1 Basic Computations

It's good to know how vectors and matrix vector products and derivatives work.

- $\begin{bmatrix}1\\2\end{bmatrix} + \begin{bmatrix}3\\2\end{bmatrix} =$
- $\begin{bmatrix}1\\2\\3\end{bmatrix} + \begin{bmatrix}2\\3\\1\end{bmatrix} + \begin{bmatrix}4\\2\\0\end{bmatrix} =$
- $\begin{bmatrix}1 & 1\\1 & 0\end{bmatrix} \cdot \begin{bmatrix}1 & 2\\3 & 4\end{bmatrix} =$
- $\begin{bmatrix}1 & 2\\3 & 4\end{bmatrix} \cdot \begin{bmatrix}1 & 1\\1 & 0\end{bmatrix} =$
- $\frac{\partial}{\partial x}(2 + 4x^2 + e^x) =$
- For vector $\mathbf{x} = [x_1, x_2, x_3]^\top \in \mathbb{R}^{3\times 1}$, $\frac{\partial}{\partial x_2}(45^2 x_1 + x_2^2 - 2025x_3) =$

2 Linear Transformations

(a) Compute the column and null space of the linear transformation $A = \begin{bmatrix}1 & 2 & 3\\-1 & 4 & 2\\0 & 6 & 5\end{bmatrix}$.

Express your answer as span of some vectors.

(b) For two linear transformations T_1 and T_2, is $T_1(T_2(\mathbf{v})) = T_2(T_1(\mathbf{v}))$ always true for all $\mathbf{v}$?
Explain why, and assume there are no issues with domain/range stuff.

(c) If two linear transformations T_1 and T_2 satisfy $T_1(T_2(\mathbf{v})) = 0$ for all $\mathbf{v}$, does one of T_1 or T_2 have to be the linear transformation that maps all vectors to $\mathbf{0}$?
Again, assume there are no issues with domain/range stuff.

3 Least Squares, Projection

(a) Compute x such that the L2 norm $\|Ax - b\|_2$ is minimized, where $A = \begin{bmatrix} 1 & 2 & 0 \\ -1 & 4 & 6 \\ 1 & 2 & 0 \end{bmatrix}$, $b = \begin{bmatrix} 3 \\ -1 \\ 5 \end{bmatrix}$.

(b) Using the previous part, find the projection of b onto the plane spanned by $v_1 = \begin{bmatrix} 1 \\ -1 \\ 1 \end{bmatrix}$ and $v_2 = \begin{bmatrix} 2 \\ 4 \\ 2 \end{bmatrix}$.

(c) From the above, what is the distance from b to $\text{span}(\{v_1, v_2\})$?

4 Perceptron

You want to predict if Saathvik will like the recommended anime. You get two recommendations from respected otaku, George and Henry, to read a script you have and rate it on a scale of 1 to 4. The critics are not perfect; here are five data points including the critics' scores and the performance of the anime:

#	Anime Name	G	H	Like?
1	Dragon Ball Z	1	1	-
2	Death Note	3	2	+
3	Sword Art Online	2	4	+
4	One Piece	3	4	+
5	Hunter x Hunter	2	3	-

1. First, you would like to examine the linear separability of the data. Plot the data on a 2D plane; label favorable anime with '+' and non-favorable anime with '-', and determine if the data are linearly separable.

2. Now you decide to use a perceptron to classify your data. Suppose you directly use the scores given above as features, together with a bias feature. That is $f_0 = 1$, $f_1 =$ score given by George, and $f_2 =$ score given by Henry.

   Run one pass through the data with the perceptron algorithm, filling out the table below. Go through the data points in order, e.g., using data point #1 at step 1.

Step	Weights	Score	Correct?
1	[-1, 0, 0]	$-1 \cdot 1 + 0 \cdot 1 + 0 \cdot 1 = -1$	Yes
2			
3			
4			
5			

   Final weights:

3. Have weights been learned that separate the data?

4. More generally, irrespective of the training data, you want to know if your features are powerful enough to allow you to handle a range of scenarios. Circle the scenarios for which a perceptron using the features above can indeed perfectly classify anime that are favorable according to the given rules:

   (a) Your reviewers are awesome: if the total of their scores is more than 8, then the anime will definitely be favorable, and otherwise it won't be.

   (b) Your reviewers are horror critics. Your anime will be favorable if and only if each reviewer gives either a score of 2 or a score of 3.

   (c) Your reviewers have weird but different tastes. Your anime will be profitable if and only if both reviewers agree.

(Credits to CS 188 Fall 2024 for this question.)