Click to prove you're human

The set of MCQ's focuses on the Introduction to Machine learning Week 6 Solutions, emphasizing data-driven disciplines and machine learning concepts. The course aims to introduce basic mathematically motivated concepts in machine learning, covering different learning paradigms, algorithms, and architectures. A single perceptron can comput the XOR function, true. For usin a gradient descent algorithm, we need to go through five steps: first, calculate the error between the actual value and the predicted value; second, repeat this process until we find the best weights of network; third, pass an input through the network and get values from output layer; fourth, initialize random values for weight and bias; fifth, adjust these values so that the error is minimizd. The correct answer is d) 2, 3, 4, 5, 1. The backpropagation learning algorithm applied to a two-layer neural network find a locally optimal solution which may be globally optimal. This means that while the algorithm can find a good solution, it may not always find the best possible one. The other options are incorrect: a) always finds the globally optimal solution is incorrect because the algorithm may not always converge; c) never finds the globally optimal solution is also incorrect because the algorithm can find a locally optimal solution that is close to the global optimum; b) finds a globally optimal solution which is never correct because the algorithm may not always converge. Machine learning is the field of study that gives computers the capability to learn without being explicitly programmed. This is one of the most exciting technologies that one would have ever come across, making it more similar to humans: The ability to learn. Machine learning is actively being used today, perhaps in many more places than one would expect. To solve the assignments and quizzes, try to attempt them by yourself first, but if you get stuck somewhere then feel free to browse the code. You can check the complete week-wise solutions for Coursera Machine Learning All Weeks Solutions assignments and quizzes taught by Andrew Ng. After completing this course, you will get a broad idea of Machine learning algorithms.Machine Learning Week 4 assignment quiz Solution ========== Looking forwrd to seeing all the solushuns to the assigment quizz for Coursera Machine Learnng Week 4. ========== Paraphrased text here ========== Join our groups and don’t forget to subscribe. Your support is crucial in helping me continue providing quality content like this FAQ about Andrew Ng’s Machine Learning course. FAQs Is Andrew Ng's Machine Learning course good? This course is considered the best for supervised machine learning, thanks to Andrew Ng Sir’s ability to explain complex concepts with ease and provide excellent examples. His teaching style makes it easy to understand even the most challenging topics in Supervised ML. How do I get answers to Coursera assignments? To find any questions answered on Coursera assignments, use "Ctrl+F" on your computer or click on the three dots in your browser and select "Find". For mobile users, this option is easily accessible. You can also search for random questions by using these options. How long does it take to finish Coursera Machine Learning? This specialization requires approximately 3 months of study with around 75 hours of materials to complete. I completed the course in just 3 weeks and spent an additional week reviewing the entire course. It's worth noting that everyone's learning pace is different. How do you submit assignments on Coursera Machine Learning? To submit a programming assignment, follow these steps: open the assignment page, read the instructions, and download any starter files if needed. Complete the coding tasks in your local coding environment, referencing the starter files and instructions as needed. You can also watch video lectures and complete review questions for associated topics to improve your understanding. I recently completed Andrew Ng’s Machine Learning course on Coursera. The course included weekly quizzes and assignments. Here are my solutions to these assignments for reference only. Please note that it's essential to solve the assignments on your own before referring to these solutions. However, if you get stuck, feel free to use them as a reference. In this exercise, you'll implement regularized linear regression to study models with different bias-variance properties. Before starting, watch the video lectures and complete the review questions for associated topics. The course includes the following files: - ex5.m - An Octave/MATLAB script that guides you through the exercise - ex5data1.mat - A dataset for your reference - submit.m - A submission script to send your solutions to the servers - featureNormalize.m - A function for normalizing features - fmincg.m - A function minimization routine similar to fminunc - plotFit.m - A script to plot a polynomial fit - trainLinearReg.m - A script that trains linear regression using your cost function - linearRegCostFunction.m - A function to compute the cost and gradient for regularized linear regression - learningCurve.m - A script to generate a learning curve - polyFeatures.m - A script to map data into polynomial feature space - validationCurve.m - A script to generate a cross-validation curve Additionally, you can access free IOT/ML tutorials on YouTube. linearRegCostFunction.m : This function computes the cost and gradient for regularized linear regression with multiple variables. It returns the cost in J and the gradient in grad when given values of X, y, theta, and lambda. ========== grad = grad(:); learningCurve.m : function [error_train, error_val] = learningCurve(X, y, Xval, yval, lambda) %LEARNINGCURVE Generates the train and cross validation set errors needed %to plot a learning curve [error_train, error_val] = ... LEARNINGCURVE(X, y, Xval, yval, lambda) returns the train and %cross validation set errors for a learning curve. In particular, %it returns two vectors of the same length - error_train and %error_val. Then, error_train(i) contains the training error for %i examples (and similarly for error_val(i)). % % In this function, you will compute the train and test errors for % dataset sizes from 1 up to m. In practice, when working with larger % datasets, you might want to do this in larger intervals. % % Number of training examples m = size(X, 1); % You need to return these values correctly error_train = zeros(m, 1); error_val = zeros(m, 1); ========== YOUR CODE HERE ========== % Instructions: Fill in this function to return training errors in % error_train and the cross validation errors in error_val. % % i.e., error_train(i) and error_val(i) should give you the errors % obtained after training on i examples. % % Note: You should evaluate the training error on the first i training % examples (i.e., X(1:i, :) and y(1:i)). % % For the cross-validation error, you should instead evaluate on % the _entire_ cross validation set (Xval and yval). % % Note: If you are using your cost function (linearRegCostFunction) % to compute the training and cross validation error, you should % call the function with the lambda argument set to 0. % % Do note that you will still need to use lambda when running % the training to obtain the theta parameters. % % Hint: You can loop over the examples with the following: % % for i = 1:m % % Compute train/cross validation errors using training examples % % X(1:i, :) and y(1:i), storing the result in % % error_train(i) and error_val(i) % % ... % % end % % ---------- Sample Solution ---------- % DIMENSIONS: error_train = m x 1 error_val = m x 1 for i = 1:m Xtrain = X(1:i,:); ytrain = y(1:i); theta = trainLinearReg(Xtrain, ytrain, lambda); error_train(i) = linearRegCostFunction(Xtrain, ytrain, theta, 0); % for lambda = 0; error_val(i) = linearRegCostFunction(Xval, yval, theta, 0); % for lambda = 0; end % ---------- % % % % % end polyFeatures.m : function [X_poly] = polyFeatures(X, p) %POLYFEATURES Maps X (1D vector) into the p-th power % % [X_poly] = POLYFEATURES(X, p) takes a data matrix X (size m x 1) % and maps each example into its polynomial features where % X_poly(i, :) = [X(i) X(i).^2 X(i).^3 ... X(i).^p]; % % You need to return the following variables correctly. X_poly = zeros(numel(X), p); % ========== YOUR CODE HERE ========== % Instructions: Given a vector X, return a matrix X_poly where the p-th column of X contains the values of X to the p-th power. % % Here, X does not include X0 == 1 column %%%% WORKING: Using for loop %%%%% for i = 1:p X_poly(:,i) = X(:,1).^i; end %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%% X_poly(:,1:p) = X(:,1).^(1:p); % w/o for loop % % % end Check-out our free tutorials on IOT (Internet of Things): ========== for i = 1:m for j = 1:length(lambda_vec) lambda = lambda_vec(j); Xtrain = X(1:i,:); ytrain = y(1:i); % Compute train / val errors when training linear % regression with regularization parameter lambda theta = trainLinearReg(Xtrain, ytrain, lambda); error_train(j) = linearRegCostFunction(Xtrain, ytrain, theta, 0); lambda = 0; error_val(j) = linearRegCostFunction(Xval, yval, theta, 0); lambda = 0; end endPredicting outcomes in various domains is a common task in machine learning. To address this, we need to understand what type of problem it is - classification or regression. ========== Classification problems involve assigning a category or label to a new input. Examples include spam vs non-spam emails, weather forecasts (sunny, cloudy, rainy), or stock prices (higher/lower). In these cases, the goal is to predict one of several possible outcomes. Regression problems, on the other hand, aim to predict a continuous value or output variable. This can be used for forecasting temperatures, predicting the number of shares traded, or estimating future income levels. ==========The objective of predicting the number of "A" grades a student receives in their second year based on the number they received in their first year. ========== To predict the value of y, we need to determine the coefficients for linear regression. Given the training set, we can calculate and . However, specific values are not provided, so we cannot give exact values.