

Discussion Session Problems 4

02/5/2026

1. Here are a few types of gradient descent updates:

- (a) **Stochastic Gradient Descent:** Uses only single training example to calculate the gradient and update parameters.
- (b) **Batch Gradient Descent:** Calculate the gradients for the whole dataset and perform just one update at each iteration.
- (c) **Mini-batch Gradient Descent:** Mini-batch gradient is a variation of stochastic gradient descent where instead of single training example, mini-batch of samples is used. It's one of the most popular optimization algorithms.

What are the pros and cons of each one and when should you use them? Walk us through your intuition and give us examples of when you are going to use what update.

2. Our internal AI team has been training a model for very long and our biggest problem is that our model does not learn anything. Our cost function is not going down, and our gradients in a layer seem to be the same. We have been debugging our gradient update function but it seems like it is correct. Do you have any tips to help us solve this problem? Walk us through your thought process and show us on the board what is going on.
3. You are asked to build a classification model about meteorites impact with Earth (important project for human civilization). After preliminary analysis, you get 99% accuracy. Should you be happy? Why not? What can you do about it?
4. You are working on a classification problem. For validation purposes, you've randomly sampled the training data set into train and validation. You are confident that your model will work incredibly well on unseen data since your validation accuracy is high. However, you get shocked after getting poor test accuracy. What went wrong?
5. Why should features be normalized when applying regularization to logistic regression and neural nets? Any further reasons to do it for neural nets even if we are not using regularization?
6. Let's explore saliency, a measure of how important a feature is for classification. We define the saliency of the i th input feature for a given example (x, y) to be the absolute value of the partial derivative of the log likelihood of the sample prediction, with respect to that input feature $\left| \frac{\partial \text{LL}}{\partial x_i} \right|$. In the images below, we show both input images and the corresponding saliency of the input features (in this case, input features are pixels):



First consider a trained logistic regression classifier with weights θ . Like the logistic regression classifier that you wrote in your homework it predicts binary class labels. In this question we allow the values of x to be real numbers, which doesn't change the algorithm (neither training nor testing).

- (a) What is the log likelihood of a single training example (x, y) ?
- (b) Compute the saliency of feature x_i .
- (c) Show that the ratio of saliency for features i and j equals $\frac{|\theta_i|}{|\theta_j|}$.

7. Suppose a hidden layer uses ReLU activation:

$$A = \max(0, Z)$$

and during backprop you receive $\frac{\partial L}{\partial A} = dA$. Write the vectorized expression for $\frac{\partial L}{\partial Z}$.


8. Let:

$$X : (m \times n), \quad W_1 : (n \times h), \quad W_2 : (h \times K)$$

Give the shapes of:

- (a) $Z_1 = XW_1$
- (b) $A_1 = \text{ReLU}(Z_1)$
- (c) $Z_2 = A_1W_2$
- (d) $\hat{Y} = \text{softmax}(Z_2)$
- (e) $\frac{\partial L}{\partial W_2}$


9. The figure below shows decision regions of a 3-class classifier in 2D. Each color represents a predicted class.



Answer the following:

- (a) Is this classifier linear or non-linear? Explain.
- (b) Could this be produced by a single linear layer? Why or why not?
- (c) What change to a model could produce curved boundaries like these?

10. In the figure below, three linear decision boundaries divide the plane into regions for classes A and B. Training points from both class lie deep inside the region predicted as the opposite class.



- (a) What does this tell you about the model's capacity?
- (b) Would adding more training time fix this issue?
- (c) Name one model change that could help.