top}{s_k^\top y_k}$$

So, we want a direct formula of computing a symmetric H_{k+1} from H_k . That is, we want to fill in the ? term in following equation

$$H_{k+1} = H_k + \frac{s_k s_k^\top}{s_k^\top y_k} + ?$$

$$\text{such that } s_k = H_{k+1} y_k$$

With careful proofs and calculation, we get

$$H_{k+1} = (I - \rho_k s_k y_k^\top) H_k (I - \rho_k y_k s_k^\top) + \rho_k s_k s_k^\top \quad (72)$$

where $\rho_k = \frac{1}{s_k^\top y_k}$.

22.3 Proof of the method

The problem can be formalized as minimizing $\|H_{k+1} - H_k\|$ such that $H_{k+1} y_k = s_k$ and $H_{k+1}^\top = H_{k+1}$. Here $\|\cdot\|$ is defined as follows:

$$\|A\|^2 = \|W^{\frac{1}{2}} A W^{\frac{1}{2}}\|_f^2$$

where $\|A\|_f$ is defined as the square sum of all entries, $\sum_{i,j} a_{ij}^2$ and W is any matrix such that $W s_k = y_k$.

It is easy to verify that $H_{k+1} y_k = s_k$ and H_{k+1} is symmetric, as follows:

$$\begin{aligned} H_{k+1} y_k &= (I - \rho_k s_k y_k^\top) H_k (I - \rho_k y_k s_k^\top) y_k + \rho_k s_k s_k^\top y_k \\ &= \dots H_k (y_k - \rho_k s_k^\top y_k) + \rho_k s_k s_k^\top y_k \\ &= \dots H_k (y_k - y_k) + s_k \\ &= s_k \end{aligned}$$

22.4 L-BFGS Algorithm

L-BFGS algorithm tries to approximate $H_{k+1} \nabla f(x_{k+1})$ together. From 72, we can unroll the last m H_k 's. Then we will compute H_{k+1} directly from H_{k-m} .

$$\begin{aligned} H_{k+1} &= V_k^\top \dots V_{k-m}^\top H_{k-m} V_{k-m} \dots V_k \\ &\quad + \rho_{k-m} V_k^\top \dots V_{k-m}^\top s_{k-m} s_{k-m}^\top V_{k-m} \dots V_k \\ &\quad + \dots \\ &\quad + \rho_k s_k s_k^\top \end{aligned}$$

In the equation above, $V_k = I - \rho_k s_k y_k^\top$. The algorithm is

In the algorithm, $V_k = I - \rho_k y_k s_k^\top$. This algorithm needs to keep track of s_k and y_k in the last m steps and each step requires $2n$ space (n for s_k and n for y_k). So a total of $O(2mn)$ space is needed.

There are many user libraries that have already implemented this algorithm, so we can just use them for our computing.

Algorithm 2: L-BFGS

```
Require:  $H_{k-m}, s_i, y_i$   
 $q \leftarrow \nabla f_k$   
for  $i = k - 1..k - m$  do  
   $\alpha_i \leftarrow \rho_i s_i^\top q$   
   $q \leftarrow q - \alpha_i y_i$   
end for  
 $r \leftarrow H_{k-m} q$   
for  $i = k - m..k - 1$  do  
   $\beta \leftarrow \rho_i y_i^\top r$   
   $r \leftarrow r + s_i(\alpha_i - \beta)$   
end for  
return  $r$ 
```

23 Reinforcement Learning

23.1 Markov Decision Processes

A Markov Decision Process is an extension of the standard (unhidden) Markov model [1]. Each state has a collection of actions that can be performed in that particular state. These actions serve to move the system into a new state. More formally, the MDP's state transitions can be described by the transition function $T(s, a, s')$, where a is an action moving performable during the current state s , and s' is some new state.

As the name implies, all MDPs obey the *Markov property*, which holds that the probability of finding the system in a given state is dependent only on the previous state. Thus, the system's state at any given time is determined solely by the transition function and the action taken during the previous timestep:

$$P(S_t = s' | S_{t-1} = s, a_t = a) = T(s, a, s')$$

Each MDP also has a reward function $R : S \mapsto \mathbb{R}$. This reward function assigns some value $R(s)$ to being in the state $s \in S$. The goal of a Markov Decision Process is to move from the current state s to some final state in a way that a) maximizes $R(s)$ and b) maximizes R 's potential value in the future.

Given a Markov Decision Process we wish to find a *policy* – a mapping from states to actions. The policy function $\Pi : S \mapsto A$ selects the appropriate action $a \in A$ given the current state $s \in S$.

23.2 Value Iteration

The consequences of actions (i.e., rewards) and the effects of policies are not always known immediately. As such, we need some mechanisms to control and adjust policy when the rewards of the current state space are uncertain. These mechanisms are collectively referred to as *reinforcement learning* [1].

One common way of trading off present reward against future reward is by introducing a *discount rate* γ . The discount rate is between 0 and 1, and we can use it to construct a weighted estimate of future rewards:

$$\sum_{t=0}^{\infty} \gamma^t r_t$$

Here, we assume that $t = 0$ is the current time. Since $0 < \gamma < 1$, greater values of t (indicating rewards farther in the future) are given smaller weight than rewards in the nearer future.

Let $V^\Pi(s)$ be the value function for the policy Π . This function $V^\Pi : S \mapsto \mathbb{R}$ maps the application of Π to some state $s \in S$ to some reward value. Assuming the system starts in state s_0 , we would expect the system to have the value

$$V^\Pi(s) = E \left[\sum_{t=0}^{\infty} \gamma^t R(s_t) \mid s_0 = s, \Pi \right]$$

Since the probability of the system being in a given state $s' \in S$ is determined by the transition function $T(s, a, s')$, we can rewrite the formula above for some arbitrary state $s \in S$ as

$$V^\Pi(s) = R(s) + \sum_{s'} T(s, a, s') \gamma V^\Pi(s')$$

where $a = \Pi(s)$ is the action selected by the policy for the given state s .

Our goal here is to determine the optimal policy $\Pi^*(s)$. Examining the formula above, we see that $R(s)$ is unaffected by choice of policy. This makes sense because at any given state s , local reward term $R(s)$ is determined simply by virtue of the fact that the system is in state s . Thus, if we wish to find the maximum policy value function (and therefore find the optimum policy) we must find the action a that maximizes the summation term above:

$$V^{\Pi^*}(s) = R(s) + \max_a \sum_{s'} T(s, a, s') \gamma V^{\Pi^*}(s')$$

Note that this formulation assumes that the number of states is finite.

The formula above forms the basis of the *value iteration* algorithm. This operation starts with some initial policy value function guess and iteratively refines $V(s)$ until some acceptable convergence is reached:

Algorithm 3: Value Iteration

1. Initialize $V(s)$.
 2. Repeat until converged:
 - (a) for all $s \in S$:
 - i. $R(s) = R(s) + \max_a \sum_{s'} T(s, a, s') \gamma V(s')$
 - ii. $\Pi(s) = \operatorname{argmax}_a Q(s, a)$
-

Each pass of the value iteration maximizes $V(s)$ and assigns to $\Pi^*(s)$ the action a that maximizes $V(s)$. The function $Q(s, a)$ represents the potential value for $V(s)$ produced by the action $a \in A$.

23.3 Q-Learning

The example of value iteration above presumes that the transition function $T(s, a, s')$ is known. If the transition function is not known, then the function $Q(s, a)$ can be obtained through a similar process of iterative learning, the aptly-named *Q-learning*.

The naive guess for a Q-learning formula would be one that closely resembles the policy value function, such as

$$Q(s, a) = R(s) + \gamma \left[\max_{a'} Q_{old}(s, a') \right]$$

This formula aggressively replaces old values of Q , though, which is not always desirable. For better results [1], use a weighted average learning function:

$$Q(s, a) = \eta Q_{old}(s) + (1 - \eta) \gamma \left[\max_{a'} Q_{old}(s, a') \right]$$

where η is a user-selected learning parameter.

This formula can be turned into an algorithm much like the value iteration algorithm above. The only difference is that the action is selected by a user-defined function $f(s)$, which returns the appropriate policy action $\Pi(s)$ most of the time, but occasionally selects a random action to blunt the effects of sampling bias.

The Q-learning algorithm below is from Ballard's textbook [1]:

Algorithm 4: One-Step Q-Learning

1. Initialize $\Pi(s)$ to $\operatorname{argmax}_a Q(s, a)$.
 2. Repeat until Π converges:
 - (a) For each $s \in S$:
 - i. Select an action $a = f(s)$.
 - ii. $Q(s, a) = \eta Q_{old}(s) + (1 - \eta) \gamma \left[\max_{a'} Q_{old}(s, a') \right]$
 - iii. $\Pi(s) = \operatorname{argmax}_a Q(s, a)$.
-

23.4 Temporal Difference Learning

One disadvantage of value iteration is that it can take a long time for updates to later states to propagate back to earlier states. For instance, an MDP attempting to navigate a maze would see its reward function jump once it reaches the final stage N , but it would take N iterations for the effects of that jump propagate back to stage 1.

Temporal difference learning remedies this by modifying the output weights for each state with estimates of the policy value for the next state:

$$\Delta w = \eta (R(s_t) + \gamma V(s_{t+1}) - V(s_t)) \sum_{k=1}^t \lambda^{t-k} \frac{\partial V_k}{\partial w}$$

Where $0 < \lambda < 1$ and $V(s)$ is presumed to be a function of the weight w . The temporal difference learning algorithm below is adapted from [1]:

Algorithm 5: Temporal Difference Learning

1. Initialize all network weights to random values.
 2. Repeat until Π converges:
 - (a) Select current state s_t .
 - (b) Use $\Pi(s_t)$ to obtain a new state s_{t+1} and calculate $V(s_{t+1})$.
 - (c) If an actual reward is available, substitute that for estimate $V(s_{t+1})$.
 - (d) $w = w + \eta (R(s_t) + \gamma V(s_{t+1}) - V(s_t)) \sum_{k=1}^t \lambda^{t-k} \frac{\partial V_k}{\partial w}$.
-

24 Game Theory

Game theory is not like the reinforcement learning which we have just talked about. In the game theory, the condition that one is playing against someone else is considered. Furthermore, game theory is not about learning, it does not learn from data or previous instances, but tries to figure out what to do given a specific game rule.

A game could be described using a game matrix. The two players are called row player and column player. The value of each element in the matrix denotes the reward for the player. Such a matrix model could be applied to many situations, e.g., board games, bets. A good example would be the rock-paper-scissors game, which has a matrix M like

| | Rock | Paper | Scissors |
|----------|------|-------|----------|
| Rock | 0 | -1 | 1 |
| Paper | 1 | 0 | -1 |
| Scissors | -1 | 1 | 0 |

Note that here, the game has the same set of choices for both players, which is not necessary in general. Moreover, the above matrix satisfies that

$$M = -M^T,$$

which makes the game symmetric. Also, we could find that this is a zero-sum game, where one player gets his/her reward from the other, with the total amount never changed.

24.1 Definition of the Problem

Define \mathbf{p} as the row player's strategy, and \mathbf{q} to be that of the column player. \mathbf{p} and \mathbf{q} are both probability distributions. Then, the expected reward with respect to the strategies is

$$E(\mathbf{p}, \mathbf{q}) = \sum_i p_i \sum_j q_j m_{ij}.$$

If we consider the problem from the perspective of the row player (which will be the same if we choose the column player), our goal is to maximize the reward E of the row player, i.e., to find a strategy \mathbf{p} that maximizes the row value v_r and

$$v_r = \max_{\mathbf{p}} \min_{\mathbf{q}} E(\mathbf{p}, \mathbf{q}),$$

where the minimization means that our rival (the column player) does whatever that is worst for us, and the maximization chooses the strategy that is best for us given the strategy of the column player. Similarly, the column value v_c is

$$v_c = \min_{\mathbf{q}} \max_{\mathbf{p}} E(\mathbf{p}, \mathbf{q}),$$

which reflects the fact that the row player does whatever that could maximize his/her reward, and that the column player chooses a best strategy to minimize it.

24.2 Some Simplification

Before solving the optimization problem, we should first do some simplification by showing that, when assuming the column player's strategy, we do not have to consider all the variations of \mathbf{q} .

If we choose one possibility of \mathbf{q} with probability of 1, i.e., the column player only does one action, say, action j , from the available action set, we have

$$\min_{\mathbf{q}} E(\mathbf{p}, \mathbf{q}) \leq \min_j E(\mathbf{p}, j), \quad (73)$$

where $E(\mathbf{p}, j) = \sum_i p_i m_{ij}$ from its definition.

Now, let us consider the strategy of the column player using the structure of v_c , i.e., the column player assumes a \mathbf{p} at first, and then choose an action $l = \operatorname{argmin}_j E(\mathbf{p}, j)$ to make a worst result for the row player. This leads to

$$\forall \mathbf{p}, \mathbf{q}, \quad E(\mathbf{p}, \mathbf{q}) = \sum_j q_j E(\mathbf{p}, j) \geq \sum_j q_j E(\mathbf{p}, l) = E(\mathbf{p}, l),$$

which could be rewritten as

$$\min_{\mathbf{q}} E(\mathbf{p}, \mathbf{q}) \geq \min_j E(\mathbf{p}, j). \quad (74)$$

Combining Equation 73 and Equation 74, we have

$$\min_{\mathbf{q}} E(\mathbf{p}, \mathbf{q}) = \min_j E(\mathbf{p}, j),$$

which means that instead of considering all possibilities of \mathbf{q} , we could consider only one move at a time. In the context of the Rock-Paper-Scissors game, this is to say that $\mathbf{q} = (1, 0, 0)$ or $(0, 1, 0)$ or $(0, 0, 1)$.

Using this equality, the row value v_r that we want to maximize is simplified as

$$\max_{\mathbf{p}} \min_{\mathbf{q}} E(\mathbf{p}, \mathbf{q}) \Rightarrow \max_{\mathbf{p}} \min_j E(\mathbf{p}, j) = \max_{\mathbf{p}} \min_j \sum_i p_i m_{ij}.$$

Similarly, v_c could be simplified as

$$\min_{\mathbf{q}} \max_i \sum_j q_j m_{ij}.$$

24.3 Nash Equilibrium

The objective functions v_r and v_c that have been simplified as above still contain two min/max operations, which makes them too complicated to solve. In order to get rid of one of the min/max, we could introduce a variable v and adding a new constraint for compensation. If we do the same for the case of v_c , we could come up with

$$\begin{array}{ll} \max_{\mathbf{p}, v} & v \\ \text{s.t.} & \sum_i p_i m_{ij} \geq v \quad \forall j \\ & \sum_i p_i = 1 \\ & p_i \geq 0 \quad \forall i \end{array} \qquad \begin{array}{ll} \min_{\mathbf{q}, v} & v \\ \text{s.t.} & \sum_j q_j m_{ij} \leq v \quad \forall i \\ & \sum_j q_j = 1 \\ & q_j \geq 0 \quad \forall j \end{array}$$

where the left one is for v_r and right one v_c . For the left one, we could simplify a little bit more by firstly adding onto the game matrix a constant such that all of its entries are positive, and define

$$y_i = p_i/v, \quad \forall i,$$

then the model is updated as the left one below

$$\begin{array}{ll} \min_{\mathbf{y}} & \sum_i y_i \\ \text{s.t.} & \sum_i y_i m_{ij} \geq 1 \quad \forall j \\ & y_i \geq 0 \quad \forall i \end{array} \qquad \begin{array}{ll} \max_{\mathbf{x}} & \sum_j x_j \\ \text{s.t.} & \sum_j x_j m_{ij} \leq 1 \quad \forall i \\ & x_j \geq 0 \quad \forall j \end{array}$$

The model on the right is obtained by introducing $\forall j, x_j = q_j/v$ for the v_c model. What we just came up with are actually special cases of the following general models

$$\begin{array}{ll} \min_{\mathbf{y}} & \mathbf{b}^T \mathbf{y} \\ \text{s.t.} & \mathbf{A}^T \mathbf{y} \geq \mathbf{c} \\ & \mathbf{y} \geq 0 \end{array} \qquad \begin{array}{ll} \max_{\mathbf{x}} & \mathbf{c}^T \mathbf{x} \\ \text{s.t.} & \mathbf{A} \mathbf{x} \leq \mathbf{b} \\ & \mathbf{x} \geq 0 \end{array}$$

which are a pair of duals in the linear programming (LP) theory, and thus have the same optimization result. The set of strategy (\mathbf{p}, \mathbf{q}) under such condition is called Nash Equilibrium, which could be mathematically described as

$$\begin{aligned}
& (\mathbf{p}, \mathbf{q}) \\
\text{s.t. } & E(\mathbf{p}, \mathbf{q}) = \max_{\mathbf{p}'} E(\mathbf{p}', \mathbf{q}) \\
& E(\mathbf{p}, \mathbf{q}) = \max_{\mathbf{q}'} E(\mathbf{p}, \mathbf{q}')
\end{aligned}$$

When the Nash Equilibrium has been reached, both players know the strategy of each other, and they all get the best reward using their current strategies against those of the rivals. Hence neither of them need to change the current strategy.

24.4 Proof of the LP Duality

The LP duality that we used to reach the Nash Equilibrium is proved as follows. We will start from the previous linear programming model on the left and will get the other one using the primal-dual method.

Firstly, the Lagrangian of the model is calculated as

$$\begin{aligned}
L(\mathbf{y}, \boldsymbol{\lambda}, \boldsymbol{\mu}) &= \sum_i y_i + \sum_j \lambda_j (1 - \sum_i y_i m_{ij}) - \sum_i \mu_i y_i \\
&= \sum_j \lambda_j + \sum_i y_i (1 - \sum_j \lambda_j m_{ij} - \mu_i)
\end{aligned}$$

Then by defining

$$g(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \min_{\mathbf{y}} L(\mathbf{y}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \begin{cases} -\infty, & \text{if } \exists i, 1 - \sum_j \lambda_j m_{ij} - \mu_i \neq 0 \\ \sum_j \lambda_j, & \text{o.w.} \end{cases}$$

we will have the dual problem as below

$$\begin{aligned}
\max \quad & g(\boldsymbol{\lambda}, \boldsymbol{\mu}) = \sum_j \lambda_j \\
\text{s.t. } \quad & 1 - \sum_j \lambda_j m_{ij} - \mu_i = 0, \quad \forall i \\
& \lambda_j \geq 0, \quad \forall j \\
& \mu_i \geq 0, \quad \forall i
\end{aligned}$$

Due to the constraint $\forall i, \mu_i \geq 0$, we could remove μ_i in the first constraint and update it as

$$1 - \sum_j \lambda_j m_{ij} \geq 0 \Rightarrow \sum_j \lambda_j m_{ij} \leq 1$$

If we replace $\boldsymbol{\lambda}$ with \mathbf{x} , this dual is exactly the same as the linear programming model for $v_{c'}$, which proves that the previous two LP models are dualities. Note that here, we did not do the proof for general cases, which is easy to get using the variables of \mathbf{b} and \mathbf{c} , which is set to 1 in our proof.

Recall that we also applied the primal-dual in the SVM, the difference is that in SVM, we go to the domain of the dual to make the problem easier to solve, while here, we generally use the duality to prove that the two model are duals, which gives us the Nash Equilibrium.