

REVIEW QUESTIONS, Exam 2, Modeling

1. What was the N -armed bandit problem? In particular, what were the two competing goals, and why were they “competing”?

SOLUTION: In the n -armed bandit problem, we have n slot machines (or a single slot machine with n arms), where we do not know that payouts for any of the machines. Therefore, we have two goals: Find the mean payout, and maximize our “reward”.

The reason that these are competing goals is that, to maximize reward, you should always play the machine with the highest mean payout. On the other hand, if we are not certain what those mean payouts are, we could easily choose the wrong machine. To find the mean payouts of all the machines quickly, we out to play them all with equal probability.

2. In the N -armed bandit problem, how were the estimates of the payoffs, $Q_t(a)$, calculated?

SOLUTION: It was calculated as the sample mean of the actual payouts. For example, if we play machine a n_a times with payouts r_1, r_2, \dots, r_{n_a} then:

$$Q_t(a) = \frac{1}{n_a} \sum_{i=1}^{n_a} r_i$$

3. There were four “strategies” that we implemented as algorithms to solve the N -armed bandit problem. What were they? Be sure to give formulas where appropriate.

SOLUTION:

- The greedy algorithm is to always play the machine with the highest mean payout.
- The ϵ -greedy algorithm is to the play the greedy algorithm most of the time, but with probability ϵ , choose on of the other machines at random.
- The softmax strategy was to create a set of probabilities for each machine, where the highest probability corresponds to the highest mean payout, and the lowest probability corresponds to the lowest mean payout. In Matlab notation, if Q is a vector of current mean payouts from each machine (some positive, some negative), then the probability vector is computed as:

`P=exp(Q/tau)./sum(exp(Q/tau));`

where τ is a parameter that allows us to adjust the probabilities from large (all probs about equal) to very close to zero (the probability for the machine with the largest payout is close to 1).

- For the pursuit (win-stay, lose-shift) strategy, we start with all of the probabilities equal. If a^* is the machine with the highest current mean payout, then we increase

the probability of choosing that machine, while lowering the probabilities of the other. How far that adjustment goes can be made by adjusting β :

$$\begin{aligned} P_{t+1}(a^*) &= P_t(a^*) + \beta(1 - P_t(a^*)) \\ P_{t+1}(a) &= P_t(a) + \beta(0 - P_t(a)) \end{aligned}$$

Whoops- The questions 4, 5, 6 were included in 3 above.

4. Describe (in words) the greedy algorithm and the ϵ -greedy algorithm. Which is probably a better strategy?

SOLUTION: (Already given above). The ϵ -greedy strategy is better if there is uncertainty about which machine is best.

5. Describe in words the softmax strategy. Be sure to include appropriate formulas, and describe what the parameter τ does.

SOLUTION: See (3)

6. What was the pursuit strategy (or “Win-Stay, Lose-Shift”) for the N -armed bandit? Again, include appropriate formulas and describe what β does.

SOLUTION: See (3)

7. Matlab Questions:

- (a) What’s the difference between a script file and a function?

SOLUTION: Matlab will treat a script file as if it were executing the commands as you typed them. A function has an input, an output, then a rule for changing the input into the output. A function can be (and is usually) called from a script to simplify the script.

- (b) What does the following code fragment produce?

```
Q=[1 3 2 1 3];
idx=find(Q==max(Q));
```

SOLUTION: The solution is [2, 5]. These are the indices where the elements of vector Q are maximum.

- (c) What is the difference between `x=rand;` and `x=randn;`

SOLUTION: `rand` returns a random number that is uniformly distributed in the interval $[0, 1]$. `randn` returns a number that is normally distributed with zero mean and unit standard deviation.

- (d) What will P be:

```
x=[0.3, 0.1, 0.2, 0.4];
P=cumsum(x);
```

SOLUTION: `cumsum` is Matlab for “cumulative sum”. The vector P would be a vector of “partial sums”:

$$P = [0.3, 0.4, 0.6, 1.0]$$

(e) What is the Matlab code that will:

- i. Plot $x^2 - 3x$ using 500 points, for $x \in [-1, 4]$

SOLUTION:

```
x=linspace(-1,4,500);
y=x.^2-3*x;
plot(x,y);
```

- ii. Compute the variance of data in a vector \mathbf{x} (possibly varying in length). You can't use `var`!

SOLUTION: Assuming x is a column vector (if it isn't, we can make it into one):

```
p=length(x);
Xm=x-mean(x);
varX=(1/(p-1))*Xm'*Xm;
```

- iii. Compute the covariance of data in a vector \mathbf{x} , and \mathbf{y} of the same, but possibly varying length. You can't use `cov`!

SOLUTION:

```
p=length(x);
Xm=x-mean(x); Ym=y-mean(y);
covarXY=(1/(p-1))*Xm'*Y;
```

8. Give mathematical formulas for the sample mean and sample variance. Give the formulas for covariance and correlation.

SOLUTION: Assuming the data is x_1, \dots, x_p , (and y_1, \dots, y_p if necessary):

$$\bar{x} = \frac{1}{p} \sum_{i=1}^p x_i \quad s_x^2 = \frac{1}{p-1} \sum_{i=1}^p (x_i - \bar{x})^2$$

The covariance and correlation:

$$s_{xy}^2 = \frac{1}{p-1} \sum_{i=1}^p (x_i - \bar{x})(y_i - \bar{y})$$

$$r_{xy} = \frac{s_{xy}^2}{s_x s_y} = \frac{\frac{1}{p-1} \sum_{i=1}^p (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\frac{1}{p-1} \sum_{i=1}^p (x_i - \bar{x})^2} \sqrt{\frac{1}{p-1} \sum_{i=1}^p (y_i - \bar{y})^2}}$$

Notice in the last computation, if \mathbf{Xm} and \mathbf{Ym} were vectors with mean subtracted data in them, then the correlation is:

$$\frac{\mathbf{Xm}^T \mathbf{Ym}}{\|\mathbf{Xm}\| \|\mathbf{Ym}\|} = \cos(\theta)$$

where θ is the angle between the vectors.

9. What is the definition of the covariance matrix to X (say that X has p vectors in \mathbb{R}^n , and X is $n \times p$). You can define it by saying what the (i, j) th term of the covariance matrix represents.

SOLUTION: We think of having n rows of data, each with p points. Therefore, the covariance matrix will be $n \times n$, with the (i, j) entry being the covariance between the data in Row i with Row j . Notice that the (i, i) entries will be the variance of the data in row i , and because the covariance of x, y is the same as the covariance between y, x , the matrix will be symmetric.

10. Find the orthogonal projection of the vector $\mathbf{x} = [1, 0, 2]^T$ to the plane defined by:

$$G = \left\{ \alpha_1 \begin{bmatrix} 1 \\ 3 \\ -2 \end{bmatrix} + \alpha_2 \begin{bmatrix} 3 \\ -1 \\ 0 \end{bmatrix} \text{ such that } \alpha_1, \alpha_2 \in \mathbb{R} \right\}$$

Determine the distance from \mathbf{x} to the plane G .

SOLUTION: Note that the vectors defining the plane are orthogonal:

$$\text{Proj}_G(\mathbf{x}) = \frac{1+0-4}{1+9+4} \begin{bmatrix} 1 \\ 3 \\ -2 \end{bmatrix} + \frac{3+0+0}{9+1} \begin{bmatrix} 3 \\ -1 \\ 0 \end{bmatrix} \approx \begin{bmatrix} 0.686 \\ -0.943 \\ 0.429 \end{bmatrix}$$

(Sorry that the numbers didn't work out better- That's why I'm allowing you to have a calculator on the exam). The distance is given by:

$$\|\mathbf{x} - \text{Proj}_G(\mathbf{x})\| \approx 1.86$$

11. If $[\mathbf{x}]_{\mathcal{B}} = (3, -1)^T$, and $\mathcal{B} = \left\{ \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 3 \\ 1 \\ 1 \end{bmatrix} \right\}$, what was \mathbf{x} (in the standard basis)?

SOLUTION: (This question is mostly about what the notation means):

$$\mathbf{x} = 3 \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} - 1 \begin{bmatrix} 3 \\ 1 \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ -1 \\ 2 \end{bmatrix}$$

12. If $\mathbf{x} = (3, -1)^T$, and $\mathcal{B} = \left\{ \begin{bmatrix} 6 \\ 1 \end{bmatrix}, \begin{bmatrix} 1 \\ -2 \end{bmatrix} \right\}$, what is $[\mathbf{x}]_{\mathcal{B}}$?

SOLUTION: Since the columns are not orthogonal, we'll have to do row reduction:

$$\left[\begin{array}{cc|c} 6 & 1 & 3 \\ 1 & -2 & -1 \end{array} \right] \rightarrow \left[\begin{array}{cc|c} 1 & -2 & -1 \\ 6 & 1 & 3 \end{array} \right] \rightarrow \left[\begin{array}{cc|c} 1 & -2 & -1 \\ 0 & 13 & 9 \end{array} \right] \rightarrow \left[\begin{array}{cc|c} 1 & 0 & 5/13 \\ 0 & 1 & 9/13 \end{array} \right]$$

Alternatively, since the columns make up an invertible matrix A :

$$A[x]_B = \mathbf{x} \quad \Rightarrow \quad [x]_B = A^{-1}\mathbf{x} = \frac{-1}{13} \begin{bmatrix} -2 & -1 \\ -1 & 6 \end{bmatrix} \begin{bmatrix} 3 \\ -1 \end{bmatrix} = \begin{bmatrix} 5/13 \\ 9/13 \end{bmatrix}$$

13. Let $\mathbf{a} = [1, 3]^T$. Find a square matrix A so that $A\mathbf{x}$ is the orthogonal projection of \mathbf{x} onto the span of \mathbf{a} .

SOLUTION: The projection formula is the following, which we can write as a function of \mathbf{x} :

$$\text{Proj}_{\mathbf{a}}(\mathbf{x}) = \frac{\mathbf{a}^T \mathbf{x}}{\mathbf{a}^T \mathbf{a}} \mathbf{a} = \mathbf{a} \frac{\mathbf{a}^T \mathbf{x}}{\mathbf{a}^T \mathbf{a}} = \frac{\mathbf{a} \mathbf{a}^T}{\mathbf{a}^T \mathbf{a}} \mathbf{x} = A\mathbf{x}$$

where

$$A = \frac{\mathbf{a} \mathbf{a}^T}{\mathbf{a}^T \mathbf{a}} = \frac{1}{10} \begin{bmatrix} 1 & 3 \\ 3 & 9 \end{bmatrix}$$

14. Determine the projection matrix that takes a vector \mathbf{x} and outputs the projection of \mathbf{x} onto the plane whose normal vector is $[1, 1, 1]^T$.

SOLUTION: Multiple ways of getting a basis for the plane. One way is we want the set (a, b, c) so that the dot product with $[1, 1, 1]^T$ is zero:

$$a + b + c = 0$$

From which we see that we have two free variables, and a basis for the null space:

$$\begin{array}{rcl} a & = & -b - c \\ b & = & b \\ c & = & c \end{array} \quad \Rightarrow \quad b \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} + c \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix}$$

From here, there are two ways to proceed. Method 1 would be to write down the formula for the projection:

$$A(A^T A)^{-1} A^T \quad \text{where } A = \begin{bmatrix} -1 & -1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}$$

If you actually compute this, we end up with:

$$\frac{1}{3} \begin{bmatrix} 2 & -1 & -1 \\ -1 & 2 & -1 \\ -1 & -1 & 2 \end{bmatrix}$$

However, that would take a long time by hand. Alternatively, we could have tried to get an orthonormal set of vectors for the plane. If we call the vectors \mathbf{u}, \mathbf{v} , then we could get:

$$\text{New } \mathbf{v} = \mathbf{v} - \text{Proj}_{\mathbf{u}}(\mathbf{v}) = \begin{bmatrix} -1/2 \\ -1/2 \\ 1 \end{bmatrix}$$

Now let the matrix U be given below. The projection matrix is formed by taking UU^T .

$$U = \begin{bmatrix} -1/\sqrt{2} & -1/\sqrt{6} \\ 0 & -1/\sqrt{6} \\ 1/\sqrt{2} & 2/\sqrt{6} \end{bmatrix}$$

15. Find (by hand) the eigenvectors and eigenvalues of the matrix A :

$$A = \begin{bmatrix} 5 & -1 \\ 3 & 1 \end{bmatrix}, \quad A = \begin{bmatrix} -2 & 1 \\ 1 & -2 \end{bmatrix}$$

SOLUTIONS: For the first matrix, we should find:

$$\lambda_1 = 4, \mathbf{v}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad \lambda_2 = 2, \mathbf{v}_1 = \begin{bmatrix} 1 \\ 3 \end{bmatrix}$$

For the second matrix, we should find:

$$\lambda_1 = -3, \mathbf{v}_1 = \begin{bmatrix} 1 \\ -1 \end{bmatrix}, \quad \lambda_2 = -1, \mathbf{v}_1 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

16. (Referring to the previous exercise) We could've predicted that the eigenvalues of the second matrix would be real, and that the eigenvectors would be orthogonal. Why?

SOLUTION: Because the matrix is symmetric, the Spectral Theorem applies, which says we'll get real eigenvalues and orthogonal eigenvectors.

17. Compute the SVD of the matrix A below.

$$A = \begin{bmatrix} 1 & 0 \\ 2 & 0 \\ 0 & 0 \end{bmatrix}$$

SOLUTION: We see that $A^T A = \begin{bmatrix} 5 & 0 \\ 0 & 0 \end{bmatrix}$ which is already in diagonalized form- We have $\lambda_1 = 5, \mathbf{v}_1 = [1, 0]^T$ and $\lambda_2 = 0$ with $\mathbf{v}_2 = [0, 1]^T$. We also know that

$$\mathbf{u}_1 = \frac{1}{\sigma_1} A^T \mathbf{v}_1 = \frac{1}{\sqrt{5}} \begin{bmatrix} 1 \\ 2 \\ 0 \end{bmatrix}$$

The other vectors in U correspond to $\lambda = 0$, so we'll look for the null space of A^T (or you can look for the null space of AA^T):

$$\begin{bmatrix} 1 & 2 & 0 \\ 0 & 0 & 0 \end{bmatrix} \Rightarrow \begin{array}{lcl} x_1 & = & -2x_2 \\ x_2 & = & x_2 \\ x_3 & = & x_3 \end{array} \Rightarrow \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

Therefore, the matrices U , Σ , V are given by:

$$U = \begin{bmatrix} 1/\sqrt{5} & -2/\sqrt{5} & 0 \\ 2/\sqrt{5} & 1/\sqrt{5} & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \Sigma = \begin{bmatrix} \sqrt{5} & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad V = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

18. Compute the orthogonal projector to the span of \mathbf{x} , if $\mathbf{x} = [1, 1, 1]^T$.

SOLUTION: Whoops- Put in way too many of these, but the formula is the same as always:

$$\frac{\mathbf{x}\mathbf{x}^T}{\mathbf{x}^T\mathbf{x}} = \frac{1}{3} \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

19. Let

$$U = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & 1 \\ 0 & 0 \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} 1 \\ 3 \\ 2 \end{bmatrix}$$

Find $[\mathbf{x}]_U$. Find the projection of \mathbf{x} into the subspace spanned by the columns of U . Find the distance between \mathbf{x} and the subspace spanned by the columns of U .

SOLUTION: Sorry, the first question doesn't make sense, since \mathbf{x} is not in the span of the columns of U . However, we can project \mathbf{x} into the column space of U , which is the $x - y$ plane: $[1, 3, 0]^T$. If you didn't recognize that, the computations should have given you the right answer anyway. The distance is therefore 2.

20. Show that $\text{Null}(A) \perp \text{Row}(A)$.

SOLUTION: We'll let \mathbf{x} be an arbitrary vector in the null space of A , and then we'll show that \mathbf{x} is orthogonal to all the rows of A .

If we let $\mathbf{r}_1, \dots, \mathbf{r}_m$ denote the rows of A , then if \mathbf{x} is in the null space:

$$A\mathbf{x} = \mathbf{0} \quad \Rightarrow \quad \begin{bmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \\ \vdots \\ \mathbf{r}_m \end{bmatrix} \mathbf{x} = \mathbf{0} \quad \Rightarrow \quad \mathbf{r}_i \mathbf{x} = 0$$

for each i . Therefore, the dot product between \mathbf{x} and every row of A is zero, so \mathbf{x} is perpendicular to the row space.

21. Show that, if X is invertible, then $X^{-1}AX$ and A have the same eigenvalues.

SOLUTION: Whenever we solve things about eigenvalues, we have three equations to choose from:

$$A\mathbf{v} = \lambda\mathbf{v} \quad (A - \lambda I)\mathbf{v} = \mathbf{0} \quad \det(A - \lambda I) = 0$$

In this case, it might be easiest to use the derivative, since the derivative of a product is the product of the derivatives:

$$\begin{aligned}\det(X^{-1}AX - \lambda I) &= \det(X^{-1}AX - \lambda X^{-1}X) = \det(X^{-1}(A - \lambda I)X) = \\ &= \det(X^{-1})\det(A - \lambda I)\det(X) = \det(A - \lambda I)\end{aligned}$$

since $\det(X^{-1})\det(X) = 1$. Therefore, the characteristic equation for $X^{-1}AX$ and A are the same (so they'll have the same eigenvalues).

22. How do we “double-center” a matrix of data?

SOLUTION: There were several ways to do this.

Method 1: (From the Matlab script in the notes):

- Find the mean of the rows of A .
- Subtract the mean from each of the rows.
- Find the mean of columns of the matrix from the last step.
- Subtract the mean from each of the columns.

Method 2: (Again from a Matlab script in the notes):

Same instructions as before, but do the columns first, then the rows.

Method 3: This was done in class.

- Find the mean of the rows, \bar{r} , the mean of the columns, \bar{c} and the grand mean, g of matrix A (in any order).
- Subtract \bar{r} from each row of A .
- Subtract \bar{c} from each column of the previous step.
- Add back in the grand mean to all values of the matrix in the previous step.

23. True or False, and give a short reason:

- (a) If the rank of A is 3, the dimension of the row space is 3.

SOLUTION: True. If the rank is 3, the dimension of the column space is 3, and the dimensions of the row space and column space are equal.

- (b) If the correlation coefficient between two sets of data is 1, then the data sets are the same.

SOLUTION: False. It means that there is a linear relationship between the data sets (positive would give positive slope).

- (c) If the correlation coefficient between two sets of data is 0, then there is no functional relationship between the two sets of data.

SOLUTION: False. It means there is no *linear* relationship between the data sets. There could be a nonlinear relationship.

- (d) If U is a 4×2 matrix, then $U^T U = I$.

SOLUTION: If U has orthonormal columns, then this is correct (I think we were supposed to assume that).

- (e) If U is a 4×2 matrix, then $U U^T = I$.

SOLUTION: Even if U has orthonormal columns, this is generally false. The matrix $U U^T$ would project a vector \mathbf{x} into the column space of U . Note that if \mathbf{x} is in the column space already, then it would be true that $U U^T \mathbf{x} = \mathbf{x}$.

- (f) If A is not invertible, then $\lambda = 0$ is an eigenvalue of A .

SOLUTION: True, if A is $n \times n$. If A is not invertible, then $\det(A) = 0$, or

$$\det(A - 0 I) = 0$$

is true.

- (g) Let

$$A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 2 & 0 \end{bmatrix}$$

Then the rank of $A A^T$ is 2.

SOLUTION: If A has rank 2, then both $A A^T$ and $A^T A$ also have rank 2. We used this in the construction of the SVD.

24. Let $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ be the normalized eigenvectors of $A^T A$, where A is $m \times n$.

- (a) Show that if λ_i is a non-zero eigenvalue of $A^T A$, then it is also a non-zero eigenvalue of $A A^T$.

SOLUTION:

$$A^T A \mathbf{v}_i = \lambda_i \mathbf{v}_i \quad \Rightarrow \quad A A^T A \mathbf{v}_i = \lambda_i A \mathbf{v}_i$$

We note here that since $\lambda_i \neq 0$, then $A \mathbf{v}_i \neq \mathbf{0}$. We then see that $\mathbf{u}_i = A \mathbf{v}_i$, with

$$A A^T \mathbf{u}_i = \lambda_i \mathbf{u}_i$$

so λ_i is a non-zero eigenvalue of $A A^T$.

- (b) True or false? The eigenvectors form an orthogonal basis of \mathbb{R}^n .

SOLUTION: The set of eigenvectors of $A^T A$ do form an orthogonal basis of \mathbb{R}^n , which is a result of the Spectral Theorem (since $A^T A$ is symmetric).

- (c) Show that, if $\mathbf{x} \in \mathbb{R}^n$, then the i^{th} coordinate of \mathbf{x} (with respect to the eigenvector basis) is $\mathbf{x}^T \mathbf{v}_i$.

SOLUTION: If we write

$$\mathbf{x} = \alpha_1 \mathbf{v}_1 + \cdots + \alpha_n \mathbf{v}_n$$

Take the dot product of both sides with \mathbf{v}_i to get:

$$\mathbf{v}_i^T \mathbf{x} = \alpha_i \mathbf{v}_i^T \mathbf{v}_i$$

And since the vectors have length 1, this simplifies to give us our result.

- (d) Let $\alpha_1, \dots, \alpha_n$ be the coordinates of \mathbf{x} with respect to $\mathbf{v}_1, \dots, \mathbf{v}_n$.

Show that

$$\|\mathbf{x}\|_2 = \alpha_1^2 + \alpha_2^2 + \cdots + \alpha_n^2$$

I'll allow you to show it just using just two vectors, $\mathbf{v}_1, \mathbf{v}_2$.

SOLUTION: Using two vectors,

$$\begin{aligned} \mathbf{x}^T \mathbf{x} &= (\alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2)^T (\alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2) = (\alpha_1 \mathbf{v}_1^T + \alpha_2 \mathbf{v}_2^T) (\alpha_1 \mathbf{v}_1 + \alpha_2 \mathbf{v}_2) = \\ &\alpha_1^2 \mathbf{v}_1^T \mathbf{v}_1 + \alpha_1 \alpha_2 \mathbf{v}_1^T \mathbf{v}_2 + \alpha_2 \alpha_1 \mathbf{v}_2^T \mathbf{v}_1 + \alpha_2^2 \mathbf{v}_2^T \mathbf{v}_2 \end{aligned}$$

Since the vectors are orthonormal, this simplifies to:

$$\|\mathbf{x}\|^2 = \alpha_1^2 + \alpha_2^2$$

- (e) Show that $A\mathbf{v}_i \perp A\mathbf{v}_j$

SOLUTION: To show this, we show that the dot product is zero. Note that we already know that the \mathbf{v}' s are orthonormal.

$$A\mathbf{v}_i \cdot A\mathbf{v}_j = (A\mathbf{v}_i)^T A\mathbf{v}_j = \mathbf{v}_i^T A^T A\mathbf{v}_j = \lambda_j \mathbf{v}_i^T \mathbf{v}_j = 0$$

- (f) Show that $A\mathbf{v}_i$ is an eigenvector of AA^T .

SOLUTION: This is basically the same proof as (a):

$$AA^T(A\mathbf{v}_i) = A(A^T A\mathbf{v}_i) = A\lambda_i \mathbf{v}_i = \lambda_i(A\mathbf{v}_i)$$

25. Show that, for the line of best fit, the normal equations produce the same equations as minimizing an appropriate error function. To be more specific, set the data as $(x_1, t_1), \dots, (x_p, t_p)$ and define the error function first. Minimize the error function to find the system of equations in m, b . Show this system is the same you get using the normal equations.

SOLUTION: We did this in class, and also on pg 38-39 of Chapter 3 (the last two pages of the chapter).

26. Given data:

$$\begin{array}{c|ccc} x & -1 & 0 & 1 \\ \hline y & 2 & 1 & 1 \end{array}$$

(a) Give the matrix equation for the *line of best fit*.

$$\begin{bmatrix} -1 & 1 \\ 0 & 1 \\ 1 & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}$$

(b) Compute the normal equations.

SOLUTION: Given $A\mathbf{x} = \mathbf{b}$, the normal equations are $A^T A\mathbf{x} = A^T \mathbf{b}$. In this case,

$$\begin{bmatrix} 2 & 0 \\ 0 & 3 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} -1 \\ 4 \end{bmatrix}$$

(c) Solve the normal equations for the slope and intercept.

SOLUTION: For the line: $y = mx + b$, we get $m = -1/2$ and $b = 4/3$.

27. Use the data in Exercise (26) to find the parabola of best fit: $y = ax^2 + bx + c$. (NOTE: Will you only get a least squares solution, or an actual solution to the appropriate matrix equation?)

SOLUTION: Plug the data into our equation:

$$\begin{array}{rcl} (-1)^2 a + (-1)b + c & = & 2 \\ (0)^2 a + (0)b + c & = & 1 \\ (1)^2 a + (1)b + c & = & 1 \end{array} \Rightarrow \begin{bmatrix} 1 & -1 & 1 \\ 0 & 0 & 1 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 2 \\ 1 \\ 1 \end{bmatrix}$$

The matrix is invertible, so there is exactly one solution to this equation. The parabola will go through all of the data points.

28. Let $\mathbf{x} = [1, 2, 1]^T$. Find the matrix $\mathbf{x}\mathbf{x}^T$, its eigenvalues, and eigenvectors. (Also, think about what happens in the general case, where a matrix is defined by $\mathbf{x}\mathbf{x}^T$). HINT: SVD

SOLUTION: If you're not sure how to start, the hint was there to get you to think about the eigenvalues of $\mathbf{x}\mathbf{x}^T$ versus $\mathbf{x}^T \mathbf{x}$ (they are the same). Therefore, the non-zero eigenvalue is $\lambda_1 = \|\mathbf{x}\|^2$, which we prove below:

$$(\mathbf{x}\mathbf{x}^T)\mathbf{x} = \mathbf{x}(\mathbf{x}^T \mathbf{x}) = \|\mathbf{x}\|^2 \mathbf{x}$$

so therefore, the non-zero eigenvalue is $\lambda_1 = \|\mathbf{x}\|^2$ and the eigenvector is \mathbf{x} (which we normalize).

Now, if \mathbf{x} is in \mathbb{R}^3 , then the other two eigenvalues are both zero, and we just try to find two vectors orthogonal to \mathbf{x} . We can also find the null space of $\mathbf{x}\mathbf{x}^T$:

$$\begin{bmatrix} 1 & 2 & 1 \\ 2 & 4 & 2 \\ 1 & 2 & 1 \end{bmatrix} \rightarrow \begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \Rightarrow \left\{ \begin{bmatrix} -2 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} \right\}$$

NOTICE: That the vectors we found for E_0 were both orthogonal to the original vector, but that they are not orthogonal to each other. That's because we can find an infinite number of vectors for the two dimensional eigenspace- If we wanted orthogonal vectors, we could get them. For example, in place of $\mathbf{v}_3 = [-1, 0, 1]$, we could take:

$$\text{New } \mathbf{v}_3 = \mathbf{v}_3 - \text{Proj}_{\mathbf{v}_2}(\mathbf{v}_3)$$

For the in-class exam, I won't ask you to do that.

29. Suppose \mathbf{x} is a vector containing n real numbers, and we understand that $m\mathbf{x} + b$ is Matlab-style notation (so we can add a vector to a scalar, done component-wise).

- (a) Find the mean of $\mathbf{y} = m\mathbf{x} + b$ in terms of the mean of \mathbf{x} .

SOLUTION: Let \bar{x} be the mean of x_1, x_2, \dots, x_n . We want to find the mean of

$$\{mx_1 + b, mx_2 + b, \dots, mx_n + b\}$$

The mean is:

$$\frac{1}{p} \sum_{i=1}^p mx_i + b = m \left(\frac{1}{p} \sum_{i=1}^p x_i \right) + b \left(\frac{1}{p} \sum_{i=1}^p 1 \right) = m\bar{x} + b$$

- (b) Show that, for fixed constants a, b , $\text{Cov}(\mathbf{x} + a, \mathbf{y} + b) = \text{Cov}(\mathbf{x}, \mathbf{y})$

SOLUTION: Given that our data looks like the following, mean subtract it. The mean for \mathbf{x} is \bar{x} (same for y). The data looks like:

$$\{x_1 + a, x_2 + a, \dots, x_n + a\}$$

Subtract the mean:

$$\{(x_1 + a) - (\bar{x} + a), (x_2 + a) - (\bar{x} + a), \dots, (x_n + a) - (\bar{x} + a)\}$$

That leaves us with the original data set, so the covariance of this set is the same as the original:

$$\{x_1 - \bar{x}, x_2 - \bar{x}, \dots, x_n - \bar{x}\}$$

This is the original mean subtracted data. The same thing happens in \mathbf{y} , so the covariance will be the same as the original.

- (c) If $\mathbf{y} = m\mathbf{x} + b$, then find the covariance and correlation coefficient between \mathbf{x} and \mathbf{y} .

SOLUTION: We find that the covariance is m times the variance:

$$S_{xy}^2 = \text{Cov}(\mathbf{x}, m\mathbf{x} + b) = \frac{1}{p-1} \sum_{i=1}^p (x_i - \bar{x})(m(x_i - \bar{x})) = mS_x^2$$

In the formula for the correlation, we'll also need the variance of \mathbf{y} :

$$S_y^2 = \frac{1}{p-1} \sum_{i=1}^p (m(x_i - \bar{x}))^2 = m^2 S_x^2$$

Now, the correlation is given by:

$$r_{xy} = \frac{S_{xy}^2}{S_x S_y} = \frac{m S_x^2}{S_x |m| S_x} = \frac{m}{|m|}$$

Notice that

$$\frac{m}{|m|} = \begin{cases} 1 & \text{if } m > 0 \\ -1 & \text{if } m < 0 \end{cases}$$

30. Suppose we have a subspace W spanned by an orthonormal set of non-zero vectors, $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$, each is in \mathbb{R}^{1000} . If a vector \mathbf{x} is in W , then there is a low dimensional (three dimensional in fact) representation of \mathbf{x} . What is it?

SOLUTION: The low dimensional representation is the representation in \mathbb{R}^3 given by coordinates of \mathbf{x} :

$$[\mathbf{x}]_V = V^T \mathbf{x} = (\mathbf{v}_1^T \mathbf{x}, \mathbf{v}_2^T \mathbf{x}, \mathbf{v}_3^T \mathbf{x})$$

31. Let the matrix A be defined below.

$$A = \begin{bmatrix} 1 & 1 \\ 2 & 1 \\ 3 & 1 \end{bmatrix}$$

TYPO: The numerical values of the SVD of A are pretty bad. We'll answer this in the abstract.

- (a) Find the psuedoinverse of A

SOLUTION: Let the reduced SVD of A be given by $A = \hat{U} \hat{\Sigma} \hat{V}^T$. Then the pseudo inverse of A is given by:

$$A^\dagger = \hat{V} \hat{\Sigma}^{-1} \hat{U}^T$$

- (b) Using the A from the previous exercise, consider the vector $[-1, 0, 1]^T$. Is the vector in the column space of A ? If so, provide its coordinates with respect to the columns of A (for the basis).

SOLUTION: If we row reduce the augmented matrix, if the system is consistent, then the vector is in the subspace (it turns out that this is true, and the coordinates are $1, -2$).

- (c) What happens if we try to project $[1, -2, 1]^T$ into the column space of A ? Explain in terms of fundamental subspaces.

SOLUTION: This question doesn't have enough information in it. If I add the fact that the given vector \mathbf{x} is in the null space of A^T , does that help?

Yes- The projection is to the origin.

32. Consider the underdetermined "system of equations": $x+3y+4z = 1$. In matrix-vector form $A\mathbf{x} = \mathbf{b}$, write the matrix A first.

- (a) What is the dimension of each of the four fundamental subspaces?

SOLUTION: Since A is 1×3 , the domain is \mathbb{R}^3 and the range is \mathbb{R} . Since we have only one row, the dimension of the row space is 1 (as is the dimension of the column space). The dimension of the null space of A is 2, and the dimension of the null space of A^T is zero (meaning zero is the only vector in that space).

- (b) Find bases for the four fundamental subspaces.

SOLUTION:

$$\text{Row}(A) = \text{span} \left\{ \begin{bmatrix} 1 \\ 3 \\ 4 \end{bmatrix} \right\} \quad \text{Null}(A) = \text{span} \left\{ \begin{bmatrix} -3 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} -4 \\ 0 \\ 1 \end{bmatrix} \right\}$$

The column space is spanned by the "vector": $\begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$ (and the null space of A^T is 0).

- (c) Find a solution with at least 2 zeros (the slash command in Matlab looks for answers with the most zeros).

SOLUTION: The set of all solutions proceeds by finding one solution to $x + 3y + 4z = 1$, then add on the vectors from the null space. Therefore, all solutions:

$$\mathbf{x} = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + s \begin{bmatrix} -3 \\ 1 \\ 0 \end{bmatrix} + t \begin{bmatrix} -4 \\ 0 \\ 1 \end{bmatrix}$$

In this case, one solution with 2 zeros would be $[1, 0, 0]^T$.

- (d) Find a 3×3 matrix P so that given a vector \mathbf{x} , $P\mathbf{x}$ is the projection of \mathbf{x} into the row space of A .

SOLUTION: I've asked this multiple times already- The solution is to construct the matrix:

$$\frac{\mathbf{xx}^T}{\mathbf{x}^T \mathbf{x}}$$

33. (SVD) Given that the SVD of a matrix was given in Matlab as:

```
>> [U,S,V]=svd(A)
U =
   -0.4346   -0.3010    0.7745    0.3326   -0.1000
   -0.1933   -0.3934    0.1103   -0.8886   -0.0777
    0.5484    0.5071    0.6045   -0.2605   -0.0944
    0.6715   -0.6841    0.0061    0.1770   -0.2231
    0.1488   -0.1720    0.1502   -0.0217    0.9619
S =
    5.72         0         0
         0     2.89         0
         0         0         0
         0         0         0
         0         0         0
V =
    0.2321   -0.9483    0.2166
   -0.2770    0.1490    0.9493
    0.9324    0.2803    0.2281
```

- (a) Which columns form a basis for the null space of A ? For the column space of A ? For the row space of A ?

SOLUTION: From the matrix S , we can surmise that A is 5×3 . The null space and row space are subspaces of \mathbb{R}^3 , and the rank is 2 (the number of non-zero singular values). Therefore, the null space is spanned by the last column of V , and the column space is spanned by the first 2 columns of U . The row space is spanned by the first 2 columns of V .

- (b) How do we “normalize” the singular values? In this case, what are they (numerically)?

SOLUTION: $\frac{5.72}{5.72+2.89}, \frac{2.89}{5.72+2.89}$, which is $[0.6643, 0.3357]$.

- (c) What is the rank of A ?

SOLUTION: The rank is 2.

- (d) How would you compute the pseudo-inverse of A (do not actually do it):

SOLUTION: Use the reduced form of the SVD, which \hat{U} formed from the first two columns of U , $\hat{\Sigma}$ is the 2×2 diagonal matrix, and \hat{V} is from the first two columns of V . Given that, the pseudo-inverse is given below:

$$A^\dagger = \hat{V} \hat{\Sigma}^{-1} \hat{U}^T$$

- (e) Let B be formed using the first two columns of U . Would the matrix $B^T B$ have any special meaning? Would BB^T ?

SOLUTION: $B^T B$ would be a 2×2 identity matrix. BB^T is the projection matrix to the first two columns of U .

34. Suppose $Q = [-0.5, 0, 0.5, 1.0]$. Use the softmax selection technique with $\tau = 0.1$ to compute the probabilities.

SOLUTION:

$$e^{10Q} = [e^{-5}, 1, e^5, e^{10}]$$

When we make these probabilities by dividing by the sum, the first two values are basically zero. The next one is about 0.01 and the last is about 0.99.

35. If $Q_1 < Q_2 < Q_3 < Q_4$ for 4 machines, how do the probabilities change (under softmax) as $\tau \rightarrow 0$? As $\tau \rightarrow 1$?

SOLUTION: TYPO: The upper limit should be ∞ , not 1. Let's compute the probabilities (we've done this before for two numbers).

$$P_1 = \lim_{\tau \rightarrow 0} \frac{e^{Q_1/\tau}}{e^{Q_1/\tau} + e^{Q_2/\tau} + e^{Q_3/\tau} + e^{Q_4/\tau}}$$

Dividing by $e^{Q_1/\tau}$, this biggest probability becomes:

$$P_1 = \lim_{\tau \rightarrow 0} \frac{1}{1 + e^{(Q_2-Q_1)/\tau} + e^{(Q_3-Q_1)/\tau} + e^{(Q_4-Q_1)/\tau}} = \frac{1}{1 + 0 + 0 + 0} = 1$$

(These expressions all go to zero since $Q_i - Q_1$ is negative). For the others, we do a similar computation and divide by the same amount. For example, the second probability becomes:

$$P_2 = \lim_{\tau \rightarrow 0} \frac{e^{(Q_2-Q_1)/\tau}}{1 + e^{(Q_2-Q_1)/\tau} + e^{(Q_3-Q_1)/\tau} + e^{(Q_4-Q_1)/\tau}} = \frac{0}{1 + 0 + 0 + 0} = 0$$

Now going in the other direction, $\tau \rightarrow \infty$, we note that $Q_i/\tau \rightarrow 0$, so that

$$P_1 = \lim_{\tau \rightarrow \infty} \frac{e^{Q_1/\tau}}{e^{Q_1/\tau} + e^{Q_2/\tau} + e^{Q_3/\tau} + e^{Q_4/\tau}} = \frac{1}{1 + 1 + 1 + 1} = \frac{1}{4}$$

And the others are calculated the same way.

36. What is the win-stay, lose-shift (or pursuit) strategy? What are the update rules?

SOLUTION: The idea behind the win-stay, lose-shift strategy is that, if we "win" with the currently selected machine, then the probability of choosing that machine should increase (and the rest decrease). For example, if we won with machine a , then

$$(P_a)_{\text{new}} = P_a + \beta(1 - P_a)$$

and all the remaining probabilities are decreased by the same amount:

$$(P_j)_{\text{new}} = P_j + \beta(0 - P_j)$$

37. Suppose we play with three machines, and machine 3 is chosen and gives a big payout (enough to make $Q_t(3)$ the maximum). Update the probabilities for win-stay, lose-shift, if they are: $P_1 = 0.3, P_2 = 0.5, P_3 = 0.2$ and $\beta = 0.3$.

SOLUTION: P_3 increases, the rest decrease:

$$(P_3)_{\text{new}} = 0.2 + 0.3(1 - 0.2) = 0.44$$

$$(P_1)_{\text{new}} = 0.3 + 0.3(0 - 0.3) = 0.21$$

$$(P_2)_{\text{new}} = 0.5 + 0.3(0 - 0.5) = 0.35$$

(Note that they still sum to 1).

38. Suppose we have a genetic algorithm with 4 chromosomes, and current fitness values $[-1, 1, 2, 3]$. We want to construct the probability of choosing these chromosomes for mating. Calculate the probabilities (if possible) using the current ordering, if we use:

- (a) Normalized fitness values: SOLUTION: Since the first value is negative, we cannot use the normalized fitness values.
- (b) Rank order, normalized: SOLUTION: The order is already from smallest to largest, so the probabilities (in the given ordering) is:

$$\frac{1}{10}, \frac{2}{10}, \frac{3}{10}, \frac{4}{10}$$

- (c) Softmax with $\tau = 1$. SOLUTION: Use a calculator to get the approximate values:

$$[0.012, 0.089, 0.242, 0.656]$$