

CS 589 Midterm review

Summer 2025

Gehao Zhang

k -Nearest Neighbor (k -NN) Algorithm

- Supervised Learning Overview:
 - Collect a dataset of images and labels.
 - Use Machine Learning to train an image classifier.
 - Evaluate the classifier on a withheld set of test images.
- k -Nearest Neighbor (k -NN) Classifier:
 - Process:
 - For each test image, find the k nearest images in the training set.
 - The labels of these k nearest images are used to determine the label of the test image.
 - Example Dataset: CIFAR-10 with 10 labels, 50,000 training images, and 10,000 test images.
 - Distance Metrics:
 - Common choices are L_1 (Manhattan) and L_2 (Euclidean) distances.
 - Hyperparameters:
 - Important to choose the right distance metric and the value of k .
 - Cross-validation is used to tune these hyperparameters effectively.

Linear Classification

- Linear Classification:
 - Parametric Approach:
 - Images are represented as arrays of numbers.
 - Use a linear function to map these arrays to class scores.
 - Example Calculation:
 - Given weights (W) and image data, calculate class scores.
 - Use these scores in a loss function to optimize W .
- Loss Functions:
 - Multiclass SVM Loss:
 - Calculate the loss based on the margin between the correct class score and other class scores.
 - Softmax Classifier:
 - Use the softmax function to convert scores to probabilities.
 - Minimize the negative log likelihood of the correct class.
- Regularization:
 - Prevent the model from having too much flexibility.
 - Helps avoid overfitting by adding a penalty term to the loss function.

Example Question with Answer

- **Question:** Given the following training examples and their scores calculated with some weights W :

Example	Class Label	Scores (Cat, Frog, Car)
1	Cat	(3.2, 5.1, -1.7)
2	Frog	(4.9, 1.3, 2.0)
3	Car	(-3.1, 2.5, 2.2)

1. Calculate the Multiclass SVM loss for each example.
2. Compute the total training loss as the mean of the individual losses.

Example Question with Answer

- Answer:
 - Multiclass SVM Loss Calculation:
 - Example 1 (cat):
 - $Loss_1 = \sum_{j \neq y_i} \max(0, s_j - s_{y_i} + 1)$
 $= \max(0, 3.1 - 3.2 + 1) + \max(0, -1.7 - 3.2 + 1) = 2.9$
 - Example 2 (Frog):
 - $Loss_2 = \max(0, 4.9 - 1.3 + 1) + \max(0, 2.0 - 1.3 + 1) = 5.3$
 - Example 3 (Car):
 - $Loss_3 = \max(0, -3.1 - 2.2 + 1) + \max(0, 2.5 - 2.2 + 1) = 1.3$
 - Total Training Loss:
 - $L = \frac{1}{N} \sum_{i=1}^N Loss_i = \frac{1}{3} (2.9 + 5.3 + 1.3) = 3.17$

Neural Networks

- Question: You are training a neural network with the following configuration:
- Input layer with 3072 neurons (corresponding to $32 \times 32 \times 3$ images).
- Hidden layer with 100 neurons using ReLU activation.
- Output layer with 10 neurons (one for each class).
- Explain the importance of using an activation function like ReLU in the hidden layer.
- Describe what could happen if we use a zero-centered activation function, such as tanh, in the hidden layer.
- Given a small batch of activations at a layer during training, describe how Batch Normalization would be applied.

Neural Networks

Answer:

1. Importance of ReLU Activation:

- ReLU (Rectified Linear Unit) computes $\max(0, x)$, which introduces non-linearity into the network. This non-linearity allows the network to learn complex patterns and decision boundaries that a linear classifier cannot. ReLU is computationally efficient and helps networks converge faster compared to sigmoid or tanh functions.

2. Using Tanh Activation:

- Tanh squashes input values to the range $[-1, 1]$, making it zero-centered, which is beneficial for having balanced gradients. However, tanh can still suffer from the vanishing gradient problem when neurons become saturated, leading to very small gradients and slow learning.

Neural Networks

Answer:

3. Applying Batch Normalization:

- Batch Normalization normalizes the activations of each layer by scaling and shifting them to have zero mean and unit variance. This is done by computing the empirical mean and variance of the activations for each batch and then normalizing each activation:
- $$y = \gamma \frac{x - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta$$
- where μ is the mean, σ^2 is the variance of the batch, and ϵ is a small constant to prevent division by zero.
- This normalization helps improve gradient flow, allows for higher learning rates, and acts as a regularizer, reducing the need for dropout.

Decision Trees

- Overview:
 - Definition: A decision tree is a model that extracts classification rules to classify instances using IF-ELSE statements.
 - Interpretability: Highly interpretable by humans, mimics human decision-making.
- Key Concepts:
 - Components:
 - Root node
 - Decision nodes
 - Leaf/terminal nodes
 - Process:
 - Perform tests on attributes of an instance.
 - Make a decision/prediction about the class of the instance.

Decision Trees

- Information Gain:
 - Entropy: Measures the homogeneity of a set of instances.
 - Calculation:
 - Entropy: $I(D) = - \sum p_i \log_2 p_i$
 - Information Gain: $Gain_A(D) = I(D) - Info_A(D)$
 - Example: Testing different attributes like weather conditions in predicting if a person will play tennis.
- Information Gain Ratio:
 - Adjusts Information Gain to account for attributes with many possible values.
 - $GainRatio(A, D) = \frac{Gain_A(D)}{SplitInfo(A)}$
- Gini Impurity:
 - Definition: Measures the impurity of a dataset.
 - $Gini(D) = 1 - \sum(p_i^2)$
 - Example: Testing attributes like weather to determine the Gini impurity for predicting if a person will play tennis.

Decision Trees

- Question: Given the following dataset for predicting whether a person will buy a computer based on three attributes: Student, Age, and Credit_Score, construct a decision tree using the ID3 algorithm and calculate the Information Gain for each attribute. Determine which attribute should be tested first.

Instance	Student	Age	CreditScore	WillBuyComputer
1	yes	young	regular	yes
2	yes	young	excellent	no
3	no	young	excellent	no
4	no	older	regular	no
5	yes	older	excellent	yes

Decision Trees

Answer:

1. Calculate the entropy of the dataset.
 - $I(D) = -(0.6 \log_2(0.6) + 0.4 \log_2 0.4) = 0.9710$
2. Calculate the Information Gain for the attribute 'Student'.
 - Split by "Student"
 - $I(stuednt = Yes) = -\left(\frac{2}{3} \log_2 \frac{2}{3} + \frac{1}{3} \log_2 \frac{1}{3}\right) = 0.918$
 - $I(stuednt = No) = -\left(\frac{2}{2} \log_2 \frac{2}{2} + 0\right) = 0$
 - $Info_{student}(D) = \frac{3}{5} \times 0.918 + \frac{2}{5} \times 0 = 0.5508$
 - $Gain_{student}(D) = I(D) - Info_{student}(D) = 0.9710 - 0.5508 = 0.4202$

Decision Trees

Answer:

3. Based on the calculated Information Gain, determine which attribute should be tested first.
 - Repeat similar calculations for “Age” and “CreditScore”.
 - Compare the information gains.
 - The attribute with the highest information gain should be tested first.
 - In this case, “student” has an information gain of 0.4202 should be tested first.

Naive Bayes Classifier

- Definition: A probabilistic classifier based on applying Bayes' theorem with strong (naive) independence assumptions between features.
- Formula: $P(y_i|x_1, x_2, \dots, x_n) = \frac{P(y_i)P(x_1, x_2, \dots, x_n|y_i)}{P(x_1, x_2, \dots, x_n)}$
- Simplified to: $P(y_i|x_1, x_2, \dots, x_n) \propto P(y_i) \prod_{k=1}^n P(x_k|y_i)$

Naive Bayes Classifier

- Application Example:
 - Spam Detection: Given a new email, calculate the probability of it being spam or not spam based on the occurrence of certain words (features).
- Laplace Smoothing:
 - Problem: Zero probabilities for unseen words in the training set.
 - Solution: Add one to each count (Laplace smoothing).
- Naive Bayes with Numeric Attributes:
 - Discretization: Convert continuous variables to categorical (e.g., temperature ranges).
 - Assumption of Distribution: Assume a specific distribution (e.g., Gaussian) for continuous variables and estimate parameters from the training data.

Naive Bayes Classifier

- Question: Imagine you are a spam filter and you have to decide if an email is spam or not spam. You are given the following information about emails:
 1. Most spam emails contain words like "win", "free", and "prize".
 2. Non-spam emails usually contain words like "meeting", "schedule", and "project".
- You receive a new email that contains the words "free" and "meeting". Based on your knowledge, what is the most likely classification for this email? Explain your reasoning.

Naive Bayes Classifier

Answer:

1. Identify Features: The email contains the words "free" and "meeting".
2. Consider Known Information:
 - The word "free" is commonly found in spam emails.
 - The word "meeting" is commonly found in non-spam emails.
3. Reasoning:
 - Since the email contains a word ("free") that is highly associated with spam, it might lean towards being classified as spam.
 - However, the presence of the word "meeting," which is associated with non-spam, suggests that it could be a legitimate email.

4. Conclusion:

- Given that "free" is a strong indicator of spam and "meeting" is a strong indicator of non-spam, the classification could go either way.
- If the filter relies more heavily on words like "free" to classify spam, it might classify this email as spam.
- Conversely, if it gives more weight to "meeting" as a non-spam indicator, it might classify it as non-spam.

Therefore, the most likely classification depends on how the filter weighs these words. If the filter considers "free" as a stronger indicator, the email will be classified as spam. If "meeting" is weighed more heavily, it will be classified as not spam.

Ensemble Methods

- Wisdom of Crowds:
 - Concept: Knowledge that emerges from collective decisions is often more accurate than that of any individual, even experts.
 - Four Conditions:
 - Diversity of Opinion: Each individual has private information.
 - Independence: Individuals' opinions are not influenced by those around them.
 - Decentralization: Decisions are made based on local knowledge without central control.
 - Aggregation: Mechanism to combine private judgments into a collective decision.
- Ensemble Learning in Machine Learning:
 - Definition: Combining multiple models to improve overall performance.
 - Example: Jelly Beans in a Jar experiment showed collective guesses are more accurate.

Ensemble Methods

- Multiple Models:
 - Train multiple models independently.
 - Majority Vote: Aggregate predictions to make final decision.
- Conditions for Effective Ensembles:
 - Accuracy: Individual classifiers should have an error rate less than 50%.
 - Diversity: Errors made by classifiers should be independent.

Ensemble Methods

- Bagging (Bootstrap Aggregating):
 - Create multiple bootstrap datasets by sampling with replacement.
 - Train a model on each dataset.
 - Combine models using majority voting to reduce variance.
- Random Forests:
 - Train multiple decision trees on different bootstrap datasets.
 - Use random subsets of features for splitting nodes.
 - Combine predictions using majority voting.
- Boosting:
 - Sequentially train weak classifiers.
 - Each classifier focuses on instances misclassified by previous classifiers.
 - Combine predictions to reduce bias and variance.

Ensemble Methods

- Question: Imagine you are explaining ensemble methods to a friend using a real-life scenario. You decide to use the example of making a group decision on where to eat dinner. Explain how the principles of ensemble learning (diversity, independence, decentralization, and aggregation) can help the group make a better decision.

Ensemble Methods

Answer:

- Diversity of Opinion:
 - Each person in the group suggests different restaurants based on their unique tastes and preferences. This diversity ensures a wide range of options is considered.
- Independence:
 - Everyone makes their suggestion without being influenced by others. This independence ensures that each opinion is based on personal preference rather than groupthink.
- Decentralization:
 - There is no single person deciding for the group. Instead, each person contributes their suggestion, allowing for a more democratic process.
- Aggregation:
 - The group uses a method to combine all suggestions, such as voting or ranking the restaurants. This aggregation ensures that the final decision reflects the collective preference of the group.

Ensemble Methods

- Real-Life Application:
 - Each friend writes down their top three restaurant choices independently (independence).
 - All choices are collected, ensuring a variety of options (diversity).
 - No single person has control over the decision (decentralization).
 - The group votes on the options, and the restaurant with the most votes is chosen (aggregation).
- This process mirrors ensemble learning, where diverse and independent models are combined to make a more accurate and reliable prediction. In this case, the group's collective decision is likely to be more satisfying to everyone compared to a decision made by a single individual.