

# Summary Notes

## EE425X - Machine Learning: A Signal Processing Perspective

Namrata Vaswani  
Iowa State University

These notes are work in progress. For the latest version at any time, see the link [https://www.dropbox.com/s/kwle0pf09mp3631/Summary\\_Notes\\_2.pdf?dl=0](https://www.dropbox.com/s/kwle0pf09mp3631/Summary_Notes_2.pdf?dl=0)

### Acknowledgement

Some of these notes are based on material presented in this Stanford Machine Learning class <http://cs229.stanford.edu/syllabus.html>.

Rest of the notes are based on material I have learned over the years from various sources.

### Contents

<b>1</b>	<b>Data Simulation Basics</b>	<b>4</b>
1.1	Why simulate – why test code on simulated data first? . . . . .	4
1.2	Monte Carlo Estimation of Expected values . . . . .	5
1.3	Understanding a multivariate Gaussian distribution . . . . .	5
1.4	How to simulate . . . . .	5
1.4.1	Notes . . . . .	5
1.5	Estimation error computations: how to know your code is correct . . . . .	5
<b>2</b>	<b>Cross Validation for Supervised Learning: training and test datasets</b>	<b>6</b>
2.1	Simple Cross Validation . . . . .	6
2.2	Leave-one-out cross validation: any general task . . . . .	6
2.2.1	Leave-one-out cross validation for Linear regression . . . . .	7
<b>3</b>	<b>Gradient Descent (GD) and Stochastic GD</b>	<b>7</b>
3.1	Setting the step size $\mu$ . . . . .	8
3.2	Extra material: Stochastic GD (SGD) and Stochastic Mini-Batch GD ** . . . . .	8
3.2.1	Good videos on stochastic and mini-batch SGD . . . . .	8
<b>4</b>	<b>Supervised Learning, Unsupervised Learning, Model-based Learning: basics</b>	<b>8</b>
4.1	Examples . . . . .	9
<b>5</b>	<b>Supervised Learning: Linear Regression</b>	<b>9</b>
5.1	Model . . . . .	9
5.2	Including a nonzero mean in the model . . . . .	10
5.3	Learning $\theta$ : minimize squared loss . . . . .	10
5.3.1	Understanding squared loss: Maximum Likelihood Estimation under i.i.d. Gaussian model ** . . . . .	10
5.4	Solutions for minimizing squared loss: also called Least Squares (LS) Estimation . . . . .	10
5.5	Least Squares, Regularized LS, Recursive LS and more . . . . .	11

<b>6</b>	<b>Supervised Learning: Two class and Multi-class Logistic Regression: Summary</b>	<b>11</b>
6.1	Two class Logistic Regression . . . . .	11
6.2	Multi-class Logistic Regression . . . . .	12
6.3	One-vs-rest (OVR) or One-vs-all use of two-class logistic regression . . . . .	13
<b>7</b>	<b>Supervised Learning: Two-class Logistic Regression Details</b>	<b>13</b>
7.1	Two-class Model and its probabilistic interpretation . . . . .	14
7.1.1	Use of bias term in Logistic regression . . . . .	14
7.2	Learning $\theta$ : Maximum Likelihood Estimation . . . . .	14
<b>8</b>	<b>Supervised Learning: Generative Learning (a.k.a. Bayesian models/learning)</b>	<b>15</b>
8.1	Learning $\theta$ : Maximum Likelihood Estimation . . . . .	15
8.2	Generative Learning: Gaussian Discriminant Analysis (GDA) . . . . .	15
8.2.1	Gaussian Discriminant Analysis (GDA) with Naive Bayes assumption and equal co- variances . . . . .	16
8.2.2	Notes . . . . .	17
8.3	Generative Learning: Spam Filter – an example of Discrete-valued or Categorical Features . .	17
8.3.1	Simple Spam Filter: entries of $x$ are binary . . . . .	17
8.3.2	Naive Bayes assumption . . . . .	18
8.3.3	Simple Spam filter with Naive Bayes . . . . .	18
8.3.4	Spam Filter: entries of $z$ count the number of times a word occurs . . . . .	19
8.3.5	Spam filter: model 3 . . . . .	19
8.3.6	More ideas and related problems . . . . .	20
8.4	Modeling Categorical data ** . . . . .	20
8.5	Dealing with discrete-valued and real-valued data together in the Generative Learning (a.k.a. Bayesian) framework ** . . . . .	20
<b>9</b>	<b>Supervised Learning: Linear Classifiers and Support Vector Machines (SVMs)</b>	<b>21</b>
9.1	Linear Classifiers ** . . . . .	21
9.2	Support Vector Machines (SVMs): Motivation and main idea . . . . .	21
9.2.1	Notation change . . . . .	22
9.2.2	Using SVM for classification . . . . .	22
9.2.3	Goal . . . . .	22
9.3	Simplifications to obtain a convex optimization problem ** . . . . .	22
9.4	Final primal problem . . . . .	22
9.5	Simplification using duality: needed to develop Kernel SVMs . . . . .	23
9.6	Kernel SVM ** . . . . .	23
9.7	Soft margin SVMs . . . . .	24
<b>10</b>	<b>Deviation: Introduction to Langrange duality to undertand how to derive the dual program **</b>	<b>24</b>
<b>11</b>	<b>Least Squares (LS): all kinds</b>	<b>24</b>
11.1	Regularized LS . . . . .	24
11.2	Use of prior knowledge and Recursive LS . . . . .	24
<b>12</b>	<b>Unsupervised Learning: PCA</b>	<b>25</b>
12.1	What is PCA . . . . .	25
12.2	Why PCA . . . . .	25
12.3	Practical information: How to implement PCA . . . . .	26
12.3.1	Reduced dimension, $r$ , is specified . . . . .	26
12.3.2	Use of PCA . . . . .	26

12.3.3	Deciding $r$ . . . . .	27
12.4	Theory – computational - 1 . . . . .	28
12.5	Theory – computational 2 . . . . .	29
12.6	Statistical optimality . . . . .	29
12.6.1	Statistical Optimality . . . . .	29
12.7	Delete ** . . . . .	31
12.7.1	Optimality from a computational viewpoint . . . . .	31
12.8	Proofs of the above claims ** . . . . .	31
12.8.1	Equivalence between statistical and computational claims 1 . . . . .	31
12.8.2	Proof of computational claim 1: Frobenius norm optimality . . . . .	32
12.8.3	Equivalence between statistical and computational claims 2 . . . . .	32
12.8.4	Proof of computational claim 2: induced 2 norm optimality . . . . .	33
12.9	PCA de-correlates the data: what does it mean ** . . . . .	34
<b>13</b>	<b>Unsupervised Learning: Clustering</b> . . . . .	<b>34</b>
13.1	Problem . . . . .	34
13.2	k-means clustering . . . . .	34
13.3	Alternating Minimization (AltMin) algorithm . . . . .	35
13.4	Probabilistic Model (Generative Model) for Clustering: Gaussian Mixture Model (GMM) . . . . .	36
13.5	EM algorithm for MLE for Gaussian Mixture Model ** . . . . .	36
13.6	Final EM algorithm for Gaussian Mixture Models ** . . . . .	37
13.7	General EM ** . . . . .	38
<b>14</b>	<b>Supervised Learning: Deep Learning / Neural Networks: basic idea and training</b> . . . . .	<b>39</b>
14.1	Multi-layer Perceptron (MLP) or Feed-forward Network . . . . .	40
14.2	Training an MLP: Back-propagation . . . . .	40
14.3	Understanding CNN's at a top level with visuals . . . . .	41
14.4	Aside: understanding convolution . . . . .	41
14.5	Convolutional Neural Network (CNN or ConvNet): basic idea . . . . .	41
14.6	Different types of layer . . . . .	41
14.6.1	Fully connected (FC) layer: . . . . .	42
14.6.2	Convolutional (CN) layer: . . . . .	42
14.6.3	Pooling layer: . . . . .	42
14.7	Why CN + pooling instead of FC . . . . .	43
14.8	Recurrent Neural Network (RNN): basic idea . . . . .	43
14.9	Different NN architectures and when to use each . . . . .	43
<b>15</b>	<b>Reliability of an output **</b> . . . . .	<b>43</b>
<b>16</b>	<b>Bias-Variance Tradeoff</b> . . . . .	<b>43</b>
16.1	What is it? . . . . .	43
16.2	Bias-Variance Tradeoff for Linear Regression . . . . .	44
<b>17</b>	<b>Picking reduced dimension <math>r</math> to minimize test-MSE (optimize Bias-Variance Tradeoff)</b> . . . . .	<b>46</b>
17.1	Approximating Test-MSE in practice . . . . .	46
17.2	How to reduce test data MSE / generalization error a.k.a. Regularization a.k.a. Feature Selection . . . . .	46
17.3	Feature normalization . . . . .	47
17.4	Feature normalization vs PCA . . . . .	48

<b>18 Learning Theory</b>	<b>48</b>
18.1 Two probability results we use . . . . .	49
18.2 Misclassification (zero-one) loss . . . . .	49
18.3 Tradeoffs . . . . .	50
18.4 Connections to what we have learnt . . . . .	50
18.5 VC dimension . . . . .	51

## Readings for Spring 2023 EE 425 course

### Read Sections

- Sec. 4 Sec. 5 (except 5.3.1): Linear Regression
- Sec. 3 (except 3.2): Gradient Descent, setting its step size. Skip SGD.
- Sec. 1.2: Monte Carlo estimation
- Sec. 1.1: Simulating data to check if algorithm coded correctly, improve algorithm
- Sec. 1.5: Computing errors in general (always use normalized errors), computing errors for simulated data (separate training and test datasets). computing errors on real data which is limited (either split as training and testing or use cross validation)
- Sec. 2: Cross Validation
- Sec. 6.1: Logistic Regression
- Sec. 8 (parts of 8.1 to 8.3): Generative Models and Bayesian Learning: GDA and spam filter
- Sec. 9.2, 9.4, 9.5, 9.7: SVMs
- Sec. 12 : 12.1, 12.2, 12.3, 12.3.3: PCA (Unsupervised learning)
- Sec. 13: 13.2, 13.3: Clustering (Unsupervised learning)
- Sec. 16: 16.1, 16.2: Bias-Variance tradeoff
- Sec. 17: Picking reduced dimension  $r$  to minimize test-MSE (optimize Bias-Variance Tradeoff)

## Notation

In these notes  $'$  and  $T$  are both used to denote vector or matrix transpose. At a few other places too MATLAB notation is used.

## Math basics

Probability - See file 322-recap- and cs229-  
 Linear Algebra - See cs229-linalg

## 1 Data Simulation Basics

### 1.1 Why simulate – why test code on simulated data first?

Consider the house price prediction based on house features example. In the previous section, we said we use linear regression to model the data and predict the price. If I write code to learn  $\theta$  and apply it to a real dataset directly and suppose it does not "work too well" (my code does not give very good predictions on test data). How do I know if (i) the linear regression model is wrong or (ii) my learning approach (normal equations or GD) is wrong, or (iii) there is a code bug (I have an extra factor of 2 at some place by mistake)?

A partial fix to the above problem to above it is to first simulate data using the linear regression model. So when I test my learning code on this data I know what is true  $\theta$  I am looking for. Then I can try to fix (ii) and (iii) issues. Once these are fixed, then we try the same code on real data and then maybe compare with another model to check which one is better.

## 1.2 Monte Carlo Estimation of Expected values

Suppose the goal is to compute/approximate  $\mathbb{E}[z]$  (expected value or average). Suppose also we know the probability distribution of  $z$ ,  $p(z)$ , and we have a way to generate (simulate) independent realizations of  $z$ . Then we can approximate

$$\mathbb{E}[z] \approx \frac{1}{m} \sum_{i=1}^m z^{(i)}, \quad z^{(i)} \stackrel{\text{i.i.d.}}{\sim} p(z)$$

We use this to approximate prediction error in ML among other things.

## 1.3 Understanding a multivariate Gaussian distribution

See these two files <http://cs229.stanford.edu/section/gaussians.pdf>

<http://cs229.stanford.edu/notes2020fall/notes2020fall/cs229-prob.pdf>

## 1.4 How to simulate

Task: Generate your own data to simulate the linear regression model  $y = \boldsymbol{\theta}^T \mathbf{x} + e$ .

**Recall:**  $n = d + 1$ , thus  $\boldsymbol{\theta}$  is a vector of length  $n$ . The  $\mathbf{x}^{(i)}$ s are of length  $d$  with 1 appended as first entry.

As an example, suppose  $d = 5$ , then  $n = 6$ . Can set  $\boldsymbol{\theta} = [10, 1, -1, -3, 4, 2]^\top$ .

Generate  $m$  such independent training data vectors. Also generate another set of  $m_{test}$  independent data vectors for the testing step.

Let  $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$  denote the Gaussian distribution with mean  $\boldsymbol{\mu}$  and covariance  $\boldsymbol{\Sigma}$ .

Do this as follows:

(1) Fix  $\boldsymbol{\theta}$  once as an  $n$  length vector.

(2) Training data generate:

For  $i = 1$  to  $m$ , generate  $\mathbf{x}^{(i)} \sim \mathcal{N}(0, I_d)$ ,  $e^{(i)} \sim \mathcal{N}(0, \sigma_e^2)$  and  $y^{(i)} = \boldsymbol{\theta}^T [1, \mathbf{x}^{(i)}] + e^{(i)}$ .

(3) Test data generate:

In a different Loop, repeat (2) for  $i = 1$  to  $m_{test}$ . Use this data in the testing step and not for training.

No-for-loop version of (2):

$\mathbf{X} = [\text{ones}(m, 1), \text{randn}(m, d)]$ ,  $\mathbf{e} = \sigma_e * \text{randn}(m, 1)$ ,  $\mathbf{y} = \mathbf{X}\boldsymbol{\theta} + \mathbf{e}$ .

### 1.4.1 Notes

*For generating the data, I have suggested using the Gaussian distribution just as an example. But you do not have to use the Gaussian. You could also use any other distribution, e.g., the uniform distribution.*

On the other hand, for the error  $e^{(i)}$ , we have assumed that it is Gaussian in our model (that is why the squared loss is justified). So for generating error, you *do have to* use the Gaussian distribution.

## 1.5 Estimation error computations: how to know your code is correct

- Always compute normalized errors: if code works well error should be less than 1. When increase  $m$  (number of samples), while keeping everything else fixed, this error should reduce.

- If data were simulated, then one can compute (normalized) parameter estimation error  $\frac{\|\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\|_2^2}{\|\boldsymbol{\theta}\|_2^2}$

- In all cases, one can compute (normalized) prediction error / training loss:  $\frac{(1/m) \sum_{i=1}^m (\mathbf{y}^{(i)} - \hat{\mathbf{y}}^{(i)})^2}{(1/m) \sum_{i=1}^m (\mathbf{y}^{(i)})^2}$

- Normalized Test MSE / Generalization Error computation:

- Simulated data: Generate an independent set of test data
- Real data: Two options

- \* Hold-out Cross Validation: Split available data into disjoint training and test sets, use most data for training, rest for testing.
- \* Or, use Leave-One-Out Cross-Validation, see Sec. 2
- For each test data point, compute  $\hat{\mathbf{y}}_{test} = h_{\hat{\theta}}(\mathbf{x}_{test})$
- Compute normalized average squared error  $\frac{\sum_{i=1}^{m_{test}} (\mathbf{y}_{test}^{(i)} - \hat{\mathbf{y}}_{test}^{(i)})^2}{\sum_{i=1}^{m_{test}} (\mathbf{y}_{test}^{(i)})^2}$
- For different problems, squared error can be replaced by other error measures.

## 2 Cross Validation for Supervised Learning: training and test datasets

In case of supervised learning, we need to define training and test datasets.

### 2.1 Simple Cross Validation

Simplest way to define training and test datasets is to use a certain fraction of the data for training and the rest for testing. A common rule of thumb is 80-20 or 70-30. In this case training and test data are disjoint which makes it easier to impose statistical independence assumptions. Also this is easiest to code in and the code runs fastest.

The simplest type of cross-validation is leave-one-out-cross-validation.

### 2.2 Leave-one-out cross validation: any general task

Use of distinct subsets is easiest to code in but wastes a lot of data. The other extreme solution is to use leave-one-out cross-validation. This means we loop over all the data multiple times. In the  $i$ -th loop, we use all but the  $i$ -th data point for training and use the  $i$ -th one for testing. Compute error on this  $i$ -th one and store it. At the end of this loop, average the stored errors from each loop to compute the average generalization error / predictor error.

- for  $i = 1$  to  $m_{tot}$ 
  - at iteration  $i$ , define the training data matrix  $\mathbf{X}$  and vector  $\mathbf{y}$  as follows: use *all* data except the  $i$ -th

$$\mathbf{X} = \begin{bmatrix} -(\mathbf{x}^{(1)})^T - \\ -(\mathbf{x}^{(2)})^T - \\ \vdots \\ -(\mathbf{x}^{(i-1)})^T - \\ -(\mathbf{x}^{(i+1)})^T - \\ \vdots \\ -(\mathbf{x}^{(m)})^T - \end{bmatrix}$$

$$\mathbf{y} = \begin{bmatrix} \mathbf{y}^{(1)} \\ \mathbf{y}^{(2)} \\ \vdots \\ \mathbf{y}^{(i-1)} \\ \mathbf{y}^{(i+1)} \\ \vdots \\ \mathbf{y}^{(m)} \end{bmatrix}$$

- Learn the model parameters use  $\mathbf{X}, \mathbf{y}$  as the training data
- Use the  $i$ -th data points for computing the error  $err(i)$ .

end for

Compute

$$TestMSE = \frac{1}{m_{tot}} \sum_{i=1}^{m_{tot}} err(i), NormalizedTestMSE = \text{as needed for the problem}$$

### 2.2.1 Leave-one-out cross validation for Linear regression

Given  $m$  training data points  $\mathbf{x}^{(i)}, \mathbf{y}^{(i)}, i = 1, 2, \dots, m$ . Suppose the goal is to evaluate the validity of a linear regression model on this data. We will compute Normalized-Test-MSE as follows.

- for  $i = 1$  to  $m_{tot}$ 
  - at iteration  $i$ , define the training data matrix  $\mathbf{X}$  and vector  $\mathbf{y}$  as follows: use *all* data except the  $i$ -th. This is defined above.

$$\mathbf{X} = \begin{bmatrix} -(\mathbf{x}^{(1)})^T - \\ -(\mathbf{x}^{(2)})^T - \\ \vdots \\ -(\mathbf{x}^{(i-1)})^T - \\ -(\mathbf{x}^{(i+1)})^T - \\ \vdots \\ -(\mathbf{x}^{(m)})^T - \end{bmatrix}, \mathbf{y} = \begin{bmatrix} \mathbf{y}^{(1)} \\ \mathbf{y}^{(2)} \\ \vdots \\ \mathbf{y}^{(i-1)} \\ \mathbf{y}^{(i+1)} \\ \vdots \\ \mathbf{y}^{(m)} \end{bmatrix}$$

- Compute  $\hat{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$
- $\hat{\mathbf{y}}^{(i)} = \hat{\theta}^T \mathbf{x}^{(i)}$
- Compute  $err(i) = (\mathbf{y}^{(i)} - \hat{\mathbf{y}}^{(i)})^2$

end for

Compute

$$TestMSE = \frac{1}{m_{tot}} \sum_{i=1}^m err(i), NormalizedTestMSE = TestMSE / \left( \sum_{i=1}^{m_{tot}} (\mathbf{y}^{(i)})^2 \right)$$

## 3 Gradient Descent (GD) and Stochastic GD

The GD approach is an iterative algorithm to find a local minimizer of a cost function. Which minimizer is found depends on how one initializes. It does not always converge and of course local minimizer may not be a global minimizer either. But if cost function is convex, it will converge to a global minimizer. Moreover if minimizer is also unique (cost is strictly convex), then the only correct solution can be found. An example of this is the squared loss for linear regression. It is convex always. It is strictly convex if  $m \geq n$  and  $\mathbf{X}$  has full rank  $n$ .

The GD Algorithm to minimize any cost  $J(\theta)$  is as follows. Recall

$$J(\theta) := \sum_{i=1}^m cost(\theta, \mathbf{x}^{(i)}, y^{(i)})$$

where *cost* is squared error (as in previous section) or can be something else (as in later sections on Logistic Regression).

1. Initialize  $\hat{\theta}$  as the zero vector  
(or anything else: this choice is problem dependent, for linear regression zero vector is good)
  2. For  $t = 1$  to  $Tmax$  do  
(use a large value for  $Tmax$ )
    - (a)  $\hat{\theta}_{old} \leftarrow \hat{\theta}$
    - (b)  $\hat{\theta} \leftarrow \hat{\theta} - \mu \nabla_{\theta} J(\hat{\theta})$
    - (c) Exit loop if  $\|\hat{\theta} - \hat{\theta}_{old}\|_2 / \|\hat{\theta}_{old}\|_2 \leq \epsilon$  with  $\epsilon$  a very small tolerance, e.g., set  $\epsilon = 10^{-6}$  (this value actually also depends on noise level)
- End for

### 3.1 Setting the step size $\mu$

- When using  $J(\theta)$  as above (it is the sum of  $m$  terms), then use  $\mu = c/m$  with  $c$  being a number between 0 and 1.
- Rule of thumb: Reduce  $c$  if the cost seems to not decrease at all or starts increasing. Increase it if the cost decreases but very slowly with iteration.
  - Detailed: start with  $c = 0.5$ , check if error goes down with iterations, if not it is too large. Use a smaller  $c$  say  $c = 0.05$ . Check again. If still does not go down, use an even smaller step size.
- As you reduce  $c$ , you will need MORE total number of iterations.
- Good coding practice: Always use a  $T_{max}$ : maximum number of iterations and an exit criterion.
- Extra info: *Variable step size: Can also vary  $\mu$  with iterations, start with a larger value but reduce with iterations.*

### 3.2 Extra material: Stochastic GD (SGD) and Stochastic Mini-Batch GD \*\*

$J(\theta)$  is typically an average of  $m$  terms; in fact it always is under our assumption of different training data points being i.i.d. As a result, its gradient is also a sum of  $m$  terms divided by  $m$ .

If  $m$  is large, computing the full gradient at each iteration can be expensive. Also, sometimes not all data is available immediately.

Stochastic GD idea: sum over a subset of the  $m$  gradients at each iteration. Pick this subset randomly or use other strategies.

Details: see videos mentioned next.

#### 3.2.1 Good videos on stochastic and mini-batch SGD

Lectures 17.1 - 17.3 of Andrew Ng's videos available at

[https://www.youtube.com/watch?v=PPLop4L2eGk&list=PLLssT5z\\_DsK-h9vYZkQkYNWcItqh1RJLN](https://www.youtube.com/watch?v=PPLop4L2eGk&list=PLLssT5z_DsK-h9vYZkQkYNWcItqh1RJLN)

## 4 Supervised Learning, Unsupervised Learning, Model-based Learning: basics

Given observed data (or features of the observed data) or other “input”  $\mathbf{x}$ , the goal is to predict some function of  $\mathbf{x}$  that is denoted by  $y$  (often called “output”). In what we have talked so far  $\mathbf{x}$  is an  $n \times 1$  vector and  $y$  is a scalar. For example,  $\mathbf{x}$  can be the feature vector of key attributes of a house, while  $y$  can be its price. In this case both are real valued. Or  $y$  can be a binary decision about whether a buyer buys the house or not.



In learning, we first decide a modeling strategy to model the input-output relationship; then come up with an algorithm to “learn” parameters given the training data (which is a set of  $m$  input-output pairs in case of supervised learning). All of this is done so that the “learnt model” can be used to predict  $y$  (get  $\hat{y}$ ) for a new query  $\mathbf{x}$ .

“Predict” is often also called “estimate” (if  $y$  is real-valued) and it is also called “detect” or “classify” (if  $y$  is binary/discrete-valued).

Learning algorithms can be supervised or unsupervised.

In supervised learning, we are provided with “training data” that allows us to “learn” the parameters used by the model that our algorithm relies on. Goal is to predict  $y$  using observed data or features  $\mathbf{x}$ .

$\mathbf{x}$  is  $n \times 1$ ,  $y$  is a scalar. We use  $\theta$  to denote the set of parameters used by our assumed model.

In many settings, the assumed model that predicts  $y$  is denoted  $h_\theta(\mathbf{x})$ . Since the model is never perfect, we assume that the “true” output  $y$  satisfies

$$y = h_\theta(\mathbf{x}) + e$$

where  $e$  is the *modeling error or noise*. This is typically modeled as a random variable with a probability density function (PDF), typically zero mean Gaussian and independent and identically distributed (i.i.d.) in each new sample.

Training data consists of  $m$  input/output pairs  $\{\mathbf{x}^{(i)}, y^{(i)}\}, i = 1, 2, \dots, m$ . The modeling error / noise  $e$  We use these to “learn”  $\theta$ . Once that is done, we can predict  $y$  from  $\mathbf{x}$  using the above equation.

## 4.1 Examples

Examples of supervised learning include: linear regression, classification - logistic regression or GDA or naive-bayes/spam-filter, SVMs, deep learning.

In case of unsupervised learning, we are just provided data and we need to “make sense” of it. This can mean different things. In case of PCA this means finding a lower dimensional subspace to represent the data. In case of clustering, this means partition the data into  $K$  disjoint “clusters”;  $K$  itself may be known or unknown.

In Model-based Learning, the observed data model is provided (usually the Physics of the problem defines this). Examples include least squares (LS) problems, Bayesian LS, Bayesian estimation, or structured data recovery problems such as sparse recovery (compressive sensing) or low-rank matrix recovery. Standard LS is usually the oversampled data setting. Bayesian LS or structure models are used in case of undersampled data - some data is missing, corrupted, or deliberately undersampled.

The model learning (regression coefficients’ vector learning) step of linear regression is an instance of LS. If limited training data is available, and we model the regression vector as being sparse, this can become an instance of sparse recovery.

Unsupervised deep learning, which is a newish area, in which one uses a single training data (that is a sequence of vectors or images or image Fourier transforms) to both learn the deep model parameters and predict the output - falls in the second or third category. Used in accelerated dynamic MRI applications.

## 5 Supervised Learning: Linear Regression

In the setting we have talked about in class,  $\mathbf{x}$  is a real-valued  $n \times 1$  vector and  $y$  is a real-valued scalar <sup>1</sup>. The parameter vector  $\theta$  is also an  $n \times 1$  vector

### 5.1 Model

In linear regression,  $h_\theta(\mathbf{x})$  is a linear function of  $\mathbf{x}$ .

$$h_\theta(\mathbf{x}) = \theta^T \mathbf{x}.$$

---

<sup>1</sup>In more general settings  $y$  can also be a real-valued vector (this will not be discussed in our class).

(I sometimes may use ' for transpose – MATLAB notation)

## 5.2 Including a nonzero mean in the model

To include a non-zero mean in the model, one can replace both  $\mathbf{x}$  and  $\theta$  by  $n + 1$  length vectors as follows. Let

$$\tilde{\theta} = \begin{bmatrix} \theta_0 \\ \theta \end{bmatrix}$$

and

$$\tilde{\mathbf{x}} = \begin{bmatrix} 1 \\ \mathbf{x} \end{bmatrix}$$

and we let

$$h_{\tilde{\theta}}(\tilde{\mathbf{x}}) = \tilde{\theta}^T \tilde{\mathbf{x}}.$$

With this model, we do everything explained above using  $\tilde{\mathbf{x}}$  and  $\tilde{\theta}$  to replace  $\theta$  and  $\mathbf{x}$  respectively.

Here and below, we let  $\theta$  be an  $n = d + 1$  length vector and same for  $\mathbf{x}^{(i)}$  (append 1 as its first entry).

## 5.3 Learning $\theta$ : minimize squared loss

Here and below, we let  $\theta$  be an  $n = d + 1$  length vector and same for  $\mathbf{x}^{(i)}$  (append 1 as its first entry).

The most common approach to learn  $\theta$  is to assume a squared error loss and try to minimize it, i.e., find

$$\arg \min_{\theta} J(\theta) := \sum_{i=1}^m (y^{(i)} - \theta^T \mathbf{x}^{(i)})^2.$$

Define an  $m \times n$  input data matrix  $\mathbf{X}$  with  $[1, (\mathbf{x}^{(i)})^T]$  as its rows, i.e., let

$$\mathbf{X} = \begin{bmatrix} 1, & -(\mathbf{x}^{(1)})^T - \\ 1, & (\mathbf{x}^{(2)})^T - \\ \vdots, & \vdots \\ 1, & -(\mathbf{x}^{(m)})^T - \end{bmatrix} \quad (1)$$

and define an  $m \times 1$  vector  $\mathbf{y}$  with  $y^{(i)}$  as its columns.

Then,  $J(\theta)$  can be expressed more compactly as

$$J(\theta) := \sum_{i=1}^m (y^{(i)} - \theta^T \mathbf{x}^{(i)})^2 = \|\mathbf{y} - \mathbf{X}\theta\|_2^2$$

### 5.3.1 Understanding squared loss: Maximum Likelihood Estimation under i.i.d. Gaussian model \*\*

The above can be motivated as Maximum Likelihood Estimation: maximize  $p(\mathbf{y}; \mathbf{X}, \theta)$  under the model  $y^{(i)} = h_{\theta}(\mathbf{x}^{(i)}) + e^{(i)}$ ,  $i = 1, 2, \dots, m$  with  $e^{(i)}$  independent identically distributed (i.i.d.) standard Gaussian. Here the randomness is only in the noise  $e$ .

## 5.4 Solutions for minimizing squared loss: also called Least Squares (LS) Estimation

1. We can get a closed form solution by taking the derivative of  $J(\theta)$  w.r.t.  $\theta$  and setting it to zero. When  $\mathbf{X}$  is full rank  $n$  (a necessary condition for this is  $m \geq n$ ), this simplifies to

$$\hat{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$$

*Details:*  $\nabla J(\theta) = 2\mathbf{X}^\top(\mathbf{y} - \mathbf{X}\theta)$ . Set this equal to zero and solve for  $\hat{\theta}$ . Reason: the gradient is zero at every minimizer, maximizer or stationary point of a differentiable cost function. gradient = 0 is a necessary condition for a point to be a minimizer. in this case, since  $J(\cdot)$  is strongly convex it has exactly one stationary point which is THE minimizer. Thus setting  $\nabla J(\theta) = 0$  helps us find THE minimizer.

- For an  $m \times n$  matrix with  $m \geq n$ ,  $(\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top$  is the pseudo-inverse of  $\mathbf{X}$ , denoted  $\mathbf{X}^\dagger$ . Thus, we can also write the above solution as

$$\hat{\theta} = \mathbf{X}^\dagger \mathbf{y}$$

*Extra details:* The pseudo-inverse is computed by first computing the singular value decomposition (SVD) of the matrix  $\mathbf{X}$ . Suppose the SVD of  $\mathbf{X} = U\Sigma V^\top$  where  $U$  is an  $m \times m$  unitary matrix.  $V$  is an  $n \times n$  unitary matrix and  $\Sigma$  is diagonal with non-negative entries. Then  $\mathbf{X}^\dagger = V\Sigma^\dagger U^\top$ . For a rectangular diagonal matrix such as  $\Sigma$ , we get the pseudo-inverse by taking the reciprocal of each non-zero element on the diagonal, leaving the zeros in place, and then transposing the matrix.

If algorithms for computing matrix inverse or matrix pseudo-inverse were "exact" (and not iterative) then the above two approaches would return the exact same solution. But they are not. Thus, when  $\mathbf{X}$  is "well-conditioned", both approaches above return same solution, but not otherwise.

- For large sized problems (where  $n$  is large), using Gradient Descent (GD) is a better idea. Since problem is convex, GD should, in principle, converge to above solution starting from any initialization, and should converge pretty quickly. We explain this in Sec. 3.
- Approximate but even faster solution:* Stochastic GD (S-GD) can be used. Advantages: faster per iteration; needs lesser memory; and is useful to get a fully streaming algorithm. But no easy guarantees on whether it will converge and to what.

## 5.5 Least Squares, Regularized LS, Recursive LS and more

See Sec. 11.

# 6 Supervised Learning: Two class and Multi-class Logistic Regression: Summary

## 6.1 Two class Logistic Regression

Let  $\theta$  be a  $n = d + 1$  length vector with the first entry corresponding to the **bias** term. Replace  $\mathbf{x}$  by  $[1, \mathbf{x}^\top]^\top$  (this it makes it a  $n = d + 1$  length vector

$$h_\theta(\mathbf{x}) = g(\theta^\top \mathbf{x}), \quad g(z) := \frac{1}{1 + e^{-z}}$$

$g(z) = \frac{1}{1 + \exp(-z)} = \frac{\exp(z)}{\exp(z) + 1}$  is called the sigmoid function,  $0 \leq g(z) \leq 1$  for all  $z$ .

- Let  $\theta$  be a  $n = d + 1$  length vector with the first entry corresponding to the **bias** term. Replace  $\mathbf{x}$  by  $[1, \mathbf{x}^\top]^\top$  (this it makes it a  $n = d + 1$  length vector
- Given training data  $\mathbf{X}, \mathbf{y}$ , use cost func as negative log likelihood. Assume independence of different data pairs.

$$J(\theta) := -\log p(\mathbf{y}|\mathbf{X}; \theta) = -\log \left( \prod_{i=1}^m p(\mathbf{y}^{(i)}; \mathbf{x}^{(i)}, \theta) \right) = -\log \left( \prod_{i=1}^m h_\theta(\mathbf{x}^{(i)})^{y^{(i)}} (1 - h_\theta(\mathbf{x}^{(i)}))^{1-y^{(i)}} \right)$$

- We can estimate  $\theta$  by minimizing  $J(\theta)$  by GD. After simplification the gradient has the following expression:

$$\nabla J(\theta) = \sum_{i=1}^m (g(\theta^\top \mathbf{x}^{(i)}) - \mathbf{y}^{(i)}) \mathbf{x}^{(i)}$$

– Use algorithm given in Sec. 3 with above gradient

- Classification: Given a query,  $\mathbf{x}$ , predict  $y$  as

$$\hat{y} = \text{round}(g(\theta^\top \mathbf{x}))$$

here  $\text{round}(\cdot)$  rounds off its input to either 0 or 1.

Or,

$\hat{y} = 1$  if  $g(\theta^\top \mathbf{x}) > 0.5$  and  $\hat{y} = 0$  otherwise

Or, equivalently, using the fact that  $g(z)$  is an increasing function and  $g(0) = 0.5$ ,

$\hat{y} = 1$  if  $\theta^\top \mathbf{x} > 0$  and  $\hat{y} = 0$  otherwise

- **Note: do not use the above model for simulating data on which to test the algorithm: for simulating data, better to use the GDA generative model.**

## 6.2 Multi-class Logistic Regression

Multi-class logistic regression (LR) is known by a variety of other names, including polytomous LR, Multinomial LR, softmax regression, multinomial logit (mlogit), the maximum entropy (MaxEnt) classifier, and the conditional maximum entropy model.

We now have  $K > 2$  classes. Recall that, for  $K = 2$ , the model is

$$\Pr(y = 1) = \frac{1}{1 + \exp(-\theta^\top \mathbf{x})} = \frac{\exp(\theta^\top \mathbf{x})}{1 + \exp(\theta^\top \mathbf{x})}$$

and

$$\Pr(y = 0) = 1 - \frac{\exp(\theta^\top \mathbf{x})}{1 + \exp(\theta^\top \mathbf{x})}$$

To extend this to  $K$  classes labeled  $k = 0, 1, \dots, K - 1$ , instead of a single "regression vector"  $\theta$ , we now use  $K - 1$  "regression" vectors  $\beta_k, k = 1, 2, \dots, K - 1$  and use the model:

$$\Pr(y = k) = \frac{\exp(\beta_k^\top \mathbf{x})}{1 + \sum_{k'=1}^{K-1} \exp(\beta_{k'}^\top \mathbf{x})}, \quad k = 1, 2, \dots, K - 1$$

and

$$\Pr(y = 0) = \frac{1}{1 + \sum_{k'=1}^{K-1} \exp(\beta_{k'}^\top \mathbf{x})}$$

If we let  $\beta_0 = \mathbf{0}$  the  $y = 0$  case can be interpreted as

$$\Pr(y = 0) = \frac{\exp(\beta_0^\top \mathbf{x})}{1 + \sum_{k'=1}^{K-1} \exp(\beta_{k'}^\top \mathbf{x})}, \quad \text{with } \beta_0 = \mathbf{0}$$

With setting  $\beta_0 = \mathbf{0}$ , observe that  $1 = \exp(\beta_0^\top \mathbf{x})$ . Replace this in both numerator and denominator of above expressions. Then, the above can be simplified to

$$\Pr(y = k) = \frac{\exp(\beta_k^\top \mathbf{x})}{\sum_{k'=0}^{K-1} \exp(\beta_{k'}^\top \mathbf{x})}, \quad k = 1, 2, \dots, K - 1$$

- Learning the parameters:

- Need to estimate  $\theta := \{\beta_1, \beta_2, \dots, \beta_{K-1}\}$ .
  - Solution 1: define the joint log-likelihood function over all  $m$  training data points and maximize it (minimize its negative).
  - This solution often does not work too well (too many parameters and one may start becoming too large).
  - Solution 2: add regularization on the  $\beta_k$ 's, e.g. add  $\lambda \sum_{k=1}^{K-1} \|\beta_k\|^2$  to the cost function and then minimize it
  - Basic GD may not work easily in this complicated setting. To explore this case, see what Python or MATLAB's built in toolboxes do for your data.
- Classification: find the  $k$  for which  $\Pr(y = k)$  is highest. Equivalently, taking log of the probabilities, using  $\beta_0 = \mathbf{0}$  (zero vector), and ignoring the denominator, this translates to

$$\hat{y} = \arg \max_{k=0, \dots, K-1} \beta_k^\top \mathbf{x}$$

### 6.3 One-vs-rest (OVR) or One-vs-all use of two-class logistic regression

There are two ways to deal with multiple classes using Log Reg. Say we have  $K$  classes,  $k = 0, 1, \dots, K-1$ .

1. Option 1: use what I have above
2. Option 2: use two class log reg in one-vs-rest or OVR mode. Idea of OVR:
  - First learn a  $\theta_0$  for separating classes into 0 and everything else (combine classes 1,2,3, ... K-1 into the “not-zero” class)
  - Then learn a  $\theta_1$  for separating classes into 1 and everything else (combine classes 0,2,3, ... K-1 into the “not-one” class)
  - Then learn a  $\theta_2$  for separating classes into 2 and everything else (combine classes 0,1,3, ... K-1 into the “not-two” class)
  - And so on.
  - Output of above set of steps will be  $K$  regression vectors  $\theta_0, \theta_1, \dots, \theta_{K-1}$
  - Classification on a query  $\mathbf{x}$ :

$$\hat{y} = \arg \max_{k=0,1,2,\dots,K-1} \theta_k^\top \mathbf{x}$$

3. More details: see the video below
4. Also see [https://en.wikipedia.org/wiki/Multiclass\\_classification](https://en.wikipedia.org/wiki/Multiclass_classification)

Video on One-vs-Rest (or One-vs-All) use of two-class Log Reg:

Lecture 6.7 of

[https://www.youtube.com/watch?v=PPLop4L2eGk&list=PLLssT5z\\_DsK-h9vYZkQkYNWcItqhlR.JLN](https://www.youtube.com/watch?v=PPLop4L2eGk&list=PLLssT5z_DsK-h9vYZkQkYNWcItqhlR.JLN)

## 7 Supervised Learning: Two-class Logistic Regression Details

In this case,  $\mathbf{x}$  is still a real-valued  $d \times 1$  vector but now  $y$  is a binary scalar. The goal is classification.

## 7.1 Two-class Model and its probabilistic interpretation

This assumes that  $\Pr(y = 1; \mathbf{x}, \theta) = h_\theta(\mathbf{x})$  with

$$h_\theta(\mathbf{x}) = g(\theta^T \mathbf{x}), \quad g(z) := \frac{1}{1 + e^{-z}}$$

$g(\cdot)$  is called the sigmoid function, it takes values between zero and one for all values of  $z$ . Thus, it can be used to model a probability. Said another way,

$$p(y; \mathbf{x}, \theta) = h_\theta(\mathbf{x})^y (1 - h_\theta(\mathbf{x}))^{1-y}$$

The prediction is

$$\hat{y} = \arg \max_{y=0,1} p(y; \mathbf{x}, \theta)$$

Thus

$$\hat{y} = 1 \text{ if } h_\theta(\mathbf{x}) > 1 - h_\theta(\mathbf{x}),$$

and  $\hat{y} = 0$  otherwise.

### 7.1.1 Use of bias term in Logistic regression

Introduce the bias term exactly as we did in case of linear regression. Make both  $\theta$  and  $\mathbf{x}$   $(n + 1)$ -th length vectors; set the first entry of  $\mathbf{x}$  equal to 1. The first entry of  $\theta$  is then the bias term.

## 7.2 Learning $\theta$ : Maximum Likelihood Estimation

Again define  $\mathbf{y}$  and  $\mathbf{X}$  as before from training data

Use Maximum Likelihood Estimation again: assume i.i.d. training data points  $y^{(i)}$  (recall that this was assumed also in linear regression – it was imposed by letting the  $w^{(i)}$ 's be i.i.d.).

Thus, we minimize the negative log likelihood,

$$J(\theta) := -\log p(\mathbf{y}|\mathbf{X}; \theta) = -\log \left( \prod_{i=1}^m p(\mathbf{y}_i; \mathbf{x}^{(i)}, \theta) \right) = -\log \left( \prod_{i=1}^m h_\theta(\mathbf{x}^{(i)})^{y^{(i)}} (1 - h_\theta(\mathbf{x}^{(i)}))^{1-y^{(i)}} \right)$$

We can simplify this a lot as follows.

$$\begin{aligned} J(\theta) &= -\sum_{i=1}^m \log \left( \left( \frac{1}{1 + \exp(\theta' \mathbf{x}^{(i)})} \right)^{y^{(i)}} \left( \frac{\exp(\theta' \mathbf{x}^{(i)})}{1 + \exp(\theta' \mathbf{x}^{(i)})} \right)^{1-y^{(i)}} \right) \\ &= -\sum_{i=1}^m \log \left( \left( \frac{1}{1 + \exp(\theta' \mathbf{x}^{(i)})} \right)^{y^{(i)}} \left( \frac{1}{1 + \exp(-\theta' \mathbf{x}^{(i)})} \right)^{1-y^{(i)}} \right) \end{aligned}$$

The only difference between the first and second terms inside  $\log(\cdot)$  is the sign of  $\theta' \mathbf{x}^{(i)}$  in the denominator and the power it is raised to. One is raised to the power  $y^{(i)}$ , the other is raised to the power  $(1 - y^{(i)})$ . Here  $y^{(i)}$  is either 0 or 1. So,  $2y^{(i)} - 1$  is either  $-1$  (when  $y^{(i)} = 0$ ) or  $+1$  (when  $y^{(i)} = 1$ ). Thus, when  $y^{(i)}$  takes only values 0 or 1,

$$\left( \frac{1}{1 + \exp(\theta' \mathbf{x}^{(i)})} \right)^{y^{(i)}} \left( \frac{1}{1 + \exp(-\theta' \mathbf{x}^{(i)})} \right)^{1-y^{(i)}} = \frac{1}{1 + \exp((\theta' \mathbf{x}^{(i)})(2y^{(i)} - 1))}$$

The reason the last equality is true is because when  $y^{(i)} = 1$ ,  $2y^{(i)} - 1 = 1$ , but when  $y^{(i)} = 0$ ,  $2y^{(i)} - 1 = -1$ .

We now simply get

$$J(\theta) = - \sum_{i=1}^m \log \left( \frac{1}{1 + \exp((\theta' \mathbf{x}^{(i)})(2y^{(i)} - 1))} \right) = \sum_{i=1}^m \log \left( 1 + \exp \left( (\theta' \mathbf{x}^{(i)})(2y^{(i)} - 1) \right) \right)$$

We can find  $\theta$  by minimizing  $J(\theta)$  by GD.

It is possible to show that  $J(\theta)$  is convex. Argument: weighted sum convex functions is convex when the weights are positive, here the weights are just 1;  $(\theta' \mathbf{x})(2y - 1)$  is an affine function of  $\theta$  (and hence is both convex and concave; the logistic function  $\log(1 + \exp(-z))$  can be shown to be convex (see next line); composition of a convex function and an affine function is convex. To show  $\log(1 + \exp(-z))$  is convex, since it is twice differentiable, we can compute the second derivative and show that it is  $-1/(1 + \exp(-z))^2 < 0$  for any  $z$ , thus it is convex everywhere.

## 8 Supervised Learning: Generative Learning (a.k.a. Bayesian models/learning)

For logistic or linear regression we just assumed a probabilistic model on how  $y$  is generated from  $\mathbf{x}$ , with  $\mathbf{x}$  being deterministic.

In Generative Learning, we assume a “generative model”: we first put a prior probabilistic model on  $y$ , and then assume a probabilistic model on how  $\mathbf{x}$  was generated from  $y$ . We then compute the probability (or probability density function in case  $y$  is real-valued) of  $y$  taking a certain value given  $\mathbf{x}$  using Bayes rule. Mathematically, we assume that we are given

$$p(\mathbf{x}|y; \theta), p(y)$$

and we use these to obtain the prediction as follows

$$\hat{y} = \arg \max_y p(y|\mathbf{x}; \theta) := \arg \max_y \frac{p(\mathbf{x}|y; \theta)p(y; \theta)}{p(\mathbf{x}; \theta)} = \arg \max_y p(\mathbf{x}|y; \theta)p(y; \theta)$$

This use of Bayes rule is called **Maximum A Posteriori (MAP)** detection or estimation in other literature. The overall approach is often called Bayesian modeling or physics-based modeling.

### 8.1 Learning $\theta$ : Maximum Likelihood Estimation

Estimate  $\theta$ : define  $\mathbf{y}$ ,  $\mathbf{X}$  as before from training data. Also assume training data points are independent:  $\{\mathbf{x}^{(i)}, y^{(i)}\}$  are mutually independent for different  $i$ . Define the cost function

$$J(\theta) := \Pr(\mathbf{y}, \mathbf{X}; \theta) = \prod_{i=1}^m p(y^{(i)}, \mathbf{x}^{(i)}; \theta)$$

or, usually its logarithm, and maximize it over  $\theta$ .

### 8.2 Generative Learning: Gaussian Discriminant Analysis (GDA)

This is one type of generative model and learning algorithm for the setting  $\mathbf{x}$  real-valued  $n \times 1$  vector and  $y$  binary scalar. Thus  $y$  can take two values 0 or 1. It assumes  $\mathbf{x}$  is Gaussian given  $y$  and  $y$  itself is Bernoulli, i.e.,

$$p(\mathbf{x}|y; \theta) = \mathcal{N}(\mathbf{x}; \mu_y, \Sigma_y), \quad p(y) = \phi^y(1 - \phi)^{1-y}$$

Notice in this case  $\theta = \{\mu_0, \mu_1, \Sigma_0, \Sigma_1, \phi\}$ . Parameters are still learnt by MLE

$$\max_{\theta} J(\theta) := \Pr(\mathbf{y}, \mathbf{X}; \theta) = \prod_{i=1}^m p(y^{(i)}, \mathbf{x}^{(i)}; \theta) \text{ s.t. } 0 \leq \phi \leq 1$$

Notice that, without extra assumptions, we have  $2n + 2n^2 + 1$  parameters. Training data are each  $n$  length vectors  $\mathbf{x}^{(i)}$ , thus we can say we have  $mn$  training data scalars. We need  $mn$  significantly larger than  $2n + 2n^2 + 1$  for training/learning to be accurate. We will need  $m$  growing at least linearly with  $n$  to be able to learn anything useful.

But the point of Bayesian (generative) modeling is that we should be able to use a smaller  $m$  and still train well.

When enough training data is not available, we need to simplify our model so that there are fewer parameters. As explained later, this will increase model bias, but will reduce the variance in parameter estimation.

A common model simplification is to assume that the different entries of each  $\mathbf{x}^{(i)}$  are independent conditioned on the class label  $y^{(i)}$ . This is called the *Naive Bayes assumption*.

### 8.2.1 Gaussian Discriminant Analysis (GDA) with Naive Bayes assumption and equal covariances

A common model simplification is to assume that the different entries of each  $\mathbf{x}^{(i)}$  are independent conditioned on the class label  $y^{(i)}$ . This is called the *Naive Bayes assumption*.

In the Gaussian case, this translates to assuming that  $\Sigma_0, \Sigma_1$  are *diagonal*. With the diagonal assumption, we now have only  $2n + 2n + 1$  parameters which is much more manageable. A second commonly used simplification is to assume the same covariance under both classes, i.e., that  $\Sigma_0 = \Sigma_1 = \Sigma$  and  $\Sigma$  is diagonal. With this assumption too, we have the following simpler model

$$\prod_{i=1}^m \left( \left( \prod_{j=1}^n \mathcal{N}(\mathbf{x}_j; (\mu_{y_i})_j, \sigma_j^2) \right) \cdot \phi^{y_i} (1 - \phi)^{1-y_i} \right)$$

Under the above assumption, the Max Likelihood Estimates (MLE) of the model parameters,  $\theta := \{\phi, \mu_0, \mu_1, \Sigma\}$  i.e., the value of the model parameters that solve

$$\arg \max_{\theta} J(\theta), \quad J(\theta) := \Pr(\mathbf{y}, \mathbf{X}; \theta) = \prod_{i=1}^m p(y^{(i)}, \mathbf{x}^{(i)}; \theta) \text{ s.t. } 0 \leq \phi \leq 1$$

are computed as follows:

$$\begin{aligned} \hat{\phi} &= \frac{1}{m} \sum_{i=1}^m \mathbf{1}(y^{(i)} = 1) \\ \hat{\mu}_0 &= \frac{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 0) \mathbf{x}^{(i)}}{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 0)} \\ \hat{\mu}_1 &= \frac{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 1) \mathbf{x}^{(i)}}{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 1)} \\ \hat{\sigma}_j^2 &= \frac{1}{m} \sum_{i=1}^m (\mathbf{x}^{(i)} - \mu_{y^{(i)}})_j^2, \quad j = 1, 2, \dots, n \end{aligned}$$

while setting all non-diagonal entries of  $\hat{\Sigma}$  to be zero. Here  $\mathbf{1}$  denotes the *indicator function* of the statement in paranthesis. Thus,  $\mathbf{1}(y^{(i)} = 0)$  equals one if  $y^{(i)} = 0$  and it equals zero otherwise.

Notice that the above is equivalent to learning the parameters for *each feature* independently, it can also



be rewritten as follows: for each  $j = 1, 2, \dots, n$ , compute

$$\begin{aligned}
 (\hat{\mu}_0)_j &= \frac{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 0)(\mathbf{x}^{(i)})_j}{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 0)} \\
 (\hat{\mu}_1)_j &= \frac{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 1)(\mathbf{x}^{(i)})_j}{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 1)} \\
 \hat{\sigma}_j^2 &= \frac{1}{m} \sum_{i=1}^m (\mathbf{x}^{(i)} - \mu_{y^{(i)}})_j^2, \quad j = 1, 2, \dots, n
 \end{aligned}$$

In this Naive Bayes' setting, the selecting of  $\hat{y}$  to solve

$$\hat{y} = \arg \max_y p(\mathbf{x}|y; \hat{\theta})p(y; \hat{\theta})$$

simplifies to the following test (just simplify the expressions):

$$\hat{y} = 1 \text{ if } \sum_{j=1}^n \frac{(\mathbf{x} - \hat{\mu}_1)_j^2}{\hat{\sigma}_j^2} < \sum_j \frac{(\mathbf{x} - \hat{\mu}_0)_j^2}{\hat{\sigma}_j^2}$$

### 8.2.2 Notes

1. In the Gaussian case, we get a simple closed form as above. In many cases of Bayesian modeling also, this is possible. In general when this is not possible, *do not* ever work directly with probabilities. *Always work with logarithms of the probabilities.* Otherwise you will run into numerical problems while coding. Remember: under the naive Bayes assumption,  $p(\mathbf{x}|y; \hat{\theta}) = \prod_{j=1}^n p(\mathbf{x}_j|y; \hat{\theta})$  if  $p(\cdot)$  is a PMF, it is a product of  $n$  real numbers all less than one. Even if  $p(\cdot)$  is a PDF (which could be more than one), very often it is actually less than one only. Product  $n$  numbers less than one can become very small very soon.
2. When working with real data, e.g., images, there may be a certain region that is black in *all* the images. For these pixels the estimate of the variance will be zero meaning  $\sigma_j^{-1} = \infty$  resulting in code bugs.
  - (a) Fix 1: do not use these directions.
  - (b) Fix 2: Better fix to deal with similar issues where in some directions variance is very small (may not be zero): use PCA to reduce the dimensionality of the data. For classification there is no need to use all  $n$  features.
  - (c) Fix 3: add a small value to replace the zeros: the added value should be *much smaller* than any of the *important directions' variances*. This is easy in the zero/nonzero case but in practice ill-conditioning of  $\Sigma$  may make this hard to do. Use of PCA and a smaller dimension is thus a much better fix.

## 8.3 Generative Learning: Spam Filter – an example of Discrete-valued or Categorical Features

In applications such as spam email detector (or filter) design, one typically models  $\mathbf{x}$  as a discrete-valued vector given  $y$ .

$y = 0$  means the email is not-spam,  $y = 1$  means it is spam.

### 8.3.1 Simple Spam Filter: entries of $\mathbf{x}$ are binary

In the simplest version,  $\mathbf{x}$  is an  $n \times 1$  binary vector with  $n$  being equal to the size of the English dictionary. We say  $\mathbf{x}_j = 1$  if the  $j$ -th dictionary word is in the email and  $\mathbf{x}_j = 0$  otherwise. This means  $n$  is really large. Also it is not counting how many times a word occurred.

Consider first the simplest model where  $\mathbf{x}$  is a binary vector. As before, we specify

$$p(\mathbf{x}|y; \theta), p(y)$$

and we predict

$$\hat{y} = \arg \max_y p(\mathbf{x}|y; \theta)p(y; \theta).$$

Notice now that  $\mathbf{x}$  can take a total of  $2^n$  possible values. We also need to specify the prior  $p(y)$ . Thus, in this most general case, the number of parameters equals  $2^n + 1$ .

Training data are each  $n$  length vectors  $\mathbf{x}^{(i)}$ , thus we can say we have  $mn$  training data scalars. We need  $mn$  significantly larger than  $2^n + 1$  for training/learning to be accurate. Here we will need  $m$  to grow linearly with  $2^{n-1}$ : this can be very large and is not practical.

### 8.3.2 Naive Bayes assumption

In both the above examples and especially the second one, the required  $m$  can be very large for accurate training/learning. Thus we add a further modeling assumption called “Naive Bayes” in ML literature. Others would call it “*conditional independence of different entries of a feature vector (the different  $\mathbf{x}_j$ ’s,  $j = 1, 2, \dots, n$ ) given  $y$* ”. Mathematically, we are assuming

$$p(\mathbf{x}|y; \theta) = \prod_{j=1}^n p(\mathbf{x}_j|y; \theta)$$

This may not be a very realistic assumption, but it significantly reduces the number of parameters required by the model.

In the Gaussian case, this implies that  $\Sigma_0, \Sigma_1$  are **diagonal matrices**. Thus, the number of parameters becomes  $2n + 2n + 1$  which is much more tractable. In the spam filter case, this means we have  $n + 1$  parameters.

So now the number of training samples  $m$  does not even need to grow with  $n$ .

### 8.3.3 Simple Spam filter with Naive Bayes

$\mathbf{x}$  is a binary vector,  $y$  is a scalar. With using Naive Bayes, in the Simple Spam Filter case, we can now define

$$\psi_{j,y} := p(\mathbf{x}_j = 1|y), j = 1, 2, \dots, n; \text{ and } \phi := p(y = 1)$$

With this we have just  $n + 1$  parameters to learn instead of  $2^n + 1$ .

We can again learn the parameters by MLE:

$$\max_{\theta} J(\theta) := \prod_{i=1}^m p(y^{(i)}, \mathbf{x}^{(i)}; \theta) = \prod_{i=1}^m \prod_{j=1}^n p(\mathbf{x}_j^{(i)}|y^{(i)}; \theta)p(y^{(i)}; \theta) = \prod_{i=1}^m \phi^{y^{(i)}} (1-\phi)^{(1-y^{(i)})} \left( \prod_{j=1}^n \psi_{j,y}^{\mathbf{x}_j^{(i)}} (1-\psi_{j,y})^{(1-\mathbf{x}_j^{(i)})} \right)$$

s.t. constraints that

$$0 \leq \psi_{j,y} \leq 1, 0 \leq \phi \leq 1$$

Can again get closed form simple expressions for the MLE:

$$\hat{\psi}_{j,0} =$$

count the number of training data points for which  $y^{(i)} = 0$  and  $\mathbf{x}_j^{(i)} = 1$  and divide by the total number of training data points for which  $y^{(i)} = 0$ .

Similarly

$$\hat{\psi}_{j,1} =$$

count the number of training data points for which  $y^{(i)} = 1$  and  $\mathbf{x}^{(i)} = 1$  and divide by the total number of training data points for which  $y^{(i)} = 1$ . and

$$\hat{\phi} =$$

count the number of training data points for which  $y^{(i)} = 1$  and divide by  $m$

Use of  $n$  to be dictionary size makes it extremely large causing the algorithm to be very slow. Also, instead of letting  $n$  be the size of the English dictionary, we can let  $n$  be the size of the vocabulary (set of all words in all training data). But this has the disadvantage that it does not tell you how to deal with an unseen word. Options include (i) ignore unseen words (can be problematic in case the unseen words are the only reason an email is obviously spam); or (ii) in the model, assume a small nonzero probability for an unseen word (this means: increase the vocabulary size by 1, this probability cannot be learned easily, you just have to make up a “reasonable” value for it).

### 8.3.4 Spam Filter: entries of $\mathbf{z}$ count the number of times a word occurs

The simplest spam filter explained above is not using word-counts (the number of times a word occurs), but just checking whether a word is present or not. A better model is to consider a different feature vector  $\mathbf{z}$  with  $z_j$  being the number of times word  $j$  occurs in the vocabulary (or dictionary).

$d$ : number of words in any email – to use this model, we can use blank-space as one “word”.

$n$ : number of words in dictionary

$m$ : number of training emails

Treating each word occurrence as independent (this is the naive Bayes assumption, it is actually not true since certain pairs of words are much more likely occur together, but simplifies our modeling), and assuming as before that  $\psi_{j,y} = \Pr(\text{word } j \text{ is present in the email} | y)$ , the feature vector  $\mathbf{z}$  would be modeled by what is called a “multinomial distribution” as follows. This assumes that there are a total of  $d$  words in the sample, i.e., that  $\sum_j z_j = d$  for all training emails. *To use this model, we can use blank-space as one “word”.*

$$p(\mathbf{z}|y) = \binom{d}{z_1, z_2, \dots, z_n} \prod_{j=1}^n \psi_{j,y}^{z_j}$$

This is called the multinomial distribution with parameters  $\psi_{j,y}$ .

**MLE** The joint likelihood of  $m$  independent training samples can be expressed as follows:

$$J(\theta) = \prod_{i=1}^m p(y^{(i)}) p(\mathbf{z}^{(i)} | y^{(i)}) = \prod_{i=1}^m \phi^{y^{(i)}} (1 - \phi)^{(1-y^{(i)})} \binom{d}{z_1^{(i)}, z_2^{(i)}, \dots, z_n^{(i)}} \prod_{j=1}^n \psi_{j,y^{(i)}}^{z_j^{(i)}}$$

Learning parameters: we need to maximize the above over  $\phi, \psi_{j,y}$  subject to the constraints that

$$0 \leq \psi_{j,y} \leq 1, \quad 0 \leq \phi \leq 1$$

It can be shown that we get

$$\hat{\psi}_{j,0} = \frac{\sum_{i=1}^m z_j^{(i)} \mathbb{1}(y^{(i)} = 0)}{\sum_{j=1}^n \sum_{i=1}^m z_j^{(i)} \mathbb{1}(y^{(i)} = 0)}$$

In words this is the total number of times word  $j$  occurs in **all** emails that are not spam (for which  $y^{(i)} = 0$ ) divided by the total number of words in **all** emails that are not spam. We can similarly compute

$$\hat{\psi}_{j,1}$$

### 8.3.5 Spam filter: model 3

A third possible model is as follows. Let  $j$  referred to the  $j$ -th word in the email and let  $\mathbf{x}_j$  be the location of the  $j$ -th word in the dictionary. So then this becomes “categorical data”. We explain below how to deal with categorical data. In this case the length of the feature vector can be different for different emails. It will be  $n^{(i)}$  if the email has  $n^{(i)}$  distinct words.

### 8.3.6 More ideas and related problems

A more advanced document model models co-occurrences of words.

Similar ideas are also used to classify blogs or webpages into various categories as well.

### 8.4 Modeling Categorical data \*\*

In some problems, the different features may be “categorical” instead of binary, which means feature  $j$  takes one of  $K_j$  possible values. For example  $\mathbf{x}_j$  could be color of the front door of a house in case of the house price example and we assume  $K_j = 5$  possible colors for example. There is no ordering to which color is preferred, thus the integer labels are arbitrary; they do not have a numerical meaning.

The model would then require us to learn  $\psi_{j,y,k} := \Pr(\mathbf{x}_j = k|y)$  for  $k = 1, 2, \dots, K_j$ ,  $y = 0, 1$ ,  $j = 1, 2, \dots, n$ .

Our data model then is as follows

$$p(\mathbf{x}|y) = \prod_{j=1}^n \prod_{k=1}^{K_j} \psi_{j,y,k}^{[\mathbf{x}_j==k]}$$

where  $[x == k]$  takes the value 1 if  $x = k$  and zero otherwise (MATLAB notation). Thus for MLE we need to maximize

$$\prod_i \left( \prod_{j=1}^n \prod_{k=1}^{K_j} (\psi_{j,y^{(i)},k})^{[\mathbf{x}_j^{(i)}==k]} \phi^{y^{(i)}} (1 - \phi)^{1-y^{(i)}} \right)$$

s.t. sum to one constraints on all the probabilities.

Then MLE estimates are given by

$$\hat{\psi}_{j,0,k} =$$

counts the number of times  $y^{(i)} = 0$  and  $\mathbf{x}_j^{(i)} = k$  in the training data and divides this by the number of times  $y^{(i)} = 0$  (number of points from class zero in the training data). Do this for every value of  $k = 1, 2, \dots, K_j$  and for every feature  $\mathbf{x}_j$ ,  $j = 1, 2, \dots, n$ .

Do the same for class label 1.

### 8.5 Dealing with discrete-valued and real-valued data together in the Generative Learning (a.k.a. Bayesian) framework \*\*

Often some of the features can be categorical and some of them can be real-valued. Once we impose the Naive Bayes assumption (conditioned on the class label, different features are independent), this is easy to deal with. For simplicity suppose that the first  $r$  features are real-valued and the rest  $n - r$  are categorical with  $k_j$  categories. Also assume the real-valued feature follow a Gaussian distribution with mean  $\mu_0$  or  $\mu_1$  (depending on the class label) and covariance matrix  $\Sigma$ . By Naive Bayes,  $\Sigma$  is diagonal.

Then, joint likelihood becomes

$$\prod_{i=1}^m \left( \left( \prod_{j=1}^r \mathcal{N}(\mathbf{x}_j^{(i)}; (\mu_{y^{(i)}})_j, \sigma_j^2) \cdot \prod_{j=r+1}^n \prod_{k=1}^{K_j} \psi_{j,y^{(i)},k}^{[\mathbf{x}_j^{(i)}==k]} \right) \cdot \phi^{y^{(i)}} (1 - \phi)^{1-y^{(i)}} \right)$$

s.t. sum to one constraints on  $\phi$  and on all the  $\psi_{j,y,k}$ 's.

The MLE estimates are computed as follows.

$$\hat{\phi} = \frac{1}{m} \sum_{i=1}^m \mathbf{1}(y^{(i)} = 1)$$

For  $j = 1, 2, \dots, r$ ,

$$\begin{aligned}
 (\hat{\mu}_0)_j &= \frac{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 0)(\mathbf{x}^{(i)})_j}{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 0)} \\
 (\hat{\mu}_1)_j &= \frac{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 1)(\mathbf{x}^{(i)})_j}{\sum_{i=1}^m \mathbf{1}(y^{(i)} = 1)} \\
 \hat{\sigma}_j^2 &= \frac{1}{m} \sum_{i=1}^m (\mathbf{x}^{(i)} - \mu_{y^{(i)}})_j^2, \quad j = 1, 2, \dots, n
 \end{aligned}$$

For  $j = r + 1, r + 2, \dots, n$ ,

$$\hat{\psi}_{j,0,k} =$$

counts the number of times  $y^{(i)} = 0$  and  $\mathbf{x}_j^{(i)} = k$  in the training data and divides this by the number of times  $y^{(i)} = 0$  (number of points from class zero in the training data). Do this for every value of  $k = 1, 2, \dots, K_j$  and for every feature  $\mathbf{x}_j, j = 1, 2, \dots, n$ .

## 9 Supervised Learning: Linear Classifiers and Support Vector Machines (SVMs)

### 9.1 Linear Classifiers \*\*

Both logistic regression and GDA with Naive Bayes and equal covariances result in a linear classifier, i.e., one can simplify the classification rule in both cases to get the following:

$$\hat{y} = 1 \quad \text{if } \mathbf{w}^T \mathbf{x} + b > 0$$

and equals 0 otherwise.

Proof for GDA:

GDA decision rule is:  $\hat{y} = 1$  if

$$(\mathbf{x} - \mu_1)^T \Sigma_1^{-1} (\mathbf{x} - \mu_1) < (\mathbf{x} - \mu_0)^T \Sigma_0^{-1} (\mathbf{x} - \mu_0)$$

With  $\Sigma_1 = \Sigma_0 = \Sigma$  and naive Bayes ( $\Sigma$  is diagonal), this simplifies to checking if

$$\sum_j \frac{(\mathbf{x} - \mu_1)_j^2}{\sigma_j^2} < \sum_j \frac{(\mathbf{x} - \mu_0)_j^2}{\sigma_j^2}$$

Simplifying the above, we equivalently need

$$\mathbf{w}^T \mathbf{x} + b > 0, \quad \text{with } \mathbf{w}_j = \frac{(\mu_1 - \mu_0)_j^2}{\sigma_j^2}, \quad b = 0.5 \sum_j \frac{((\mu_1 - \mu_0)_j)(\mu_1 + \mu_0)_j}{\sigma_j^2}$$

In fact we can also get an expression for  $\mathbf{w}$  and  $b$  even if we just have  $\Sigma_1 = \Sigma_0$ . We will get  $\mathbf{w} = \Sigma^{-1}(\mu_1 - \mu_0)$  and  $b = \mu_1^T \Sigma^{-1} \mu_1 - \mu_0^T \Sigma^{-1} \mu_0$ . ?? check.

### 9.2 Support Vector Machines (SVMs): Motivation and main idea

Both the classifiers we talked about so far are linear classifiers as explained above. Consider logistic regression. The classification decision would be much more reliable if  $\theta^T \mathbf{x}$  were either much larger or much smaller than zero.

The idea of SVM is this: we do not assume any data model here. Instead we try to look for the “separating hyperplane”, equivalently, a vector  $\theta$  so that the “margin” from the decision boundary is maximized for *all* training data points. Visual explanation in class or see cs229-notes-3.

### 9.2.1 Notation change

Instead of a single vector  $\theta$  with the first entry used for the bias term, in case of SVMs, we use a weight vector  $\mathbf{w}$  which is the same length as the data and a scalar  $b$ . Also, instead of labeling the two classes as 0 and 1, we label them as  $-1$  and  $+1$  because this simplifies some of the writing.

Margin: the distance of a data point from the separating hyperplane. Margin for the  $i$ -th training data point is computed as

$$y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b)$$

### 9.2.2 Using SVM for classification

Suppose first that the optimal choice of  $\mathbf{w}, b$  is available. Then we classify as follows

$$\hat{y} = \text{sign}(\mathbf{w}^T \mathbf{x} + b)$$

If the term  $\hat{y} > 0$ , then the class is  $+1$  else the class is  $-1$ .

### 9.2.3 Goal

The goal in case of SVMs is to find  $\mathbf{w}, b$  that maximize the worst-case margin defined by

$$\min_{i=1,2,\dots,m} y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b)$$

By multiplying  $\mathbf{w}, b$  by a scalar we could keep increasing the margin, but that will not improve classification. Thus, we need to impose the constraint that  $\|\mathbf{w}\|_2 = 1$ .

## 9.3 Simplifications to obtain a convex optimization problem \*\*

Writing a slightly different way, we need to solve

$$\max_{\mathbf{w}, b} \gamma \text{ s.t. } y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b) \geq \gamma, \quad i = 1, 2, \dots, m \text{ and } \|\mathbf{w}\|_2 = 1$$

This is not a convex optimization problem yet. So we try to simplify further. The above is equivalent to dividing everything by  $\|\mathbf{w}\|_2$ . Doing this gives

$$\max_{\mathbf{w}, b} \frac{\gamma}{\|\mathbf{w}\|_2} \text{ subject to } y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b) \geq \gamma, \quad i = 1, 2, \dots, m$$

This is still not convex. Another point to notice is this: we could pick  $\|\mathbf{w}\|_2$  to be anything, so equivalently, we could fix the value of the margin  $\gamma$  to 1, and nothing will change. This gives

$$\max_{\mathbf{w}, b} \frac{1}{\|\mathbf{w}\|_2} \text{ subject to } y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b) \geq 1, \quad i = 1, 2, \dots, m$$

This is not convex either but now a simple reformulation gives us a convex problem. Maximizing  $1/\|\mathbf{w}\|_2$  is equivalent to minimizing  $\|\mathbf{w}\|_2$  which is the same as minimizing  $\|\mathbf{w}\|_2^2$ . This gives

$$\min_{\mathbf{w}, b} \|\mathbf{w}\|_2^2 \text{ subject to } y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b) \geq 1, \quad i = 1, 2, \dots, m$$

This is now a convex optimization problem since the cost is a convex function and the inequality constraints are linear. In particular it is what is called a Quadratic Program or QP.

## 9.4 Final primal problem

$$\min_{\mathbf{w}, b} \|\mathbf{w}\|_2^2 \text{ subject to } y^{(i)}(\mathbf{w}^T \mathbf{x}^{(i)} + b) \geq 1, \quad i = 1, 2, \dots, m$$

This is now a convex optimization problem since the cost is a convex function and the inequality constraints are linear. In particular it is what is called a Quadratic Program or QP.

## 9.5 Simplification using duality: needed to develop Kernel SVMs

By using *Lagrange duality*, it is possible to show that we can also compute the optimal  $\mathbf{w}, b$  as follows.

1. Solve the following “dual problem”: optimize over the Lagrange multipliers  $\alpha_i$

$$\max_{\alpha} \left( \sum_{i=1}^m \alpha_i - 0.5 \sum_{i=1}^m \sum_{j=1}^m \alpha_i \alpha_j y^{(i)} y^{(j)} \langle \mathbf{x}^{(i)}, \mathbf{x}^{(j)} \rangle \right) \text{ s.t. } \sum_{i=1}^m \alpha_i y^{(i)} = 0, \alpha_i \geq 0, i = 1, 2, \dots, m$$

(here  $\langle \mathbf{x}_1, \mathbf{x}_2 \rangle = \mathbf{x}_1^T \mathbf{x}_2$ )

2. Obtain

$$\hat{\mathbf{w}} = \sum_i \hat{\alpha}_i y^{(i)} \mathbf{x}^{(i)}$$

$$\hat{b} = -0.5 \left( \max_{i: y^{(i)} = -1} \hat{\mathbf{w}}^T \mathbf{x}^{(i)} + \min_{i: y^{(i)} = 1} \hat{\mathbf{w}}^T \mathbf{x}^{(i)} \right)$$

Notice the dependence on feature vectors is only through inner products

Classification:

1. Classification also only uses inner products

$$\hat{\mathbf{w}}^T \mathbf{x} + \hat{b} = \sum_i \alpha_i y^{(i)} \langle \mathbf{x}^{(i)}, \mathbf{x} \rangle + \hat{b}$$

The above is useful in two settings: (i) if  $n \gg m$ , i.e. the original data lies in a much higher dimensional space compared to the available number of data points (this often happens in classification problems), then the above dual is much less expensive to solve.

(ii) Notice that in the above problem, all dependence on the feature vectors  $\mathbf{x}^{(i)}$  is through the inner product  $\langle \mathbf{x}^{(i)}, \mathbf{x}^{(j)} \rangle$ . This means we could consider higher dimensional feature vectors, i.e., convert  $\mathbf{x}$  to  $\phi(\mathbf{x})$  by considering features of the form  $\mathbf{x}_j^2, \mathbf{x}_j \mathbf{x}_k$ .

For many such high dimensional feature mappings,  $\phi(\mathbf{x}^{(i)})$  is very expensive to compute for all training data  $i = 1, 2, \dots, m$ . But just computing the inner product, defined as the “kernel product”

$$K(\mathbf{x}_1, \mathbf{x}_2) := \langle \phi(\mathbf{x}_1), \phi(\mathbf{x}_2) \rangle$$

is much less expensive. An example is  $\phi(x)$  obtained as all pairwise products of entries of  $\mathbf{x}$ . For this, computing  $\phi(\mathbf{x}^{(i)})$  takes order  $n^2$  time, thus computing the new feature vector for all  $i$  takes order  $n^2 m$  time. But computing one kernel product can be done as

$$K(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{x}_1^T \mathbf{x}_2)^2$$

This only takes order  $n$  time. There are  $m(m+1)/2$  total products to compute. Thus, the time needed is of order  $nm^2$ . Since typically,  $m \ll n$ , this is much quicker.

(iii) It is also possible to define kernel products for cases where the actual feature mapping is infinite dimensional. Gaussian kernel is an example.

## 9.6 Kernel SVM \*\*

A very large number of kernels can be defined. The purpose is for datasets which are not linearly separable in the original feature space, it is possible they are in a higher dimensional space.

Kernel product: is basically some measure of similarity between two data points. Any useful measure of similarity can be used to define a “kernel” (kernel product), one may not even need to specify the underlying feature mapping.

This “kernel trick” can be used for many other learning algorithms as well. Anytime all computation depends on inner products between the features, this can be used.

## 9.7 Soft margin SVMs

See page 19 of ML-cs229-noted-3

## 10 Deviation: Introduction to Lagrange duality to understand how to derive the dual program \*\*

To be updated later (to make this course 425/525).

See ML-cs229-notes3

## 11 Least Squares (LS): all kinds

### 11.1 Regularized LS

The opt problem that needs to be solved to learn the parameter vector  $\theta$  for Lin Reg is called the “Least Squares” problem. This involves solving

$$\min_{\theta} \|\mathbf{y} - \mathbf{X}\theta\|^2$$

The above is called *Ordinary Least Squares* or *Ordinary Linear Regression*. This problem also occurs in various other domains. Many examples in communications and in signal processing. Thus it is useful to understand its various modifications.

We may claim that no house feature is way too important, and thus, no entry of  $\theta$  should be too large. In this case we solve

$$\min_{\theta} \|\mathbf{y} - \mathbf{X}\theta\|^2 + \lambda\|\theta\|^2$$

The above is also called *Ridge Regression* or *ell-2 regularization*

ell-1 regularized LS: this ensures the vector  $\theta$  is sparse (has many zero entries), larger  $\lambda$  makes it sparser. Use if only a subset of all the features in the feature vector matter but do not know which ones.

$$\min_{\theta} \|\mathbf{y} - \mathbf{X}\theta\|^2 + \lambda\|\theta\|_1$$

The above is also called *Lasso* or *ell-1 regularization*.

### 11.2 Use of prior knowledge and Recursive LS

In many settings, we may have prior knowledge that  $\theta$  should be close to  $\theta_0$ . For example, we may have  $\theta_0$  from houses sold a year ago and their features. In such cases, we solve

$$\min_{\theta} \|\mathbf{y} - \mathbf{X}\theta\|^2 + \lambda\|\theta - \theta_0\|^2$$

Recursive LS: consider house price example again and suppose we want to update our model each week. Instead of starting from scratch, can we just “update” the model: save time/computation?

Turns out we can using recursive LS: suppose we update with each new house sold, then we get this

$$\begin{aligned}\hat{\theta}_0 &= \theta_0, P_0 = \lambda\mathbf{I}, \\ K_t &= P_{t-1}\mathbf{x}^{(i)}(\lambda + \mathbf{x}^{(i)\top}P_{t-1}\mathbf{x}^{(i)})^{-1}, \\ P_t &= (\mathbf{I} - K_t\mathbf{x}^{(i)\top})P_{t-1}, \\ \hat{\theta}_t &= \hat{\theta}_{t-1} + K_t(\mathbf{y}^{(i)} - \mathbf{x}^{(i)\top}\hat{\theta}_{t-1}),\end{aligned}$$

Suppose we update after a set of  $m_t$  new houses sold (so  $\mathbf{x}^{(i)}$  is an  $n \times m_t$  matrix and  $\mathbf{y}^{(i)}$  is a  $m_t \times 1$  vector), then we get this

$$\hat{\theta}_0 = \theta_0, P_0 = \mathbf{I},$$



$$\begin{aligned}
K_t &= P_{t-1} \mathbf{x}^{(i)} (\lambda \mathbf{I} + \mathbf{x}^{(i)\top} P_{t-1} \mathbf{x}^{(i)})^{-1}, \\
P_t &= (\mathbf{I} - K_t \mathbf{x}^{(i)\top}) P_{t-1}, \\
\hat{\theta}_t &= \hat{\theta}_{t-1} + K_t (\mathbf{y}^{(i)} - \mathbf{x}^{(i)\top} \hat{\theta}_{t-1}),
\end{aligned}$$

## 12 Unsupervised Learning: PCA

In unsupervised learning, there is no *labeled* training data to learn parameters from. So no “output”  $y^{(i)}$  is available for “input”  $\mathbf{x}^{(i)}$ . PCA is an unsupervised learning technique that is used for dimension reduction. Given data vectors in  $\mathfrak{R}^n$ , if they approximately lie in a lower dimensional subspace, how do we find that subspace?

### 12.1 What is PCA

As before we assume that we are given  $m$  data vectors (usually called feature vectors)  $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}, \dots, \mathbf{x}^{(m)}$  each in  $\mathfrak{R}^n$ , and these are stacked as the rows of a matrix

$$\mathbf{X} = \begin{bmatrix} -(\mathbf{x}^{(1)})^T - \\ -(\mathbf{x}^{(2)})^T - \\ \vdots \\ -(\mathbf{x}^{(m)})^T - \end{bmatrix} \quad (2)$$

Given a reduced dimension  $r$ , the goal is find an  $r$  dimensional representation of the data vectors so that maximum information (variance) is retained.

### 12.2 Why PCA

Notice each data vector is  $n$ -length.  $n$  may be very large can, e.g., can be equal to the image size if one uses all the pixels as “features”. So we try to reduce the dimension to  $r < n$ . This is helpful in the following ways:

- Less storage needed to save the data.
- Analysis of the reduced-dimension dataset, e.g., regression or classification, will be faster;
- If the original data is noisy, the data analysis task, e.g., regression or classification, can also be more accurate in the reduced dimensional space. This point is discussed more in the Bias-Variance tradeoff and Learnin Theory section. When you reduce the dimension from  $n$  to an  $r < n$ , you increase the bias in your model. But you decrease the parameter estimation error variance. The error variance is proportional to the “noise-level” in the data, so  $\sigma_e^2$  in case of linear regression.
  - You may notice this for certain real datasets with PCA implemented carefully. For simulated data, to observe this, you will have to deliberately generate noisy simulated data.
  - In some other cases, there may be zero variance along some directions and when the data covariance matrix is computed, it ends up being rank deficient. Hence it cannot be inverted, but inversion is needed for example, for GDA based classification. This is actually the perfect application where PCA resolve the problem and may in fact help improve classification accuracy.
- Finally: PCA helps to get subspace coefficients which are uncorrelated. This is needed in some applications and is convenient in others. For example, for GDA things means that the covariance matrix of the new data vectors is diagonal.
  - For this, we need not reduce the dimension though, we can obtain uncorrelated variables in  $n$  dimensions too by using the full SVD.

## 12.3 Practical information: How to implement PCA

### 12.3.1 Reduced dimension, $r$ , is specified

**INPUT:**  $m \times n$  feature vectors' matrix  $\mathbf{X}$ , chosen rank  $r$ .

1. Compute  $\hat{\boldsymbol{\mu}} = \frac{1}{m} \sum_i \mathbf{x}^{(i)}$
2. For  $i = 1, 2, \dots, m$ , let  $\mathbf{z}^{(i)} = \mathbf{x}^{(i)} - \hat{\boldsymbol{\mu}}$  and let

$$\mathbf{Z} = \begin{bmatrix} -(\mathbf{z}^{(1)})^T - \\ -(\mathbf{z}^{(2)})^T - \\ \vdots \\ -(\mathbf{z}^{(m)})^T - \end{bmatrix} \quad (3)$$

OR: compute directly

$$\mathbf{Z} = \mathbf{X} - \mathbf{1}_m (\mathbf{1}_m^T \mathbf{X} / m)$$

where  $\mathbf{1}_m = [1 \ 1 \ \dots \ 1]^T$  is an  $m$ -length vector of ones.

3. Compute the singular value decomposition (SVD) of  $\mathbf{Z}$  and set  $\mathbf{V}$  equal to the  $r$  right singular vectors with the largest singular values (called “top  $r$ ” right singular vectors), i.e., compute

$$\mathbf{Z} \stackrel{SVD}{=} \mathbf{U}_{full} \mathbf{S}_{full} \mathbf{V}_{full}^T$$

where  $\mathbf{U}_{full}, \mathbf{V}_{full}$  are unitary matrices and  $\mathbf{S}_{full}$  is an  $m \times n$  diagonal matrix with non-negative entries (singular values) arranged in decreasing order of magnitude.

Set  $\mathbf{V} = \mathbf{V}_{full}(:, 1:r)$  in MATLAB notation (set  $\mathbf{V}$  equal to first  $r$  columns of  $\mathbf{V}_{full}$ ).

Thus  $\mathbf{V}$  is the  $n \times r$  matrix whose columns span the computed *principal subspace*.

- We can also obtain  $\mathbf{V}$  as the eigenvectors with the  $r$  largest eigenvalues (called top  $r$  eigenvectors) of  $\mathbf{Z}^T \mathbf{Z}$ .
4. Project the original data  $\mathbf{X}$  to  $\text{range}(\mathbf{V})$ : compute  $\mathbf{b}^{(i)} = \mathbf{V}^T \mathbf{x}^{(i)}$  for each  $i$  or equivalently, compute

$$\mathbf{B} = \mathbf{XV}$$

Thus  $\mathbf{B}$  is  $m \times r$ . These are the new feature vectors in the reduced dimensional space.

**OUTPUT:**  $m \times r$  matrix  $\mathbf{B}$ : reduced dim feature vectors;  $n \times r$  matrix  $\mathbf{V}$ : principal subspace.

- If the rank  $r$  approximation of  $\mathbf{X}$  is needed, this is obtained as the  $m \times n$  matrix

$$\mathbf{L} = \mathbf{BV}^T = \mathbf{XVV}^T$$

### 12.3.2 Use of PCA

PCA is a dimension reduction technique.

It is one way to do feature selection if the assumption that “features having the most variation are the most important ones” is true. In addition, of course PCA helps de-correlate the features.

A typical setting where this is useful is when the raw feature vector is all image pixels. There are too many pixels and pixels are highly correlated.

*PCA is not very useful if features are first variance normalized and then PCA is applied.*

### 12.3.3 Deciding $r$

So far we have assumed that the desired lower-dimension  $r$  is given. However, there is no one correct way of deciding  $r$  in practice. Two common approaches:

1. A common heuristic is the 99% or some other high-enough percent heuristic: retain all eigenvectors so the that variance in the reduced dimensional space is at least 99% of the total variance of the data. In other words, find the smallest value  $r$  so that

$$\sum_{i=1}^r \sigma_i^2 \geq 0.99 \sum_{i=1}^{\min(m,n)} \sigma_i^2.$$

There is nothing “special” about 99%, we could also use another percentage.

2. An alternative approach is to pick the  $r$  that is best for the final application for which PCA is being used as a pre-processing step. For example, if PCA is applied to reduce the dimension of the input training data used for linear regression, then we use cross-validation to compute the value of  $r$  that minimizes the test-MSE. Ditto if PCA is being used before GDA or Logistic Regression.

Details for the second approach. Consider Linear Regression. Given  $\mathbf{X}$ ,  $\mathbf{y}$  and  $\mathbf{X}_{test}$ ,  $\mathbf{y}_{test}$  (split all available data into 10% testing and 90% training). OR BETTER: use cross validation (most efficient but more difficult to explain here).

1. Compute  $\mathbf{Z} = \mathbf{X} - (\mathbf{1}_m \mathbf{1}_m^T \mathbf{X} / m)$
2. Compute SVD of  $\mathbf{Z}$

$$\mathbf{Z} \stackrel{SVD}{=} \mathbf{U}_{full} \mathbf{S}_{full} \mathbf{V}_{full}^T$$

where  $\mathbf{U}_{full}$ ,  $\mathbf{V}_{full}$  are unitary matrices and  $\mathbf{S}_{full}$  is an  $m \times n$  diagonal matrix with non-negative entries (singular values) arranged in decreasing order of magnitude.

3. Loop over  $r$  from 1 to  $n$

- (a) Set  $\mathbf{V} = \mathbf{V}_{full}(:, 1 : r)$  in MATLAB notation (set  $\mathbf{V}$  equal to first  $r$  columns of  $\mathbf{V}_{full}$ ).
- (b) Compute  $\mathbf{B} = \mathbf{XV}$
- (c) Compute  $\hat{\theta}$  using  $\mathbf{B}$  (instead of  $\mathbf{X}$ ) and make sure to learn the bias term as well as explained earlier. Let  $\tilde{\mathbf{B}} = [\mathbf{1}_m, \mathbf{B}]$ . Compute

$$\hat{\theta} = (\tilde{\mathbf{B}}^T \tilde{\mathbf{B}})^{-1} \tilde{\mathbf{B}}^T \mathbf{y}$$

- (d) Compute test error
  - i. Compute  $\mathbf{B}_{test} = \mathbf{X}_{test} \mathbf{V}$
  - ii. Obtain the vector of predictions for test data,  $\hat{\mathbf{y}}_{test} = \mathbf{B}_{test} \hat{\theta}$
  - iii. Compute  $NormalizedTestMSE(r) = \|\mathbf{y}_{test} - \hat{\mathbf{y}}_{test}\|_2^2 / \|\mathbf{y}_{test}\|_2^2$  on the test data as explained before

*I think the algebra in the above is correct, but please verify to be sure.*

4. Pick  $r$  for which  $NormalizedTestMSE(r)$  is the smallest.

## 12.4 Theory – computational - 1

PCA solves

$$\min_L \|\mathbf{X} - \mathbf{L}\|_F$$

Set  $\mathbf{L} = \mathbf{B}\mathbf{V}^\top$  and simplify, then this becomes

$$\min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \|\mathbf{X} - \mathbf{B}\mathbf{V}^\top\|_F^2 = \min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \sum_i \|\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)}\|^2$$

Now we can solve for  $\mathbf{b}^{(i)}$  in closed form:  $\mathbf{b}^{(i)} = \mathbf{V}^\top \mathbf{x}^{(i)}$  and substitute it in. Thus we get

$$\min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \|\mathbf{X} - \mathbf{B}\mathbf{V}^\top\|_F^2 = \min_{\mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \sum_i \|(\mathbf{I} - \mathbf{V}\mathbf{V}^\top)\mathbf{x}^{(i)}\|^2$$

Since

$$\|(\mathbf{I} - \mathbf{V}\mathbf{V}^\top)\mathbf{x}\|^2 = \|\mathbf{x}\|^2 + \|\mathbf{V}^\top \mathbf{x}\|^2 - 2\|\mathbf{V}^\top \mathbf{x}\|^2 = \|\mathbf{x}\|^2 - \|\mathbf{V}^\top \mathbf{x}\|^2$$

the min simplifies to

$$\max_{\mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \sum_i \|\mathbf{V}^\top \mathbf{x}^{(i)}\|^2$$

Since

$$\sum_i \|\mathbf{V}^\top \mathbf{x}^{(i)}\|^2 = \sum_i (\mathbf{x}^{(i)\top} \mathbf{V}\mathbf{V}^\top \mathbf{x}^{(i)}) = \text{trace}(\sum_i \mathbf{x}^{(i)\top} \mathbf{V}\mathbf{V}^\top \mathbf{x}^{(i)}) = \text{trace}(\mathbf{V}^\top (\sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)\top}) \mathbf{V})$$

thus this further simplifies to

$$\max_{\mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \text{trace}(\mathbf{V}^\top (\sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)\top}) \mathbf{V})$$

Let EVD of  $\hat{\Sigma} = (\sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)\top})/m = \mathbf{V}\Lambda\mathbf{V}^\top$  with  $\mathbf{U}$  orthonormal. Suppose  $r = 2$ . This then simplifies to

$$\max_{\mathbf{v}_2 \perp \mathbf{v}_1, \|\mathbf{v}_2\|=1} [\max_{\mathbf{v}_1: \|\mathbf{v}_1\|=1} \mathbf{v}_1^\top (\sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)\top}) \mathbf{v}_1 + \mathbf{v}_2^\top (\sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)\top}) \mathbf{v}_2]$$

The first term is maximized by the top eigenvector,  $\mathbf{v}_1$  of  $\hat{\Sigma} = (\sum_i \mathbf{x}^{(i)} \mathbf{x}^{(i)\top})/m$  and its maximum value equals the top eigenvalue,  $\lambda_1$ . Thus,

$$\max_{\mathbf{V}: \mathbf{V}^\top \mathbf{V} = \mathbf{I}} \text{trace}(\mathbf{V}^\top \hat{\Sigma} \mathbf{V}) = \lambda_1 + \max_{\mathbf{v}_2 \perp \mathbf{v}_1, \|\mathbf{v}_2\|=1} \mathbf{v}_2^\top \mathbf{U}\Lambda\mathbf{U}^\top \mathbf{v}_2$$

Thus, this is solved by  $\mathbf{v}_2 =$  the second eigenvector of  $\hat{\Sigma}$ . Proceeding similarly

$$\hat{\mathbf{V}} = \mathbf{V}_{1:r}$$

solves the maximization.

Eigenvectors of  $\hat{\Sigma}$  are right singular vectors of  $\mathbf{X}$ .

Then setting  $\hat{\mathbf{B}} = \mathbf{X}\hat{\mathbf{V}}^\top$  proves that

$$\hat{\mathbf{L}} = \mathbf{X}\mathbf{V}_{1:r}\mathbf{V}_{1:r}^\top = \mathbf{U}_{1:r}\mathbf{S}_{r \times r}\mathbf{V}_{1:r}^\top$$

( $r$ -SVD of  $\mathbf{X}$ ) solves

$$\min_L \|\mathbf{X} - \mathbf{L}\|_F$$

## 12.5 Theory – computational 2

PCA also solves

$$\min_{\mathbf{L} \text{ rank } r} \|\mathbf{X} - \mathbf{L}\|_2^2 = \min_{\mathbf{L} \text{ rank } r} \lambda_{\max}((\mathbf{X} - \mathbf{L})^T(\mathbf{X} - \mathbf{L}))$$

By Weyl-type inequality for singular values,  $\sigma_i(M) \leq \sigma_i(M_2) + \|M - M_2\|$  for any singular value  $i$  and any two matrices  $M, M_2$ . Thus,

$$\|\mathbf{X} - \mathbf{L}\|_2 \geq \sigma_i(\mathbf{X}) - \sigma_i(\mathbf{L})$$

Since  $\mathbf{L}$  is rank  $r$ , it has only  $r$  nonzero singular values. Thus  $\sigma_{r+1}(\mathbf{L}) = 0$ . And so

$$\|\mathbf{X} - \mathbf{L}\|_2 \geq \sigma_{r+1}(\mathbf{X}).$$

Since this is true for any rank  $r$  matrix  $\mathbf{L}$ , it is also true for the minimizer, i.e.

$$\min_{\mathbf{L}: \text{rank } r} \|\mathbf{X} - \mathbf{L}\|_2 \geq \sigma_{r+1}(\mathbf{X})$$

Now if we can find a specific matrix  $\mathbf{L}$  for which  $\|\mathbf{X} - \mathbf{L}\|_2 = \sigma_{r+1}(\mathbf{X})$  that will be the minimizer (since the minimum value cannot be any smaller than this).

Set

$$\hat{\mathbf{L}} = \sum_{i=1}^r \sigma_i u_i v_i^T$$

( $r$ -SVD of  $\mathbf{X}$ ), then  $\mathbf{X} - \hat{\mathbf{L}} = \sum_{i=r+1}^{\min(m,n)} \sigma_i u_i v_i^T$  and so

$$\|\mathbf{X} - \hat{\mathbf{L}}\|_2 = \sigma_{r+1}$$

here  $\sigma_i = \sigma_i(\mathbf{X})$ .

## 12.6 Statistical optimality

### 12.6.1 Statistical Optimality

First assume everything is zero mean.

1. PCA finds the subspace  $\mathbf{V}$  ( $\mathbf{V}$  is a matrix with orthonormal columns that define the subspace) and the projected random vector  $\mathbf{b}$  so that the expected value of the squared 2-norm of the reconstruction error  $\mathbf{x} - \mathbf{V}\mathbf{b}$  is minimized. Thus, it solves

$$\min_{\mathbf{b} \in \mathbb{R}^r, \mathbf{V} \in \mathbb{R}^{n \times r}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|\mathbf{x} - \mathbf{V}\mathbf{b}\|_2^2]$$

If we minimize over  $\mathbf{b}$  first as a function of  $\mathbf{V}$ , then this is a standard Least Squares problem whose solution is

$$\mathbf{b} = (\mathbf{V}^T \mathbf{V})^{-1} \mathbf{V}^T \mathbf{x} = \mathbf{V}^T \mathbf{x} \quad \text{since } \mathbf{V}^T \mathbf{V} = \mathbf{I} \text{ in this case}$$

Thus, we need to solve

$$\min_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|\mathbf{x} - \mathbf{V}\mathbf{V}^T \mathbf{x}\|_2^2]$$

Since  $\mathbf{V}^T \mathbf{V} = \mathbf{I}$ , thus  $\|\mathbf{V}\mathbf{V}^T \mathbf{x}\|_2^2 = \|\mathbf{V}^T \mathbf{x}\|_2^2$  and so  $\|\mathbf{x} - \mathbf{V}\mathbf{V}^T \mathbf{x}\|_2^2 = \|\mathbf{x}\|_2^2 - \|\mathbf{V}^T \mathbf{x}\|_2^2$ . With this, the above is also equivalent to

$$\max_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|\mathbf{V}^T \mathbf{x}\|_2^2]$$

Using property of trace, this is further equivalent to

$$\max_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \text{trace}(\mathbf{V}^T \mathbb{E}[\mathbf{x}\mathbf{x}^T] \mathbf{V})$$

This is another commonly use definition for PCA: PCA finds the  $r$  directions of largest variance of the data. Here  $\mathbb{E}[\mathbf{x}\mathbf{x}^T]$  is the covariance matrix.

2. PCA can also be understood as minimizing the worst-case (largest) expected reconstruction error in any direction: for direction  $\mathbf{w}$ , this is  $|\mathbf{w}^T(\mathbf{x} - \mathbf{V}\mathbf{b})|$

$$\min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \max_{\mathbf{w}: \|\mathbf{w}\|_2 = 1} \mathbb{E}[(\mathbf{w}^T(\mathbf{x} - \mathbf{V}\mathbf{b}))^2]$$

Using the properties of trace, this is equivalent to

$$\min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \max_{\mathbf{w}: \|\mathbf{w}\|_2 = 1} \mathbf{w}^T \mathbb{E}[(\mathbf{x} - \mathbf{V}\mathbf{b})(\mathbf{x} - \mathbf{V}\mathbf{b})^T] \mathbf{w} = \min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \lambda_{\max}(\mathbb{E}[(\mathbf{x} - \mathbf{V}\mathbf{b})(\mathbf{x} - \mathbf{V}\mathbf{b})^T])$$

The last equality follows using the variational definition of the maximum eigenvalue.

In the nonzero mean case, we do either of the above for  $(\mathbf{x} - \mu)$ .

## 12.7 Delete \*\*

*Claim: All of the below optimization problems are solved by the PCA procedure above*

### 12.7.1 Optimality from a computational viewpoint

This assumes that we are given a data matrix  $\mathbf{X}$  (or  $\mathbf{Z}$  after subtracting the empirically computed mean). The goal is to find a rank  $r$  approximation to  $\mathbf{X}$ , denoted by  $\mathbf{L}$ , that minimizes either the Frobenius norm of the error or the induced 2-norm of the error. Thus PCA on  $\mathbf{X}$  solves

1.

$$\min_{\mathbf{L} \text{ rank } r} \|\mathbf{X} - \mathbf{L}\|_F^2$$

Any rank  $r$  matrix  $\mathbf{L}$  can be expressed as  $\mathbf{L} = \mathbf{B}\mathbf{V}^T$  where  $\mathbf{V}$  is an  $n \times r$  matrix with orthonormal columns ( $\mathbf{V}^T\mathbf{V} = \mathbf{I}$ ) and  $\mathbf{B}$  is  $m \times r$ . Thus, above is equivalent to

$$\min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^T\mathbf{V}=\mathbf{I}} \|\mathbf{X} - \mathbf{B}\mathbf{V}^T\|_F^2 = \min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^T\mathbf{V}=\mathbf{I}} \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)}\|_2^2$$

The last equality follows by expressing  $\mathbf{X}$  and  $\mathbf{B}$  in terms of their rows.

With the above, the minimizations over the different  $\mathbf{b}^{(i)}$ s (as a function of  $\mathbf{V}$ ) are decoupled and each is a simple LS problem that is solved by  $\mathbf{b}^{(i)} = (\mathbf{V}^T\mathbf{V})^{-1}\mathbf{V}^T\mathbf{x}^{(i)} = \mathbf{V}^T\mathbf{x}^{(i)}$ . Substituting this for  $\mathbf{b}^{(i)}$ , we need to solve

$$\min_{\mathbf{V}: \mathbf{V}^T\mathbf{V}=\mathbf{I}} \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mathbf{V}\mathbf{V}^T\mathbf{x}^{(i)}\|_2^2$$

2. It also solves

$$\min_{\mathbf{L} \text{ rank } r} \|\mathbf{X} - \mathbf{L}\|_2^2 = \min_{\mathbf{L} \text{ rank } r} \lambda_{\max}((\mathbf{X} - \mathbf{L})^T(\mathbf{X} - \mathbf{L}))$$

Once again writing  $\mathbf{L} = \mathbf{B}\mathbf{V}^T$ , the above is equivalent to

$$\min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^T\mathbf{V}=\mathbf{I}} \lambda_{\max}((\mathbf{X} - \mathbf{B}\mathbf{V}^T)^T(\mathbf{X} - \mathbf{B}\mathbf{V}^T))$$

Here again, expressing  $\mathbf{X}$  and  $\mathbf{B}$  in terms of their rows,  $(\mathbf{X} - \mathbf{B}\mathbf{V}^T)^T(\mathbf{X} - \mathbf{B}\mathbf{V}^T) = \sum_{i=1}^m (\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})(\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})^T$  and so, the above is further equivalent to

$$\min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^T\mathbf{V}=\mathbf{I}} \lambda_{\max}\left(\sum_{i=1}^m (\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})(\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})^T\right)$$

## 12.8 Proofs of the above claims \*\*

First we prove that the statistical optimality claims are equivalent to the computational claims given above if we use an empirical approximation to the expected values. Next we prove that the  $r$ -SVD solution solves each of the two computational claims above.

### 12.8.1 Equivalence between statistical and computational claims 1

Using the empirical approximation

$$\mathbb{E}[\|\mathbf{x} - \mathbf{V}\mathbf{V}^T\mathbf{x}\|^2] \approx \frac{1}{m} \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mathbf{V}\mathbf{V}^T\mathbf{x}^{(i)}\|_2^2$$

and thus the first claims from the statistical and computational viewpoints are equivalent, i.e.

$$\min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|\mathbf{x} - \mathbf{V}\mathbf{b}\|_2^2] = \min_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|\mathbf{x} - \mathbf{V}\mathbf{V}^T \mathbf{x}\|_2^2]$$

is equivalent to

$$\min_{\mathbf{L}: \text{rank } r} \|\mathbf{X} - \mathbf{L}\|_F^2 = \min_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mathbf{V}\mathbf{V}^T \mathbf{x}^{(i)}\|_2^2$$

### 12.8.2 Proof of computational claim 1: Frobenius norm optimality

Goal is to find a rank  $r$  matrix  $\mathbf{L}$  that solves

$$\min_{\mathbf{L}: \text{rank } r} \|\mathbf{X} - \mathbf{L}\|_F^2$$

As explained above, this is equivalent to  $\min_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mathbf{V}\mathbf{V}^T \mathbf{x}^{(i)}\|_2^2$ . This is further equivalent to

$$\max_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \sum_{i=1}^m \|\mathbf{V}^T \mathbf{x}^{(i)}\|_2^2$$

Using  $\sum_{i=1}^m \|\mathbf{V}^T \mathbf{x}^{(i)}\|_2^2 = \|\mathbf{X}\mathbf{V}\|_F^2 = \sum_{j=1}^r \|\mathbf{X}\mathbf{v}_j\|_2^2 = \sum_{j=1}^r \mathbf{v}_j^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_j$ , the above is equivalent to

$$\max_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \sum_{j=1}^r \mathbf{v}_j^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_j$$

If we did not have the constraint  $\mathbf{V}^T \mathbf{V} = \mathbf{I}$  we now have  $r$  decoupled problems to solve. The constraint itself can be rewritten as

$$\{\mathbf{V} : \mathbf{V}^T \mathbf{V} = \mathbf{I}\} = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r \text{ unit 2-norm and } \mathbf{v}_2 \perp \mathbf{v}_1, \mathbf{v}_3 \perp \text{span}(\mathbf{v}_1, \mathbf{v}_2), \text{ and so on}\}$$

To write things simply, suppose  $r = 3$  so that  $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3]$ . Then,

$$\begin{aligned} & \max_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \sum_{j=1}^r \sum_{i=1}^m \mathbf{v}_j^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_j \\ &= \max_{\mathbf{v}_3 \perp \text{span}(\mathbf{v}_1, \mathbf{v}_2): \|\mathbf{v}_3\|_2=1} \left( \max_{\mathbf{v}_2 \perp \mathbf{v}_1: \|\mathbf{v}_2\|_2=1} \left( \left( \max_{\mathbf{v}_1: \|\mathbf{v}_1\|_2=1} \sum_{i=1}^m \mathbf{v}_1^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_1 \right) + \mathbf{v}_2^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_2 \right) + \mathbf{v}_3^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_3 \right) \end{aligned}$$

Consider the max over  $\mathbf{v}_1$ . By definition of the first eigenvector (eigenvector with largest eigenvalue), the maximizer of  $\max_{\mathbf{v}_1: \|\mathbf{v}_1\|_2=1} \mathbf{v}_1^T (\mathbf{X}^T \mathbf{X}) \mathbf{v}_1$  is given by the first eigenvector of  $\mathbf{X}^T \mathbf{X}$ . This is also the first right singular vector of  $\mathbf{X}$ .

Now consider the max over  $\mathbf{v}_2 \perp \mathbf{v}_1$ . This can be simplified by writing  $\mathbf{v}_2 = (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^T) \mathbf{z}$  and maximizing over  $\mathbf{z}$ . With this, and with expressing  $\mathbf{X}^T \mathbf{X}$  in terms of its EVD, it can be shown that the maximizer will be the second eigenvector (eigenvector with second largest eigenvalue). This is the second right singular vector of  $\mathbf{X}$ .

*The proof of the second eigenvector claim is not provided; need to fill this in*

For max over  $\mathbf{v}_3$ , we can express  $\mathbf{v}_3 = (\mathbf{I} - \mathbf{v}_1 \mathbf{v}_1^T - \mathbf{v}_2 \mathbf{v}_2^T) \mathbf{z}$  and repeat the same argument as above.

### 12.8.3 Equivalence between statistical and computational claims 2

Similarly, using the empirical approximation

$$\mathbb{E}[\mathbf{w}^T (\mathbf{x} - \mathbf{V}\mathbf{b})(\mathbf{x} - \mathbf{V}\mathbf{b})^T \mathbf{w}] \approx \sum_{i=1}^m \mathbf{w}^T (\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)}) (\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})^T \mathbf{w}$$



the second claims from the statistical and computational viewpoints are also equivalent, i.e.

$$\min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \lambda_{\max}(\mathbb{E}[(\mathbf{x} - \mathbf{V}\mathbf{b})(\mathbf{x} - \mathbf{V}\mathbf{b})^T]) = \min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \max_{\mathbf{w}: \|\mathbf{w}\|_2 = 1} \mathbb{E}[(\mathbf{w}^T(\mathbf{x} - \mathbf{V}\mathbf{b}))^2]$$

This is further equivalent to

$$\min_{\mathbf{L}: \text{rank } r} \|\mathbf{X} - \mathbf{L}\|_2^2 = \min_{\mathbf{B}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \lambda_{\max}\left(\sum_{i=1}^m (\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})(\mathbf{x}^{(i)} - \mathbf{V}\mathbf{b}^{(i)})^T\right)$$

#### 12.8.4 Proof of computational claim 2: induced 2 norm optimality

Goal is to find a rank  $r$  matrix  $\mathbf{L}$  that solves

$$\min_{\mathbf{L}: \text{rank } r} \|\mathbf{X} - \mathbf{L}\|_2$$

By Weyl-type inequality for singular values,  $\sigma_i(M) \leq \sigma_i(M_2) + \|M - M_2\|$  for any singular value  $i$  and any two matrices  $M, M_2$ .

Thus,  $\|\mathbf{X} - \mathbf{L}\|_2 \geq \sigma_i(\mathbf{X}) - \sigma_i(\mathbf{L})$ .

Since  $\mathbf{L}$  is rank  $r$ , it has only  $r$  nonzero singular values. Thus  $\sigma_{r+1}(\mathbf{L}) = 0$ . And so  $\|\mathbf{X} - \mathbf{L}\|_2 \geq \sigma_{r+1}(\mathbf{X})$ . Since this is true for any rank  $r$  matrix  $\mathbf{L}$ , it is also true for the minimizer, i.e.

$$\min_{\mathbf{L}: \text{rank } r} \|\mathbf{X} - \mathbf{L}\|_2 \geq \sigma_{r+1}(\mathbf{X})$$

Now if we can find a specific matrix  $\mathbf{L}$  for which  $\|\mathbf{X} - \mathbf{L}\|_2 = \sigma_{r+1}(\mathbf{X})$  that will be the minimizer (since the minimum value cannot be any smaller than this).

If we let  $\hat{\mathbf{L}} = \sum_{i=1}^r \sigma_i \mathbf{u}_i \mathbf{v}_i^T$  ( $r$ -SVD of  $\mathbf{X}$ ), then  $\mathbf{X} - \hat{\mathbf{L}} = \sum_{i=r+1}^{\min(m,n)} \sigma_i \mathbf{u}_i \mathbf{v}_i^T$  and so  $\|\mathbf{X} - \hat{\mathbf{L}}\|_2 = \sigma_{r+1}$ ; here  $\sigma_i = \sigma_i(\mathbf{X})$ .

## 12.9 PCA de-correlates the data: what does it mean \*\*

This claim depends on how the principal subspace is defined. In (statistical) theory, we are finding  $\mathbf{V}$  and  $\mathbf{b}$  that solves  $\min_{\mathbf{b}, \mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|(\mathbf{x} - \mu) - \mathbf{V}\mathbf{b} + \mathbf{V}\mathbf{V}^T \mu\|_2^2]$ , i.e., it minimizes the expected value of the reconstruction error. Since the minimizer over  $\mathbf{b}$  is  $\mathbf{b} = \mathbf{V}^T \mathbf{x}$ , we are actually finding  $\mathbf{V}$  that solves  $\min_{\mathbf{V}: \mathbf{V}^T \mathbf{V} = \mathbf{I}} \mathbb{E}[\|(\mathbf{x} - \mu) - \mathbf{V}\mathbf{V}^T(\mathbf{x} - \mu)\|_2^2]$  and then setting  $\mathbf{b} = \mathbf{V}^T \mathbf{x}$ . If we can find this  $\mathbf{V}$ , then the computed lower dimensional r.v.  $\mathbf{b}$  satisfies

$$\mathbb{E}[(\mathbf{b} - \mathbb{E}[\mathbf{b}])(\mathbf{b} - \mathbb{E}[\mathbf{b}])^T] \text{ is diagonal}$$

i.e.

$$\mathbb{E}[(\mathbf{b} - \mathbb{E}[\mathbf{b}])_j(\mathbf{b} - \mathbb{E}[\mathbf{b}])_k] = 0$$

for  $j \neq k$ .

*Proof: assume everything is zero mean for ease of writing. We have  $\mathbf{b} = \mathbf{V}^T \mathbf{x}$  so  $\mathbf{b}_j = \mathbf{v}'_j \mathbf{x}$ , thus,  $\mathbb{E}[\mathbf{b}_j \mathbf{b}_k] = \mathbf{v}'_j \mathbb{E}[\mathbf{x} \mathbf{x}'] \mathbf{v}_k$ . As explained in previous section,  $\mathbf{V}$  is the matrix of top  $r$  eigenvectors of  $\mathbb{E}[\mathbf{x} \mathbf{x}']$ , i.e., that  $\mathbb{E}[\mathbf{x} \mathbf{x}'] = \mathbf{V} \Sigma \mathbf{V}' + \mathbf{V}_\perp \Sigma_\perp \mathbf{V}'_\perp$ . Hence,  $\mathbb{E}[\mathbf{b}_j \mathbf{b}_k] = \mathbf{v}'_j (\mathbf{V} \Sigma \mathbf{V}' + \mathbf{V}_\perp \Sigma_\perp \mathbf{V}'_\perp) \mathbf{v}_k = \mathbf{v}'_j (\mathbf{V} \Sigma \mathbf{V}') \mathbf{v}_k = 0$ .*

In practice, we cannot find above but only its data-based (empirical) approximation. Hence in practice, we are always finding  $\mathbf{V}$  as the top  $r$  right singular vectors of  $\mathbf{Z} = \mathbf{X} - \mathbf{1}_m \mathbf{1}_m^T \mathbf{X} / m$  (recall: where  $\mathbf{1}_m = [1 \ 1 \ \dots \ 1]^T$  is an  $m$ -length vector of ones). With this choice of  $\mathbf{V}$ , “uncorrelated” means the following:

$$\sum_{i=1}^m (\mathbf{b}^{(i)})_j (\mathbf{b}^{(i)})_k = 0$$

for  $j \neq k$ . In other words, the columns of the matrix  $\mathbf{B}$  are mutually orthogonal.

*Proof: same basic idea as above.*

## 13 Unsupervised Learning: Clustering

Use ML-cs229- notes on Clustering to see the figures.

### 13.1 Problem

Given unlabeled data/features  $\mathbf{x}^{(i)}$ ,  $i = 1, 2, \dots, m$ , the goal is to partition the dataset into “cohesive” clusters (all points in same cluster are “close” while those in different clusters are “far”). Suppose we want to partition into  $k$  clusters. Then the goal can also be stated as: for each  $i$ , find the class label  $y^{(i)}$  (this can take values from  $\{1, 2, \dots, k\}$ ).

### 13.2 k-means clustering

The goal is to find the class labels  $y^{(i)}$  for each data point and the cluster centers so that the following cost is minimized

$$J(\mu_j, j = 1, 2, \dots, k, \mathbf{y}) = \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mu_{y^{(i)}}\|_2^2$$

The original k-means clustering algorithm provides an Alternating-Minimization (AltMin) algorithm to minimize the above cost.

1. Initialize cluster centers  $\hat{\mu}_1, \hat{\mu}_2, \dots, \hat{\mu}_k$ . Can do random init.
2. Repeat
  - (a) For each  $i = 1, 2, \dots, m$ , find the class labels

$$\hat{y}^{(i)} = \arg \min_{j=1,2,\dots,k} \|\mathbf{x}^{(i)} - \hat{\mu}_j\|_2^2.$$

(b) Update cluster centers: for each  $j = 1, 2, \dots, k$ , compute

$$\hat{\mu}_j = \frac{\sum_{i=1}^m \mathbb{1}(\hat{y}^{(i)} == j) \mathbf{x}^{(i)}}{\sum_{i=1}^m \mathbb{1}(\hat{y}^{(i)} == j)}$$

(the above is a solution to  $\arg \min_{\mu_j, j=1,2,\dots,k} \sum_{i=1}^m \|\mathbf{x}^{(i)} - \mu_{yhat^{(i)}}\|_2^2$ )

until cluster center estimates do not change much, i.e., until  $\max_j (\|\hat{\mu}_j^{(t+1)} - \hat{\mu}_j^{(t)}\| / \|\hat{\mu}_j^{(t)}\|) < threshold$  where  $threshold = 0.001$  or some small fraction.

3. IMPROVED VERSION: repeat above algorithm  $N$  times with different random init's each time. For each repeat, compute the cost function value  $J(\hat{\mu}_j, j = 1, 2, \dots, k, \hat{y})$ . Keep the output of the repeat with the smallest cost.

IMPROVEMENT 2: one can replace the regular Euclidean distance to the cluster center by any other distance that is more relevant to the application. As an example, if it is known that different features are likely to have significantly different variances, one could init with  $\hat{\Sigma}_j = \mathbf{I}$ , replace  $\arg \min_{j=1,2,\dots,k} \|\mathbf{x}^{(i)} - \hat{\mu}_j\|_2$  in step 2a by

$$\arg \min_{j=1,2,\dots,k} (\mathbf{x}^{(i)} - \hat{\mu}_j)^T \hat{\Sigma}_j^{-1} (\mathbf{x}^{(i)} - \hat{\mu}_j)$$

and in step 2b, update

$$\hat{\Sigma}_j = \frac{\sum_{i=1}^m \mathbb{1}(\hat{y}^{(i)} == j) (\mathbf{x}^{(i)} - \hat{\mu}_j)(\mathbf{x}^{(i)} - \hat{\mu}_j)^T}{\sum_{i=1}^m \mathbb{1}(\hat{y}^{(i)} == j)}$$

### 13.3 Alternating Minimization (AltMin) algorithm

AltMin is another approach to solve an optimization problem. It is a better one to use than Gradient Descent when the variable to be minimized over,  $\mathbf{x}$ , can be split into two subsets of variables  $\mathbf{x} = [\mathbf{x}_1, \mathbf{x}_2]$  such that minimizing over  $\mathbf{x}_1$  keeping  $\mathbf{x}_2$  fixed and vice versa is either closed form or otherwise easy to do.

AltMin proceeds as follows

1. Randomly initialize  $\mathbf{x}_1$  to  $\hat{\mathbf{x}}_1$ .

2. Repeat

(a) Minimize over  $\mathbf{x}_2$  keeping  $\mathbf{x}_1$  fixed at  $\hat{\mathbf{x}}_1$ , i.e., compute

$$\hat{\mathbf{x}}_2 = \arg \min_{\mathbf{x}_2} J(\hat{\mathbf{x}}_1, \mathbf{x}_2)$$

(b) Minimize over  $\mathbf{x}_1$  keeping  $\mathbf{x}_2$  fixed at  $\hat{\mathbf{x}}_2$ , i.e., compute

$$\hat{\mathbf{x}}_1 = \arg \min_{\mathbf{x}_1} J(\mathbf{x}_1, \hat{\mathbf{x}}_2)$$

until “convergence” i.e. estimates of  $\hat{\mathbf{x}}_1$  do not change much from previous to current iteration.

Like Grad Desc, AltMin also only converges to the local minimum of the cost function. Thus to make it work better one can run it  $N$  times with different random init's and pick the solutions that result in the smallest cost function value.

### 13.4 Probabilistic Model (Generative Model) for Clustering: Gaussian Mixture Model (GMM)

The Gaussian Mixture Model or GMM is a common way to specify a clustering problem.  $\mathbf{x}$  follows the GMM with  $k$  components means that

$$p(\mathbf{x}; \theta) = \sum_{j=1}^k \mathcal{N}(\mathbf{x}; \mu_j, \Sigma_j) \phi_j$$

where  $\phi_j$ 's are the mixture weights (probability of  $x$  coming from the  $j$ -th class in the mixture) and thus  $\sum_{j=1}^k \phi_j = 1$ . We often refer to  $j$  as the class labels.

GMM: means that  $\mathbf{x}$  is generated from class  $j$  with probability  $\phi_j$ , and given that it is generated from class  $j$ , it follows a Gaussian distribution with mean  $\mu_j$  and covariance  $\Sigma_j$ .

The model assumed by Gaussian discriminant analysis (GDA) covered earlier and in HW 2 is also GMM. See Sec 5.2. Except there we had labeled data (for each training data point, the class label was available), so it was easy to “learn” the model parameters. As a result, the learning of  $\phi_j$ 's was decoupled from the learning of the mean and covariances. As a result we in fact got closed form expressions for the MLE.

Here we do not have class labels and thus the learning problem is more difficult. We need to use Gradient Descent or some other iterative approach.

### 13.5 EM algorithm for MLE for Gaussian Mixture Model \*\*

A popular approach for MLE for the GMM is the EM algorithm.

EM algorithm is another approach (besides AltMin and Grad Descent) to solve a Maximum Likelihood estimation (MLE) problem. Like AltMin it is useful for certain types of problems in which, by using some tricks, part of the problem can be made non-iterative (closed form) or can otherwise be simplified.

Consider Maximum Likelihood Estimation over i.i.d. data samples  $x^{(i)}$  coming from a distribution  $p(x; \theta)$ . The goal is to maximum the likelihood or equivalently its log, i.e.,

$$\max_{\theta} \ell(\theta) := \sum_{i=1}^m \log p(x^{(i)}; \theta)$$

For simple cases like Gaussian, this is easy to do (can get a closed form expression). But in other problems it is not easy, the cost function is not convex for example. In certain such settings, the EM algorithm helps simplify. Consider the Gaussian mixture model with  $k$  components. Then

$$p(x; \theta) = \sum_{j=1}^k \mathcal{N}(x; \mu_j, \Sigma_j) \phi_j$$

where  $\phi_j$ 's are the mixture weights (probability of  $x$  coming from the  $j$ -th class in the mixture) and thus  $\sum_{j=1}^k \phi_j = 1$ . We often refer to  $j$  as the class labels. Thus,

$$\ell(\theta) = \sum_{i=1}^m \log \sum_{j=1}^k \phi_j \mathcal{N}(x; \mu_j, \Sigma_j) \quad \text{s.t.} \quad \sum_j \phi_j = 1$$

and the parameters  $\theta$  are

$$\theta = \{\mu_j, \Sigma_j, \phi_j\}, j = 1, 2, \dots, k$$

with  $\mu_j$  being  $n \times 1$ ,  $\phi_j$  is a scalar and  $\Sigma_j$  is  $n \times n$ . Recall  $\mathcal{N}(\mathbf{x}; \mu, \Sigma) := C \exp(-(\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu) + \log(\det(\Sigma)))$  where  $C$  is a constant (contains the  $1/\sqrt{2\pi}$  etc terms).

$\ell(\theta)$  is a pretty messy expression to even compute the gradient of: notice it involves  $\log \sum_j \phi_j \exp(\dots)$ . And Grad Descent will not work well. Compare this with the expression we had for GDA: since  $y^i$ 's were

known, there was no log of sums of weighted exponentials. We instead just had sum of log of exponentials which simplified easily.

For the above problem, the following is an easier approach – somehow try to approximate the “complete data likelihood”. Since class labels are not available, we call this “incomplete data”. The idea proceeds as follows.

We will try to lower bound  $\ell(\theta)$  and then maximize the lower bound by Alternating-Minimization. To do this we will first multiply divide the above expression by  $q_i(j)$  which are such that  $q_i(j) \geq 0$  and  $\sum_{j=1}^k q_i(j) = 1$ . So one can think of them as some probability distribution over the class labels.

$$\begin{aligned} \ell(\theta) &= \sum_{i=1}^m \log \sum_{j=1}^k \phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j) \\ &= \sum_{i=1}^m \log \sum_{j=1}^k \phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j) \frac{q_i(j)}{q_i(j)} \\ &= \sum_{i=1}^m \log \sum_{j=1}^k q_i(j) \frac{\phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)}{q_i(j)} \\ &\geq \sum_{i=1}^m \sum_{j=1}^k q_i(j) \log \frac{\phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)}{q_i(j)} \end{aligned}$$

with equality holding if and only if  $q_i(j) = c \phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)$ . The last step used log-sum (Jensen)’s inequality. Since the  $q_i(j)$ ’s sum to 1, we get that

$$c = \frac{1}{\sum_{j=1}^k \phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)}.$$

In all of above we also needed to keep the constraint  $\sum_j \phi_j = 1$  but we skipped it for ease of writing.

Now we will use AltMin to MAXIMIZE the above LOWER BOUND ON  $\ell(\theta)$ .

By the above approach, we have increased the number of unknowns - now the unknowns include  $q_i(j)$ ’s and  $\theta$ . BUT, the maximization over  $\theta$  given  $q_i(j)$ ’s is easy (closed form, as we will see below); and the maximization over  $q_i(j)$ ’s given  $\theta$  fixed is also easy (the lower bound is maximized when the inequality holds with equality, i.e., for  $q_i(j) = c \phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)$  with  $c$  specified above).

### 13.6 Final EM algorithm for Gaussian Mixture Models \*\*

Thus, AltMin over  $q_i(j)$ ’s and  $\theta$  is our FINAL ALGORITHM. It is called “EM algorithm for Gaussian Mixture Models”. This is summarized next.

- Initialize the values of  $\theta$ . Recall  $\theta = \{\phi_j, \mu_j, \Sigma_j\}, j = 1, 2, \dots, k$ . Can be random or something else.
- Iterate the following two steps until “convergence” (some reasonable stopping criterion holds):
  - E-step (maximize lower bound on  $\ell(\theta)$  over  $q_i(j)$ ’s, holding  $\theta$  fixed): this is obtained by

$$q_i(j) = \frac{\phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)}{\sum_{j=1}^k \phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)}, \quad j = 1, 2, \dots, k, i = 1, 2, \dots, m.$$

- M-step (maximize lower bound on  $\ell(\theta)$  over  $\theta$  holding  $q_i(j)$  fixed at the above value:

$$\max_{\theta} \sum_{i=1}^m \sum_{j=1}^k q_i(j) \log \frac{\phi_j \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)}{q_i(j)}$$

Since  $q_i(j)$ 's are held fixed, this is equivalent to

$$\max_{\theta} \ell_2(\theta) := \sum_{j=1}^k \sum_{i=1}^m q_i(j) [\log \phi_j + \log \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)] \quad \text{s.t.} \quad \sum_j \phi_j = 1$$

Notice we have converted log of weighted sums of exponentials to “weighted sum of log of exponentials”, so the form of above expression is similar to that for GDA.

This is easy, it is almost like GDA because it can be separated out over subsets of variables:

\* Solve for  $\phi_j$ 's

$$\max_{\phi_1, \phi_2, \dots, \phi_k} \sum_{j=1}^k q_i(j) \log \phi_j \quad \text{s.t.} \quad \sum_j \phi_j = 1$$

This is solved by  $\hat{\phi}_j = \frac{1}{m} \sum_{i=1}^m q_i(j)$ . This <sup>2</sup>

\* For each  $j = 1, 2, \dots, k$ ,

$$\max_{\mu_j, \Sigma_j} \sum_{j=1}^k \sum_{i=1}^m q_i(j) \log \mathcal{N}(x^{(i)}; \mu_j, \Sigma_j)$$

This is solved by

$$\hat{\mu}_j = c \sum_{i=1}^m q_i(j) x^{(i)}$$

and

$$\hat{\Sigma}_j = c \sum_{i=1}^m q_i(j) (x^{(i)} - \hat{\mu}_j)(x^{(i)} - \hat{\mu}_j)^T$$

with  $c = \frac{1}{\sum_{i=1}^m q_i(j)}$ .

- Repeat with  $N$  different initializations. Pick the one for which we get the largest value of  $\ell(\theta)$

The “E step” is the “Expectation step”. Reason it is called this is the following. We have data that is incomplete. We do not know the values of the missing entries, and so we do not know the complete data likelihood (in this case this corresponds to the likelihood func for GDA). But given estimates of the parameters from the previous iteration, we can compute what is called the posterior expectation of the complete data likelihood.

- The  $q_i(j)$ 's are called the posterior probabilities of  $y^{(i)} = j$  given observed data, and
- $\ell_2(\theta)$  is called the posterior Expectation of the complete data likelihood.

Thus “E step” computes this posterior expectation given previous estimate of parameters, while “M step” maximizes this posterior expectation to find a new value of the parameters. Hence the same EM algorithm.

### 13.7 General EM \*\*

In the above writing, I have written the above out for the Gaussian Mixture Model first to make it easier to understand. More generally, you could introduce any “missing data” variables  $y^{(i)}$  and use a process similar

<sup>2</sup>we actually need  $c = \frac{1}{\sum_{j=1}^k \sum_{i=1}^m q_i(j)}$ ; not hard to see that this simplifies to  $1/m$ .

to the above. If the labels are discrete-valued, then we can write things as

$$\begin{aligned}
\ell(\theta) &= \sum_{i=1}^m \log \sum_{j=1}^k p(x^{(i)}, y^{(i)} = j; \theta) \quad \text{s.t.} \quad \sum_{j=1}^k q_i(y^{(i)} = j) = 1 \\
&= \sum_{i=1}^m \log \sum_{j=1}^k p(x^{(i)}, y^{(i)} = j; \theta) \frac{q_i(y^{(i)} = j)}{q_i(y^{(i)} = j)} \quad \text{s.t.} \quad \sum_{j=1}^k q_i(y^{(i)} = j) = 1 \\
&= \sum_{i=1}^m \log \sum_{j=1}^k q_i(y^{(i)} = j) \frac{p(x^{(i)}, y^{(i)} = j; \theta)}{q_i(y^{(i)} = j)} \quad \text{s.t.} \quad \sum_{j=1}^k q_i(y^{(i)} = j) = 1 \\
&\geq \sum_{i=1}^m \sum_{j=1}^k q_i(y^{(i)} = j) \log \frac{p(x^{(i)}, y^{(i)} = j; \theta)}{q_i(y^{(i)} = j)} \quad \text{s.t.} \quad \sum_{j=1}^k q_i(y^{(i)} = j) = 1
\end{aligned}$$

We now use Alternating-Maximization to maximize the above lower bound over  $q_i(y^{(i)} = j)$ 's keeping  $\theta$  fixed and vice-versa.

We know that the RHS cannot larger than its upper bound, thus if we can find a value that helps achieve the upper bound we are done. The following is easy to see, if  $\sum_j g_j = 1$ , then,  $\sum_j g_j \log c = \log c = \log \sum_j g_j c$ . Thus the lower bound is maximized when

$$q_i(y^{(i)} = j) = cp(x^{(i)}, y^{(i)} = j; \theta)$$

Since the above sums to one over  $j$ , we need  $c = 1/p(x^{(i)}; \theta)$  and so

$$q_i(y^{(i)} = j) = p(y^{(i)} = j|x^{(i)}; \theta)$$

Thus the general EM algorithm proceeds as follows.

- E step: keeping  $\theta$  fixed at its previous value, compute

$$q_i(y^{(i)} = j) = p(y^{(i)} = j|x^{(i)}; \theta), j = 1, 2, \dots, k, i = 1, 2, \dots, m$$

- M step: keeping  $q_i(y^{(i)} = j)$  fixed at above value, compute

$$\begin{aligned}
&\arg \max_{\theta} \sum_{i=1}^m \sum_{j=1}^k q_i(y^{(i)} = j) \log \frac{p(x^{(i)}, y^{(i)} = j; \theta)}{q_i(y^{(i)} = j)} \\
&= \arg \max_{\theta} \sum_{i=1}^m \sum_{j=1}^k q_i(y^{(i)} = j) \log p(x^{(i)}, y^{(i)} = j; \theta) \\
&= \arg \max_{\theta} \sum_{j=1}^k \sum_{i=1}^m q_i(y^{(i)} = j) \log p(x^{(i)}, y^{(i)} = j; \theta)
\end{aligned}$$

This is often easier if the parameters  $\theta$  have a mixture-model type form and the maximization can be separated out

In more general settings  $y^{(i)}$  may be real-valued (continuous r.v.'s). In these cases, the summation over  $j$  gets replaced by integration, but a lot of the essential approach remains the same. A log-sum inequality exists for integrals also.

## 14 Supervised Learning: Deep Learning / Neural Networks: basic idea and training

This is another Supervised Learning approach.

## 14.1 Multi-layer Perceptron (MLP) or Feed-forward Network

The simplest neural net is the Feed-forward Network also called the Multilayer Perceptron or MLP. Each neuron receives a weighted sum of the outputs of the neurons of the previous layer, and applies a nonlinear “activation function” on this. Thus neuron  $j$  in layer  $k$  receives  $o^{k-1}$  as input. It then computes

$$z_j^k = \sum_{i=1}^{r_k} w_{ij}^k o_i^{k-1}$$

and then outputs

$$o_j^k = g(z_j^k)$$

Here  $g(z)$  is an element-wise nonlinearity. It could be the sigmoid function  $1/(1 + e^{-z})$  or the Rectified Linear Unit (ReLU) function  $\max(z, 0)$  or the tanh function.

Vectorizing the above, it can be expressed as

$$\mathbf{o}^k = g_{vec}(\mathbf{z}^k), \quad \mathbf{z}^k = \mathbf{W}^k \mathbf{o}^{k-1}$$

or equivalently,

$$\mathbf{z}^k = \mathbf{W}^k g_{vec}(\mathbf{z}^{k-1})$$

Here  $g_{vec}(\mathbf{z})$  applies  $g(\cdot)$  to each entry of the vector  $\mathbf{z}$ .

The first layer takes the input  $\mathbf{x}$  as input thus

$$\mathbf{z}^1 = \mathbf{x}$$

and computes  $g_{vec}(\mathbf{z})$ . Suppose the NN has 10 layers. The final (output) layer has only one neuron which outputs

$$\hat{y} = g(z_1^{10})$$

Thus, the NN is

$$\hat{y} = g(\mathbf{W}^{10} g_{vec}(\mathbf{W}^9 g_{vec}(\mathbf{W}^8 \dots g_{vec}(\mathbf{x}))))$$

where  $\mathbf{W}^{10}$  is a row vector (instead of a matrix).

## 14.2 Training an MLP: Back-propagation

Given training data  $(\mathbf{x}^{(i)}, y^{(i)})$ ,  $i = 1, 2, \dots, m$ , we use gradient descent or stochastic / mini-batch GD to train. For all of these, the first task is to define the cost function  $E(\hat{y}(\mathbf{x}), y)$  and to compute its gradient w.r.t. to each weight in each layer, i.e., compute

$$\frac{\partial E}{\partial w_{ij}^k}$$

This computation requires careful application of chain rule of differentiation. This leads to the following algorithm: consider a 10 layer NN and let  $r_k$  denote the number of neurons in layer  $k$

- For a given input  $\mathbf{x}$ , compute the outputs of all the layers and  $\hat{y}$ .
- Compute the following intermediate quantity:

$$\delta_i^k := \frac{\partial E}{\partial z_i^k}$$

using the following backward recursion (back-propagation)



– compute

$$\delta_1^{10} = \frac{\partial E}{\partial \hat{y}}(\hat{y}, y) \cdot g'(z_1^{10})$$

here  $g'(z) = \frac{\partial g(z)}{\partial z}$

– for each  $k = 9, 8, \dots, 1$ , compute the following for each  $i = 1, 2, \dots, r_k$ :

$$\delta_i^k = g'(z_i^k) \sum_{j=1}^{r_{k+1}} w_{ij}^{k+1} \delta_j^{k+1}$$

vectorized computation in MATLAB (or do similar in Python):  $\boldsymbol{\delta}^k = g'(z^k) \cdot * (\mathbf{W}^{k+1} \boldsymbol{\delta}^{k+1})$

• Compute

$$\frac{\partial E}{\partial w_{ij}^k} = \delta_j^k \cdot o_i^{k-1}$$

this can be vectorized too.

The above gives us  $\frac{\partial E}{\partial w_{ij}^k}(\hat{y}(\mathbf{x}), y)$  for one input-output pair  $\mathbf{x}, y$ . The gradient w.r.t. the cost function that uses all the training data is thus

$$\frac{1}{m} \sum_{i=1}^m \frac{\partial E}{\partial w_{ij}^k}(\hat{y}(\mathbf{x}^{(i)}), y^{(i)})$$

### 14.3 Understanding CNN's at a top level with visuals

See <https://towardsdatascience.com/beginners-guide-to-understanding-convolutional-neural-networks->

### 14.4 Aside: understanding convolution

1D convolution:

2D convolution:

types of filters: low-pass, high-pass, band-pass,

low-pass filter, e.g., averaging filter: reduces noise but also blurs edges

high-pass filter, e.g., edge detector: nonzero response only at edges, zero response in parts of image that are constant (piecewise constant image)

### 14.5 Convolutional Neural Network (CNN or ConvNet): basic idea

In case of MLP, each layer is fully-connected (FC), this means that all (or most) weights of the weight matrix are nonzero.

In case of CNN, this is not true. CNN consists of

- a convolutional (CN) layer followed by
- a pooling layer (pool),
- this sequence repeated a few times,
- finally 2-3 fully-connected (FC) layers.

### 14.6 Different types of layer

Let  $\mathbf{x}$  be the input to the layer and  $\mathbf{z}$  the output.

### 14.6.1 Fully connected (FC) layer:

$\mathbf{x}, \mathbf{z}$  are represented as vectors

- $\mathbf{x}$  is  $n \times 1$ ,  $\mathbf{z}$  is  $n' \times 1$ ,  $n'$  can be  $>, =, < n$
- parameters:  $\mathbf{W}$
- number of parameters:  $n \cdot n'$
- apply nonlinearity, then weights' matrix

$$\mathbf{z} = \mathbf{W} g_{vec}(\mathbf{z})$$

### 14.6.2 Convolutional (CN) layer:

$\mathbf{x}, \mathbf{z}$  are represented as 3D tensors

- $\mathbf{x}$  is  $n_1 \times n_2 \times K$ ,  $\mathbf{z}$  is  $n'_1 \times n'_2 \times K'$  with  $n'_1 \approx n_1$ ,  $n'_2 \approx n_2$
- parameters:  $h_k(u, v)$ ,  $k = 1, 2, \dots, K'$ ,  $u = 0, 1, \dots, \ell_h$ ,  $v = 0, 1, \dots, \ell_h$
- number of parameters:  $\ell_h^2 \cdot K'$
- apply nonlinearity, convolution over the first 2 dimensions, sum over the third

$$z(a, b, k) = \sum_w \left( \sum_u \sum_v h_k(u, v) g_{vec}(\mathbf{x}(a + u, b + v, w)) \right)$$

- Details:
  - dealing with edges: zero pad, or reduce dimension of output
  - stride: above is stride = 1 (moving filter window by 1 pixel at a time) – default option
  - stride = 2, move filter window by 2 pixels
  - stride =  $\ell_h$ : non-overlapping window

### 14.6.3 Pooling layer:

$\mathbf{x}, \mathbf{z}$  are represented as 3D tensors

- $\mathbf{x}, \mathbf{z}$  are represented as 3D tensors,  $\mathbf{x}$  is  $n_1 \times n_2 \times K$ ,  $\mathbf{z}$  is  $n'_1 \times n'_2 \times K$
- if using a  $\ell \times \ell$  non-overlapping window for pooling, then  $n'_1 = n_1/\ell$ ,  $n'_2 = n_2/\ell$
- parameters: none
- number of parameters: zero
- applied on each  $k = 1, 2, \dots, K$  "image" or "activation map" separately.
- helps reduce the output dimension (first 2 dimensions)
- max pooling over a  $\ell \times \ell$  size window: keep the largest magnitude pixel value and throw the rest
- average pooling:
- stride =  $\ell$ : non-overlapping window – default option

## 14.7 Why CN + pooling instead of FC

- Local connectivity : useful when inputs are images
- Much fewer parameters (parameter sharing):  $\ell_h^2 \cdot K'$ , even with  $\ell_h = 11, K' = 96$  (first layer of Alex Net), number of parameters is around 1000  
but if we vectorize even a  $100 \times 100$  image and the next layer has more neurons than input layer (typically this is used), then there are more than 10000 parameters in an FC

## 14.8 Recurrent Neural Network (RNN): basic idea

To Do

## 14.9 Different NN architectures and when to use each

see the other handout NeuralNets-intro.pdf

## 15 Reliability of an output \*\*

Linear regression predicts a real-valued scalar where as everything we learnt after that predicts a class label (solves a classification problem). Given a query, we can always obtain a prediction. But the other important question to answer is : how reliable is the prediction we obtained. The answer to this question depends on

- the problem itself, e.g., in case of classification by GDA, if the two class means are very close, it is not easy to distinguish the classes. More precisely what matters is how close the class means along a given direction compared to the standard deviation along that direction. Practically this means the following: it easier to distinguish dog pictures from human pictures than from cat pictures

Similarly for a regression problem, the amount of modeling error  $e$  or its variance decides how good the prediction is.

- the number of training data points, and how well the training and test data match (this decides quality of learnt model). In most of what we learn, it is assumed that training and test data are generated from the same distribution, but in real life this may not be true.
- the specific query: if the query image is of a fluffy cat that may look dog-like, then it is hard to reliably provide a correct classification.
- the last problem can be partly addressed by changing the learning algorithm (the assumed model on the data).

## 16 Bias-Variance Tradeoff

### 16.1 What is it?

Consider a generative model: suppose that  $y, \mathbf{x}$  satisfy

$$y = f(\mathbf{x}) + e, \quad \mathbb{E}[e] = 0, \mathbb{E}[e^2] = \sigma^2$$

with  $e$  being zero mean “modeling error”/“noise” that is independent of  $\mathbf{x}$ , it is also independent for each data point  $y_i$ .

We do not know  $f(\cdot)$ .

We try to “model” it as  $\hat{y} = \hat{f}(x) = h_{\hat{\theta}}(\mathbf{x})$ , e.g., in linear regression,  $\hat{f}(x) = h_{\hat{\theta}}(\mathbf{x}) = \hat{\theta}^T \mathbf{x}$  with  $\hat{\theta}$  estimated by Maximum Likelihood estimation (MLE) as described earlier using training data

$$(y^i, \mathbf{x}^i), i = 1, 2, \dots, m$$

In logistic regression,  $\hat{f}(x) = h_{\hat{\theta}}(\mathbf{x}) = g(\hat{\theta}^T \mathbf{x})$  with  $g(\cdot)$  being the sigmoid function. MLE uses “training data”

The question is how good is my learnt model (in terms of mean squared error on test data), i.e., for a test query  $\mathbf{x}$ , what is  $\mathbb{E}[(y - \hat{f}(\mathbf{x}))^2]$  and what can we do to improve it?

We define Test-MSE as

$$\text{Test-MSE} := \mathbb{E}[(y - \hat{y})^2] = \mathbb{E}[(y - \hat{f}(\mathbf{x}))^2] = \mathbb{E}[(f(\mathbf{x}) + e - \hat{f}(\mathbf{x}))^2]$$

over test data, i.e.,  $\mathbb{E}[\cdot] \equiv \mathbb{E}_{\text{test data}}[\cdot]$ . Usually  $\mathbf{x}$  is treated as a constant, so then the expected value is over the distribution of the noise  $e$ .

Since  $e$  is test-data noise, it is independent of  $\hat{f}(\mathbf{x}) = h_{\hat{\theta}}(\mathbf{x})$  since  $\hat{\theta}$  was estimated using training data. Also, by assumption,  $e$  is independent of  $f(\mathbf{x})$ . Thus, we have

$$\begin{aligned} \text{Test-MSE} &= \mathbb{E}[e^2] + \mathbb{E}[(f(\mathbf{x}) - \hat{f}(\mathbf{x}))^2] \\ &= \mathbb{E}[e^2] + (\mathbb{E}[f(\mathbf{x}) - \hat{f}(\mathbf{x})])^2 + \text{Variance}[f(\mathbf{x}) - \hat{f}(\mathbf{x})] \\ &= \sigma^2 + \text{Bias}^2 + \text{Variance} \end{aligned}$$

where  $\text{Variance}(Z) := E[(Z - E[Z])^2]$ . The first term,  $\sigma^2$ , depends on how noisy the data is. The second two terms depend on the “assumed model” and how well its parameters are estimated.

## 16.2 Bias-Variance Tradeoff for Linear Regression

THERE IS A MISTAKE HERE: NEED TO FIX THESE NOTES to all for a constant ”mean” term in linear regression.

Suppose that  $y$  truly satisfies the following model

$$y = f(x) + e, \quad f(x) := \theta_{\text{full}}^T \mathbf{x}_{\text{full}}, \quad \mathbb{E}[e] = 0, \quad \text{Var}[e] = \sigma^2$$

where  $\mathbf{x}_{\text{full}}$  is the full  $n_{\text{full}}$  length ”feature” or data. Also, the noise  $e$  is i.i.d. across various samples.

For the sake of tractability (reducing variance), when “modeling”  $y$ , we throw away some of the features to get an  $n$ -length “feature” vector  $\mathbf{x}$ ; and refer to the dropped part of  $\mathbf{x}_{\text{full}}$  as  $\mathbf{x}_{\text{drop}}$ . This is a  $n_{\text{drop}}$  length vector. Thus  $n_{\text{full}} = n + n_{\text{drop}}$ .

$$\mathbf{x}_{\text{full}}^T = [\mathbf{x}^T, \mathbf{x}_{\text{drop}}^T]$$

Thus,  $y$  can be rewritten as

$$y = \theta^T \mathbf{x} + \underbrace{\theta_{\text{drop}}^T \mathbf{x}_{\text{drop}}}_{\mu} + e.$$

Thus, the ”linear regression” model for  $y$  is  $\theta^T \mathbf{x}$ . Assume as before we are given training data  $\{y^{(i)}, \mathbf{x}^{(i)}\}, i = 1, 2, \dots, m$ , and define the matrix  $\mathbf{X}$  and the vector  $\mathbf{y}$  as before.

We use MLE under this model to get the MLE estimate

$$\hat{\theta} := (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}, \quad \text{with } \mathbf{y} := \mathbf{X} \theta + \mathbf{X}_{\text{drop}} \theta_{\text{drop}} + \mathbf{e}$$

where  $\mathbf{X}_{\text{drop}} := [\mathbf{x}_{\text{drop}}^{(1)}; \mathbf{x}_{\text{drop}}^{(2)}; \dots; \mathbf{x}_{\text{drop}}^{(m)}]^T$  is an  $m \times n_{\text{drop}}$  matrix with the dropped parts of each training data vector as its rows. Recall that  $\mathbf{e}$  (training data noise vector) is independent of  $\mathbf{X}$  (training data).

Thus, for a query  $\mathbf{x}_{\text{full}}^{\text{tst}}$ , with extracted features  $\mathbf{x}^{\text{tst}}$ , we predict

$$\hat{y}^{\text{tst}} = \mathbf{x}^{\text{tst}T} \hat{\theta}$$

The true output,  $y^{\text{tst}}$ , for the query satisfies

$$y^{\text{tst}} = \theta^T \mathbf{x}^{\text{tst}} + \theta_{\text{drop}}^T \mathbf{x}_{\text{drop}}^{\text{tst}} + e^{\text{tst}}$$

Observe that  $\hat{\theta}$  satisfies

$$\hat{\theta} := (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T (\mathbf{X} \theta + \mathbf{X}_{\text{drop}} \theta_{\text{drop}} + \mathbf{e}) = \theta + (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X}_{\text{drop}} \theta_{\text{drop}} + (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{e}$$

So,

$$\begin{aligned} \hat{y}^{\text{tst}} - \mathbf{y}^{\text{tst}} &= \mathbf{x}^{\text{tst}T} \hat{\theta} - \mathbf{x}^{\text{tst}T} \theta - \mathbf{x}_{\text{drop}}^{\text{tst}T} \theta_{\text{drop}} - e^{\text{tst}} \\ &= \underbrace{\mathbf{x}^{\text{tst}T} ((\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X}_{\text{drop}} \theta_{\text{drop}} + (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{e}) - \mathbf{x}_{\text{drop}}^{\text{tst}T} \theta_{\text{drop}} - e^{\text{tst}}}_Z \\ &= Z - e^{\text{tst}} \end{aligned}$$

Recall that  $\mathbb{E}[\cdot]$  is expected value over test-data noise  $e^{\text{tst}}$ . As before (since  $e^{\text{tst}}$  independent of everything in  $Z$ ),

$$\begin{aligned} \text{TestMSE} &:= \mathbb{E}[(\hat{y}^{\text{tst}} - \mathbf{y}^{\text{tst}})^2] = \mathbb{E}[(e^{\text{tst}})^2] + \mathbb{E}[Z^2] \\ &= \mathbb{E}[(e^{\text{tst}})^2] + \text{Bias}(Z)^2 + \text{Variance}(Z) \end{aligned}$$

with

$$\begin{aligned} \text{Bias}(Z) &:= \mathbb{E}[\mathbf{x}^{\text{tst}T} ((\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X}_{\text{drop}} \theta_{\text{drop}} + (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{e}) - \mathbf{x}_{\text{drop}}^{\text{tst}T} \theta_{\text{drop}}] \\ &= \mathbf{x}^{\text{tst}T} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X}_{\text{drop}} \theta_{\text{drop}} - \mathbf{x}_{\text{drop}}^{\text{tst}T} \theta_{\text{drop}} \\ &= [\mathbf{x}^{\text{tst}T} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{X}_{\text{drop}} - \mathbf{x}_{\text{drop}}^{\text{tst}T}] \theta_{\text{drop}} \end{aligned}$$

The second row follows because  $\mathbf{e}$  is independent of  $\mathbf{X}$ ,  $\mathbf{X}_{\text{drop}}$  (training data noise is independent of training data features) and  $\mathbf{e}$  is independent of  $\mathbf{x}^{\text{tst}}$  and  $\mathbf{e}$  is zero mean.

Since the training data noise vector  $\mathbf{e}$  is zero mean and i.i.d., we have  $\mathbb{E}[\mathbf{e}\mathbf{e}^T] = \sigma^2 \mathbf{I}$ . Using this,

$$\begin{aligned} \text{Variance}(Z) &:= \mathbb{E}[(Z - \text{Bias}(Z))^2] \\ &= \mathbb{E}[(\mathbf{x}^{\text{tst}T} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{e})^2] \\ &= \mathbb{E}[(\mathbf{x}^{\text{tst}T} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{e})(\mathbf{e}^T \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}^{\text{tst}})] \\ &= \mathbf{x}^{\text{tst}T} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbb{E}[\mathbf{e}\mathbf{e}^T] \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}^{\text{tst}} \\ &= \mathbf{x}^{\text{tst}T} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{x}^{\text{tst}} \sigma^2 \end{aligned}$$

This is upper bounded by

$$\text{Variance}(Z) \leq \frac{\|\mathbf{x}^{\text{tst}}\|^2}{\lambda_{\min}(\mathbf{X}^T \mathbf{X})} \sigma^2$$

1. Variance: Thus the variance is smaller if the minimum eigenvalue of  $(\mathbf{X}^T \mathbf{X})$  is larger.

- Notice that  $\mathbf{X}$  is an  $m \times n$  matrix. Consider  $\mathbf{X}^T \mathbf{X}$  which is  $n \times n$ . If  $m < n$ , its minimum eigenvalue is zero. For  $m > n$ , its minimum eigenvalue is proportional to  $m/n$ . The reason is, for a fixed  $m$ , by the interlacing theorem for eigenvalues, as you increase  $n$ , the smallest eigenvalue can only decrease<sup>3</sup>.

Thus for a fixed available amount of training data,  $m$ , the Variance will either be reduced, or stay the same, if we use a smaller  $n$ .

2. Bias: Consider the Bias. This is some complicated function of  $\theta_{\text{drop}}$ . We can simplify it by making more assumptions, but that is not needed. What is clear is that Bias will become larger as  $n_{\text{drop}}$  is made larger (or equivalently, as  $n = n_{\text{full}} - n_{\text{drop}}$  is made smaller). Thus, as we reduce  $n$ , we increase  $n_{\text{drop}}$ , and so we increase Bias.

Thus with reducing  $n$ , Variance decreases but Bias increases.

<sup>3</sup>To understand this with a concrete example, if the training data vectors were i.i.d. standard Gaussian, then with high probability, the minimum eigenvalue is lower bounded by  $(\sqrt{m} - C\sqrt{n})^2$ . Reference: Vershynin tutorial.

## 17 Picking reduced dimension $r$ to minimize test-MSE (optimize Bias-Variance Tradeoff)

As noted in previous section,

$$\text{Test-MSE} := \mathbb{E}[(y - \hat{y})^2] = \mathbb{E}[(y - \hat{f}(\mathbf{x}))^2] = \mathbb{E}[(f(\mathbf{x}) + e - \hat{f}(\mathbf{x}))^2]$$

over test data, i.e.,  $\mathbb{E}[\cdot] \equiv \mathbb{E}_{\text{test data}}[\cdot]$ . Usually  $\mathbf{x}$  is treated as a constant, so then the expected value is over the distribution of the noise  $e$ .

Since  $e$  is test-data noise, it is independent of  $\hat{f}(\mathbf{x}) = h_{\hat{\theta}}(\mathbf{x})$  since  $\hat{\theta}$  was estimated using training data. Also, by assumption,  $e$  it is independent of  $f(\mathbf{x})$ . Thus, we have

$$\begin{aligned} \text{Test-MSE} &= \mathbb{E}[e^2] + \mathbb{E}[(f(\mathbf{x}) - \hat{f}(\mathbf{x}))^2] \\ &= \mathbb{E}[e^2] + (\mathbb{E}[f(\mathbf{x}) - \hat{f}(\mathbf{x})])^2 + \text{Variance}[f(\mathbf{x}) - \hat{f}(\mathbf{x})] \\ &= \sigma^2 + \text{Bias}^2 + \text{Variance} \end{aligned}$$

where  $\text{Variance}(Z) := E[(Z - E[Z])^2]$ . The first term,  $\sigma^2$ , depends on how noisy the data is. The second two terms depend on the “assumed model” and how well its parameters are estimated.

### 17.1 Approximating Test-MSE in practice

While we can write things as above, it is *not* actually possible to compute the above decomposition for test data.

**Simple approach:** All one can do is the following: for a given model, one can approximate Test-MSE using the following approach

- Split available training data into training and test data: in other words do not use all  $m$  data points to train, split them so that  $m = m_{\text{train}} + m_{\text{test}}$ .
- Use the  $m_{\text{train}}$  data points to train, i.e. to estimate  $\theta$  for the assumed model
- Approximate Test-MSE by  $\frac{1}{m_{\text{test}}} \sum_{j=1}^{m_{\text{test}}} (y_j - h_{\hat{\theta}}(\mathbf{x}_j))^2$

The above is one way to do what is called “Cross-Validation”. Typically, one uses  $m_{\text{test}} = 0.25m$  and  $m_{\text{train}} = 0.75m$  or similar.

**Leave-one-out Cross-Validation:** Do above with  $m_{\text{test}} = 1$ , and  $m_{\text{train}} = m - 1$ , but repeat the procedure  $m$  times to compute the average error (Test-MSE). See Cross Validation section earlier.

### 17.2 How to reduce test data MSE / generalization error a.k.a. Regularization a.k.a. Feature Selection

Try to reduce variance, while hopefully not increasing the bias too much:

1. Regularization on  $\theta$ : use domain knowledge such as assuming entries of  $\theta$  are in decreasing order of magnitude. If this is known, then you just use the first  $r$  features. Vary  $r$  to pick which option reduces test data error the most. This was an extra credit part of HW 2 (or 1b).
2. Regularization on  $\theta$  – model  $\theta$  as being sparse (only a few features are important but we do not know which ones), this is a generalization of the assumption used in item 1. In this case, we add  $\|\theta\|_1$  into the cost function for learning  $\theta$ .

In general, if the cost function is  $J(\theta)$ , then this approach says we replace it by  $J(\theta) + \lambda\|\theta\|_1$

In case of linear regression this becomes

$$\min_{\theta} \|\mathbf{y} - \mathbf{X}\theta\|^2 + \lambda\|\theta\|_1$$

**Do not try to write your own GD code for this. Use cvxopt or cvxpy etc.**  
**Vary  $\lambda$  as multiples of  $m$ , so  $\lambda = \beta m$ , with  $\beta = 0.01, 0.1, 1, 10$  etc**

3. Regularization on  $\theta$  – use the assumption that all features are important and no feature is too important (no weight can be too large): then use the  $\ell_2$  regularization, i.e. add  $\|\theta\|^2$  into the cost function for learning  $\theta$ .

In general, if the cost function is  $J(\theta)$ , then this approach says we replace it by  $J(\theta) + \lambda\|\theta\|^2$

In case of linear regression this becomes

$$\min_{\theta} \|\mathbf{y} - \mathbf{X}\theta\|^2 + \lambda\|\theta\|^2$$

**This cost has closed form solution, given next.**

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

4. Regularization on  $\theta$  – suppose prior knowledge is available that  $\theta$  is close to a given vector  $\theta_0$ , then add the  $\|\theta - \theta_0\|^2$  into the cost function for finding  $\theta$ . See Sec. 11.2.
5. Regularization on input features' matrix  $\mathbf{X}$  using PCA: PCA is a dimension reduction technique. It is used for feature selection under the assumption that “features having the most variation are the most important ones”. Also of course PCA helps de-correlate the features. A typical setting where this is useful is when the raw feature vector is all image pixels. There are too many pixels and pixels are highly correlated. Also, see PCA section. *PCA does not do much if features are first “normalized” and then PCA is applied, see Sec. 17.3 and 17.4.*
6. In case of Clustering (say assuming Gaussian Mixture Model), reducing  $k$  (number of classes) will reduce the variance, increasing it will reduce the bias.
7. Naive Bayes assumption described earlier is another way to reduce the number of parameters. This is used in generative or Bayesian learning settings where the feature vector itself is modeled probabilistically.

With each new intervention, compute the Test-MSE as explained above, see if it gets reduced or not. In general as you keep reducing the number of features that are being used, variance reduces but bias increases. Initially when you begin reducing the number, the reduction in variance is significant but the increase in bias is negligible. So Test MSE will decrease. This will happen until the bias starts increasing (Test MSE starts increasing). Use the number of features that results in the smallest Test MSE.

### 17.3 Feature normalization

When ML features have vastly different ranges of values, e.g., in the house price example, one house feature can be number of bedrooms which takes values say between 1 to 5, while a second feature is the house areas in square feet, this could be a few 100 square feet. The two features have vastly different ranges.

If these were not normalized, then the matrix  $\mathbf{X}$  would be ill-conditioned. This means that it will have a very high condition number. This, in turn, will imply inaccurate estimation if the closed form solution is used (due to numerical errors in inverting an ill conditioned  $\mathbf{X}^T \mathbf{X}$ ). If GD is used, the same can happen unless the step size is chosen to be very small. In this case, it will converge to the correct solution but will need too many iterations.

One solution is to normalize each feature: compute the  $j$ -th feature's mean  $\mu_j$  and standard deviation  $\sigma_j$ ; replace  $\mathbf{x}_j$  by  $(\mathbf{x}_j - \mu_j)/\sigma_j$ .

## 17.4 Feature normalization vs PCA

Feature normalization is used in ML settings like the above where different features have different ranges because they are measured using very different measuring tools or sensors; and consequently, larger variance does not mean the feature is more important. In the house features example, number of bedrooms is as or important than house area.

PCA on the other hand is a “dimension reduction” technique that is used when all entries of the original data vector are acquired using the same type of sensor / measuring tool. For example, image pixels. All CCD sensors are the same type in a camera. In such settings, it may be a valid assumption to claim that pixels with larger variance are more relevant features. When this assumption is true, then the PCA can be used as a feature selection technique.

However PCA is not very useful if it is applied after feature normalization. Reason: feature norm is ensuring the new features all have unit variance. It can still be useful to de-correlate.

## 18 Learning Theory

based on Andrew Ng's *cs229-notes-4*

Hypothesis refers to a hypothesized model on the input output data. So suppose we assume that  $y = \text{sign}(\theta^T \mathbf{x})$  then  $\text{sign}(\theta^T \mathbf{x})$  is the hypothesis.

Hypothesis class  $\mathcal{H}$  is the set of all hypothesis from a certain class, e.g., set of all linear classifiers is

$$\mathcal{H} = \{h : h = \text{sign}(\theta^T \mathbf{x}), \forall \theta \in \mathbb{R}^n\}$$

Empirical Risk (Training Error) is denoted by  $\hat{\varepsilon}(h)$ . For a loss function  $\text{loss}(y, \hat{y})$ , it is computed as

$$\hat{\varepsilon}(h) = \frac{1}{m} \sum_{i=1}^m \text{loss}(y^i, h(\mathbf{x}^i))$$

Generalization Error (Test data Error) is denoted by  $\varepsilon(h)$ . It is the expected loss for a query sample

$$\varepsilon(h) = \mathbb{E}[\text{loss}(y, h(\mathbf{x}))]$$

Assumptions

- Training and Test data are generated from the same probability distribution
- All training samples as well as test sample are mutually independent. The two assumptions combined mean that all training and test data are i.i.d. (independent identically distributed).

Under the above assumptions, by using a law of large numbers' result, it can usually be argued that  $\hat{\varepsilon}(h)$  converges to  $\varepsilon(h)$  in probability as  $m$  goes to infinity. We will look at a simple zero-one loss function and actually work this out.

Empirical Risk Minimization (ERM) means minimize the empirical risk over all hypotheses from a certain class, i.e., try to find

$$\hat{h} = \arg \min_{h \in \mathcal{H}} \hat{\varepsilon}(h)$$

Ideally what we would like to find is the hypothesis that minimizes the generalization error

$$h^* = \arg \min_{h \in \mathcal{H}} \varepsilon(h)$$

Thus the minimum generalization error is  $\varepsilon(h^*)$ , i.e.,

$$\varepsilon(h^*) = \min_{h \in \mathcal{H}} \varepsilon(h)$$

We would like to use  $h^*$  but we cannot compute it. We instead use  $\hat{h}$  on the test data too. The question is how much worse is this? I.e., how much worse is  $\varepsilon(\hat{h})$  compared to  $\varepsilon(h^*) = \min_{h \in \mathcal{H}} \varepsilon(h)$ ? We work this out for the zero-one loss.



## 18.1 Two probability results we use

**Lemma 18.1** (Hoeffding inequality for Bernoulli r.v.'s). *Let  $Z_1, Z_2, \dots, Z_m$  be iid Bernoulli random variables with parameter  $\phi$ . Let  $\hat{\phi} := \frac{1}{m} \sum_{i=1}^m Z_i$  be the empirical mean of these random variables. Then*

$$\Pr(|\hat{\phi} - \phi| > \gamma) \leq 2e^{-2\gamma^2 m}$$

**Lemma 18.2** (Union Bound). *For any  $K$  events  $A_1, A_2, \dots, A_K$*

$$\Pr(A_1 \cup A_2 \cdots \cup A_K) \leq \sum_{k=1}^K \Pr(A_k)$$

and thus

$$\Pr(A_1^c \cap A_2^c \cdots \cap A_K^c) \geq 1 - \sum_{k=1}^K \Pr(A_k)$$

## 18.2 Misclassification (zero-one) loss

Let us specialize to the misclassification (zero-one) loss:

$$\text{loss}(y, \hat{y}) = \mathbb{1}(y \neq \hat{y}) = \mathbb{1}(y \neq h(\mathbf{x}))$$

For this, the empirical risk for a given  $h$  is

$$\hat{\varepsilon}(h) = \frac{1}{m} \sum_{i=1}^m \mathbb{1}(y^i \neq h(\mathbf{x}^i))$$

while

$$\varepsilon(h) = \mathbb{E}[\mathbb{1}(y \neq h(\mathbf{x}))] = \Pr(y \neq h(\mathbf{x}))$$

Define the r.v.  $Z_i = \mathbb{1}(y^i \neq h(\mathbf{x}^i))$ . Clearly the  $Z_i$  are Bernoulli with probability of one equal to  $\varepsilon(h)$ . We will use the Hoeffding inequality

Using Hoeffding inequality, this means that, for a given hypothesis  $h$ ,

$$\Pr(|\hat{\varepsilon}(h) - \varepsilon(h)| > \gamma) \leq 2e^{-2\gamma^2 m}$$

Suppose for a moment that the size of the hypothesis class is  $k$ , i.e.,

$$|\mathcal{H}| = k.$$

Then, by Union Bound,

$$\Pr(|\hat{\varepsilon}(h) - \varepsilon(h)| > \gamma, \text{ for some } h \in \mathcal{H}) \leq 2ke^{-2\gamma^2 m}$$

or equivalently

$$\Pr(|\hat{\varepsilon}(h) - \varepsilon(h)| < \gamma, \text{ for ALL } h \in \mathcal{H}) \geq 1 - 2ke^{-2\gamma^2 m}$$

For this probability to be at least  $1 - \delta$ , we need to set  $\gamma = 2\sqrt{\frac{1}{2m} \log \frac{2k}{\delta}}$ . With this we can rewrite things as follows

With probability at least  $1 - \delta$ , for all  $h \in \mathcal{H}$ , with  $k = |\mathcal{H}|$ ,

$$|\hat{\varepsilon}(h) - \varepsilon(h)| \leq \gamma := \sqrt{\frac{1}{2m} \log \frac{2k}{\delta}}$$

Since above is true for all  $h$ , it is true for  $\hat{h}$  and  $h^*$  too. So with the above probability,

$$\varepsilon(\hat{h}) \leq \hat{\varepsilon}(\hat{h}) + \gamma \leq \hat{\varepsilon}(h^*) + \gamma \leq \varepsilon(h^*) + \gamma + \gamma = \min_{h \in \mathcal{H}} \varepsilon(h) + 2\gamma$$

The first and third inequalities used the Hoeffding bound from above; the second used the fact that  $\hat{\varepsilon}(\hat{h})$  is the minimum over all  $h$ .

Thus, substituting for  $\gamma$ , we can write the following theorem

**Theorem 18.3.** Consider a hypothesis class  $\mathcal{H}$  with  $|\mathcal{H}| = k$ . And consider zero-one loss. Suppose that  $\hat{h}$  minimizes the empirical risk.

With probability  $\geq 1 - \delta$ ,

$$\varepsilon(\hat{h}) \leq \left( \min_{h \in \mathcal{H}} \varepsilon(h) \right) + 2\sqrt{\frac{1}{2m} \log\left(\frac{2k}{\delta}\right)}$$

?? add corollary for lower bound on  $m$ .

This theorem works for a finite hypothesis class, but not for an infinite one, for example it does not work for the class of all linear classifiers. However since everything is done on a computer (say one on which each real number is represented by 64 bits), we can assume that there are really only a finite number of degrees of freedom. Consider a hypothesis class with  $d$  parameters all of which are real numbers. Then the total number of possible options for hypotheses is

$$k = \underbrace{2^{64} * 2^{64} \dots * 2^{64}}_{d \text{ times}} = 2^{64d}$$

We have the following corollary

**Corollary 18.4.** Consider a hypothesis class  $\mathcal{H}_d$  with  $d$  real-valued parameters. Assume implementation on a 64-bit computer so that  $k = |\mathcal{H}| = 2^{64d}$ . Consider zero-one loss. Suppose that  $\hat{h}$  minimizes the empirical risk. Also set  $\delta = 2^{-d}$ . Then, we can claim the following

With probability  $\geq 1 - 2^{-d}$ ,

$$\varepsilon(\hat{h}) \leq \left( \min_{h \in \mathcal{H}_d} \varepsilon(h) \right) + C\sqrt{\frac{65d+1}{2m}}$$

**Remark 18.5.** In most of this class, we have assumed  $d = n + 1$  parameters. This is true for linear regression and logistic regression. For GDA,  $d > n$ .

### 18.3 Tradeoffs

First note: we are discussing these tradeoffs using an "upper bound" on  $\varepsilon(\hat{h})$ . This upper bound may or may not be tight. The discussion is valid only when it is tight.

If we increase  $d$  (for a fixed  $m$ ),

- the size of the hypotheses class on a 64-bit computer is  $|\mathcal{H}_d| = 2^{64d}$  as argued above. Thus, increasing  $d$ , increases its size. This means, in the first term from the above corollary,  $\min_{h \in \mathcal{H}_d} \varepsilon(h)$ , we are taking a minimum over a larger set. Minimizing over a larger set implies the minimum value is the same or smaller (cannot be larger). Thus the first term can only *decrease* or *stay the same* if we increase  $d$ . This term does not depend on  $m$  at all.
- but the second term,  $\sqrt{\frac{65d+1}{2m}}$  clearly increases linearly with  $d$

Increasing  $d$  and hence the size of the hypotheses class is analogous to *reducing the bias*. Typically, up to a certain increase in  $d$ , the first term will decrease (bias will decrease). But when  $m$  is fixed, this also *increases the variance*.

In fact at the very least, we need  $m > d$  to even just get the RHS to be smaller than one. Notice  $\varepsilon(h)$  is a probability so has to be less than one,

### 18.4 Connections to what we have learnt

In case of the 0-1 loss, the empirical risk is computationally not possible to minimize because one will literally have to check all  $2^{64d}$  possible hypotheses values and compute the empirical risk for each of them to find the minimize. It is what is called a "combinatorial optimization" problem and cannot be done in any reasonable time.

Also, as noted above, we are discussing these tradeoffs using an "upper bound" on  $\varepsilon(\hat{h})$ . This upper bound may or may not be tight. The discussion is valid only when it is tight.

For the above two reasons, the Tradeoffs' discussion does not apply to Generative Learning Models like Gaussian Discriminant Analysis where we are placing many more assumptions on our data. For instance, in GDA, once we impose  $\Sigma_0 = \Sigma_1 = \Sigma$  and  $\Sigma$  diagonal, we have basically decoupled the learning of each scalar parameter for each feature. In this case, as seen in your Home it is possible to get good parameter estimates and thus good classification even with  $m < d$ , in this case  $d = 2n + n + 1$ .

It "somewhat" applies to SVMs, though not directly since there we aren't trying to find a  $\mathbf{w}, b$  to minimize empirical risk (training error), instead we are trying to find  $w, b$  to maximize the worst-case margin.

## 18.5 VC dimension

We will not talk about VC dimension (last two pages of Andrew Ng's cs229 notes-4) in this course.